Probabilistic Sea-Level Projections from Ice Sheet and Earth System Models 3: Performance, Optimization and Uncertainty Quantification

Performance

Goals:

Improve time-to-solution and scalability of BISICLES and MALI on new and existing high performance computing architectures

Progress:

Performance Portability:

- In order to run effectively on new and emerging architectures (e.g. GPUs), MALI has adopted the Kokkos programming model for on-node parallelism. The results show a speedup over traditional MPI-only simulations for the finite element assembly process of the 4km-20km Greenland ice sheet on multiple supercomputing architectures without architecture dependent code optimizations.

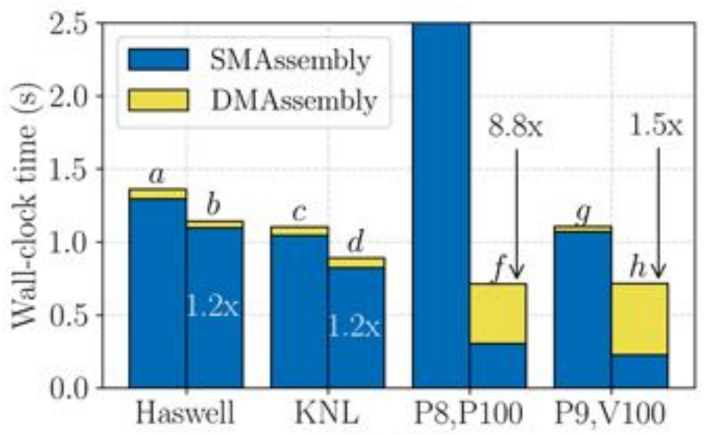

(a,c,d,g) MPI-only (b,d) MPI+OpenMP (f,h) MPI+GPU

- As a Chombo-based application, BISICLES will be an early adopter of Chombo's PROTO framework for performance portability, currently under development as a part of the ECP project.

Improvements:

- Memoization is used in MALI to store static quantities and avoid unnecessary recomputations during initialization and assembly.
- Sacado automatic differentiation is used with static data types to improve runtime during assembly.

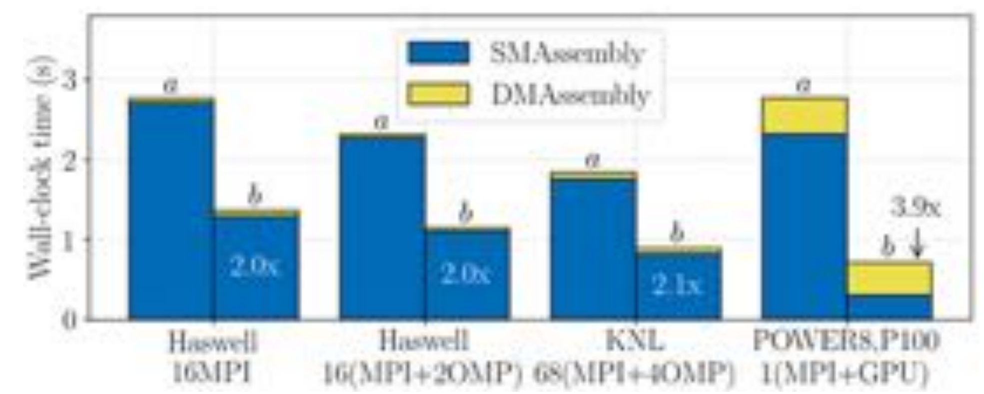

Finite element assembly, 4km-20km Greenland ice sheet
(a) No memoization (b) Memoization

