Exceptional service in the national interest





Performance & Performance Portability of the Albany/FELIX Finite Element Land-Ice Solver

I. Tezaur¹, J. Watkins¹, R. Tuminaro¹, I. Demeshko²

¹ Sandia National Laboratories, Livermore, CA, USA.

² Los Alamos National Laboratory, Los Alamos, NM, USA.

SIAM GS 2017 Erlangen, Germany September 11-14, 2017





Sandia National Laboratories is a multi-mission laboratory managed and operated by National Technology and Engineering Solutions of Sandia, LLC., a wholly owned subsidiary of Honeywell International, Inc., for the U.S. Department of Energy's National Nuclear Security Administration under contract DE-NA0003525.

SAND2017-9362C



Outline

- 1. Background.
 - Motivation.
 - PISCEES project for land-ice modeling.
 - Albany/FELIX "First-Order Stokes" land-ice model.
- 2. Finite Element Assembly.
 - Performance-portability via Kokkos.
- 3. Linear Solve.
 - AMG preconditioning.
- 4. Summary & future work.







Outline

- 1. Background.
 - Motivation.
 - PISCEES project for land-ice modeling.
 - Albany/FELIX "First-Order Stokes" land-ice model.
- 2. Finite Element Assembly.
 - Performance-portability via Kokkos.
- 3. Linear Solve.
 - AMG preconditioning.
- 4. Summary & future work.





Motivation



- Scientific models (e.g. climate models) need more *computational power* to achieve *higher resolutions*.
- High performance computing (HPC) architectures are becoming increasingly more *heterogeneous* in a move towards *exascale*.
- Climate models need to adapt to execute *correctly* & *efficiently* on new HPC architectures with drastically *different memory models*.





MPI+X Programming Model



- HPC architectures are rapidly changing, but *trends* remain the same.
 - Computations are cheap, memory transfer is expensive.
 - *Single core cycle* time has improved but stagnated.
 - Increased *computational power* achieved through *manycore architectures*.
 - → MPI-only is not enough to exploit emerging massively parallel architectures.

Year

1980s

Today

Memory

Access Time

~100 ns

~50-100 ns

Approach: MPI+X Programming Model

- MPI: inter-node parallelism.
- X: intra-node parallelism.
 - \rightarrow *Examples:* X = OpenMP, CUDA, Pthreads, etc.



Single Core

Cycle Time

~100 ns

~1 ns



Outline

1. Background.

- Motivation.
- PISCEES project for land-ice modeling.
- Albany/FELIX "First-Order Stokes" land-ice model.
- 2. Finite Element Assembly.
 - Performance-portability via Kokkos.
- 3. Linear Solve.
 - AMG preconditioning.
- 4. Summary & future work.





PISCEES Project for Land-Ice Modeling





"PISCEES" = Predicting Ice Sheet Climate Evolution at Extreme Scales 5 year SciDAC3 project (2012-2017).

<u>Sandia's Role in the PISCEES Project</u>: to develop and support a robust and scalable land ice solver based on the "First-Order" (FO) Stokes equations \rightarrow Albany/FELIX*

Requirements for Albany/FELIX:

- Unstructured grid meshes.
- Scalable, fast and robust.
- Verified and validated.
- Portable to new architecture machines.
- Advanced analysis capabilities: deterministic inversion, calibration, uncertainty quantification.

As part of **ACME** *DOE Earth System Model*, solver will provide actionable predictions of 21st century sea-level change (including uncertainty bounds).





Outline

1. Background.

- Motivation.
- PISCEES project for land-ice modeling.
- Albany/FELIX "First-Order Stokes" land-ice model.
- 2. Finite Element Assembly.
 - Performance-portability via Kokkos.
- 3. Linear Solve.
 - AMG preconditioning.
- 4. Summary & future work.





$\begin{cases} -\nabla \cdot (2\mu\dot{\epsilon}_{1}) = -\rho g \frac{\partial s}{\partial x} \\ -\nabla \cdot (2\mu\dot{\epsilon}_{2}) = -\rho g \frac{\partial s}{\partial x} \end{cases}$ $\mu = \frac{1}{2} A^{-\frac{1}{n}} \left(\frac{1}{2} \sum_{ij} \dot{\epsilon}_{ij}^{2} \right)^{\left(\frac{1}{2n} - \frac{1}{2}\right)}$

Algorithmic choices for Albany/FELIX:

- **3D unstructured grid** FEM discretization.
- *Newton method* nonlinear solver with automatic differentiation Jacobians.
- Preconditioned *Krylov iterative linear* solvers.
- Advanced analysis capabilities: deterministic inversion, calibration, UQ.

