FASTMath: Frameworks, Algorithms and Scalable Technologies for Mathematics

Lori Diachin, LLNL

Institute Director

The FASTMath project brings leading edge computational mathematics technologies to the SciDAC Program

Develop advanced numerical techniques for DOE applications

- Eight focused topical areas based on application needs
- High level synergistic techniques

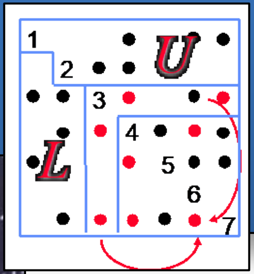

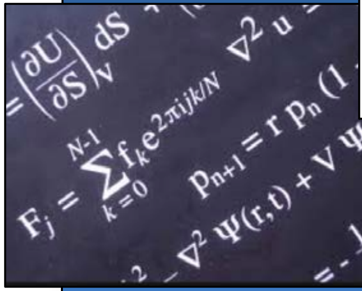

Deploy high-performance software on DOE supercomputers

- Algorithmic and implementation scalability
- Performance portability
- Interoperability of libraries

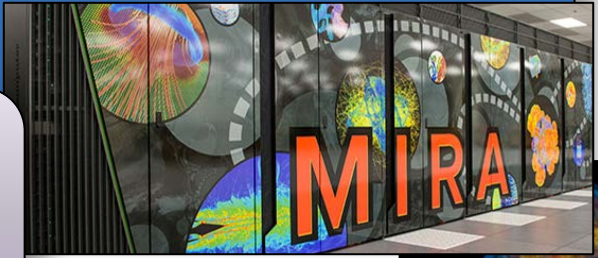

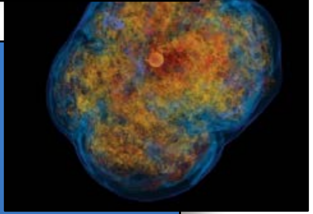

FASTMath Objective: Reduce the barriers facing application scientists

Demonstrate basic research technologies from applied mathematics

- Build from existing connections with basic research
- Focus on research results that are most likely to meet application needs

Engage and support of the computational science community

- Publications and presentations in highly visible venues
- Team tutorials
- Workforce pipeline and training
- Web presence

The SciDAC-4 FASTMath Institute leverages and builds on the successes of SciDAC-3 to meet application needs

FASTMATH 3 + QUEST

+ Numerical Optimization

+ Data Analytics

=

FASTMATH 4

FASTMath brings together an exceptional team of researchers and software library capabilities

Our team comprises over 50 researchers from 5 national laboratories and 5 universities

BERKELEY LAB

Argonne NATIONAL LABORATORY

OAK RIDGE National Laboratory

Sandia National Laboratories

Boulder

USC

MIT

Rensselaer

SMU

Our software has 100s of person years of development behind it

SuperLU

sundials

PETSc

mfem

ZOLTAN

PUMi Parallel Unstructured Mesh Infrastruc

hypre high performance preconditioners

PEXSI

AMReX

UQTk

Trilinos

Albany

DAKOTA

U.S. DEPARTMENT OF ENERGY | Office of Science

For more information contact: Lori Diachin, LLNL
diachin2@llnl.gov

FASTMATH

The FASTMath team has a proven record of advancing application simulation codes

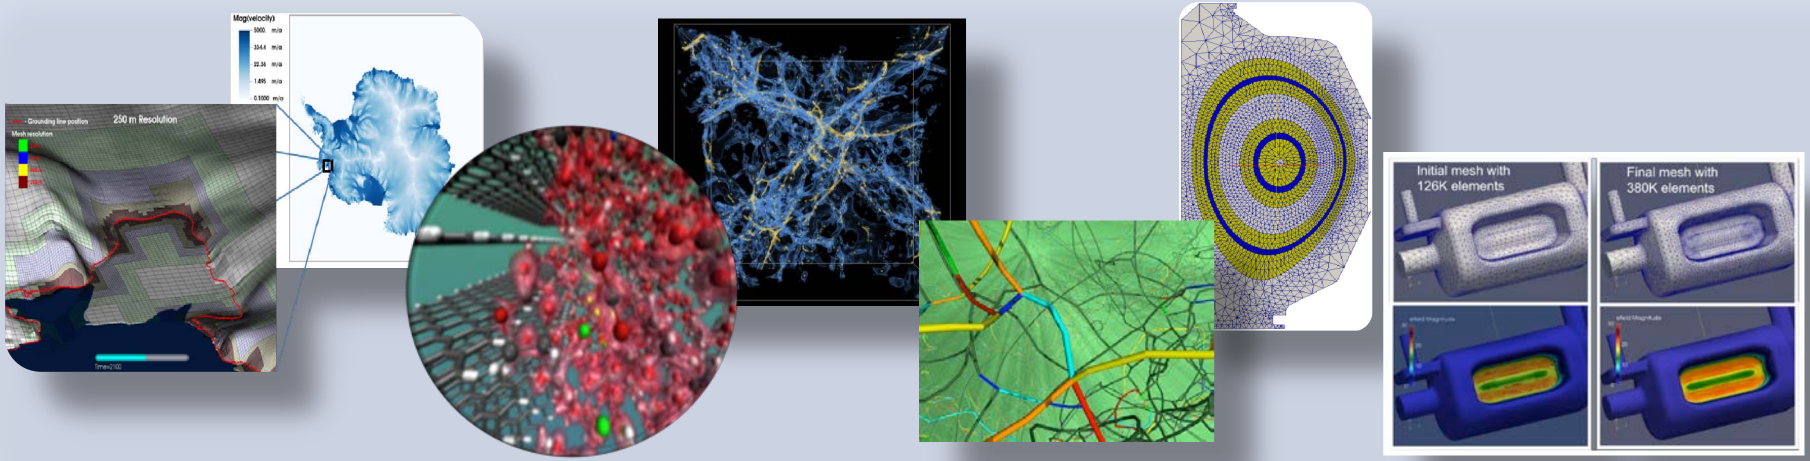

Next Generation Application Codes

- Created unique DOE capabilities in ice sheet modeling; first simulations in Antarctic that capture dynamics of warm water forcing
- First ever, self consistent solution of continuum gyrokinetic system in edge plasmas
- Unprecedented resolution for Nyx cosmology code

Faster Time to Solution

- New eigensolvers 2X faster for quantum chemistry software
- Parallel PEXSI software enabled electronic structure calculations with 10,000 atoms (compared to 1000's)
- Accelerated nonlinear solver enabled largest dislocation dynamics simulation with ParaDiS

More Robust Simulations

- Dramatically decreased time to generate meshes for fusion tokamak codes
- Adaptive mesh refinement and discretizations to resolve ELM disruptions in tokamaks
- Order of magnitude improvement in accuracy of integral calculations in material chemistry

