Update on the *Albany/FELIX* First Order Stokes Solver and the *CISM-Albany* and *MPAS-Albany* Dycores

Irina Kalashnikova, Mauro Perego, Andy Salinger,
Ray Tuminaro, Steve Price, Matt Hoffman

Sandia National Laboratories*

CESM Land Ice Working Group Meeting
February 2-3, 2015

NCAR - MESA Lab
Boulder, CO

SAND2015-0478 C

*Sandia is a multiprogram laboratory operated by Sandia corporation, a Lockheed Martin Company, for the U.S. Department of Energy under contract DE-AC04-94AL85000.

Sandia's Role in the PISCEES Project: *Albany/FELIX, CISM-Albany, MPAS-Albany*

SciDAC Scientific Discovery through Advanced Computing

PISCEES = "Predicting Climate and Evolution at Extreme Scales"
(SciDAC application partnership b/w DOE's BER + ASCR Divisions, began June 2012, 5 years).

- **Sandia's Role in PISCEES**: to **develop** and **support** a **production-ready** robust & scalable unstructured grid finite element land ice dycore based on the "first-order" (FO) Stokes physics.

[Computational Science Flavor]

Albany/FELIX
(Part I of talk)

Steady state **finite element** solver for momentum balance (FO Stokes PDEs)

Implemented in a Sandia (open-source) parallel C++ component-based FE code: Albany

Trilinos: Solvers, preconditioners, automatic differentiation, FE library, performance-portable kernels, meshes, …

DAKOTA: Parameter estimation, UQ, optimization, Bayesian inference.

[Climate Science/ Applied Flavor]

CISM-Albany and ***MPAS-Albany***
(Part II of talk)

Dynamic solver for ice sheet **evolution** PDEs (thickness & temperature advection-diffusion PDEs)

Implemented by writing **interfaces** b/w *Albany* and *CISM/MPAS* codes

***Albany/FELIX* (steady):** FO Stokes PDEs (stress-velocity solve) → / ← ***CISM/MPAS* (dynamic):** ice sheet evolution PDEs (thickness/temp. evolution)

Production code for long-term use in ACME.

Recap of 2012-14 Progress

*Use of **Trilinos** components has enabled **rapid** development of Albany/FELIX!*

2012	2013	2014
• Implement FO Stokes PDEs and relevant BCs in *Albany* code → *Albany/ FELIX* solver is born. • Verify *Albany/FELIX* on MMS and canonical benchmark problems. • Preliminary performance (robustness and scalability) studies.	• Import GIS/AIS data (β, temperature,…) into *Albany/FELIX* in various mesh formats (structured hex & tet, unstructured). • Couple *Albany/FELIX* to *MPAS* and *CISM* codes. • Convergence/performance studies on GIS. • Deterministic inversion for initialization (in *LifeV*). • Bayesian calibration for initialization.	• Implementation of adjoints for deterministic inversion in *Albany*. • Scalability studies on large-scale GIS and AIS problems. • Performance portability in *Albany/FELIX* MiniApp. • Continued maturation of *CISM-Albany* and *MPAS-Albany*. • *GMD, ICCS, SISC* papers submitted/in progress.

Recap of 2012-14 Progress

*Use of **Trilinos** components has enabled **rapid** development of Albany/FELIX!*

2012

- Implement FO Stokes PDEs and relevant BCs in *Albany* code → *Albany/ FELIX* solver is born.
- Verify *Albany/FELIX* on MMS and canonical benchmark problems.
- Preliminary performance (robustness and scalability) studies.

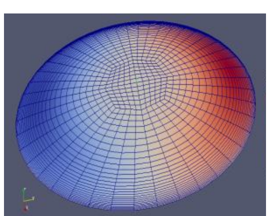

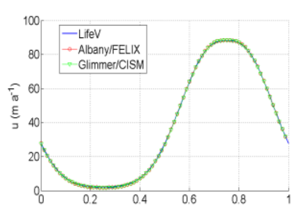

2013

- Import GIS/AIS data (β, temperature,…) into *Albany/FELIX* in various mesh formats (structured hex & tet, unstructured).
- Couple *Albany/FELIX* to *MPAS* and *CISM* codes.
- Convergence/performance studies on GIS.
- Deterministic inversion for initialization (in *LifeV*).
- Bayesian calibration for initialization.

2014

- Implementation of adjoints for deterministic inversion in *Albany*.
- Scalability studies on large-scale GIS and AIS problems.
- Performance portability in *Albany/FELIX* MiniApp.
- Continued maturation of *CISM-Albany* and *MPAS-Albany*.
- *GMD, ICCS, SISC* papers submitted/in progress.

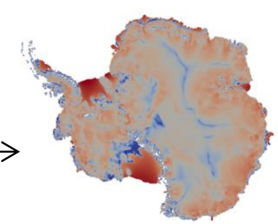

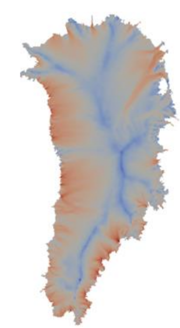

Sandia National Laboratories

Part I: The *Albany/FELIX* First Order (FO) Stokes Solver

- **PDEs:** first-order (FO) Stokes PDEs with Glen's law viscosity.
- **BCs:** stress-free, basal, floating ice, kinematic.
- **Discretization:** unstructured grid finite element method (FEM).
- **Meshes:** structured hex, structured tet, unstructured tet.
- **Nonlinear solver:** full Newton with analytic (automatic differentiation) derivatives and homotopy continuation.
- **Iterative linear solver**: CG or GMRES with ILU or AMG preconditioner.
- **Advanced analysis capabilities**: sensitivities, UQ, responses, adjoint-based optimization.

$$\begin{cases} -\nabla \cdot (2\mu \dot{\epsilon}_1) = -\rho g \dfrac{\partial s}{\partial x} \\ -\nabla \cdot (2\mu \dot{\epsilon}_2) = -\rho g \dfrac{\partial s}{\partial y} \end{cases}$$

$$\mu = \frac{1}{2} A^{-\frac{1}{n}} \left(\frac{1}{2} \sum_{ij} \dot{\epsilon}_{ij}{}^2 + \gamma \right)^{\left(\frac{1}{2n} - \frac{1}{2}\right)}$$

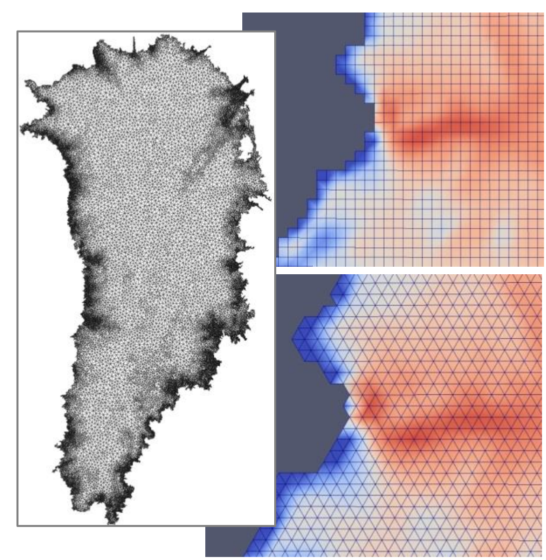

Part I: The *Albany/FELIX* First Order (FO) Stokes Solver

- **PDEs:** first-order (FO) Stokes PDEs with Glen's law viscosity.
- **BCs:** stress-free, basal, floating ice, kinematic.
- **Discretization:** unstructured grid finite element method (FEM).
- **Meshes:** structured hex, structured tet, unstructured tet.
- **Nonlinear solver:** full Newton with analytic (automatic differentiation) derivatives and homotopy continuation.
- **Iterative linear solver**: CG or GMRES with ILU or AMG preconditioner.
- **Advanced analysis capabilities**: sensitivities, UQ, responses, adjoint-based optimization.

