Report from Dagstuhl Seminar 16372

Uncertainty Quantification and High Performance Computing

Edited by

Vincent Heuveline[1], Michael Schick[2], Clayton Webster[3], and Peter Zaspel[4]

1 **HITS & Universität Heidelberg, DE,** vincent.heuveline@h-its.org
2 **Robert Bosch GmbH – Stuttgart, DE,** michael.schick3@de.bosch.com
3 **Oak Ridge National Laboratory, US,** webstercg@ornl.gov
4 **HITS & Universität Heidelberg, DE,** peter.zaspel@uni-heidelberg.de

Abstract

High performance computing is a key technology to solve large-scale real-world simulation problems on parallel computers. Simulations for a fixed, deterministic set of parameters are current state of the art. However, there is a growing demand in methods to appropriately cope with uncertainties in those input parameters. This is addressed in the developing research field of uncertainty quantification. Here, Monte-Carlo methods are easy to parallelize and thus fit well for parallel computing. However, their weak approximation capabilities lead to inaccurate results. The Dagstuhl Seminar 16372 "Uncertainty Quantification and High Performance Computing" brought together experts in the fields of uncertainty quantification and high performance computing. Discussions on the latest numerical techniques beyond pure Monte-Carlo and with strong approximation capabilities were fostered. This has been put in context of real-world problems on parallel computers.

Seminar September 11–16, 2016 – http://www.dagstuhl.de/16372
1998 ACM Subject Classification D.1.3 Concurrent Programming, G.1.2 Approximation, G.3 Probability and Statistics
Keywords and phrases high performance computing, parallelization, stochastic modeling, uncertainty quantification
Digital Object Identifier 10.4230/DagRep.6.9.59

1 Executive Summary

Vincent Heuveline
Michael Schick
Clayton Webster
Peter Zaspel

Topics

Uncertainty quantification (UQ) aims at approximating measures for the impact of uncertainties in e.g. simulation parameters or simulation domains. By this way, it is of great importance for both academic research and industrial development. In uncertainty quantification, one distinguishes between classical forward uncertainty propagation and more involved inference, optimization or control problems under uncertainties. Forward uncertainty

Uncertainty Quantification and High Performance Computing, *Dagstuhl Reports*, Vol. 6, Issue 9, pp. 59–73
Editors: Vincent Heuveline, Michael Schick, Clayton Webster, and Peter Zaspel

Dagstuhl Reports
Schloss Dagstuhl – Leibniz-Zentrum für Informatik, Dagstuhl Publishing, Germany

propagation is concerned with deterministic numerical models for e.g. engineering problems, in which parts of the input data (domain, parameters, ...) might be affected by uncertainties, i.e. they have a random nature. Randomness is usually characterized by random fields that replace the originally deterministic inputs. In Bayesian inference, parameters of a system shall be derived for given measurements. Since the measurements are assumed to be affected by some (stochastic) error, this inference approach tries to derive probabilities under which a given parameter leads to the observed measurements. In some sense, Bayesian inference complements classical inverse problems in a stochastic sense. Other fields of interest for a similar uncertainty analysis are optimization and control.

High performance computing (HPC) is an interdisciplinary research field in computer science, mathematics and engineering. Its aim is to develop hardware, algorithmic approaches and software to solve (usually) mathematically formulated problems on large clusters of interconnected computers. The dominant part of the involved research is done in parallel computing. From a hardware perspective, HPC or parallel computing requires to develop computing technologies that can e.g. solve several problems at the same time at high performance and low power. Moreover, hardware developments in HPC often aim at improving network communication technologies, which are necessary to let a (potentially) large set of computers solve a single problem in a distributed way. From an algorithmic perspective, methods known from numerical mathematics and data processing are adapted such that they can run in a distributed way on different computers. Here, a key notion is (parallel) scalability which describes the ability to improve the performance or throughput of a given method by increasing the number of used computers. Most algorithmic developments shall improve this scalability for numerical methods. Research in software aims at defining appropriate programming models for parallel algorithms, providing efficient management layers for the underlying hardware and implementing the proposed parallel algorithms in real software.

