INTERNATIONAL JOURNAL FOR NUMERICAL METHODS IN FLUIDS, VOL. 24, 519–530 (1997)

A LOCAL MESH REFINEMENT ALGORITHM APPLIED TO TURBULENT FLOW

PETER EMVIN* AND LARS DAVIDSON

Thermo and Fluid Dynamics, Chalmers University of Technology, S-412 96 Gothenburg, Sweden

SUMMARY

This paper presents a local mesh refinement procedure based on a discretization over internal interfaces where the averaging is performed on the coarse side. It is implemented in a multigrid environment but can optionally be used without it. The discretization for the convective terms in the velocity and the temperature equation is the QUICK scheme, while the HYBRID-UPWIND scheme is used in the turbulence equations. The turbulence model used is a two-layer k–ε model. We have applied this formulation on a backward-facing step at $Re = 800$ and on a three-dimensional turbulent ventilated enclosure, where we have resolved a geometrically complex inlet consisting of 84 nozzles. In both cases the concept of local mesh refinements was found to be an efficient and accurate solution strategy.

KEY WORDS: local mesh refinement; multigrid; finite volume; SIMPLEC; k–ε; buoyancy

1. INTRODUCTION

In fluid dynamics a large number of scales are present that must be resolved. Thus regions of high shear must have a high resolution, while a smoothly varying flow allows the use of a rather coarse grid. Resolution of a boundary layer using structured grids is usually achieved by stretching the mesh close to the wall. However, to resolve convex corners, small inlets and other local features, a large region must often be refined using structured grids. Local mesh refinements are thus a useful supplement to a structured grid formulation to avoid unnecessary global refinements.

Local mesh refinements can be done in various ways but usually end in some internal boundary treatment. One choice is to use an overlap region where information is transferred between the zones. This increases the flexibility in constructing composite grids but needs careful treatment of the overlap region.[1] The second choice is to use only an internal interface without overlap and to either project boundary conditions on this interface or discretize over it. For staggered grids a conservative interpolation of velocities has often been used to obtain the internal boundary conditions for the fine grid.[2,3] This is also a suitable formulation for a multigrid implementation in accordance with the MLAT (multilevel adaptive technique) of Brandt.[4] For collocated grids the same procedure is possible but is not as easily implemented, since the location of the nodes is not on the interface and collocated algorithms normally use some kind of pressure dissipation PWIM (pressure-weighted interpolation method) to achieve control of spurious pressure oscillations.[5,6] In the present work we have therefore chosen to evaluate fluxes over the interfaces and introduce them conservatively into each subdomain.

*Peter Emvin changed his name from Peter Johansson in September 1995.

CCC 0271–2091/97/050519–12

Received February 1996
Revised September 1996

We have chosen not to use adaptive refinements but to specify the zones to be refined interactively. The refinements are specified as structured subgrids in order to use the same structured solver on each subdomain. This is usually not a limitation, since the regions that must be refined are often well approximated by only a few zones. To identify the regions to be refined, we make a preliminary simulation on a coarse mesh using the multigrid method and simply plot the multigrid residual (not volume integrated), which is proportional to the truncation error.[4] Assuming second-order convergence and a truncation error constant in the domain, we obtain the mesh size as a continuous field. We next specify a locally refined grid that well approximates our mesh size estimate and then make our final simulation on this refined grid. This procedure was shown to give the expected accuracy, which is demonstrated for the backward-facing step.

The code used is PEC-FLTBB,[7] which is a finite volume multigrid SIMPLEC-BFC code based on the work in References 6 and 8, extended by the present local mesh refinement formulation.

2. BASIC EQUATIONS

The conservation equations for stationary incompressible turbulent flow using the k–ε model can be written as

$$\nabla \cdot (\rho \mathbf{U}\Phi - \Gamma \nabla \Phi) = S. \tag{1}$$

Table I lists the source terms and diffusion coefficients for the different variables. Near walls the standard k–ε model is modified by a two-layer approach where the k-equation is solved and the turbulent length scale is prescribed using a mixing length approach.[9]

3. DISCRETIZATION

The discretization is based on a second-order finite volume formulation with a piecewise constant approximation of fluxes over the faces and of source terms in the volumes. The convective fluxes are evaluated by the QUICK scheme,[10] where a quadratic approximation of the convective flux is made by two nodes upstream and one node downstream. The diffusive flux is approximated by a second-order central difference stencil. To suppress pressure oscillations, the PWIM described in Reference 6

Table I. Parameters in general transport equation. The coefficients in the model are $C_\mu = 0{\cdot}09$, $C_{\varepsilon 1} = 1{\cdot}44$, $C_{\varepsilon 2} = 1{\cdot}92$, $\sigma_T = 0{\cdot}9$, $\sigma_k = 1{\cdot}0$ and $\sigma_\varepsilon = 1{\cdot}3$

