Space-time window reconstruction inside finite volume simulations

- Reducing storage for scientific simulations

M.Sc.Dipl.Eng. Anton Alin-Adrian alin.anton@cs.upt.ro **Prof.Dr.Eng. Cretu Vladimir-Ioan** vladimir.cretu@cs.upt.ro	“Politehnica” University of Timisoara RO-300223, Timisoara, Romania

Abstract

The scope of this work is to reduce the storage requirements for scientific simulation data by reconstructing space-time windows defined inside the solution frame, with increased or lower mesh resolutions.

The price for handling large datasets of scientific simulations is very restrictive for most computational science and engineering users. It is necessary to develop new approaches on the handling of simulation data and provide solutions for visualization on commodity computers.

The mainstream approaches from the literature can be grouped according to Table 1.

1	Remote access to dataware centers [1]	Centralised, hard to replicate, vulnerable location
2	Specialized (parallel) filesystems [2]	Equipment-dependant, immobile
3	Databases [3]	Not really adapted to time-varying multidimensional data
4	Floating Point Compression [4,5,6,7]	Not practical for very large output or phenomenons that locally appear as random variances
5	Distributed cache prefetchers [8]	Network-dependant, suitable for LAN of commodity computers

Table 1 – Mainstream Approaches from Literature

Rows 4 and 5 in Table 1 present the implementations that make use of commodity hardware, and thus have great potential for unlocking the chains which bind computational scientists and engineers to expensive and limited resources. However the user is still bounded to high-speed network or Internet connections, and has limited improvements on the storage and transport capabilitites.

We have successfully experimented with the extraction of window-defined boundary conditions on specific time intervals, using the OpenFOAM case format, and we have used the extracted information to reconstruct the space-time subdomain contained within the configured window box. In this case, we have successfully reconstructed the velocity vector fields from the *pitzDaily* test case, according to user-defined space and time subdomains. The space and time subdomain is defined as a rectangular cuboid of interest.

The extracted points are interpolated such that meshes with a different resolution can be used during the reconstruction process. The accuracy of the reconstruction, the time required to recompute the values and the gain in storage requirements are evaluated and discussed. Different interpolation techniques and various density levels for the extracted points, are also taken into account.

In Figure 1 (a) the velocity field from the *pitzDaily* test case is presented at the final timestamp 1000. The user-defined rectangular cuboid is a bounding box of coordinates (0.075 -0.02 -0.0005) (0.15 0.01 0.0005). The boundary conditions required for reconstructing the contained subdomain are extracted and stored. The stored data is then interpolated for a new submesh resolution and the same *simpleFoam* solver is used to reconstruct the contained, missing subfield, as obtained in Figure 1 (b).

The procedure can be applied to any number of timestamps, which, in turn, can be reconstructed in parallel when the necessary hardware is available.

The slight color scale difference in Figures 1 (a) and (b) shows that the reconstruction has been done in different, unconnected case directories and the accuracy of the reconstruction is taken into account and discussed.

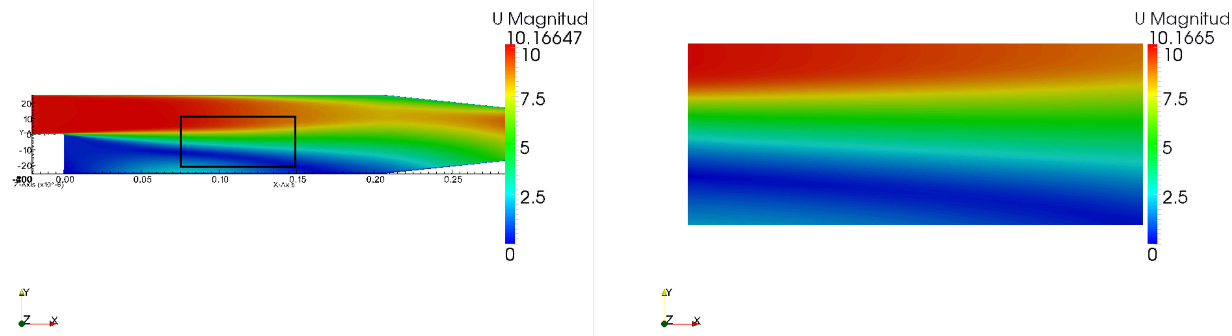

(a) – Boundary conditions are extracted for a rectangular cuboid subdomain (bounding box window)

(b) – The boundary conditions are interpolated and used to reconstruct the missing velocity fields contained within the subdomain

Figure 1: The space-time window is reconstructed using the interpolated boundary conditions

This is work in progress, which has been studied in different contexts, and is yet to be published. By only saving the necessary boundary conditions and the original geometry, the storage requirements are drastically reduced on behalf of the computational power needed to reconstruct the subfields.

Key words: space-time window, subdomain reconstruction, scientific data handling

Acknowledgements

The authors acknowledge Dr. Eng. Albert Ruprecht from the University of Stuttgart, IHS, and Prof. Dr. Eng. Romeo-Susan Resiga, from the "Politehnica" University of Timisoara, for invaluable support and research directions. We also thank Dr. Eng. Sebastian Muntean for his support.

This work is supported by the strategic grant POSDRU 6/1.5/S/13 project id 6998 of the Ministry of Labour, Family and Social Protection and co-financed by the European Social Fund.

References

[1] Jim Gray, David T. Liu, Maria Nieto-Santisteban, Alexander Szalay, David J. DeWitt, and Gerd Heber, *Scientific data management in the coming decade*, SIGMOD Rec. 34 (2005), no. 4, 34-41

[2] Jaechun No, Rajeev Thakur, and Alok Choudhary, *Integrating parallel file I/O and database support for high-performance scientific data management*, Supercomputing '00: Proceedings of the 2000 ACM/IEEE conference on Supercomputing (CDROM) (Washington, DC, USA), IEEE Computer Society, 2000, p. 57

[3] Randal Burns, Susan B. Davidson, Yannis Ioannidis, Miron Livny, and Jignesh M. Patel, *Scientific data management: An orphan in the database community?*, ICDE'08: Proceedings of the 2008 IEEE 24th International Conference on Data Engineering (Washington, DC, USA), IEEE Computer Society, 2008, p.9

[4] M. Burtscher and P. Ratanaworabhan, "FPC: A High-Speed Compressor for Double-Precision Floating-Point Data", *IEEE Transactions on Computers*, pp. 18-31, January, 2008

[5] P. Lindstrom and M. Isenburg, "Fast and efficient compression of floating-point data", *IEEE Transactions on Vizualization and Computer Graphics*, 12(5):1245-1250, 2006

[6] P. Ratanaworabhan, J. Ke, and M. Burtscher, Fast lossless compression of scientific floating-point data, *Data Compression Conference*, 0:133-142, 2006

[7] J. Wilson, Wavelet-based lossy compression of turbulence data, *Data Compression Conference*, 0:578, 2000

[8] S.S. Vazhkudai, X. Ma, V.W. Freeh, J.W. Strickland, N. Tammineedi, T. Simon, and S.L. Scott, "Constructing collaborative desktop storage caches for large scientific datasets", *Trans. Storage,* 2(3):221-254, 2006