PECOS
Predictive Engineering and Computational Sciences

THE QUESO PARALLEL LIBRARY FOR UNCERTAINTY QUANTIFICATION IN LARGE-SCALE MODELS

Christopher Simmons, Robert Moser, Omar Ghattas

Introduction

The QUESO library is a collection of parallel statistical algorithms and object-oriented programming constructs supporting research into the uncertainty quantification of mathematical models.

- Quantification of Uncertainty for Estimation, Simulation and Optimization
- Available at https://github.com/libqueso/queso
- Development began in 2008 under the PSAAP program
- Since 2011, QUESO development has been part of the SCIDac3 program as part of the QUEST center
- QUESO is used by Dakota to solve the inverse problem
- QUESO has traditionally focused on the inverse problem but has forward propagation capabilities as well

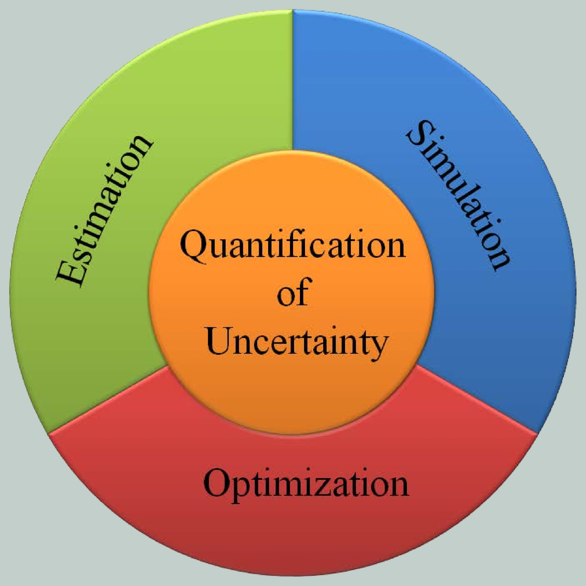

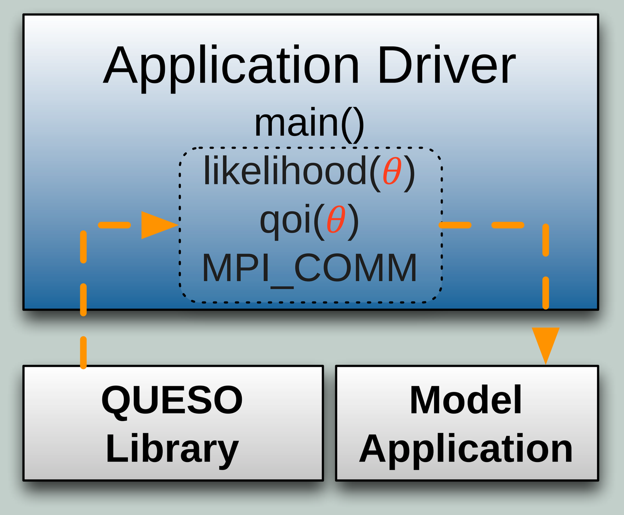

Why use QUESO?

- Other solutions available: R, PyMC, emcee, MICA, etc
- QUESO solves the same problem but has significantly more CS&E capabilities
  - Has been designed to be used with large forward problems
  - Has been used with over 10k cores
  - Support for finite and infinite dimensional problems
  - Can sample multimodal distributions
  - Can leverage Dakota for forward propagation (Dakota can use QUESO for the inverse problem)
  - Emulation capabilities being developed