Equation	Φ	Γ	S
Continuity	1	0	0
Momentum	U	$\mu + \mu_t$	$-\partial p/\partial x$
Momentum	V	$\mu + \mu_t$	$-\partial p/\partial y$
Momentum	W	$\mu + \mu_t$	$-\partial p/\partial z - g(\rho - \rho_0)$
Turbulent kinetic energy	k	$\mu + \mu_t/\sigma_k$	$P_k + G_k - \rho\varepsilon$
Dissipation of k	ε	$\mu + \mu_t/\sigma_\varepsilon$	$\varepsilon[C_{\varepsilon 1}(P_k + C_{\varepsilon 3}G_k) - C_{\varepsilon 2}\rho\varepsilon]/k$
Temperature	T	$\mu/Pr + \mu_t/\sigma_T$	

$P_k = \mu_t(\partial U_j/\partial x_i)(\partial U_j/\partial x_i + \partial U_i/\partial x_j)$, $G_k = -(\mu_t/\sigma_T)g\beta\partial T/\partial z$, $\mu_t = C_\mu \rho k^2/\varepsilon$, $C_{\varepsilon 3} = \tanh[|W|/\max(|U|, |V|)]$.

is used, which was shown to have stable behaviour in conjunction with the SIMPLE method. All fluxes are then assembled conservatively and, together with the source terms, we obtained the nodal residuals.

The HYBRID-UPWIND scheme[11] is used for the turbulent quantities, since it is difficult to ensure the positivity of k and ε with a non-monotone scheme. With this discretization the convergence of the velocities and pressure field is expected to be second-order, as the influence of the artificial viscosity in the turbulence equations is expected to be small.[12]

3.1. Discretization of an internal interface

Two different approaches were tested. In the first approach, coarse ghost cells were assembled on the fine side of an interface by merging four fine cells together. Using these coarse ghost cells, coarse interface fluxes were evaluated. These fluxes were then split conservatively to give the contribution to the corresponding fine cells at the other side of the interface. This was found to be an unstable discretization, where the velocities at the fine nodes close to the interface oscillated wildly. This phenomenon was also reported by Eriksson and Rai[13], who suggest that if fine ghost cells are assembled on the coarse side, it will be a stable discretization when a 2:1 relation is used, i.e. one coarse cell corresponds to two fine cells at the interface (in 2D). When adopting this technique, we achieved a stable result.

Our interface discretization is implemented with the second-order QUICK scheme and the first-order HYBRID scheme. The QUICK scheme and the pressure dissipation PWIM have a four-node support. This requires careful treatment of the cell faces at the interface F and of the cell faces on F^+ and F^- next to interface F. We thus need three rows of ghost cells on the coarse side; see Figure 1.

For each coarse cell face on the interface F and F^+ we create corresponding fine faces by splitting each coarse face into four corresponding fine faces. Fluxes on these fine faces are then evaluated as in the remainder of the domain, since the weights for the quadratic polynomial in the QUICK scheme are achieved by Lagrange polynomials using the fine ghost cells on the coarse side of the interface together with the first three rows on the fine side of the interface. The coefficients in the pressure dissipation for the interfaces are evaluated using a stencil approximating $\partial^3/\partial x^3$ over a uniform mesh with a sudden expansion over two cells of a factor two. Care is also taken so that the scaling factor in the pressure dissipation flux changes smoothly over the three incorporated faces near the internal boundary. The fine grid interface fluxes on F, F^+ and F^- are then assembled conservatively to nodal residuals.

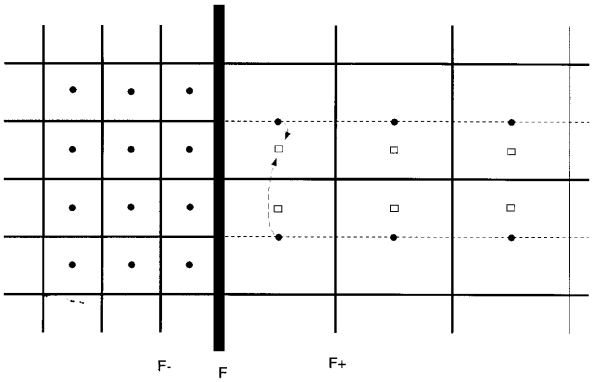

Figure 1. Interface stencil. Dots denote nodes and rectangles denote ghost nodes interpolated from coarse nodes

The near-wall modification of the turbulence model is introduced in each subdomain if the subdomain boundary coincides with a physical boundary. With this formulation we have not noticed wiggles or any other problems near the interface and have achieved the expected accuracy (i.e second-order convergence).

