ER25914: Uncertainty Quantification for Large-Scale Ice Sheet Modeling
Final Report (9/01/09–8/31/14)

PI: Omar Ghattas
Institute for Computational Sciences and Engineering
The University of Texas at Austin

1. Project Summary

In the past several decades, satellite observations of the Antarctic ice sheet have revealed considerable acceleration of the flow and thinning of the ice in a number of outflow glaciers, due likely to the warming of the ocean. This has raised the concern that portions of the West Antarctic ice sheet could destabilize and collapse, leading to as much as a half meter or more of sea level rise by the end of the 21st century. There is now considerable interest among climate scientists in creating high fidelity models of the ice sheet flow, and coupling them to global climate models to simulate sea level rise under various warming scenarios. However, progress toward this goal has been stymied due to the large uncertainties in, and the significant complexities of, ice sheet models; indeed, previous IPCC reports left out ice sheet flow contributions to sea level rise due to these uncertainties and complexities.

The end goal of this project has been to create systematic, rigorous, and scalable algorithms for quantifying uncertainties in advanced, high-fidelity ice sheet models. These uncertainties reflect our incomplete knowledge of rheological laws, basal topography, and basal boundary conditions, including geothermal heat fluxes, and can be quantified by assimilating observational data into the model: this is fundamentally an inverse problem. Bayesian inference provides a systematic framework for incorporating uncertainties in observations, models, and prior knowledge to quantify the uncertainties in the model parameters. However, solution of this Bayesian inverse problem for expensive forward models (such as advanced ice sheet models) and for large numbers of parameters (as occurs when spatially-varying fields such as the basal boundary condition are discretized) presents significant computational difficulties.

To this end, we have developed a new generation parallel, adaptive-mesh, adjoint-enabled, full-Stokes ice sheet simulator capable of scaling to petascale systems with $O(10^5)$ cores. By augmenting the forward ice sheet simulator with adjoint-based gradient and Hessian information, our code provides capabilities for conducting sensitivity analysis, finding inverse solutions, and quantifying uncertainties at a cost—measured in linearized Stokes solves—that is independent of the parameter dimension, which can be in the million for our target problems. The challenges to achieving this goal include: (i) the wide range of length scales present, from $O(10^2)$ meters grid resolution necessary in transition zones to $O(10^6)$ meters in overall dimensions; (ii) the highly nonlinear, anisotropic, temperature- and strain-rate-dependent rheologies characterizing polar ice sheets; (iii) the large-scale and nonlinear nature of the inverse problems faced when inferring uncertain ice sheet model parameters from rapidly-growing volumes of observational data; and (iv) the need to rigorously quantify uncertainties in the parameter inversion and propagate those uncertainties into predictions of quantities of interest, such as ice mass flux into the ocean. Specific research issues and tasks that had to be addressed to overcome these challenges included:

- Creation of a state-of-the-art 3D full-Stokes ice sheet simulator that incorporates the full 3D mass and momentum balance equations with nonlinear shear-thinning rheology, mesh adaptivity, error estimation, high-order and mass-conservative discretization, inexact Newton-Krylov nonlinear solver, and linear-complexity parallel multilevel linear solver capable of scaling to systems with at least $O(10^5)$ cores. Adaptive mesh refinement is required to resolve flow transitions and high velocity gradient regions. The linear systems that arise at each Newton step are extremely ill-conditioned, due to the adaptive mesh refinement, high-aspect ratio domain (3 orders of magnitude), highly-variable effective viscosity (3 orders of magnitude), and highly-variable basal friction (9 orders of magnitude).

- Endowment of the ice sheet simulator with full adjoint solution capabilities, to permit rigorous error estimation for adaptive mesh refinement, fast sensitivity analysis, and inexpensive gradients and Hessians for use in inverse algorithms and uncertainty quantification. One challenge is that a 4th order anisotropic tensor viscosity emerges in the adjoint Stokes operator, despite the act that the forward nonlinear Stokes operator has a scalar and isotropic viscosity. This is due to the form of the rheological law (in particular, the dependence of the viscosity on strain rate).

- Development of scalable inexact matrix-free Newton-CG methods for solution of deterministic ice sheet inverse problems, capitalizing on adjoint solutions to compute gradient information and Hessian-vector products at little additional cost beyond the forward solution. Each Hessian-vector product computation requires solution of a pair of linearized forward/adjoint equations. These methods represent the state-of-the-art in optimization and offer a powerful way to globalize the (nonlinear) inverse iterations (which need to be globalized due to the nonconvexity of the problem) as well as truncated CG iterations in early Newton steps to prevent "oversolving."