Challenges

In UQ, (partial) differential equations with random data are approximately solved by either intrusive or non-intrusive methods. An intrusive technique simultaneously discretizes stochastic and physical space with the classical example of stochastic Galerkin approaches. This method delivers favorable properties such as small errors with fewer number of equations and potentially small overall run-time. To achieve that, it requires to re-discretize and re-implement existing deterministic PDE solvers. On the other hand, non-intrusive techniques (e.g. (quasi-)Monte Carlo, multi-level Monte Carlo, stochastic collocation, ...) reuse existing solvers / simulation tools and generate a series of deterministic solutions which are used to approximate stochastic moments. It is thereby possible to perform uncertainty quantification analysis even for very complex large-scale applications for which a re-implementation of existing solvers is no option. The non-intrusive approach is connected to a rather extreme computational effort, with at least hundreds, thousands or even more deterministic problems that have to be solved. While a single real-world forward uncertainty propagation problem is already extremely computational intensive, even on a larger parallel computer, inference, optimization and control under uncertainties often go beyond the limits of currently available parallel computers.

In HPC, we have to distinguish methods that are intrinsically (often also called embarrassingly) parallel and those that have to exchange data to compute a result. That

is, embarrassingly parallel algorithms are able to independently compute on completely decoupled parts of a given problem. A prominent example in UQ are Monte-Carlo-type methods. The other extreme are approaches that require to exchange a lot of data in order to solve a given problem. Here, prominent examples are adaptive and multi-level methods in general and stochastic Galerkin methods. Both method types tend to have excellent approximation properties, but require a considerable effort in parallel algorithms to be scalable on parallel computers. Scalability considerations might become even more important on the next generation of the largest parallel computers, which are expected to be available at the beginning of the next decade. These parallel Exascale computers will be able to process on the exaFLOP level, thus they will be able to issue 10 18 floating-point instructions within a second. Technological limitations in chip production will force computing centers to install systems with a parallel processor count which is by orders of magnitude higher than in current systems. Current parallel algorithms might not be prepared for this next step.

The Dagstuhl Seminar on "Uncertainty Quantification and High Performance Computing", brought together experts from UQ and HPC to discuss some of the following challenging questions:

- How can real-world forward uncertainty problems or even inference, control and optimization under uncertainties be made tractable by high performance computing?
- What types of numerical uncertainty quantification approaches are able to scale on current or future parallel computers, without sticking to pure Monte Carlo methods?
- Might adaptivity, model reduction or similar techniques improve existing uncertainty quantification approaches, without breaking their parallel performance?
- Can we efficiently use Exascale computing for large-scale uncertainty quantification problems without being affected by performance, scalability and resilience problems?
- Does current research in uncertainty quantification fit the needs of industrial users? Would industrial users be willing and able to use HPC systems to solve uncertainty quantification problems?

Seminar outcome

Several presentations covered Bayesion inference / inversion (Ghattas, Marzouk, Najm, Peters), where seismology is an extremely computationally expensive problem that can only be solved by the largest parallel computers (Ghattas). While the parallelization is crucial, the numerical methods have to be adapted as well, such that fast convergence is achieved (Ghattas, Marzouk, Peters). The very computationally intensive optimization under uncertainties (Benner) becomes tractable by the use of tensor approximation methods (Benner, Osedelets). Tensor approximation methods as well as hierarchical matrices (Börm, Zaspel) are optimal complexity numerical methods for a series of applications in UQ. However their large-scale parallelization is still subject to research.

A series of talks considered mesh-free approximation methods (Rieger, Teckentrup, Zaspel) with examples in Gaussian process regression (Teckentrup) and kernel-based methods. It was possible to see that these methods have provable error bounds (Rieger, Teckentrup) and can be scaled on parallel computers (Rieger, Zaspel). Moreover these methods even fit well for inference (Teckentrup). Sparse grid techniques were considered as example for classical approximation methods for higher-dimensional problems (Stoyanov, Peters, Harbrecht, Pflüger). Here, recent developments in adaptivity and optimal convergence were discussed. Sparse grid techniques are usually considered in a non-intrusive setting such

that parallel scalability is often guaranteed. Compressed sensing promises to reduce the amount of simulations in a non-intrusive framework (Dexter). Quasi-Monte Carlo methods are under investigation for optimal convergence (Nuyens). The latter methods are of high interest for excellent parallel scalability on parallel computers due to the full decoupling of all deterministic PDE solves while keeping convergence orders beyond classical Monte Carlo methods.