4. SOLUTION PROCEDURE: FULL MULTIGRID ENVIRONMENT

The solution procedure is based on the multigrid concept.[4] It reduces the amount of CPU time significantly, since the overall work of a multigrid method is linearly dependent on the size of the problem. This should be compared with the approximately quadratic dependence of standard methods such as SIMPLE or time-stepping.[4]

This specific multigrid method is a nested FMG FAS V-cycle approach using bilinear restriction and prolongation. The multigrid is applied to the whole equation system and uses the SIMPLEC method as a smoother. The internal smoothing of each transport equation is performed by the TDMA. To avoid the problem of applying a Gauss–Sidel method to a non-diagonally dominant matrix, the second-order QUICK scheme is implemented in a deferred correction manner with the first-order HYBRID-UPWIND scheme.[3] For more details on this specific multigrid method see References 6 and 8.

It has here been extended to handle the locally refined zones. There are two obvious choices of how to incorporate these refinements in a multigrid environment. The first is in accordance with the MLAT concept,[4] where all the refined zones with the same grid density constitute a grid level. In Figure 2 this means that (a) is the finest grid level, (b) is the second finest and so on. This method has asymptotic optimal behaviour, which means that a problem should converge algebraically to the level of the truncation error with the CPU time linearly dependent on the number of nodes.

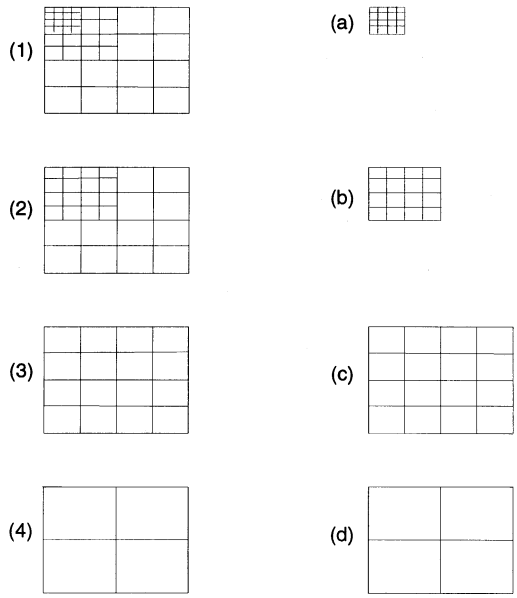

Figure 2. Multigrid structure and data structure

The second choice is to let the refined zones together with the parts of the coarser meshes that are not refined constitute a grid level. This corresponds in Figure 2 to (1) being the finest grid level, (2) the second finest and so on. This method has the penalty that more memory is needed, because all non-refined parts are stored more than once. It also requires more work at each MG cycle, since the non-refined zones are smoothed more than once. This would mean significantly more work for the case of a small refined zone with many levels of refinement. The penalty should not be too great for problems with a low number of refinement levels and large refined zones.

In Reference 14 it was shown that the work using the second method was twice as great compared with the first using six levels of refinements on a tandem wing pair. The gain, on the other hand, is that extra smoothing is performed over the interfaces that should increase the robustness, which is important when calculating complex flows.

We have chosen a blend between the two methods mentioned earlier. We store only the non-refined parts once but, when we smooth the equations, we use the global grids, where each grid level is composed of a refined part and the rest of the next coarser grid. In Figure 2 this means that the smoothing is performed on grid (1)–(4) while the storage is made only on grid (a)–(d). In our calculations we estimate the additional work compared with the MLAT on a typical engineering flow case to be about 1·5, which is a reasonable price for gained robustness. This also gives a natural option for using the refinement technique without using the multigrid, which can be nice in geometrically complex cases.

5. RESULTS

We have chosen two test cases. The first case is the backward-facing step with $Re = UH/v = 800$ ($H = 1, U = 1, v = 1/800$) and the second is a three-dimensional turbulent ventilated room where the flow is partially temperature-driven. In the first case we make successively local mesh refinements for the purpose of showing the consistency of the implementation as well as the efficiency of this algorithm in terms of CPU time and memory requirements. Our calculations are compared with the benchmark solutions of Gartling[15] and the agreement is very good.

The second case constitutes a 'real' three-dimensional application including turbulence, strong buoyancy and a geometrically complex inlet diffuser consisting of 84 small nozzles. We use our local grid refinement technique to refine the mesh at the inlet region so that each jet can be represented by four cells. We have focused on the behaviours of the interface discretization and solution procedure, but a detailed comparison of the influence of the different inlet boundary conditions and measurements is given in Reference 16.