- Development of scalable Bayesian inversion methods for uncertainty quantification of ice sheet model parameters from data, by building on our deterministic ice sheet inversion methods that exploit adjoint information (gradients, Hessians) to directly speed up MCMC sampling methods by employing a local Gaussian approximation to serve as a proposal for MCMC. A critical issue is that the inverse of the Hessian (of the parameter to observable map) shows up as the covariance of this local Gaussian, and straightforward construction of this Hessian is intractable (every column requires a forward Stokes solve). Instead, me must build a low rank approximation of the data misfit Hessian (motivated by the ill-posedness of the inverse problem, i.e., the data inform just a low-dimensional subspace of the parameter space), and combine that with a Sherman-Morrison-Woodbury inverse of the total Hessian to express the inverse Hessian at a cost (measured in linearized Stokes solves) that scales only with the amount of information contained in the data—not the with data nor parameter directions.

- Creation of algorithms and implementations for automated generation of high-quality coarse hexahedral meshes of polar ice sheets from available ice sheet geometry datasets. These course meshes serve as initial meshes to be refined by AMR.

- Employment of the deterministic and Bayesian inversion methods described above to infer uncertain basal friction, ice rheology, and geothermal heat flux fields that are consistent with satellite observational data.

The research issues listed above constitute a grand challenge of the highest order, involving many of the frontier mathematical and computational difficulties encountered in computational science today. However, these challenges must be attacked and overcome in order to rigorously assess and quantify the uncertainties in contributions of ice sheets to sea level rise in the coming years. Over the course of our project, we have systematically addressed each of these issues, in the process developing the most scalable, most efficient, and highest fidelity ice sheet model today, and the first capable of rigorous uncertainty quantification at large scale. The next section summarizes our specific contributions.

2. Specific Contributions

The sections below summarize our technical contributions over the course of the project. These are organized under three main thrusts: (1) **forward solver:** a new state-of-the-art parallel adaptive 3D nonlinear full Stokes ice sheet flow simulator; (2) **inverse solver:** a new adjoint-based inexact Newton method for deterministic inverse problems governed by the 3D nonlinear full Stokes ice flow model; and (3) **uncertainty quantification:** a novel Hessian-based Bayesian method for quantifying uncertainties in the inverse solution and propagating them forward into predictions of quantities of interest.

2.1. The forward solver: a parallel adaptive 3D nonlinear full-Stokes ice sheet flow model

We have developed a parallel adaptive 3D nonlinear full Stokes ice sheet dynamics simulator that incorporates state-of-the-art numerical algorithms. Its components are detailed below.

- **Meshing of full continental-scale ice sheet geometries.** We developed an automated software tool for generating meshes of polar ice sheets from the SeaRISE dataset. This dataset includes data on bed topography, thickness, ice surface elevation, surface temperature, accumulation, basal heat flux, and balance velocity. From the geometry (topography, thickness, elevation) data, we infer the lateral boundaries of the ice sheet, with an option to include ice shelves. From the lateral boundaries, we generate a 2D triangular mesh, with bounds on triangle size to satisfy aspect ratio restrictions. We preprocess this mesh by collapsing opposing triangles into quadrilaterals where possible; remaining triangles are then split into three quads, resulting in an all-quad mesh. From this 2D quad mesh, we then extrude vertically to create a 3D hexahedral mesh, which serves as the initial coarse mesh for ice sheet simulations. This mesh can then be further adapted to capture features that require finer resolution, such as near grounding lines or regions of high velocity or temperature gradients. The basis for adaptation is to treat each hexahedron as a zero-level octree, which can be refined to an arbitrary number of levels. For example, our initial mesh of Antarctica contains 47,000 hexahedra of size 25 km or less, as dictated by geometry. Further refinement adapts the mesh in the vicinity of the grounding line to a resolution of about 2 km. This serves as the initial coarse mesh, which can be refined dynamically.

