SAND2011-6807P

Sandia National Laboratories

Quantification of Uncertainty in Extreme Scale Computations (QUEST)

Habib N. Najm*, Bert J. Debusschere, Michael S. Eldred
Sandia National Laboratories
*P.O. Box 969, MS 9051; Livermore, CA 94551-0969
Tel/Fax: (925) 294-2054/2595; Email: hnnajm@sandia.gov

David M. Higdon, James R. Gattiker
Los Alamos National Laboratory

Roger G. Ghanem
University of Southern California

Omar M. Knio
The Johns Hopkins University

Youssef M. Marzouk
Massachusetts Institute of Technology

Omar Ghattas, Robert Moser, Ernesto Prudencio
The University of Texas at Austin

We propose a SciDAC Institute focused on Quantification of Uncertainty in Extreme Scale Computations (QUEST). The goal of this institute is to provide modeling, algorithmic, and general uncertainty quantification (UQ) expertise, together with software tools, to other SciDAC Institutes, SciDAC applications, and Office of Science projects in general—thereby enabling and guiding a broad range of UQ activities in their respective contexts. We will establish strong ties with other projects, seek to understand their uncertainty quantification needs, and provide them with both proactive and as-needed services. We will conduct tutorials and hands-on training workshops for SciDAC researchers, providing familiarity with our suite of existing and planned software tools. We will also work individually with SciDAC researchers on their extreme-scale problems of interest, guiding their understanding and use of our UQ tools in particular applications and providing requisite UQ solutions. QUEST member institutions come equipped with a suite of software tools that have been designed for use in a broad variety of application contexts, and that have been optimized over a number of years to perform well in demanding computational environments. Starting from these tools, and targeting further developments in pursuit of greater robustness, performance, and versatility, we will deliver a suite of general UQ tools for extreme-scale computations. We will demonstrate these tools on benchmark problems that are representative of a range of problem classes (e.g., prediction and inversion for PDE, ODE, and DAE models) while presenting varied challenges, from nonlinearity and high dimensionality to missing data and fault tolerance.

In the context of UQ methodologies, QUEST will work on a broad range of activities that are critical to meaningful quantification of uncertainty in large scale scientific computing. These include: (1) statistical inference for characterizing uncertainty in model inputs; (2) adaptive methods for reduced-dimension and reduced-order stochastic representations; (3) methods for handling missing UQ samples from application code failures; (4) fault tolerance and inference of UQ information given partial data; (5) methods for forward and inverse propagation of uncertainty in computational models that are solver- and architecture-aware; and (6) statistical means for validating/comparing/selecting competing models and characterizing uncertainty arising from model inadequacy. We will also work with SciDAC Math/CS Institutes to include efficient hooks for UQ, and for the analysis and visualization of UQ results, in new solvers and software.

QUEST capabilities will be demonstrated and honed for extreme-scale computations in the context of an advanced and established software suite including DAKOTA/UQTk (SNL), QUESO (UT), and GPMSA (LANL), with robust front end interfaces including JAGUAR (SNL).

Contents

1 Executive Summary

Our proposed SciDAC institute "Quantification of Uncertainty in Extreme Scale Computations (QUEST)" is focused on uncertainty quantification (UQ) in large-scale scientific computations. It is a collaboration among six institutions with a history of in-depth collaborations on the development, implementation, and use of UQ algorithms/software in challenging high-performance computing (HPC) environments. QUEST members are the lead-developers of UQ theory, methods, and software in the technical community. They have worked in *all* aspects of the UQ problem, and have a solid grasp of the challenges and opportunities in this area. They have developed and maintain *the* leading UQ software products that have been applied in HPC environments, with challenging scientific application codes, including climate, geophysics, and combustion.

The vision of QUEST encompasses all aspects of UQ in HPC. This includes the well-founded setup of the UQ problem; characterization of the input space given available data/information; local and global sensitivity analysis; adaptive dimensionality and order reduction; forward and inverse propagation of uncertainty; handling of application code failures, missing data, and hardware/software fault tolerance; and model inadequacy, comparison, validation, selection, and averaging. The nature of the UQ problem requires the seamless combination of data, models, and information across this whole landscape in a manner that provides a self-consistent quantification of requisite uncertainties in predictions from computational models. Accordingly, our UQ methods and tools span an interdisciplinary space across applied math, information theory, and statistics.

We see our key products as being: (1) delivering expertise, advice, and state of the art UQ tools to SciDAC and Office of Science (SC) math/CS/application institutes/projects utilizing extreme scale computations on advanced computational architectures; and (2) shepherding forward our extensive repertoire of UQ theory, algorithms, and software, and enhancing their robustness/effectiveness for relevant benchmark problems in extreme-scale computational settings.

We will set up a management structure encompassing both external and internal interactions. We will manage external interactions with SciDAC/SC projects, including both proactive and reactive mechanisms, to adequately address their UQ needs. This will include tutorials and short courses on various aspects of the UQ problem, and our associated software tools. It also includes setting up a committee that conducts periodic reviews of our interactions with individual SciDAC/SC projects. We will also maintain an informative, and up to date, web presence that provides information on our available tools, and requisite installation and use information on a range of hardware/software environments. This will be a central location which provides a self-consistent front end to SciDAC/SC users, directing them optionally to talk to us directly about their needs or to the repositories of downloadable inter-operable tools at each of our member institutes. Finally, we will also work directly with other SciDAC math/CS institutes to advise them on incorporating UQ-awareness in their ongoing development of algorithms and software for extreme-scale computations.

Our internal management mechanisms will aim at facilitating the sharing of expertise and collaboration among our member institutes, and delivering effectively on our targets and milestones. We aim to work together, in a manner that benefits from extensive synergies among member institutes. Our members have a range of strengths spanning the UQ specturm. These strengths will be shared and brought to bear on the set of software solutions made available to SciDAC/SC customers. We will fund 20% of an administrative person at SNL that keeps track of our progress across the various institutes, and monitors our path forward on milestones and all other goals and obligations.

We will measure our success by the degree to which we meet the needs of SciDAC/SC math, CS, and applications institutes and projects. We will develop mechanisms for quantifying our responsiveness and effectiveness; and ensure attention towards improving our performance. We also see our success as measured by our overall impact on the scientific community, by providing a well-characterized and well-documented, robust and effective suite of UQ theory, methods, and software tools.

2 Narrative

2.1 Background and Recent Accomplishments

2.1.1 Background

Perhaps the central challenge facing the field of computational science and engineering today is: *how do we quantify uncertainties in the predictions of our large-scale simulations, given limitations in observational data, computational resources, and our understanding of physical processes?* For many societal grand challenges, the "single point" deterministic predictions delivered by most contemporary large-scale simulations of complex systems are just a first step: to be of value for decision-making (design, control, allocation of resources, policy-making, etc.), they must be accompanied by the degree of confidence we have in the predictions. This is particularly true for many of the problems in the DOE Office of Science portfolio for which large-scale simulations are essential for decision-making, including: mitigation of global climate change, siting of nuclear waste repositories, monitoring of subsurface contaminants, control of carbon sequestration processes, design of clean combustion and coal gasification systems, management of the nuclear fuel cycle, design of new nano-structured materials and energy storage systems, control of fusion energy systems, and design of particle accelerators, to name just a few.

The urgent need for uncertainty quantification (UQ) in critical scientific and societal large-scale simulations is now widely recognized across numerous scientific communities. This recognition is echoed by recent reports from a number of blue ribbon panels commissioned by Federal agencies [1, 4, 41, 69, 171, 172] as well as reports from scientific meetings and workshops, including cross-cutting [40] as well as discipline-specific reports in such fields as biology [199], basic energy sciences [92], nuclear energy systems [165], fusion energy [201], accelerator design [218], climate science [210], national security [163], energy and the environment [200], and alternative and renewable energy [39]. The awareness of UQ has reached even the popular press, as exemplified by a recent article in the New York Times on the recent M9.0 NE Japan earthquake and resulting damage to the Fukushima Daiichi nuclear facility, which contrasted Japan's "deterministic" approach to assessing tsunami and earthquake hazards, with the "probabilistic" approach taken in the U.S. [173].

UQ arises in three fundamental ways in large-scale simulations. In *stochastic inverse problems*, we seek to estimate probability densities for uncertain model inputs,* given noisy observations or measurements. In *stochastic forward problems*, we wish to propagate input uncertainties through a forward model to yield stochastic predictions of output quantities of interest (QoIs). Finally, given the ability to infer uncertain inputs from data and propagate them forward through a model, we can perform *stochastic optimization*, which solves a problem in optimal design or control under uncertainty, using statistics of the model output QoIs as objectives or constraints.

The central challenge of quantifying uncertainties in SciDAC-class simulations is that (1) they are so large and complex that even a single deterministic solution is extremely computationally challenging, and can require hours or days (or more) on petascale systems, and (2) they are characterized by high dimensional uncertainty spaces arising from uncertain volume or surface fields (e.g., initial conditions of ocean models, heterogeneous permeability of porous media models, basal boundary conditions of ice sheet models, as-built surface geometry of accelerator cavities) or complex submodel parametrizations (e.g., in combustion chemistry or systems biology models). So, while solution of stochastic inverse, forward, or optimization problems can be carried out today for inexpensive models with a modest number of uncertain parameters, *these tasks are computationally prohibitive for complex systems characterized by large-scale simulations and high-dimensional input spaces using standard Monte Carlo-type methods.* Moreover, the expected thousand-fold increase in peak performance promised by the arrival of the exascale era near the end of the decade will, by itself, do little to overcome the curse of dimensionality, *unless this performance increase is accompanied by the use of advanced UQ methods.* Fortunately, recently developed methods for systematic dimensionality reduction can be exploited to make UQ tractable for SciDAC-class applications. The challenge is to robustify, scale-up, and deploy these methods into critical SciDAC applications.

*We use the term *inputs* here to describe general model uncertainties, which might include model parameters, boundary conditions, initial conditions, sources, and so on.

Bringing advanced UQ methods and software to the Office of Science application portfolio can be done only through a sustained, concerted, and focused UQ effort, undertaken by leading authorities on UQ for large-scale simulation—in other words, a dedicated SciDAC Institute. We stress again that UQ will not be solved by simply throwing more processors at simulation codes to run larger ensembles. A partnership of application scientists, UQ experts, computer scientists, and applied mathematicians is needed to develop and deploy advanced application-aware UQ algorithms and software that can make UQ tractable for SciDAC-class simulations. Accordingly, we propose to create the *Institute for Quantification of Uncertainty in Extreme Scale Computations* (QUEST). The Institute brings together many of the leading authorities on UQ modeling, algorithms, software, applications, and education in this country. The PIs have a long history of notable contributions to UQ methods, but many of the methods have not been deployed within large-scale SciDAC application codes, due to the technical and sociological challenges of bringing UQ to large-scale complex simulations, and due to the focus in SciDAC-1 and -2 on core deterministic solvers.

The time is now ripe to launch QUEST, which will provide UQ expertise, algorithms, software, and training to SciDAC applications and Institutes—and to the broader computational science & engineering (CS&E) communities. QUEST will facilitate rigorous and rational analysis of the overall predictive fidelity of extreme scale scientific simulations, given limitations in data, physical understanding, and computational capabilities. We will establish strong ties with other centers, seeking to understand their UQ needs, and provide both proactive and as-needed services to them. We will conduct hands-on training workshops for SciDAC researchers, providing familiarity with our suite of existing and planned software tools. We will also work individually with them on their extreme-scale problems of interest, guiding their understanding and use of our UQ tools.

The extreme challenges associated with carrying out UQ for SciDAC-class simulations mean that UQ cannot be an afterthought or a side activity. UQ must be a central part of the development of the next generation of extreme-scale simulation technologies. The dominant large-scale computing paradigm that has reigned supreme for the last 20 years—the flat memory message passing model—is coming to an end, and every large-scale scientific simulation code must be re-engineered to exploit emerging manycore and accelerated architectures. With SciDAC-3 poised to ignite this transformation, there is a once-in-a-generation opportunity to take advantage of this sea-change to embed UQ thinking into the refactoring and redesign of the next generation of codes. *We believe that the creation of QUEST offers the best opportunity to catalyze a transformation of large-scale SciDAC codes, from their current single-point deterministic simulations—which do not account for uncertainties in inputs, models, numerics, and data—to simulations that rigorously and systematically quantify uncertainties in their output predictions*.

2.1.2 Team Expertise and Accomplishments

The QUEST project team combines a unique set of skills, capabilities, experience, and accomplishments as evidenced by the following:

PIs have played leading roles in the development of UQ methods and algorithms. These include intrusive methods [22, 67, 68, 74, 97, 99, 100, 104–106, 112, 136, 137, 140, 142, 187, 188, 197]; non-intrusive methods [17, 21, 68, 98, 105, 111, 138, 186]; methods for parameter inference, stochastic identification, and inverse problems [11–16, 23, 25, 31–33, 35–38, 43–45, 47, 61, 66, 71, 87, 88, 90, 91, 102, 110, 149, 155–157, 181, 217]; stochastic solvers [17, 70, 90, 139, 152, 174]; methods for epistemic and mixed uncertainties [78, 80, 125]; algorithms for optimization under uncertainty [30, 77, 79, 81, 83, 85, 161]; and optimal experimental design [123]. QUEST will thus build on a rich and substantial mathematical, statistical, and algorithmic foundation.

The team has a substantial track record in building complex UQ software and software libraries. Prior contributions in which the PIs have played a leading role include: *DAKOTA* [ME/SNL], *UQTk* [BD,HN/SNL; OK/JHU; RG/USC], *GPMSA* [DH,JG/LANL], and *QUESO* [EP/UT]. Many of these software tools were the result of collaborative efforts by large development teams, in some cases involving multiple institutions. Team members are also working to "lower the bar" to make it easier for users to deploy massive ensembles on some of today's largest supercomputers [82, 84]. Other analysis tools, such as GPMSA, provide methods and algorithms for data analysis, digesting output from an ensemble of forward model runs and physical observations, carying out analyses and making inferences.

PIs on the project team have substantial experience managing large teams and activities. Two of the PIs are currently serving as directors of large centers at universities and have had experience directing or serving on projects of similar size or larger than that proposed here [OG,RM/UT].

The PIs have a track record of attracting substantial support for UQ research through the regular peer-review process. The PIs lead [*xx*] current projects that focus on UQ, of which [*yy*] are sponsored by DOE, and [*zz*] are collaborative involving more than one institution. The latter reflect that fruitful collaborations among the PIs have been sustained over many years. QUEST will capitalize on these and prior investments, thereby accelerating the development of multipurpose tools that are essential for addressing the extreme-scale UQ challenges critical to Office of Science applications.

The team brings substantial experience in application areas relevant to the Office of Science. These include applications in: geophysics and climate [11–16, 23, 25, 47, 59, 87, 90]; chemical sciences [143, 144, 167, 187, 188]; groundwater flow [28, 90, 119, 145, 149, 162]; cosmology [117, 118, 135]; particle accelerators [15, 131, 132, 146, 148, 184]; and materials [189, 203, 212]. We believe that our experience in a broad scope of application areas will provide a solid foundation for seamless interactions with SciDAC applications and with the CS&E community at large.

Team members have substantial expertise designing, implementing, optimizing, and scaling up parallel algorithms on leading edge supercomputers. Recent examples of this work include [11, 12, 14–16, 18, 19, 24, 34, 36–38, 46, 48–53, 56, 82, 84, 87, 90, 115, 147, 150, 154, 164, 168, 176–180, 183, 185, 193, 194, 198, 205–207, 211, 216].

The PIs have a sustained record of broad dissemination and outreach. This includes publications in peer-reviewed journals (highlighted in biosketches), monographs, and conference presentations, and the release of open source software (highlighted above). In addition, the PIs have organized a large number of symposia at national and international meetings, focused workshops, summer schools, and tutorials. Selected work is highlighted in the PIs' biosketches under synergistic activities.

Prior contributions by the PIs have had a large impact on the user community. Selected metrics include: (1) DAKOTA has gained wide acceptance in the user community. Its current worldwide user base features over 7000 site registrations, with over 3500 distribution downloads occurring over the last year. A primary focus of the DAKOTA effort over the last 10 years has been on UQ algorithms. (2) GPMSA has been employed by industrial partners working with LANL, as well as other national laboratories and academic research efforts. (3) The QUESO library is being employed in the Bayesian analysis of various mathematical models by research teams in UT (turbulence, chemistry, heat conduction, earthquakes), LANL, and SNL (nuclear engineering, peridynamics).

