PRACE project results : performing calculations on full Tier0 supercomputers with mesh adaptation and FEM very large linear systems resolution

- ADMOS 2015 -

H. Digonnet*, L. Silva* and T. Coupez*

* Institut du Calcul Intensif (ICI)
Ecole Centrale de Nantes
1 Rue de la Noë, 44300 Nantes, France
e-mail: hugues.digonnet@ec-nantes.fr, web page: http://www.ec-nantes.fr

ABSTRACT

In this paper we will present works and results obtained during the one year PRACE project Cim128Ki [1]. This project aims to make computation on full Tier0 supercomputers or at least up to 131,072 cores.

The goal of such a project was to validate the scalability used to develop our application at such a large scale. The development context is a finite element formulation using an implicit scheme in time discretization leading to solve at the end very large linear systems. Another main axe in our strategy is doing mesh adaptation to reduce the size of the space discretization keeping the precision of the simulation unchanged. The main idea behind this to combine the benefits of every numerical technique rather than choosing one neglected others. This has become crucial has the power given by 10^5 cores allow us to deal with very large problems containing several billion unknowns.

To illustrate we want to use up to 10^5 cores but keeping anisotropic mesh adaptation [2] that could reduce the number of unknowns to solve the problem by a factor 10^3. To purchase in that way of reducing CPU time for larger problems we have implemented a parallel multigrid solver using PETSc framework [3] to reduce the algorithmic complexity for solving linear system and again reduce the number of operations done to solve the system by an other factor 10^3. At the end as we combine all these improvements we are able to reduce the CPU time by a factor 10^{11} to be compare to “only” 10^5 if we “only” take full advantage of Tier0 supercomputers.

We will first present improvements done to make us able to use more than 10^5 cores where the details become a bottleneck, like using the MPI_Alltoall function; memory and IO management; but also keeping in mind that the size of local data hosted by a core has the same order of the number of cores used.

Then we will present some parallel performances obtain during this PRACE campaign in term of hard and weak speed-up for the two main CPU consuming steps that are mesh adaptation and linear system resolution. 2d respectively 3d “biggest” runs will also be presented using two different Tier0 supercomputer (Curie : Bullx Intel/InfiniBand with 4GB/core and JuQUEEN IBM BlueGENE/Q with 1GB/core) leading to solve a 100 billion unknowns system respectively 50 billion unknowns using 64,536 cores on Curie and 262,144 cores on JuQUEEN in only some hundreds of seconds.

Finally we will present some more reasonable (using about one thousand cores) but more realistic (using complex and real data given by big 3d tomographic image or complex object) simulations.

REFERENCES

[1] H. Digonnet, M. Henri, *Cim128Ki–Enable UpScaling to 128Ki cores and Perform Full Process Simulation at 4Ki cores using CimLib,* Final Report of the 6th PRACE regular call, *2013.*

[2] H. Digonnet, L. Silva and T. Coupez, *Massivelly parallel computation on anisotropic meshes, In proceedings of 6th ADMOS conference, 2013*.

[3] S. Balay and all, *PETSc user manual,* Argonne National Laboratory, ANL-95/11 - Revision 3.1, 2010.