- **High-order accurate, mass-conserving finite element discretization.** We have developed a high-order-accurate finite element discretization of the full Stokes equations. Our previous discretization was based on trilinear finite elements for both velocity and pressure, which is a popular choice for incompressible flows due to simplicity of implementation and ease of preconditioning. However the equal order of approximation for both velocity and pressure does not guarantee satisfaction of the LBB stability condition; thus, some form of pressure stabilization is usually enforced to ensure stability. Unfortunately, pressure stabilization leads to a loss of local mass conservation, which can have a deleterious effect on the accuracy of the solver, particularly for non-Newtonian flows such as ice sheets. Taylor-Hood elements (quadratic in velocity, linear in pressure) do satisfy LBB stability; however, while they improve on equal-order elements, they still do not conserve mass at the element level. Therefore we developed a discretization that uses continuous hexahedral elements of arbitrarily high order for the velocity, coupled with a discontinuous pressure approximation of two orders lower than the velocity approximation. Not only is this high-order element pair significantly more accurate per degree of freedom, it is also guaranteed to be stable, is more robust to element deformation, and most importantly, guarantees discrete element-wise mass conservation, an important property for simulation over long times of the ice sheet. While the *approximation* properties of this element pair are superior to standard discretizations, a significant challenge is to design *solvers* that can deliver optimal performance for such a complex discretization. Our work on designing scalable solvers is discussed next.

- **Improved parallel scalable linear solver and preconditioner.** To go along with our new discretization of the full Stokes equations, we designed a new solver and preconditioner for the linearized variable-viscosity Stokes equations that arise at each Newton iteration of our ice sheet simulator. First, we implemented a block upper triangular preconditioner, consisting of the viscous operator in the (1,1) block, the discrete gradient in the (1,2) block, and a Schur complement approximation in the (2,2) block (more about this below). To "invert" this preconditioner, we need to invert the (1,1) viscous block and the (2,2) Schur complement. The standard approach of approximating the inverse of the (1,1) viscous block with one V-cycle of algebraic multigrid (AMG) does not work with our high order finite element approximations. Instead, we partition the nodes of each high order element into a grid of trilinear hexahedra, and then employ a custom parallel smoothed aggregation algebraic multigrid with anisotropic smoothing to approximate the inverse of this low order discretization. We have contributed software for our custom AMG with anisotropic smoothing to PETSc so that others who solve PDEs on high aspect ratio 2.5D domains can benefit from our method. This approach yields a spectrally-equivalent approximation, and thereby retains the optimality of multigrid—i.e., the number of iterations needed for convergence is independent of problem size. The final component of the preconditioner is the approximation of the Schur complement operator in the (2,2) block. Testing revealed that the performance of our previous "inverse viscosity-scaled mass matrix" approximation of the Schur complement, which had been excellent for blocky domains, deteriorated for realistic ice sheet geometries, which are very shallow. The reason for this is that the Schur complement is no longer spectrally equivalent to a mass matrix, since the basal friction boundary condition provides a boundary term in the Schur complement that effectively acts like a stiffness matrix. For thin geometries, this can dominate. Therefore, we needed a more sophisticated preconditioner. We turned to a so-called BFB^T preconditioner, which approximates the action of the inverse of the Schur complement by forming pseudo-inverses of the discrete div and grad operators. Testing revealed essentially optimal algorithmic and parallel performance for realistic ice sheet geometries (including full continental Antarctic meshes) and basal boundary conditions: convergence independent of mesh refinement, element aspect ratio, and occurrence of nonconforming element faces, and only very weak dependence on finite element polynomial order and core count. With our parallel ice flow solver, we were able to carry out solutions of full Stokes ice sheet flow at the scale of the entire Antarctic continent, which is a first in the field of glaciology (the full Stokes model is the highest fidelity model available for ice flow). Finally, our nonlinear Stokes high-order-accurate parallel adaptive discretization and solver framework is at the core of the implicit solver presented in our SC15 paper *An Extreme-Scale Implicit Solver for Complex PDEs: Highly Heterogeneous Flow in Earths Mantle*. This solver weak-scaled with 97% parallel efficiency to the 1.6 million cores of the Sequoia BG/Q system, making it (by far) the most scalable implicit solver today. For this is received the 2015 Gordon Bell Prize.