5.1. *Backward-facing step, Re = 800*

The backward-facing step is an extensively used test case (see Figure 3). It is a model problem for channel flow with a sudden expansion. At the step a recirculation is formed and for $Re > 400$ a second separation bubble at the ceiling is present. For $Re = 1200$ the flow undergoes a transition to turbulent flow and becomes fully turbulent for $Re > 6600$.[17] We have chosen $Re = 800$ because it is close to the transition but stable,[18] which means that it is a good test on both the solver and the discretization, and because the benchmark solutions of Gartling are very accurate. Unfortunately the two-dimensionality in the measurements is lost for $Re > 400$, so we cannot make comparisons with them, but it is not in the scope of this paper to simulate the step itself as accurately as possible but to evaluate our discretization.

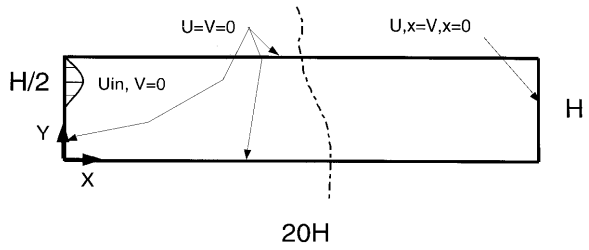

Figure 3. Backward-facing step configuration

A parabolic profile with zero vertical velocity is set at the inlet. The inlet height is half the channel height H and the channel length is set to $20H$. This gives a fully developed velocity profile at the outlet so that zero-gradient conditions are justified there. No-slip conditions are used at all walls.

The calculations were performed on a sequence of uniform grids (50×10, 100×20, 200×40, 400×80, 800×160) and on two grids with local grid refinements. The grid distribution in the two latter meshes is based on a truncation error estimate by Brandt.[4] He states that the balancing source term appearing in the multigrid is closely related to the truncation error. Thus simply plotting the truncation error estimate and assuming a second-order behaviour of the discretization gives the local relative mesh density. The level of the grid density was chosen such that the first refined mesh (Grefin1 in Table II) had 32,000 nodes and the second one (Grefin2 in Table II) 128,000 nodes. This is the same as the two finest global meshes. The distribution of grid points in these two refined grids can be seen in Figure 4. Here $X_{\mathrm{ksi}} = \partial X/\partial i$, which is equivalent to the mesh size in the X-direction.

Table II gives the results of the calculations; as can be seen, the required CPU time is linearly dependent on the grid density. The convergence criterion was set approximately a factor of four lower for each finer mesh to account for the reduced truncation error. We used five to six F-cycles on each level in the FMG process and a rough estimate is that all calculations are iteratively converged to a level at least a magnitude below the level of truncation. To reduce the algebraic error below the truncation error in one cycle, a cubic projection of the full approximation would be required. As we use linear interpolation to project the initial guess on the next finer grid in the FMG algorithm, we need at least two cycles to reduce the algebraic error to the level of the truncation error.

This means that the CPU times could not have been reduced by more than a factor of three or so using a higher convergence criterion. As a cubic projection is very expensive and complex on a distorted mesh, we have not considered that.

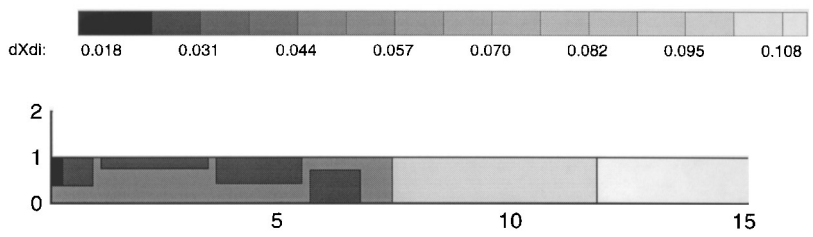

Figure 4. Grid distribution plot of patched grids

Table II. CPU times of multigrid and single-grid calculations together with endpoint co-ordinates of lower and upper recirculation

	CPU_{FMG}	CPU_{SG}	$X2_{low}$	$X1_{high}$	$X2_{high}$	DOF
10 × 50	190	195	3·60	2·1	6·5	500
20 × 100	450	1290	5·50	4·2	9·94	2000
40 × 200	1200	12150	5·96	4·66	10·38	8000
80 × 400	5000	$\sim 2 \times 10^5$	6·06	4·80	10·46	32000
160 × 800	20000	$\sim 3 \times 10^6$	6·09	4·84	10·48	128000
Grefin1	8000	$\sim 3 \times 10^6$	6·09	4·84	10·46	32000
Grefin2	32000	$\sim 5 \times 10^7$	6·10	4·85	10·48	128000
FEM			6·10	4·85	10·48	

Table II also gives the CPU times (or estimated CPU times) for the corresponding single-grid calculations. For the patched grids the size of a uniform mesh that should give equivalent accuracy is first estimated and then the CPU time for the corresponding single-grid calculation is estimated. As can be seen, the multigrid local mesh refinement method is very powerful on such large problems. From the data of the attachment points of the recirculation bubbles and the velocity profiles in Table II and Figures 5 and 6 we can see that they converge quadratically as expected. We can also see that the patched meshes have an error that is approximately 25% of that on a uniform mesh with the same number of nodes. A small exception is the reattachment point of the upper recirculation ($X2_{high}$), which has a somewhat worse behaviour.