U.S. DEPARTMENT OF ENERGY | Office of Science

For more information contact: Lori Diachin, LLNL
diachin2@llnl.gov

FASTMATH

FASTMath is focused on eight core technology areas

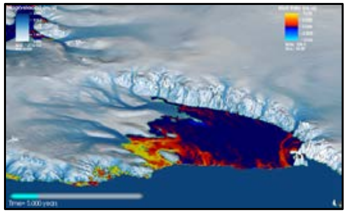

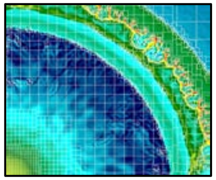

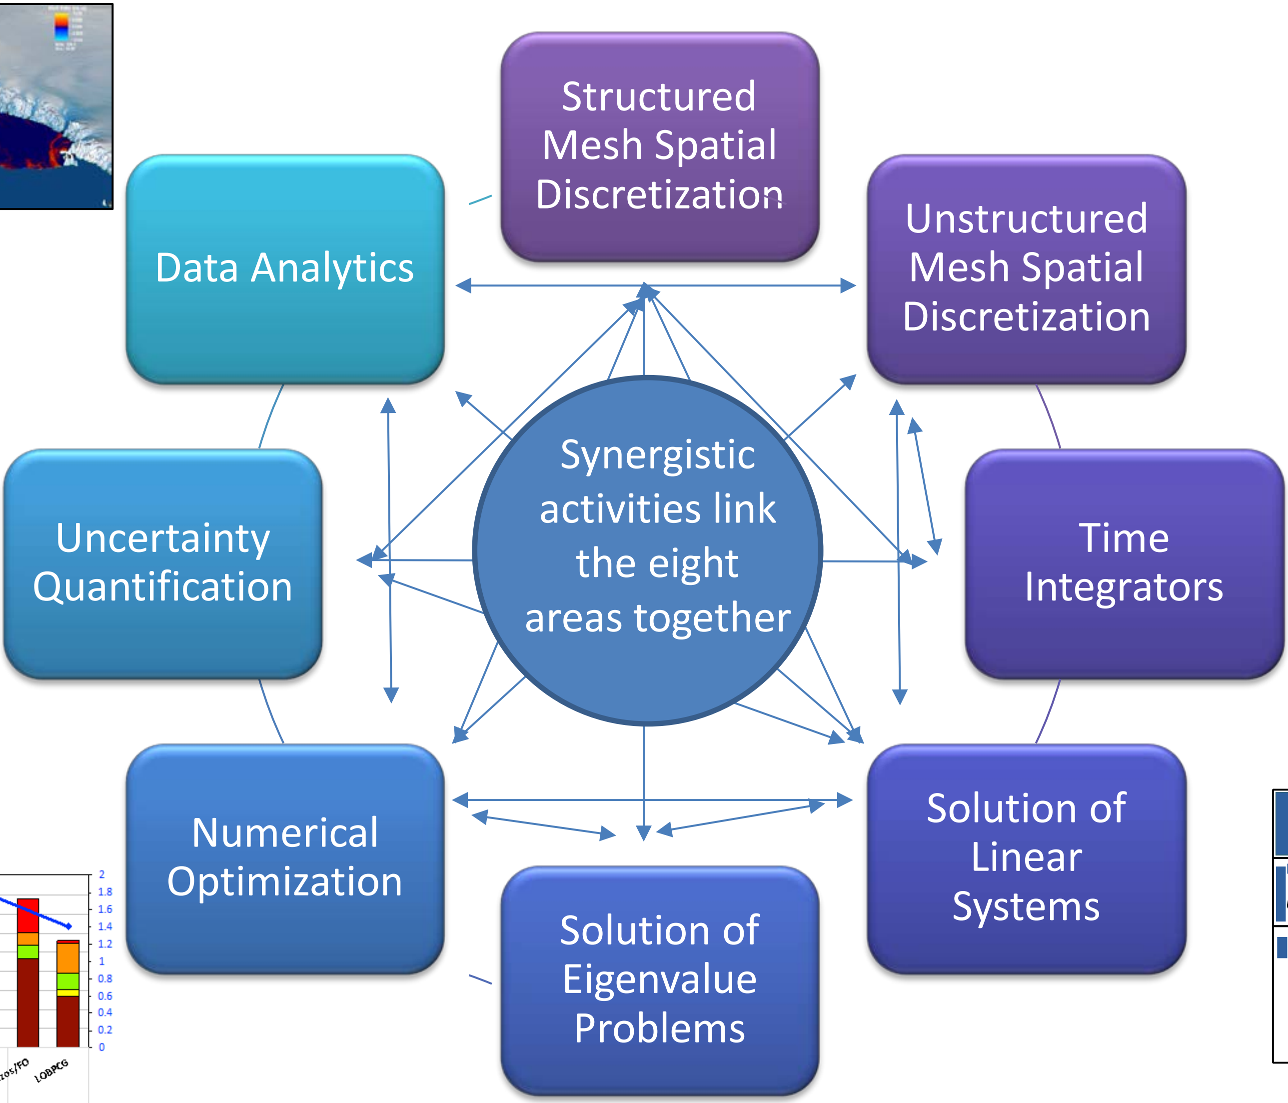

Structured Mesh Spatial Discretization

Unstructured Mesh Spatial Discretization

Time Integrators

Solution of Linear Systems

Solution of Eigenvalue Problems

Numerical Optimization

Uncertainty Quantification

Data Analytics

Synergistic activities link the eight areas together

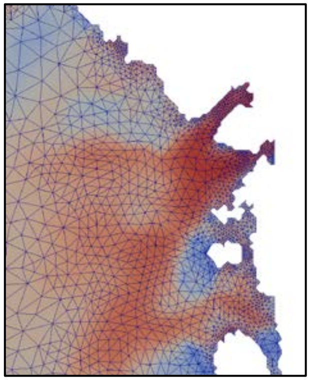

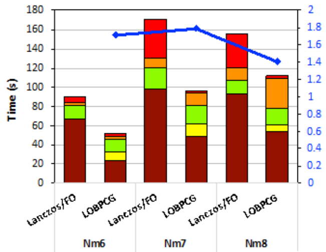

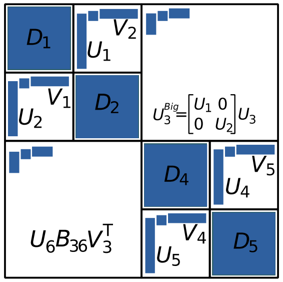

U.S. DEPARTMENT OF ENERGY | Office of Science

FASTMATH

Synergystic activities will result in new capabilities or higher efficiencies

New Capabilities

- Optimization under uncertainty
- Discrete and multi-objective optimization for data analytics
- In situ simulation on unstructured meshes

Higher Efficiency

- Leverage multiple right hand sides from optimization and UQ ensembles in linear and nonlinear solvers
- Adaptivity in the spatial and stochastic space in UQ on unstructured grids;
- Dynamic UQ load balancing
- In situ simulation on unstructured meshes

FASTMath is actively engaged with 18 SciDAC-4 application partnerships

BER (5)

- Structured grid AMR
- Unstructured grid AMR
- Time integration
- Linear/Nonlinear solvers, Preconditioners
- Optimization
- Verification / UQ

FES (5)

- Unstructured meshes
- Discretization technologies
- Iterative Linear solvers
- UQ

HEP (3)

- Direct solvers
- Structured Grid AMR
- Optimization
- Sensitivity Analysis
- Inference and machine learning

NP (2)

- Structured grid AMR
- Eigenvalue problems
- Inference and Machine Learning

BES (2)

- Nonlinear and tensor eigenvalue problems
- Linear solvers and Preconditioners

NE (1)

- UQ

Eight core technology areas: Structured Mesh

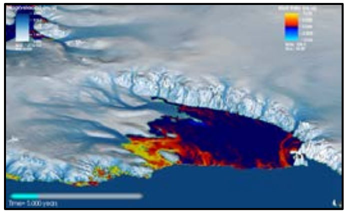

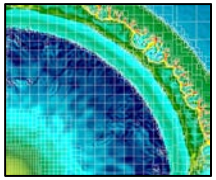

Structured Mesh Spatial Discretization

Data Analytics

Unstructured Mesh Spatial Discretization

Uncertainty Quantification

Synergistic activities link the eight areas together

Time Integrators

Numerical Optimization

Solution of Linear Systems

Solution of Eigenvalue Problems

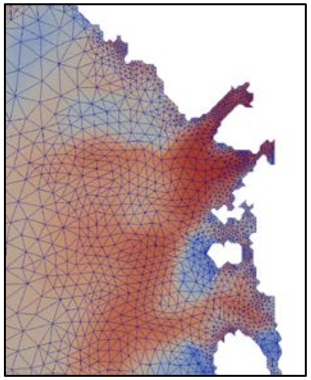

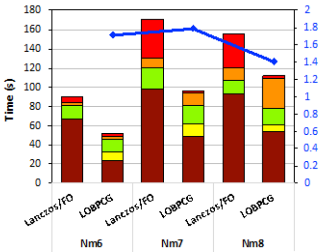

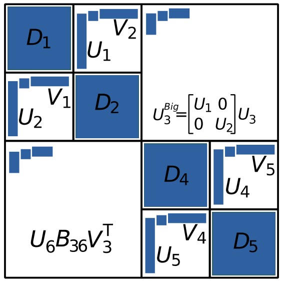

For more information contact: Ann Almgren (LBNL), asalmgren@lbl.gov

Structured Mesh Discretization Tools

- **Goal:** Provide software support for efficient parallel solution of a large variety of problems in science and engineering using block-structured adaptive mesh approaches
- **Software frameworks:** AMReX, Chombo
- **FASTMath Tasks**

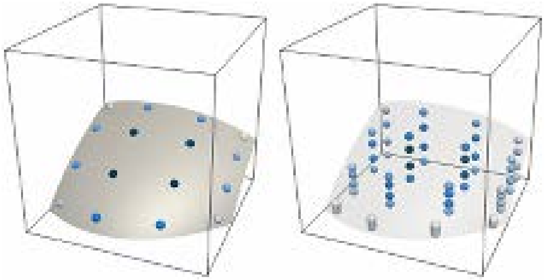

- Adaptive Mesh Refinement (AMR)
- Higher-order interior spatial discretizations
- Higher-order time-stepping
- Higher-order tools for interface dynamics
- Particle dynamics and particle-mesh operations
- Mapped multi-block domains
- Dynamic load balancing
- Interoperability with other solvers, e.g. those in SUNDIALS, PETSc, hypre

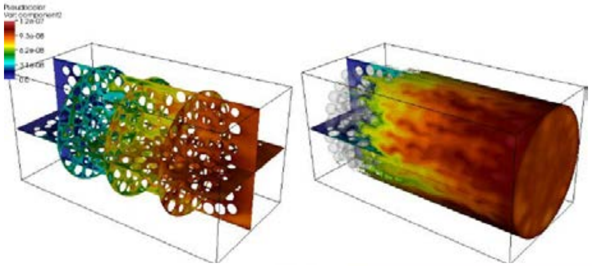

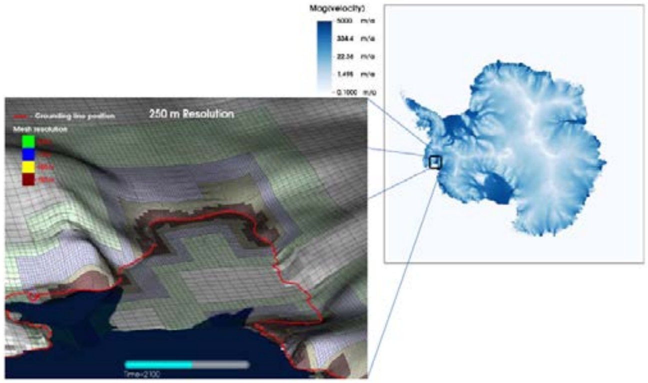

- **Applications:** Accelerator modeling, fusion, astrophysics, ice sheet, manufacturing, combustion, cosmology
- **For more information:** Ann Almgren(ASAlmgren@lbl.gov)

Robust large-scale inversion capability for ProSPect ice sheet modeling

Partnership Application:

Probabilistic Sea-Level Projection from Ice Sheet and Earth System Models (ProSPect)

***ProSPect* will allow for probabilistic projections of sea-level change from ice sheet, ocean, and Earth system models by addressing existing model deficiencies**

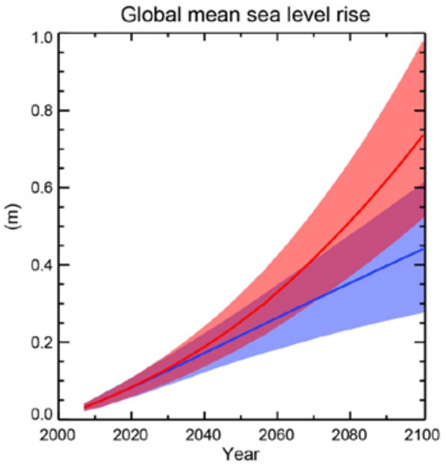

Focus areas:

- Ice sheet & ocean model physics
- Coupling between ice sheet, ocean, & Earth system models
- **AMR and Nonlinear solvers for fast, efficient computation**
- **Ice sheet model optimization; coupled ice sheet & Earth system model initialization**
- **Quantification of uncertainty in sea-level projections**
- **Ice sheet model performance on next-generation architectures**

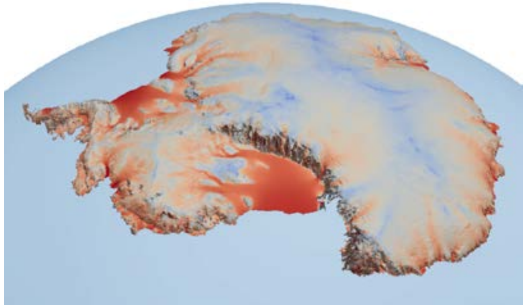

Our structured grid capabilities are being used to help model ice sheets and sea level rise

- The Antarctic Ice Sheet (particularly in West Antarctica) is believed to be vulnerable to collapse driven by warm-water incursion under ice shelves, which causes a loss of buttressing, subsequent grounding-line retreat, and large (up to 4m) contributions to sea level rise.
- Very fine (finer than 1km) spatial resolution needed to resolve ice dynamics around grounding lines (the point at which grounded ice begins to float)
- FASTMath-supported Chombo-based BISICLES ice sheet model (part of the BER ProSPect SciDAC partnership) uses adaptive mesh refinement (AMR) to enable sufficiently-resolved modeling of full-continent ice sheet response to climate forcing.

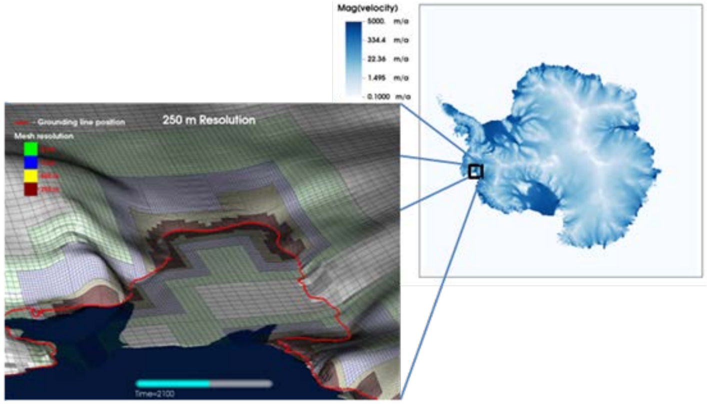

(right) Computed ice velocity for Antarctica; (left) meshing and grounding line location for the Pine Island Glacier. The **grounding line** is the location at which the ice flowing into the ocean separates from the bedrock and begins to float.

FASTMath support

- FASTMath-supported Chombo framework provides mathematical and software infrastructure needed for a scalable and efficient AMR model.
- BISICLES solves a nonlinear coupled viscous tensor solve – FASTMath solver technologies and support resulted in greatly improved solver efficiency and robustness.
- In SciDAC-4, FASTMath will continue to provide solver and framework improvements, while also contributing toward better discretization of grounding lines (locations where grounded ice begins to float) using multifluid discretizations.
  - Correct and accurate treatment of grounding lines is key to correctly modeling marine ice sheets like West Antarctica.