- **Parallel algorithms for adaptive mesh refinement.** Many problems in science and engineering modeled by partial differential equations are characterized by a wide range of spatial scales and strong localization in the solutions of the PDEs. For such problems, of which continental ice sheet flow is an outstanding example, adaptive mesh refinement and coarsening (AMR) is essential and can reduce the number of unknowns by several orders of magnitude and more. Unfortunately, parallel AMR methods impose significant overhead, in particular on highly parallel systems, due to their need for frequent readaptation and repartitioning of the mesh over the course of the simulation. Because of the complex data structures and frequent communication and load balancing, scaling dynamic AMR to petascale systems has long been considered a challenge. To meet this challenge, we developed the *p4est* software library, which features an octree-based AMR method that makes use of recursive encoding schemes to enable hierarchical refinement, along with Morton-coded space filling curves to manage adaptivity in parallel. To represent general geometries, it merges multiple adaptive octrees to create a "forest" of octrees. *p4est* is based on new parallel scalable algorithms that strictly adhere to fully distributed storage. It extends the scalability of single-octree algorithms to general geometries and arbitrarily-high order discretizations, and is the most scalable AMR method and software library available today (with demonstrated excellent scalability to the 1.6 million cores of LLNL's Sequoia BG/Q system). We designed extensions of *p4est* to 2.5D geometries, as occur in many geophysical flows such as ice sheets, oceans, and the atmosphere, and in plate and shell structural mechanics. The library was released in open source form (www.p4est.org) and has also been adopted by the most widely-used finite element software library, *deal.II*, as its AMR engine. Moreover, *PETSc*'s unstructured mesh format has been extended to interface with *p4est*'s hanging-node data structure, thereby opening the way for *PETSc*'s huge base of users to benefit from *p4est*'s parallel AMR capabilities. The significance of the advances in the state-of-the art of parallel AMR offered by *p4est* was recognized by the selection of the SC10 paper *Extreme-scale AMR* as a 2010 Gordon Bell Prize Finalist. Moreover, p4est is the underlying AMR library for the implicit solver presented in the SC15 paper *An Extreme-Scale Implicit Solver for Complex PDEs: Highly Heterogeneous Flow in Earth's Mantle*, which was awarded the 2015 Gordon Bell Prize, as mentioned above. Finally, it received the Best Poster Award at SC09 and SC14, and powered the plate tectonics adaptive simulations that were featured on the cover of *Science* in the August 27, 2010 issue.

In conclusion, we believe that our solver described above is the most advanced (parallel, scalable, efficient, high-order-accurate, adaptive, high fidelity) ice sheet flow solver available today.

2.2. The inverse problem: Adjoint- and Newton-based inversion methods for inferring basal friction, uncertain rheology, and geothermal heat flux in 3D full Stokes ice sheet flow

In this project we addressed inverse ice sheet flow problems governed by the forward model described above. In the "forward" ice flow problem, all inputs (boundary conditions, rheology parameters, geometry, initial conditions) to the ice sheet model are specified, and the model is run to predict the outputs (velocities and pressures) throughout the domain. However, it is rarely the case that all inputs are known, and this is especially the case for polar ice sheet models such as those of Antarctica. In particular, the rheology parameters (i.e., those characterizing Glen's flow law) are uncertain because laboratory ice (on which experiments to infer these parameters are commonly conducted) differs from field ice; moreover, the basal boundary conditions (parameters characterizing the friction law, bed topography, and geothermal heat flux) are uncertain, because they cannot be observed directly. This uncertainty in parameters describing the basal boundary conditions and rheology is perhaps the greatest barrier to progress in modeling polar ice sheet flow and projection of consequent sea level rise.

To address this problem, we have formulated several inverse problems for the uncertain parameters: the goal is to estimate the rheology parameters (specifically, the spatially-varying parameter field in Glen's law) and basal boundary condition parameters (specifically, the basal "friction" field characterizing the tradeoff between no-slip and full-slip in the basal boundary condition as well as the geothermal heat flux) that minimize the L_2 norm misfit of model predictions with measurements of surface velocity. This is in general an "ill-posed" inverse problem, since multiple sets of parameters are consistent with the data, which are typically very sparse. We have overcome this ill-posedness by judicious choice of a "regularization" functional, which penalizes unwanted features in the parameter fields.

Using the method of Lagrange multipliers, we impose the governing forward full Stokes ice flow equations as constraints on the inversion, both for a linear and a nonlinear strain-rate-dependent rheology, and for linear and nonlinear basal slip boundary conditions. Using tools from variational calculus, infinite-dimensional optimality conditions are derived by seeking stationarity of the Lagrangian functional with respect to the Lagrange multipliers (also known as the adjoint variables), the state variables, and the inversion parameter fields (the Glen's law rheology parameter field, the basal friction parameter field, and the geothermal heat flux). This results in an optimality system consisting of the original forward ice sheet model; a so-called adjoint ice sheet model with similar structure to the forward model, except that its operator is the adjoint of the linearized forward model and it is driven by a source term stemming from the data misfit; and finally a partial differential equation written on the domain for the uncertain rheology parameter field, and written on the basal boundary for the uncertain basal friction and geothermal heat flux parameter fields. A significant challenge is that the adjoint Stokes flow problem contains a 4th order anisotropic tensor effective viscosity, despite the fact that the forward ice flow model's effective viscosity is isotropic and scalar; this is a consequence of the power law nature of the dependence of the viscosity on strain rate. The forward solver was extended to handle this difficulty. The discretizations of the adjoint velocity and pressure fields follow those of their forward counterparts, while the uncertain rheology and basal parameter fields are discretized following the velocity field or one order lower.