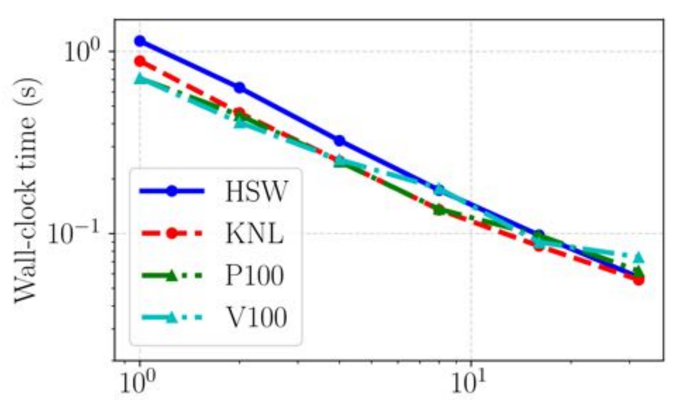

Finite element assembly strong scalability, 4km-20km Greenland ice sheet

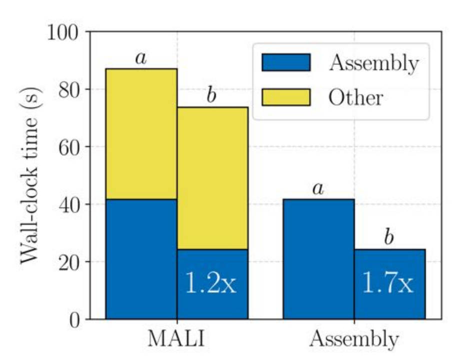

MALI, MISMIP test, single core

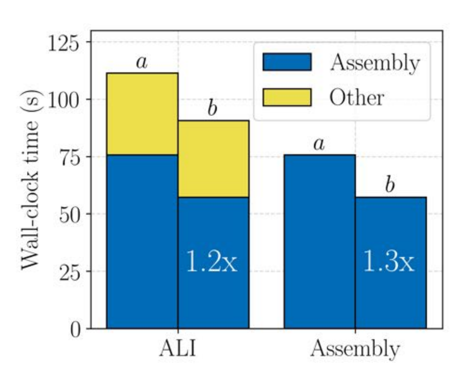

Basal friction initialization, coarse Greenland ice sheet, single core

Scalability:

- In order to scale on next generation supercomputers, MALI has focused on testing and studying the scalability of the finite element assembly process on various architectures including Intel Haswell/Skylake/KNL, IBM POWER8/POWER9, ARM Cavium ThunderX2 and NVIDIA P100/V100 GPUs. These studies have to led to various bug fixes, improvements and insights for future architectures while preparing MALI for exascale computing.
- BISICLES is also undergoing scaling and performance optimization for NERSC's Cori platform after the retirement of the Edison platform.

Linear Solvers:

- In order to improve convergence and robustness of the linear solvers it is important to detect runtime geometry features like icebergs and ice hinges. In collaboration with FASTMath we developed an efficient and scalable graph algorithm for detecting such features. This work (Ian Bogle et al, Proceedings of ACM, 2019) lead to the **ICPP best paper award.**

Optimization

Goals:

Several processes critical for the accurate modeling of ice sheets (e.g., internal rheological and basal properties) cannot be observed directly and therefore must be inferred by combining observational data and ice sheet models. This results in an optimization problem constrained by ice sheet model equations.

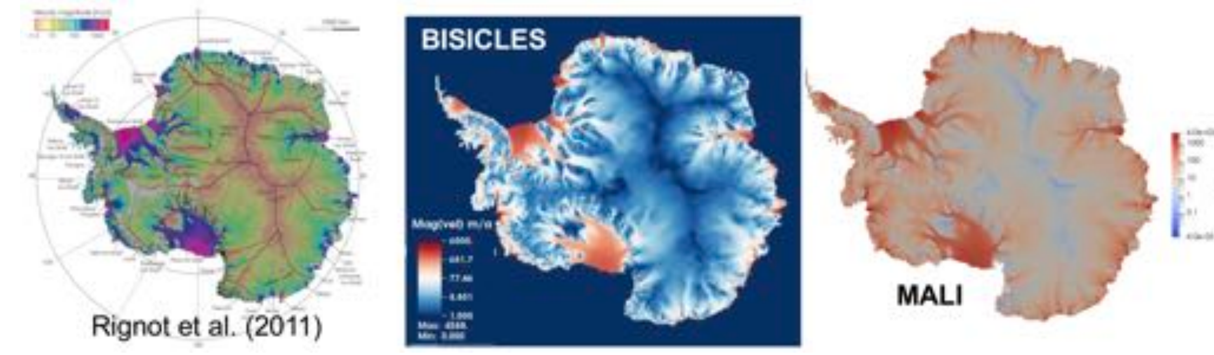

Left: Observed Antarctic ice sheet surface velocities and modeled surface velocities from BISICLES and MALI. Model-based velocities shown here rely on the "basic" optimization schemes discussion below.

MALI and BISICLES already employ several modern optimization methods (see below). However, *ProSPect* aims to offer and perfect new methods that not only provide a best match to the observed ice sheet *state* but also to *tendencies*, while simultaneously enforcing that the initial state is self-consistent with the model physics (momentum, temperature, hydrology). This is critically important when generating initial conditions targeted for use with coupled Earth System Models like E3SM.

