Quantification of Uncertainty in Extreme Scale Computations (QUEST)

www.quest-scidac.org

H. Najm[1], B. Debusschere[1], M. Eldred[1], **R. Ghanem**[2], **O. Ghattas**[3], R. Moser[3], E. Prudencio[3], **D. Higdon**[4], J. Gattiker[4], **O. Knio**[5], **Y. Marzouk**[6]

[1]Sandia National Laboratories, Livermore, CA & Albuquerque, NM
[2]University of Southern California, Los Angeles, CA
[3]University of Texas, Austin, TX
[4]Los Alamos National Laboratory, Los Alamos, NM
[5]Duke University, Durham, NC
[6]Massachusetts Institute of Technology, Cambridge, MA

SciDAC PI Meeting, 10–12 Sep 2012, Rockville, MD

Outline

1. Introduction
2. QUEST Overview
3. Technical Progress
4. Closure

Introduction – Motivation

Why Uncertainty Quantification (UQ) ?

- Assessment of confidence in computational predictions
- Validation and comparison of scientific/engineering models
- Design optimization
- Use of computational predictions for decision-support
- Assimilation of observational data and model construction

Why UQ in SciDAC ?

- Explore model response over range of parameter variation
- Enhanced understanding extracted from computations
- Particularly important given *cost* of SciDAC computations

QUEST Goals

1. Advance the state of the art in UQ theory, methods, and software, addressing UQ challenges with extreme scale computational problems
   - High-dimensionality
   - Nonlinearity
   - Sparse data
2. Provide expertise, advice, and state of the art UQ algorithms and software tools to SciDAC projects
   - UQ software products
   - SciDAC partnerships
   - Outreach: UQ tutorials, summer school, web

Scope

The scope of QUEST covers a range of UQ activities including:

- UQ problem setup
- Characterization of the input space
- Local and global sensitivity analysis
- Adaptive stochastic dimensionality and order reduction
- Forward and Inverse UQ
- Fault tolerant UQ methods
- Model comparison and validation

Key Elements of our UQ strategy

- Probabilistic framework
  - Uncertainty is represented using probability theory
- Parameter Estimation, Model Calibration
  - Experimental measurements
  - Regression, Bayesian Inference
- Forward propagation of uncertainty
  - Polynomial Chaos (PC) Stochastic Galerkin methods
    - Intrusive/non-intrusive
  - Stochastic Collocation methods
- Model comparison, selection, and validation
- Model averaging
- Experimental design and uncertainty management

Team Expertise and Capabilities

Institution	Expertise	Tools
SNL	Forward and inverse UQ methods, design under uncertainty	DAKOTA UQTK
USC	Intrusive UQ methods probabilistic modeling	
Duke	Sparse adaptive forward UQ methods	
UT	Large scale inverse problems validation, inverse UQ	QUESO
LANL	Gaussian process modeling, inverse UQ	GPMSA
MIT	Calibration, adaptive sampling, inverse UQ, experimental design	

QUEST UQ Software tools

DAKOTA

- Optimization and calibration
- Non-intrusive UQ
- Global Sensitivity Analysis
- > 10K registered downloads

QUESO

- Bayesian Inference
- Parallel MultiChain MCMC
- Bayesian Model Analysis
- Model Calibration

GPMSA

- Bayesian Inference
- Gaussian Process Emulation
- Model Calibration
- Model discrepancy analysis

UQTk

- Intrusive PC UQ
- Non-intrusive sampling
- Customized sparse PCE
- Random fields

QUEST Partnerships

DOE	Project Title	Lead PI	QUEST
NNSA	Parallel Dislocation Simulator (ParaDiS)	T. Arsenlis LLNL	Najm SNL
FES	Center for Edge Plasma Physics Simulation (EPSI)	C.S. Chang Princeton	Moser UT
FES	Plasma Surface Interactions: Bridging from the Surface to the Micron Frontier	B. Wirth ORNL	Higdon LANL
BER	Predicting Ice Sheet & Climate Evolution at Extreme Scales (PISCEES)	P. Jones LANL	Eldred, Ghattas SNL, UT
BER	Multiscale Methods for Accurate, Efficient & Scale-Aware Earth System Modeling	B. Collins LBNL	Debusschere SNL
NP	Nuclear Computational Low Energy Initiative (NUCLEI)	J. Carlson LANL	Higdon LANL
HEP	Computation-Driven Discovery for the Dark Universe	S. Habib ANL	Higdon LANL
HEP	Community Project for Accelerator Science & Simulation (ComPASS)	P. Spentzouris FNAL	Prudencio UT