Albany/FELIX implemented in open-source* multi-physics FE Trilinos-based code:

* https://github.com/gahansen/Albany.







Ice sheet



First-Order (FO) Stokes Model

• Ice velocities given by the "First-Order" Stokes PDEs with nonlinear viscosity:

$$\begin{cases} -\nabla \cdot (2\mu \dot{\boldsymbol{\epsilon}}_{1}) = -\rho g \frac{\partial s}{\partial x} \\ -\nabla \cdot (2\mu \dot{\boldsymbol{\epsilon}}_{2}) = -\rho g \frac{\partial s}{\partial y} \end{cases}$$

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

$$\dot{\epsilon}_1^T = (2\dot{\epsilon}_{11} + \dot{\epsilon}_{22}, \dot{\epsilon}_{12}, \dot{\epsilon}_{13})$$
$$\dot{\epsilon}_2^T = (2\dot{\epsilon}_{12}, \dot{\epsilon}_{11} + 2\dot{\epsilon}_{22}, \dot{\epsilon}_{23})$$
$$\dot{\epsilon}_{ij} = \frac{1}{2} \left(\frac{\partial u_i}{\partial x_j} + \frac{\partial u_j}{\partial x_i} \right)$$

Algorithmic choices for Albany/FELIX:

- **3D unstructured grid** FEM discretization.
- Newton method nonlinear solver with automatic differentiation Jacobians.
- Preconditioned *Krylov iterative linear* solvers.
- *Advanced analysis capabilities*: deterministic inversion, calibration, UQ.

Albany/FELIX implemented in open-source* multi-physics FE Trilinos-based code:



* https://github.com/gahansen/Albany. **Finite Element Assembly





Outline

- 1. Background.
 - Motivation.
 - PISCEES project for land-ice modeling.
 - Albany/FELIX "First-Order Stokes" land-ice model.
- 2. Finite Element Assembly.
 - Performance-portability via Kokkos.
 - FEA = 50% CPU-time

- 3. Linear Solve.
 - AMG preconditioning.
- 4. Summary & future work.





Albany/FELIX Finite Element Assembly (FEA)

- *Gather operation* extracts solution values out of global solution vector.
- Physics *evaluator* functions operate on *workset* of elements, store evaluated quantities in local field arrays.
- FEA relies on *template based generic programming* + *automatic differentiation* for Jacobians, tangents, etc.
- *Scatter operation* adds local residual, Jacobian to global residual, Jacobian.

Performance-portability: focus on FEA.



| Problem Type | % CPU time for FEA | |
|--------------|--------------------|--|
| Implicit | 50% | |
| Explicit | 99% | |

Albany/FELIX Finite Element Assembly (FEA)

- *Gather operation* extracts solution values out of global solution vector.
- Physics *evaluator* functions operate on *workset* of elements, store evaluated quantities in local field arrays.
- FEA relies on *template based generic programming* + *automatic differentiation* for Jacobians, tangents, etc.
- Scatter operation adds local residual, Jacobian to global residual, Jacobian.

Performance-portability: focus on FEA.

- MPI-only FEA:
 - Each MPI process has workset of cells & computes nested parallel for loops.



| Problem Type | % CPU time for FEA | |
|--------------|--------------------|--|
| Implicit | 50% | |
| Explicit | 99% | |

Albany/FELIX Finite Element Assembly (FEA)

- *Gather operation* extracts solution values out of global solution vector.
- Physics *evaluator* functions operate on *workset* of elements, store evaluated quantities in local field arrays.
- FEA relies on *template based generic programming* + *automatic differentiation* for Jacobians, tangents, etc.
- Scatter operation adds local residual, Jacobian to global residual, Jacobian.

Performance-portability: focus on FEA.

- MPI-only FEA:
 - Each MPI process has workset of cells & computes nested parallel for loops.
- MPI+X FEA:
 - Each MPI process has workset of cells.
 - Multi-dimensional parallelism with +X (X=OpenMP, CUDA) for nested parallel for loops.



| Problem Type | % CPU time for FEA | |
|--------------|--------------------|--|
| Implicit | 50% | |
| Explicit | 99% | |

Performance-portability via Kokkos

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

- In Albany/FELIX, we achieve performance-portability via *Kokkos*.
 - *Kokkos:* C++ library and programming model that provides performance portability across multiple computing architectures.

 \rightarrow *Examples:* Multicore CPU, NVIDIA GPU, Intel Xeon Phi, and more.