$$\begin{cases} -\nabla \cdot (2\mu \dot{\epsilon}_1) = -\rho g \dfrac{\partial s}{\partial x} \\ -\nabla \cdot (2\mu \dot{\epsilon}_2) = -\rho g \dfrac{\partial s}{\partial y} \end{cases}$$

$$\mu = \frac{1}{2} A^{-\frac{1}{n}} \left(\frac{1}{2} \sum_{ij} \dot{\epsilon}_{ij}^{\,2} + \gamma \right)^{\left(\frac{1}{2n} - \frac{1}{2}\right)}$$

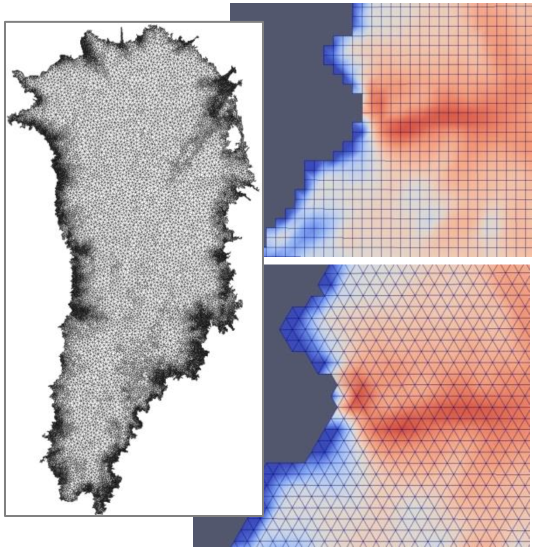

Highlights of Recent Work:

- Built-in adjoints for inversion.
- AIS scaling studies (CG vs. GMRES, ILU vs. new AMG preconditioner aggressive semi-coarsening).
- Performance-portable kernels.

Courtesy of: *M. Perego (SNL)*

Implementation of Adjoints in *Albany/FELIX* for Deterministic Inversion

2-8 km unstructured tet mesh, 10 layers

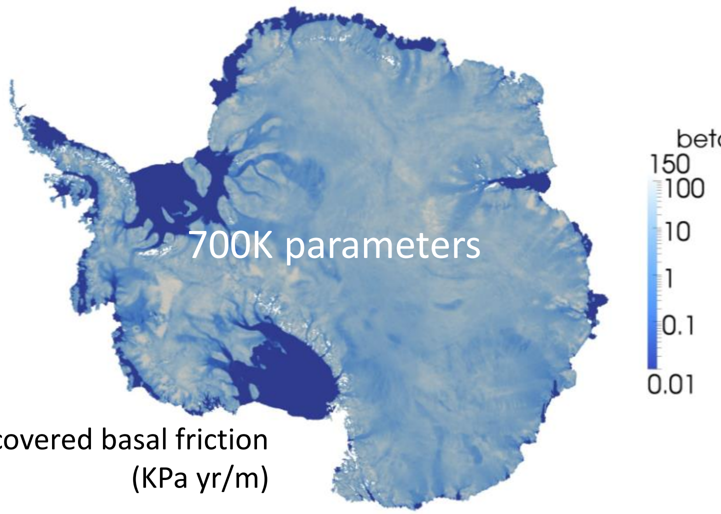

Recovered basal friction (KPa yr/m)

Objective Functional* to be minimized:

$$\mathcal{J}(\mathbf{u}(\beta),\ \beta) = \int_{\Sigma} \frac{1}{\sigma_u^2} |\mathbf{u} - \mathbf{u}^{obs}|^2\, ds + \ \alpha \int_{\Sigma} |\nabla \beta|^2\, ds$$

Subject to: FO Stokes PDEs.

Software details:

- Adjoints and derivatives w.r.t. parameters are computed using automatic differentiation *(Sacado).*
- Reduced gradient based optimization performed using *ROL* (*Rapid Optimization Library*, part of *Trilinos*).
- Optimization Algorithm: Limited-Memory BFGS

*Perego, Price, Stadler, *JGR*, 2014

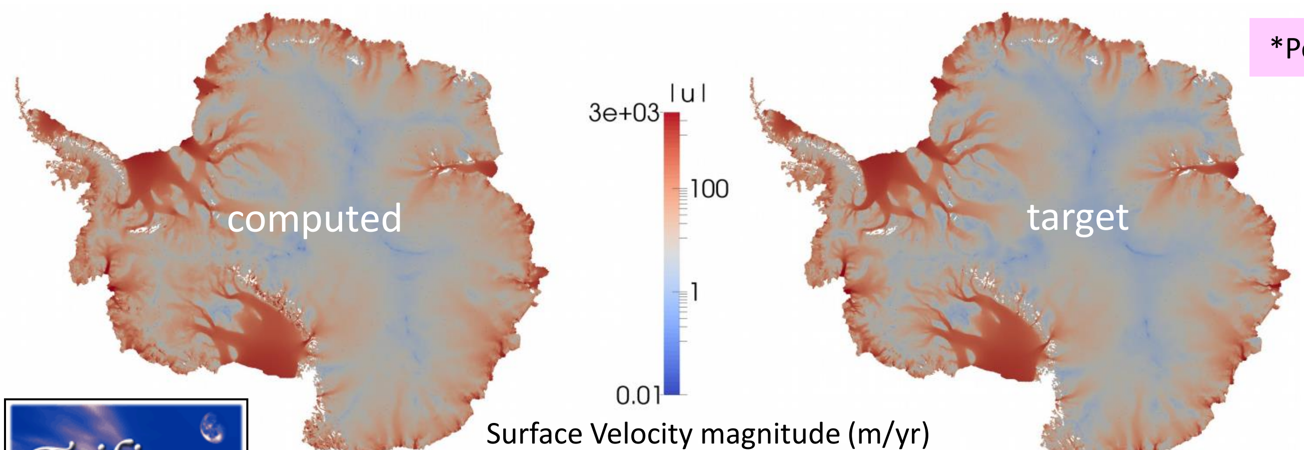

Surface Velocity magnitude (m/yr)

Geometry and fields:
(Conford, Martin, et al, *in prep.*)
BEDMAP2 (Fretwell et al. 2013)
Temperature (Pattyn, 2010)

***Collaborators**: E. Phipps, D. Ridzal, D. Kouri (SNL)*

***Collaborators**: R. Tuminaro (SNL)*

Albany/FELIX Weak Scaling on a Moderate-Resolution AIS Problem

- Weak scaling study on AIS problem (8km w/ 5 layers → 2km with 20 layers).
- Initialized with realistic basal friction (from deterministic inversion) and temperature field from BEDMAP2 (*thanks to **D. Martin**!*)
- **Iterative linear solver**: GMRES.
- **Preconditioner**: ILU vs. new AMG based on aggressive semi-coarsening (Kalashnikova et al *GMD* 2014, Kalashnikova et al *ICCS* 2015, Tuminaro et al *SISC* 2015).