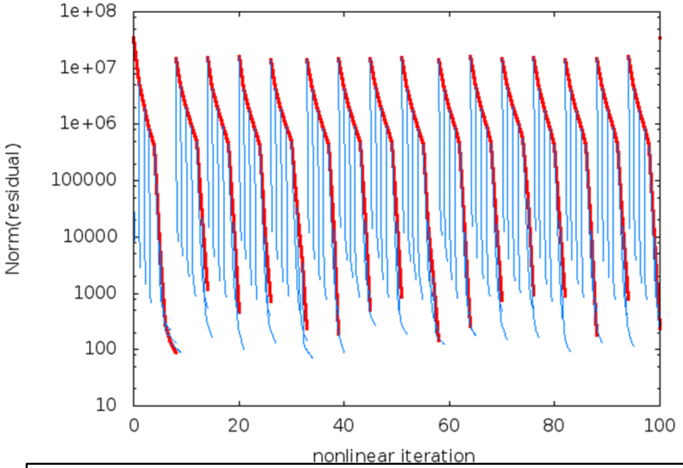

Solver convergence for full-continent Antarctica ice velocity solves. Red is outer nonlinear residual, blue is inner linear-solver residuals. SciDAC3 FASTMath led to better solver performance and robustness.

Results to Date

- ## Thwaites Glacier Instability
  - Waibel, Hulbe, Jackson, and Martin (2018). *Geophysical Research Letters*, **45**, 809-816.
  - BISICLES used to examine the processes regulating basin-wide ice mass loss for Thwaites Glacier, part of the West Antarctic Ice Sheet (WAIS)
  - Finds that rates of mass loss are especially sensitive near the point the system transitions into a regime of self-sustained retreat where the effects of anomalous forcing and larger driving stresses are sustained by the dynamics of a retreating ice sheet.

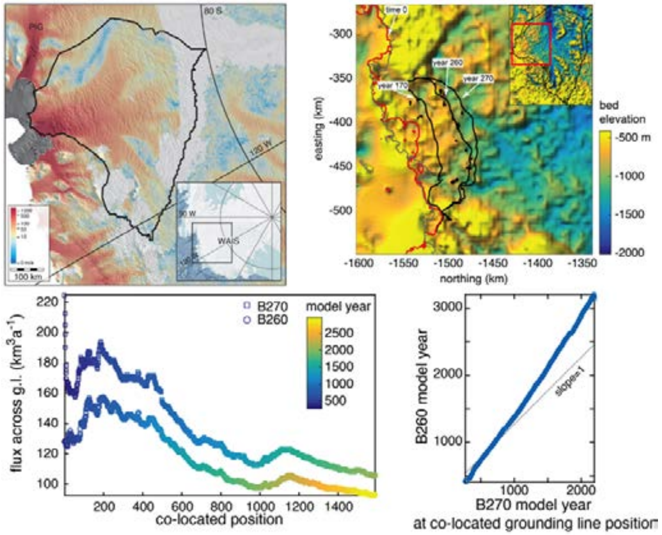

Top left: *Location and surface velocities of Thwaites basin, West Antarctica.* ***Top right:*** *Position of grounding line where melt rate reaches 13 m/yr (model year 260) and 13.5 m/yr (model year 270).* ***Bottom left & right:*** *By model year 270 mass loss accelerates and remains high for duration of experiment (from Waibel et al., 2018).*

- ## Antarctic Vulnerability
  - Martin, et al, submitted to *Nature Climate Change*.
  - Used BISICLES to study Antarctic vulnerability to regional ice shelf collapse.
  - Finds high vulnerability of Western Antarctic Ice Sheet to loss of any of its ice shelves
  - Limited vulnerability elsewhere.

Left: *Ice sheet grounding-line evolution for different ice shelf collapse chronologies.*
Right: *Antarctic vulnerability (in mm Sea Level Equivalent (SLE) to regional ice shelf collapse for each of 14 sectors.*

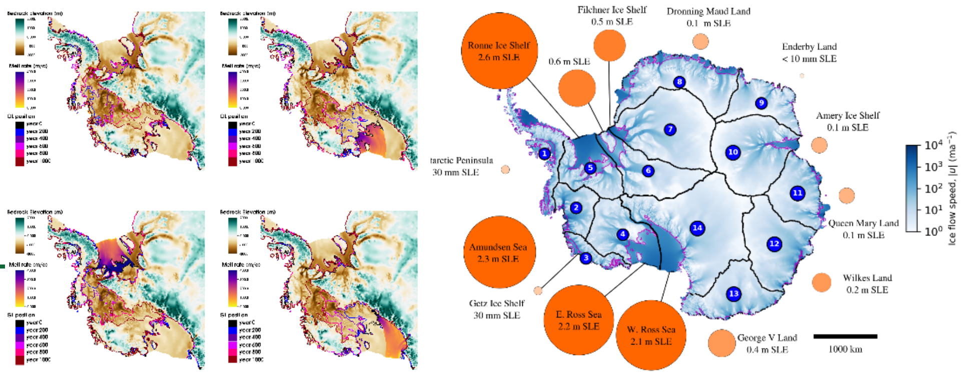

Eight core technology areas: Unstructured Meshes

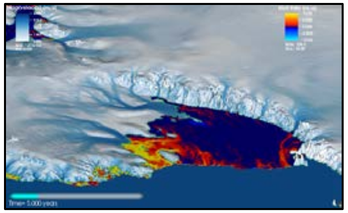

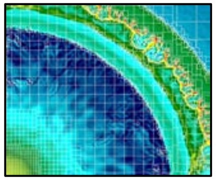

Structured Mesh Spatial Discretization

Data Analytics

Unstructured Mesh Spatial Discretization

Uncertainty Quantification

Synergistic activities link the eight areas together

Time Integrators

Numerical Optimization

Solution of Linear Systems

Solution of Eigenvalue Problems

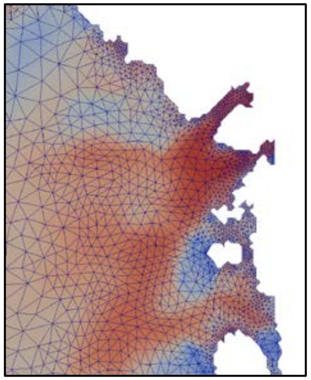

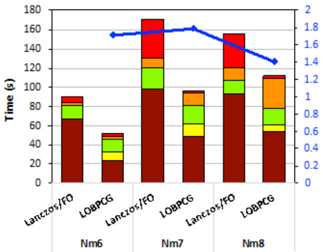

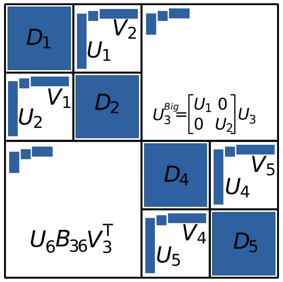

For more information contact: Mark Shephard (RPI), shephard@rpi.edu

Unstructured Mesh Discretization tools

- **Goal:** Scalable and performant unstructured mesh tools to support complex geometry applications
- **Software tools:** Albany, EnGPar, MeshAdapt, MFEM, PHASTA, PUMI, Zoltan/Zoltan2
- **FASTMath Tasks:**

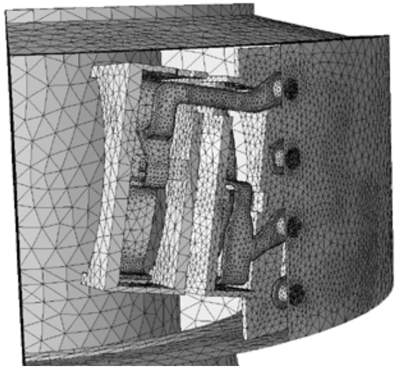

Mesh Adaptation

- Conforming and non-conforming adaptation
- High-order curved mesh adaptation
- Anisotropic adaptivity
- Adaptation in UQ space
- Parallel meshes with particles methods
- Combined mesh motion and adaptation
- Effective in-memory integration methods

Scalable & Performant

- Performant on latest many core and hybrid systems nodes
- Robust and fast dynamic multi-criteria load balancing
- Architecture-aware load balancing
- Performance portability using latest tools
- In situ visualization and problem modification

Discretization Technologies

- High-order FEM
- Stabilized FEM for complex flow problems
- Non-linear coupled mechanics
- Error estimation procedures
- Conservative field transfer
- Scalable and performant on the latest systems

- **Applications:** Fusion, atmospheric modeling, ice sheet melting
- **For more information:** Mark Shephard (shephard@scorec.rpi.edu)

FASTMath unstructured mesh tools are providing the foundation for four fusion simulation partnerships

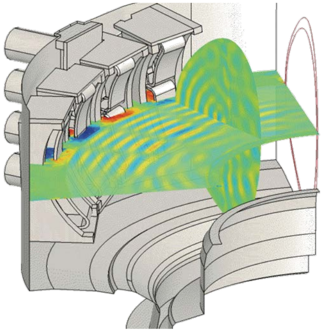

High-Order Finite Elements for RF Simulations (Bonoli)

- Solve on graded unstructured meshes for geometrically complex domains
- Scale to >100M unknowns; use high order for accuracy
- MFEM finite element library ideally suited to address these needs
- Frequency domain EM solver developed in MFEM
- EM solver integrated into the RF SciDAC simulation workflow

Geometry and Meshing Developments (Chang, Jardin, Bonoli)

- 3D geometry and meshing including CAD antenna modules
- Specific feature field following mesh generation for XGC1
- 2D to 3D and mesh adaptivity for M3D-C1
- Geometry/meshing for RF simulations
- Curved mesh adaptation

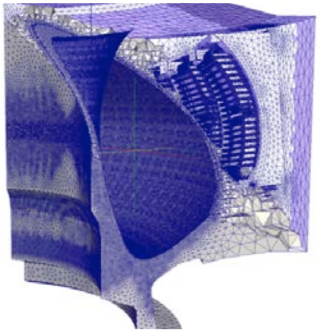

Particle support on unstructured meshes (Chang, Wirth)

- Mesh to particle and particle to mesh adaptation
- Particle migration and loading
- Parallel solve using parallel PETSc solves on each plane
- Dynamic load balancing approach defined
- Needed for XGC Edge plasma and GITR Impurity Transport code

MPI Task placement to reduce communication in E3SM HOMME atmospheric modeling code

- **E3SM: Energy Exascale Earth System Model**
- **HOMME: High-Order Method Modeling Environment for atmospheric modeling (PI: Mark Taylor, SNL)**
  - Uses unstructured quadrilateral meshes on the sphere (e.g., cubed-sphere mesh)
  - Tasks are vertical columns of elements in the atmosphere
- **Aim: scale with very small number (O(1)) of columns per core**
  - Task placement in network is important to keep communication cost low
  - HOMME's default is space-filling curve (SFC) for partitioning and linear task placement