- Provides *automatic access* to OpenMP, CUDA, Pthreads, etc.
- Designed to work with the *MPI+X* programming model.
- Abstracts *data layouts* for optimal performance ("array of strucs" vs. struct of arrays", locality).

With *Kokkos*, you write an algorithm once, and just change a template parameter to get the optimal data layout for your hardware.

 \rightarrow Allows researcher to focus on *application development* for large heterogeneous architectures.







• *MPI-only* nested for loop:

for (int cell=0; cell<numCells; ++cell)
for (int node=0; node<numNodes; ++node)
for (int qp=0; qp<numQPs; ++qp)
compute A; MPI process n</pre>





 Multi-dimensional parallelism for nested for loops via Kokkos:

for (int cell=0; cell<numCells; ++cell) for (int node=0; node<numNodes; ++nod for (int qp=0; qp<numQPs; ++qp)

compute A;

/IPI process *n*

Thread 1 computes A for (cell,node,qp)=(0,0,0)

Thread 2 computes A for (cell,node,qp)=(0,0,1)

Thread N computes A for (cell,node,qp)=(numCells,numNodes,numQPs)

computeA_Policy range({0,0,0},{(int)numCells,(int)numNodes,(int)numQPs}); Kokkos::Experimental::md_parallel_for<ExecutionSpace>(range,*this);



 Multi-dimensional parallelism for nested for loops via Kokkos:

for (int cell=0; cell<numCells; ++cell) for (int node=0; node<numNodes; ++node for (int qp=0; qp<numQPs; ++qp)

ompute A;

/IPI process n

Thread 1 computes A for (cell,node,qp)=(0,0,0)

Thread 2 computes A for (cell,node,qp)=(0,0,1)

Thread N computes A for (cell,node,qp)=(numCells,numNodes,numQPs)

computeA_Policy range({0,0,0},{(int)numCells,(int)numNodes,(int)numQPs});
Kokkos::Experimental::md_parallel_for<ExecutionSpace>(range,*this);

ExecutionSpace defined at compile time, e.g.
 typedef Kokkos::OpenMP ExecutionSpace; //MPI+OpenMP
 typedef Kokkos::CUDA ExecutionSpace; //MPI+CUDA
 typedef Kokkos::Serial ExecutionSpace; //MPI-only



 Multi-dimensional parallelism for nested for loops via Kokkos:

for (int cell=0; cell<numCells; ++cell) for (int node=0; node<numNodes; ++node for (int qp=0; qp<numQPs; ++qp)

ite A;

/IPI process n

Thread 1 computes A for (cell,node,qp)=(0,0,0)

Thread 2 computes A for (cell,node,qp)=(0,0,1)

Thread N computes A for (cell,node,qp)=(numCells,numNodes,numQPs)

computeA_Policy range({0,0,0},{(int)numCells,(int)numNodes,(int)numQPs});
Kokkos::Experimental::md_parallel_for<ExecutionSpace>(range,*this);

- ExecutionSpace defined at *compile time*, e.g. typedef Kokkos::OpenMP ExecutionSpace; //MPI+OpenMP typedef Kokkos::CUDA ExecutionSpace; //MPI+CUDA typedef Kokkos::Serial ExecutionSpace; //MPI-only
- Atomics used to scatter local data to global data structures Kokkos::atomic_fetch_add



 Multi-dimensional parallelism for nested for loops via Kokkos:

for (int cell=0; cell<numCells; ++cell) for (int node=0; node<numNodes; ++node for (int qp=0; qp<numQPs; ++qp)

1

Thread 1 computes A for (cell,node,qp)=(0,0,0)

Thread 2 computes A for (cell,node,qp)=(0,0,1)

Thread N computes A for (cell,node,qp)=(numCells,numNodes,numQPs)

computeA_Policy range({0,0,0},{(int)numCells,(int)numNodes,(int)numQPs});
Kokkos::Experimental::md_parallel_for<ExecutionSpace>(range,*this);

- ExecutionSpace defined at *compile time*, e.g. typedef Kokkos::OpenMP ExecutionSpace; //MPI+OpenMP typedef Kokkos::CUDA ExecutionSpace; //MPI+CUDA typedef Kokkos::Serial ExecutionSpace; //MPI-only
- Atomics used to scatter local data to global data structures Kokkos::atomic_fetch_add
- For MPI+CUDA, data transfer from host to device handled by **CUDA UVM***.





