Journal of Computational Physics 296 (2015) 348–368

Contents lists available at ScienceDirect

Journal of Computational Physics

www.elsevier.com/locate/jcp

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

Tobin Isaac [a], Noemi Petra [b,*], Georg Stadler [c], Omar Ghattas [a,d,e]

[a] Institute for Computational Engineering & Sciences, The University of Texas at Austin, Austin, TX, USA
[b] Applied Mathematics, School of Natural Sciences, University of California, Merced, CA, USA
[c] Courant Institute of Mathematical Sciences, New York University, New York, NY, USA
[d] Department of Mechanical Engineering, The University of Texas at Austin, Austin, TX, USA
[e] Jackson School of Geosciences, The University of Texas at Austin, Austin, TX, USA

ARTICLE INFO

Article history:
Received 6 October 2014
Received in revised form 21 April 2015
Accepted 30 April 2015
Available online 7 May 2015

Keywords:
Uncertainty quantification
Inverse problems
Bayesian inference
Data-to-prediction
Low-rank approximation
Adjoint-based Hessian
Nonlinear Stokes equations
Inexact Newton–Krylov method
Preconditioning
Ice sheet flow modeling
Antarctic ice sheet

ABSTRACT

The majority of research on efficient and scalable algorithms in computational science and engineering has focused on the *forward problem*: given parameter inputs, solve the governing equations to determine output quantities of interest. In contrast, here we consider the broader question: given a (large-scale) model containing uncertain parameters, (possibly) noisy observational data, and a prediction quantity of interest, how do we construct efficient and scalable algorithms to (1) infer the model parameters from the data (the *deterministic inverse problem*), (2) quantify the uncertainty in the inferred parameters (the *Bayesian inference problem*), and (3) propagate the resulting uncertain parameters through the model to issue predictions with quantified uncertainties (the *forward uncertainty propagation problem*)?

We present efficient and scalable algorithms for this end-to-end, data-to-prediction process under the Gaussian approximation and in the context of modeling the flow of the Antarctic ice sheet and its effect on loss of grounded ice to the ocean. The ice is modeled as a viscous, incompressible, creeping, shear-thinning fluid. The observational data come from satellite measurements of surface ice flow velocity, and the uncertain parameter field to be inferred is the basal sliding parameter, represented by a heterogeneous coefficient in a Robin boundary condition at the base of the ice sheet. The prediction quantity of interest is the present-day ice mass flux from the Antarctic continent to the ocean.

We show that the work required for executing this data-to-prediction process—measured in number of forward (and adjoint) ice sheet model solves—is independent of the state dimension, parameter dimension, data dimension, and the number of processor cores. The key to achieving this dimension independence is to exploit the fact that, despite their large size, the observational data typically provide only sparse information on model parameters. This property can be exploited to construct a low rank approximation of the linearized parameter-to-observable map via randomized SVD methods and adjoint-based actions of Hessians of the data misfit functional.

* Corresponding author.
E-mail addresses: tisaac@ices.utexas.edu (T. Isaac), npetra@ucmerced.edu (N. Petra), stadler@cims.nyu.edu (G. Stadler), omar@ices.utexas.edu (O. Ghattas).

http://dx.doi.org/10.1016/j.jcp.2015.04.047
0021-9991/

1. Introduction

The future mass balance of the polar ice sheets will be critical to climate in the coming century, yet there is much uncertainty surrounding even their current mass balance. The current rate of ice sheet mass loss was recently estimated at roughly 200 billion metric tons per year in [1] using data from various sources, including radar and laser altimetry, gravimetric observations, and surface mass balance calculations of regional climate models.[1] Moreover, this mass loss has been observed to be accelerating [2]. Driven by increased warming, collapse of even a small portion of one of these ice sheets has the potential to greatly accelerate this figure. Indeed, recent evidence suggests that sea level rose abruptly at the end of the last interglacial period (118,000 years ago) by 5–6 m; the likely cause is catastrophic collapse of an ice sheet driven by warming oceans [3]. Based on a conservative estimate of a half meter of sea level rise, the Organization for Economic Cooperation and Development estimates that the 136 largest port cities, with 150 million inhabitants and $35 trillion worth of assets, will be at risk from coastal flooding by 2070 [4].