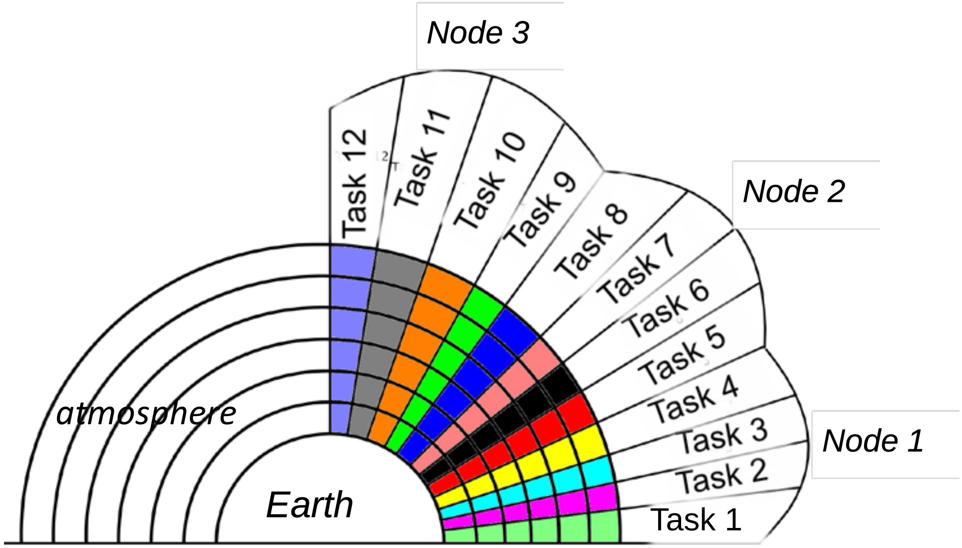

Zoltan2's geometric task placement reduces HOMME's communication time at large scale

- **Geometric task placement assigns communicating tasks to "nearby" cores in the node allocation**

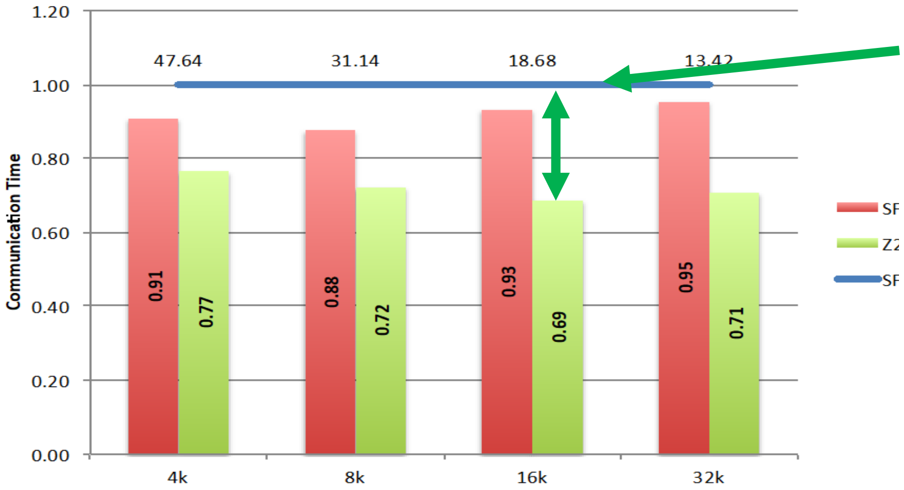

31% reduction relative to HOMME's default task placement

SFC partition + Zoltan2 task placement

Zoltan2 partition & task placement

SFC partition & task placement (HOMME Default)

- **E3SM coupler and other components may benefit from task mapping**
  - Reduce communication between/within components
  - Included in E3SM Phase 2 proposal
- **Geometric mapping on Aries networks with adaptive routing require new coordinate transformations to represent "proximity"**

Option to use Zoltan2 as TPL has been accepted in HOMME code base

Eight core technology areas: Time Integrators

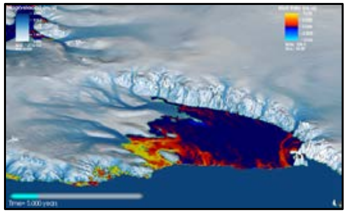

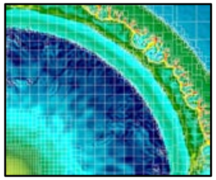

Structured Mesh Spatial Discretization

Data Analytics

Unstructured Mesh Spatial Discretization

Uncertainty Quantification

Synergistic activities link the eight areas together

Time Integrators

Numerical Optimization

Solution of Linear Systems

Solution of Eigenvalue Problems

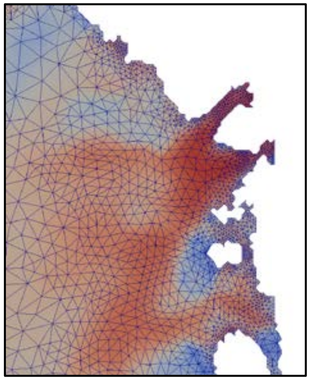

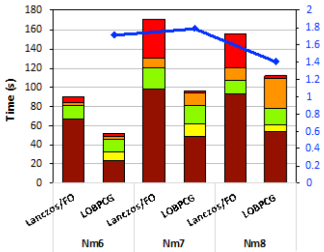

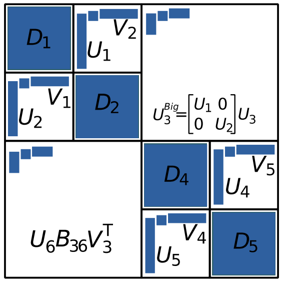

For more information contact: Carol Woodward (LLNL), cswoodward@llnl.gov

Time Integrators

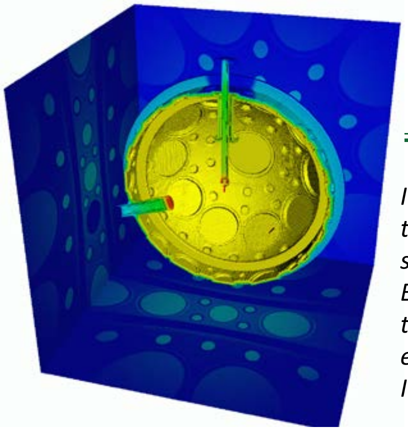

Intensity in NOVA test chamber; solution of Boltzmann transport equation using IDA

- **Goal:** Efficient and robust time integration methods and software for stiff implicit, explicit, and multirate systems in DOE applications
- **Software tools:** SUNDIALS, PETSc, SDC in AMReX
- **FASTMath Tasks**

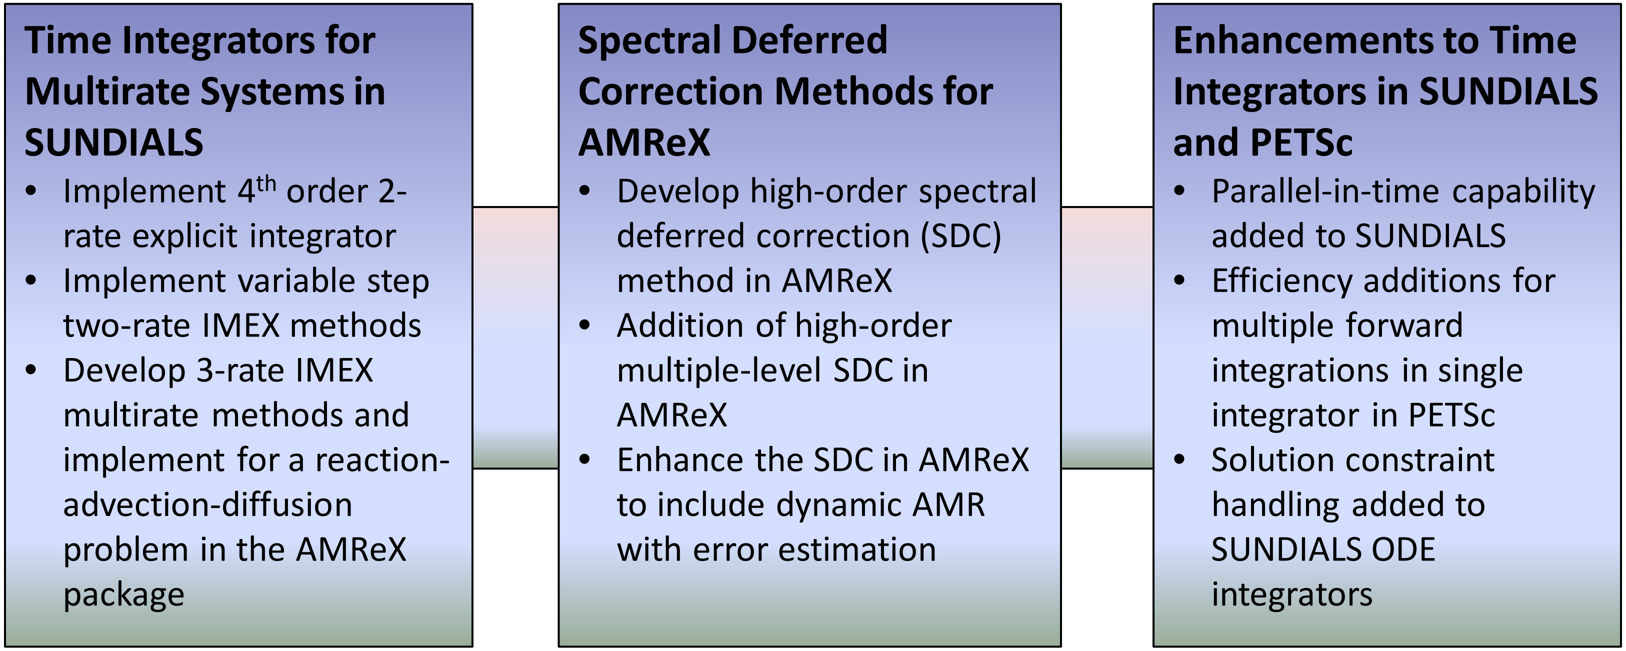

Time Integrators for Multirate Systems in SUNDIALS

- Implement 4th order 2-rate explicit integrator
- Implement variable step two-rate IMEX methods
- Develop 3-rate IMEX multirate methods and implement for a reaction-advection-diffusion problem in the AMReX package

Spectral Deferred Correction Methods for AMReX

- Develop high-order spectral deferred correction (SDC) method in AMReX
- Addition of high-order multiple-level SDC in AMReX
- Enhance the SDC in AMReX to include dynamic AMR with error estimation

Enhancements to Time Integrators in SUNDIALS and PETSc

- Parallel-in-time capability added to SUNDIALS
- Efficiency additions for multiple forward integrations in single integrator in PETSc
- Solution constraint handling added to SUNDIALS ODE integrators

- **Applications:** Atmospheric dynamics, ice sheets, combustion, power grid, cosmology, subsurface flow
- **For more information:** Carol Woodward (woodward6@llnl.gov)

The FASTMath time integrator team is working to improve integration in the DOE E3SM climate code

We are working with two application partnerships to address time integration needs in the Energy Exascale Earth System Model (E3SM)

Partnership applications:

"Assessing and Improving the Numerical Solution of Atmospheric Physics in ACME," PI: Hui Wan (PNNL)

Goal is to identify and address temporal convergence bottlenecks in the climate physics part of E3SM