Computer Architectures

Performance-portability of FEA in Albany has been tested across *multiple architectures*: Intel Sandy Bridge, IBM Power8, Keplar/Pascal GPUs, KNL Xeon Phi

- **Ride** (SNL) used for verification, performance tests 12 nodes (dual-Power8 (16 cores) + P100 quad-GPU)
- Bowman (SNL) used for verification
 10 nodes (Intel Xeon Phi KNL (68 cores))
- **Cori** (NERSC) used for verification, performance tests 9688 nodes (Intel Xeon Phi KNL (68 cores))
- **Summit** (ORLCF) is ultimate GPU target 4600 nodes (dual-Power9 + 6 NVIDIA Volta)









Computer Architectures

Performance-portability of FEA in Albany has been tested across *multiple architectures*: Intel Sandy Bridge, IBM Power8, Keplar/Pascal GPUs, KNL Xeon Phi

- **Ride** (SNL) used for verification, performance tests 12 nodes (dual-Power8 (16 cores) + P100 quad-GPU)
- Bowman (SNL) used for verification
 10 nodes (Intel Xeon Phi KNL (68 cores))
- **Cori** (NERSC) used for verification, performance tests 9688 nodes (Intel Xeon Phi KNL (68 cores))
- **Summit** (ORLCF) is ultimate GPU target 4600 nodes (dual-Power9 + 6 NVIDIA Volta)









Platforms utilized here.





<u>A single node comparison</u>: 16MPI vs. 4(MPI+GPU) [*Left*], 68MPI vs. 68(MPI+40MP) [*Right*] (GIS 4km-20km unstructured mesh with 1.51M tet elements)

Blue (evaluateFields): mostly Residual/Jacobian computation + Gather/Scatter (Albany/FELIX); **Yellow** (GlobalAssembly): mostly communication + Trilinos operations.

GIS Kokkos 16 MPI vs. 4(MPI+GPU) Results



<u>GlobalAssembly (yellow)</u>: mostly communication + Trilinos operations

- 8x slow-down is not reasonable.
- Most slow-downs are in Trilinos (Tpetra)
 → Trilinos packages are currently being reworked using CUDA-aware MPI*.

Summary: speed-ups on GPU are not yet as expected but may be improved by introducing padding, removing CUDA UVM and unnecessary data movement, switching to CUDA-aware MPI.

evaluateFields (blue): mostly residual/Jacobian computation + Gather/Scatter

- evaluateFields<Jacobian> much faster than evaluateFields<Residual> b/c there is more work in computing Jacobian.
- 2x speedup not much considering 4 GPUs (desirable speedup: 10x or more).
 - Data movement is lagging speedup can be improved by removing CUDA
 UVM, data padding to prevent data misalignment.
 - A few kernels still for boundary conditions still need to be ported to Kokkos.

* With CUDA-aware MPI, GPU buffers can be passed directly to MPI w/o staging using cudaMemcpy.

GIS Kokkos 68MPI vs. 68(MPI+40MP) Results



GlobalAssembly (yellow): mostly communication + Trilinos operations

 1.5x slowdown is in Jacobian assembly not Residual assembly → Trilinos team is investigating this now.

Summary: besides the global Jacobian assembly, results are promising. More studies are needed once we profile nonlinear/linear solvers.

evaluateFields (blue): mostly residual/Jacobian computation + Gather/Scatter

- 1.2x speedup from hardware threads is reasonable (there are 2 VPUs/KNL core, so 2x speedup is ideal but will be limited by L1 cache size in core for bandwidth bound operation)
- Once the nonlinear/linear solver is included, more OpenMP threads will likely be used (e.g. 4(MPI+68OMP)) to improve speedup.
 - More OpenMP threads on cores in FEA reduces performance because it takes away from coarser grain MPI parallelism

GIS MPI+Device Weak Scalability





- Scalability of **GlobalAssembly + evaluateFields** is studied.
- Weak Scalability is for GIS 4km-20km and 1km-7km (1.51M and 14.4M elements) tet meshes.
- Weak scalability is comparable for P100 and KNL.
 - KNL performs **better** because of heavy use of MPI.
- Optimizations/profiling required for strong scalability of FELIX; we have demonstrated strong scalability for other applications in Albany (I. Demeshko *et al* 2017).



Outline

- 1. Background.
 - Motivation.
 - PISCEES project for land-ice modeling.
 - Albany/FELIX "First-Order Stokes" land-ice model.
- 2. Finite Element Assembly.
 - Performance-portability via Kokkos.
 - Linear Solve = 50% CPU-time

- 3. Linear Solve.
 - AMG preconditioning.
- 4. Summary & future work.













Scalability results are *not* acceptable!





Scalability results are *not* acceptable!

Why is scalability so **bad** for out-of-the-box preconditioners?

