How GPU Computational Power Can Alter the Established Comparison of Explicit to Implicit CFD Simulations?

Mohamed hassanine Aissa, Tom Verstraete
Von Karman Institute for Fluid Dynamics, Sint-Genesius-Rode, 1640, Belgium
aissa@vki.ac.be, Verstraete@vki.ac.be

Cornelis Vuik
Delft University of Technology, 2628 CD Delft, the Netherlands
c.vuik@tudelft.nl

Abstract

The design optimization of turbomachinery components is based on CFD simulations that are used as a metric to evaluate the design's performance. The convergence and the performance of these simulations are largely depending on the time integration, which is about computing the change of state from one time instance to the next one. This change can be calculated explicitly as a combination of computed residual values or implicitly as a solution of a linear system of equations.

The difference between explicit and implicit time stepping considering stability and performance has been well covered in the literature but not extended to consider modern high performance computing systems such as Graphics Processing Units (GPU).

In this work we present the GPU implementation of the two time-stepping methods, highlighting the different challenges on the programming approach. The aim is to extend the classical comparison between explicit and implicit time stepping to take into account the performance of each method on GPU. Numerical experiments of a turbine and a compressor optimization have shown that slow converging explicit solvers can be two orders of magnitude faster on GPUs, while the performance gain of implicit solvers on GPUs suffers from inherently serial parts like the Incomplete LU (ILU) preconditioner.

Finally, our findings prove, that the choice between explicit and implicit time stepping relies mainly on the stability of explicit solvers and the efficiency of ILU implementation on GPU [1,2].

References

1. Li. Ruipeng and Y. Saad. GPU-accelerated preconditioned iterative linear solvers. The Journal of Supercomputing 63.2 (2013): 443-466.
2. E. Chow and A. Patel. Fine-Grained Parallel Incomplete LU Factorization. SIAM Journal on Scientific Computing, 37.2 (2015), C169-C193.