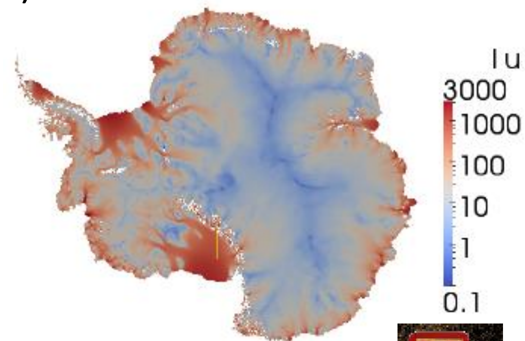

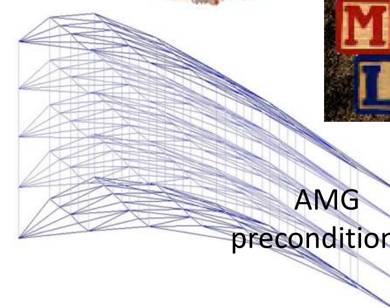

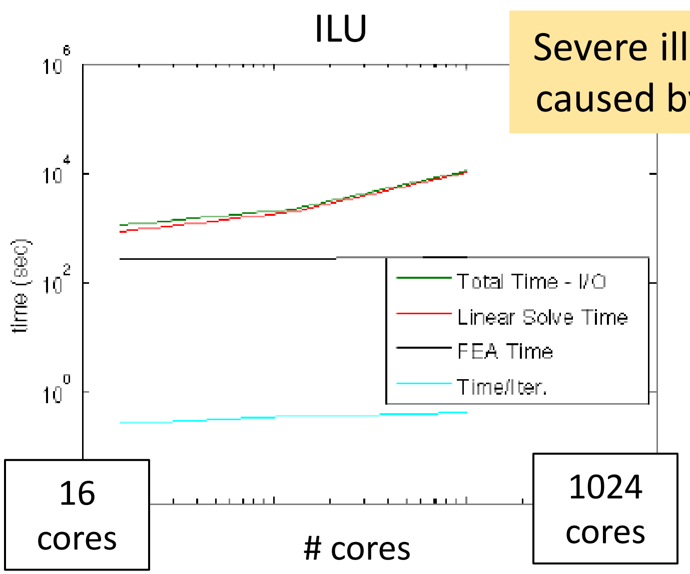

Severe ill-conditioning caused by ice shelves!

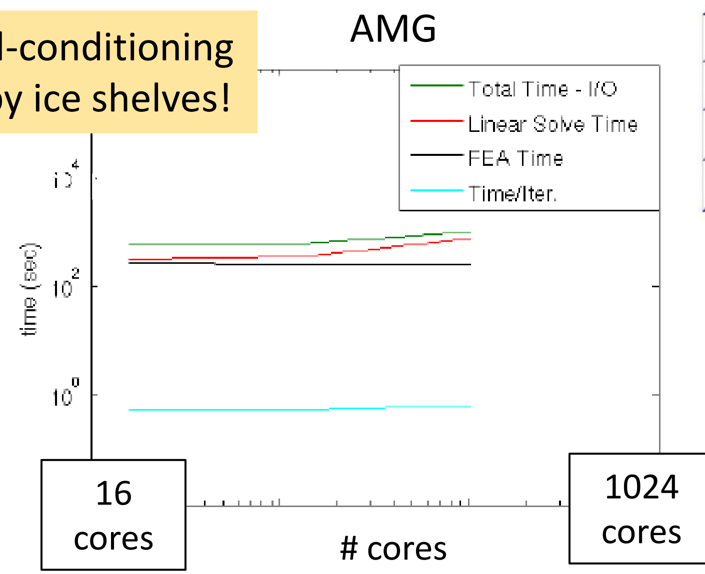

(vertical > horizontal coupling)
+
Neumann BCs
=
nearly singular submatrix associated with vertical lines

GMRES less sensitive than CG to rounding errors from ill-conditioning [also minimizes different norm].

AMG preconditioner less sensitive than ILU to ill-conditioning.

Performance-Portability of *Albany/FELIX*

We need to be able to run *Albany/FELIX* on ***new architecture machines*** (hybrid systems) and ***manycore devices*** (multi-core CPU, NVIDIA GPU, Intel Xeon Phi, etc.) .

- ***Kokkos***: *Trilinos* library that provides performance portability across diverse devises with different memory models.
- With *Kokkos*, you write an algorithm once, and just change a template parameter to get the optimal data layout for your hardware.
- **Work in progress**: converting *Albany* to use *Kokkos* kernels
  - *Albany/FELIX* MiniApp for FE Assembly using *Kokkos.*
  - *Albany* branch based on *Tpetra* released Oct. 2014 on `github`.

Albany/FELIX MiniApp 20km GIS

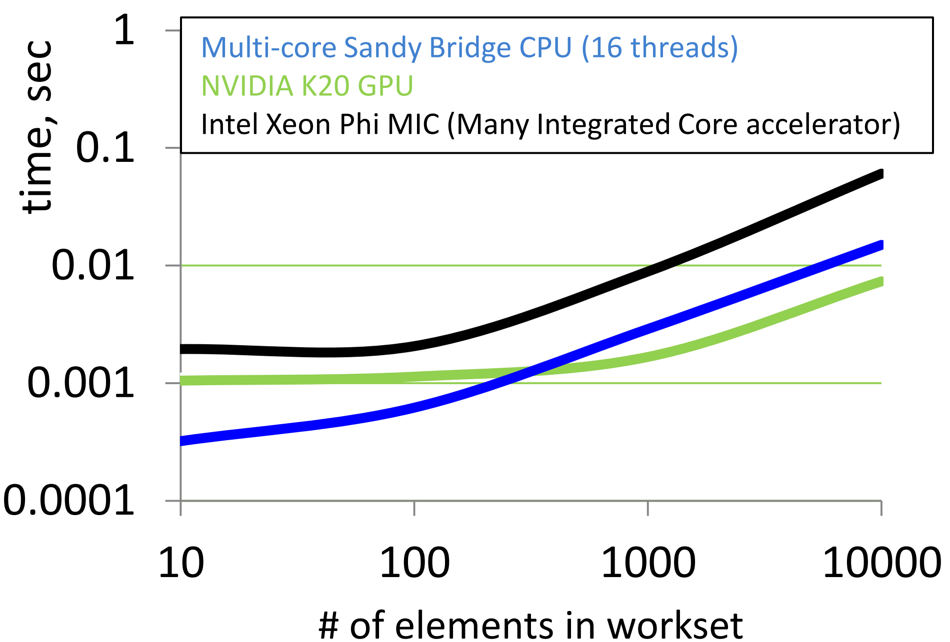

Collaborator: *I. Demeshko (SNL)*

Sandia National Laboratories

Part II: *CISM-Albany* & *MPAS-Albany* Interfaces

CISM-Albany

(I. Kalashnikova)

simple_glide

[cmake]

CISM (Fortran)
Thickness evolution, temperature solve, coupling to CESM

C++/Fortran Interface, Mesh Conversion

Albany/FELIX (C++)
velocity solve

MPAS-Albany

(M. Perego, M. Hoffman)

LandIce_model

[Makefile]

MPAS Land-Ice (Fortran)
Thickness evolution, temperature solve, coupling to DOE-ESM

C++/Fortran Interface, Mesh Conversion

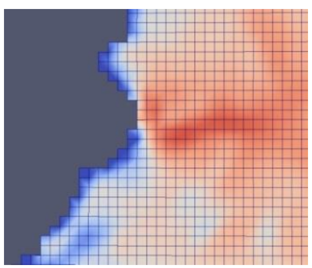

Structured hex meshes (rectangles extruded to hexes).

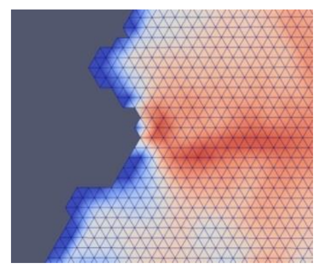

Tet meshes (dual of hexaganonal mesh, extruded to tets).

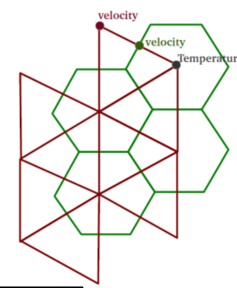

- Both available on `github`.
- Both use forward Euler time-stepping*.
- Same BCs available (free-surface, basal friction, floating ice, kinematic Dirichlet) in *CISM-Albany* and *MPAS-Albany*.

Production codes for long-term use in ESMs through ACME project!

*Some differences in evolution schemes: FV upwinding, flux-corrected transport for MPAS, incremental remap for CISM.

Sandia National Laboratories

New BCs : Floating Ice (*CISM-Albany*) and Kinematic Lateral (*CISM-,MPAS-Albany*)

- **Floating ice lateral BCs** (for ice shelves): assumes ice is in hydrostatic equilibrium with water/air around it.
- **Kinematic lateral BCs**: values of ice velocities are specified on lateral boundaries.