Solving the resulting discretized inverse problem constitutes a major challenge, since the inverse problem is high nonlinear and ill-posed. To this end, we have created a state-of-the-art numerical inversion method. Motivated by the mesh-independent convergence (and hence insensitivity to problem size) of Newton's method for a wide class of nonlinear operator equations, we have chosen to incorporate a Newton solver for the inversion, despite the more complex implementation required. Explicitly forming the Hessian matrix required by Newton's method at each inversion iteration is intractable—it would require as many forward linearized Stokes solves as there are parameters, which for our discretized inversion fields, number from the tens of thousands up to one million for full continental-scale inversions. Instead, we solve the Newton system using conjugate gradients, which requires only the action of the Hessian (of the parameter-to-observable map) on a given vector at each CG iteration. This in turn can be formed matrix-free and on-the-fly, by solving one forward and one adjoint linearized Stokes-like problem. Preconditioning the Hessian by the regularization operator results in a "compact perturbation of the identity" structure (similar to a Fredholm second kind integral operator), which leads to mesh-independent convergence of conjugate gradients. Employing Eisenstat-Walker type inexactness in the CG iterations (i.e., the iterations are aggressively truncated when we are far from the solution of the inverse problem to prevent "oversolving," without spoiling the asymptotic quadratic convergence rate), we typically find that the average number of CG solves per Newton iteration—and thus forward/adjoint Stokes solves—is kept small, often to a handful. Finally, we globalize Newton's method with a backtracking Armijo line search.

To test our Newton-based inversion method against contemporary inverse methods and examine its effectiveness and scalability, we created a series of synthetic observed surface velocity datasets and inverted for uncertain basal and rheology parameter fields from these observations. The numerical inversion tests show that with our inexact adjoint-based Newton-CG method, the cost of solving the inverse problem (measured in number of forward and adjoint Stokes solves) is independent of the number of inversion parameters. This is an important finding for uncertainty quantification in ice sheet inverse problems, where the uncertain parameters are discretized fields and thus of very high dimension. To compare against the most popular method used in geophysical inverse problems, we implemented the nonlinear conjugate gradient method and compared its performance to our inexact Newton-CG method. The comparison shows that our inexact Newton method is an order of magnitude more efficient than the nonlinear CG method for moderately sized problems (there is no comparison in performance for large-scale problems; Newton is the only feasible choice as nonlinear CG takes an unacceptably large number of iterations). Finally, numerical studies show that our regularized inverse method is capable of accurately reconstructing smooth features of the basal friction parameter field as well as the Glen's law rheology exponent parameter.

Employing the Newton-based inversion capability described above, we achieved a milestone in ice sheet flow inverse modeling by successfully carrying out the first full Antarctic inversion for basal friction using the highest fidelity model available, the 3D full Stokes model as implemented in our forward ice flow solver. The inversion employs surface velocity observational data, both synthetic (generated from a known basal friction field this is used to assess the accuracy of the inversion) and real (obtained from satellite interferometric synthetic aperture radar). The inverse problem is solved for the basal friction field, an uncertain coefficient in a Robin boundary condition that balances tangential traction with tangential velocity at the base of the ice sheet. This field is discretized on the same mesh as the velocity and pressure fields. To solve the inverse problem, we employ H^1 Tikhonov regularization as well as total variation regularization, along with the inexact Newton CG method described above, which entails computing the action of the Hessian on a vector at each CG iteration by solving a pair of linearized forward/adjoint Stokes problems, and a regularization-based preconditioner. Experiments with noisy synthetic data for an inverse problem characterized by 4M velocity/pressure unknowns and 400,000 basal friction parameters indicate convergence in a modest number of inner iterations (85) and outer (18) iterations, and the ability to reconstruct the basic structure of the basal friction field. Inversions with InSAR data (observations of the ice sheet surface velocity field) produce a map of the basal friction field that displays low basal friction deep into the interior of the ice sheet, particularly in West Antarctica, indicating potential weakness of the ice sheet in these regions.

