

Multilevel Methods for Forward and Inverse Ice Sheet Modeling

Tobin Isaac

Institute for Computational Engineering & Sciences
The University of Texas at Austin

SIAM CSE 2015
Salt Lake City, Utah
 $\frac{7}{2}$

1 Hybrid 2D/3D Adaptive Mesh Refinement

2 Anisotropic multigrid

My collaborators on this work are Georg Stadler, Noémi Petra, Johann Rudi, Carsten Burstedde and Omar Ghattas.


These slides can be found at www.ices.utexas.edu/~tisaac/slides.

We are interested in scalable, accurate modeling of the polar ice sheets:

- Scalable discretization and solution of the equations of ice dynamics (conservation of mass, momentum) in 3D (see our recent submission, (Isaac et al., 2014b)).
 - High-order, inf-sup stable discretizations: $\mathbb{Q}_k \times \mathbb{Q}_{k-2}$
 - Adaptively refined meshes
 - We develop fast solvers for the Stokes form of the equations, but preconditioning methods I present also apply for hydrostatic (Blatter-Pattyn)
- Scalable deterministic and statistical inversion for parameters from surface observations (see our recent submission, (Isaac et al., 2014a)).
 - Please stay for Noémi Petra's talk.

Mesh adaptivity and partitioning

In the beginning. . . (CSE13)



Conformal hexahedral mesh, 75:1
aspect ratio limit: 100K hexahedra



We used the p4est¹ library for
parallel adaptive mesh refinement.

- Fast routines for isotropic local refinement from a coarse conformal hexahedral mesh.
- 3D isotropic refinement → aspect ratio of discretization must be enforced in the coarse mesh
- Coarse mesh is duplicated meta-data for each process.

¹p4est.org



Mesh partitioning via space-filling curve (SFC)

Well-shaped partitions can be maintained for arbitrarily refined meshes by linearly partitioning a space-filling curve. . .



Mesh partitioning via space-filling curve (SFC)

. . . but this approach does not keep columns of elements together.

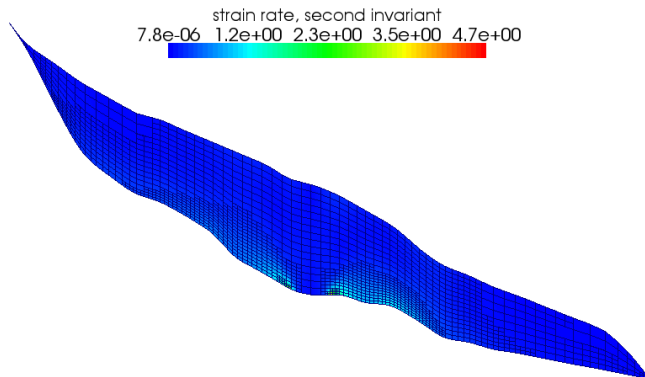


a periodic slab with variable mesh refinement, partitioned between 7 processes using a Morton-ordering space-filling curve (standard isotropic octree refinement)

One could. . .


. . . extrude a 2D mesh into a uniform vertical profile

- but more vertical resolution is needed at margins and the grounding line than in the bulk of the ice sheet or in the ice shelves.



Arolla glacier simulation with Dirichlet/Neumann transitions, second invariant refinement criterion.

A hybrid AMR scheme



- A p4est forest-of-quadtrees to manage columns,
- with each column stored as a flat, linear binary tree of layers,
- which guarantees column integrity.

An extension to p4est: hybrid routines have the prefix "p6est_", reproduce most of the standard p4est API, are documented on the website¹.

¹p4est.github.io/api


Comparison with isotropic octrees

Some good qualities of isotropic AMR are inherited

- Well-shaped partitions quickly generated by quadtree SFC,
 - using weighted partitioning algorithm already present in p4est (Burstedde et al., 2011).
- In-place, single-sweep vertical refinement and vertical coarsening algorithms: fast.

Some good qualities of isotropic AMR are lost


- Partitioning has coarser granularity
 - Load-balancing can theoretically be poorer,



- but only if there are columns with $N_{\text{column}} \gtrsim N_{\text{total}}/P$ layers.


Refinement trade-offs vs. isotropic AMR

Gain local control of element H:W aspect ratios:




Lose local horizontal refinement granularity:

- Individual layers cannot be horizontally refined: the mesh size grows quickly with horizontal refinement.



■ needs refinement ■ refinement unnecessary

Antarctic coarse mesh comparison



Coarse mesh is now topologically more complex (allows holes), but 27% as big.

Other applications

Well-suited for other climate and earth systems models.