Fig: Convergence of atmospheric physics is less than expected first order

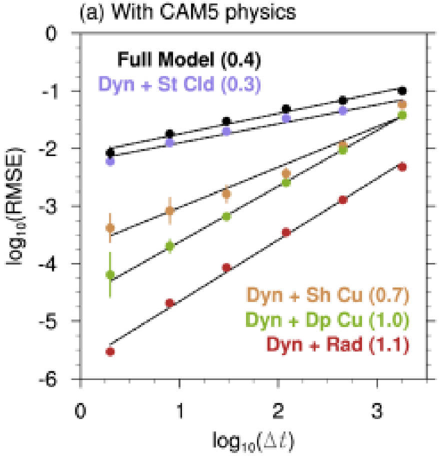

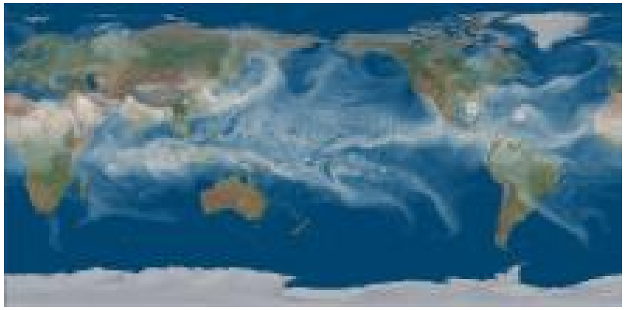

https://climatemodeling.science.energy.gov/projects/energy-earth-systemmodel

"A Non-hydrostatic Variable Resolution Atmospheric Model in ACME," PI: Mark Taylor (SNL)

Goal is to develop effective numerical methods and implementation for nonhydrostatic atmospheric dynamics

We are applying rigorous analysis and modern integration methods to atmospheric simulations

- **For climate physics:**
  - Developing error analysis for current schemes and analyzing when assumptions break in model system
  - Analyzing effects of limiters on solution accuracy
  - Developing alternative model formulations
  - Developing new schemes that are more robust

- **For atmospheric dynamics we are developing implicit/explicit (IMEX) integration methods**

$$\frac{d\vec{y}}{dt} = \vec{f}^E(t, \vec{y}) + \vec{f}^I(t, \vec{y})$$

  - Implicit in the vertical direction; explicit in the horizontal directions
  - Interfaced SUNDIALS ARKode package to Tempest (UC Davis) nonhydrostatic code for testing various splittings and methods,
  - Developed solver approach for implicit portion
  - Interfaced ARKode to E3SM; now testing most promising methods on new nonhydrostatic code

Progress to date in both areas is promising

For physics:

- Identified singularity in one model and verified resulting breakdown of convergence
- Rederived parameterization
- Implemented new model in E3SM framework
- No adjustments to data to enforce consistency required
- First order convergence is recovered

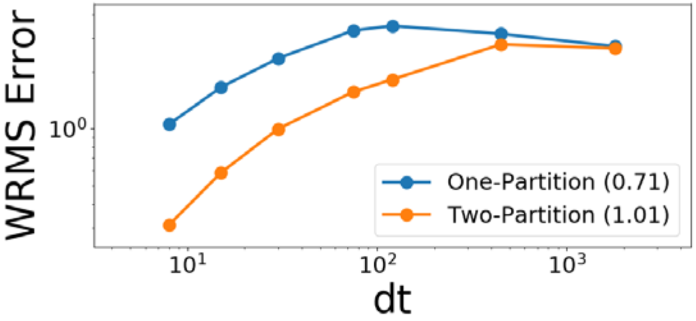

For dynamics

- We tested 252 combinations of methods, splittings, and solvers
- Identified highly efficient methods
- Acceptable solutions with steps as high as 300s; iterating nonlinear solve could raise this to 400s
- now applying ARKode to test most effective approaches in E3SM

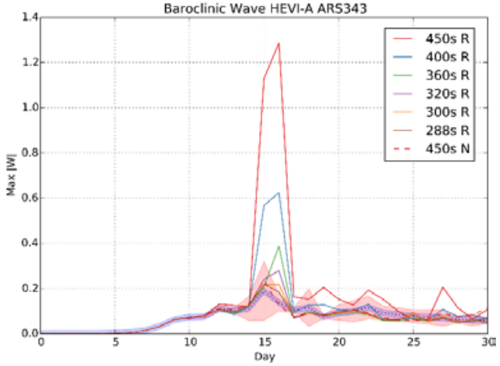

Maximum vertical velocity over 30 days with fastest IMEX method using various step sizes and solvers showing region of acceptable solution.

Eight core technology areas: Linear Systems

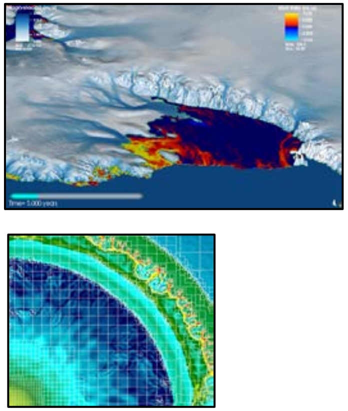

Structured Mesh Spatial Discretization

Data Analytics

Unstructured Mesh Spatial Discretization

Synergistic activities link the eight areas together

Uncertainty Quantification

Time Integrators

Numerical Optimization

Solution of Linear Systems

Solution of Eigenvalue Problems

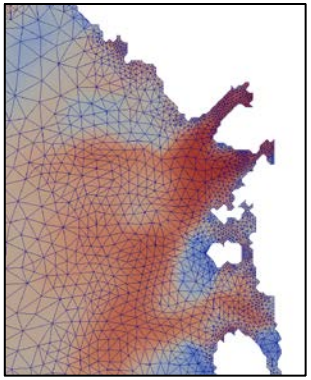

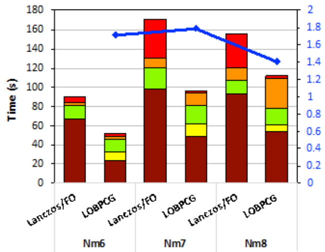

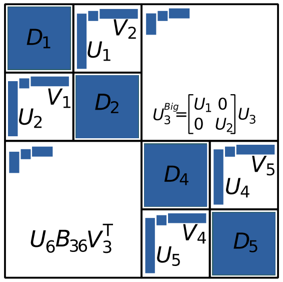

For more information contact: Ulrike Yang (LLNL), yang11@llnl.gov

FASTMath Linear Solvers

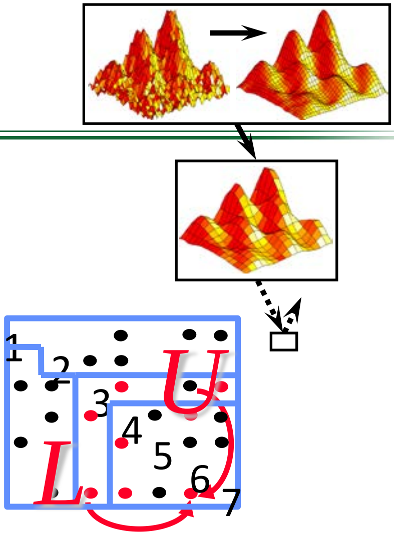

- **Goal:** Providing efficient direct and iterative linear solvers and preconditioners for large scale parallel computers.
- **Associated software packages:** hypre; PETSc; MueLu, ShyLU (Trilinos); KokkosKernels; SuperLU; STRUMPACK; symPACK
- **FASTMath Tasks:**

Efficient kernels and solvers

- On-node and GPU kernels for iterative solvers
- Develop efficient linear solvers for systems with multiple right hand sides
- Improve scalability for hierarchical low-rank preconditioners
- Design and deploy local discrete convolution methods (LDCMs)

Develop multigrid methods

- Design and implement semi-structured multigrid methods that effectively use grid structures
- Implement efficient Schur complement-based multilevel algorithms for strongly coupled multi-physics systems
- Develop fast hybrid geometric/algebraic multigrid solvers for elliptic and hyperbolic systems

Deploy highly efficient direct linear solvers

- Investigate and implement butterfly-based data-sparse methods
- Decrease communication in symmetric indefinite solvers
- Improve structural and data sparsity through scalable ordering and symbolic factorization algorithms

- **Applications**: climate, astrophysics, fusion, ice sheet, subsurface flow, power grid
- **For more information**: Ulrike Yang (yang11@llnl.gov)

New algebraic multigrid solvers developed for nonsymmetric problems

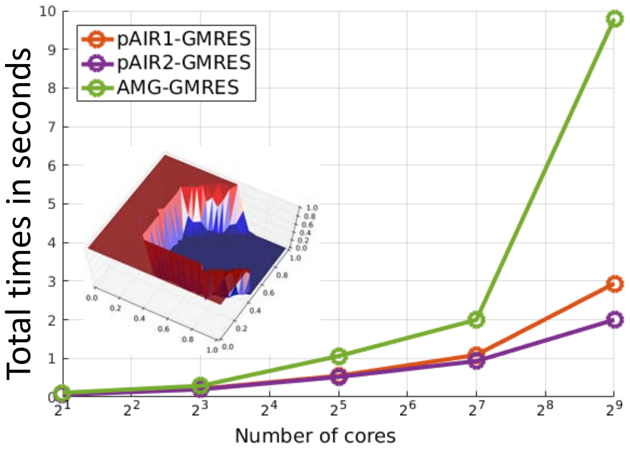

Weak scalability for a nonsymmetric problem (solution pictured) comparing pAIR and AMG

Scientific Achievement

Developed new algebraic multigrid method, pAIR (parallel Approximate Ideal Restrictions), capable of solving highly non-symmetric problems that conventional AMG methods are unable to solve or only solve poorly; integrated the new methods in hypre.

Significance and Impact

Nonsymmetric advection dominated problems, occurring e.g. in transport, often present a challenge for algebraic multigrid methods. The new AMG method enables significantly faster solution.

Research Details

pAIR is a localized multigrid reduction method that approximates the ideal restriction operator from a local neighborhood and uses a low-complexity interpolation operator. It was implemented in hypre and tested on various nonsymmetric problems, leading to speedups of up to 5 and solving some problems that could not be solved before.

Work was performed at LLNL
For more information contact Ulrike Yang: yang11@llnl.gov

FASTMath and Fusion teams participated in Dungeon Session at Intel in Hillsboro, OR, Sept. 19-21, 2017

- **Fusion team: Center for Tokamak Transient Simulations**
  - Steve Jardin (PI), Jin Chen, PPPL
  - Two-fluid simulation codes to solve MHD equations: M3D-C1, NIMROD
  - SuperLU_DIST and STRUMPACK used as Block-Jacobi preconditioner
  - Learn and understand Vtune functionalities
- **Improvements**
  - KNL-optimized SuperLU_DIST improves significantly
    - Factorization of each block up to 2.3x faster
    - PETSc preconditioned GMRES 12% faster
  - Build SuperLU_DIST with craype-mic-knl leads to 10% improvement in PETSc
  - Integrate STRUMPACK into PETSc, used in M3D-C1 as a preconditioner
    - Resolved an interface issue with arguments passing between STRUMPACK and PETSc
  - Tuning of parameters for PETSc GMRES led to 1.1x to 5.2x speedup in the solve time
    - Optimize restart number
    - Relax GMRES to reuse 2 previous directions
  - Vtune points to inefficiency in “extend-add” routine in STRUMPACK. Loops in this routine are improved 10%.