2.3. Uncertainty quantification and Bayesian inverse problems

Building on our (deterministic) inversion methods described above, we created scalable algorithms for estimating uncertainties in the solution of ice sheet inverse problems (in particular, inversion for unknown basal friction from observations of surface velocities) within the framework of Bayesian inference. The standard approach to exploring the posterior probability distribution function (pdf) is based on sampling using a Markov chain Monte Carlo (MCMC) method. The use of conventional MCMC methods becomes intractable for large-scale ice sheet problems due to the expense of solving the forward model (3D Stokes flow with nonlinear rheology) and the high dimensionality of the uncertain parameters (the friction in the basal boundary condition is a field, which when discretized leads to as many as one million parameters). For such high-dimensional parameter spaces, MCMC sampling methods can have extremely low acceptance probabilities, resulting in slow convergence. The challenge for MCMC methods is to construct proposal distributions that simultaneously provide a good representation of the target density (hence increasing the acceptance rate) while being inexpensive to construct and to sample from.

To overcome this challenge, we developed a Gaussian proposal for MCMC that adapts to the local pdf curvature by taking the covariance matrix to be the inverse Hessian of the negative log posterior, which amounts to the regularized data misfit functional. Thus the method is analogous to a Newton method for optimization, but with a stochastic component. This not only allows us to exploit the connections between deterministic optimization methods for minimizing the cost functional and statistical methods for sampling the posterior, but also allows us to guide the samples using problem structure. To construct the local Gaussian approximation efficiently, we exploit the rapidly decaying spectral property of the Hessian of the data misfit, which stems from the fact that the observational data typically provide only sparse information on model parameters (due to the smoothing properties of the parameter-to-observable map). This allows us to construct a low rank approximation (using matrix-free randomized SVD methods), enabling the algorithm to draw a sample at a cost independent of the parameter and data dimensions. For the largest problems (with one million basal friction field inversion parameters and full Antarctica geometry), even this low-rank approximation (which leads to an effective compression of the parameter dimension by three orders of magnitude) is not sufficient to permit repeated evaluation of the Hessian at each sample point; instead, we freeze the Hessian and its low-rank-based covariance approximation at the maximum a posteriori (MAP) point, and use that as a basis for sampling, or even simply construct a Laplace approximation (a Gaussian based on this covariance at the MAP point) to characterize uncertainty around the MAP. Finally, having inferred the basal friction with quantified uncertainty, we propagate the mean and variance of the basal friction through a linearized parameter-to-prediction map to estimate the mean and variance of the ice mass flux from Antarctica to the ocean.

Our JCP paper reporting this research (published in 2015) shows that the work required for executing this data-to-prediction-with-quantified-uncertainties process (again, measured in the number of linearized forward and adjoint ice sheet model solves) is independent of the state dimension, parameter dimension, data dimension, and number of processor cores. As testament to the significance of the inversion and uncertainty quantification results, we cite comments from the reviews of the JCP paper. One reviewer wrote:

> *Overall, I think this is a fantastic, unique, and very important paper. It is well motivated, well organized, clearly written, and presents (again, to my knowledge) novel results not only within the ice sheet modeling community, but probably within the geophysical and/or earth system modeling communities in general... While there are still further advances to be made, it represents a singular achievement and I believe that it will become a seminal paper to reference in the field of large-scale, geophysical model inversion and uncertainty quantification (both currently "bleeding edge" topics of interest in large scale, computationally intensive geophysical and earth system model applications)."*

A second reviewer wrote:

> *The manuscript represents in my view a very important contribution to the topic of UQ in ice sheet prediction. The approach developed is comprehensive, rigorous, and is urgently needed to put UQ in climate prediction on a more solid basis.*

3. UT Personnel Involved:

- Don Blankenship (research scientist, UTIG)
- Tan Bui-Thanh (research scientist, ICES; currently Assistant Professor, Aerospace Engineering & Engineering Mechanics, UT-Austin)
- Carsten Burstedde (research scientist, ICES; currently Associate Professor, Mathematics, University of Bonn)
- Omar Ghattas (professor, ICES and Departments of Geological Sciences and Mechanical Engineering)
- Charles Jackson (research scientist, UTIG)
- Tobin Isaac (Ph.D. student, ICES; currently postdoctoral researcher, University of Chicago)
- James Martin (Ph.D. student, ICES)
- Noemi Petra (postdoctoral researcher, ICES; currently Assistant Professor, Applied Mathematics, UC-Merced)
- Georg Stadler (research scientist, ICES; currently, Associate Professor, Courant Instite, NYU)
- Alice Zhu (Ph.D. student, ICES)