- NUMA: Non-hydrostatic Unified Model of the Atmosphere² (Giraldo et al., 2013) is using p6est for partitioning (adaptivity in progress). Initial reports of scalability to 100K processes on Mira BG/Q.



numa
Non-hydrostatic Unified Model
of the Atmosphere

²faculty.nps.edu/fxgiraldo/projects/NUMA/Introduction_to_NUMA.html

Multilevel solvers for 3D ice sheet equations

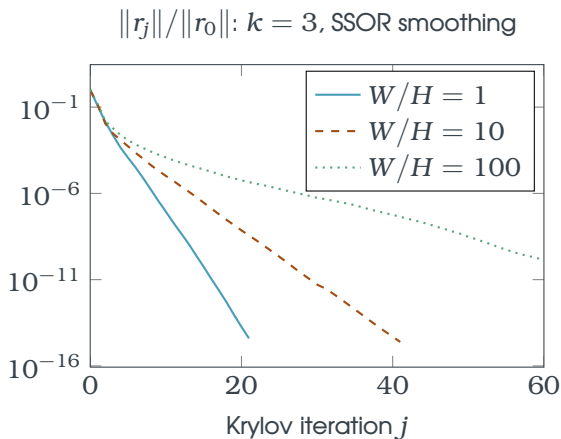
Overview

- Textbook multigrid efficiency demonstrated for hydrostatic equations (Brown et al., 2013):
 - Full-multigrid-like grid continuation to obtain initial guess for inexact Newton-Krylov, preconditioned by geometric multigrid V-cycle, demonstrated on topologically Cartesian meshes.
 - Geometric multigrid is much more efficient in terms of memory operations: intergrid / coarse-grid operations can be computed matrix-free, setup does not require Galerkin projection ($A_{\text{coarse}} \leftarrow PAP^T$).
- Our recent submission for solving Stokes equations (Isaac et al., 2014b):
 - Inexact Newton-Krylov preconditioned by smoothed aggregation algebraic multigrid (SA-AMG), demonstrated on Antarctic ice sheet.
 - Galerkin projection is more robust in some cases; applies to discretizations without a mesh hierarchy available.
- AMG+GMG (Sundar et al., 2012) (see Johann Rudi's poster at the Monday poster session) offers the chance to combine the good qualities of both.
 - An implementation for the hybrid AMR scheme is in development.

Multilevel solvers for 3D ice sheet equations

Common element: non-local smoothing

The large $W:H$ aspect ratio of ice sheet discretizations negatively affects the convergence factor of local smoothers: (point-block) Jacobi, SSOR




periodic slab test problem geometry, (1,1)-block of Stokes system

Incomplete factorization smoothing

Summary of geometric multigrid theory


$(2d + 1)$ -point stencil Laplacian: if the dofs are ordered first in the strong direction, GMG smoothed by lexically ordered ILU on each level converges independent of ε .

- The ordering is key: tridiagonal solves on tightly coupled subsystems.



Incomplete factorization smoothing

In practice with standard SA-AMG



$$\mathbf{u} \cdot \mathbf{n} = 0; (\mathbf{I} - \mathbf{n} \cdot \mathbf{n})(\boldsymbol{\sigma} + \beta \mathbf{u}) = \mathbf{0}$$


Convergence factor, standard SA-AMG (randomized MIS)

ℓ	$\beta = 1^0$	1^{-2}	1^{-4}	1^{-8}
0	0.14	0.14	0.57	0.63
1	0.17	0.27	0.75	0.78
2	0.20	0.51	0.82	0.83

(1,1)-block of Stokes system, ℓ levels of refinement, column-ordered dofs on the fine grid

Smoother/aggregate incompatibility

Strong β masks inefficient smoothing: without it, smoother/aggregate incompatibility is exposed.



- Column structure is not preserved.
- Discretization does not neatly separate into tightly coupled groups.
- “Smoothed” aggregates actually introduce a lot of high-frequency error.

ILU smoother is no longer effective on coarse grids.

Smoother/aggregate incompatibility

Strong β masks inefficient smoothing: without it, smoother/aggregate incompatibility is exposed.




- Column structure is not preserved.
- Discretization does not neatly separate into tightly coupled groups.
- “Smoothed” aggregates actually introduce a lot of high-frequency error.

ILU smoother is no longer effective on coarse grids.

Smoother/aggregate incompatibility

Strong β masks inefficient smoothing: without it, smoother/aggregate incompatibility is exposed.




- Column structure is not preserved.
- Discretization does not neatly separate into tightly coupled groups.
- “Smoothed” aggregates actually introduce a lot of high-frequency error.

ILU smoother is no longer effective on coarse grids.