Outreach Activities

- Website
  - www.quest-scidac.org
  - Production version will be publicly accessible soon
- UQ Tutorials in workshops/conferences
  - SAMSI UQ workshop, Raleigh, NC; Sep 7-9, 2011
  - SIAM Conference on UQ, Raleigh, NC; Apr 2-5, 2012
- UQ Summer School
  - USC, LA; Aug 22-24, 2012
- UQ Tools Tutorial
  - Hands-on practice with UQ software tools
  - SNL, Livermore, CA; Oct 22-23, 2012.
  - Announcements went out in late July
  - http://cadmus.usc.edu/Quest-Tutorial
    - Some openings still available

SNL: Software: DAKOTA – dakota.sandia.gov

M. Eldred, J. Jakeman

- Development of interfaces: QUESO–**DAKOTA**–GPMSA
  - Ongoing
  - DAKOTA interfaces to both
  - C++ GPMSA implementation using QUESO components
- Stochastic collocation
  - Nodal or hierarchical interpolation on structured grids
  - Interpolants may be local or global
    - value-based or gradient-enhanced
  - Automated refinement
    - uniform, dimension-adaptive, or locally-adaptive
  - Hierarchical surplus error estimates for values and gradients applied to QoI (e.g., response covariance)
- Compressive sensing: basis pursuit and basis denoising

DAKOTA: Application in Nuclear Reactor Modeling

M. Eldred

- Work with CASL energy innovation hub
- PCE/SC with uniform/adaptive refinement vs LHS