- 1. Ice sheet geometries have **bad aspect ratios** $(dx \gg dz)$.
- 2. Ice shelves give rise to severely ill-conditioned matrices.
- 3. Islands and hinged peninsulas lead to solver failures.





Scalability results are *not* acceptable!

Why is scalability so **bad** for out-of-the-box preconditioners?

- 1. Ice sheet geometries have **bad aspect ratios** $(dx \gg dz)$.
- 2. Ice shelves give rise to severely ill-conditioned matrices.
- 3. Islands and hinged peninsulas lead to solver failures.

We **<u>mitigate</u>** these difficulties through the development of:

- New AMG* preconditioner based on semi-coarsening.
- Island/hinge removal algorithm.

* Algebraic Multi-Grid.



Outline

- 1. Background.
 - Motivation.
 - PISCEES project for land-ice modeling.
 - Albany/FELIX "First-Order Stokes" land-ice model.
- 2. Finite Element Assembly.
 - Performance-portability via Kokkos.
 - FEA = 50% CPU-time

- 3. Linear Solve.
 - AMG preconditioning.
- 4. Summary & future work.





Scalability via Algebraic Multi-Grid Preconditioning with Semi-Coarsening



Bad aspect ratios $(dx \gg dz)$ ruin classical AMG convergence rates!

- relatively small horizontal coupling terms, hard to smooth horizontal errors
- \Rightarrow Solvers (AMG and ILU) must take aspect ratios into account

We developed a **new AMG solver** based on aggressive **semi-coarsening** (available in *ML/MueLu* packages of *Trilinos*)





See (Tezaur *et al.,* 2015), (Tuminaro *et al.,* 2016). Scaling studies (next slides): New AMG preconditioner vs. ILU

Greenland Controlled Weak Scalability Study



- 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 AMG preconditioner developed by R. Tuminaro based on semi-coarsening (coarsening in z-direction only).
- *Significant improvement* in scalability with new AMG preconditioner over ILU preconditioner!

Greenland Controlled Weak Scalability Study



Moderate Resolution Antarctica Weak Scaling Study



Antarctica is fundamentally different than Greenland: AIS contains large ice shelves (floating extensions of land ice).

- Along ice shelf front: open-ocean BC (Neumann).
- Along ice shelf base: zero traction BC (Neumann).

 \Rightarrow For vertical grid lines that lie within ice shelves, top and bottom BCs resemble Neumann BCs so sub-matrix associated with one of these lines is almost* singular.



⇒ Ice shelves give rise to severe illconditioning of linear systems!





*Completely singular in the presence of islands and some ice tongues.

Moderate Resolution Antarctica Weak Scaling Study

- Weak scaling study on Antarctic problem (8km w/ 5 layers \rightarrow 2km w/ 20 layers).
- Initialized with realistic basal friction (from deterministic inversion) and temperature field from BEDMAP2.
- Iterative linear solver: GMRES.
- **Preconditioner**: ILU vs. new AMG based on aggressive semi-coarsening.







Outline

- 1. Background.
 - Motivation.
 - PISCEES project for land-ice modeling.
 - Albany/FELIX "First-Order Stokes" land-ice model.
- 2. Finite Element Assembly.
 - Performance-portability via Kokkos.
- 3. Linear Solve.
 - AMG preconditioning.
- 4. Summary & future work.





Summary & Conclusions



- A *performance portable* implementation of the FEA in the *FELIX land-ice* model was created using *Kokkos* within the Albany code base.
 - With this implementation, the *same code* can run on devices with drastically *different memory models* (many-core CPU, NVIDIA GPU, Intel Xeon Phi, etc.).
 - Performance studies show that *further optimization* is needed to fully utilize all resources.

More on *performance-portability* of Albany using *Kokkos* can be found here: <u>https://github.com/gahansen/Albany/wiki/Albany-</u> <u>performance-on-next-generation-platforms</u>

Summary & Conclusions



- A *performance portable* implementation of the FEA in the *FELIX land-ice* model was created using *Kokkos* within the Albany code base.
 - With this implementation, the *same code* can run on devices with drastically *different memory models* (many-core CPU, NVIDIA GPU, Intel Xeon Phi, etc.).
 - Performance studies show that *further optimization* is needed to fully utilize all resources.

More on *performance-portability* of Albany using *Kokkos* can be found here: <u>https://github.com/gahansen/Albany/wiki/Albany-</u> performance-on-next-generation-platforms

- **Scalable, fast** and **robust** linear solve is achieved via algebraic multigrid (AMG) preconditioner that takes advantage of layered nature of meshes.
 - Performance portability of linear solve is work in progress.