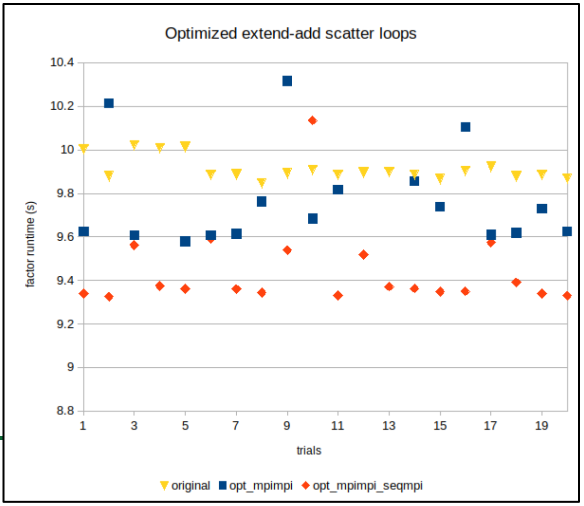

Eight core technology areas: Eigenvalue Problems

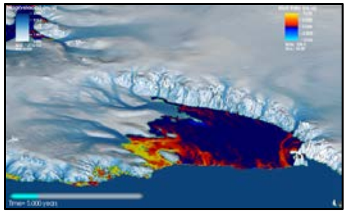

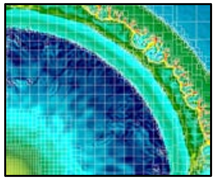

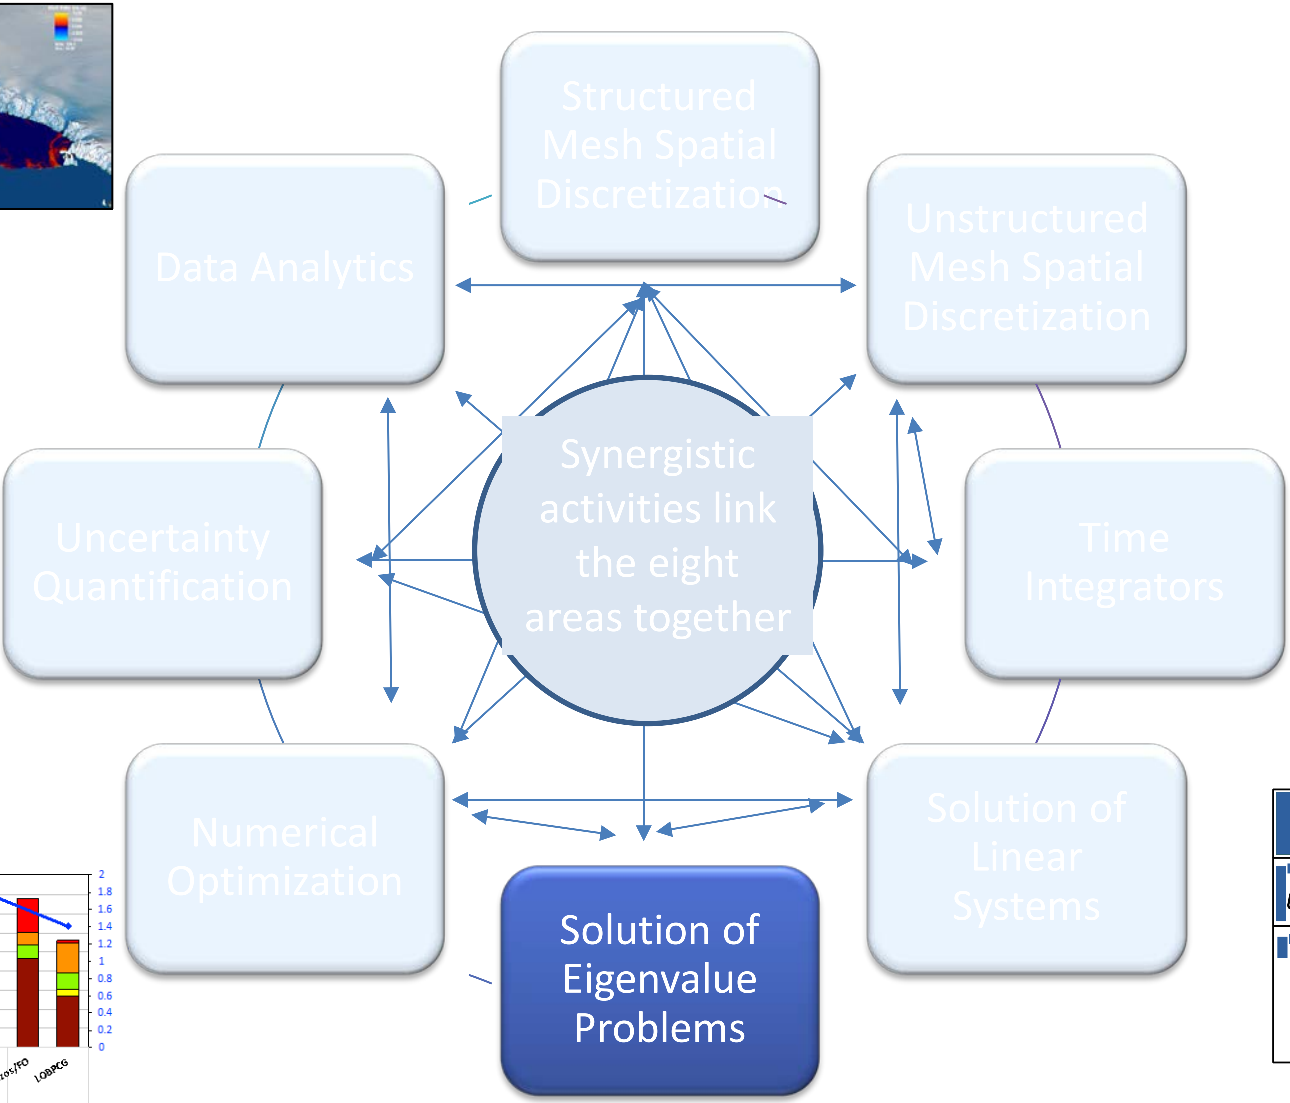

Structured Mesh Spatial Discretization

Unstructured Mesh Spatial Discretization

Data Analytics

Synergistic activities link the eight areas together

Uncertainty Quantification

Time Integrators

Numerical Optimization

Solution of Linear Systems

Solution of Eigenvalue Problems

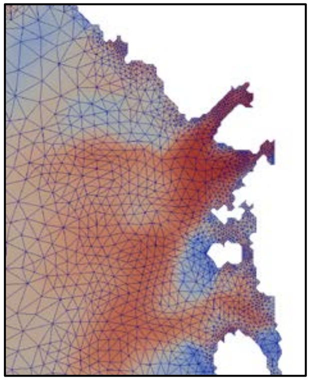

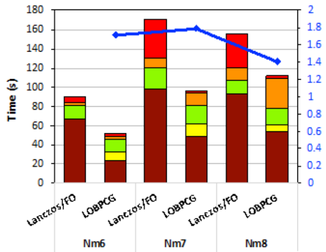

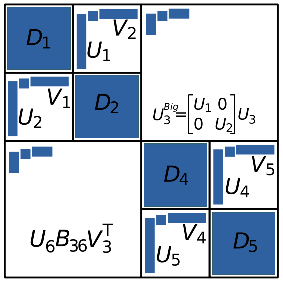

U.S. DEPARTMENT OF
ENERGY

Office o
Science

For more information contact: Chao Yang (LBNL), cyang@lbl.gov

FASTMath Eigenvalue Toolset

- **Goal:** Scalable tools for linear and nonlinear eigenvalue problems in DOE applications
- **Software tools:** FASTEig (a collection of eigensolvers for a variety of eigenvalue problems)
- **FASTMath Tasks:**

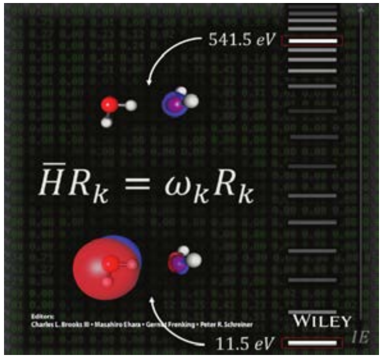

Efficient algorithms for eigenvalue computation

- Spectrum slicing for many eigenpairs
- Pade approximation and linearization, compact rational Krylov for nonlinear problems
- Algorithms for tensor eigenvalue problems
- Structure preserving algorithms for linear response and spectroscopy

Scalable implementation for distributed many-cores and GPUs

- Increase concurrency though multilevel algorithms and randomization
- Load balancing
- Overlap computation with communication

Deploy solvers in software packages used in DOE SciDAC and other applications

- Linear eigensolver for Quantum Espresso, PARSEC, NWChemEX
- Structure eigensolver for BerkeleyGW
- Linear response eigensolver for Chronus-Q
- Solvers for iTensor

- **Applications:** Catalysis, topological materials, nuclear structures, chemistry
- **For more information:** Chao Yang (CYang@lbl.gov)

Eigensolver for Electronic Structure of Catalytic Materials

- **BES Partnership (PI: Martin Head-Gordon)**

  Advancing Catalysis Modeling: From Atomistic Chemistry to Whole System Simulation

- **Challenge:**
  - Large-scale electronic structure calculation via hybrid functional density functional theory (DFT)
  - Ab initio molecular dynamics
- **Expected impact:**
  - Accurate binding energy, Gibbs free energy, activation barrier calculations
  - Elucidate reaction pathway of CO2 reduction on catalytic metal surfaces

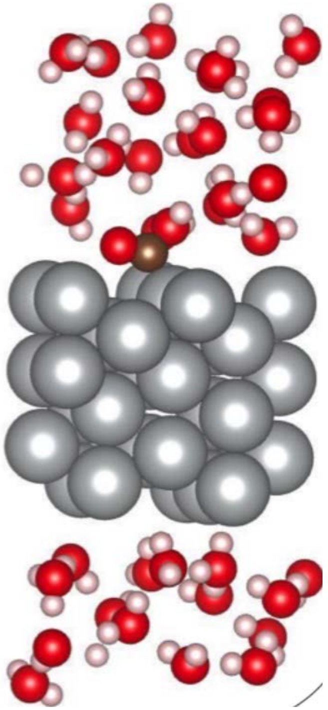

FASTMath Team: L. Lin, X. Li, E. Ng and C. Yang

FASTMath Approaches

Nonlinear Eigenvalue Problem: $[\boldsymbol{H_o}+\boldsymbol{V_X}(\boldsymbol{Z})]\boldsymbol{Z} = \boldsymbol{Z}\boldsymbol{\Lambda}$

- Instead of an SVD, construct low-rank approximation to the (non-local) exchange potential V_X through interpolative separable density fitting

$$\varphi_i(\mathbf{r})\psi_j(\mathbf{r}) \approx \sum_{\mu=1}^{N_\mu} \zeta_\mu(\mathbf{r}) \boxed{\varphi_i(\hat{\mathbf{r}}_\mu)\psi_j(\hat{\mathbf{r}}_\mu)} \longleftarrow \text{free!}$$

- Use adaptive compressed exchange (ACE) to further reduce the cost of applying V_X to the occupied orbitals in an iterative diagonalization procedure

- Robust and efficient convergence acceleration without constructing the full Hamiltonian or density matrix

The use of acceleration techniques significantly reduces computational cost

- Two orders of magnitude improvement in computational cost of a 1000 atom silicon cluster benchmark