$n = 4$, smooth, mild anisotropy

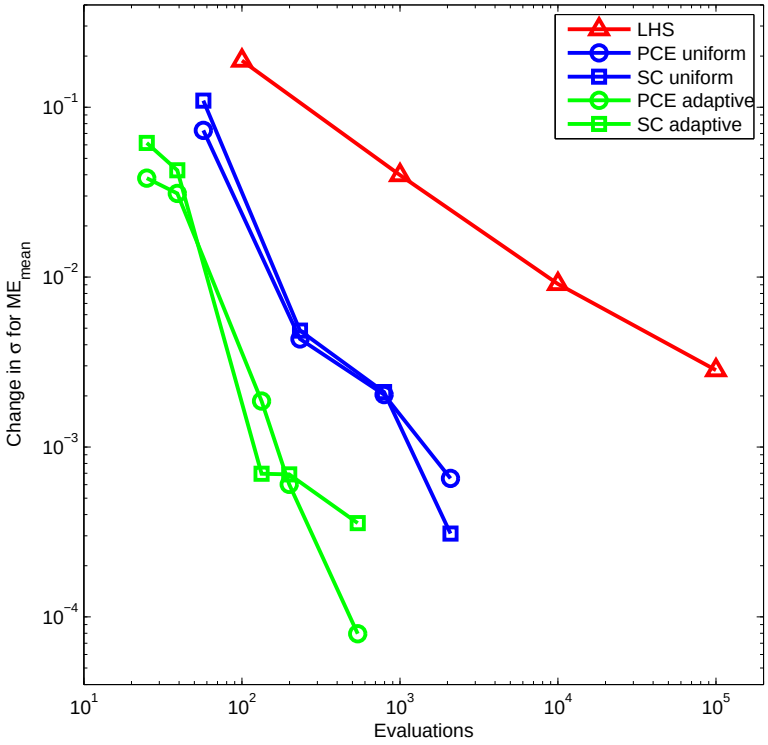

$n = 10$, discontinuous, high anisotropy

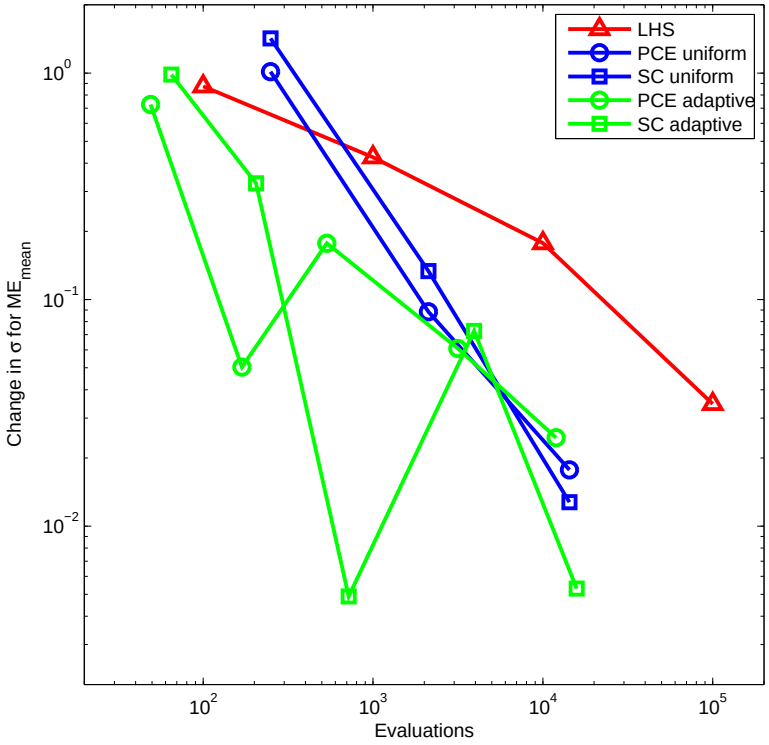

SNL: Software: UQTk – www.sandia.gov/UQToolkit

B. Debusschere, C. Safta, K. Sargsyan

- Version 1.0 published under the GNU LGPL
  - Intrusive PC functionality
- New release targeted for Fall 2012
  - Intrusive and non-intrusive utilities
  - User-specified multi-index capabilities
    - Flexible efficient sparse tensor representations
    - Effective for high-dimensional systems
  - Random fields:
    - Covariance matrix estimation (many samples)
    - Karhunen-Loève expansions (KLEs)
  - Matlab version
- Example/benchmark problems, tutorial materials

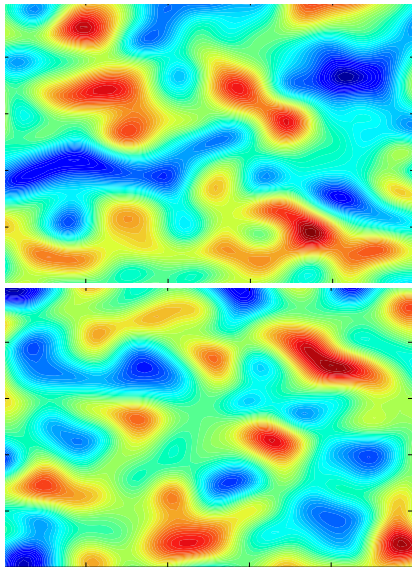

SNL: Algorithms: Gradients & Sparsity

M. Eldred, J. Jakeman, K. Sargsyan, C. Safta, B. Debusschere, H. Najm

- Hierarchical interpolation with generalized sparse grids
  - Gradient-enhancement
  - Error indicators leverage both value and gradient surpluses
- Building Sparse PC representations
  - Compressed Sensing (CS) – ℓ_1 regularization
    - cross validation, tolerances for model choice
  - Bayesian Compressed Sensing (BCS) – Laplace priors
  - BCS/CS comparisons on Genz functions – 5-10d
    - Similar convergence with no. of samples
    - Slightly higher accuracy with CS
    - BCS: $\mathcal{O}(100)\times$ reduction in no. of PCE terms

discovery of sparse signals:

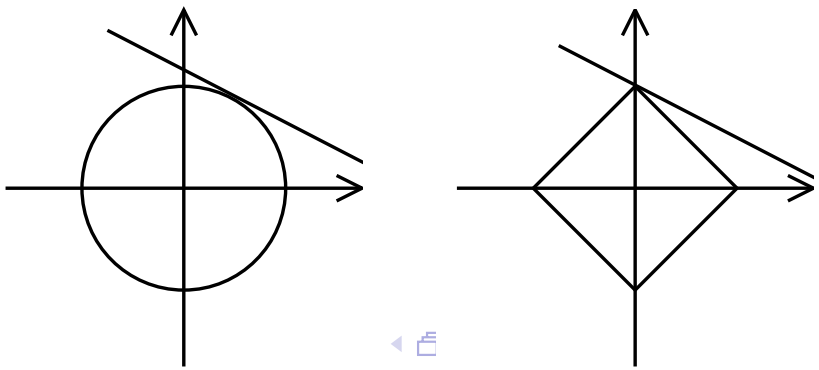

SNL: Algorithms: Missing Data

H. Najm, B. Debusschere, C. Safta, K. Sargsyan, K. Chowdhary

Context

- Missing/failed measurements or computational samples
- Partial specification of uncertain information
  - Error bars vs. joint PDF
- Processed data products

Imputation methods

- Existing data $\Rightarrow$ probabilistic prediction of missing data

Data Free Inference (DFI) algorithm

- Given information $\Rightarrow$ probabilistic models of missing data
  - Application in chemical ignition
  - Extension to processed data products

LANL: GPMSA & BART Developments

D. Higdon, J. Gattiker

- New release of GPMSA for sensitivity analysis and computer model calibration using Bayesian methods
  - Tutorial material
  - Range of sample problems
    - sensitivity, calibration, & multivariate output

- Prototype parallel implementation of the Bayesian additive regression tree (BART) for HPC.
  - linear scaling up to ∼50p
  - tests with higher proc counts in progress