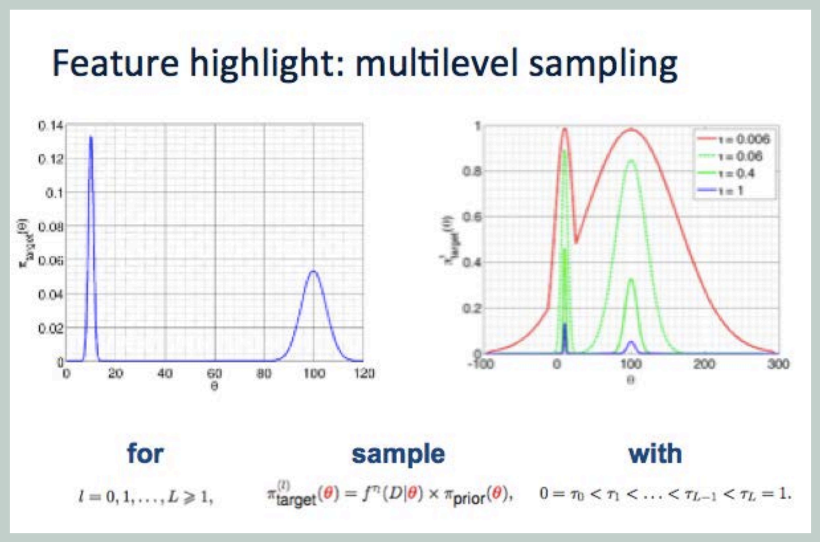

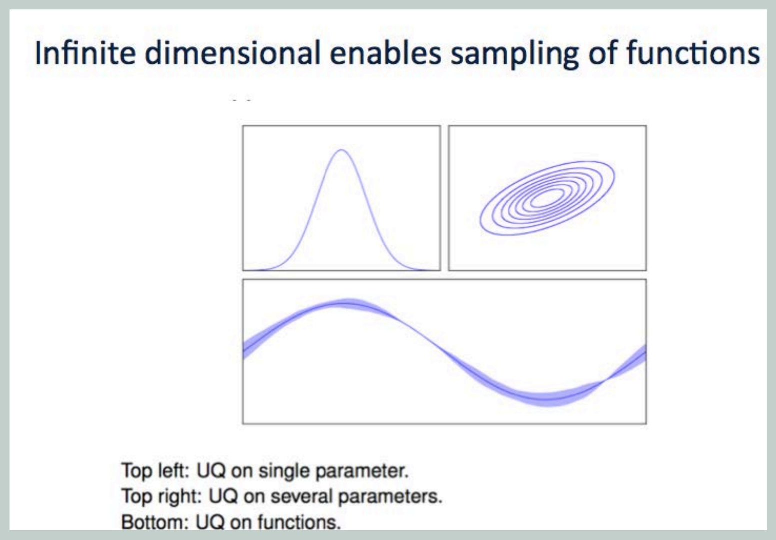

QUEST Impact on QUESO

- PAST
  - Documentation and testing
  - Trilinos integration
  - Dakota integration
  - New example problems
- PRESENT
  - GPMSA
  - More Dakota integration
  - Software quality and usability improvements
  - User community development
  - Infinite dimensional UQ
- FUTURE
  - Further emulation development
  - Continued software engineering improvement
  - Additional options for vector/matrix classes to increase user base
- Opportunity to be adopted as THE community code for uncertainty quantification

HPC Relevance

- Parallel Sampling
- Automatic load-balancing for homogenous systems
- Future load-balancing for heterogeneous systems
- Fault tolerance

Sponsors

- 2008 – 2014: DOE NNSA, PSAAP Program
- 2010 – 2011: DOE SNL-NM, Peridynamics Program
- 2010 – 2012: KAUST, AEA2 Program
- 2011 – 2013: AFOSR, RTC, DDDAS Program
- 2011 – 2015: DOE SC, SciDAC3 Program
- 2012 – 2013: DOE LANL and ORNL, CASL Project
- 2012 – 2014: KAUST, AEA3 Program

Acknowledgements

- E. Prudencio: original author of QUESO
- D. McDougall, K. Schulz, K. Estacio-Hiroms and the rest of the PECOS development team
- DOE laboratories: B. Adams (SNL), M. Eldred (SNL), J. Gattiker (LANL), D. Higdon (LANL), L. P. Swiler (SNL), B. Williams (LANL).