UQ work conducted by PIs has received substantial recognition. These include institutional awards as well as national and international awards. Selected awards include: (1) Presidential Early Career Award for Scientists and Engineers [BD/SNL]; (2) IASSAR Senior Research Award and USACM Computational Structural Mechanics Award [RG/USC]; (3) Friedrich Wilhem Bessel Award [OK/JHU]; (4) DOE Early Career Research Award [YM/MIT]; (5) DOE Award of Excellence for UQ work in Science Based Stockpile Stewardship [ME/SNL]; (6) 2003 Gordon Bell Prize for Special Achievement [OG/UT]; (7) SC2002 Best Paper Award [OG/UT].

As made clear by the activities described above, we believe that the present team is uniquely positioned to address the extreme-scale UQ challenges that are central to SciDAC applications.

2.2 Proposed Research and Tasks

QUEST will provide state-of-the-art algorithms, software, and expertise to support a broad range of UQ capabilities that are needed for application to large-scale computational science problems of interest to SciDAC and

DOE-SC. The algorithmic and software challenges in UQ go well beyond the propagation of uncertainties from model inputs to output QoIs. They include: (1) accurate estimation of the uncertainty in the inputs; (2) dimensionality reduction and reduced order representations of the uncertain inputs; (3) robust treatment of system non-linearity and/or discontinuous behavior over parametric space; (4) ensuring robustness, stability, and scalability of application codes as functions of uncertain parameters; (5) fault-tolerance in UQ tools allowing for hardware and software failures; (6) efficient forward and inverse propagation of uncertainty in high dimensional spaces; and (7) effective analysis of UQ results. Moreover, computational scientists typically face additional problems in the application of UQ. These include (a) model calibration, validation, comparison, and selection under uncertainty; and (b) optimization, design, and decision support in the presence of uncertainties.

In addition to providing tools and expertise to address the UQ challenges discussed above, QUEST will pursue research to advance the state of the art on a number of algorithmic/software research fronts, as discussed below.

2.2.1 UQ algorithms

A central theme underlying our algorithmic UQ work is our use of a *probabilistic* representation of uncertain quantities, as either random variables (RVs) or processes. Probability theory provides a rich mathematical structure for characterizing an RV, e.g., by using its moments, quantiles, and/or density, by enumerating its possible values, or by describing its topological properties/geometry with respect to some Hilbertian basis. The latter approach lies behind the spectacular development, over the past 20 years, of polynomial chaos (PC) methods for UQ [105, 130, 166]. Specifically, it can be shown that any random variable with finite variance on a given probability space can be represented using a PC expansion (PCE), in terms of a basis set relative to some L_2 space with a suitably chosen measure. Examples include Hermite polynomial functions relative to Gaussian measure, and Legendre polynomial functions or wavelets relative to uniform bounded measure. These PC representations capitalize on the function analytic structure of RVs, thus permitting a more versatile yet rigorous treatment of continuity, proximity, and approximation as they pertain to RVs and processes. This mathematical structure proves important as it brings UQ analysis into the same realm as mainstream deterministic analysis. We note that one key utility of representing RVs using PCEs can be seen from considerations of forward propagation of uncertainty from model parameters/inputs to its output QoIs. While this can be done using Monte Carlo (MC) sampling of input probability density functions (PDFs), followed by binning the output statistics to construct associated PDFs, this is generally a very expensive proposition given the poor convergence rate of MC/quasi-MC methods. PC methods, by presuming a given smoothness/structure of the solution, enable the use of alternate, more efficient, means of propagation of uncertainty, as will be further detailed below. Another key utility is the ability to characterize random variables in a non-parametric fashion, thus bypassing the need to abide by standard labelled PDFs.

This is not to say that we will use PC methods exclusively; far from it. An effective UQ strategy requires flexible handling of information. We will employ PCEs as well as densities, and other means of description of RVs, as outlined in more detail below, choosing the optimal representation for the job at hand.

With this brief preamble on the general probabilistic setting, we proceed to a discussion of our range of areas of UQ algorithmic focus, briefly outlining the state of the art, challenges, and our proposed developments.

2.2.1.1 Inverse problems and characterization of uncertainty in model inputs

Characterization of uncertainty in model inputs generally involves the use of available information and/or data in order to construct a probabilistic model for the uncertain inputs/parameters. This belongs to the general class of model-based statistical inverse problems. The challenge is that inverse problems are generally "ill-posed"—that is, the data inform a limited subspace of the parameters. A convenient goal then is to describe the joint probability density function (PDF) over the inputs. Both available measurement data and/or other specified information can be used to constrain or define the parametric PDF. This can be done using Bayesian inference, combining prior information with available data in order to arrive at the posterior density. We capitalize on a range of existing tools that implement Bayesian inference in various contexts. This includes implementations of Markov chain Monte

Carlo (MCMC) methods in DAKOTA, QUESO, GPMSA, UQTk, and other software tools, as outlined in §2.2.2. Here, we discuss specific algorithmic aspects of inference relevant to this proposal.

Surrogates: Surrogates are key to efficient Bayesian inference, as they allow feasible execution times for long MCMC chains. The construction of the surrogate can itself be a non-trivial challenge, e.g., when the original function exhibits non-smooth behavior, or when it is high dimensional. Surrogates are commonly constructed employing either Gaussian Process (GP) or PC models. GPs have been used to emulate the behavior of complex computer model output since the late '80s [192]. More recently such models have been adapted to account for high dimensionality [26, 63, 121, 126]. They have been used in parameter estimation for computer model calibration and inverse problems [27, 120, 128, 195]. While such tools have proven useful for black box UQ when the ensemble size is small (< 1000), these methods are not directly applicable when the ensemble size is larger (10^5–10^7). In these cases alternative response surfaces need to be employed (e.g., [63]), or GP algorithms need to be adapted to deal with large ensemble sizes that could be produced by extreme-scale computing. We will investigate the use of sparse representations [126, 191] and integrated nested Laplace approximations [124] in this regard. PCEs have been used since the early '90s for representation of RVs for purposes of propagation of uncertainty in computational models [105]. Their utility as surrogates for purposes of inference, and the associated performance gains, were more recently recognized [157]. They have since been used successfully for accelerating Bayesian inference for parameter estimation [102, 155–157]. The construction of a PC surrogate requires the propagation of PC-formulated uncertainty through the forward model being used for parameter inference, with convenient distributions spanning the parameter range of interest (e.g., the prior support). Accordingly, model nonlinearity and high dimensionality are relevant challenges. Plans for addressing them are outlined below in the nonlinearity, dimensionality-reduction, and forward UQ sections (2.2.1.2–2.2.1.3).

MCMC methods and large-scale inverse problems: MCMC algorithms are often the method of choice for sampling posterior densities. Our tools offer state-of-the-art implementations including annealing for handling difficult/multimodal posteriors, adaptivity, and utilization of concurrent chains [60, 61, 116, 181, 182]. In particular, adaptive MCMC (AMCMC) methods use a local multivariate normal approximation to the posterior, constructed from past chain samples, as a proposal distribution for efficient and robust sampling. However, this approach is challenged by the need for a good initial set of samples, which is non-trivial in difficult posteriors. We will target the use of the analytical Hessian of the likelihood function, given an underlying PC surrogate, to construct efficient proposal distributions without reliance on an initial set of samples.

As a complement to the surrogate-based approaches proposed above, which explicitly construct a reduced representation of outputs QoIs in parameter space, we will pursue methods that effect this reduction *implicitly*. These methods, which we term *stochastic Newton* [90, 153], sample directly in the full parameter space, but adapt to the structure of the posterior density through MCMC proposals that employ local covariance matrix approximations derived from inverses of Hessians of likelihood functions, and thus implicitly detect and exploit low-dimensional manifolds for which the observations are most informative. This is particularly important for infinite-dimensional inverse problems (such as inversion for boundary condition sources or PDE coefficients). Stochastic Newton can be viewed as a natural analog to deterministic Newton methods that have been so successful for optimization problems governed by PDEs, in the sense that both exploit the Hessian to build a local model in parameter space. Newton takes a step based on the quadratic model, while stochastic Newton samples from the associated local Gaussian approximation as a proposal step for MCMC. We will capitalize on a large body of our prior work on deterministic Newton methods for inverse problems [11–16, 23, 25, 35–38, 43–45, 47, 87, 90, 110, 149, 217], which generally exhibit convergence independent of the parameter dimension, data dimension, state dimension, and number of processors. Prototype results with stochastic Newton show much faster convergence (several orders of magnitude) compared to "black-box" MCMC methods and the ability to scale to high dimensions. Our proposed algorithmic work here will focus on extension of the core stochastic Newton algorithms to multi-chain variants (to take advantage of coarse granularity parallelism), to effect low-rank approximations of the Hessians of likelihood functions (to exploit their nature as compact operators), and to reuse Hessian information to build

Hessian-informed Gaussian process proposal densities (to better handle multimodal distributions).

Absence of Data: Inference relies on availability of data. Frequently however, some information on uncertain parameters, such as nominal values and marginal bounds, is available, but not the raw data of the original experiments. While this information is useful, what is actually needed is the joint PDF on the uncertain parameters, with its specific correlation structure. This correlation can be *crucial* for reliable estimation of model output uncertainties [167]. We have recently developed a "data free inference" (DFI) procedure that deals with this challenge, by discovering the joint parametric density that is consistent with the *given* information, in the absence of raw data [29]. While this PDF is not necessarily the posterior density corresponding to the absent data, it is by construction the consensus density among all data sets that are consistent with the given nominals, bounds, data range, and fit-model structure. This procedure has been shown to work well with nonlinear exponential-decay models, and with a single-step Arrhenius-rate chemical kinetics model. We will pursue further evaluations with more complex/higher-dimensional models, covering a range of model structures. We will also pursue necessary developments targeting a range of possible constraints, such as conditional bounds and error bars in the data space.

2.2.1.2 Reduced dimension/order stochastic representations

In high-dimensional complex models, UQ feasibility relies on reduction of the dimensionality and order of stochastic representations of the input space and the solution. Paths along this direction include efficient local/global sensitivity analysis methods [54, 57], dimensionality reduction methods [21, 74, 197], and optimal representations of uncertain inputs [66, 78, 102, 108], as well as the development of efficient probabilistic representations that are adapted to particular physical settings including available data, relevant physics, and quantities of interest [65, 81]. We will adapt state-of-the-art methods to the needs of QUEST and its customers. These will include:

Sparse approximations: Current team expertise in sparse approximations includes various quadrature and cubature techniques as well as low-rank approximations using separated representations and hierarchical models [73, 93, 169]. Recent work by team members on symmetrized approximations has demonstrated complexity that scales linearly with stochastic dimension. This approach computes a low-dimensional manifold on which a PC representation is shown to exist that has the same probability measure as the solution to the problem, while discarding its L_2 convergence to the solution. We will add this approach to our conventional sparse approximations arsenal, providing optimal sparse representations depending on the goal at hand.

Adaptation of probabilistic measures: The expertise of the team in the adaptation of probabilistic measures is centered around the detection of probability measures with respect to which the characterization of the solution (rather than the parameters) has low complexity. This is typically achieved by relying on a low-fidelity approximation of the solution to estimate its probability measure, which is then used to develop the stochastic analysis of the forward problem [75]. We will work on extending this approach from its original setting (elliptic PDEs) to other more general PDE and ODE benchmark problem models.

Adaptive spatio-temporal-stochastic refinement: Adaptive stochastic refinement techniques include both h- and p-refinement strategies that can involve either dimension-adaptive (global grids defined from anisotropic tensor/sparse or generalized sparse constructions) or domain-adaptive (localized region refinement based on multi-element PCE or hierarchical surplus-driven local refinement of Lagrange interpolants) approaches [10, 81, 96, 151, 170, 208]. By further introducing adjoint-gradient enhancement into this picture, we will enrich the basis within our adaptive machinery to include higher-order regression-based PCE and Hermite interpolants (local cubic spline or global) [134, 175]. Thus, through preferential refinement in the most important regions of the stochastic input space and through the leveraging of adjoint sensitivity information with respect to the random variables, we will develop core UQ algorithms that are more scalable than currrently available methods.

Broadly, these methodologies will allow us to address the curse of dimensionality and provide high-quality uncertainty estimates in high-dimensional problems.

2.2.1.3 Handling strong nonlinearity and discontinuities

Many physical systems exhibit bifurcative behavior at certain critical parameter values. Classic examples include the transition from laminar to turbulent flow in fluid systems, and the transition from conductive to convective heat transfer in Rayleigh-Bénard flow. Generally, when the observable of interest in the computational model exhibits strong nonlinear/discontinuous response with respect to the uncertain parameters, particular attention is needed in the choice of the stochastic representation. Clearly, expansions in terms of global polynomials are not adequate in this context. Two options, not mutually exclusive, are available for dealing with this challenge. These are: (1) employing a multiblock decomposition of the stochastic space with local compact-support PC constructions on each block/element, and (2) enrichment of the conventional polynomial basis with discontinuous functions or wavelets. Team members have extensive expertise in both approaches [112, 140, 142, 144]. In general, both approaches must be pursued carefully in the context of high-dimensionality. Sparse tensorized representations are necessary for efficient PC approximation, but partitioning of the high-dimensional stochastic space must be done only in those dimensions where the response discontinuity exists. We have relied in previous work on the use of unidirectional variance measures for anisotropic stochastic mesh refinement. On the other hand, incorporation of general robust adaptive global basis enrichment is work in progress, requiring means of efficient detection and localization of discontinuities [20].

2.2.1.4 Forward propagation of uncertainty in computational models

Forward propagation of uncertainty, from inputs to output QoIs, is a major challenge. Here is where the advantage of PC methods comes strongly to the fore, as already mentioned above.

There are two general classes of PC methods for forward propagation of uncertainty. *Intrusive* methods involve a reformulation, via Galerkin projection, of the original model, arriving at a new set of governing equations for the PC coefficients of the model outputs [22, 67, 68, 75, 86, 97, 99, 100, 104–106, 112, 136, 137, 140–142, 144, 159, 160, 187, 188, 197, 209, 213–215]. This approach can lead to substantial computational savings, requiring a single direct solve, albeit of a larger set of equations. Its chief drawback is that it generally requires develpment of new solvers/software for the reformulated model, although alternate means around this have been explored [101, 107]. In contrast, *non-intrusive* methods rely on black-box sampling of the forward model to provide numerical estimation of the Galerkin projection integrals for the model output PC coefficients [17, 21, 68, 98, 103, 111, 138, 169, 170, 186]. The key advantage here is that no new solvers are necessary; however now the burden is on efficient estimation of integrals over the stochastic space. These can be evaluated in fact using MC, but this disregards the presumed smoothness in the solution. Alternatively, they can be evaluated using sparse quadrature formulae. When the dimensionality is high, cubature rules or advanced adaptive anisotropic sparse quadrature formulae [93, 96, 108, 151, 169, 170] are needed to mitigate the curse of dimensionality. It is noteworthy that sparse quadrature formulae are ill-suited for handling *noisy* integrands, because of their use of negative weights. Thus, when the forward model is intrinsically noisy, non-intrusive estimation of model output PC coefficients is best accomplished via dense quadrature formulae or by regression/Bayesian procedures using the samples as data.

Also of interest in the non-intrusive context are algorithms for identifying sparse PC representations of model outputs via *compressive sensing* (CS) [42, 58, 72, 204]. Projection approaches generally require many more sample points than basis functions for acceptable accuracy; this requirement limits their efficiency with bases of large cardinality, as in high dimensions or for strongly nonlinear functions. CS approaches, on the other hand, have shown significant early promise for generating sparse polynomial surrogates in high dimensions [73, 158] with a minimal number of model runs. These approaches may also offer additional benefits in terms of fault tolerance (§2.2.1.5) and a flexible choice of bases for capturing irregular or discontinuous model responses (§2.2.1.3).

The QUEST team has had a leadership role over the years in development and honing of all these methods, and has existing tools that cover the range of options outlined above. We will continue to refine these algorithms, targeting improved robustness and efficiency, with particular attention to extreme-scale computational problems. In the intrusive context, we will work with other SciDAC institutes (§**??**) on incorporating intrusive UQ constructions in their development of future solvers. In the non-intrusive context, we will target improvements in adaptive sampling strategies, relying on error measures in quantities of interest to direct the selection of additional samples.