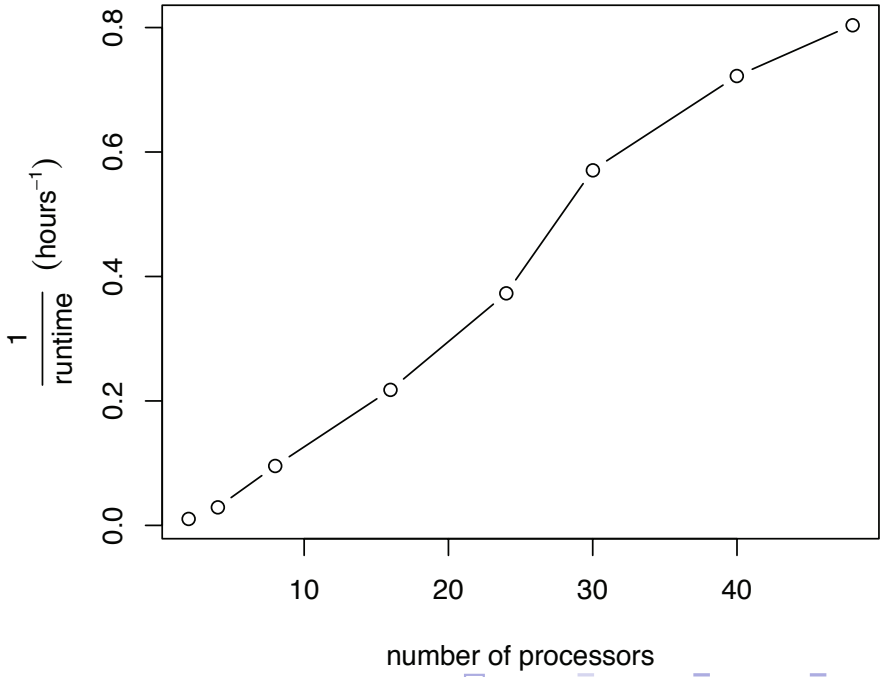

UT-Austin: Scalable Parallel Algorithms for Extreme-Scale Stochastic Inverse Problems

T. Bui-Thanh, O. Ghattas, J. Martin, G. Stadler (also funded by AFOSR and NSF)

Stochastic Inverse Probs:

- PDEs & high-dim parameter spaces (from discretized fields)
- Current methods are prohibitive
- Challenges:
  - appropriate choice of prior
  - consistent discretizations (guarantee convergence to infinite-dim problem)
  - scalable parallel MCMC algorithms

Recent accomplishments:

- Consistent discretizations via appropriate mass matrix weightings
- Prior defined by inverse of elliptic operator; carried out by multigrid
- Low rank approximation of Hessian enables sampling of Gaussianized posterior in dimension-independent number of forward solves
- Scaling to 1M parameters and 100K processor cores