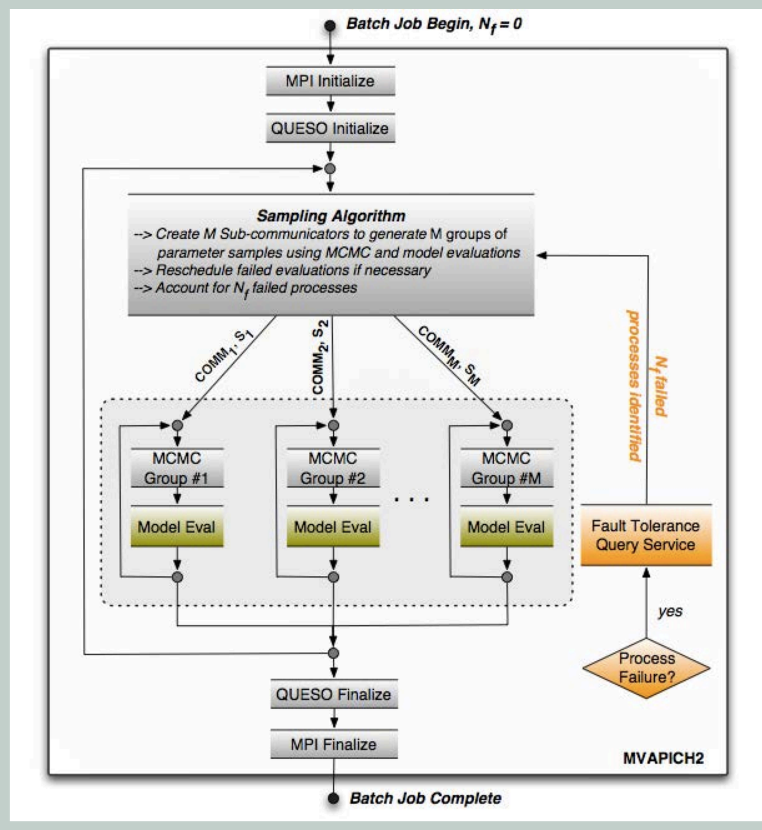

E. E. Prudencio and K. W. Schulz. The parallel C++ statistical library QUESO: Quantification of Uncertainty for Estimation, Simulation and Optimization. In M. Alexander et al., editors, Euro-Par 2011 Workshops, Part I, volume 7155 of Lecture Notes in Computer Science, pages 398-407. Springer-Verlag, Berlin Heidelberg, 2012.

New UQ algorithmic developments that are planned for future incorporation into QUESO

Ultimate goal: Achieve total scalability for forward and inverse UQ analysis, defined by making the cost of UQ independent of state, parameter, and data dimensions and number of processor cores.

Key tools to ensure scalability of UQ computations

- Derive all algorithmic steps at the infinite-dimensional level
- Discretize (in a consistent way)
- Linear the parameter-to-observable map and the parameter-to-predict map
- Independence of state, parameter, and data dimensions provably follows when using the correct algorithms, as described below
- Independence of cost from the number of processor cores follows when a scalable forward PDE solver is employed.

Inverse Problem

Given (possibly) noisy data, a model, and a prior distribution on parameters, infer the posterior distribution of the parameters.

- First, compute the maximum a posteriori (MAP) point by solving a deterministic optimization problem that is equivalent to a particular Tikhonov-regularized, weighted inverse problem. Use of an inexact-Newton-conjugate gradient method, preconditioned by the prior covariance, in conjunction with adjoint-based gradients and Hessian-vector products, guarantees that the cost (measured in forward PDE solves) is independent of the state, parameter, and data dimensions.
- Second, construct the inverse of the Hessian (of the negative log posterior) at the MAP point. This is made tractable by a low rank approximation (computed by randomized SVD) of the (prior-preconditioned) Hessian of the data misfit term, combined with a Sherman-Morrison-Woodbury inverse. Compactness of the data misfit Hessian, the smoothing property of the prior, and the properties of the randomized SVD algorithm guarantee that the inverse of the Hessian can be constructed at a cost (measured in forward PDE solves) that is independent of the state, parameter, and data dimensions.
- Linearizing the parameter-to-observable map, the posterior is then given by a Gaussian with mean equal to the MAP point, and covariance equal to the inverse Hessian. The overall cost of inverse UQ analysis depends only on the *true information content* of the data, measured by the number of modes of the parameter field that are informed by the data, i.e., the dimension of the dominant spectrum of the prior-preconditioned data misfit Hessian, which is typically much less than (and independent of) the parameter and data dimensions.

Forward Problem

Given a model and input parameter distribution, propagate uncertain parameters through the model to yield distributions of prediction quantities of interest.

- Linearizing the parameter-to-prediction map, the mean of the prediction QoI distribution is given by the mean of the input parameter distribution; the covariance of the prediction QoI is given by the push-forward of the parameter covariance (compactly represented through a low rank approximation) through the linearized parameter-to-prediction map.
- Computing the covariance of the prediction QoI distribution requires solution of an additional adjoint PDE for each QoI; the source term of the adjoint PDE is given by the derivative of the QoI with respect to the state.
- The overall cost of forward UQ analysis depends only on the *number of QoIs*, which of course is independent of the parameter and data dimensions.

Long term goal: Can we achieve scalability (defined above) without invoking the linearization approximations?