4. Publications Based on Work Conducted Under this Grant

- H. Zhu, N. Petra, G. Stadler, T. Isaac, T.J.R. Hughes, and O. Ghattas, Inversion of geothermal heat flux in a thermomechanically coupled nonlinear Stokes ice sheet model, submitted, 2015.
- J. Rudi, A.C.I. Malossi, T. Isaac, G. Stadler, M. Gurnis, P.W.J. Staar, Y. Ineichen, C. Bekas, A. Curioni, O. Ghattas, An Extreme-Scale Implicit Solver for Complex PDEs: Highly Heterogeneous Flow in Earth's Mantle, *Proceedings of IEEE/ACM SC15*, 2015. (2015 Gordon Bell Prize)
- T. Isaac, G. Stadler, and O. Ghattas, Solution of nonlinear Stokes equations discretized by high-order finite elements on nonconforming and anisotropic meshes, with application to ice sheet dynamics, *SIAM Journal on Scientific Computing*, 37(6):B804B833, 2015.
- T. Isaac, C. Burstedde, L.C. Wilcox, and O. Ghattas, Recursive algorithms for distributed forests of octrees, *SIAM Journal on Scientific Computing*, 37(5):C497–C531, 2015.
- T. Isaac, N. Petra, G. Stadler, O. Ghattas, Scalable and efficient algorithms for the propagation of uncertainty from data through inference to prediction for large-scale problems, with application to flow of the Antarctic ice sheet, *Journal of Computational Physics*, 296(1):348–368, 2015.
- T. Bui-Thanh and O. Ghattas, A scalable MAP solver for Bayesian inverse problems with Besov priors, *Inverse Problems and Imaging*, 9(1):27–54, 2015.
- N. Petra, J. Martin, G. Stadler, and O. Ghattas, A computational framework for infinite-dimensional Bayesian inverse problems. Part II: Stochastic Newton MCMC with application to ice sheet flow inverse problems, *SIAM Journal on Scientific Computing*, 36(4):A1525–A1555, 2014.
- T. Bui-Thanh, O. Ghattas, J. Martin, and G. Stadler, A computational framework for infinite-dimensional Bayesian inverse problems. Part I: The linearized case, with applications to global seismic inversion. *SIAM Journal on Scientific Computing*, 35(6):A2494-A2523, 2013.
- T. Bui-Thanh, C. Burstedde, O. Ghattas, J. Martin, G. Stadler, and L.C. Wilcox, Extreme-scale UQ for Bayesian inverse problems governed by PDEs, *Proceedings of IEEE/ACM SC12*, 2012. (2012 Gordon Bell Prize Finalist)
- H. Sundar, G. Biros, C. Burstedde, J. Rudi, O. Ghattas, G. Stadler, Parallel geometric-algebraic multigrid on unstructured forests of octrees, *Proceedings of IEEE/ACM SC12*, 2012.
- N. Petra, H. Zhu, G. Stadler, T.J.R. Hughes, O. Ghattas, An inexact Gauss-Newton method for inversion of basal sliding and rheology parameters in a nonlinear Stokes ice sheet model, *Journal of Glaciology*, 58(211):889–903, 2012.
- T. Bui-Thanh, O. Ghattas, and D. Higdon, Adaptive Hessian-based non-stationary Gaussian process response surface method for probability density approximation with application to Bayesian solution of large-scale inverse problems, *SIAM Journal on Scientific Computing*, 34(6):A2837–A2871, 2012.
- 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, *SIAM Journal on Scientific Computing*, 34(3):A1460-A1487, 2012.

- T. Isaac, C. Burstedde, and O. Ghattas, Low-Cost Parallel Algorithms for 2:1 Octree Balance, *International Parallel and Distributed Processing Symposium (IPDPS 12)*, 426–437, IEEE Computer Society, 2012.

- C. Burstedde, L.C. Wilcox, and O. Ghattas, p4est: Scalable algorithms for parallel adaptive mesh refinement on forests of octrees, *SIAM Journal on Scientific Computing*, 33(3):1103–1133, 2011.

- H.P. Flath, L.C. Wilcox, V. Akcelik, 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.