Example: Extreme-scale Seismic Inversion

T. Bui-Thanh, O. Ghattas, J. Martin, G. Stadler

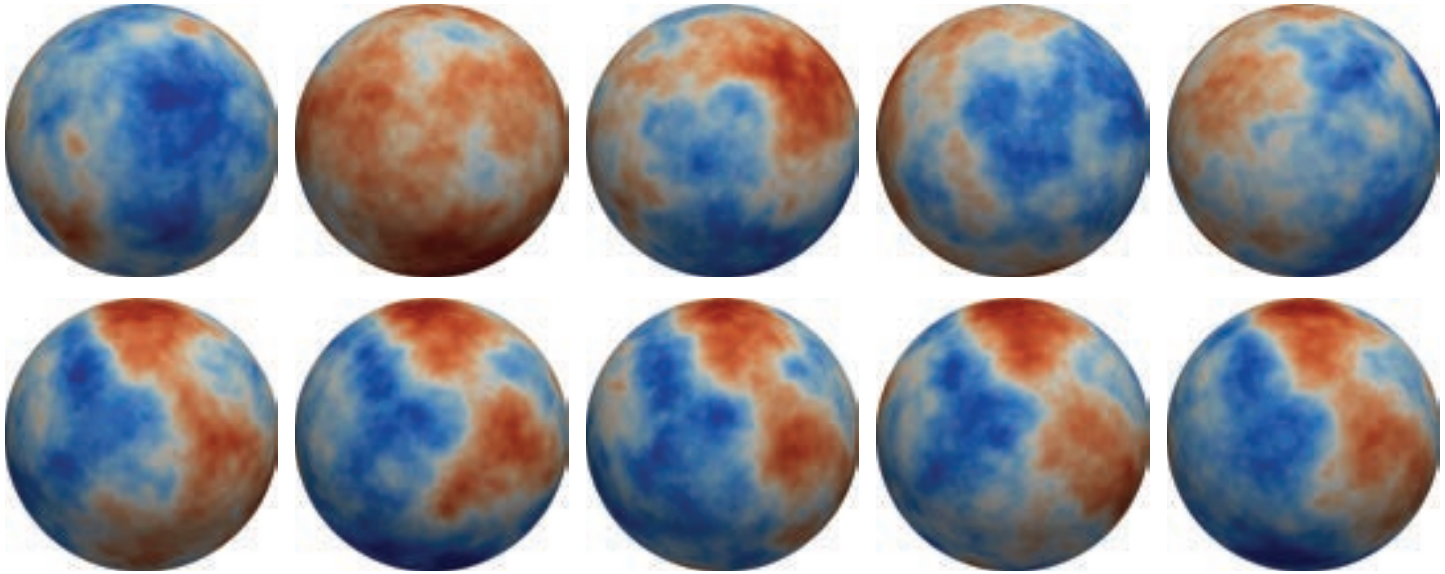

- linearized 3D global seismic inversion
- 1.07M earth model parameters
- 630M wave propagation unknowns
- 100K cores on Jaguar (ORNL)
- 2000$\times$ reduction in effective problem dimension due to low rank approx

- *Top row:* Prior samples
- *Bottom row:* Posterior samples
- Difference between rows indicates information gained from (and uncertainty reduced due to) data
- Gordon Bell Prize finalist, SC12

UT-Austin: Software: QUESO

K. C. Estacio-Hiroms, E. E. Prudencio, K. W. Schulz (also funded by NNSA)

- Improvement of QUESO-DAKOTA usability
  - Periodic output of samples
  - Output of extra information
  - Informative output summary
- Implementation of GPMSA models and algorithms
  - QUESO capabilities will be usable through DAKOTA
- Preparation of tutorial material
  - Bayesian inversion, and forward propagation of uncertainty
  - Object-oriented mapping of mathematical concepts
  - Solution of statistical inverse problems with DRAM MCMC
  - Solution of statistical forward problems with Monte Carlo
  - Use of parallel computing for statistical analysis
  - References to Bayesian analysis, MCMC, Monte Carlo, C/C++, MPI

Duke: Stochastic Preconditioning

O. Knio, A. Alexanderian, O. Le Maître

- Developed a multiscale Bayesian preconditioning approach
- Demonstrated capability to simultaneously
  - address stiffness and noise
  - represent noisy outputs w/sparse, low-order, PCEs
- Order of magnitude reduction in # of samples / replicas