Lateral boundaries identified in *CISM/MPAS*, data passed to *Albany*.

Circular-Shelf

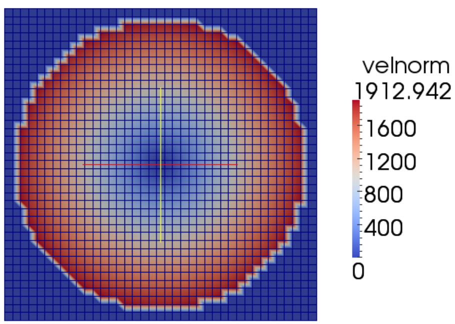

CISM-Albany

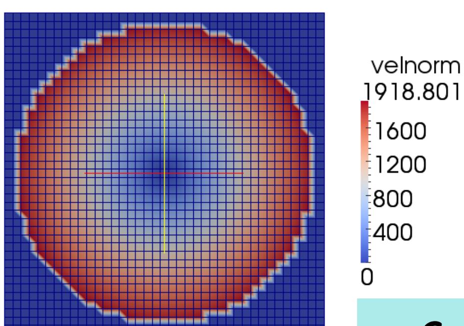

Glissade

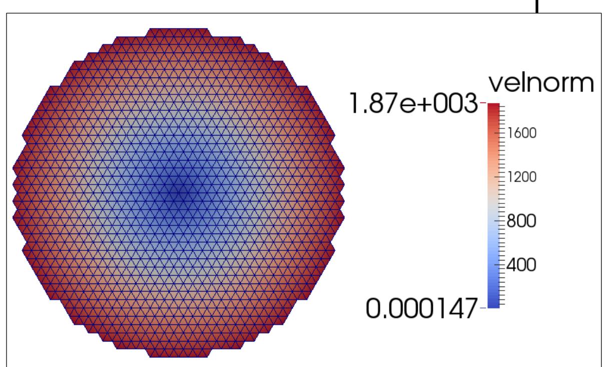

MPAS-Albany

Ross

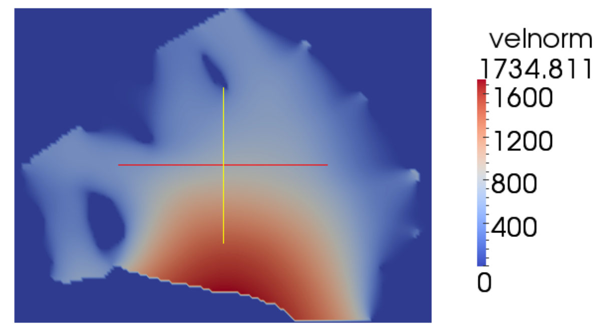

CISM-Albany

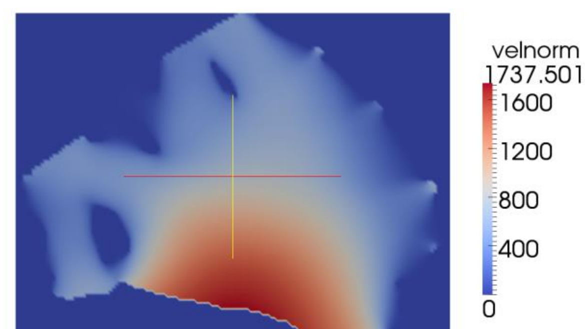

Glissade

***Collaborators**: M. Perego (SNL); S. Price, W. Lipscomb, M. Hoffman (LANL)*

Collaborators: *M. Perego (SNL), S. Price, M. Hoffman (LANL)*

Progress Towards Transient Simulations in *CISM-, MPAS-Albany*

	CISM-Albany	***MPAS-Albany***
Variable β field	X	X
Temperature-based flow factor	X	X
Upstream ∇s calculation	X	
Floating ice BCs	X	X
Kinematic lateral BCs	X	X
Upwinding*	X	X
Incremental Remap*	X	
Flux-Corrected Transport*		X

New since CESM Annual Meeting

Coming soon

Needs further testing

*In *MPAS/CISM*.

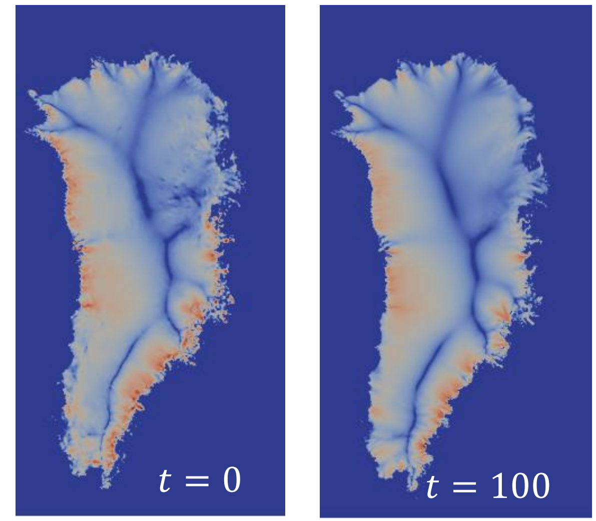

Above: **100 year** 4 km GIS transient simulation using *CISM-Albany* converged on Hopper ***out-of-the box!***

Other planned future work on evolution solvers (*MPAS/CISM*):

- Circumventing CFL restrictions for explicit advection schemes.
- RK-4 time-integrator for *MPAS-LI*.

Collaborators: *M. Perego (SNL), S. Price, M. Hoffman (LANL)*

Progress Towards Transient Simulations in *CISM-, MPAS-Albany*

	CISM-Albany	***MPAS-Albany***
Variable β field	X	X
Temperature-based flow factor	X	X
Upwind ∇s calculation	X	
Floating ice BCs	X	X
Kinematic lateral BCs	X	X
Upwinding*	X	X
Incremental Remap*	X	
Flux-Corrected Transport*		X

*In *MPAS/CISM*.

New since CESM Annual Meeting

Coming soon

Needs further testing

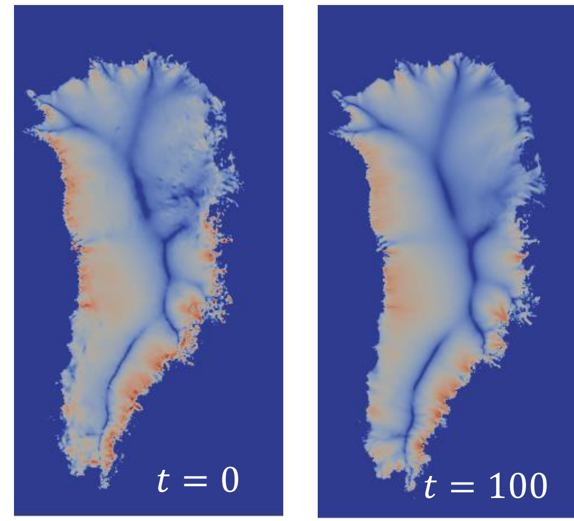

Above: **100 year** 4 km GIS transient simulation using *CISM-Albany* converged on Hopper ***out-of-the box!***

AIS transient simulations run too as of last week!

Other planned future work on evolution solvers (*MPAS/CISM*):

- Circumventing CFL restrictions for explicit advection schemes.
- RK-4 time-integrator for *MPAS-LI*.

Summary and Future Work

Summary:

- Continued maturing and ripening of *Albany/FELIX* (scalability, verification, adjoints, performance portability).
- *CISM-Albany* and *MPAS-Albany* are ready for science runs.
- Use of *Trilinos* components results in code with dozens of built in advanced analysis capabilities (sensitivity analysis, responses, UQ, …)

Progress has been made towards release of a production code supported for long-term use in ACME.