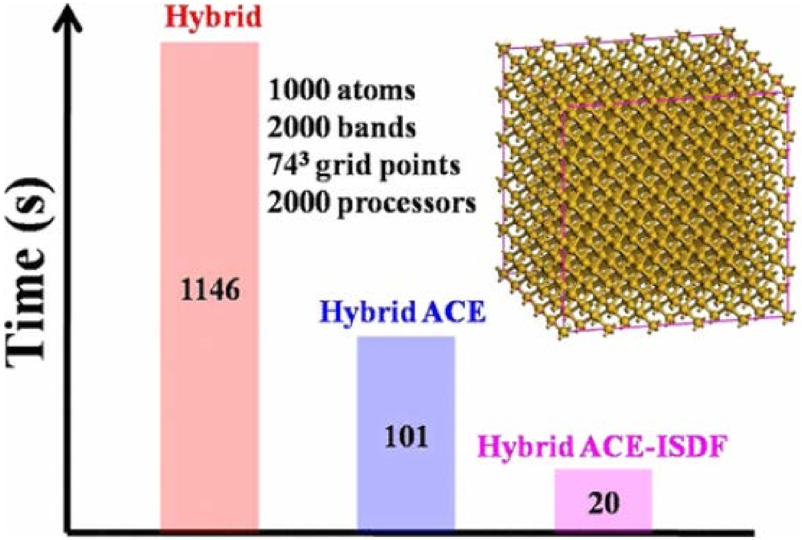

W. Hu, L Lin, C. Yang, J. Chem. Theory Comput. 13, 5420, 2017

- New acceleration technique eliminates need to perform inner/outer self-consistent field iteration; improving efficiency of hybrid DFT calculation

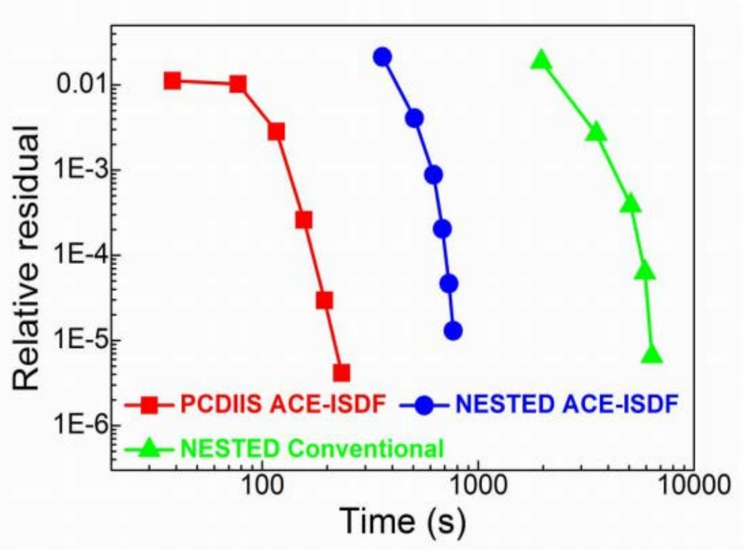

W. Hu, L. Lin, C. Yang, J. Chem. Theory Comput. 13, 5458, 2017

Eight core technology areas: Numerical Optimization

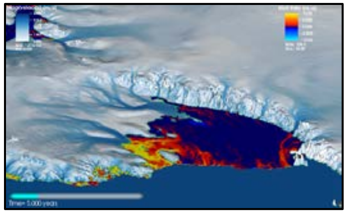

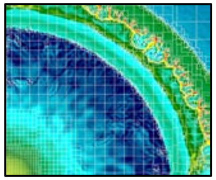

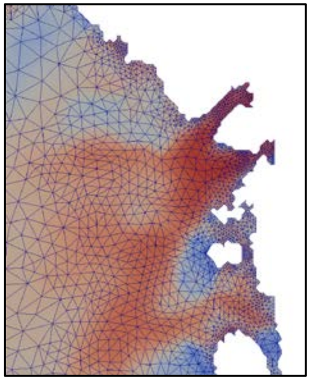

Structured Mesh Spatial Discretization

Data Analytics

Unstructured Mesh Spatial Discretization

Uncertainty Quantification

Synergistic activities link the eight areas together

Time Integrators

Numerical Optimization

Solution of Linear Systems

Solution of Eigenvalue Problems

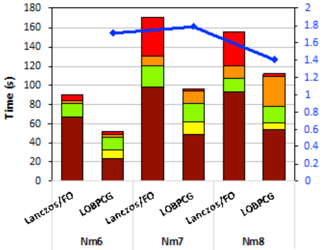

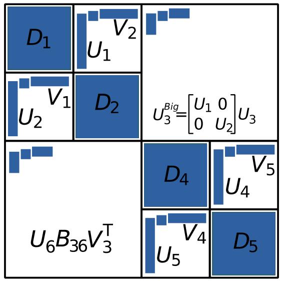

For more information contact: Todd Munson (ANL), tmunson@mcs.anl.gov

Numerical Optimization Tools

- **Goal:** Develop methods for numerical optimization problems with constraints and for sensitivity analysis using adjoint capabilities.
- **Software tools:** MATSuMoTo, MINOTAUR, ORBIT, ROL, TAO
- **FASTMath Tasks:**

Dynamic Optimization

- Deliver new capabilities for problems with PDE constraints that include:
  - Dynamics and controls
  - State and design constraints
  - Discrete variables
  - Multiple objectives
- Support a range of derivative requirements

Sensitivity Analysis

- Develop iterative sampling methods that employ sensitivity analysis and surrogate models to determine the most important parameters
- Explore multilevel approach that uses a low-fidelity model to predict parameter sensitivities

Adjoints

- Develop advanced adjoint and forward sensitivity capabilities to provide essential derivative information
- Provide methods for computing the action of second-order adjoints
- Support calculations involving many quantities of interest arising from constraints and multiple objectives

- **Applications:** nuclear physics, accelerator modeling, ice sheet modeling
- **For more information:** Todd Munson (tmunson@mcs.anl.gov)

FASTMath approach to deliver robust inversion capability for ProSPect

To find an ice sheet initial state that is consistent with present-day observations and climate forcing requires

- a robust large-scale inversion capability
- constrained by coupled physical models
- simultaneously inverting for several fields (e.g. basal friction, geothermal heat flux, bed topography)
- characterized by tens of millions of parameters

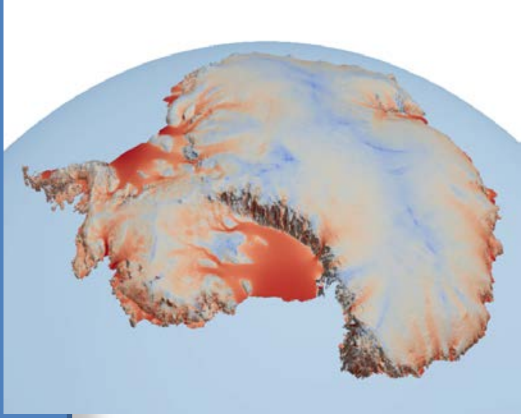

FASTMath is providing infrastructure to support large-scale numerical optimization and UQ in Albany/Trilinos

- Enable advanced optimization capabilities in Albanyinterfacing with the ROL optimization library
- Compute Hessians to enable Newton optimization methods
- Enable transient optimization
- Use new Trilinos Tpetra stack and Kokkos for performance portability

The partnership team has made significant progress to date

- Ported the computation of sensitivities using automatic differentiation to the new Trilinos Tpetra stack. This will enable performance portability.

- Initialized Antarctica ice sheet inverting simultaneously for basal friction and stiffening coefficient, see Figure.