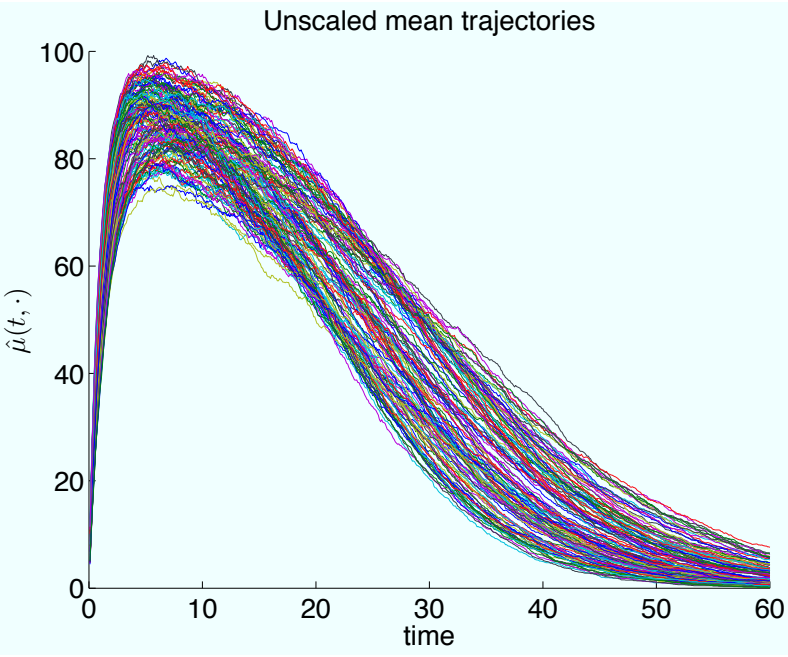

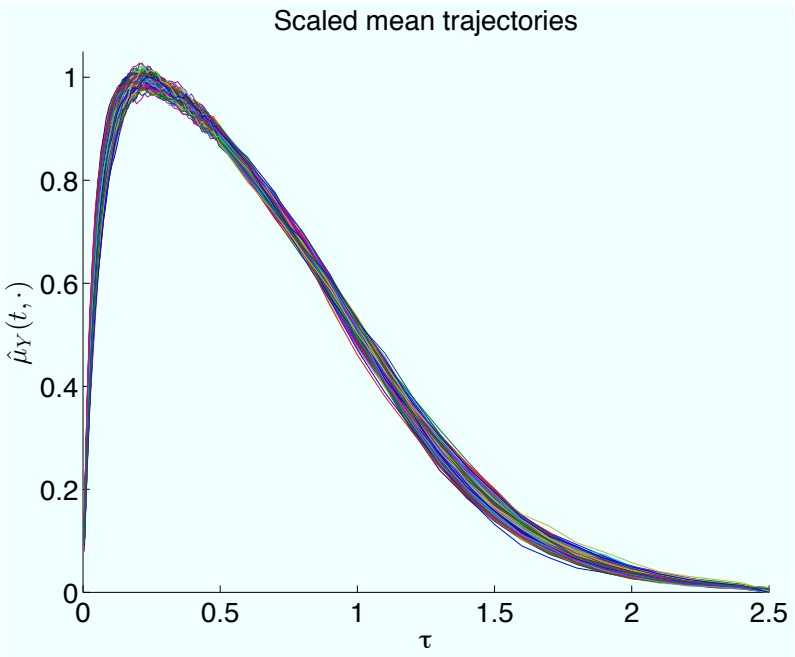

Duke-MIT: Sparse Adaptive Sampling

J. Winokur, P. Conrad, O. Knio, Y. Marzouk

- Developed a sparse adaptive pseudospectral sampling algorithm
  - accommodates arbitrary admissible stencils
  - including a maximal polynomial basis
    - without internal aliasing
- Analysis of algorithm performance based on existing Ocean General Circulation Model (OGCM) databases
- Demonstrated order-of-magnitude computational savings in simulations of the ocean circulation in the Pacific

MIT: Large-Scale Bayesian Inference

T. Moselhy, Y. Marzouk

Current state of the art

- Markov chain Monte Carlo (MCMC) sampling is the *workhorse algorithm* for Bayesian inference and prediction
- Challenges: enormous computational effort, difficult proposal design, insufficient convergence diagnostics

Inference with optimal maps

- *New approach:* find a deterministic map that *pushes forward* the prior measure to the posterior measure
- Converts inference to an optimization problem, with natural convergence diagnostics
- Outperforms MCMC in efficiency and accuracy on a variety of inference problems, with 10^3 dimensions or more

MIT: Large-Scale Bayesian Inference

T. Moselhy, Y. Marzouk

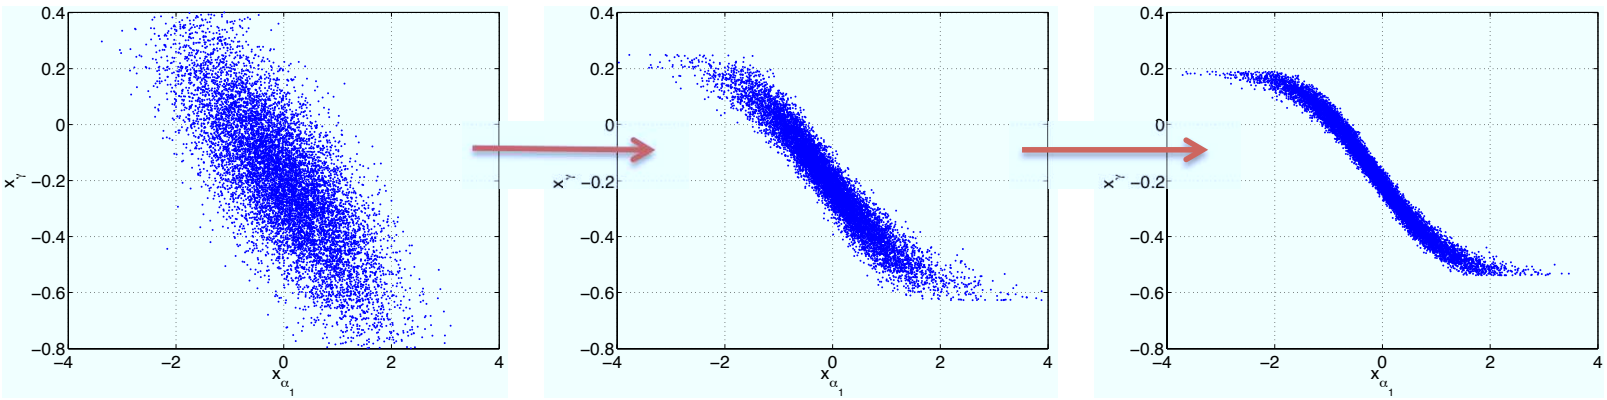

- *(above)* sequence of maps yields samples from non-Gaussian posterior in a chemical kinetic system
- Current work on map-based inference:
  - Hierarchical Bayesian models
  - Parallel algorithms for stochastic optimization
  - Sequential data assimilation (i.e., filtering and smoothing)