Ongoing/future work:

- Dynamic simulations of ice evolution for GIS & AIS problems.
- Deterministic inversion/calibration using new adjoint capabilities in *Albany/FELIX*.
- Bayesian inference/UQ.
- Porting to hybrid/new architecture machines.
- *GMD, ICCS, SISC* papers in review/preparation.
- Delivering code to users in climate community through coupling to ESMs.

Funding/Acknowledgements

Support for this work was provided through Scientific Discovery through Advanced Computing (**SciDAC**) projects funded by the U.S. Department of Energy, Office of Science (**OSCR**), Advanced Scientific Computing Research and Biological and Environmental Research (**BER**) → **PISCEES SciDAC Application Partnership.**

PISCEES team members: *W. Lipscomb, S. Price, M. Hoffman, A. Salinger, M. Perego, I. Kalashnikova, R. Tuminaro, P. Jones, K. Evans, P. Worley, M. Gunzburger, C. Jackson;*

Trilinos/Dakota collaborators*: E. Phipps, M. Eldred, J. Jakeman, L. Swiler.*

Thank you! Questions?

References

[1] M.A. Heroux *et al.* "An overview of the Trilinos project." *ACM Trans. Math. Softw.* **31**(3) (2005).

[2] A.G. Salinger *et al*. "Albany: Using Agile Components to Develop a Flexible, Generic Multiphysics Analysis Code", *Comput. Sci. Disc.* (in prep).

[3] **I. Kalashnikova**, M. Perego, A. Salinger, R. Tuminaro, S. Price. "*Albany/FELIX*: A Parallel, Scalable and Robust Finite Element Higher-Order Stokes Ice Sheet Solver Built for Advanced Analysis", *Geosci. Model Develop. Discuss.* 7 (2014) 8079-8149 (under review for *GMD).*

[4] **I. Kalashnikova**, R. Tuminaro, M. Perego, A. Salinger, S. Price. "On the scalability of the *Albany/FELIX* first-order Stokes approximation ice sheet solver for large-scale simulations of the Greenland and Antarctic ice sheets", *MSESM/ICCS15*, Reykjavik, Iceland (June 2014).

[5] R.S. Tuminaro, **I. Kalashnikova**, M. Perego, A.G. Salinger. "A Hybrid Operator Dependent Multi-Grid/Algebraic Multi-Grid Approach: Application to Ice Sheet Modeling", *SIAM J. Sci. Comput.* (in prep).

[6] M. Perego, S. Price, G. Stadler. "Optimal initial conditions for coupling ice sheet models to ESMs", *J. Geophys. Res.* **119** (2014) 1894-1917.

Appendix: Code Verification and Performance

- Implementation of PDEs + BCs (no-slip, stress-free, basal sliding, open-ocean) has been ***verified*** through MMS tests (right) and code-to-code comparisons (confined-shelf, below).

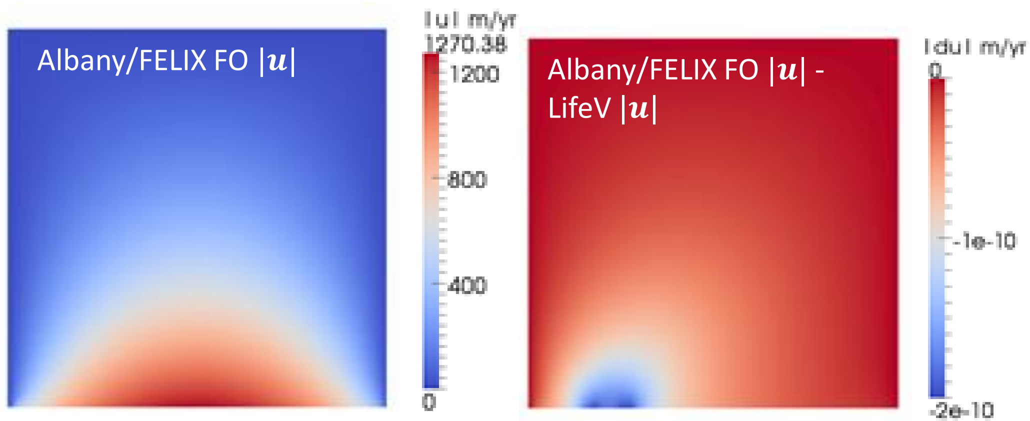

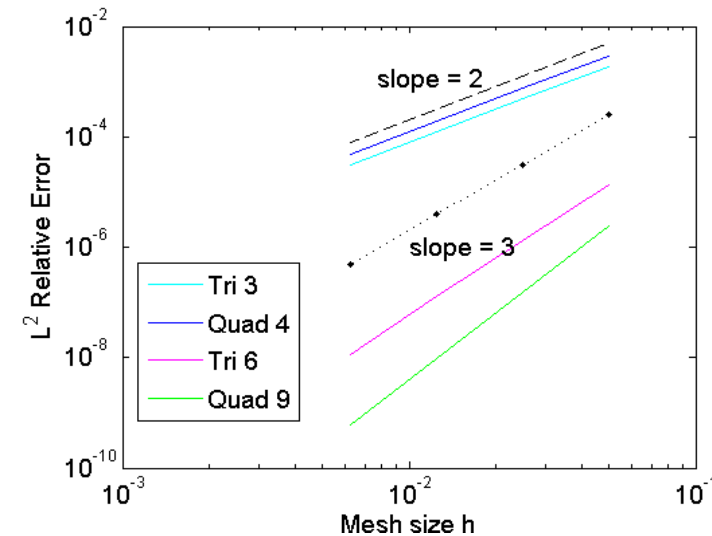

- ***Robust*** nonlinear solves (Newton converges out-of-the-box!) with ***homotopy*** continuation of γ in Glen's law viscosity:

$$\mu = \frac{1}{2} A^{-\frac{1}{n}} \left(\frac{1}{2} \sum_{ij} \dot{\epsilon}_{ij}^{\,2} + \gamma \right)^{\left(\frac{1}{2n} - \frac{1}{2}\right)}$$

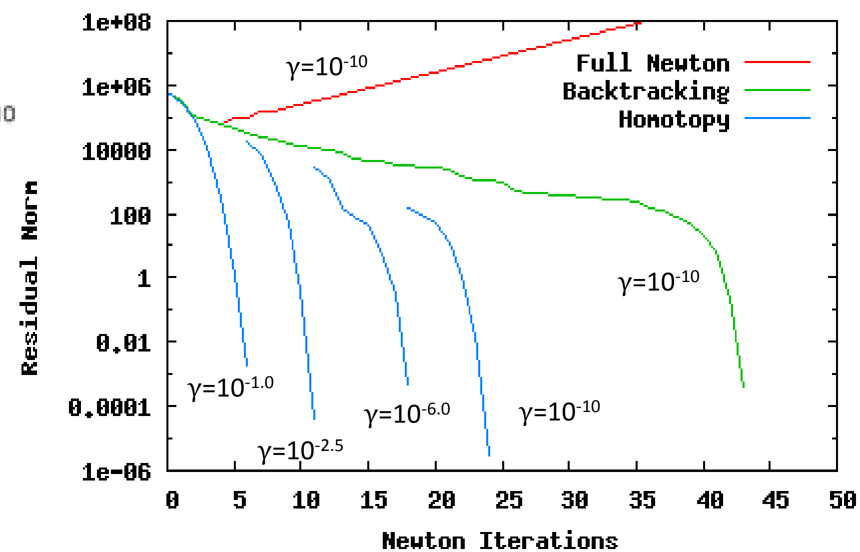

In collaboration with: R. Tuminaro (SNL)