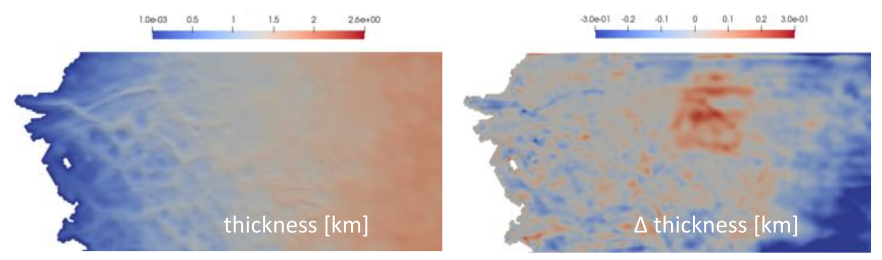

Novelty 1: Ice thickness is allowed to vary during the optimization (constrained by uncertainties) to provide another degree of freedom.

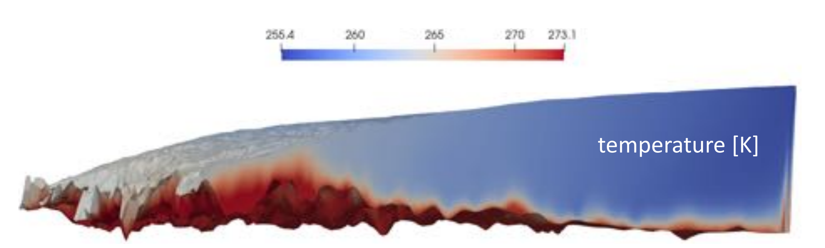

Novelty 2: Ice temperature, a strong control on rheology, is simultaneously optimized to be consistent with ice dynamics (via enthalpy solution).

Progress:

In addition to the commonly used objective functional for optimizing the ice sheet velocity field (given an assumed geometry), we have introduced new observational constraints and new parameters specifically targeting an improved match to important ice sheet model *tendencies* (like the rate of thickness change). Moreover we are enforcing that the initial state is self-consistent with the temperature.

Here (on the left and below) we present results related to the initialization of the Isunnguata Sermia glacier from Western Greenland.