Summary & Conclusions



- A *performance portable* implementation of the FEA in the *FELIX land-ice* model was created using *Kokkos* within the Albany code base.
 - With this implementation, the *same code* can run on devices with drastically *different memory models* (many-core CPU, NVIDIA GPU, Intel Xeon Phi, etc.).
 - Performance studies show that *further optimization* is needed to fully utilize all resources.

More on *performance-portability* of Albany using *Kokkos* can be found here: <u>https://github.com/gahansen/Albany/wiki/Albany-</u> performance-on-next-generation-platforms

- **Scalable, fast** and **robust** linear solve is achieved via algebraic multigrid (AMG) preconditioner that takes advantage of layered nature of meshes.
 - Performance portability of linear solve is work in progress.

Heterogeneous HPC architectures can now be utilized for land-ice research using Albany/FELIX.

Ongoing/Future Work

Finite Element Assembly (FEA):

- **Profiling** using TAU and nvprof.
- Methods for *improving performance*:
 - Reduce excess memory usage.
 - Utilize shared memory.
 - Replace CUDA UVM with manual memory transfer.
 - Parallelize over nodes, quad points, levels in addition to cells.
 - Add data padding to prevent misalignment.
 - Switch to CUDA-aware MPI.
- Large-scale runs on Cori and Summit.

Linear Solve:

- Performance-portability of *preconditioned iterative linear solve* using *Kokkos* for implicit problems in Albany (e.g., FELIX).
 - Finish converting MueLu/Ifpack2 to use Kokkos.
 - Algorithm redesign may be necessary for GPUs.
 - Considering other solvers, e.g., hierarchical solvers, Shylu (FAST-ILU, multi-threaded GS).









Funding/Acknowledgements





PISCEES team members: K. Evans, M. Gunzburger, M. Hoffman, C. Jackson, P. Jones, W. Lipscomb, M. Perego, S. Price, A. Salinger, I. Tezaur, R. Tuminaro, P. Worley.
 Trilinos/DAKOTA collaborators: M. Eldred, J. Jakeman, E. Phipps, L. Swiler.
 Computing resources: NERSC, OLCF.

References



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

[2] A. Salinger, *et al.* "Albany: Using Agile Components to Develop a Flexible, Generic Multiphysics Analysis Code", *Int. J. Multiscale Comput. Engng.* 14(4) (2016) 415-438.

[3] **I. Tezaur**, 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*. 8 (2015) 1-24.

[4] C. Edwards, C. Trott, D. Sunderland. "Kokkos: Enabling manycore performance portability through polymorphic memory access patterns", *J. Par. & Distr. Comput.* **74** (12) (2014) 3202-3216.

[5] R. Tuminaro, M. Perego, **I. Tezaur**, A. Salinger, S. Price. "A matrix dependent/algebraic multigrid approach for extruded meshes with applications to ice sheet modeling", *SIAM J. Sci. Comput.* 38(5) (2016) C504-C532.

[6] **I. Tezaur**, 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", *Procedia Computer Science*, 51 (2015) 2026-2035.

[7] I. Demeshko, J. Watkins, **I. Tezaur**, O. Guba, W. Spotz, A. Salinger, R. Pawlowski, M. Heroux. "Towards performance-portability of the Albany finite element analysis code using the Kokkos library", submitted to *J. HPC Appl.*

[8] S. Price, M. Hoffman, J. Bonin, T. Neumann, I. Howat, J. Guerber, **I. Tezaur**, J. Saba, J. Lanaerts, D. Chambers, W. Lipscomb, M. Perego, A. Salinger, R. Tuminaro. "An ice sheet model validation framework for the Greenland ice sheet", *Geosci. Model Dev.* 10 (2017) 255-270

Appendix: Parallelism on Modern Hardware Sandia Laboratories

| Year | Memory Access Time | Single Core Cycle Time |
|-------|--------------------|------------------------|
| 1980s | ~100 ns | ~100 ns |
| Today | ~50-100 ns | ~1 ns |

- *Memory access time* has remained the *same*.
- *Single core* performance has *improved* but *stagnated*.
- Computations are cheap, memory transfer is expensive.
- *More performance* from *multicore/manycore* processors.

Appendix:

FO-Stokes model is implemented within Albany, Sandia open-source* parallel, C++, multi-physics finite element code → Albany/FELIX**.

FO-Stokes model is implemented within Albany, Sandia open-source^{*} parallel, C++, multi-physics finite element code \rightarrow **Albany/FELIX**^{**}.