Appendix: *Albany/FELIX* GIS Controlled Weak Scaling Study

New ML preconditioner

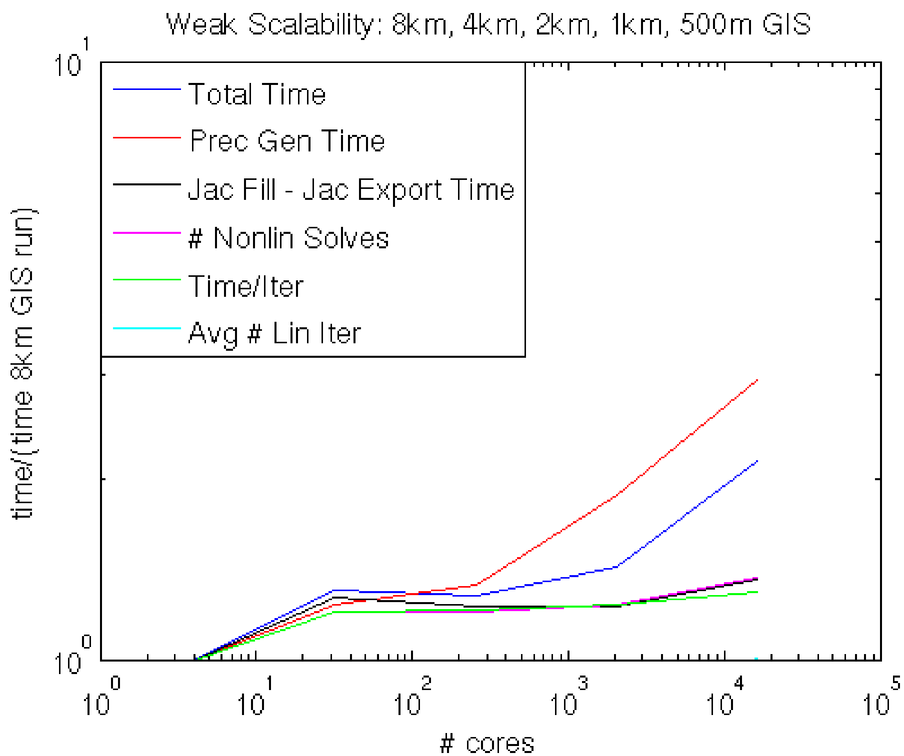

4 cores 334K dofs 8 km GIS, 5 vertical layers	→ $\times 8^4$ scale up	16,384 cores **1.12B dofs(!)** 0.5 km GIS, 80 vertical layers

- Weak scaling study with fixed dataset, 4 mesh bisections.
- ~70-80K dofs/core.
- ***Conjugate Gradient (CG) iterative method*** for linear solves (faster convergence than GMRES).
- ***New algebraic multigrid preconditioner (ML)*** developed by R. Tuminaro based on ***semi-coarsening*** (coarsening in z-direction only).
- ***Significant improvement*** in scalability with new ML preconditioner over ILU preconditioner!

Sandia National Laboratories

Appendix: *Albany/FELIX* GIS Controlled Weak Scaling Study

In collaboration with:
R. Tuminaro (SNL)

New ML preconditioner

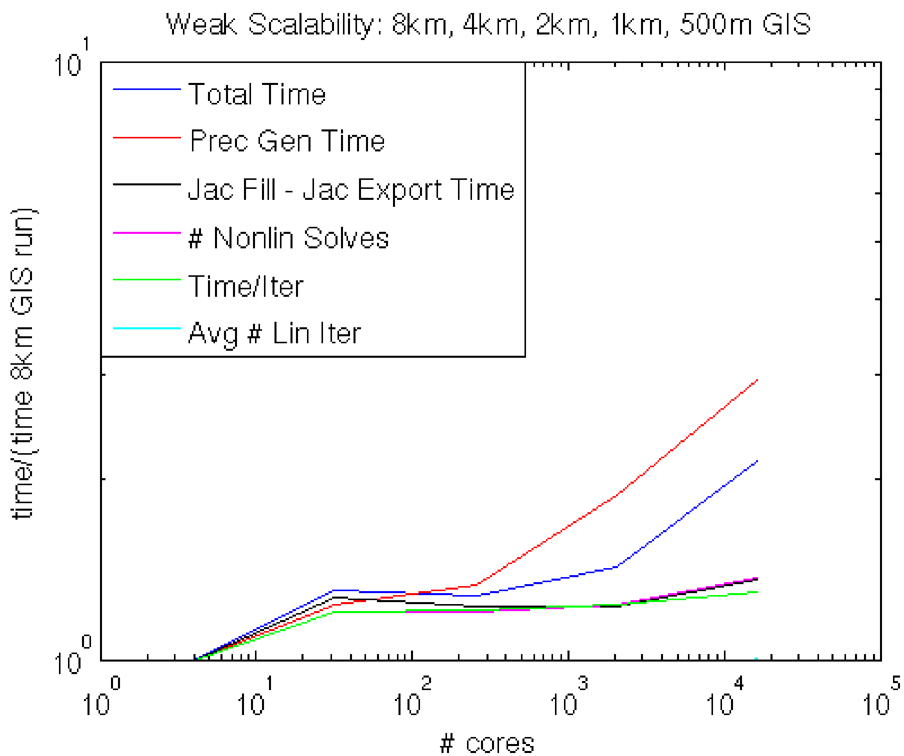

ILU preconditioner

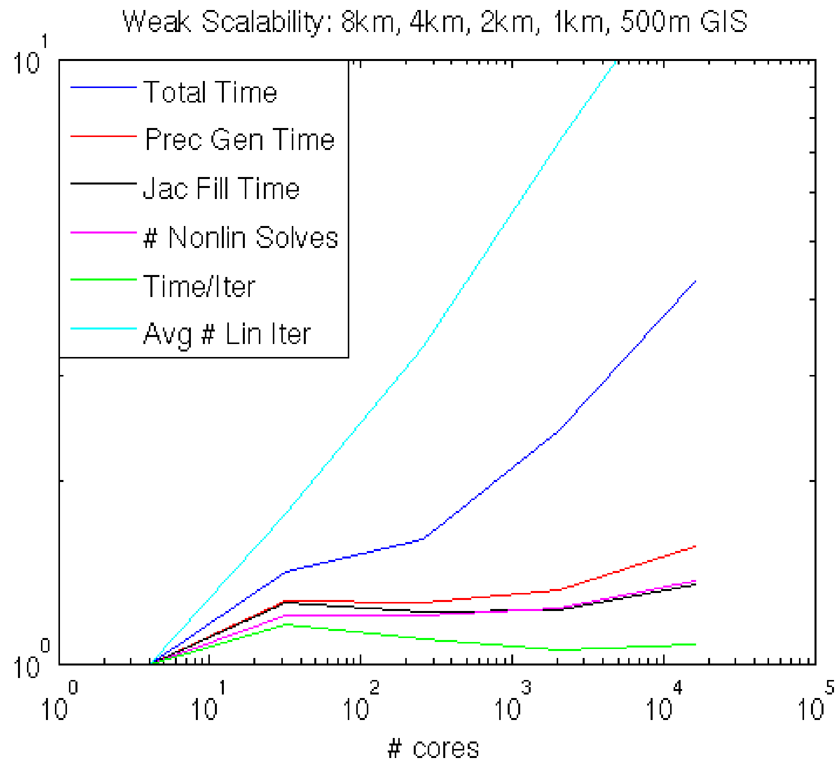

4 cores 334K dofs 8 km GIS, 5 vertical layers	→ $\times 8^4$ scale up	16,384 cores **1.12B dofs(!)** 0.5 km GIS, 80 vertical layers

- ***Significant improvement*** in scalability with new ML preconditioner over ILU preconditioner!

Sandia National Laboratories

Collaborators: *R. Tuminaro (SNL)*

Appendix: *Albany/FELIX* Strong Scaling On a Fine-Resolution GIS Problem

- Strong scaling study on 1km GIS with 40 vertical layers (143M dofs, hex elements).
- Initialized with realistic basal friction (from deterministic inversion) and temperature fields → interpolated from coarser to fine mesh.
- **Iterative linear solver**: CG.
- **Preconditioner**: ILU vs. new AMG (based on aggressive semi-coarsening).

