11th World Congress on Computational Mechanics (WCCM XI)
5th European Conference on Computational Mechanics (ECCM V)
6th European Conference on Computational Fluid Dynamics (ECFD VI)
July 20–25, 2014, Barcelona, Spain

MULTI-GPU UNCERTAINTY QUANTIFICATION FOR LARGE-SCALE FLOW PROBLEMS

Peter Zaspel[1], Christian Rieger[2] and Michael Griebel[2]

[1] University of Bonn, Wegelerstraße 6, zaspel@ins.uni-bonn.de
[2] University of Bonn, Wegelerstraße 6

Key words: *High Performance Computing, Multi-GPU, Two-Phase Flows, Uncertainty Quantification.*

One big problem in simulations for real-world engineering applications in industry is the appropriate handling of small uncertainties in the involved quantities. These uncertainties include but are not limited to varying material parameters, physical constraints (e.g. temperature, gravitation) and shapes of geometrical objects. We have introduced techniques for uncertainty quantification into the field of incompressible two-phase flows with the Navier-Stokes equations. Our approach to consider uncertainties is to use non-intrusive stochastic collocation methods. We approximate the stochastic parameter space by reproducing kernel Hilbert space methods, namely radial-symmetric basis function (RBF) kernels. Depending on the smoothness in the parameter space, we can achieve higher-order algebraic or even exponential convergence rates by the choice of appropriate kernels.

The big challenge of non-intrusive stochastic collocation methods is their requirement to perform calculations of hundreds or thousands of highly resolved CFD problems to extract stochastic data such as expectation value, variance or covariance. Obviously these computations push current high performance compute (HPC) clusters at their limits. Furthermore the extraction of stochastic data out of hundreds of Gigabytes to tens of Terabytes of computed data clearly is one extreme compute task.

We now attack both problems with the massively parallel compute power of GPU clusters. As a first step, the fluid solver NaSt3DGPF, a solver for the two-phase incompressible Navier-Stokes equations, was ported to multi-GPU hardware. Furthermore, we recently implemented a non-intrusive stochastic collocation method which is fully running on clusters of GPUs and allows us to perform uncertainty analysis on our fluid simulations.

In our presentation, we will highlight the numerical methods and some of the implementation details for parallel scalability and showcase several large-scale applications on multi-GPU clusters.