Adaptivity leads to strongly improved approximations using the same amount of deterministic PDE solutions (Pflüger, Stoyanov, Webster, ...). However, a clear statement on how to parallelize adaptive schemes in an efficient way is still subject to research. The general class of multi-level schemes was also under investigation (Dodwell, Zhang), including but not being limited to multi-level Monte-Carlo and multi-level reduced basis approaches. These methods show excellent convergence properties. However their efficient and scalable parallelization is part of intensive studies, as well.

Performance considerations in the field of HPC (including future parallel computers) have been discussed (Heuveline, Legrand). Performance predictability is necessary to understand scaling behavior of parallel codes on future machines (Legrand). Parallel scalability of (elliptic) stochastic PDEs by domain decomposition has been discussed by LeMaître. His approach allows to increase parallel scalability and might show hints towards resilience.

Industrial applications were considered for the company Bosch (Schick), where intrusive and non-intrusive approaches are under investigation. High performance computing is still subject to discussion in this industrial context. One of the key applications, which is expected to become an industrial-like application, is UQ in medical engineering (Heuveline). Once introduced into the daily work cycle at hospitals, it will soon become a driving technology for our health.

Perspectives

Based on the survey and personal feedback from the invitees, the general consensus is that there is a high interest in deepening the discussions at the border of UQ and HPC. While some answers to the above questions could be given, there is still a lot more to learn, to discuss and to develop. A general wish is therefore to have similar meetings in the future.

Acknowledgements. The organizers would like to express their gratitude to all participants of the Seminar. Special thanks go to the Schloss Dagstuhl team for its extremely friendly support during the preparation phase and for the warm welcome at Schloss Dagstuhl.

2 Table of Contents

3 Overview of Talks

3.1 Optimization of Random Navier-Stokes Equations

Peter Benner (MPI – Magdeburg, DE)

We discuss the optimization and optimal control of flow problems described by the unsteady, incompressible Navier-Stokes equations. Randomness is introduced by modeling the uncertainty in the dynamic viscosity as a random variable. Using a stochastic discretization of the optimality system leads to a large-scale nonlinear system of equations in saddle point form. Nonlinearity is treated with a Picard-type iteration in which linear saddle point systems have to be solved in each iteration step. Using data compression based on separation of variables and the tensor train (TT) format, we show how these large-scale indefinite and nonsymmetric systems that typically have 10^8–10^{11} unknowns can be solved without the use of HPC technology. The key observation is that the unknown and the data can be well approximated in a new block TT format that reduces complexity by several orders of magnitude. We illustrate our findings by numerical examples.

3.2 Hierarchical tensor approximation

Steffen Börm (Universität Kiel, DE)

Joint work of Steffen Börm, Dirk Boysen, Isabelle Greff

We consider the computation of two-point correlations of the stochastic partial differential equation

$$-\Delta u(x,\omega) = f(x,\omega),$$

where x is a point in a domain D and ω is an element of a probability space. Following a result by Schwab and Todor, the two-point correlations C_u satisfy the equation

$$\Delta_x \Delta_y C_u(x,y) = C_f(x,y),$$

where C_f denotes the two-point correlations of the right-hand side. Since this is an equation in $D \times D$, the computational cost of standard discretization schemes is fairly high even if D is only a two-dimensional domain.

We propose an alternative approach: an analysis by Pentenrieder and Schwab indicates that C_u is smooth in large parts of the domain $D \times D$, so it is possible to approximate the solution by an *hp* finite element method. In order to avoid having to construct a locally refined mesh for the four- or even six-dimensional domain $D \times D$, we employ a hierarchical partition of unity in combination with suitable tensor-product functions. We introduce a recursive algorithm for constructing the sparsity pattern of the resulting system matrix. This algorithm also suggests a technique for obtaining the matrix coefficients based only on the system matrix of the original partial differential equation by using suitable inter-grid transfer operators.