Inconsistencies in the discretization of the internal interfaces often become visible as pressure wiggles. In Figure 7 it can clearly be seen that the present method does not give rise to any wiggles, even though there are several interfaces close to each other (separated by only a few calls). This indicates that the discretization is stable and the accuracy is indicated by the fact that the refinements give the expected effect.

5.2. 3D ventilated enclosure

This case was investigated both experimentally and numerically in an IEA project called Annex 20.[19] Here we study case E1; the configuration is shown in Figure 8. It is a room of height $H = 2{\cdot}5$ m, width $B = 3{\cdot}6$ m and length $L = 4{\cdot}2$ m. On the front wall an inlet with 84 nozzles is placed 0·285 m

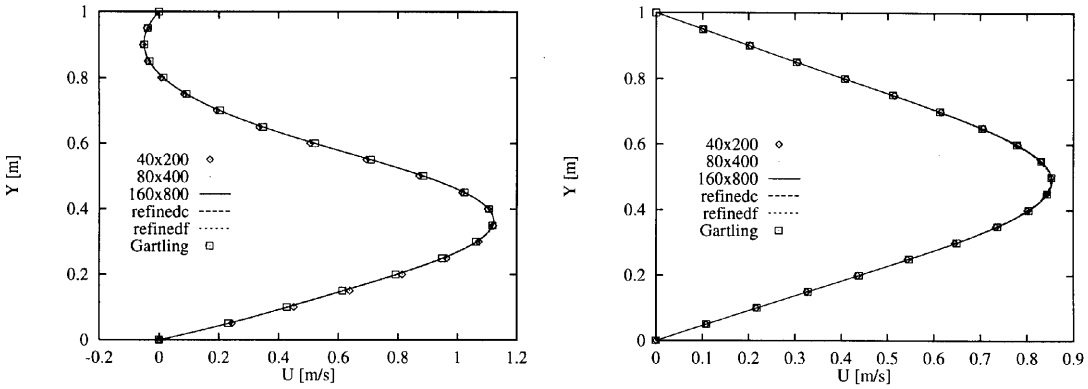

Figure 5. U-velocity profiles at $x = 7H$ and $15H$

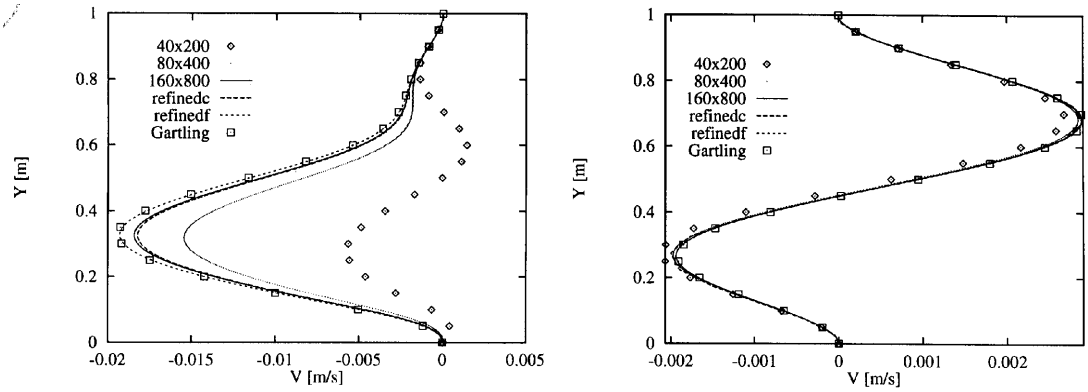

Figure 6. V-velocity profiles at $x = 7H$ and $15H$

below the ceiling and symmetrically between the side walls. The nozzles are inclined 40° towards the ceiling and their diameter is 12 mm. They are here approximated by 84 square inlet jets with the same area as the cylindrical ones. On the front wall an outlet ($H_{out} = 0{\cdot}2$ m and $B_{out} = 0{\cdot}3$ m) is placed 1·7 m from the floor and symmetrically between the side walls. On the opposite wall a window ($H_{win} = 1{\cdot}6$ m and $B_{win} = 2{\cdot}0$ m) is symmetrically located 0·7 m above the floor. The volume flux of the jet is $\dot{V} = 0{\cdot}0158$ m^3 s^{-1}, the inlet temperature is $T_{in} = 10\,^{\circ}$C and the window temperature is $T_{win} = 30\,^{\circ}$C. The walls have a constant temperature of 21 °C, except for the front wall which has a temperature of 22°C. The inlet boundary conditions of the turbulent kinetic energy and the dissipation of turbulent kinetic energy are set as $k_{in} = 0{\cdot}01(U^2 + W^2)$ and $\varepsilon_{in} = 0{\cdot}00016(U^2 + W^2)^{1{\cdot}5}/d_h$, where $d_h = 11{\cdot}5$ mm is the hydraulic diameter. At the outlet a constant mass flux is specified and a symmetry condition is used for the remaining variables. Please note that although the room has a geometrical symmetry plane of $y = 1{\cdot}8$ m, we compute the flow in

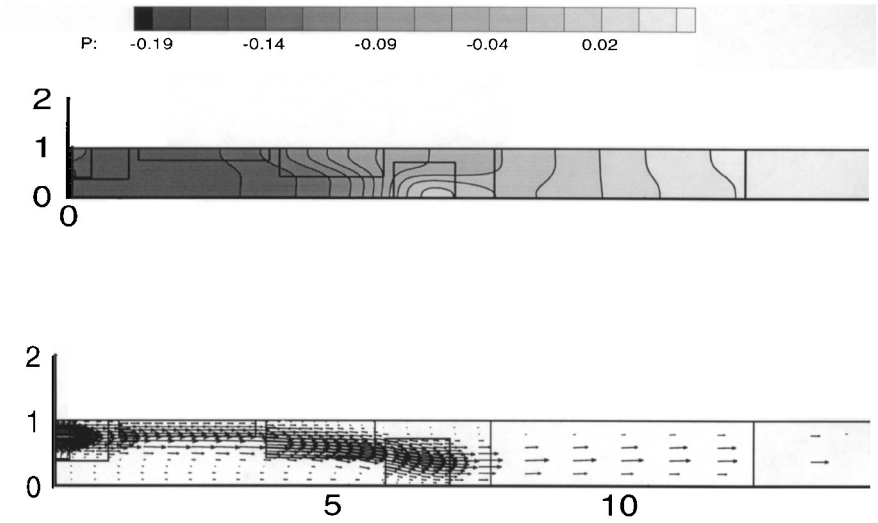

Figure 7. Pressure and velocity plots for backward-facing step, $Re = 800$

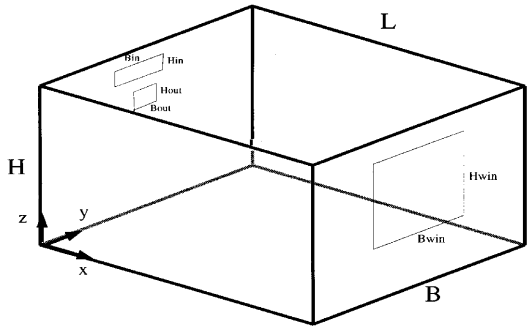

Figure 8. Three-dimensional ventilated room configuration: Annex 20, case E1

the whole room. The reason for this is that the flow has shown to be non-symmetric [6, 16, 19]. This is a very difficult flow case to simulate,[19] as the cold inlet jet should drop and hit the floor before it has reached the opposite wall.

The calculations were performed using a locally refined grid. The base mesh is a $50 \times 58 \times 34$ mesh with three levels of refinement. The first refinement level includes both inlet and outlet, while the second and third refinement levels include only the inlet region. With this mesh the inlet can be geometrically resolved and all 84 jets can be implemented. The jets in this mesh are 2×2 cells wide, but a calculation with each jet approximated by only one cell was found to give similar results in the room except in a small region ($O(d)$, where d is the jet diameter) close to the inlet, which indicates that the mesh is sufficiently fine in this region if the flow in the major part of the room is considered. The results of these calculations differ considerably from the calculations with the inlet approximated by only a single jet with the same mass flux and momentum as the sum of all nozzles.[6,16] This indicates that the physics of many small jets cannot be modelled by a single large jet, since the information on the effect of the spreading of the inlet will not be correctly accounted for. Here we