- Component-based design for rapid development of new physics & capabilities.
- Extensive use of libraries from the opensource *Trilinos* project:
 - Automatic differentiation.
 - Discretizations/meshes, mesh adaptivity.
 - Solvers, preconditioners.
 - Performance-portable kernels.
- Advanced analysis capabilities:
 - Parameter estimation.
 - Uncertainty quantification (DAKOTA).
 - Optimization.
 - Sensitivity analysis.

| Analysis Tools | Mesh |
|-------------------|--------------------|
| (DIACK-DOX) | M |
| Optimization | 1 Inlin |
| UQ (sampling | <mark>j)</mark> Pa |
| Parameter Stud | lies Load |
| Calibration | A |
| Reliability | Grid |
| Composite Physi | ics Quality |
| MultiPhysics Cou | |
| System UQ | |
| | |
| Analysis Tools | Mesh |
| (embedded) | Mes |
| Nonlinear Solve | er Geom |
| Time Integration | n Solut |
| Continuation | Check |
| Sensitivity Analy | sis |
| Stability Analys | is |
| Constrained Solv | /es |
| Optimization | Disc |
| UQ Solver | Discr |
| | Fi |
| Linear Algebra | |
| Data Structure: | s Deriv |
| Iterative Solver | s S |
| Direct Solvers | |
| Eigen Solver | |
| Preconditioner | S D |
| Multi-Level Metho | ods |

Multiphysics Code





| h Tools | Utilities |
|--------------------|--------------------------|
| Mesh I/O | Input File Parser |
| ine Meshing | Parameter List |
| Partitioning | Memory Management |
| ad Balancing | I/O Management |
| Adaptivity | Communicators |
| id Transfers | Runtime Compiler |
| ty Improvement | Architecture- |
| Search | Dependent Kernels |
| DOF map | Multi-Core |
| | Accelerators |
| h Database | |
| esh Database | Post Processing |
| metry Database | In-situ Visualization |
| ution Database | Verification |
| eckpoint/Restart | QOI Computation |
| | Model Reduction |
| | _ocal Fill |
| scretizations | <u>zaranovalli</u> |
| cretization Libra | Physics Fill |
| Field Manager | |
| ui ve ti ve Te e l | Source Terms |
| | BCs |
| Sensitivities | Material Models |
| Derivatives | Responses |
| Adjoints | Parameters |
| Propagation | |
| | |

40+ packages; 120+ libraries

Appendix: First-Order (FO) Stokes Model



- Ice behaves like a very *viscous shear-thinning fluid* (similar to lava flow).
- Quasi-static model with momentum balance given by "First-Order" Stokes PDEs: "nice" elliptic approximation* to Stokes' flow equations.

$$\begin{cases} -\nabla \cdot (2\mu \dot{\boldsymbol{\epsilon}}_1) = -\rho g \frac{\partial s}{\partial x} \\ -\nabla \cdot (2\mu \dot{\boldsymbol{\epsilon}}_2) = -\rho g \frac{\partial s}{\partial y} \end{cases}, \quad \text{in } \Omega \end{cases}$$

$$\dot{\boldsymbol{\epsilon}}_{1}^{T} = (2\dot{\boldsymbol{\epsilon}}_{11} + \dot{\boldsymbol{\epsilon}}_{22}, \dot{\boldsymbol{\epsilon}}_{12}, \dot{\boldsymbol{\epsilon}}_{13})$$
$$\dot{\boldsymbol{\epsilon}}_{2}^{T} = (2\dot{\boldsymbol{\epsilon}}_{12}, \dot{\boldsymbol{\epsilon}}_{11} + 2\dot{\boldsymbol{\epsilon}}_{22}, \dot{\boldsymbol{\epsilon}}_{23})$$
$$\dot{\boldsymbol{\epsilon}}_{ij} = \frac{1}{2} \left(\frac{\partial u_{i}}{\partial x_{j}} + \frac{\partial u_{j}}{\partial x_{i}} \right)$$

• Viscosity μ is nonlinear function given by "*Glen's law*":

$$\mu = \frac{1}{2}A(T)^{-\frac{1}{n}} \left(\frac{1}{2}\sum_{ij} \dot{\epsilon}_{ij}^{2}\right)^{\left(\frac{1}{2n} - \frac{1}{2}\right)} \qquad (n = 3)$$

- Relevant boundary conditions:
 - Stress-free BC: $2\mu \dot{\boldsymbol{\epsilon}}_i \cdot \boldsymbol{n} = 0$, on Γ_s
 - Floating ice BC: $2\mu \dot{\boldsymbol{\epsilon}}_i \cdot \boldsymbol{n} = \begin{cases} \rho g z \boldsymbol{n}, \text{ if } z > 0 \\ 0, \quad \text{ if } z < 0 \end{cases}, \text{ on } \Gamma_l$
 - Basal sliding BC:

$$2\mu \dot{\boldsymbol{\epsilon}}_i \cdot \boldsymbol{n} + \beta(x, y)u_i = 0$$
, on Γ_{β}



 $\beta(x, y) =$ basal sliding coefficient

*Assumption: aspect ratio δ is small and normals to upper/lower surfaces are almost vertical.