3.3 Towards UQ + HPC for Bayesian Inversion, with Application to Global Seismology

Omar Ghattas (University of Texas at Austin, US)

Joint work of Tan Bui, Carsten Burstedde, Pearl Flath, Omar Ghattas, James Martin, Georg Stadler, Hari Sundar, Lucas Wilcox

Inverse problems governed by acoustic, elastic, or electromagnetic wave propagation – in which we seek to reconstruct the unknown shape of a scatterer, or the unknown properties of a medium, from observations of waves that are scattered by the shape or medium – play an important role in a number of engineered or natural systems. Our goal is to address the quantification of uncertainty in the solution of the inverse problem by casting the inverse problem as one in Bayesian inference. This provides a systematic and coherent treatment of uncertainties in all components of the inverse problem, from observations to prior knowledge to the wave propagation model, yielding the uncertainty in the inferred medium/shape in a systematic and consistent manner. Unfortunately, state-of-the-art MCMC methods for characterizing the solution of Bayesian inverse problems are prohibitive when the forward problem is expensive (as in our 100–1000 wavelength target problems) and a high-dimensional parametrization is employed to describe the unknown medium (as in our target problems involving infinite-dimensional medium/shape fields, which result in millions of parameters when discretized).

The Hessian operator of the negative log posterior plays an important role in the efficient solution of Bayesian inverse problems. When the parameter-to-observable map is linearized at the MAP point (and the prior and noise are Gaussian), the posterior is a Gaussian with the inverse Hessian as its covariance operator. More generally, this geometry-aware Gaussian approximation can be used within a proposal to accelerate MCMC methods for sampling non-Gaussian posteriors, such as in the so-called stochastic Newton, Riemannian manifold, or DILI MCMC methods.

The Hessian is often the sum of a compact operator (the data misfit) and an elliptic differential operator (the inverse prior), and this invites a low-rank approximation of the (prior-preconditioned) data misfit term, leading to an effective reduction in dimensionality (often several orders of magnitude).

Here we show that the following combination of conditions leads to a class of methods whose cost (measured in forward/adjoint PDE solves) scales independent of the parameter and data dimensions and number of processor cores:

- the prior-preconditioned data misfit Hessian is compact with mesh and data independent dominant spectrum (typical of ill-posed inverse problems)
- dominant spectrum is captured in O(r) matvecs with Hessian (we use randomized SVD for the low rank approximation)
- Hessian-vector products are computed matrix-free using second-order adjoint-based methods (amounts to 2 linearized PDE solves per matvec)
- fast O(n) elliptic solvers used for prior operator applications (we use hybrid geometric/algebraic multigrid to handle heterogeneous/anisotropic priors)
- the forward and adjoint PDE solves scale well with number of cores

The cost to construct the Laplace approximation of the posterior (or the local Gaussian at every MCMC iteration when the posterior is sufficiently non-Gaussian) is overwhelmingly

dominated by O(r) linearized forward/adjoint PDE solves (to construct the low-rank approximation). Everything else is negligible linear algebra. So when the PDE forward/adjoint solver scales well, one achieves a scalable UQ method. For mildly non-Gaussian posteriors, we evaluate the Hessian and its inverse at the maximum a posteriori point and reuse it during the MCMC iterations.

For strongly non-Gaussian posteriors, the inverse Hessian formally has to be computed repeatedly, which is intractable for large-scale, high-dimensional problems, even if the number of (linearized) forward solves is independent of the parameter and data dimensions. The challenge is to find better representations of the Hessian beyond low rank (e.g. H-matrix-based) or else to derive effective preconditioners (e.g. based on its symbol).

We present applications to a Bayesian inverse problem in global seismology with up to one million earth model parameters and 630 million state variables, on up to 100,000 processor cores.

3.4 Solution of free boundary problems in the presence of geometric uncertainties

Helmut Harbrecht (Universität Basel, CH), Marc Dambrine, Michael Peters, and Benedicte Puig

The solution of Bernoulli's exterior free boundary problem is considered in case of an interior boundary which is random. Two ways are introduced to define the expectation and the deviation of the resulting annular domain. To compare both approaches, some analytical examples for a circular interior boundary are studied. Moreover, numerical experiments are performed for more general geometrical configurations. In order to numerically approximate the expectation and the deviation, a sampling method is proposed like the (quasi-) Monte Carlo quadrature. The free boundary is determined for each sample by the trial method which is a fixed-point like iteration.