2.2.1.5 Fault tolerance: handling missing UQ samples and code failure

In the batch processing of samples for UQ in extreme-scale simulations, missing samples are expected to be a common phenomenon, either due to (1) hardware failures, (2) software failures, or (3) inappropriate input conditions leading to infeasible cases. Robust approaches are needed to mitigate the effect of missing samples.

When some samples lead to infeasible cases, the sample input space needs to be adjusted to avoid those situations. For quadrature rules, this causes already computed samples to no longer be suitable. Also, the input space may become irregular, further complicating quadrature integration. To enable quadrature, irregular input spaces need to be mapped into a hypercube, e.g., using a Rosenblatt transformation [196]. New quadrature rules need to be developed, tailored to the locations of already computed samples, to avoid having to re-run all samples. Such approaches will be developed as part of QUEST.

For samples that are missing due to hardware or software failures, the original samples can be resubmitted (with adjustments to the software settings if needed). Intermediate answers can be obtained from the samples that are already available, using alternate quadrature rules or collocation approaches that are tailored to the available samples. Also, h-adaptive approaches can be set up to work around missing samples by appropriately coarsening the subdomains in the vicinity of the missing samples.

A more natural fault tolerance for missing samples due to hardware or software failures is provided by methods that rely on random sampling approaches. Point collocation and probabilistic collocation methods [122, 202] with an over-sized sample set to fit a PCE using least-squares regression can readily exclude some missing samples from the set without a significant loss in accuracy. The same holds for Galerkin projection with random sampling to evaluate the projection integrals [138, 186]. On the other hand, statistical response surface based approaches will properly account for the added uncertainty due to incomplete data [113]. Fault tolerant regression-based PCE and statistical response surface methods are implemented in the QUEST software tools (§2.2.2) and are therefore readily available to the community [8]. Further, compressive sensing approaches, as discussed in §2.2.1.4, offer an excellent combination of fault tolerance and efficiency. In CS, failed or omitted simulation information is effectively "filled in" by the choice of basis and the existence of a sparse signal representation. A second class of fault tolerant methods uses statistical approaches to backfill missing samples, e.g., through regression based on neighboring points or using imputation methods for missing data [190]. Bayesian missing data approaches allow estimation of the confidence in the backfilled samples [94, 95, 133].

As part of QUEST, we will study these various approaches for fault tolerance based on random sampling or backfilling of missing data. We will analyze the robustness and accuracy of the most promising algorithms, and make them available to the community via implementation in QUEST's software tools.

2.2.1.6 Analysis of UQ results

Uncertainty quantification in extreme-scale simulations inherently leads to the generation of massive amounts of of distributed data, which leads to significant challenges in analyzing UQ results. These challenges are most pronounced for non-intrusive approaches, particularly due to the need of first generating a suitable representation of the uncertainty from distributed quantities of interest (QoIs), and then using the resulting local representations to generate the desired statistical information.

Fundamentally, these challenges can be suitably addressed by coupling data indexing and management algorithms with the UQ analysis functionalities pertinent to the stochastic representation of the solution. For this purpose, we will rely on the *UQ Toolkit*, which provides for several functionalities, including the generation of PDFs and user-specified statistical moments. We will build on this infrastructure by providing algorithms for: (1) extracting local (derivative), partial and global sensitivity information; and (2) performing general non-linear transformations of the underlying probability measures.

2.2.1.7 Model validation

One of the key uses for UQ, both in scientific discovery and in engineering design, is model validation. The validation of a computational model is an assessment of its ability to adequately represent the processes being modeled for the specific purposes at hand. In particular, we can consider that a computational model is to be used to predict certain QoIs with some specified tolerance for error and uncertainty. The validation must thus assess the ability of the model to predict these QoIs to the required tolerance. Generally, such validation assessments must rely on comparison of model outputs to observations of the physical systems being modeled or to high fidelity "truth models" whose range of validity is well understood, and must account for the uncertainties in the observations (or truth models) as well as in the models being validated. In most cases, this assessment is complicated by the fact that the QoIs are not observable, at least at the time and under the conditions at which the predictions will be made, since otherwise predictions would not be necessary.

We have had extensive focus on model validation coupled with uncertainty [62, 109]. We will bring both this expertise and associated tools to the SciDAC community. We will also pursue a number of mathematical and algorithmic developments to advance the state of the art of UQ-aware validation processes. In particular, we will develop the following: (1) *Validation* metrics and criteria that characterize a model's ability to make the required predictions based on available data and their uncertainty. Promising approaches make use of Bayesian hypothesis testing and such metrics as mutual information and the related Kullback-Leibler divergence; (2) *Tools* for the comparison of and selection among multiple competing models based on available data and their uncertainty. One approach relies on the relative posterior probability of the models given the data (the plausibility); (3) *Algorithms* to measure the value of experimental observations as validation tests. At issue here is determining how relevant a set of measurements is to the model predictions, and the extent to which the measurements provide new information not available from data used to calibrate the model or in previous validation tests. Among the useful metrics are evaluations of the expected information gain from a set of measurements (in the sense of Shannon information theory). Such tools will also allow the *design* of experiments to be most informative [123].

2.2.1.8 Design optimization and decision support under uncertainty

Design optimization and decision support under uncertainty are necessary activities in numerous fields, including current SciDAC projects involving accelerator design, fusion pellet design, and control of subsurface contaminants. In the presence of uncertainty, optimization and decision support strategies need to comprehend the underlying uncertainties in computational predictions. Frequently, this involves statistical measures/targets of system performance, and associated probabilistic decision support machinery. Commonly used statistical performance measures include robustness (e.g., QoI variance) and/or reliability (e.g., QoI probability of failure) metrics [9, 77, 81].

Optimization strategies are facilitated by the use of PC representations. Given the PCE for the response QoI, analytic expressions can be derived for the moments of the expansions as well as for various sensitivity measures (e.g., derivatives of moments with respect to design variables [77, 78]), thereby allowing for efficient design-under-uncertainty formulations. Relative to similar approaches based on random sampling or local reliability UQ methods [9, 79, 83, 85], PC-based approaches [77, 161] provide fast convergence and mitigate algorithmic robustness issues due to nonsmoothness, multimodality, and high degrees of nonlinearity in the response QoI.

2.2.2 UQ software

In order to support a broad variety of SciDAC application projects, we will pursue ongoing development of a family of flexible and extensible object-oriented software tools for performing core uncertainty quantification (UQ) as well as higher level UQ-enabled analyses. Key objectives include both *performance*, through improving the robustness and efficiency of UQ algorithms for complex, high-dimensional problems, and *usability*, by "lowering the bar" for adoption of UQ by SciDAC application institutes. These software tools provide an essential conduit in deploying the latest emerging UQ capabilities to the applications.

2.2.2.1 Core UQ software

Tools for *forward uncertainty propagation* include DAKOTA, and UQTk:

DAKOTA (SNL) [2] provides a variety of non-intrusive algorithms for design optimization, model calibration, uncertainty quantification, global sensitivity analysis, solution verification, and parameter studies [8]. The UQ area has been a primary emphasis over the last 10 years and the DAKOTA team has developed algorithms covering aleatory, epistemic, and mixed UQ, including random and stratified sampling, local and global reliability, stochastic expansions (polynomial chaos, stochastic collocation), and epistemic methods (interval propagation, Dempster-Shafer). Alongside these iterative algorithms, a computational model abstraction provides flexible support for surrogate and multifidelity modeling, algorithm nesting, and problem formulation recasting (including variable transformations and dimension reduction). These iterator/model foundational components allow the construction of more advanced solution strategies combining multiple algorithms with models of varying fidelity, thereby providing an environment for exploration of the multi-component solution approaches proposed herein, including hybrid solution methods, design under uncertainty, mixed aleatory-epistemic UQ, and Bayesian inference and model selection. Finally, DAKOTA supports multilevel parallel computing through communicator partitioning and recursive scheduling, and can be used as either a stand-alone application or as a set of library services.

UQ Toolkit (UQTk, SNL) [7] is a lightweight C++ library, primarily offering tools for intrusive Galerkin uncertainty propagation. UQTk provides routines for evaluating algebraic expressions and transcendental functions of random variables represented with PC expansions [68]. Non-intrusive quadrature-based Galerkin projection methods are also provided. The toolkit supports a variety of standard as well as custom basis functions. UQTk is particularly well suited for algorithm prototyping as well as for educational and tutorial purposes.

Tools for *inverse problems, inference, and design* include QUESO, GPMSA, and DAKOTA:

QUESO (UT) is a MPI/C++ library that provides statistical algorithms for Bayesian inference, model calibration, model validation, and decision making under uncertainty [182]. It naturally maps, into C++ classes, the mathematical entities present in stochastic problems and solution methods, thus enabling the easy integration of new algorithms. It currently includes several variants of MCMC algorithms, Monte Carlo sampling, MCMC convergence assessment capabilities and information theoretic metrics. QUESO supports parallelism through the computation of multiple chains, as well as parallelism in the model application. New algorithms are being introduced into QUESO, such as parallel adapative multilevel algorithms (for multimodal posterior PDFs) and stochastic Newton.

GPMSA (LANL) also focuses on Bayesian inference, using a Gaussian process response surface, trained from an ensemble of forward model runs, to minimize the number of forward model calls required in the inference [121]. This software allows for global sensitivity analysis, forward propagation of uncertainty, model calibration/parameter estimation, and predictions with uncertainty. GPMSA also provides a limited form of statistical models to characterize the *model discrepancy* or *structural model error*, a term describing how simulated outputs might vary from the corresponding outputs on the physical system of interest. More generic approaches for characterizing model discrepancy will be developed. In its current state, GPMSA is a serial code. This effort will extend the software and design new methods and algorithms that are tailored for HPC environments. This will include new software approaches for response surfaces, sensitivity analyses, and parallel MCMC.

DAKOTA (SNL) provides a variety of UQ-enabled analyses, including algorithms for design and calibration under uncertainty (§2.2.1.8) that are tailored to the structure of the underlying UQ methods [30, 77, 79, 81, 83, 85].

***Foundational utilities* that are shared among higher-level tools include Pecos and Percept:**

Pecos (SNL) [5] is a utility library that provides lower level UQ services, including numerical integration drivers (LHS, quadrature, cubature, sparse grids), variable transformations (Nataf, Rosenblatt), polynomial basis functions (global or piecewise, orthogonal or interpolation, nodal or hierarchical, value-based or derivative-enhanced), and inverse FFT approaches for stochastic process modeling.

Percept (SNL) is currently part of the Trilinos STK package [6] and provides tools for verification (code and solution) and mesh modification (uniform refinement, enrichment, and element type conversion). As part of the development of a standalone Percept package in Trilinos, additional capabilities are being added for local adaptive mesh refinement, post-processing quantities of interest (QoIs), error estimation in these QoIs, and realizations of spatial random fields discretized by Karhunen-Loève expansions.

A key feature of this collection of tools is the interoperability of reusable software components, in the hierarchy defined from problem-solving environments and frameworks through algorithm libraries and utilities. Current inter-relationships include DAKOTA/QUESO/GPMSA, as well as Trilinos/DAKOTA, UQTk/Percept/Pecos. First, both QUESO and GPMSA are integrated as part of DAKOTA, with the software intent of providing general plug-and-play inference facilities using a variety of samplers, likelihoods, emulators, and post-processors. DAKOTA is likewise integrated as part of QUESO in order to provide forward UQ services following Bayesian calibration. Second, Trilinos provides the TriKOTA package, which provides a convenient delivery vehicle for utilizing DAKOTA library services when already employing other Trilinos services. Third, UQTk provides services for intrusive projection through numerical expressions to DAKOTA. Under QUEST, these inter-relationships will be further extended to provide a unified set of best-in-class tools for UQ, inference, and design. In particular, greater coupling among DAKOTA, QUESO, and GPMSA non-intrusive tools, connection of UQTk to DAKOTA/QUESO/GPMSA for leveraging intrusive tools in inverse/inference/design, and shared lower level foundations within Pecos and Percept are all anticipated.

2.2.2.2 Usability software

Tools and approaches for *lowering the bar* for adoption of UQ methods within SciDAC applications include library embedding, and JAGUAR:

Library embedding (all) eases the burden of custom simulation interface development for UQ studies through integrating input streams, eliminating the need for custom drivers and file processing, streamlining simulation flow and I/O, enabling simulation warm-starting, and simplifying parallel execution on MPP platforms.

JAGUAR (SNL) [3] is a JAVA Eclipse-based front end for DAKOTA that provides widget-based input specification, an input editor with syntax completion and embedded help, and simplified analysis wizards that streamline problem set-up for common studies. Future focus areas include simulation interfacing and UQ visualization tools.

These tools work together to make UQ more directly accessible to SciDAC application customers by automating problem definition, simulation interfacing, and computational resource management. By addressing these common barriers to the adoption of UQ technologies, the QUEST institute will enable the wide-spread adoption of UQ algorithms and UQ-enabled methodologies within the SciDAC application institutes. The impact of this will be to evolve existing HPC simulation capabilities into predictive science capabilities grounded in verification and validation. This increase in simulation confidence for SciDAC physics simulation codes will then allow greater adoption of modeling/simulation approaches within the application domains they service.

2.2.3 Institute Awareness

Overcoming the challenges of equipping complex large-scale SciDAC-class applications with UQ capabilities will be facilitated by our partnering with other SciDAC Math and Computer Science Institutes, as well as the applications themselves. Here we discuss our strategy for partnering with other Institutes to achieve our goals (and assist them in achieving theirs). Strategies for collaborating with application groups follow in the next section.

Math Institutes: UQ requires fast and scalable forward solvers; as such, advances in fast solvers created by Math Institutes will naturally benefit UQ. Beyond these straightforward benefits, however, there are significant opportunities—and a critical need—to create new *UQ-aware solvers*. First, many UQ methods lead to solutions of instances of the same problem with source terms or operators varying over parameter/stochastic space. This provides incentive for developing techniques that reuse solver components, such as preconditioners and Krylov bases, over the varying parameter space. Depending on the parametric dependence, the construction of more effective preconditioners can be amortized over a large number of solves. Second, intrusive forward UQ algorithms lead to systems of infinite or finite dimensional equations with new structure not encountered in the usual forward solver ecosystem. Third, intrusive inverse UQ methods make use of derivatives of response surfaces or posterior densities to adapt to the local surface structure; these methods require fast algorithms for approximating Hessian-vector products. Finally, stochastic reduction methods pursued within QUEST can be coupled with adaptive meshing and time-stepping to produce computational representations that optimize the use of resources in pursuit of improved validity and accuracy; the stochastic refinement must be balanced against the deterministic spatio-temporal refinement. In all these areas, new discretization, adaptivity, solver, and preconditioner issues arise due to the specialized structure of the intrusive forward and inverse solvers, and the need to solve multiple instances of related problems. This provides fertile ground for collaboration with Math Institutes.

CS Institutes: The QUEST agenda has natural links to CS Institutes in several areas. First, exploiting emerging manycore node technologies is a critical technology for UQ, since much of the three orders of magnitude increase in performance in moving from petascale to exascale systems will be delivered by large increases in parallelism at the node level. Beyond benefiting from expected advances in manycore utilization of forward solvers, there are opportunities to develop new UQ algorithms that are manycore-aware. These include the four areas mentioned above. In particular, significant speedups can be realized on manycore systems using multi-vector evaluations for multiple-instance forward problems, as is the case with polynomial chaos representations and UQ solvers. The outer (typically less sparse) stochastic discretization can be exchanged with the inner spatial discretization to also realize speedups.

2.2.4 Application Awareness

Historically, the development of theory, methods, and algorithms for UQ has been driven by application needs—spectral methods to efficiently represent uncertainty from a high-dimensional distribution [68], adjoint methods to squeeze the most information out of each forward model run [55, 64], response surface approaches trained on an ensemble of forward model runs to emulate its response at untried settings [89, 192], developing probabilistic descriptions linking the computational model to physical reality to enable parameter estimation and model calibration [114, 128], and linking high- and low-fidelity computational models to better sample parametric space [76, 127]. The success of this effort depends on maintaining strong connections to applications, not just to guide method development and research, but to ensure usability of the software generated with this effort.

Fortunately, the assembled team has a strong record in building and maintaining strong collaborative ties to applications involving large-scale computational models. The national lab contingents on this proposal have a long history of interacting with applications in stockpile stewardship. More recently, their efforts have been channelled into new, highly collaborative open science applications. QUEST will provide us with the opportunity (and mandate) to engage new SciDAC applications.