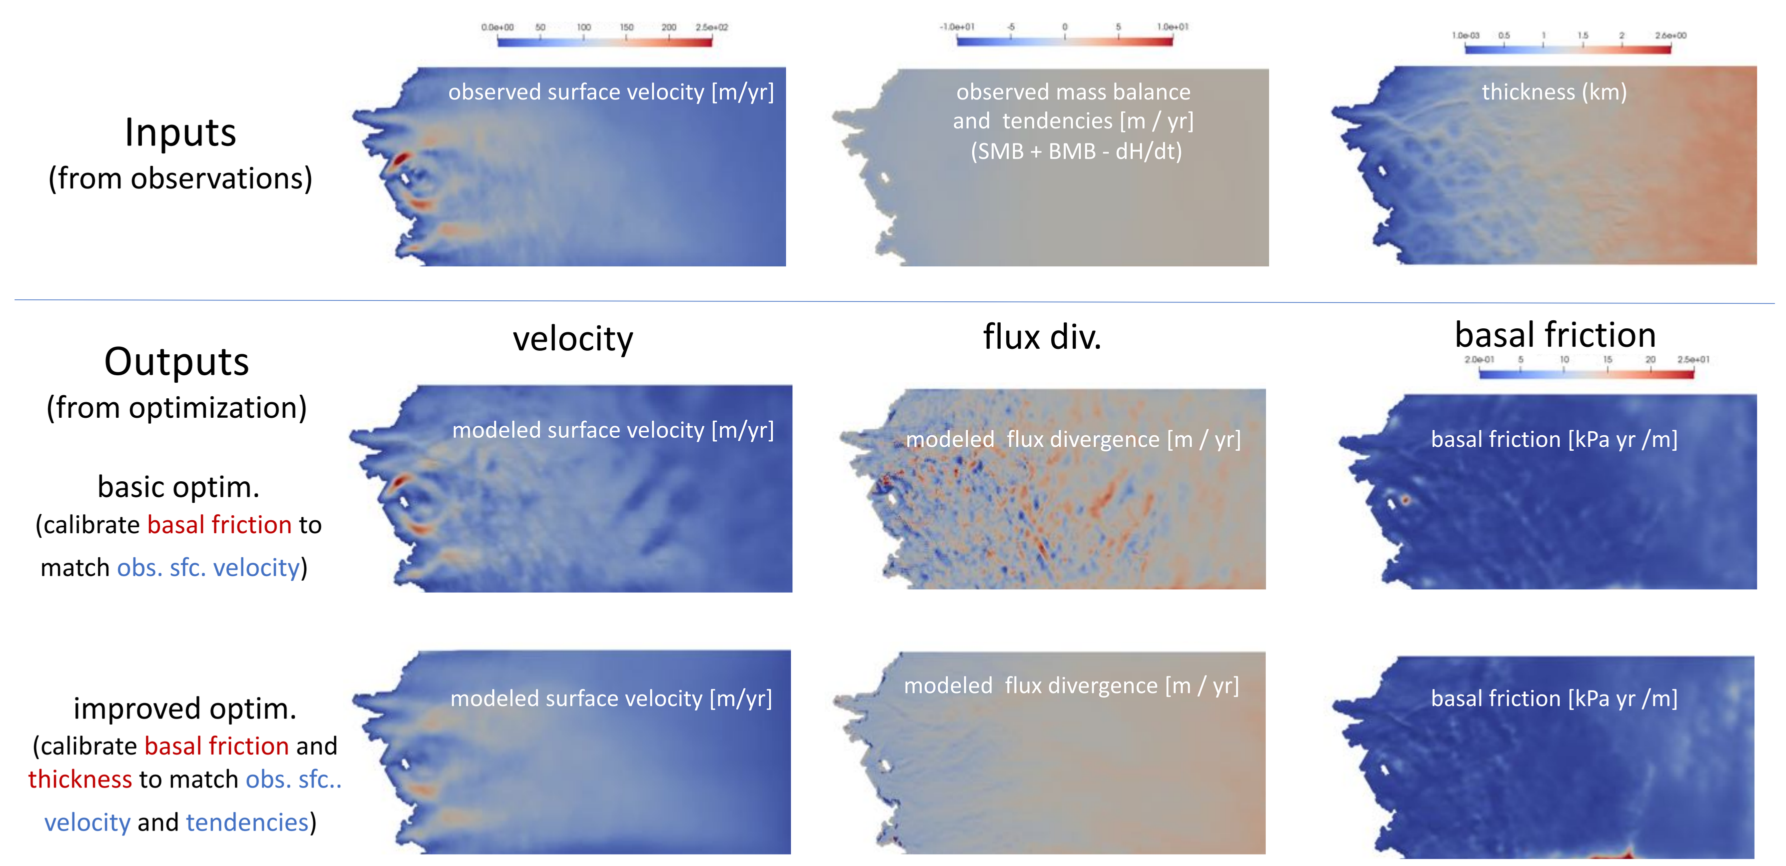

Uncertainty Quantification

Goals:

- Estimate uncertainty in projections of sea-level rise from ice sheet model simulations, with a particular focus on uncertainties resulting from high-dimensional parameter fields.

Progress:

- dimension reduction and Bayesian inference to compute posterior distribution of basal friction field
- propagated samples from this distribution through transient model
- developed suite of models with varying fidelity (Stokes, Blatter-Pattyn (BP), and Shallow Ice Approximation (SIA)) and idealized test cases to systematically study prospects and limitation of UQ algorithms.