Smoother/aggregate incompatibility

Strong β masks inefficient smoothing: without it, smoother/aggregate incompatibility is exposed.




- Column structure is not preserved.
- Discretization does not neatly separate into tightly coupled groups.
- “Smoothed” aggregates actually introduce a lot of high-frequency error.

ILU smoother is no longer effective on coarse grids.

Constructing column-preserving aggregates: DofColumns

Plugin³ for PETSc's GAMG preconditioner:

- Describe column structure of matrix A with `MatSetDofColumns(A, ...)`
- Add `-pc_gamg_type dofcol` to command line
 - Creates aggregate columns, partitions columns to make node aggregates, preserving column structure



³bitbucket.org/tisaac/DofColumns

Incomplete factorization smoothing

In practice with column preserving SA-AMG

Convergence factor, standard SA-AMG (3D randomized MIS)

ℓ	$\beta = 1^0$	1^{-2}	1^{-4}	1^{-8}
0	0.14	0.14	0.57	0.63
1	0.17	0.27	0.75	0.78
2	0.20	0.51	0.82	0.83

Convergence factor, DoFColumns SA-AMG (2D randomized MIS)

ℓ	$\beta = 1^0$	1^{-2}	1^{-4}	1^{-8}
0	0.14	0.14	0.30	0.35
1	0.17	0.20	0.39	0.47
2	0.19	0.25	0.48	0.54

(1,1)-block system, ℓ levels of refinement, column-ordered dofs on the fine grid

Incomplete factorization smoothing

Direct comparison with SSOR on Antarctic ice sheet

		SSOR, standard SA-AMG		IC(0), column-preserving SA-AMG	
$\#N_{\text{dof}}$	P	$\#N$	$\#K$	$\#N$	$\#K$
38M	1,024	8	147	7	85
268M	8,192	9	243	8	98

Comparison for solve to 10^{-12} of nonlinear Stokes problem, described in (Isaac et al., 2014a), on the Antarctic ice sheet, on the same mesh.

Incomplete factorization smoothing

One final note

Either (1,1)-block of Stokes or hydrostatic, the existence of an SPD incomplete factorization without pivoting is not guaranteed: a strategy is needed.

- PETSc's default behavior is the Manteuffel shift (`-pc_factor_shift_type positive_definite`): computing a factor of

$$\tilde{F} = F + \alpha I$$

for sufficiently large α .

- For variable diagonal sizes (if you have elements with different volumes, very different coefficients), this tends to shift high-frequency error components out of the region of the spectrum damped by the smoother.
- Manifestation: a Newton step where convergence stalls.
- Shifting "in blocks" (`-pc_factor_shift_type in_blocks`)—replacing small/negative pivots with a minimum value—may locally reduce the effectiveness of the smoother, but if this happens in only a handful of places, then an outer Krylov method can capture those poorly smoothed modes.

Summary

We have developed two tools for the multilevel discretization and solution of 3D PDEs in thin domains, and demonstrated them on the Stokes equations for ice sheet dynamics, though they can be applied in broader settings:

- `p6est`: Scalable parallel hybrid quadtree/octree AMR.
- `DofColumns`: column-preserving extension for PCGAMG.

Thank you!

References I

- Jed Brown, Barry Smith, and Aron Ahmadi. Achieving textbook multigrid efficiency for hydrostatic ice sheet flow. *SIAM Journal on Scientific Computing*, 35(2):B359–B375, 2013.
- Carsten Burstedde, Lucas C. Wilcox, and Omar Ghattas. p4est: Scalable algorithms for parallel adaptive mesh refinement on forests of octrees. *SIAM Journal on Scientific Computing*, 33(3):1103–1133, 2011. doi: 10.1137/100791634.
- Francis X Giraldo, JF Kelly, and EM Constantinescu. Implicit-explicit formulations for a 3D nonhydrostatic unified model of the atmosphere (NUMA). *SIAM Journal of Scientific Computing*, 2013.
- Tobin Isaac, Noemi Petra, Georg Stadler, and Omar 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* (submitted). <http://arxiv.org/abs/1410.1221>, 2014a.

References II

Tobin Isaac, Georg Stadler, and Omar 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* (submitted), 2014b.

<http://arxiv.org/abs/1406.6573>.

Hari Sundar, George Biros, Carsten Burstedde, Johann Rudi, Omar Ghattas, and Georg Stadler. Parallel geometric-algebraic multigrid on unstructured forests of octrees. In *SC12: Proceedings of the International Conference for High Performance Computing, Networking, Storage and Analysis*, Salt Lake City, UT, 2012. ACM/IEEE.