Initially, this effort will leverage ongoing collaborations to ensure relevance to a variety of applications. Current applications include climate (UT, LANL, SNL), subsurface flow (LANL, UT, MIT), engineering (SNL, UT, USC, MIT, JHU), combustion chemistry (SNL, MIT, JHU), electric grids (MIT, SNL, USC, JHU), atmospheric monitoring (UT, SNL, LANL), and cosmology and astrophysics (LANL). Of key importance is the variety of issues spanned by these different applications. These include forward uncertainty propagation, sensitivity analysis, inverse or calibration problems, high volumes of simulated and physical data, availability of adjoints, the need to account for structural error or model inadequacy in inverse or calibration problems, planning experimental campaigns to improve prediction accuracy or other metrics, and extrapolation. Maintaining these strong collaborations, and developing new ones, will help ensure the relevance of QUEST methods and tools.

A key element of our vision for working with SciDAC application projects is to address common barriers to the adoption of UQ technologies. As described in §2.2.2, QUEST will invest in tools that enable widespread adoption of UQ technologies within SciDAC application projects. The QUEST team has many years of experience in providing algorithm services to application teams. For UQ, we envision three primary delivery models:

Black box: For rapid interfacing to existing simulation simulation codes, DAKOTA supports approaches that employ simulation drivers to spawn instances of distinct simulation executables, providing communication through the file system based on pre/post-processing tools applied to simulation input/output files and databases. These interfaces can be developed very quickly, allowing fast turn-around in time-critical UQ studies.

Semi-intrusive: Library embedding requires a moderate level of simulation intrusion to modify top-level program flow and to provide call backs for simulation execution. In return, the primary benefit is in streamlining problem setup by eliminating the burden of custom interface creation. In addition, this approach provides opportunities for gaining efficiency, by reducing file I/O, eliminating redundant simulation processing steps, and taking advantage of opportunities for warm/hot-starting of solvers, such as with Krylov basis recycling.

Intrusive: By leveraging existing solver foundations already broadly used by SciDAC application projects, and by working with other SciDAC Math/CS Institutes, the burden of intrusive methods can be significantly reduced.

This ensemble of techniques allows application projects to weigh the benefits of computational efficiency versus integration effort and to select the approach that is most appropriate for their computational environment.

Benchmark problems: We will also work with SciDAC application projects to understand their constraints, targets, and UQ needs. This will help us define well-characterized relevant benchmark problems that will be used within QUEST as demonstrations/test-cases for development and honing of UQ algorithms and software tools. We will include benchmark problems covering a spectrum of problem types, including PDE, ODE, and DAE problems, with a range of scale and complexity. The QUEST team is also keenly aware that proof-of-principle UQ demonstrations are, as stated in the SciDAC solicitation, "no substitute for realistic, full-scale applications or data sets." We therefore have milestones directed toward realistic applications or data sets, with the goal of increasing application awareness and lowering the adoption barrier for UQ technologies.

Furthermore, many of the UQ software tools in §2.2.2 have interfaces to different applications. QUEST is planning to collect and share the application interfaces across the partners. These activities will increase the application awareness of the individual UQ toolkits, and facilitate the ease with which UQ technologies can be applied to new applications.

2.2.5 Architecture Awareness

The inherent computational demands of UQ activities can be orders of magnitude larger than the demands of single-point forward simulation. Fortunately, the broad recognition of the need for equipping large-scale simulation codes with UQ capabilities has occurred simultaneously (and not coincidentally) with the impending arrival of multi-petaflops systems next year, and a roadmap for exascale systems in the out years of this decade. In particular, the SciDAC community will benefit from new multi-petaflops systems that go on line in 2012, including OLCF's *Titan* (a 20 PF accelerated Cray/Nvidia system), and ALCF's *Mira* (a 10 PF IBM BG/Q system). Bringing UQ capabilities to realistic science applications thus will capitalize on these and follow-on hundred-petaflops and ultimately exaflop systems, providing strong incentives to exploit hardware at the largest possible scales. At the same time, next generation systems bring new architectural features, including massive concurrency at multiple levels, and complex hierarchies of memory and inter-thread communication [129]. Taking full advantage of these architectures, and simultaneously weaving UQ methodology and software into the core of new simulation codes, requires intensely *architecture-aware* algorithmic developments. Our algorithms will address key aspects of extreme-scale computational architectures as follows. First, we note that UQ for large-scale simulations immediately yields *multiple levels of parallelism*. Large-scale PDE solvers sit on top of sparse numerical linear algebra and graph-based partitioning algorithms, which will achieve scalability through both massive multithreading at the processor level (exploiting manycore/accelerated chips) and message-passing at the internode level, up to hundreds of thousands of nodes. Intrusive UQ algorithms can and should exploit the very same solvers and partitioning approaches; moreover, as discussed §2.2.3, the structure of intrusive methods permits an additional opportunity to exploit manycore/accelerated processors at a fine-grained level. Surrounding all of this is a coarser-grained outer layer of "sample-based" parallelism encompassing both non-intrusive forward UQ methods and Monte Carlo methods in inversion.

Also crucial to performance on extreme-scale architectures are *fault monitoring, fault reporting, and fault recovery*. Below we will describe interfaces between our UQ codes and lower-level software tools for detecting faults, with means for seamless and asynchronous recovery. Mathematical approaches for fault tolerance (e.g., imputation of missing UQ samples), on the other hand, are detailed in §2.2.1.5. Further, we note that extreme-scale platforms will require sophisticated *load-balancing* approaches; internode latency and bandwidth limitations, coupled with myriad hardware and software sources of non-uniformity, call for *synchronization* at all levels to be *reduced* as much as possible. And finally, UQ algorithms must address *data locality* constraints imposed by memory and machine architecture.

Our current UQ software tools are designed with many of these architectural requirements in mind, and will provide an ideal platform from which to tackle the rest. A central feature of our current software tools is an emphasis on *high-performance parallel computing*. Notable examples include:

DAKOTA provides a scalable multilevel parallelism capability [82, 84] that utilizes MPI communicator partitioning and recursive scheduling to augment fine-grained parallel simulations with multiple levels of coarse-grained concurrency. This supports either black box simulation drivers that require "tiling" of parallel simulation jobs within total queue allocations (to avoid repeated queue delays) or library embedding strategies that require modularity in communication and file I/O. In addition, DAKOTA currently provides basic services for failure capturing and recovery, restart, and fault-tolerant sampling [8].

QUESO allows for multiple mathematical models to be assessed simultaneously, with each model assessment requiring simultaneous chains, and each chain requiring multiple MPI nodes for model simulations. QUESO synchronously calls the nodes in each chain, and provides a unified analysis of chains in each model assessment.

Moving these capabilities towards next generation computational platforms is a key goal of the proposed effort. Activities along these lines will include:

Data locality and enhanced scalability: DAKOTA's multilevel parallelism capabilities will be extended to include hybrid parallelism (MPI plus multithreading), expanded concurrency in algorithm recursions (e.g., for mixed aleatory-epistemic UQ or design under uncertainty), and distributed scheduling (avoiding single points of failure). Augmenting this parallel scalability is algorithmic scalability. We will address high dimensionality in random inputs through adaptive refinement and adjoint-enhancement (see §2.2.1.2) and address high dimensionality in output response metrics through out-of-core solutions involving standardized results databases. MCMC algorithms in QUESO will also have the ability to take advantage of architectural heterogeneity to achieve better load balancing among several chains.

Fault tolerance: We will continue to build on our capabilities for adaptive, fault-tolerant UQ algorithms, as described in §2.2.1.5. It is not sufficient for a UQ algorithm to feed simulation ensembles to a separate scheduling algorithm; rather, the algorithm and scheduler must become more integrated in order to support feedback of scheduling data and simulation faults that directly influence adaptive refinement logic within fully asynchronous algorithms. QUESO will take advantage of fault-tolerant versions of MPI, as they become available.

2.2.6 Project schedule, milestones, and deliverables

The main QUEST activities involve developing new and improved UQ algorithms and capabilities, implementing these algorithms into software, and making the tools available to address UQ needs in the SciDAC/SC community. Several activities will take place on an ongoing basis throughout the project duration and will be the shared responsibility of *all* QUEST member institutions. These activities are:

- Organize yearly workshops and tutorials to gather input from and disseminate tools to the community.
- Adapt approaches to the needs of the customers.
- Organize yearly internal QUEST workshops to synchronize and promote communication among groups.
- Disseminate results through minisymposia at relevant conferences and journal publications.
- Release software as open source; maintain documentation and user guides on the QUEST website.

The schedule below outlines further key tasks on a yearly basis. The acronyms following the tasks indicate the institutions involved, with the lead institution for each task indicated in bold font.

Year 1:

1.1 Develop and publish QUEST website as the central location for dissemination of publications, software, and workshop announcements, and for soliciting UQ community input (§2.3) [**SNL**, all]

1.2 Extend core stochastic Newton algorithms to large-scale setting via low-rank approximations of the likelihood Hessian; incorporate into QUESO library (§2.2.1.1). [**UT**, MIT]

1.3 Develop adapted stochastic representations with respect to arbitrary measures; compile library of error indicators for intrusive and non-intrusive forward UQ methods (§2.2.1.2). [**USC**, SNL, JHU]

1.4 Develop optimal sequential sampling strategies for non-intrusive forward UQ, tied to error measures in quantities of interest (§2.2.1.4). [**MIT**, LANL]

1.5 Develop algorithms and software for response surface methods, such as the Gaussian process, that will effectively take advantage of multiprocessor HPC environments (§2.2.2.1). [**LANL**, UT, USC]

1.6 Improve interoperability between QUEST software to provide a unified set of best-in-class tools for UQ, inference, and design (§2.2.2.1). [**SNL**, all]

1.7 Demonstrate improved dimension scalability in forward UQ based on adaptive refinement and adjoint-enhancement of non-intrusive polynomial chaos and stochastic collocation expansions (§2.2.1.2). [**SNL**, MIT, USC, JHU]

Year 2:

2.1 Develop multi-chain variants of stochastic Newton algorithms to support coarse-granularity parallelism (§2.2.1.1) [**UT**, LANL]

2.2 Develop high-performance algorithms for expansions on symmetrized bases (§2.2.1.2). [**USC**, JHU]

2.3 Develop non-intrusive algorithms to detect and represent discontinuous model responses with surrogates, using basis enrichment and compressive sensing (§2.2.1.3). [**MIT**, SNL, USC, JHU]

2.4 Demonstrate advanced Bayesian inference in DAKOTA/QUESO/GPMSA employing adaptively refined stochastic expansions as emulator models (§2.2.2.1). [**SNL**, UT, LANL]

2.5 Develop software for carrying out global sensitivity analyses and calibration using response surface models in HPC environment (§2.2.2.1). [**LANL**, SNL]

2.6 Demonstrate UQ-enabled solver approaches for non-intrusive forward UQ methods based on solver information reuse (§2.2.3) [**SNL**, USC]

Year 3:

3.1 Extend stochastic Newton algorithms to more effectively sample multi-modal distributions via Hessian-aware Gaussian process proposals (§2.2.1.1) [**UT**, MIT]

3.2 Guide deployment of QUEST's full range of surrogate-construction methodologies (intrusive and non-intrusive) on large-scale statistical inference and inverse problems (§2.2.1.1). [**MIT**, SNL, LANL]

3.3 Develop adaptive basis representations for multiphysics applications (§2.2.1.2, §2.2.1.3). [**USC**, SNL, JHU]

3.4 Demonstrate effectiveness of fault-tolerant UQ methods for dealing with missing samples and domain adjustments (§2.2.1.5). [**SNL**]

3.5 Integrate new HPC-aware response surface software into GPMSA, and integrate this new version into DAKOTA (§2.2.2.1). [**LANL**, SNL]
3.6 Refine specification of benchmark problems through interactions with applications (§2.2.4). [**SNL**, all]
3.7 Improve parallel scalability of all algorithms (through hybrid MPI+threading, concurrency in iterator recursions, and distributed scheduling) (§2.2.5). [**SNL**, all]

Year 4:

4.1 Apply stochastic Newton algorithms to large-scale stochastic inverse problems in subsurface flow (§2.2.1.1) [**UT**]
4.2 Develop stochastic surrogates, described using PCEs with random coefficients (§2.2.1.1). [**USC**, MIT]
4.3 Develop dimensionality reduction algorithms in large-scale statistical inverse problems (§2.2.1.1), leveraging reduced-order stochastic representations for forward UQ (§2.2.1.2). [**MIT**, SNL, UT, LANL, USC]
4.4 Develop generic methods and software for specifiying error models to account for model discrepancy (§2.2.2.1). [**LANL**, SNL, USC]
4.5 Demonstrate effective resource utilization through improved parallel scalability, fault tolerance, and architecture awareness for comprehensive, end-to-end UQ analyses of large scale benchmark problems on high end DOE/SC computational platforms (§2.2.4, §2.2.5). [**SNL**, all]

Year 5:

5.1 Scale up and robustify stochastic Newton algorithms to exploit massive concurrencies provided by leading edge systems (§2.2.1.1) [**UT**]
5.2 Implement algorithms for data-driven basis enrichment and adaptation (§2.2.1.1 – §2.2.1.3). [**USC**, JHU]
5.3 Develop large-scale optimal experimental design algorithms to help domain scientists choose the most informative experimental observations for parameter inference, prediction, and model validation (§2.2.1.7). [**MIT**, UT, LANL]
5.4 Integrate model discrepancy software into GPMSA (§2.2.2.1). [**LANL**, SNL]
5.5 Based on benchmark problems and community feedback, evaluate our success measured by QUEST's effectiveness in serving SciDAC customers, and identify areas for continued improvement (§2.3). [**SNL**, all]

2.3 Management Plan

2.3.1 Working with other SciDAC Projects

QUEST will initially be driven by an internal vision of the role of UQ in the scientific process generally and computational science specifically. QUEST management and software architectures will be designed to allow this vision to adapt to evolving science, hardware, and opportunities.

While the core of QUEST technology will be self-justified and self-contained in the form of software libraries and algorithms, its ultimate and fully developed impact will be fulfilled through interaction with complementary technologies. We envision the following four capabilities, fully developed on flagship hardware components, as critical for the complete deployment of QUEST: (1) tools for solvers and optimization; (2) tools for massive and massively distributed data sets and data-driven simulations; (3) tools for scientific-UQ and decision-oriented visualization; (4) tools for adaptive discretization, including meshing and multiscale/multimodel simulations.

We expect that some of these capabilities will form the core competencies for a number of SciDAC institutes. We will work closely with these institutes to ensure interoperability and consistency of software architecture and data structures. While we have initiated communication with several SciDAC proposing teams, we adopt a robust strategy that permits us to rely on state-of-the-art capabilities in any of the four topical areas above, whether or not these are supported through SciDAC institutes. Specifically, we will rely on the latest stable releases of off-the shelf software tools for capabilities not supported through SciDAC, and on the most current releases of all SciDAC software tools.

In addition to working with SciDAC Institutes to maximize the value of our deliverables, we will also work closely with the scientific community at large to facilitate access to and use of our resources. We will specifically target potential SciDAC science applicants and make sure they are informed of the scientific significance and computational requirements of our tools. We expect to work closely with funded science applications both to provide them with usable UQ capabilities and to solicit their feedback towards enhancing QUEST algorithmic and computational infrastructure.

We will achieve our objectives for maximum outreach and impact, as indicated above, through the following organizational structure and activities.

Tutorials: Tutorials and workshops will be organized, targeted at SciDAC/SC researchers and aimed at demonstrating the relevance of our tools for applications as well as the software architecture and algorithmic elements of QUEST. The objective of these tutorials will be both to educate colleagues about QUEST and to synchronize QUEST software architecture and algorithms for interoperability with other Institutes. Administrative support for outreach has been budgeted for. A committee consisting of Habib Najm, Omar Knio, and Omar Ghattas will direct these activities.

Short courses: Scientists from QUEST will offer short courses at SciDAC, SIAM, USACM, and other conferences. These activities will ensure that the SC community at large remains informed of the availability of QUEST.

Summer school: A UQ summer school will be organized for one week each summer, providing a forum for in-depth presentations and dissemination of latest innovations in UQ research. The summer school will be designed to be financially self-sufficient.

Web presence: A web presence will be designed for a unified, simple, and user-targeted interface to the outside world. Our budget includes support for this activity. The website will be informative on the capabilities available, without being a confusing enumeration of tools. It will provide means for the SciDAC/SC scientist to understand the basic highlights of UQ and how it can advance his/her scientific understanding. It will also provide an explanation of how we can help—directing each user, through an interactive interface, to the particular QUEST tool he/she needs to use and the contact person in QUEST who can help with it. Mike Eldred will serve as the single point of contact on the QUEST team to help in directing users, if this web interaction was not sufficient.