- C. Burstedde, O. Ghattas, M. Gurnis, T. Isaac, G. Stadler, T. Warburton, L.C. Wilcox, Extreme-Scale AMR, *Proceedings of ACM/IEEE SC10*, 2010. (2010 Gordon Bell Prize Finalist)

5. Awards and Honors (to PI Ghattas, unless otherwise noted)

- *2016 SIAM Early Career Prize in Supercomputing* will be awarded to former Ph.D. student Tobin Isaac (who developed the forward/inverse ice sheet simulator) at the 2016 SIAM Conference on Parallel Processing, April 2016, Paris, France. The prize comes with an invitation to deliver a plenary lecture.

- Awarded *2015 ACM Gordon Bell Prize* for the paper: J. Rudi, A.C.I. Malossi, T. Isaac, G. Stadler, M. Gurnis, P.W.J. Staar, Y. Ineichen, C. Bekas, A. Curioni, O. Ghattas, An Extreme-Scale Implicit Solver for Complex PDEs: Highly Heterogeneous Flow in Earth's Mantle, *Proceedings of IEEE/ACM SC15*, 2015.

- Invited to serve on *National Academies Committee on Models of the World*, 2015–2016.

- 2014 IEEE/ACM SC14 *Best Poster Award* for the poster *Parallel High-Order Geometric Multigrid Methods on Adaptive Meshes for Highly Heterogeneous Nonlinear Stokes Flow Simulations of Earth's Mantle*, Johann Rudi, Hari Sundar, Tobin Isaac, Georg Stadler, Michael Gurnis, and Omar Ghattas.

- *SIAM Fellow*, 2014. "For contributions to optimization of systems governed by partial differential equations and leadership to promote computational science and engineering."

- *2012 Joseph C. Walter Excellence Award*, Jackson School of Geosciences.

- Finalist, *2012 ACM Gordon Bell Prize*, for the paper T. Bui-Thanh, C. Burstedde, O. Ghattas, J. Martin, G. Stadler, and L.C. Wilcox, "Extreme-scale UQ for Bayesian inverse problems governed by PDEs," *Proceedings of ACM/IEEE SC12*, November 2012.

- Invited to serve on the National Research Council of the National Academies *Committee on Mathematical Foundations of Verification, Validation, and Uncertainty Quantification*, 2010–2012.

- Finalist, *2010 Gordon Bell Prize*, for the paper C. Burstedde, O. Ghattas, M. Gurnis, T. Isaac, G. Stadler, T. Warburton, L.C. Wilcox, "Extreme-Scale AMR," *Proceedings of ACM/IEEE SC10*, November 2010.

- Research on multiresolution supercomputing models of mantle flow and plate tectonics was featured on the cover of August 27, 2010 issue of *Science*.

- *SC09 Best Poster Award*, for the poster ALPS: A Framework for Parallel Adaptive PDE Solution, SC09, Portland, OR, 2009.
- Invited to deliver plenary or keynote lectures at the following international conferences:
  - VECPAR 2016, Porto, Portugal, June 27–30, 2016.
  - PRACE 2016, Prague, Czech Republic, May 12, 2016.
  - Joint Conference of the German Mathematical Society (DMV) and the International Association of Applied Mathematics and Mechanics (GAMM), Technical University of Braunschweig, Germany, March 7–11, 2016.
  - 24th International Meshing Roundtable, Austin, TX, October 12–14, 2015.
  - International Conference on Scientific Computing and Differential Equations (SciCADE 2015), Potsdam, Germany, September 14–18, 2015.
  - 52nd Meeting of the Society for Natural Philosophy, Rio de Janeiro, Brazil, October 22–24, 2014.
  - SIAM Annual Meeting, Chicago, IL, July 7–14, 2014.
  - PASC14 Conference, Zurich, Switzerland, June 2–3, 2014.
  - International Conference on Scientific Computing at Extreme Scale, Shanghai Jiao Tong University, Shanghai, China, May 7–9, 2014.
  - European Conference on Numerical Mathematics and Advanced Applications (ENUMATH) 2013, Lausanne, Switzerland, August 26–30, 2013.
  - MOPTA 2013, Bethlehem, PA, August 14–16, 2013.
  - US National Congress on Computational Mechanics (USNCCM) 2013, Raleigh, NC, July 22–25, 2013.
  - Uncertainty in Computer Models 2012, Sheffield, UK, July 2–4, 2012.
  - 10th International Conference on Mathematical and Numerical Aspects of Waves Propagation (WAVES 2011), Vancouver, Canada, July 25–29, 2011.
  - 2009 Iberian-Latin American Conference on Computational Methods in Engineering, Armacao de Buzios, Brazil, November 8–11, 2009.
  - 10th LCI International Conference on High-Performance Clustered Computing, Boulder, CO, March 10–12, 2009