MIT: Optimal Experimental Design

X. Huan, Y. Marzouk

- How to choose observations or experimental conditions *optimally?*
  - Bayesian approach: maximize *expected information gain* for parameter inference, prediction, model discrimination, etc
- Key computational ingredients:
  - Surrogates for physical model describing experiments
  - Statistical estimators and stochastic optimization methods
- Recent accomplishments: **stochastic approximation** and **sample-average approximation** for optimal Bayesian design, using estimators of mutual information gradient

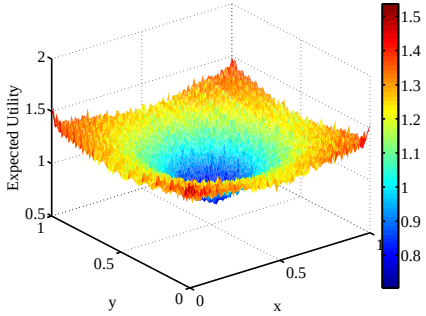

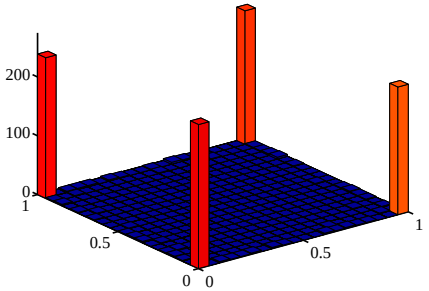

USC: Constrained & Adaptive Constructions

E. Kalligiannaki, R. Tipireddy, G. Ghanem

Develop Constrained Stochastic Representations

- Positive random variables
- More general constraints on either function values or values of nonlinear functionals of the random variables

Develop Bases Adapted to Quantity of Interest

- Scales linearly with stochastic dimension

USC: Constrained Stochastic Representations

E. Kalligiannaki, G. Ghanem

$\mathcal{I} = \{y(\omega) \in L_2(\Omega, \Sigma(H), P) : y(w)$ satisfies constraints $\forall \omega\}$

The projection of $y \in L_2$ on $\mathcal{I}$:

- Sample from prior PC expansion
- Delete realizations that do not satisfy constraints
- Recompute PC coefficients from remaining realizations

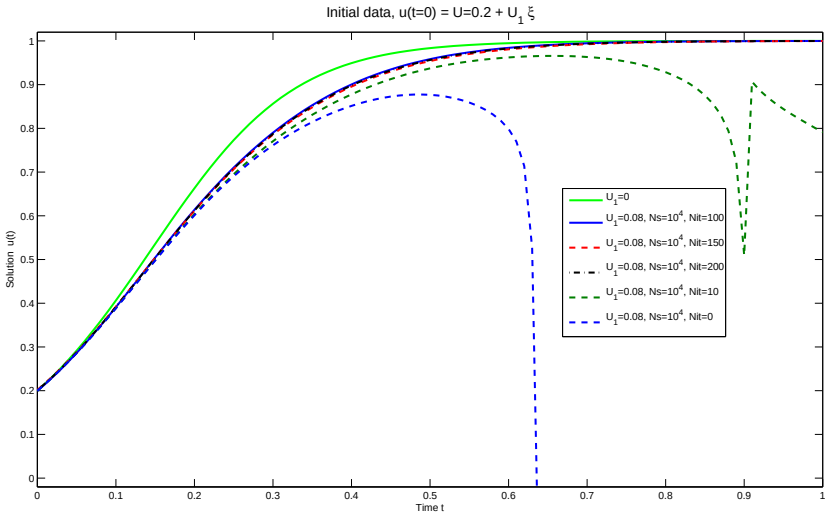

Improve Convergence of Stochastic ODE

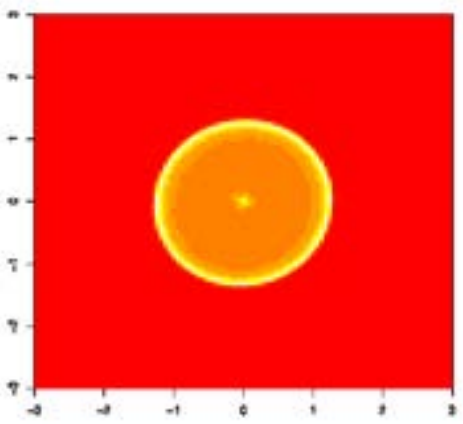

Generator for constrained populations

USC: Adaptation to Quantity of Interest (QoI)

R. Tipireddy, G. Ghanem

Expand $u(\xi)$: polynomials in $\eta = A\xi$

- with proper choice of A, the measure of the solution is concentrated along leading η_1 dimension
- A is chosen so that η_1 contains all Gaussian content of QoI