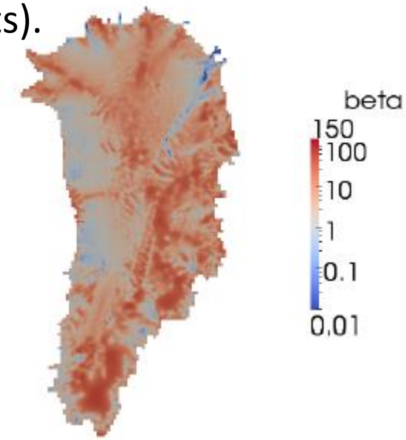

ILU

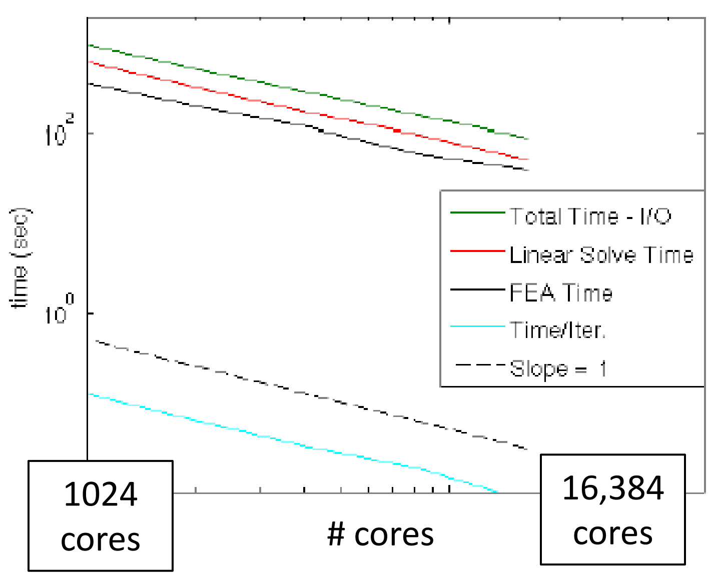

AMG

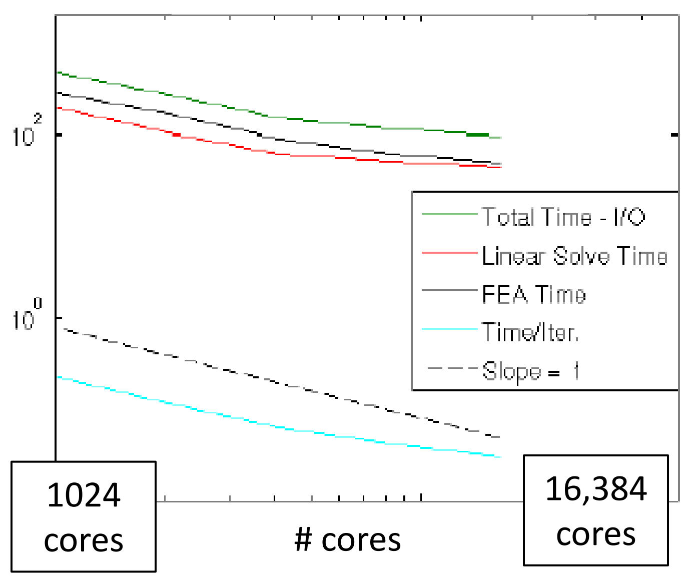

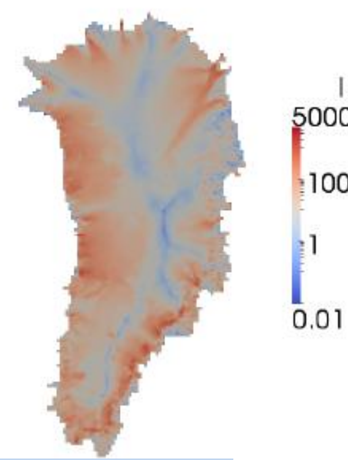

ILU preconditioner scales better than AMG but ILU-preconditioned solve is slightly slower (see ICCS 2015 paper [4]).

Sandia National Laboratories

Collaborators: *S. Price, W. Lipscomb, M. Hoffman (LANL)*

Appendix: Calculation of Surface Height Gradients in *CISM-Albany*

- In the FEM there are several ways to calculate ∇s for RHS in FO Stokes in each element e:

New to *CISM-Albany* (not available in *MPAS-Albany*)

1. $\nabla s^e = \sum_{i=1}^{\#\,nodes} s_i^e \nabla \phi_i^e$, — s passed to *Albany* from *CISM*

2. $\nabla s^e = \sum_{i=1}^{\#\,nodes} \nabla sie \phi_i^e$ — ∇s passed to *Albany* from *CISM*
   - 2(a): ∇sie calculated using central difference
   - 2(b): ∇sie calculated using forward difference

ϕ_i^e =finite element shape functions
s_i^e = values of s at node i of element e

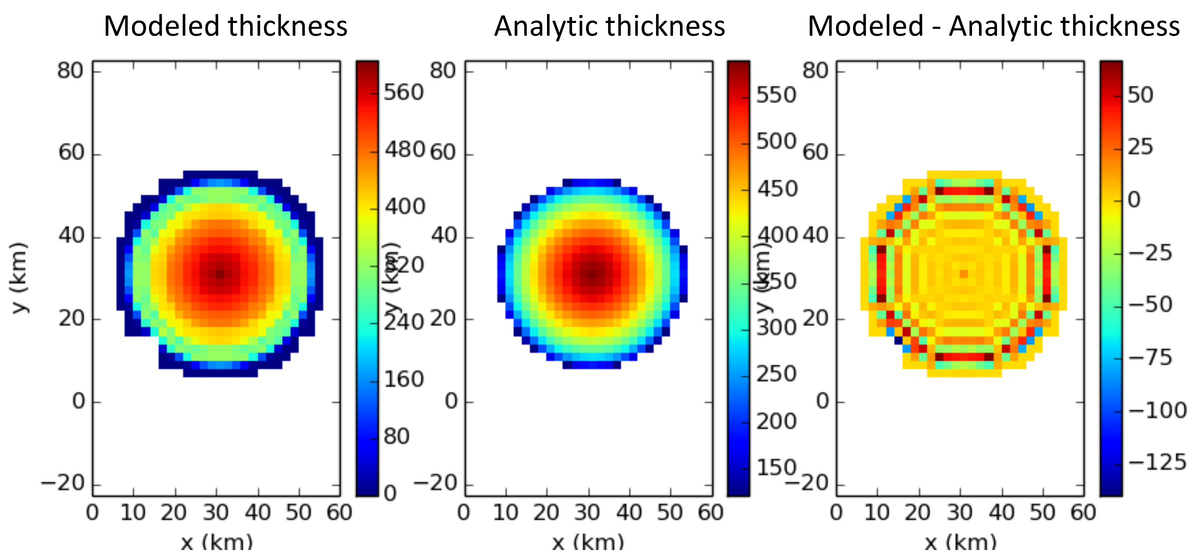

1.

Halfar test case (t = 100 years)

1 and 2(a): checkerboard pattern in thickness error.

2(b): checkerboard pattern in thickness error vanishes.

Observed by B. Lipscomb at CESM Annual Meeting 2014.