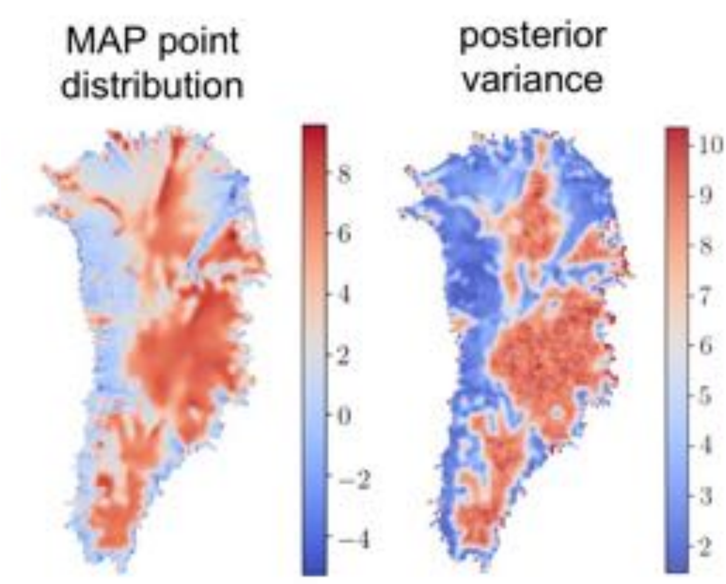

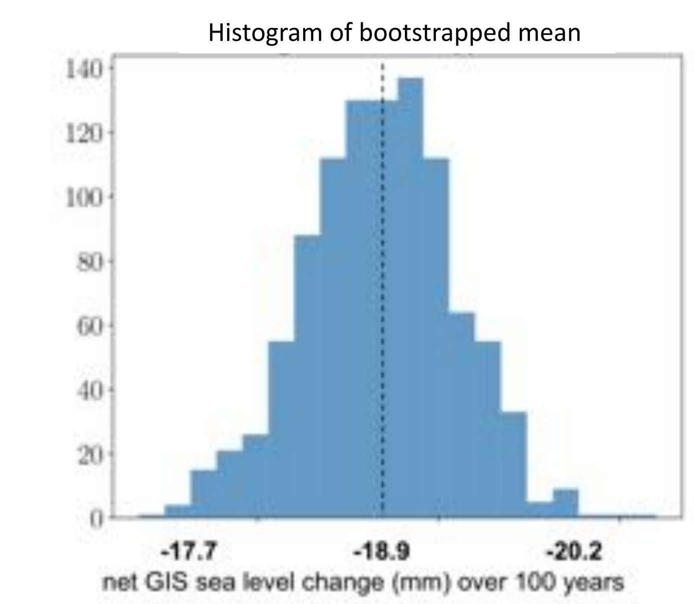

Above: (Left) Map point and sample of posterior obtained from Bayesian inversion. We use dimension reduction to identify parameter directions significantly informed by data. (Right) Mean of projected sea-level rise. Log-normal prior may be cause of bias towards mass increase.

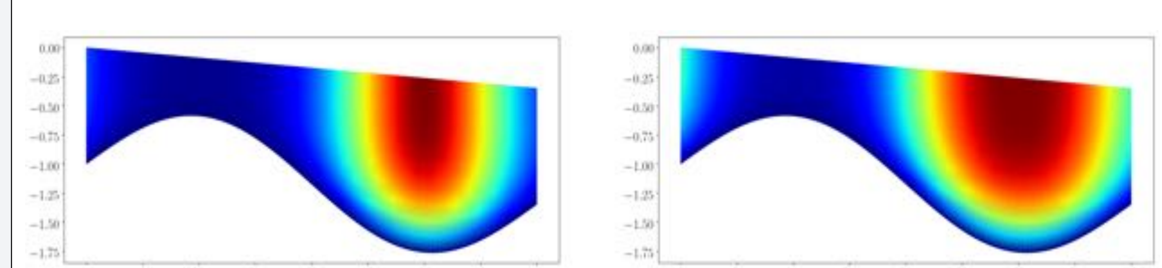

Above: Ice velocity for ISMIPHOM benchmark experiment B using SIA (left) and BP models (right). Cost is >1000 times less than the above Greenland problem.

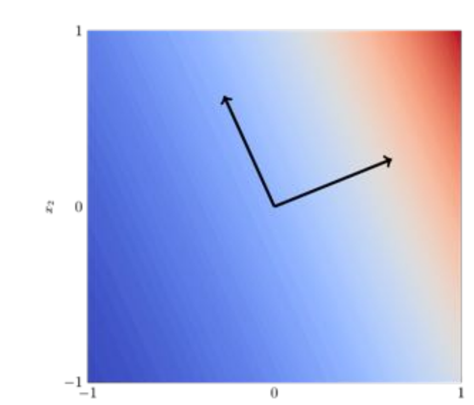

Above: We can use gradients of mass loss to determine directions that significantly impact sea-level rise.

Ongoing work:

- Use low fidelity models to study problems (such as bias above) with the large-scale, high-resolution, expensive end-to-end framework.
- Investigate multi-fidelity methods which complement a limited number of high-fidelity simulations with a large number of low-fidelity runs in a manner that balances physical prediction and statistical errors.
- Use dimension reduction, leveraging transient adjoints obtained from new model suite, to reduce cost of propagating uncertainties through transient model.

Project Members: Stephen Price (PI; LANL), Esmond Ng (PI; LBNL), Xylar Asay-Davis (LANL), Jeremy Bassis (UM), L. Bertagna (SNL), C. Book (LANL), Darin Comeau (LANL), Matt Dunlop (NYU), Katherine Evans (ORNL), Matthew Hoffman (LANL), John Jakeman (SNL), Sam Kachuck (UM), Joseph Kennedy (ORNL), Dan Martin (LBNL), Mauro Perego (SNL), Mark Petersen (LANL), Andrew Salinger (SNL), Adam Schneider (UCI), Chad Sockwell (SNL), Georg Stadler (NYU), Irina Tezaur (SNL), Raymond Tuminaro (SNL), Brian Van Straalen (LBNL), Jerry Watkins (SNL), Tong Zhang (LANL)

https://doe-prospect.github.io/