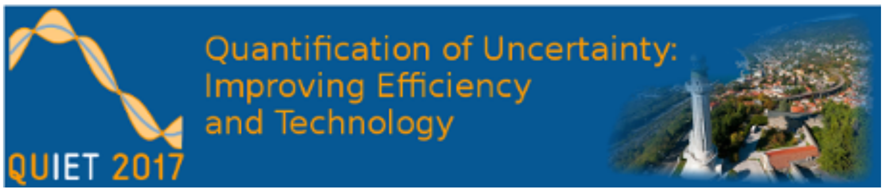

Scalable solvers for meshless methods on many-core clusters

P. Zaspel[1]

[1]University of Basel, Switzerland

Our goal is to solve large-scale stochastic collocation problems in a high-order convergent and scaling fashion. To this end, we recently discussed the radial basis function (RBF) kernel-based stochastic collocation method [1]. In this meshless method, the higher-dimensional stochastic space is sampled by (quasi-)Monte Carlo sequences, which are used as centers of radial basis functions in a collocation scheme. This non-intrusive approach combines high-order algebraic or even exponential convergence rates of spectral (sparse) tensor-product methods with good pre-asymptotic convergence of kriging, the profound stochastic framework of Gaussian process regression and parts of the simplicity of Monte Carlo methods.

Preliminary applications for this uncertainty quantification framework were (elliptic) model problems and incompressible two-phase flows with applications in chemical bubble reactors and river engineering. All solvers were parallelized to run on clusters of many-core hardware (Graphics Processing Units, GPUs) with profound scalability results.

One specific challenge of the discussed approach is the solution of a well-structured large to huge dense linear system of the type

$$\begin{pmatrix} k(\vec{y}_1,\vec{y}_1) & \dots & k(\vec{y}_1,\vec{y}_N) \\ \vdots & \ddots & \vdots \\ k(\vec{y}_N,\vec{y}_1) & \dots & k(\vec{y}_N,\vec{y}_N) \end{pmatrix} \vec{\alpha} = \begin{pmatrix} \int_\Gamma k(\vec{y}_1,\vec{y})\rho(\vec{y})d\vec{y} \\ \vdots \\ \int_\Gamma k(\vec{y}_N,\vec{y})\rho(\vec{y})d\vec{y} \end{pmatrix}$$

to compute the quadrature weights. Here, the $\vec{y}_i$ are the (quasi-)Monte Carlo samples in stochastic space and N is the sample count. Linear systems of similar type arise in Gaussian process regression and several machine learning approaches. In those cases N is the number of instances to learn from.

Classical direct factorization techniques to solve the above linear system for a large to huge kernel sample count are barely tractable, even on large parallel computers. Therefore, we discuss iterative approaches to solve such linear systems on large parallel clusters with a special emphasis on many-core hardware. To keep the iteration count small, a large-scale preconditioner with excellent strong scalability properties has been developed for GPU clusters. Moreover, we work on an optimal-complexity matrix approximation by hierarchical matrices on many-core hardware.

The presentation will cover the latest results with respect to numerical methods and applications. Performance and scalability results will be given based on studies on the Titan GPU cluster at Oak Ridge National Lab.

This work is partly based on joint work with Michael Griebel, Helmut Harbrecht and Christian Rieger.

References

[1] P. Zaspel. *Parallel RBF Kernel-Based Stochastic Collocation for Large-Scale Random PDEs.* Dissertation, Institute for Numerical Simulation, University of Bonn, Germany, 2015.