Illustrative application to Antarctic ice sheet flow

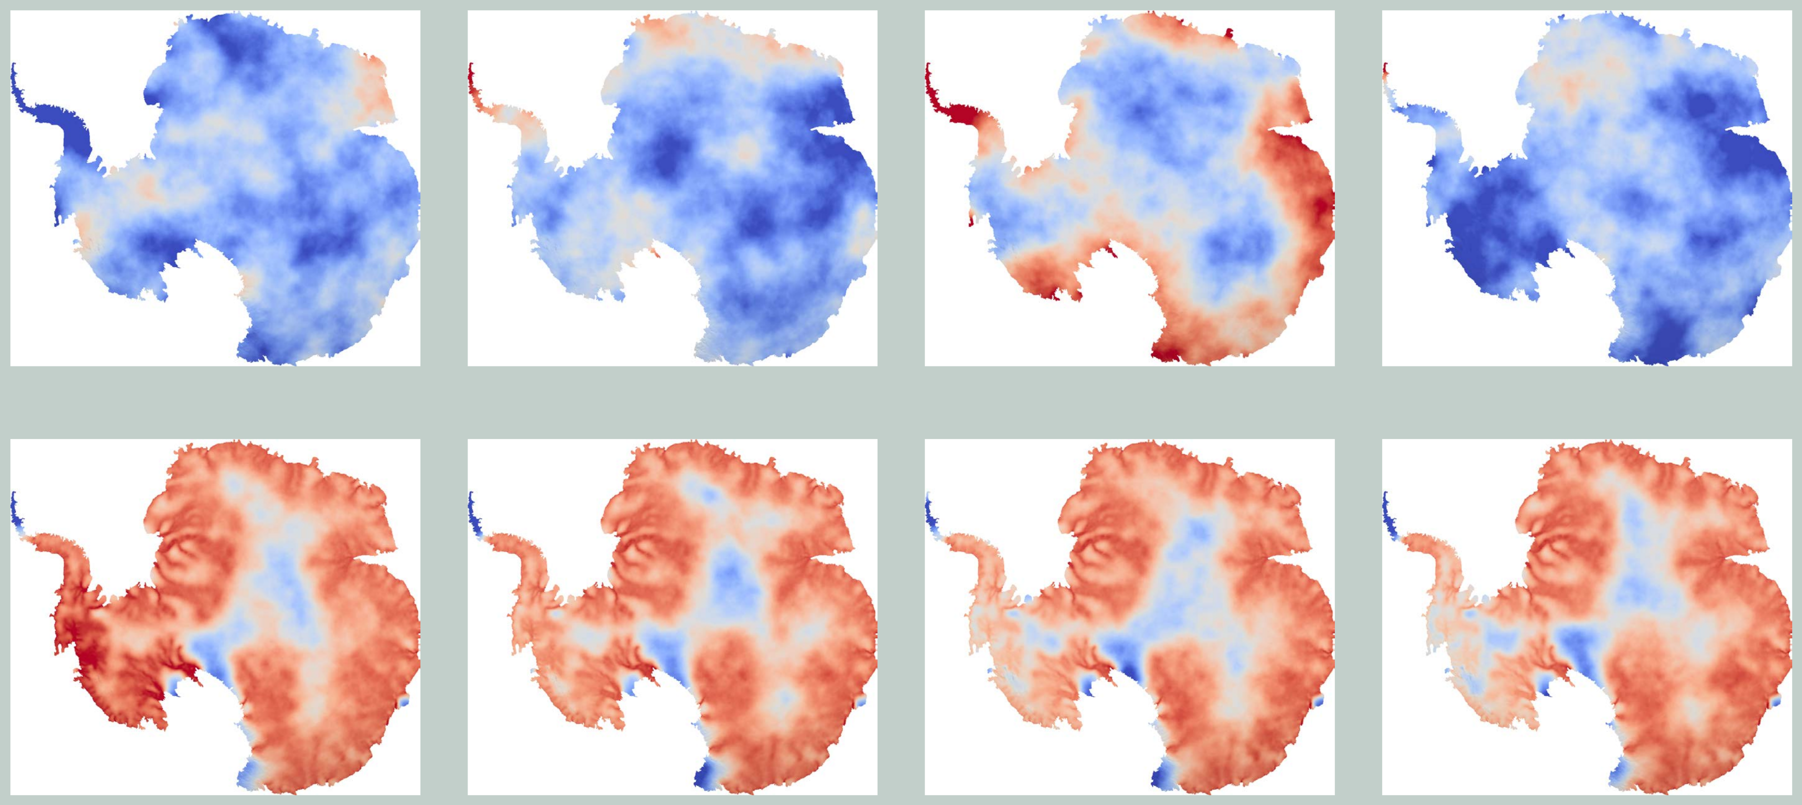

Inference of the basal friction field at the base of the ice sheet from satellite (InSAR) observations of the surface ice flow velocity Top: samples from the prior Bottom: samples from the posterior.

- Basal friction features beneath ice streams persist across posterior samples, indicating low variance in their inference
- Large variance is exhibited in samples of basal friction in interior of continent
- Gaussianized solution of Bayesian inverse problem computed in $O(10^4)$ forward solves, for $O(10^6)$ parameters and $O(10^6)$ observations.

Reference: T. Isaac, N. Petra, G. Stadler, and O. Ghattas, From data to prediction with quantified uncertainties: Scalable parallel algorithms and applications to the dynamics of the Antarctic ice sheet, in prep.

Above algorithms research funded primarily by DOE MMICCs and DOE Extreme Scale UQ with additional funding from AFOSR Comp Math, NSF ARC, and DOE SciDAC.