TOWARDS A MULTI-GPU SOLVER FOR THE THREE-DIMENSIONAL TWO-PHASE INCOMPRESSIBLE NAVIER-STOKES EQUATIONS

Peter Zaspel*, **Michael Griebel**†

*Institute for Numerical Simulation, University of Bonn
Wegelerstr. 6, D-53115 Bonn, Germany
e-mail: zaspel@ins.uni-bonn.de

†Institute for Numerical Simulation, University of Bonn
Wegelerstr. 6, D-53115 Bonn, Germany
e-mail: griebel@ins.uni-bonn.de

ABSTRACT

The use of graphics hardware for general purpose computations allows scientists to enormously speed up their numerical codes. We presently investigate the impact of this technology on our in-house CPU based finite difference solver [1] for the three-dimensional two-phase Navier-Stokes equations. To our knowledge, this is the first time, that a high order two-phase solver for the three-dimensional Navier-Stokes equations profits from the computation power of several GPUs.

As part of our project, we have implemented a GPU solver for the pressure Poisson equation based on the Jacobi preconditioned conjugate gradient method. Additionally, the reinitialization process for the level-set function used in our fluid solver is performed on the GPU. The reinitialization involves the solution of a partial differential equation of Hamilton-Jacobi type. This PDE is discretized in space by a fifth-order WENO finite difference scheme. Time discretization is done by a third-order Runge-Kutta method.

Our GPU codes were implemented in double precision using the GPU programming framework CUDA and parallelized by a domain decomposition approach with MPI. By choosing the right data exchange strategy, our multi-GPU solvers scale on large distributed memory clusters equipped with graphics hardware.

First GPU benchmarks show a major speedup factor of 16.6 for the Poisson solver and a speedup factor of 11 for the reinitialization process. These results reflect the comparison between the computation times on one GPU of an Nvidia Tesla S1070 with those of a 3.16 GHz CPU. Our multi-node/multi-GPU parallelization scales well in the number of GPUs.

In this talk, the latest results of our GPU porting project will be presented. A key focus will be on the description of our highly efficient parallelization and GPU implementation strategies.

REFERENCES

[1] R. Croce, M. Griebel, and M. A. Schweitzer, Numerical Simulation of Bubble and Droplet-Deformation by a Level Set Approach with Surface Tension in Three Dimensions, *International Journal for Numerical Methods in Fluids*, 2009, accepted