References

1 H. Harbrecht and M. Peters. *Solution of free boundary problems in the presence of geometric uncertainties.* Preprint 2015-02, Mathematisches Institut, Universität Basel, Switzerland, 2015 (to appear in Radon Series on Computational and Applied Mathematics, de Gruyter).
2 M. Dambrine, H. Harbrecht, M. Peters, and B. Puig. *On Bernoulli's free boundary problem with a random boundary.* Manuscript, 2016.

3.5 Uncertainty Quantification and High Performance Computing: Quid?

Vincent Heuveline (HITS & Universität Heidelberg)

The increasing demand on reliable results in scientific computing makes the quantification of uncertainties in mathematical models a crucial task. Including Uncertainty Quantification

to scientific computing leads for many applications to a shift of paradigm from purely deterministic problems to the stochastic models. In addition, the development of new technologies in high performance computing enables to consider new numerical methods in order to solve the challenging problems arising in Uncertainty Quantification. The talk adresses the interface between Uncertainty Quantification and High Performance Computing with a main emphasize in:

- intrusive methods in Uncertainty Quantification for systems of partial differential equations (PDEs);
- efficient accelerator and preconditioning technologies to be used on large-scale supercomputing clusters;
- open-source software-development for making the implementations accessible for the worldwide research community.

Applications in medical engineering are presented. A blood pump scenario where the inflow boundary condition, viscosity and the rotation speed are modeled as uncertain parameter is depicted. It shows up both the potential of high performance computing for uncertainty quantification but also still existing numerical challenges for real world applications.

3.6 Performance Prediction of HPC Applications: The SimGrid Project

Arnaud Legrand (INRIA – Grenoble, FR)

Simulation of HPC applications. Parallel platforms have progressively become more and more heterogeneous and complicated. After a quick presentation of typical recent supercomputers, I have presented how task-based programming and dynamic runtimes allow to efficiently exploit such architecture. Yet, evaluating performance of such complex systems is particularly challenging. I have thus presented our recent work on StarPU/SimGrid, a custom simulator that can be used to predict the performance of task-based applications running on top of StarPU to exploit hybrid (CPU+GPU) architectures. We have demonstrated the faithfulness of StarPU/SimGrid for both modern dense and sparse linear algebra solvers.

References

1 Luka Stanisic, Samuel Thibault, Arnaud Legrand, Brice Videau, Jean-François Méhaut. *Faithful Performance Prediction of a Dynamic Task-Based Runtime System for Heterogeneous Multi-Core Architectures*. Concurrency and Computation: Practice and Experience, Wiley, 2015, pp. 16.

2 Luka Stanisic, Emmanuel Agullo, Alfredo Buttari, Abdou Guermouche, Arnaud Legrand, et al.. *Fast and Accurate Simulation of Multithreaded Sparse Linear Algebra Solvers*. The 21st IEEE International Conference on Parallel and Distributed Systems, Dec. 2015, Melbourne, Australia.

3 Henri Casanova, Arnaud Giersch, Arnaud Legrand, Martin Quinson, Frédéric Suter. *Versatile, Scalable, and Accurate Simulation of Distributed Applications and Platforms*. Journal of Parallel and Distributed Computing, Elsevier, 2014, 74 (10), pp. 2899–2917.

3.7 Quasi-Monte Carlo methods for elliptic PDEs with random coefficients

Dirk Nuyens (KU Leuven, BE)

Joint work of I. Graham, Frances Y. Kuo, Dirk Nuyens, Rob Scheichl, Ian H. Sloan

I first discuss the current theory of getting dimension independent convergence in approximating high-dimensional and infinite-dimensional integrals. This is done by using weighted reproducing kernel Hilbert spaces. For several quasi-Monte Carlo methods we know function spaces and weights for which we have optimal convergence independent of the number of dimensions. Then I apply this theory to a parametrised PDE where we balance the dimension truncation error, FEM error and quadrature/cubature error. For log-normal random fields we obtain dimension-independent convergence of N^{-1} using randomly shifted lattice rules.