Appendix: *Kokkos*-ification of Finite Element Assembly (Example)





Appendix: Ice Sheet Dynamic Equations

Model for *evolution of the boundaries* (thickness evolution equation):

$$\frac{\partial H}{\partial t} = -\nabla \cdot (\overline{\boldsymbol{u}}H) + \dot{\boldsymbol{b}}$$

where \overline{u} = vertically averaged velocity, \dot{b} = surface mass balance (conservation of mass).

• Temperature equation (advection-diffusion):

$$\rho c \frac{\partial T}{\partial t} = \nabla \cdot (k \nabla T) - \rho c \boldsymbol{u} \cdot \nabla T + 2 \dot{\boldsymbol{\epsilon}} \boldsymbol{\sigma}$$

(energy balance).

- **Flow factor** A in Glen's law depends on temperature T: A = A(T).
- Ice sheet *grows/retreats* depending on thickness *H*.







Appendix: MPI+Device Scalability Study



Sandia National

lahoratories

Device Comparison, P100 vs. KNL (GIS 4km-20km mesh)

- 1. Blue: mostly Residual/Jacobian computation, Yellow: mostly communication.
- 2. KNL performs better because of heavy use of MPI
- 3. P100 performance is hindered by communication cost (worse when scaling because of lack of CUDA aware MPI)

Appendix: MPI+Device Scalability Study





- Scalability of **GlobalAssembly + evaluateFields** is studied.
- **Strong Scalability** is for GIS 4km-20km tet mesh (1.51M elements).
- Weak Scalability is for GIS 4km-20km and 1km-7km (1.51M and 14.4M elements) tet meshes.
- KNL strong scaling **under investigation** requires profiling.
- P100 results hindered by communication cost (worse when scaling b/c no CUDA-aware MPI)
 - Scalability can likely be improved by removing CUDA UVM.
- Weak scaling is comparable for P100 and KNL.
 - KNL performs **better** because of heavy use of MPI.

Appendix: Greenland Weak Scalability on *Titan*

Sandia National Laboratories

Weak scalability on Titan (16km, 8km, 4km, 2km, 1km Greenland)



Appendix: Scalability with Increasing Order Sandia Elements



- *Left:* speedup plot shows benefit of using higher orders to obtain better strong scalability for MPI+OpenMP for Aeras atmosphere dycore shallow water test case.
- *Right:* weak scalability for MPI + GPU on the Ride for Aeras atmosphere dycore shallow water test case. Efficiency drops significantly for lower order elements, but GPU is better able to maintain weak scaling for higher order p6 spectral element.

Appendix: Kokkos Range vs. MDRange Policy



Sandia

National

• Range vs. MDRange policy for shallow water test case in Aeras atmosphere dycore with p6 spectral element for MPI + GPU.

Appendix: Improved Linear Solver Performance through Hinge Removal

Islands and certain hinged peninsulas lead to **solver failures**

- We have developed an algorithm to detect/remove problematic hinged peninsulas & islands based on coloring and repeated use of connected component algorithms (Tuminaro et al., 2016).
- Solves are ~2x faster with hinges removed.
- Current implementation is MATLAB, but working on C++ implementation for integration into dycores.



| | Resolu- | ILU – | ILU – no | ML – | ML – no |
|---|------------------|----------------|------------------------------|---------------------------|---------------------------|
| | tion | hinges | hinges | hinges | hinges |
| | 8km/5 | 878 sec, | 693 sec, | 254 sec, | 220 sec, |
| | layers | 84 iter/solve | 71 iter/solve | 11 iter/solve | 9 iter/solve |
| Ī | 4km/10 | 1953 sec, | 1969 sec, | 285 sec, | 245 sec, |
| | layers | 160 iter/solve | 160 iter/solve | 13 iter/solve | 12 iter/solve |
| | 2km/20 | 10942 sec, | 5576 sec, | 482 sec, | 294 sec, |
| | layers | 710 iter/solve | 426 iter/solve | 24 iter/solve | 15 iter/solve |
| | 1km/40 layers | | 15716 sec, 881 iter/solve | 668 sec, 34 iter/solve | 378 sec, 20 iter/solve |