Partnering with other SciDAC Institutes: We expect that developments within QUEST will have tight coupling with developments at other Math, CS, Data, and Visualization Institutes. Specifically, our interest in developing error indicators and intrusive UQ will capitalize heavily on existing solver technology. Furthermore, given the various layers of parallelism that are intrinsic to UQ problems, optimality of our algorithms will be greatly influenced by manycore technology. We will identify one person from the QUEST team to be our point of contact with each of the funded SciDAC Institutes. We will request development access to their repositories and relevant subsets of QUEST will hold bi-weekly meetings to review changes to these Institutes' software architectures and data models. Decisions about propagating these changes into QUEST will be made on an ongoing basis, following discussions with our contact points at the other Institutes.

Success metrics: We will collect multiple types of information to assess our performance. Our main measures of success, as already stated, are the degree to which we *serve the needs* of the SciDAC/SC user community and our *impact* on the scientific community as a whole. The latter will be measured in terms of our publication output, invited talks, organization of technical meetings, and similar professional engagements. Success in pursuing the former—and more important—goal of serving SciDAC/SC customers will be measured in several objective and subjective manners. We will collect usage statistics on our tools, websites, and tutorial participation; measure the improvement in the performance of our tools on benchmark problems; poll our user base with surveys on the utility, accessibility, and performance improvements of our software tools; and continually modify our practices in response to this feedback and the user community's evolving needs.

Communication: The QUEST team has a diverse domain expertise, and as such is well-positioned to inform a very broad scientific community of the enabling role of QUEST in scientific discovery and predictive science. Through our network of academic and lab partners, and through our outreach program detailed above, we will have a good assessment of the competencies available to respond to various opportunities relevant to SciDAC. We will make sure, through a variety of mechanisms and channels, that QUEST capabilities and resources are sufficiently communicated. We will provide personal interactions with SciDAC science applications, providing respective teams with expert assessment of a meaningful implementation of UQ into their scientific effort.

We will provide assistance to all SciDAC science applications in defining well-posed UQ formulations relevant to their needs. These include proper choice of quantities of interest, selection of validation metrics, identification of stochastic inversion, forward UQ procedures, and analysis.

2.3.2 Internal management structure and plan

In the following we describe our plan for managing the QUEST efforts across the several partner institutions, the roles and responsibilities of each of the senior investigators, mechanisms we use to facilitate collaboration and manage software development, and our plans to mitigate risks to the success of the QUEST institute.

2.3.2.1 Management Structure and Decision-Making

QUEST brings together a world-leading team of researchers and developers of UQ algorithms and software. This expertise is spread across four Universities and two DOE Laboratories. To meet the ambitious goals of the QUEST project, the team's efforts across the six institutions will need to be coordinated, and the required collaboration will need to be fosterd. This will be accomplished by a two-level management structure as described below.

The overall activites of QUEST will be coordinated by a small management team consisting of Habib Najm as Director, Mike Eldred as lead for software development and Roger Ghanem as lead for algorithm development. This management team will be expanded to include other major activities (e.g. for user support or applications), as need arises. Each of the activity leads will in turn be responsible for coordination within their domains. The members of the management team, assisted by an administrative person at SNL funded by QUEST, will monitor progress on projects, identify efforts that need more resources, those that may need less or can be dropped and will evaluate new opportunities. Collectively they will prioritize efforts, develop plans and milestones and propose adjustments in responsibilities of QUEST project staff as necessary. The management team will confer by teleconference regularly (approximately every other week) to execute this coordination function.

One of the challenges of a distributed effort like QUEST is to enable sufficient flexibility in the execution of the various projects when the distribution of resources among the partnering institutions is not very flexible. This will be accomplished through flexible collaboration among the partner institutions. When priorities and responsibilities need to be adjusted, the management team will propose adjustments to a larger group of core investigators, consisting of PI's, co-PI's and other senior personnel at the partner institutions. These core investigators are responsible for managing the efforts at their institutions. When adjustments are proposed by the management team, they will consult with the core investigators to develop an execution plan. Required changes in staff responsibilities will often be accomplished by shifting efforts within one of the partner institutions. But when this is not possible, adjustment in responsibilities across partners will be arranged by the core investigators, essentially bringing in collaborators to an effort from the other institutions when needed. The core investigators will also confer by teleconference regularly (approximately every other month) to discuss progress and roadblocks, to encourage cross-fertilization, and to execute priority and resource adjustments.

2.3.2.2 Team Roles & Responsibilities

Following is a list of the key team members, with a brief description of their roles.

Habib Najm: PI and Director; responsible for overall coordination of the project. Expertise in uncertainty quantification and statistical inference methods development, and their application in ODE and PDE problems.

Bert Debusschere: Lead UQTk developer; responsible for algorithm and software development. Expertise in intrusive/non-intrusive UQ, statistical inference, stochastic dynamical systems, and stochastic multiscale modeling.

Michael Eldred: Lead DAKOTA and Pecos developer; responsible for software inter-operability and coordination within QUEST; expertise in adaptive stochastic methods and design under uncertainty.

Roger Ghanem: Responsible for various aspects of intrusive UQ methods and probabilistic modeling.

Omar Knio: Responsible for development of UQ methods/algorithms for sparse representation of extreme-scale data. Expertise in development of intrusive/non-intrusive/adaptive UQ methods, and UQ software.

Omar Ghattas: Responsible for development of parallel Hessian-based stochastic Newton-MCMC methods for inverse problems. Expertise in large-scale inverse problems, parallel computing, and geophysical applications.

Robert Moser: Responsible for development and testing of algorithms and process for validation under uncertainty and treatments of uncertainty arising from model inadequacy.

Ernesto Prudencio: Lead architect of the QUESO library, responsible for the research of parallel adaptive UQ algorithms for the Bayesian analysis of mathematical models, and for the integration of algorithms into QUESO.

David Higdon: Responsible for developing response-based UQ methods within GPMSA and developing new MCMC algorithms and developing collaborations with applications.

James Gattiker: Lead GPMSA developer; responsible for integration of new methods into GPMSA, adaptation of GPMSA software to the HPC environment, and integration of GPMSA with other software.

Youssef Marzouk: Responsible for methods to characterize uncertainty in model inputs, and for adaptive sampling strategies. Expertise in statistical inference, experimental design, and surrogate/reduced-order modeling.

2.3.2.3 Collaboration, Coordination & Software Management

Another challenge faced by a large collaborative effort such as QUEST is enabling day-to-day collaboration among participants across the institutions. A critical tool for this purpose will be a central QUEST-wide online collaboration space, using a software tool such as Redmine. Such systems provide an integrated suite of collaboration services, such as project wikis, documentation and publication repositories, and chat channels. Redmine and similar tools integrate with svn software and document repositories, provide project management functions such as milestone and ticket tracking, and integrate with continuous regression testing systems. Redmine is currently in use for this purpose at the PECOS Center at the University of Texas.

In addition, collaboration and coordination will be fostered through regular teleconference presentations of progress and activities by team members to the entire QUEST team. This will allow QUEST team members to share ideas and provide feedback to their colleagues. Collaboration will also be facilitated through regular visits of QUEST researchers among the various partner institutions, and an annual QUEST team workshop in which the QUEST team will gather to assess progress, coordinate efforts, and plan future developments. Each partner has budgeted travel funds to support such efforts.

2.3.2.4 Risk Management

The risks to the success of QUEST, and our strategy to mitigate them, are outlined briefly below.

UQ research risks: UQ for large-scale computation is a relatively young field, where the development of algorithms/formulations is not mature. Research and development in this area thus carries risks that some proposed approaches will not succeed. This risk will be mitigated by active management of research and development activities, to identify unpromising lines of research early and shift efforts to pursue more promising ideas. We will rely both on our expertise and on active collaboration with SciDAC Math Institutes in this regard. The track record of the team, both in innovative research and teaming activities, provides confidence in our ability to respond successfully to these anticipated challenges.

UQ software development risks: The development of sophisticated software tools carries risks that development efforts may fail by producing software that does not meet its functional or interface requirements, cannot be verified to be correct or performs too poorly to be useful. These risks will be minimized through rigorous software engineering practices, and will be mitigated by active management to identify and correct problems early in the development process. The team's extensive experience in UQ software development/maintenance, and active support of user communities, will be key in our successful software development efforts.

Hardware architecture risks: Over the next several years, disruptive changes in architecture of high performance computing hardware are expected. Such changes pose a risk to the QUEST effort due to the potential to render the developed algorithms and/or software obsolete. These risks will be mitigated by tracking expected architectural changes, and proactively adapting development efforts to the relevant characteristics of emerging new hardware. We will do this in close collaboration with other SciDAC Math/CS Institutes that have extensive expertise, and their own research goals, in responding to such architectural changes.

User adoption risks: UQ is effectively new to many SciDAC/SC computational scientists in almost all application areas. Accordingly, even though DOE/ASCR program workshops have consistently highlighted the need for UQ, there is a risk of user apathy, in that it is easiest to stick with business-as-usual. This would translate to a failure of our UQ efforts, as our key goal is providing tools for UQ to be *used* by SciDAC/SC applications researchers. We will mitigate this risk by active outreach to the SciDAC/SC community, and by providing effective and responsive mechanisms for addressing user requests, as outlined elsewhere. We will also benefit in this regard from our close collaborations with other SciDAC Institutes that have an extensive record in this context, with an established SciDAC/SC user community.

2.4 Impact

QUEST will have significant impact on SciDAC and science in general.

Impact on SciDAC Goals: QUEST will serve as a trustworthy resource that will provide UQ capability to SciDAC Institutes and applications. This capability will be in the form of software resources as well as scientific interactions and domain-specific expertise. QUEST will enable SciDAC grantees to leverage SciDAC investment while endowing their resources with a UQ dimension. SciDAC applications will be able to explore parameter space, model space, and decision space effectively, in an effort to better understand the extent to which their underlying models are representative of nature, together with guidelines on improving the model.

Impact on Science in General: QUEST will provide UQ capabilities that are comensurate with state-of-the-art computational resources. These will permit the scientific community at large to take advantage of developments in UQ methodology as they attempt to certify the value of their predictions. QUEST resources will make it easier for the scientific community to report and exchange information, in a consistent fashion, about uncertainties in data, models, and simulations. This is a critical step towards bringing statistical thinking back into the mainstream computational science enterprise.

2.5 Institutions and their Roles

QUEST is a collaboration among six institutions that have had multiple ongoing collaborations on the development of UQ methods, and UQ software deployment, in the context of large-scale computations. Their collaboration on QUEST brings together a group of researchers with a shared vision, and a range of strengths and qualifications.

The roles of the individual institutions on the QUEST collaboration are designed around the leading strengths and expertise of each instituion, while always relying on synergistic coupling of the range of strengths/expertise among the overall group. The institutions and their roles are summarized below.

Sandia National Laboratories (SNL): Sandia will lead, manage, and coordinate the project progress. It will also be responsible for development of non-intrusive forward UQ algorithms. Sandia will work on developing and serving several UQ software products, namely DAKOTA, UQTk, JAGUAR, Pecos, and Percept.

Los Alamos National Laboratory (LANL): LANL provides broad expertise in statistics, and in the application of statistical UQ methods in large scale computational problems (ocean modeling, cosmology, material science). It will be responsible for development of algorithms/software for construction and use of statistical GP surrogates. LANL will develop and serve the GPMSA software toolkit.

University of Southern California (USC): USC provides expertise in development of UQ methods for general PDE problems. It will be responsible for development of methods and software for construction of adaptive, optimized, UQ representations and intrusive UQ methods.

University of Texas at Austin (UT): UT provides expertise in UQ for large-scale inverse problems, statistical inference, and model validation. It will be responsible for development of UQ methods for inverse problems and for model assessment under uncertainty. UT will work on developing and serving the QUESO software toolkit.

Johns Hopkins University (JHU): JHU provides expertise in development of intrusive/non-intrusive UQ methods and software. It will be responsible for development/deployment of methods/software for sparse UQ representations.

Massachusetts Institute of Technology (MIT): MIT provides expertise in statistical inference, surrogate and reduced-order modeling, and stochastic optimization. It will be responsible for the development of adaptive sampling methods and software to be used in forward UQ and statistical inverse problems.

3 References

[1] The Opportunities and Challenges of Exascale Computing. The ASCAC Subcommittee on Exascale Computing , Department of Energy, November 2010. http://www.sc.doe.gov/ascr/ASCAC/Reports/Exascale-Subcommittee-Report.pdf. 4

[2] The DAKOTA Project: Large-Scale Engineering Optimization and Uncertainty Analysis. http://dakota.sandia.gov/. 13

[3] Chapter 15: Inputs to DAKOTA. http://dakota.sandia.gov/licensing/release/Users5.1.pdf. 14

[4] Mathematical Research Challenges in Optimization of Complex Systems. Department of Energy, December 2006. http://www.er.doe.gov/ascr/Research/AM/Docs/ComplexSystemsWorkshopReport.pdf. 4

[5] PECOS Developers Manual. http://dakota.sandia.gov/packages/pecos/html/index.html. 14

[6] Sierra Toolkit. http://trilinos.sandia.gov/packages/stk/. 14

[7] Uncertainty Quantification Toolkit. http://www.sandia.gov/UQToolkit/. 13

[8] B. Adams, W. Bohnhoff, K. Dalbey, J. Eddy, M. Eldred, D. Gay, K. Haskell, P. Hough, S. Lefantzi, and L. Swiler. DAKOTA, a multilevel parallel object-oriented framework for design optimization, parameter estimation, uncertainty quantification, and sensitivity analysis: Version 5.1 users manual. Technical Report SAND2010-2183, updated January 2011, Sandia National Laboratories, Albuquerque, NM, January 2011. 11, 13, 17

[9] B. M. Adams, B. J. Bichon, M. S. Eldred, B. Carnes, K. D. Copps, D. C. Neckels, M. M. Hopkins, P. K. Notz, S. R. Subia, and J. W. Wittwer. Solution-verified reliability analysis and design of bistable mems using error estimation and adaptivity. Technical Report SAND2006-6286, Sandia National Laboratories, Albuquerque, NM, October 2006. 12

[10] N. Agarwal and N. Aluru. A domain adaptive stochastic collocation approach for analysis of MEMS under uncertainties. *J. Comput. Phys.*, 228:7662–7688, 2009. 9

[11] V. Akçelik, G. Biros, and O. Ghattas. Parallel multiscale Gauss-Newton-Krylov methods for inverse wave propagation. In *Proceedings of IEEE/ACM SC2002 Conference*, Baltimore, MD, Nov. 2002. SC2002 Best Technical Paper Award. 5, 6, 8

[12] V. Akçelik, J. Bielak, G. Biros, I. Epanomeritakis, A. Fernandez, O. Ghattas, E. J. Kim, J. Lopez, D. R. O'Hallaron, T. Tu, and J. Urbanic. High resolution forward and inverse earthquake modeling on terascale computers. In *SC03: Proceedings of the International Conference for High Performance Computing, Networking, Storage, and Analysis*. ACM/IEEE, 2003. Gordon Bell Prize for Special Achievement. 6

[13] V. Akçelik, G. Biros, O. Ghattas, K. R. Long, and B. van Bloemen Waanders. A variational finite element method for source inversion for convective-diffusive transport. *Finite Elements in Analysis and Design*, 39 (8):683–705, 2003.