3.8 Bayesian Inversion for Electrical Impedance Tomography

Michael Peters (Universität Basel, CH)

Joint work of Robert Gantner, Helmut Harbrecht, Michael Peters, Markus Siebenmorgen

In this talk, we consider a Bayesian approach towards Electrical Impedance Tomography, where we are interested in computing moments, in particular the expectation, of the contour of an unknown inclusion, given noisy current measurements at the surface. By casting the forward problem into the framework of elliptic diffusion problems on random domains, we solve a suitably parametrized version by means of the domain mapping method. This straightforwardly yields parametric regularity results for the system response, which we exploit to conduct a rigorous analysis of the posterior measure, facilitating the application of sophisticated quadrature methods for the approximation of moments of quantities of interest. As an example of such a quadrature method, we consider an anisotropic sparse grid quadrature. To solve the forward problem numerically, we employ a fast boundary integral solver. Numerical examples are provided to illustrate the presented approach and validate the theoretical findings.

References

1 R. N. Gantner, M. D. Peters. *Higher Order Quasi-Monte Carlo for Baysian Shape Inversion.* Preprint 2016-18, Mathematisches Institut, Universität Basel, Switzerland, 2016.
2 A.-L. Haji-Ali, H. Harbrecht, M. Peters, and M. Siebenmorgen. *Novel results for the anisotropic sparse quadrature and their impact on random diffusion problems.* Preprint 2015-27, Mathematisches Institut, Universität Basel, Switzerland, 2015.
3 H. Harbrecht, M. Peters, and M. Siebenmorgen. *Analysis of the domain mapping method for elliptic diffusion problems on random domains.* Numer. Math., 134(4):823–856, 2016.

3.9 From Data to Uncertainty: Efficient Data-Driven Adaptive Sparse Grids for UQ

Dirk Pflüger (Universität Stuttgart, DE)

Joint work of Franzelin, Fabian; Jakeman, John; Pfander, David
Main reference F. Franzelin, D. Pflüger, "From Data to Uncertainty: An Efficient Integrated Data-Driven Sparse Grid Approach to Propagate Uncertainty", in Sparse Grids and Applications – Stuttgart 2014, LNCSE, Vol. 109, pp. 29–49, Springer, 2016.
URL http://dx.doi.org/10.1007/978-3-319-28262-6_2

We consider non-intrusive stochastic collocation for uncertainty quantification, as our applications require us to treat the underlying simulation code as a black box. We propose spatially adaptive sparse grids for both the estimation of the stochastic densities and the stochastic collocation.

With sparse grids, the numerical discretization is still possible in higher-dimensional settings, and the integrated sparse grid approach leads to fast and efficient algorithms and implementations. This allows us to start with data that is provided by measurements and to combine the estimated densities with the model function's surrogate without introducing additional sampling or approximation errors. Bayesian inference and Bayesian updating allow us to incorporate observations and to adaptively refine the surrogate based on the posterior.

Efficient and scalable algorithms for the evaluation of the surrogate function are available, which can achieve close-to-peak performance even on hybrid hardware.

3.10 Kernel methods for large scale data analysis problems arising in UQ

Christian Rieger (Universität Bonn, DE)

Many problems in uncertainty quantification (UQ) are modeled via parametric partial differential equations (PDEs). Here, the parameters often stem from a given high–dimensional space. A typical reconstruction process consists of three steps. In the first step, one has to solve the parametric PDE for a given set of parameter values. The second step is to compute some derived quantity of interest (QoI) such as a mean of the solution of the PDE for a fixed parameter. Hence, one obtains point-evaluations from a function directly mapping from the parameter space to the real numbers describing the QoI as function of the parameter. Both steps involve some numerical procedure and hence introduce a numerical error to the data. As a third step, one is often only interested in approximatively reconstructing the QoI as function from the parameter space, in order to evaluate this function for new parameter values.