Clearly, model-based projections of the evolution of the polar ice sheets will play a central role in anticipating future sea level rise. However, current ice sheet models are subject to considerable uncertainties. Indeed, ice sheet models were left out of the Intergovernmental Panel on Climate Change's 4th Assessment Report, which stated that "the uncertainty in the projections of the land ice contributions [to sea level rise] is dominated by the various uncertainties in the land ice models themselves … rather than in the temperature projections" [5].

Ice is modeled as a creeping, viscous, incompressible, shear-thinning fluid with strain-rate- and temperature-dependent viscosity. Severe mathematical and computational challenges place significant barriers on improving predictability of ice sheet flow models. These include complex and very high-aspect ratio (thin) geometry, highly nonlinear and anisotropic rheology, extremely ill-conditioned linear and nonlinear algebraic systems that arise upon discretization as a result of heterogeneous, widely-varying viscosity and basal sliding parameters, a broad range of relevant length scales (tens of meters to thousands of kilometers), localization phenomena including fracture, and complex sub-basal hydrological processes.

However, the greatest mathematical and computational challenges lie in quantifying the uncertainties in the predictions of the ice sheet models. These models are characterized by unknown or poorly constrained fields describing the basal sliding parameter (resistance to sliding at the base of the ice sheet), basal topography, geothermal heat flux, and rheology. While many of these parameter fields cannot be directly observed, they can be inferred from satellite observations, such as those of ice surface velocities or ice thickness, which leads to a severely ill-posed inverse problem whose solution is extremely challenging. Quantifying the uncertainties that result from inference of these ice sheet parameter fields from noisy data can be accomplished via the framework of Bayesian inference. Upon discretization of the unknown infinite-dimensional parameter field, the solution of the Bayesian inference problem takes the form of a very high-dimensional *posterior* probability density function (pdf) that assigns to any candidate set of parameter fields our belief (expressed as a probability) that a member of this candidate set is the "true" parameter field that gave rise to the observed data. Sampling this posterior pdf to compute, for example, the mean and covariance of the parameters presents tremendous challenges, since not only is it high-dimensional, but evaluating the pdf at any point in parameter space requires a forward ice sheet flow simulation—and millions of such evaluations may be required to obtain statistics of interest using state-of-the-art Markov chain Monte Carlo methods. Finally, the ice sheet model parameters and their associated uncertainties can be propagated through the ice sheet flow model to yield predictions of not only the mass flux of ice into the ocean, but also the confidence we have in those predictions. This amounts to solving a system of stochastic PDEs, which again is intractable when the PDEs are complex and highly nonlinear and the parameters are high-dimensional due to discretization of an infinite-dimensional field.

In summary, while one can *formulate* a data-to-prediction framework to quantify uncertainties from data to inferred model parameters to predictions with an underlying model of non-Newtonian ice sheet flow, attempting to *execute* this framework for the Antarctic ice sheet (or other large-scale complex models) is intractable for high-dimensional parameter fields using current algorithms. Yet, quantifying the uncertainties in predictions of ice sheet models is essential if these models are to play a significant role in projections of future sea level. The purpose of this paper is to present an integrated framework and efficient, scalable algorithms for carrying out this data-to-prediction process. By *scalable*, we mean that the cost—measured in number of (linearized) forward (and adjoint) solves—is independent of not only the number of processor cores, but importantly the state variable dimension, the parameter dimension, and the data dimension.

Two key ideas are needed to produce such scalable algorithms. First, we use Gaussian approximations of both the posterior pdf that results from Bayesian solution of the inverse problem of inferring ice sheet parameter fields from satellite observations of surface velocity, as well as the pdf resulting from propagating the uncertain parameter fields through the forward ice sheet model to yield predictions of present-day mass flux into the ocean. This is accomplished by linearizing the parameter-to-observable map as well as the parameter-to-prediction map around the maximum a posterior point. We have found that for ice sheet flow problems with the basal sliding parameter as the field of interest, such linearizations are satisfactory approximations for what would otherwise be an intractable problem [6].

However, even with these linearizations, computing the covariance of each of the resulting pdf's is prohibitive due to the need to solve the forward ice sheet model a number of times equal to the parameter dimension (or data dimension).

[1] This estimate is broken down into estimates for individual ice sheets with confidence intervals (-149 ± 49 Gigatonnes per year from Greenland, $+14 \pm 43$ from East Antarctica, -65 ± 26 from West Antarctica, -20 ± 15 from the Antarctic peninsula).