[14] V. Akçelik, G. Biros, A. Draganescu, O. Ghattas, J. Hill, and B. van Bloeman Waanders. Dynamic data-driven inversion for terascale simulations: Real-time identification of airborne contaminants. In *Proceedings of SC2005*, Seattle, 2005. 6

[15] V. Akçelik, G. Biros, O. Ghattas, D. Keyes, K. Ko, L.-Q. Lee, and E. Ng. Adjoint methods for electromagnetic shape optimization of the low-loss cavity for the International Linear Collider. *Journal of Physics: Conference Series*, 16:435–445, 2005. 6

[16] V. Akçelik, G. Biros, O. Ghattas, J. Hill, D. Keyes, and B. van Bloeman Waanders. Parallel PDE-constrained optimization. In M. Heroux, P. Raghaven, and H. Simon, editors, *Parallel Processing for Scientific Computing*. SIAM, 2006. 5, 6, 8

[17] A. Alexanderian, O. L. Maître, H. Najm, M. Iskandarani, and O. Knio. Multiscale stochastic preconditioners in non-intrusive spectral projection. *J. Sci. Comput.*, in press, 2010. 5, 10

[18] L. Alisic, M. Gurnis, G. Stadler, C. Burstedde, L. C. Wilcox, and O. Ghattas. Slab stress and strain rate as constraints on global mantle flow. *Geophysical Research Letters*, 37:L22308, 2010. doi: 10.1029/2010GL045312. 6

[19] J. F. Antaki, G. E. Blelloch, O. Ghattas, I. Malevic, G. L. Miller, and N. J. Walkington. A parallel dynamic mesh Lagrangian method for simulation of flows with dynamic interfaces. In *Proceedings of Supercomputing 2000, Dallas, November 4–10*, 2000. 6

[20] R. Archibald, A. Gelb, R. Saxena, and D. Xiu. Discontinuity detection in multivariate space for stochastic simulations. *J. Comp. Phys.*, 228:2676–2689, 2009. 10

[21] M. Arnst and R. Ghanem. Probabilistic equivalence and stochastic model reduction in multiscale analysis. *Computer Methods in Applied Mechanics and Engineering*, 197(43-44):3584–3592, 2008. 5, 9, 10

[22] M. Arnst and R. Ghanem. Probabilistic electromechanical modeling of nanostructures with random geometry. *Journal of Computational and Theoretical Nanoscience*, in press, 2009. 5, 10

[23] A. Askan, V. Akçelik, J. Bielak, and O. Ghattas. Full waveform inversion for seismic velocity and anelastic losses in heterogeneous structures. *Bulletin of Seismological Society of America*, 97(6):1990–2008, 2007. 5, 6, 8

[24] R. A. Bartlett, D. M. Gay, and E. T. Phipps. Automatic differentiation of C++ codes for large-scale scientific computing. In V. N. Alexandrov, G. D. van Albada, P. M. A. Sloot, and J. Dongarra, editors, *Computational Science – ICCS 2006*, volume 3994 of *Lecture Notes in Computer Science*, pages 525–532, Heidelberg, 2006. Springer. ISBN 3-540-34385-7. 6

[25] O. Bashir, K. Willcox, O. Ghattas, B. van Bloemen Waanders, and J. Hill. Hessian-based model reduction for large-scale systems with initial condition inputs. *International Journal for Numerical Methods in Engineering*, 73:844–868, 2008. 5, 6, 8

[26] M. Bayarri, J. Berger, J. Cafeo, G. Garcia-Donato, F. Liu, J. Sacks, and D. Walsh. Computer model validation with functional output. *Annals of Statistics*, 35(5):1874, 2007. 8

[27] M. J. Bayarri, J. O. Berger, R. Paulo, J. Sacks, J. A. Cafeo, J. Cavendish, C. Lin, and J. Tu. A framework for validation of computer models. *Technometrics*, 49:138–154, 2007. 8

[28] J. Bernardo, M. Bayarri, J. Berger, A. Dawid, A. Heckerman, D. a nd Smith, and M. West. Markov chain Monte Carlo-based approaches for inference in computationally intensive inverse problems. In *Bayesian Statistics 7: Proceedings of the Seventh Valencia International M eeting*, page 181. Citeseer, 2003. ISBN 0198526156. 6

[29] R. Berry, H. Najm, B. Debusschere, H. Adalsteinsson, and Y. Marzouk. Data-free inference of the joint distribution of uncertain model parameters. *Journal of Computational Physics*, 2011. submitted. 9

[30] B. J. Bichon, S. Mahadevan, and M. S. Eldred. Reliability-based design optimization using efficient global reliability analysis. In *Proc. 11th AIAA Non-Deterministic Approaches Conference*, number AIAA-2009-2261, Palm Springs, CA, 2009. 5, 13

[31] L. Biegler, G. Biros, O. Ghattas, Y. Marzouk, M. Heinkenschloss, D. Keyes, B. Mallick, , L. Tenorio, B. van Bloemen Waanders, and K. Willcox, editors. *Large-scale Inverse Problems and Quantification of Uncertainty*. Wiley, 2011. 5

[32] L. T. Biegler, O. Ghattas, M. Heinkenschloss, and B. van Bloemen Waanders, editors. *Large-Scale PDE-Constrained Optimization*. Lecture Notes in Computational Science and Engineering, Vol. 30. Springer-Verlag, Heidelberg, 2003.

[33] L. T. Biegler, O. Ghattas, M. Heinkenschloss, D. Keyes, and B. van Bloemen Waanders, editors. *Real-Time PDE-Constrained Optimization*. SIAM, 2007. 5

[34] J. Bielak, O. Ghattas, and E. Kim. Parallel octree-based finite element method for large-scale earthquake ground motion simulations. *Computer Modeling in Engineering and Sciences*, 10:99–112, 2005. 6

[35] G. Biros and O. Ghattas. Inexactness issues in the Lagrange-Newton-Krylov-Schur method for PDE-constrained optimization. In L. T. Biegler, O. Ghattas, M. Heinkenschloss, and B. van Bloemen Waanders, editors, *Large-Scale PDE-Constrained Optimization*, volume 30 of *Lecture Notes in Computational Science and Engineering*, pages 93–114. Springer, 2003. 5, 8

[36] G. Biros and O. Ghattas. Parallel Lagrange–Newton–Krylov–Schur methods for PDE–constrained optimization. Part II: The Lagrange–Newton solver and its application to optimal control of steady viscous flows. *SIAM Journal on Scientific Computing*, 27(2):714–739, 2005. 6

[37] G. Biros and O. Ghattas. Parallel Lagrange–Newton–Krylov–Schur methods for PDE–constrained optimization. Part I: The Krylov–Schur solver. *SIAM Journal on Scientific Computing*, 27(2):687–713, 2005.

[38] G. Biros and O. Ghattas. Parallel Newton-Krylov algorithms for PDE–constrained optimization. In *Proceedings of SC99*, Portland, Oregon, 1999. 5, 6, 8

[39] G. Bothun, S. Hammond, and S. Picataggio. Computational research needs in alternative and renewable energy. Workshop Report, Department of Energy, September 2007. `http://www.er.doe.gov/ascr/ProgramDocuments/Docs/CRNAREWorkshopReport.pdf`. 4

[40] D. Brown, P. Messina, P. Beckman, D. Keyes, and J. Vetter. Crosscutting Technologies for Computing at the Exa scale. Workshop Report, Department of Energy, January 2009. `http://extremecomputing.labworks.org/crosscut/CrosscutWSFinalReptDraft02.pdf`. 4

[41] D. Brown et al. Applied Mathematics at the U.S. Department of Energy: Past, Present and a view to the Future. Department of Energy, May 2008. `http://www.sc.doe.gov/ascr/ProgramDocuments/Docs/Brown_Report_May_08.pdf`. 4

[42] A. M. Bruckstein, D. L. Donoho, and M. Elad. From sparse solutions of systems of equations to sparse modeling of signals and images. *SIAM Review*, 51(1):34–81, 2009. 10

[43] T. Bui-Thanh, K. Willcox, O. Ghattas, and B. Bloemen Waanders. Goal-oriented model-constrained optimization for reduction of large-scale systems. *Journal of Computational Physics*, 224:880–896, 2007. 5, 8

[44] T. Bui-Thanh, K. Willcox, and O. Ghattas. Parametric reduced-order models for probabilistic analysis of unsteady aerodynamic applications. *AIAA Journal*, 46:2520–2529, 2008.

[45] T. Bui-Thanh, K. Willcox, and O. Ghattas. Model reduction for large-scale systems with high-dimensional parametric input space. *SIAM Journal on Scientific Computing*, 30:3270–3288, 2008. 5, 8

[46] T. Bui-Thanh, L. C. Wilcox, C. Burstedde, and O. Ghattas. A discontinuous galerkin method for electromagnetic wave propagation on h-non-conforming adapted meshes. *in preparation*, 2011. 6

[47] C. Burstedde and O. Ghattas. Algorithmic strategies for full waveform inversion: 1D experiments. *Geophysics*, 74(6):WCC37–WCC46, 2009. doi: 10.1190/1.3237116. 5, 6, 8

[48] C. Burstedde, O. Ghattas, M. Gurnis, E. Tan, T. Tu, G. Stadler, L. C. Wilcox, and S. Zhong. Scalable adaptive mantle convection simulation on petascale supercomputers. In *SC08: Proceedings of the International Conference for High Performance Computing, Networking, Storage, and Analysis*. ACM/IEEE, 2008. ISBN 978-1-4244-2835-9. Gordon Bell Prize finalist. 6

[49] C. Burstedde, O. Ghattas, G. Stadler, T. Tu, and L. C. Wilcox. Towards adaptive mesh PDE simulations on petascale computers. In *Proceedings of Teragrid '08*, 2008. *Winner, NSF TeraGrid Capability Computing Challenge*.

[50] C. Burstedde, M. Burtscher, O. Ghattas, G. Stadler, T. Tu, and L. C. Wilcox. ALPS: A framework for parallel adaptive PDE solution. *Journal of Physics: Conference Series*, 180:012009, 2009. doi: 10.1088/1742-6596/180/1/012009.

[51] C. Burstedde, O. Ghattas, G. Stadler, T. Tu, and L. C. Wilcox. Parallel scalable adjoint-based adaptive solution for variable-viscosity Stokes flows. *Computer Methods in Applied Mechanics and Engineering*, 198:1691–1700, 2009.

[52] C. Burstedde, O. Ghattas, M. Gurnis, T. Isaac, G. Stadler, T. Warburton, and L. C. Wilcox. Extreme-scale AMR. In *SC10: Proceedings of the International Conference for High Performance Computing, Networking, Storage, and Analysis*. ACM/IEEE, 2010. Gordon Bell Prize finalist.

[53] C. Burstedde, L. C. Wilcox, and O. Ghattas. `p4est`: Scalable algorithms for parallel adaptive mesh refinement on forests of octrees. Accepted for publication in *SIAM Journal on Scientific Computing*, 2011. 6

[54] D. Cacuci, editor. *Sensitivity and Uncertainty Analysis Theory*, volume 1. Chapman & Hall/CRC, Boca Raton, FL, 2003. 9

[55] D. Cacuci. The forward and adjoint methods of sensitivity analysis. *Uncertainty Analysis*, 282, 1988. 15

[56] M. Calaf, C. Meneveau, and J. Meyers. Large eddy simulation study of fully developed wind-turbine array boundary layers. *Phys. Fluids*, 22:015110, 2010. 6

[57] F. Campolongo, A. Saltelli, T. Sørensen, and S. Tarantola. Hitchhiker's guide to sensitivity analysis. In A. Saltelli, . K. Chan, and E. Scott, editors, *Sensitivity Analysis*. Wiley, Chicester, 2000. 9

[58] E. J. Candès, J. Romberg, and T. Tao. Robust uncertainty principles: Exact signal reconstruction from highly incomplete frequency information. *IEEE Transactions on Information Theory*, 52(2):489–509, 2006. 10

[59] P. Challenor, D. McNeall, and J. Gattiker. Assessing the probability of rare climate events. *The Oxford handbook of applied Bayesian analysis*, page 403, 2010. 6

[60] S. H. Cheung and J. L. Beck. New Bayesian updating methodology for model validation and robust predictions of a target system based on hierarchical subsystem tests. *Computer Methods in Applied Mechanics and Engineering*, 2010, accepted. 8

[61] S. H. Cheung and E. E. Prudencio. Adaptive multi-level stochastic simulation algorithm for Bayesian model class updating, assessment, averaging and validation. *in preparation*, 2011. 5, 8

[62] S. H. Cheung, T. A. Oliver, E. E. Prudencio, S. Prudhomme, and R. D. Moser. Bayesian uncertainty analysis with applications to turbulence modeling. *Reliability Engineering and System Safety (RESS) Journal, Special Issue on Quantification of Margins and Uncertainties (QMU)*, accepted, 2010. 12

[63] H. Chipman, E. George, and R. McCulloch. BART: Bayesian additive regression trees. *The Annals of Applied Statistics*, 4(1):266–298, 2010. ISSN 1932-6157. 8

[64] S. Collis, R. Bartlett, T. Smith, M. Heinkenschloss, L. Wilcox, J. Hill, O. Ghattas, M. Berggren, V. Akcelik, C. Ober, et al. Sensitivity technologies for large scale simulation. Technical report, Sandia National Laboratories, 2005. 15

[65] S. Das. *Model, Identification & Analysis of Complex Stochastic Systems: Applications in Stochastic Partial Differential Equations and Multiscale Mechanics*. PhD thesis, University of Southern California, 2008. 9

[66] S. Das, R. Ghanem, and J. Spall. Asymptotic sampling distribution for polynomial chaos representation from data: a maximum entropy and Fisher information matrix approach. *SIAM J. Sci. Comp.*, 30(3):2207–2234, 2008. 5, 9

[67] B. Debusschere, H. Najm, A. Matta, O. Knio, R. Ghanem, and O. Le Maître. Protein labeling reactions in electrochemical microchannel flow: Numerical simulation and uncertainty propagation. *Physics of Fluids*, 15(8):2238–2250, 2003. 5, 10

[68] B. Debusschere, H. Najm, P. Pébay, O. Knio, R. Ghanem, and O. Le Maître. Numerical challenges in the use of polynomial chaos representations for stochastic processes. *SIAM J. Sci. Comp.*, 26(2):698–719, 2004. 5, 10, 13, 15

[69] J. Decker et al. Exascale workshop panel meeting report. Department of Energy, January 2010. `http://www.er.doe.gov/ascr/ProgramDocuments/Docs/TrivelpieceExascaleWorkshop.pdf`. 4

[70] C. Descelliers, R. Ghanem, and C. Soize. Polynomial chaos representation of a stochastic preconditioner. *International Journal of Numerical Methods in Engineering*, 64(5):618–634, 2005. 5

[71] C. Descelliers, R. Ghanem, and C. Soize. Maximum likelihood estimation of stochastic chaos representation from experimental data. *International Journal for Numerical Methods in Engineering*, 66(6):978–1001, 2006. 5

[72] D. L. Donoho. Compressed sensing. *IEEE Transactions on Information Theory*, 52(4):1289–1306, 2006. 10

[73] A. Doostan and H. Owhadi. A non-adapted sparse approximation of PDEs with stochastic inputs. *Journal of Computational Physics*, 2011. 9, 10

[74] A. Doostan, R. Ghanem, and J. Red-Horse. Stochastic model reduction for chaos representations. *Computer Methods in Applied Mechanics and Engineering*, 196(37-40):3951–3966, 2007. 5, 9

[75] A. Doostan, R. G. Ghanem, and J. Red-Horse. Stochastic model reduction for chaos representations. *Computer Methods in Applied Mechanics and Engineering*, 196:3951–3966, 2007. 9, 10

[76] Y. Efendiev, A. Datta-Gupta, X. Ma, and B. Mallick. Efficient sampling techniques for uncertainty quantification in history matching using nonlinear error models and ensemble level upscaling techniques. *Water Resources Research*, 45(11):W11414, 2009. ISSN 0043-1397. 15

[77] M. Eldred. Design Under Uncertainty Employing Stochastic Expansion Methods. *International Journal for Uncertainty Quantification*, 1(2):119–146, 2011. 5, 12, 13

[78] M. Eldred, L. Swiler, and G. Tang. Mixed Aleatory-Epistemic Uncertainty Quantification with Stochastic Expansions and Optimization-Based Interval Estimation. *Reliability Engineering and System Safety*, to appear, 2011. 5, 9, 12

[79] M. S. Eldred and B. J. Bichon. Second-order reliability formulations in DAKOTA/UQ. In *Proceedings of the 47th AIAA/ASME/ASCE/AHS/ASC Structures, Structural Dynamics and Materials Conference*, number AIAA-2006-1828, Newport, RI, May 1–4 2006. 5, 12, 13

[80] M. S. Eldred and L. P. Swiler. Efficient algorithms for mixed aleatory-epistemic uncertainty quantification with application to radiation-hardened electronics. part i: Algorithms and benchmark results. Technical Report SAND2009-5805, Sandia National Laboratories, Albuquerque, NM, September 2009. 5

[81] M. S. Eldred and L. P. Swiler. Towards goal-oriented stochastic design employing adaptive collocation methods. In *Proc. 13th AIAA/ISSMO Multidisciplinary Analysis and Optimization Conference*, number AIAA-2010-9125, Fort Worth, TX, 2010. 5, 9, 12, 13