In this talk, we focus on the third step. We make use of the fact that the function mapping the parameter space to the QoI is typically a smooth function of the parameter, see [1]. We present different regularization techniques which aim at saving numerical costs and solving the approximation problem up to the numerical evaluation error level stemming from the first tow steps above. To this end, we propose an error balancing strategy where we compare the numerical evaluation error of the quantity interest and the approximation error which

stems from the fact that we use only finitely many data values. Such a balancing requires an a priori error analysis in order to determine accuracy for the numerical evaluation which is needed in the first two steps. We present an error analysis based on sampling inequalities, see [2]. For the approximation we use reproducing kernels of certain problem-adapted Hilbert space, see [2]. For recent numerical examples using kernel based-methods, see also [3].

This talk is based on joint works with M. Griebel (Bonn), T. Hangelbroek (Hawaii), F. Narcowich (Texas A&M), J. Ward (Texas A&M), and P. Zaspel (Heidelberg).

References

1 A. Cohen, R. DeVore, and C. Schwab. *Analytic regularity and polynomial approximation of parametric and stochastic elliptic pde's.* Analysis and Applications, 09(01):11–47, 2011.
2 M. Griebel and C. Rieger. *Reproducing kernel Hilbert spaces for parametric partial differential equations.* 2015. To appear in SIAM/ASA Journal on Uncertainty Quantification, DOI: 10.1137/15M1026870, also available as INS Preprint No. 1511.
3 P. Zaspel. *Parallel RBF Kernel-Based Stochastic Collocation for Large-Scale Random PDEs.* Dissertation, Institut für Numerische Simulation, Universität Bonn, 2015.

3.11 A Dynamically Adaptive Sparse Grids Method for Quasi-Optimal Interpolation of Multidimensional Functions

Miroslav Stoyanov (Oak Ridge National Laboratory, US) and Clayton Webster (Oak Ridge National Laboratory, US)

In this work we develop a dynamically adaptive sparse grids (SG) method for quasi-optimal interpolation of multidimensional analytic functions defined over a product of one dimensional bounded domains. The goal of such approach is to construct an interpolant in space that corresponds to the "best M-terms" based on sharp a priori estimate of polynomial coefficients. In the past, SG methods have been successful in achieving this, with a traditional construction that relies on the solution to a Knapsack problem: only the most profitable hierarchical surpluses are added to the SG. However, this approach requires additional sharp estimates related to the size of the analytic region and the norm of the interpolation operator, i.e., the Lebesgue constant. Instead, we present an iterative SG procedure that adaptively refines an estimate of the region and accounts for the effects of the Lebesgue constant. Our approach does not require any a priori knowledge of the analyticity or operator norm, is easily generalized to both affine and non-affine analytic functions, and can be applied to sparse grids built from one dimensional rules with arbitrary growth of the number of nodes. In several numerical examples, we utilize our dynamically adaptive SG to interpolate quantities of interest related to the solutions of parametrized elliptic and hyperbolic PDEs, and compare the performance of our quasi-optimal interpolant to several alternative SG schemes.

References

1 M. Stoyanov, C. Webster, *A Dynamically Adaptive Sparse Grid Method for Quasi-Optimal Interpolation of Multidimensional Analytic Functions*, Computers & Mathematics with Applications, Vol. 71, Num. 11, pp. 2449–2465, 2016.

3.12 Gaussian process regression in Bayesian inverse problems

Aretha Teckentrup (University of Warwick – Coventry, GB) and Andrew Stuart

Main reference A. M. Stuart, A. L. Teckentrup, "Posterior Consistency for Gaussian Process Approximations of Bayesian Posterior Distributions", arXiv:1603.02004v2 [math.NA], 2016.
URL https://arxiv.org/abs/1603.02004v2

A major challenge in the application of sampling methods to large scale inverse problems, is the high computational cost associated with solving the forward model for a given set of input parameters. To overcome this difficulty, we consider using a surrogate model that approximates the solution of the forward model at a much lower computational cost. We focus in particular on Gaussian process emulators, and analyse the error in the posterior distribution resulting from this approximation.