Sandia National Laboratories

Collaborators: *S. Price, W. Lipscomb, M. Hoffman (LANL)*

Appendix: Calculation of Surface Height Gradients in *CISM-Albany*

- In the FEM there are several ways to calculate ∇s for RHS in FO Stokes in each element e:

New to *CISM-Albany* (not available in *MPAS-Albany*)

1. $\nabla s^e = \sum_{i=1}^{\# nodes} s_i^e \nabla \phi_i^e$, — s passed to *Albany* from *CISM*
2. $\nabla s^e = \sum_{i=1}^{\# nodes} \nabla sie \phi_i^e$ — ∇s passed to *Albany* from *CISM*
   - 2(a): ∇sie calculated using central difference
   - 2(b): ∇sie calculated using forward difference

ϕ_i^e =finite element shape functions
s_i^e = values of s at node i of element e

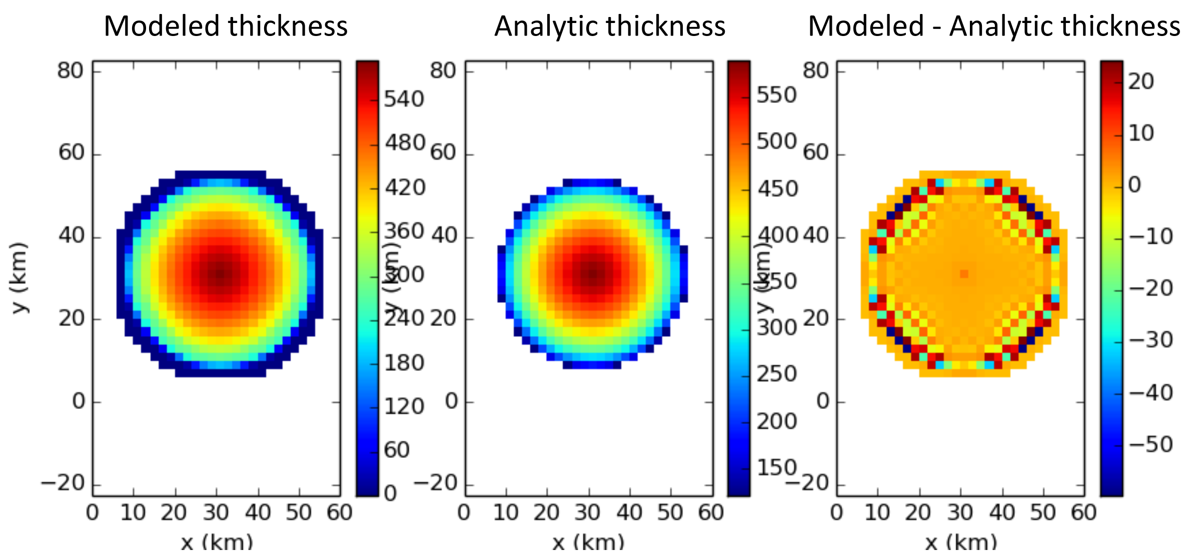

2(a).

Halfar test case (t = 100 years)

1 and 2(a): checkerboard pattern in thickness error.

2(b): checkerboard pattern in thickness error vanishes.

Observed by B. Lipscomb at CESM Annual Meeting 2014.

Collaborators: *S. Price, W. Lipscomb, M. Hoffman (LANL)*

Appendix: Calculation of Surface Height Gradients in *CISM-Albany*

- In the FEM there are several ways to calculate ∇s for RHS in FO Stokes in each element e:

New to *CISM-Albany* (not available in *MPAS-Albany*)

1. $\nabla s^e = \sum_{i=1}^{\#\,nodes} s_i^e \nabla \phi_i^e$, — s passed to *Albany* from *CISM*

2. $\nabla s^e = \sum_{i=1}^{\#\,nodes} \nabla sie \phi_i^e$ — ∇s passed to *Albany* from *CISM*
   - 2(a): ∇sie calculated using central difference
   - 2(b): ∇sie calculated using forward difference

ϕ_i^e =finite element shape functions

s_i^e = values of s at node i of element e

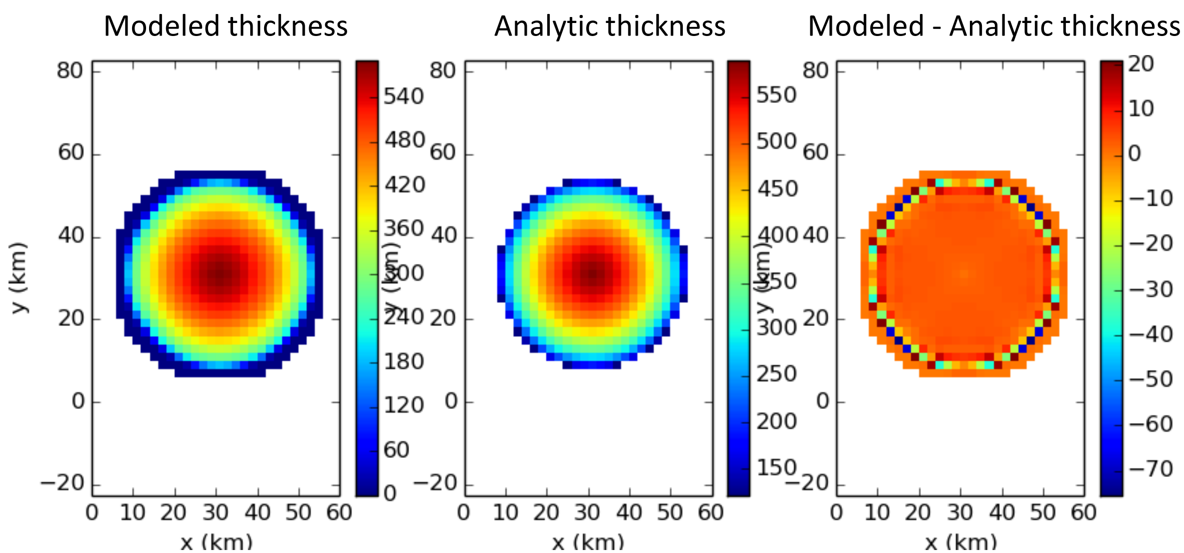

2(b).

Halfar test case (t = 100 years)

1 and 2(a): checkerboard pattern in thickness error.

2(b): checkerboard pattern in thickness error vanishes.

Observed by B. Lipscomb at CESM Annual Meeting 2014.

Sandia National Laboratories

Appendix: Bayesian Inversion/UQ

Difficulty in UQ: "Curse of Dimensionality"
The β-field inversion problem has $O(20{,}000)$ dimensions!

- **Step 1:** Model reduction (from $O(20{,}000)$ parameters to $O(5)$ parameters) using ***Karhunen-Loeve Expansion*** (or ***eigenvectors of Hessian***, in future) of basal sliding field:

$$log\big(\beta(\omega)\big) = \bar{\beta} + \sum_{k=1}^{K} \sqrt{\lambda_k}\boldsymbol{\phi}_k \xi_k(\omega)$$

- **Step 2:** ***Polynomial Chaos Expansion (PCE)*** emulator for mismatch over surface velocity discrepancy.

- **Step 3:** ***Markov Chain Monte Carlo (MCMC)*** calibration using PCE emulator.

With:
J. Jakeman,
M. Eldred (SNL)

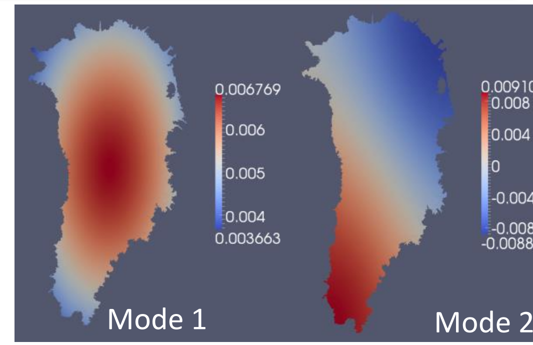

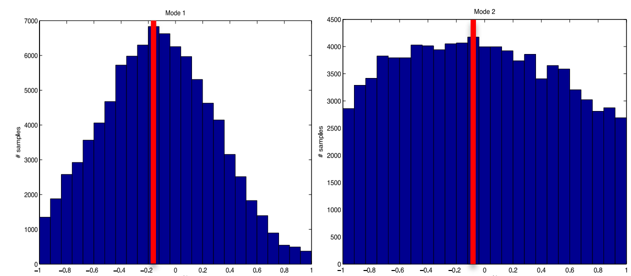

Posterior Distributions of 1st 2 KLE coefficients