[82] M. S. Eldred, W. E. Hart, B. D. Schimel, and B. G. van Bloemen Waanders. Multilevel parallelism for optimization on MP computers: Theory and experiment. In *Proc. 8th AIAA/USAF/NASA/ISSMO Symposium on Multidisciplinary Analysis and Optimization*, number AIAA-2000-4818, Long Beach, CA, 2000. 5, 6, 17

[83] M. S. Eldred, A. A. Giunta, S. F. Wojtkiewicz, Jr., and T. G. Trucano. Formulations for surrogate-based optimization under uncertainty. In *Proc. 9th AIAA/ISSMO Symposium on Multidisciplinary Analysis and Optimization*, number AIAA-2002-5585, Atlanta, GA, 2002. 5, 12, 13

[84] M. S. Eldred, A. A. Giunta, and B. G. van Bloemen Waanders. Multilevel parallel optimization using massively parallel structural dynamics. *Structural and Multidisciplinary Optimization*, 27(1-2):97–109, 2004. 5, 6, 17

[85] M. S. Eldred, H. Agarwal, V. M. Perez, S. F. Wojtkiewicz, Jr., and J. E. Renaud. Investigation of reliability method formulations in DAKOTA/UQ. *Structure & Infrastructure Engineering: Maintenance, Management, Life-Cycle Design & Performance*, 3(3):199–213, 2007. 5, 12, 13

[86] H. Elman, C. Miller, E. T. Phipps, and R. S. Tuminaro. Assessment of collocation and galerkin approaches to linear diffusion equations with random data. *International Journal for Uncertainty Quantification*, 1(1): 19–33, Jan 2011. 10

[87] I. Epanomeritakis, V. Akçelik, O. Ghattas, and J. Bielak. A Newton-CG method for large-scale three-dimensional elastic full-waveform seismic inversion. *Inverse Problems*, 24(3):034015 (26pp), 2008. doi: 10.1088/0266-5611/24/3/034015. 5, 6, 8

[88] B. Faverjon and R. Ghanem. Stochastic inversion in acoustic scattering. *JASA*, 119(6):3577–3588, 2006. 5

[89] W. Fendt and B. Wandelt. Pico: Parameters for the Impatient Cosmologist. *The Astrophysical Journal*, 654 (1):2–11, 2007. 15

[90] H. P. Flath, L. C. Wilcox, V. Akçelik, J. Hill, B. van Bloemen Waanders, and O. Ghattas. Fast algorithms for Bayesian uncertainty quantification in large-scale linear inverse problems based on low-rank partial Hessian approximations. *SIAM Journal on Scientific Computing*, 33(1):407–432, 2011. doi: 10.1137/090780717. 5, 6, 8

[91] D. Galbally, K. Fidkowski, K. Willcox, and O. Ghattas. Nonlinear model reduction for uncertainty quantification in large-scale inverse problems. *International Journal For Numerical Methods in Engineering*, 81: 1581–1608, 2010. 5

[92] G. Galli and T. Dunning. Discovery in basic energy sciences: The role of computing at the extreme scales. Workshop Report, Department of Energy, August 2009. `http://www.er.doe.gov/ascr/ProgramDocuments/Docs/BESExascaleReport.pdf`. 4

[93] B. Ganapathysubramanian and N. Zabaras. Sparse grid collocation schemes for stochastic natural convection problems. *J. Comput. Phys.*, 225(1):652–685, 2007. 9, 10

[94] A. Gelman, J. B. Carlin, H. S. Stern, and D. B. Rubin. *Bayesian Data Analysis*. Chapman & Hall CRC, 2 edition, 2003. 11

[95] A. Gelman, I. V. Mechelen, G. Verbeke, D. Heitjan, and M. Meulders. Multiple imputation for model checking: Completed–data plots with missing and latent data. *Biometrics*, 61:74–85, 2005. 11

[96] T. Gerstner and M. Griebel. Dimension-adaptive tensor product quadrature. *Computing*, 71:65–87, 2003. 9, 10

[97] R. Ghanem. Probabilistic characterization of transport in heterogeneous media. *Comput. Methods Appl. Mech. Engrg.*, 158:199–220, 1998. 5, 10

[98] R. Ghanem. Hybrid Stochastic Finite Elements and Generalized Monte Carlo Simulation. *ASME J. Appl. Mech.*, 65:1004–1009, 1998. 5, 10

[99] R. Ghanem. Ingredients for a general purpose stochastic finite element formulation. *Comput. Methods Appl. Mech. Engrg.*, 168:19–34, 1999. 5, 10

[100] R. Ghanem. Stochastic finite elements for heterogeneous media with multiple random non-gaussian properties. *ASCE J. Eng. Mech.*, 125:26–40, 1999. 5, 10

[101] R. Ghanem and J. Abras. A general purpose library for stochastic finite element computations. In *Second MIT Conference on Computational Mechanics*, Cambridge, MA, June 17-20 2003. 10

[102] R. Ghanem and R. Doostan. On the construction and analysis of stochastic models: Characterization and propagation of the errors associated with limited data. *Journal of Computational Physics*, 217(1):63–81, 2006. 5, 8, 9

[103] R. Ghanem and D. Ghiocel. A new implementation of the spectral stochastic finite element method for stochastic constitutive relations. In *Proceedings of the ASCE* 12^{th} *Engineering Mechanics Conference*, May 17-20 1998. 10

[104] R. Ghanem and D. Ghosh. An efficient characterization for the random eigenvalue problem in a polynomial chaos decomposition. *International Journal of Numerical methods in Engineering*, 72:486–504, 2007. 5, 10

[105] R. Ghanem and P. Spanos. *Stochastic Finite Elements: A Spectral Approach*. Springer Verlag, New York, 1991. 5, 7, 8

[106] R. Ghanem and P. Spanos. A stochastic Galerkin expansion for nonlinear random vibration analysis. *Probabilistic Engineering Mechanics*, 8:255–264, 1993. 5, 10

[107] R. Ghanem, J. Abras, S. Wojtkiewicz, and G. Reese. Integration of stochastic finite element capabilities into analysis software. In *Seventh US National Congress on Computational Mechanics*, Albuquerque, NM, July 27-30 2003. 10

[108] R. Ghanem, J. Red-Horse, A. Benjamin, A. Doostan, and Y. Zou. Stochastic process model for material properties under incomplete information. In *48th AIAA/ASME/ASCE/AHS/ASC Structures, Structural Dynamics, and Materials Conference*, Honolulu, Hawaii, April 23-26 2007. 9, 10

[109] R. Ghanem, A. Doostan, and J. Red-Horse. A probabilistic construction of model validation. *Computer Methods in Applied Mechanics and Engineering*, 197(43-44):3584–3592, 2008. 12

[110] O. Ghattas and J.-H. Bark. Optimal control of two- and three-dimensional incompressible Navier-Stokes flows. *Journal of Computational Physics*, 136:231–244, 1997. 5, 8

[111] D. Ghiocel and R. Ghanem. Stochastic Finite-Element Analysis of Seismic Soil-Structure Interaction. *J. Eng. Mech.*, 128(1):66–77, 2002. 5, 10

[112] D. Ghosh and R. Ghanem. Stochastic convergence acceleration through basis enrichment of polynomial chaos expansions. *International Journal of Numerical methods in Engineering*, 73:162–184, 2008. 5, 10

[113] A. Giunta, L. Swiler, S. Brown, M. Eldred, M. Richards, , and E. Cyr. The Surfpack software library for surrogate modeling of sparse irregularly spaced multidimensional data. In *Proc. 11th AIAA/ISSMO Multidisciplinary Analysis and Optimization Conference*, number AIAA-2006-7049, Portsmouth, VA, 2006. 11

[114] M. Goldstein and J. C. Rougier. Reified Bayesian modelling and inference for physical systems (with discussion). *Journal of Statistical Planning and Inference*, 139:1221–1239, 2009. 15

[115] J. Graham, K. Bai, C. Meneveau, and J. Katz. Les modeling and experimental measurement of boundary layer flow over multi-scale, fractal canopies. *in: Proceedings of DLES-8, editor: H. Kuerten*, .:Springer, to appear, 2011. 6

[116] H. Haario, M. Laine, A. Mira, and E. Saksman. DRAM: Efficient adaptive MCMC. *Stat. Comput.*, 16: 339–354, 2006. 8

[117] S. Habib, K. Heitmann, D. Higdon, C. Nakhleh, and B. Williams. Cosmic calibration: Constraints from the matter power spectrum and the cosmic microwave background. *Physical Review D*, 76(8):083503, 2007. 6

[118] K. Heitmann, D. Higdon, M. White, S. Habib, B. Williams, E. Lawrence, and C. Wagner. The Coyote Universe. II. Cosmological Models and Precision Emulation of the Nonlinear Matter Power Spectrum. *The Astrophysical Journal*, 705:156, 2009. 6

[119] D. Higdon, H. Lee, and Z. Bi. A Bayesian approach to characterizing uncertainty in inverse problems using coarse and fine scale information. *IEEE Transactions in Signal Processing*, 50:389–399, 2002. 6

[120] D. Higdon, M. Kennedy, J. Cavendish, J. Cafeo, and R. D. Ryne. Combining field observations and simulations for calibration and prediction. *SIAM Journal of Scientific Computing*, 26:448–466, 2004. 8

[121] D. Higdon, J. Gattiker, B. Williams, and M. Rightley. Computer Model Calibration Using High-Dimensional Output. *Journal of the American Statistical Association*, 103(482):570–583, 2008. 8, 13

[122] S. Hosder, R. Walters, and R. Perez. A non-intrusive polynomial chaos method for uncertainty propagation in cfd simulations. In *Paper AIAA 2006-0891, 44th AIAA Aerospace Sciences Meeting and Exhibit*, Reno, NV, Jan 2006. 11

[123] X. Huan and Y. Marzouk. Optimal Bayesian experimental design for combustion kinetics. In *AIAA-2011-0513*, 2011. 5, 12

[124] J. Illian and H. Rue. A toolbox for fitting complex spatial point process models using integrated Laplace transformation (INLA). Technical report, Technical Report 6, Department of Mathematical Sciences, Norwegian University of Science and Technology, 2010. 8

[125] J. Jakeman, M. S. Eldred, and D. Xiu. A numerical approach for quantification of epistemic uncertainty. *Journal of Computational Physics*, 229(12), 2010. 5

[126] C. Kaufman, M. Schervish, and D. Nychka. Covariance tapering for likelihood-based estimation in large spatial data sets. *Journal of the American Statistical Association*, 103(484):1545–1555, 2008. ISSN 0162-1459. 8

[127] M. Kennedy and A. O'Hagan. Predicting the output from a complex computer code when fast approximations are available. *Biometrika*, 87:1–13, 2000. 15

[128] M. Kennedy and A. O'Hagan. Bayesian calibration of computer models. *Journal of the Royal Statistical Society. Series B (Statistical Methodology)*, 63(3):425–464, 2001. ISSN 1369-7412. 8, 15

[129] D. E. Keyes. Exaflop/s: The why and how. *Comptes Rendus Mecanique*, 339:70–77, 2011. 16

[130] O. Knio and O. Le Maître. Uncertainty propagation in cfd using polynomial chaos decomposition. *Fluid Dynamics Research*, 38(9):616 – 40, September 2006. 7

[131] K. Ko, N. Folwell, L. Ge, A. Guetz, V. Ivanov, A. Kabel, M. Kowalski, L. Lee, Z. Li, C. Ng, E. Prudencio, G. Schussman, R. Uplenchwar, L. Xiao, E. Ng, W. Gao, X. Li, C. Yang, P. Husbands, A. Pinar, D. Bailey, D. Gunter, L. Diachin, D. Brown, D. Quinlan, R. Vuduc, P. Knupp, K. Devine, J. Kraftcheck, O. Ghattas, V. Akcelik, D. Keyes, M. Shephard, E. Seol, A. Brewer, G. Golub, K. Ma, H. Yu, Z. Bai, B. Liao, T. Tautges, and H. Kim. Impact of SciDAC on accelerator projects across the Office of Science through electromagnetic modeling. *Journal of Physics: Conference Series*, 16:195–204, 2005. 6

[132] K. Ko, N. Folwell, L. Ge, A. Guetz, L. Lee, Z. Li, C. Ng, E. E. Prudencio, G. Schussman, R. Uplenchwar, and L. Xiao. Advances in electromagnetic modeling through high performance computing. *Physica C*, 441: 258–262, 2006. 6

[133] A. Kong, J. Liu, and W. Wong. Sequential imputations and bayesian missing data problems. *Journal of the American Statistical Association*, 89(425):278–288, 1994. 11

[134] C. C. Lalescu. Two hierarchies of spline interpolations. practical algorithms for multivariate higher order splines. *CoRR*, abs/0905.3564, 2009. 9

[135] E. Lawrence, K. Heitmann, M. White, D. Higdon, C. Wagner, S. Habib, and B. Williams. The Coyote Universe III: Simulation Suite and Precision Emulator for the Nonlinear Matter Power Spectrum. *Arxiv preprint arXiv:0912.4490*, 2009. 6

[136] O. Le Maître and O. Knio. A stochastic particle-mesh scheme for uncertainty propagation in vortical flows. *J. Comp. Phys.*, 226:645–671, 2007. 5, 10

[137] O. Le Maître, O. Knio, H. Najm, and R. Ghanem. A stochastic projection method for fluid flow. I. Basic formulation. *J. Comput. Phyics*, 173:481–511, 2001. 5, 10

[138] O. Le Maître, M. Reagan, H. Najm, R. Ghanem, and O. Knio. A stochastic projection method for fluid flow II. Random process. *J. Comput. Phys.*, 181:9–44, 2002. 5, 10, 11

[139] O. Le Maître, O. Knio, B. Debusschere, H. Najm, and R. Ghanem. A multigrid solver for two-dimensional stochastic diffusion equations. *Comp. Meth. App. Mech. and Eng.*, 192:4723–4744, 2003. 5

[140] O. Le Maître, R. Ghanem, O. Knio, and H. Najm. Uncertainty propagation using Wiener-Haar expansions. *J. Comput. Phys.*, 197(1):28–57, 2004. 5, 10

[141] O. Le Maître, H. Najm, R. Ghanem, and O. Knio. Multi-resolution analysis of Wiener-type uncertainty propagation schemes. *J. Comput. Phys.*, 197:502–531, 2004.

[142] O. Le Maître, H. Najm, R. Ghanem, and O. Knio. Multi-resolution analysis of Wiener-type uncertainty propagation schemes. *J. Comput. Phys.*, 197(2):502–531, 2004. 5, 10

[143] O. Le Maître, M. Reagan, B. Debusschere, H. Najm, R. Ghanem, and O. Knio. Natural convection in a closed cavity under stochastic, non-Boussinesq conditions. *SIAM J. Sci. Comput.*, 26(2):375–394, 2004. 6

[144] O. Le Maître, H. Najm, P. Pébay, R. Ghanem, and O. Knio. Multi-resolution-analysis scheme for uncertainty quantification in chemical systems. *SIAM J. Sci. Comput.*, 29(2):864–889, 2007. 6, 10

[145] H. Lee, D. Higdon, Z. Bi, M. Ferreira, and M. West. Markov random field models for high-dimensional parameters in simulations of fluid flow in porous media. *Technometrics*, 44(3):230–241, 2002. ISSN 0040-1706. 6

[146] L.-Q. Lee, L. Ge, Z. Li, C. Ng, G. Schussman, K. Ko, E. Ng, W. Gao, P. Husbands, X. Li, C. Yang, V. Akcelik, O. Ghattas, D. Keyes, T. Tautges, H. Kim, J. Craftcheck, P. Knupp, L. F. Diachin, K.-L. Ma, Z. Bai, and G. Golub. Achievements in ISICs/SAPP collaborations for electromagnetic modeling of accelerators. *Journal of Physics: Conference Series*, 16:205–209, 2005. 6

[147] Y. Li, E. Perlman, M. Wan, Y. Yang, C. Meneveau, R. Burns, S. Chen, A. Szalay, and G. Eyink. A public turbulence database cluster and applications to study Lagrangian evolution of velocity increments in turbulence. *J. Turbulence*, 9:No. 31, 2008. 6

[148] Z. Li, V. Akcelik, A. Candel, S. Chen, L. Ge, A. Kabel, L. Lee, C. Ng, E. E. Prudencio, G. Schussman, R. Uplenchwar, L. Xiao, and K. Ko. Towards simulation of electromagnetic and beam physics at the petascale. *Proceedings of PAC 2007*, 2007. 6

[149] C. Lieberman, K. Willcox, and O. Ghattas. Parameter and state model reduction for large-scale statistical inverse problems. *SIAM Journal on Scientific Computing*, 32(5):2523–2542, 2010. doi: 10.1137/090775622. 5, 6, 8

[150] K.-L. Ma, A. Stompel, J. Bielak, O. Ghattas, and E. Kim. Visualizing very large-scale earthquake simulations. In *Proceedings of ACM/IEEE SC2003*, 2003. 6

[151] X. Ma and N. Zabaras. An adaptive hierarchical sparse grid collocation algorithm for the solution of stochastic differential equations. *J. Comput. Phys.*, 228:3084–3113, 2009. 9, 10

[152] O. L. Maître, L. Mathelin, O. Knio, and M. Hussaini. Asynchronous time integration for polynomial chaos expansion of uncertain periodic dynamics. *Discrete and Continuous Dynamical Systems*, 28:199–226, 2010. 5

[153] J. Martin, L. C. Wilcox, C. Burstedde, and O. Ghattas. A stochastic Newton MCMC method for large scale statistical inverse problems with application to seismic inversion. In preparation. 8

[154] Y. M. Marzouk and A. F. Ghoniem. K-means clustering for optimal partitioning and dynamic load balancing of parallel hierarchical N-body simulations. *J. Comp. Phys.*, 207:493–528, 2005. 6

[155] Y. M. Marzouk and H. N. Najm. Dimensionality reduction and polynomial chaos acceleration of Bayesian inference in inverse problems. *J. Comp. Phys.*, 228(6):1862–1902, 2009. 5, 8

[156] Y. M. Marzouk and D. Xiu. A stochastic collocation approach to Bayesian inference in inverse problems. *Comm. Comput. Phys.*, 6(4):826–847, 2009.