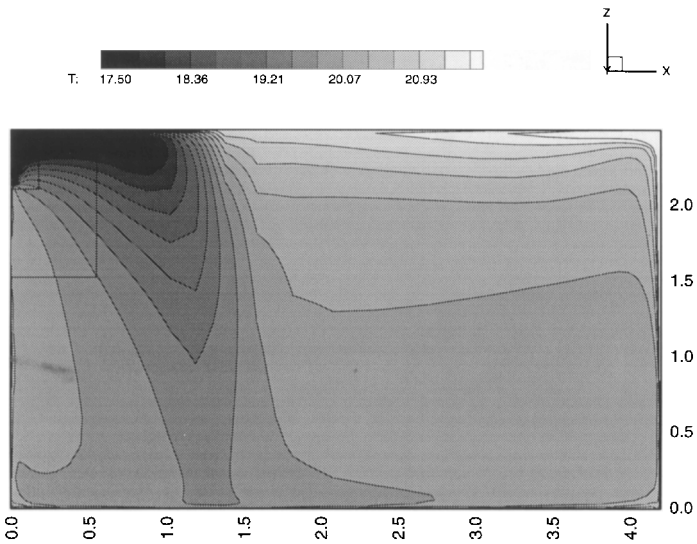

Figure 9. Temperature plot of Annex 20, E1 at $y = 1{\cdot}8$ m

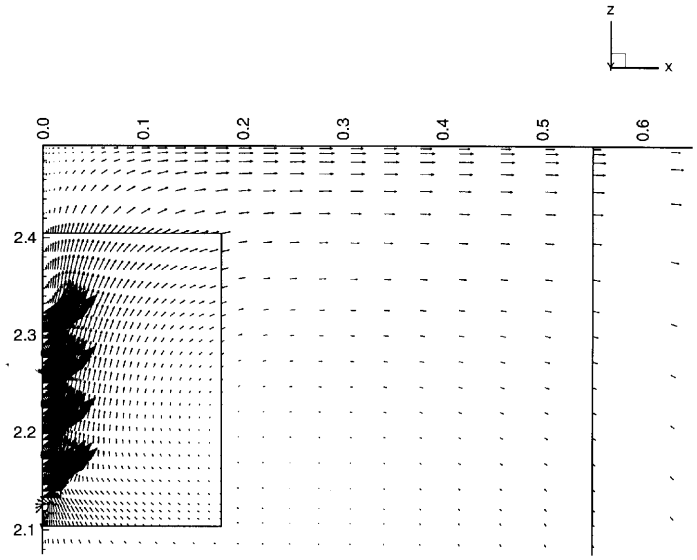

Figure 10. Velocity plot of inlet region in Annex 20, E1 at $y = 1{\cdot}8$ m

focus mainly on the numerical accuracy and the behaviour of the solution procedure. For a more detailed analysis of this simulation compared with other inlet approximations and measurements see Reference 16.

We here use the local mesh refinement strategy without the multigrid, since the small inlet jets cannot be represented on a coarse mesh. We attempted with limited success to represent the inlet jets on coarser meshes in the multigrid by decoupling the momentum and the mass flux inlet boundary condition.

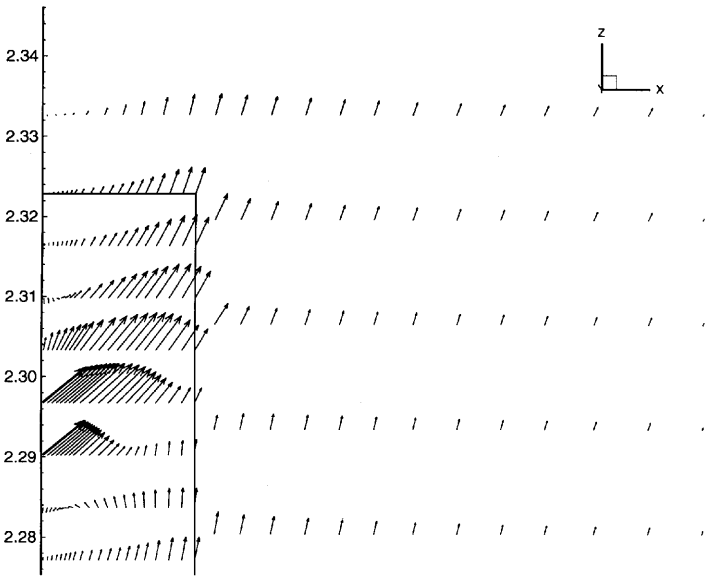

Figure 11. Velocity plot of one of the inlet nozzles in Annex 20, E1 at $y = 1{\cdot}8$ m

The calculation converged nicely in approximately 8000 iterations, which took about 400 CPU hours on a SUN-Sparc10 machine. The mesh size dependence on the base mesh in the major part of the room was analysed in a similar flow case in Reference 6 and was shown to be sufficiently fine. As in the 2D laminar problem, no wiggles or other unphysical events could be seen over the internal interfaces; see Figures 9–11.

6. CONCLUSIONS

1. The refinement procedure was shown to be accurate and robust.
2. The number of cycles in the multigrid was close to optimal, but somewhat greater work on each cycle is needed to increase the robustness.
3. The refinement procedure was shown to be robust when using a low-Reynolds k–ε turbulence model. No wiggles or unphysical events appeared.