3.13 Scalable hierarchical methods on many-core hardware – Fast matrix approximations in kernel-based collocation

Peter Zaspel (HITS & Universität Heidelberg)

It is well-known that future parallel hardware architectures will have a constantly growing number of parallel processing units. Nowadays, many-core processors (GPUs, Xeon Phi) give a first insight into the degree of parallelism that we expect to see in the future. However, it is usually said that the extreme parallelism can only be effectively used, if we apply it to "simple" algorithms. On the other hand, current optimal methods for approximation in the field of uncertainty quantification (multi-level / multi-index Monte Carlo, hierarchical matrices, ...) use complex hierarchical / tree constructions to achieve optimal complexities and approximation results. Seemingly, we have two contradicting development directions: (1) simple, very parallel algorithms; (2) complex, optimal algorithms.

This presentation shall shed light on problems and opportunities we face and have on current many-core processors if we use them to execute hierarchical algorithms. We base our discussion on our recent work in the field of radial basis function (RBF) kernel-based stochastic collocation. This non-intrusive approximation method combines high-order algebraic or even exponential convergence rates of spectral (sparse) tensor-product methods with optimal pre-asymptotic convergence of kriging and the profound stochastic framework of Gaussian process regression. Our recent applications for this approach were (elliptic) model problems and incompressible two-phase flows.

One important part of the kernel-based stochastic collocation is the solution of large to huge dense linear systems with Vandermonde-type matrices. This presentation will discuss the efficient parallel and optimal-complexity solution of these kind of linear systems by iterative solvers and fast matrix-approximations by H-matrices on many-core hardware.

Current limitations and opportunities will be highlighted.

3.14 A multilevel reduced-basis method for parameterized partial differential equations

Guannan Zhang (Oak Ridge National Laboratory, US)

Joint work of Miroslav Stoyanov and Clayton Webster

An important approximation scheme for alleviating the overall computational complexity of solving parameterized PDEs is known as multilevel methods, which have been successfully used in the Monte Carlo and collocation setting. In this effort, we propose to improve the multilevel methods with the use of reduced-basis (RB) techniques for constructing the spatial-temporal model hierarchy of PDEs. Instead of approximating the solution manifold of the PDE, the key ingredient is to build approximate manifolds of first-order differences of PDE solutions on consecutive levels. To this end, we utilize a hierarchical finite element (FE) framework to formulate an easy-to-solve variational FE system for the first-order differences. Moreover, by deriving a posteriori error estimates for the RB solutions, we also intend to develop a greedy-type adaptive strategy in order to construct a good set of snapshots. The main advantage of our approach lies in the fact that the manifold of the first-order differences becomes progressively linear as the physical level increases. Thus, much fewer expensive snapshots are required to achieve a prescribed accuracy, resulting in significant reduction of the offline computational cost of greedy algorithms. Furthermore, our approach combines the advantages of both multilevel Monte Carlo and multilevel collocation methods, in the sense that it can generate snapshots anywhere in the parameter domain but also features fast convergence.

Participants

- Peter Benner
MPI – Magdeburg, DE
- Steffen Börm
Universität Kiel, DE
- Nick Dexter
University of Tennessee – Knoxville, US
- Tim Dodwell
University of Exeter, GB
- Omar Ghattas
University of Texas at Austin, US
- Helmut Harbrecht
Universität Basel, CH
- Vincent Heuveline
HITS &
Universität Heidelberg, DE
- Olivier Le Maitre
LIMSI – Orsay, FR
- Arnaud Legrand
INRIA – Grenoble, FR
- Youssef M. Marzouk
MIT – Cambridge, US
- Habib Najm
Sandia Nat. Labs – Livermore, US
- Dirk Nuyens
KU Leuven, BE
- Ivan Oseledets
Skoltech – Moscow, RU
- Michael Peters
Universität Basel, CH
- Dirk Pflüger
Universität Stuttgart, DE
- Christian Rieger
Universität Bonn, DE
- Michael Schick
Robert Bosch GmbH – Stuttgart, DE
- Miroslav Stoyanov
Oak Ridge National Lab., US
- Aretha Teckentrup
University of Warwick – Coventry, GB
- Clayton Webster
Oak Ridge National Lab., US
- Peter Zaspel
HITS &
Universität Heidelberg, DE
- Guannan Zhang
Oak Ridge National Lab., US