[157] Y. M. Marzouk, H. N. Najm, and L. A. Rahn. Stochastic spectral methods for efficient Bayesian solution of inverse problems. *J. Comp. Phys.*, 224:560–586, 2007. 5, 8

[158] L. Mathelin and K. A. Gallivan. A compressed sensing approach for partial differential equations with random input data. *Computer Methods in Applied Mechanics and Engineering*, 2011. 10

[159] L. Mathelin and O. Le Maître. Dual-based a posteriori error estimate for stochastic finite element methods. *Comm. App. Math. Comp. Sci.*, 2(1):83–115, 2007. 10

[160] H. Matthies and C. Bucher. Finite elements for stochastic media problems. *Comput. Methods Appl. Mech. Engrg.*, 168:3–17, 1999. 10

[161] K. Maute, G. Weickum, and M. S. Eldred. A reduced-order stochastic finite element approach for design optimization under uncertainty. *Structural Safety*, 31(6):450–459, 2009. 5, 12

[162] S. McKenna, J. Ray, Y. M. Marzouk, and B. van Bloemen Waanders. Truncated multi-Gaussian fields and effective conductance of binary media. *Advances in Water Resources*, 34:617–626, 2011. 6

[163] P. Messina and A. Bishop. Scientific Grand Challenges in National Security: The Role of Computing at the Extreme Scale. Workshop Report, Department of Energy, October 2009. `http://www.er.doe.gov/ascr/ProgramDocuments/Docs/NNSAGrandChallengesReport.pdf`. 4

[164] J. Meyers and C. Meneveau. Large eddy simulations of large wind-turbine arrays in the atmospheric boundary layer, paper aiaa 2010-827. In *48th AIAA Aerospace Sciences Meeting*, January, Orlando, Florida 2010. 6

[165] E. Moniz and R. Rosner. Science based nuclear energy systems enabled by advanced modeling and simulation at the extreme scale. Workshop Report, Department of Energy, May 2009. `http://www.er.doe.gov/ascr/ProgramDocuments/Docs/SC-NEWorkshopReport.pdf`. 4

[166] H. Najm. Uncertainty Quantification and Polynomial Chaos Techniques in Computational Fluid Dynamics. *Annual Review of Fluid Mechanics*, 41(1):35–52, 2009. 7

[167] H. Najm, B. Debusschere, Y. Marzouk, S. Widmer, and O. Le Maître. Uncertainty Quantification in Chemical Systems. *Int. J. Num. Meth. Eng.*, 80:789–814, 2009. 6, 9

[168] H. Najm, M. Valorani, D. Goussis, and J. Prager. Analysis of Methane-Air Edge Flame Structure. *Combustion Theory and Modelling*, 14(2):257–294, 2010. 6

[169] F. Nobile, R. Tempone, and C. Webster. A sparse grid stochastic collocation method for partial differential equations with random input data. *SIAM J. Num. Anal.*, 46(5):2309–2345, 2008. 9, 10

[170] F. Nobile, R. Tempone, and C. Webster. An anisotropic sparse grid stochastic collocation method for partial differential equations with random input data. *SIAM J. Num. Anal.*, 46(5):2411–2442, 2008. 9, 10

[171] J. Oden et al. Simulation-Based Engineering Science: Revolutionizing Engineering Science through Simulation. A Report of the NSF Blue Ribbon Panel on Simulation-Based Engineering Science (SBES), 2006. http://www.nsf.gov/pubs/reports/sbes_final_report.pdf. 4

[172] J. T. Oden, O. Ghattas, et al. Cyber Science and Engineering: A Report of the NSF Advisory Committee for Cyberinfrastructure Task Force on Grand Challenges. National Science Foundation, 2011. http://www.nsf.gov/od/oci/taskforces/. 4

[173] N. Onishi and J. Glanz. Japanese Rules for Nuclear Plants Relied on Old Science. New York Times, March 26, 2011. http://www.nytimes.com/2011/03/27/world/asia/27nuke.html. 4

[174] M. Pellissetti and R. Ghanem. Iterative solution of systems of linear equations arising in the context of the stochastic fem. *Journal of Advances in Engineering Software*, 31:607–616, 2000. 5

[175] R. A. Perez. *Uncertainty Analysis of Computational Fluid Dynamics Via Polynomial Chaos*. PhD thesis, Virginia Polytechnic Institute and State University, 2008. 9

[176] E. Phipps, R. Bartlett, D. Gay, and R. Hoekstra. Large-scale transient sensitivity analysis of a radiation-damaged bipolar junction transistor via AD. In C. Bischof, M. Bücker, P. Hovland, U. Naumann, and J.Utke, editors, *Advances in Automatic Differentiation*, Lecture Notes in Computational Science and Engineering, 2008. 6

[177] J. Prager, H. Najm, M. Valorani, and D. Goussis. Skeletal Mechanism Generation with CSP and Validation for Premixed n-Heptane Flames. *Proc. Comb. Inst.*, 32(1):509–517, 2008.

[178] J. Prager, H. Najm, M. Valorani, and D. Goussis. Structure of n-Heptane/Air Triple Flames in Partially-Premixed Mixing Layers. *Combustion and Flame*, 2011. in press.

[179] E. E. Prudencio and X. C. Cai. Robust multilevel restricted Schwarz preconditioners and applications. *Domain Decomposition Methods in Science and Engineering XVI, O. B. Widlund and D. E. Keyes, eds., vol. 55 of Lecture Notes in Computational Science and Engineering, Springer-Verlag*, pages 155–162, 2006.

[180] E. E. Prudencio and X. C. Cai. Parallel multilevel restricted Schwarz preconditioners with pollution removing for PDE-constrained optimization. *SIAM J. Sci. Comp.*, 29:964–985, 2007. 6

[181] E. E. Prudencio and S. H. Cheung. Parallel adaptive multi-level sampling algorithms for the Bayesian analysis of mathematical models. *to be submitted*, May 2011. 5, 8

[182] E. E. Prudencio and K. Schulz. The parallel C++ statistical library 'QUESO': Quantification of Uncertainty for Estimation, Simulation and Optimization. *in preparation*, 2011. 8, 13

[183] E. E. Prudencio, R. Byrd, and X. C. Cai. Parallel full space SQP Lagrange-Newton-Krylov-Schwarz algorithms for PDE-constrained optimization problems. *SIAM J. Sci. Comp.*, 27:1305–1328, 2006. 6

[184] E. E. Prudencio, A. Candel, L. Ge, A. Kabel, K. Ko, L. Lee, Z. Li, C. Ng, and G. Schussman. Parallel 3D finite element numerical modeling of DC electron guns. *SLAC-PUB-13097*, 2008. 6

[185] J. Ray, R. Armstrong, C. Safta, B. Debusschere, B. Allan, and H. Najm. Computational frameworks for advanced combustion simulations. In T. Echekki and N. Mastorakos, editors, *Turbulent Combustion Modeling: Advances, New Trends and Perspectives*, pages 409–437. Springer-Verlag, Berlin,, 2011. 6

[186] M. Reagan, H. Najm, R. Ghanem, and O. Knio. Uncertainty quantification in reacting flow simulations through non-intrusive spectral projection. *Combustion and Flame*, 132:545–555, 2003. 5, 10, 11

[187] M. Reagan, H. Najm, B. Debusschere, O. Le Maître, O. Knio, and R. Ghanem. Spectral stochastic uncertainty quantification in chemical systems. *Combustion Theory & Modeling*, 8:607–632, 2004. 5, 6, 10

[188] M. Reagan, H. Najm, P. Pébay, O. Knio, and R. Ghanem. Quantifying uncertainty in chemical systems modeling. *Int. J. Chem. Kin.*, 37(6):368–382, 2005. 5, 6, 10

[189] F. Rizzi, M. Salloum, Y. Marzouk, R.-G. Xu, M. Falk, T. Weihs, G. Fritz, and O. Knio. Bayesian inference of atomic diffusivity in a binary ni/al system based on molecular dynamics. *Multiscale Modelling & Simulation*, 9:486–512, 2011. 6

[190] D. Rubin. Multiple imputation after 18+ years. *Journal of the American Statistical Association*, 91(434), 1996. 11

[191] H. Rue, L. Held, and E. Corporation. *Gaussian Markov random fields: theory and applications*, volume 734. Chapman & Hall/CRC Boca Raton:, 2005. 8

[192] J. Sacks, W. J. Welch, T. J. Mitchell, and H. P. Wynn. Design and analysis of computer experiments (with discussion). *Statistical Science*, 4:409–423, 1989. 8, 15

[193] C. Safta, J. Ray, and H. Najm. A High-Order Low-Mach Number AMR Construction for Chemically Reacting Flows. *J. Comput. Phys.*, 229(24):9299–9322, 2010. 6

[194] A. Salinger, E. Burroughs, R. Pawlowski, E. Phipps, and L. Romero. Bifurcation tracking algorithms and software for large scale applications. *Int J Bifurcat Chaos*, 15(3):1015–1032, Jan 2005. 6

[195] B. Sanso, C. Forest, and D. Zantedeschi. Inferring climate system properties using a computer model. *Bayesian Analysis*, 3(1):1–38, 2008. 8

[196] K. Sargsyan, C. Safta, B. Debusschere, and H. Najm. Uncertainty quantification given discontinuous model response and a limited number of model runs. *SIAM J. Scientific Computing*, submitted, 2010. 11

[197] C. Soize and R. Ghanem. Reduced chaos decomposition with random coefficients of vector-valued random variables and random fields. *Computer Methods in Applied Mechanics and Engineering*, in press, 2009. 5, 9, 10

[198] G. Stadler, M. Gurnis, C. Burstedde, L. C. Wilcox, L. Alisic, and O. Ghattas. The dynamics of plate tectonics and mantle flow: From local to global scales. *Science*, 329(5995):1033–1038, 2010. doi: 10.1126/science.1191223. 6

[199] R. Stevens and M. Ellisman. Opportunities in biology at the extreme scale of computing. Workshop Report, Department of Energy, August 2009. `http://www.er.doe.gov/ascr/ProgramDocuments/Docs/BiologyReport.pdf`. 4

[200] R. Stevens, T. Zacharia, and H. Simon. Modeling and simulation at the exascale for energy and the environment. Report on the Advanced Scientific Computing Research Town Hall Meetings on Simulation and Modeling at the Exascale for Energy, Ecological Sustainability and Global Security (E3), Department of Energy, 2008. `http://www.sc.doe.gov/ascr/ProgramDocuments/Docs/TownHall.pdf`. 4

[201] W. Tang and D. Keyes. Fusion energy sciences and the role of computing at the extreme scale. Workshop Report, Department of Energy, March 2009. `http://www.er.doe.gov/ascr/ProgramDocuments/Docs/FusionReport.pdf`. 4

[202] M. Tatang. *Direct Incorporation of Uncertainty in Chemical and Environmental Engineering Systems*. PhD thesis, Mass. Inst. Tech., 1995. 11

[203] J. TenCate, D. Pasqualini, S. Habib, K. Heitmann, D. Higdon, and P. Johnson. Nonlinear and nonequilibrium dynamics in geomaterials. *Physical review letters*, 93(6):65501, 2004. ISSN 1079-7114. 6

[204] J. A. Tropp and A. C. Gilbert. Signal recovery from random measurements via orthogonal matching pursuit. *IEEE Transactions on Information Theory*, 53(12):4655–4666, 2007. 10

[205] T. Tu, D. R. O'Hallaron, and O. Ghattas. Scalable parallel octree meshing for terascale applications. In *SC '05: Proceedings of the International Conference for High Performance Computing, Networking, Storage, and Analysis*. ACM/IEEE, 2005. ISBN 1-59593-061-2. doi: http://dx.doi.org/10.1109/SC.2005.61. 6

[206] T. Tu, H. Yu, L. Ramirez-Guzman, J. Bielak, O. Ghattas, K.-L. Ma, and D. R. O'Hallaron. From mesh generation to scientific visualization: An end-to-end approach to parallel supercomputing. In *Proceedings of SC2006*, Tampa, FL, 2006.

[207] M. Wan, Z. Xiao, C. Meneveau, G. Eyink., and S. Chen. Dissipation-energy flux correlations as evidence for the Lagrangian energy cascade in turbulence. *Phys. Fluids*, 22:061702, 2010. 6

[208] X. Wan and G. Karniadakis. Multi-Element Generalized Polynomial Chaos for Arbitrary Probability Measures. *SIAM J. Sci. Comput.*, 28(3):901–928, 2006. 9

[209] X. Wan, D. Xiu, and G. Karniadakis. Stochastic solutions for the two-dimensional advection-diffusion equation. *SIAM J. Sci. Comput.*, 26(2):578–590, 2004. 10

[210] W. Washington. Challenges in climate change science and the role of computing at the extreme scale. Workshop Report, Department of Energy, November 2008. `http://www.er.doe.gov/ascr/ProgramDocuments/Docs/ClimateReport.pdf`. 4

[211] L. C. Wilcox, G. Stadler, C. Burstedde, and O. Ghattas. A high-order discontinuous Galerkin method for wave propagation through coupled elastic-acoustic media. *Journal of Computational Physics*, 229(24): 9373–9396, 2010. 6

[212] B. Williams, D. Higdon, J. Gattiker, L. Moore, M. McKay, and S. Keller-McNulty. Combining experimental data and computer simulations, with an application to flyer plate experiments. *Bayesian Analysis*, 1(4):765–792, 2006. 6

[213] D. Xiu and G. Karniadakis. Modeling uncertainty in steady state diffusion problems via generalized polynomial chaos. *Computer Methods in Applied Mechanics and Engineering*, 191:4927–4948, 2002. 10

[214] D. Xiu and G. Karniadakis. The Wiener-Askey polynomial chaos for stochastic differential equations. *SIAM J. Sci. Comp.*, 24(2):619–644, 2002.

[215] D. Xiu and G. Karniadakis. A new stochastic approach to transient heat conduction modeling with uncertainty. *Int. J. of Heat and Mass Transfer*, 46:4681–4693, 2003. 10

[216] H. Yang, E. E. Prudencio, and X. C. Cai. Fully implicit Lagrange-Newton-Krylov-Schwarz algorithms for boundary control of unsteady incompressible flows. *submitted*, 2011. 6

[217] S. Yang, G. Stadler, R. Moser, and O. Ghattas. A shape Hessian-based boundary roughness analysis of Navier–Stokes flow. *SIAM Journal on Applied Mathematics*, 71(1):333–355, 2011. doi: 10.1137/100796789. 5, 8

[218] G. Young, D. Dean, and M. Savage. Forefront questions in nuclear science and the role of computing at the extreme scale. Workshop Report, Department of Energy, January 2009. http://www.er.doe.gov/ascr/ProgramDocuments/Docs/NPReport.pdf. 4