Very large scale unsteady simulations: massively parallel unstructured mesh adaptation and multigrid solvers

Hugues Digonnet*, **Thierry Coupez*** and **Luisa Silva***

* Ecole Centrale de Nantes
High Performance Computing Institute
1 rue de la Noë, 44321 Nantes Cedex 3, France
e-mail: hugues.digonnet@ec-nantes.fr - web page: https://ici.ec-nantes.fr

ABSTRACT

In this paper, we present the work performed on three main optimizations done on a finite element based software to enable simulations on Tier0 supercomputers, which are related to parallel computing, mesh adaptation, and multigrid solvers. They allowed us to: increase the computational power using a very large number of cores; reduce the problem size, for a fixed and given accepted error, by using a heterogeneous mesh size or metrics field; reduce the number of FLOPS by lowering the algorithmic complexity when solving very large linear systems. Combining all these points enabled running extremely large scale simulations to obtain the most accurate results on complex data structures given by real geometrical or physical complexities.

Firstly, scalability of our application has been tested, by considering two main steps: parallel mesh adaptation [1] to build the partitioned mesh in order to discretize the calculation space; parallel multigrid solvers [3] using the PETSc [2] framework, as we deal with implicit discretizations, both in space and in time. Mesh adaptation is done using an error estimator on the finest mesh and so this latter is not only an uniform mesh subdivision of a coarser one and the partition of each multigrid level mesh may be different from the others. Then, code optimizations have been implemented and include the suppression of almost all the calls to the MPI Alltoall function, but also a very strict memory management, especially when building interpolation and restriction operators between two grid levels. Some discussion/comments are done on the management of input/output files, MPI processes mapping on the hardware and the double precision limitation for floating points. Parallel performances are presented up to 262,144 cores for both mesh adaptation and linear system resolution. Hard and weak speed-ups are given on both type of hardwares: Intel Xeon/ InfinyBand network and IBM BlueGene/Q hardwares. "Largest" runs in 2d and 3d are described, showing that we have been able to use almost all the cores available on Tier0 supercomputers.

Finally, more realistic unsteady simulations performed on a number of the cores up to several thousands are shown. The global scalability of the methodology is analyzed on multiphase flow computations [4]. The main test case consists in the direct simulation of bubbly flows, where the computational domain is a cube with a number of bubbles of up to 10 thousand.

REFERENCES

[1] Digonnet, H., Coupez, T., Laure, P. and Silva, L. Massively parallel anisotropic mesh adaptation. *Int. J. High Perf. Comp. Appl.* (2017).

[2] S. Balay. *PETSc users manual*, Argonne National Laboratory, ANL-95/11 Revision 3.1 (2010).

[3] Silva, L., Coupez, T. and Digonnet, H. Massively parallel mesh adaptation and linear system solution for multiphase flows. *Int. J. Comp. Fl. Dyn.*, Vol. **30** (6), pp. 431–436 (2016).

[4] Coupez, T., Silva, L. and Hachem, E. Implicit boundary and adaptive anisotropic meshing, *New Challenges in Grid Generation and Adaptivity for Scientific Computing*, Vol. **5**, pp. 1–18 (2015).