ACKNOWLEDGEMENTS

The support of the Swedish Research Council for Engineering Sciences (TFR) is gratefully acknowledged.

APPENDIX: NOMENCLATURE

g	gravitational acceleration
h	mesh parameter
k	turbulent kinetic energy
P	pressure
Re	Reynolds number
s	source term in discretized equation
S	source term in each transport equation
T	temperature
$\mathbf{U}$	velocity vector

Greek letters

Γ	diffusion coefficient
ε	dissipation of turbulent kinetic energy
μ	dynamic viscosity
μ_t	turbulent dynamic viscosity
ν	kinematic viscosity
ρ	density
ρ_0	reference density
∇	gradient operator

REFERENCES

1. J. Y. Tu and L. Fuchs, 'Overlapping grids and multigrid methods for three-dimensional unsteady flow calculations in IC engines', *Int. j. numer. methods fluids*, **15**, 693–714 (1992).
2. Y. Li, L. Fuchs and S. Holmberg, 'An evaluation of a computer code for predicing indoor airflow and heat transfer', *Proc. 12th AIV Conf.*, Ottawa, 1991, pp. 123–136.

3. S. C. Caruso, J. H. Ferziger and J. Oliger, 'Adaptive grid techniques for elliptic fluid-flow problems', *Rep. TF-23*, Thermosciences Division, Department of Mechanical Engineering, Stanford University, 1985.
4. A. Brandt, *Multigrid Techniques: 1984 Guide with Applicaiton to Fluid Dynamics*, Computational Fluid Dynamics Lecture Notes at Von Karman Institute, 1984.
5. C. M. Rhie and W. L. Chow, 'Numerical study of the turbulent flow past an airfoil with trailing edge separation', *AIAA J.*, **21**, 1525–1532 (1983).
6. P. Johansson and L. Davidson, 'Modified collocated SIMPLEC algorithm applied to buoyancy-affected turbulent flow using a multigrid solution procedure', *Numer. Heat Transfer B*, **28**, 39–57 (1995).
7. P. Emvin, 'PEC-FLTBB, a CFD code for turbulent buoyant flow based on a local mesh refinement full multigrid strategy using boundary fitted coordinates', *Rep. 96/7*, Chalmers University of Technology, Gothenburg, 1995.
8. P. Johansson and L. Davidson, 'A full multigrid method applied to turbulent flow using the SIMPLEC algorithm together with a collocated arrangement', *Multigrid Methods IV*, Birkhäuser, 1994, pp. 245–256.
9. H. C. Chen and V. C. Patel, 'Near-wall turbulence models for complex flows including separation', *AIAA J.*, **26**, 641–640 (1988).
10. B. P. Leonard, 'A stable and accurate convective modeling procedure based on quadratic upstream interpolation', *Comput. Methods Appl. Mech. Eng.*, **19**, 59–98 (1979).
11. S. V. Patankar, *Numerical Heat Transfer and Fluid Flow*, Hemisphere, Washington, DC, 1980.
12. F. S. Lien, 'Computational modelling of 3D flow in complex ducts and passages', *Ph.D. Thesis*, University of Manchester, 1992.
13. L. E. Eriksson and M. M. Rai, 'A stability analysis of various patched grid interface conditions for hyperbolic equations', *Rep. 144*, Flygtekniska Försöksanstalten, 1988.
14. D. J. Mavripilis, *Multigrid Techniques for Unstructured Meshes*, Computational Fluid Dynamics Lecture Notes at Von Karman Institute, 1995.
15. D. Gartling, 'A test problem for outflow boundary conditions—flow over a backwards-facing step', *Int. j. numer. methods fluids*, **11**, 953–967 (1990).
16. P. Emvin and L. Davidson, 'A numerical comparison of three inlet approximations of the diffuser in case El Annex20', *5th Int. Conf. on Air Distributions in Rooms, Roomvent '96*, Tokyo, Vol. 1, 1996, pp. 219–226.
17. B. F. Armaly, F. Durst, J. F. C. Pereira and B. Schönung, 'Experimental and theoretical investigation of backwards-facing step flow', *J. Fluid Mech.*, **127**, 473–496 (1983).
18. P. M. Gresho, D. K. Gartling, J. R. Torczynski, K. A. Cliffe, T. J. Garatt, A. Spence, K. H. Winters and J. W. Goodrich, 'Is the viscous incompressible two-dimensional flow over a backwards-facing step at $Re = 800$ stable?', *Int. j. number. methods fluids*, **17**, 501–541 (1993).
19. A. D. Lemaire (ed.), *Annex 20, Room Air and Flow Contaminant Flow, Evaluation of Computational Methods, Subtask 1 Summary Report*, TNO Building and Construction Research, Delft, 1993.