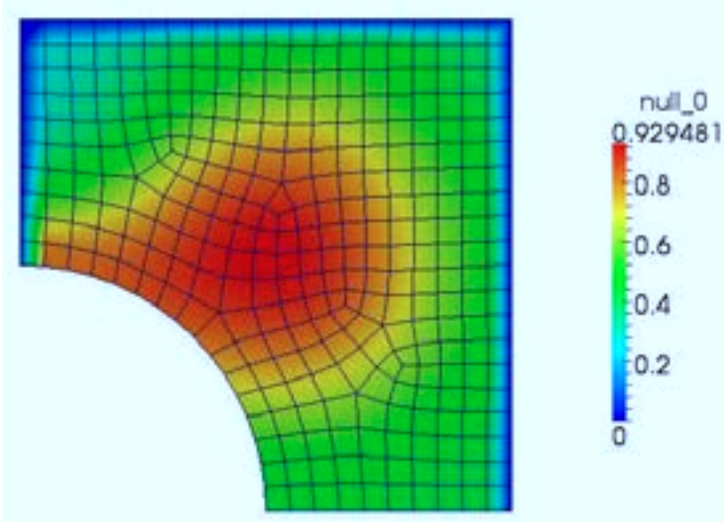

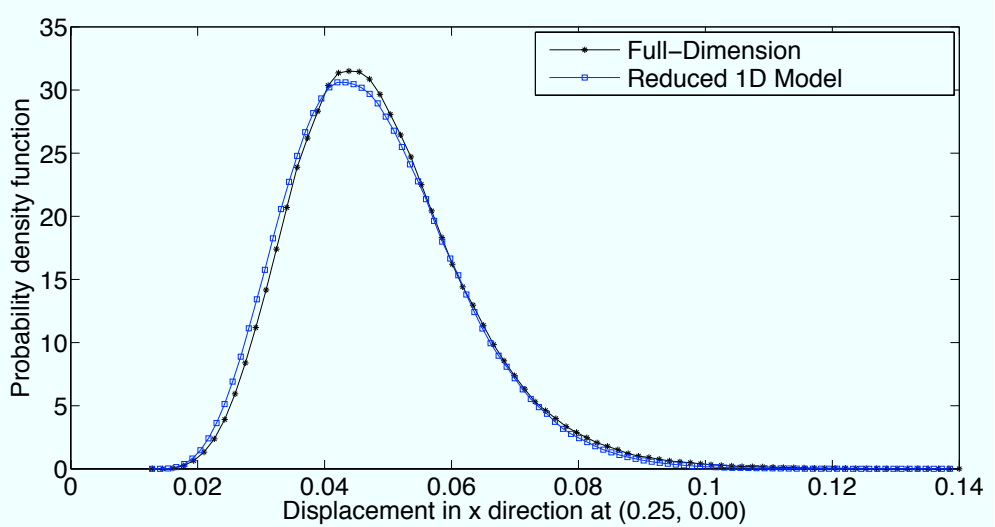

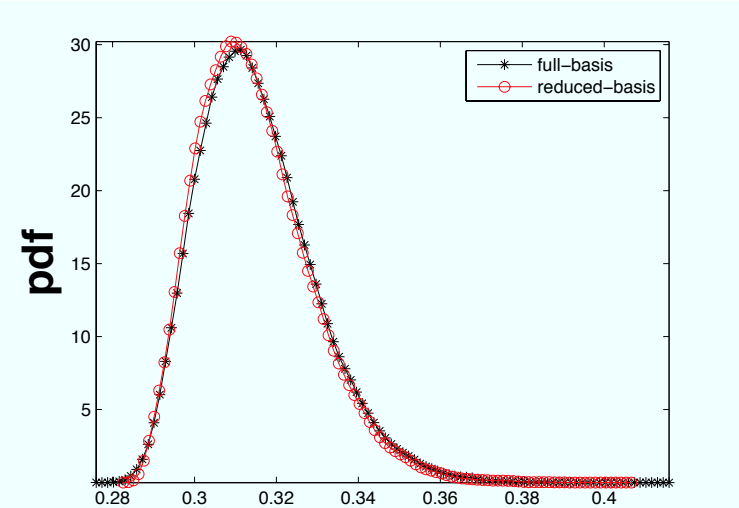

Von Mises stress at (0.2245,0.1101)

Closure

- Work on UQ software and algorithms development
  - Computational efficiency
  - Functionality, usability, scalability
  - Adaptivity, sparsity, preconditioning
  - Reduced-order, low-rank
  - Convergence, stability
  - Partial information, missing data
- *Robustifying* algorithms for large-scale applications
- Software integration well along the way
- Outreach via web, tutorials, and summer school
- SciDAC partnership activities getting off the ground