Inverted for 2.6M parameters

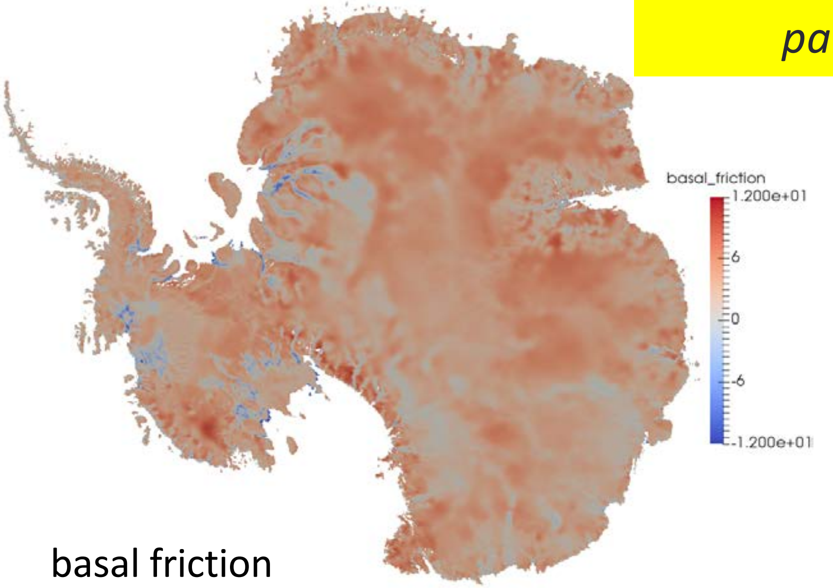

basal friction

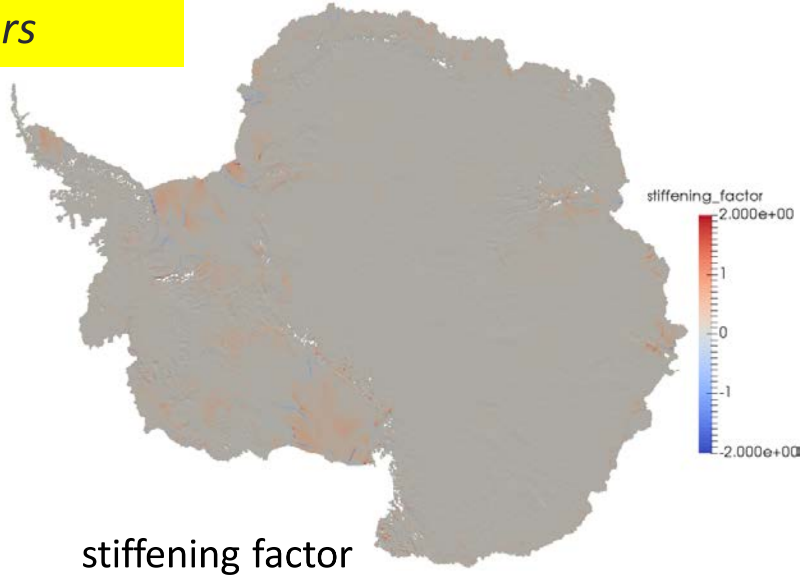

stiffening factor

Eight core technology areas: Uncertainty Quantification

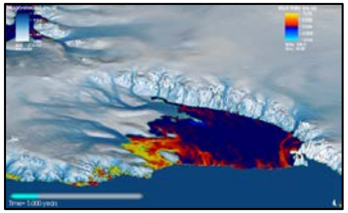

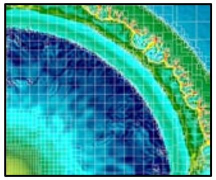

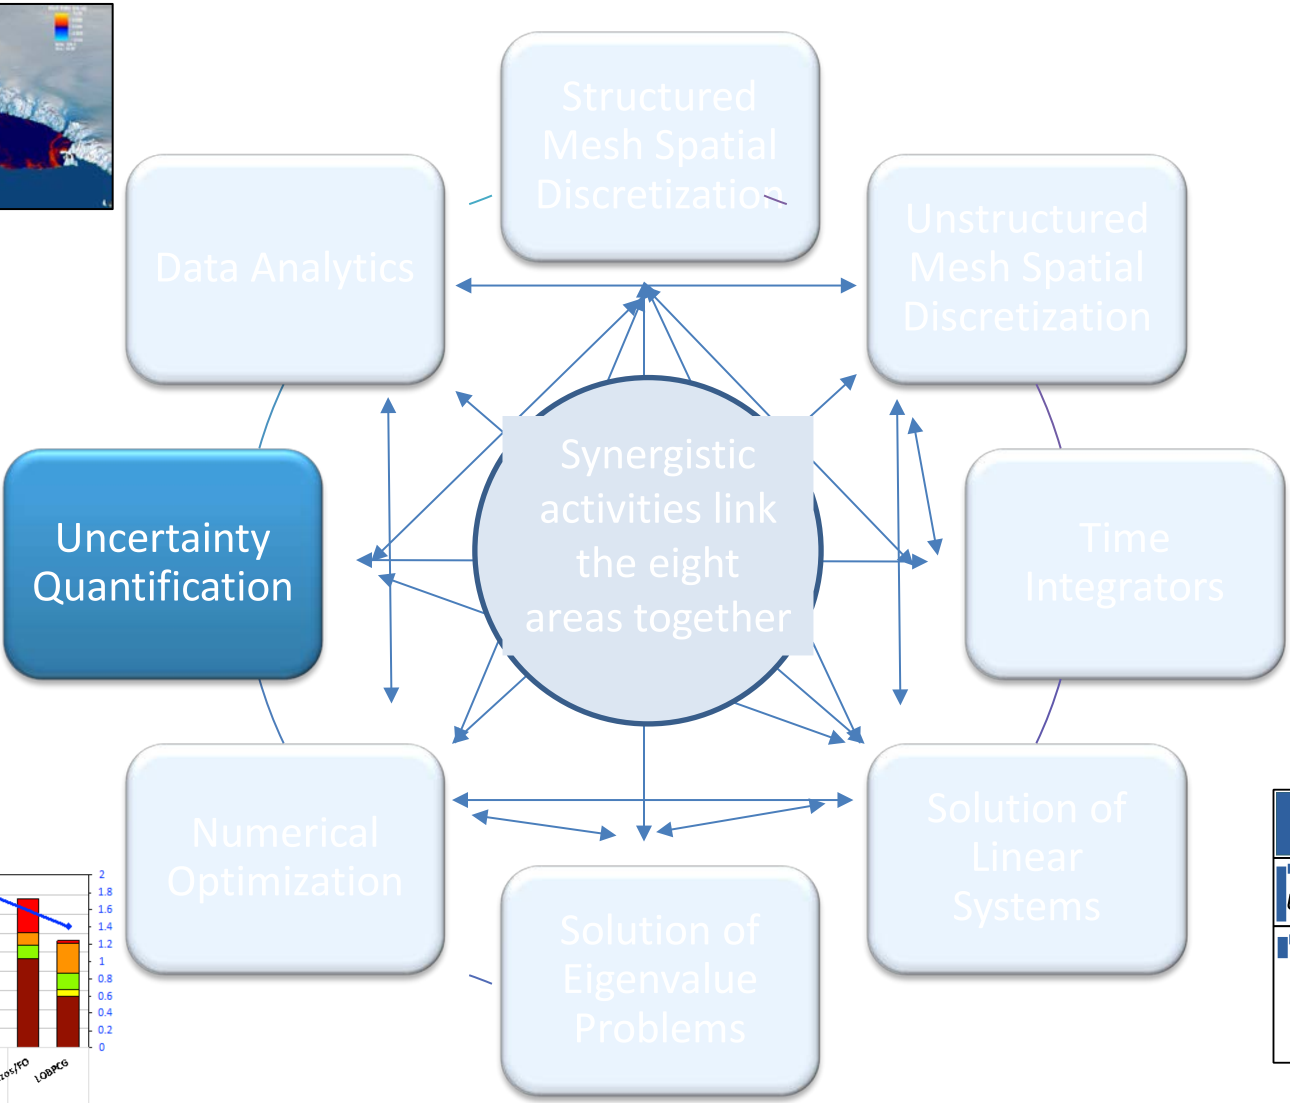

Structured Mesh Spatial Discretization

Data Analytics

Unstructured Mesh Spatial Discretization

Uncertainty Quantification

Synergistic activities link the eight areas together

Time Integrators

Numerical Optimization

Solution of Linear Systems

Solution of Eigenvalue Problems

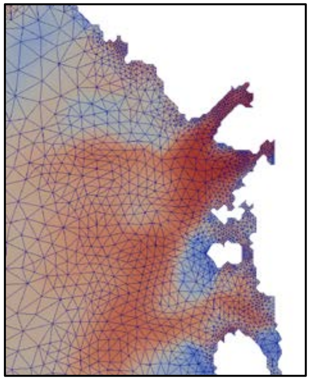

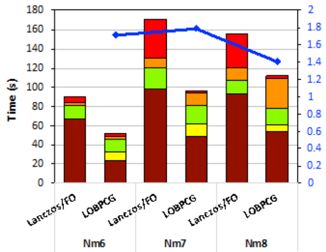

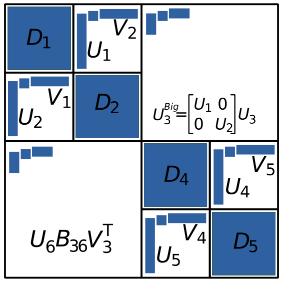

For more information contact: Habib Najm (SNL), hnnajm@sandia.gov

Uncertainty Quantification

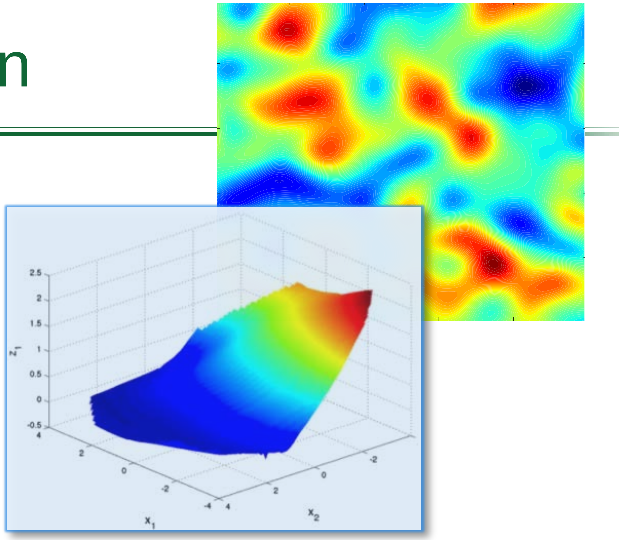

- **Goal:** Provide robust and efficient capabilities for uncertainty quantification in large-scale computations of physical systems.
- **Software tools:** DAKOTA, UQTk
- **FASTMath Tasks:**

Deploy advanced algorithms in UQ tools

- Functional low rank decompositions on mixed spaces
- MCMC methods focused on likelihood informed subspaces
- Adaptive basis
- Multilevel multifidelity
- Regression and classification
- Model error

Advance UQ methods

- Adaptive low rank sparse constructions
- Parallel SDE sampling on manifolds
- Multilevel UQ based on sparse low rank emulators
- Online local surrogate refinement coupled with data driven dimension reduction

Tools for optimization under uncertainty

- Recursive trust-region model management for optimization across deep model hierarchies
- Advanced methods for probability of failure estimation
- Reliability based OUU for design in the presence of rare events

- **Applications:** ice sheet modeling, fusion, wind energy, climate, plasma, nuclear fuels
- **For more information:** Habib N. Najm (hnnajm@sandia.gov)

ProSPect ice sheet modeling is also leveraging FASTMath UQ tools

UQ Overall approach

- Identify and characterize major sources of model uncertainty and assign distributions based upon prior belief.
- Update the conservative prior descriptions of uncertainty using observational data
- Quantify uncertainty in sea level projections by propagating posterior distributions through land-ice model

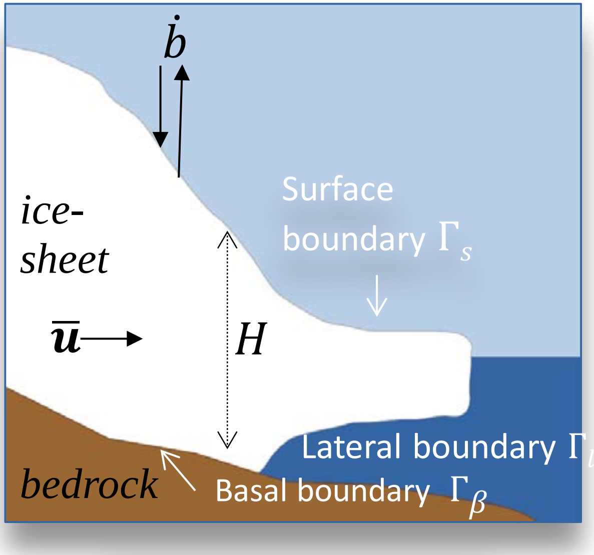

Challenge

- Model has large numbers of uncertainties $O(10^6)$
- Model is computationally intensive requiring approximately $O(10^3)$ core hours for moderately resolved meshes
- Number of simulations required for uncertainty quantification grows strongly with dimension

Proposed solution

- Use adjoint-based gradients to facilitate dimension reduction.
- Find parameter directions most informed by data that cause largest change from prior to posterior distribution

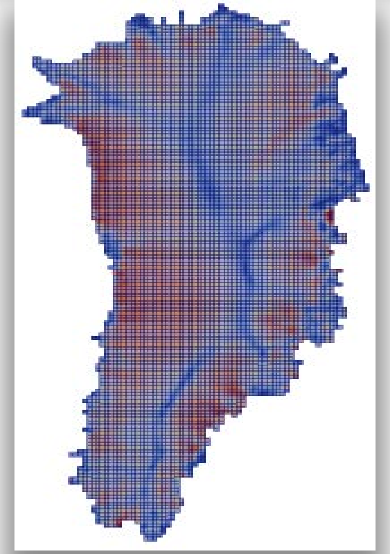

Friction is represented by a finite representation of a log-normal random field at the mesh points on land-ice boundary

Uncertainty quantification for ProSPect: Results to date

Focus to date

- Our current focus is on uncertain initial condition which are a major source of diverging projections of ice sheet mass loss & sea-level rise
- If initial conditions are inconsistent with observations and/or climate forcing this can lead to non-physical transients.

Approach to date

- Use gradient based optimization to find basal friction field consistent with observational data (constraints).
- Use Bayesian inference to obtain a statistical description of the parameters that are consistent with the data
- Use Monte Carlo sampling from posterior to estimate sea-level change

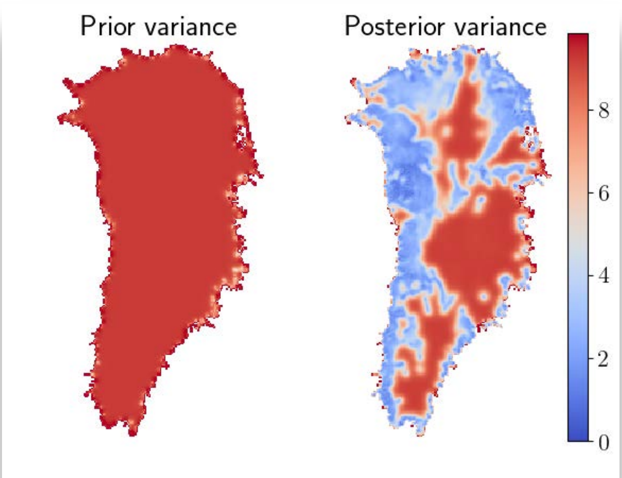

Prior and posterior distributions of friction over Greenland

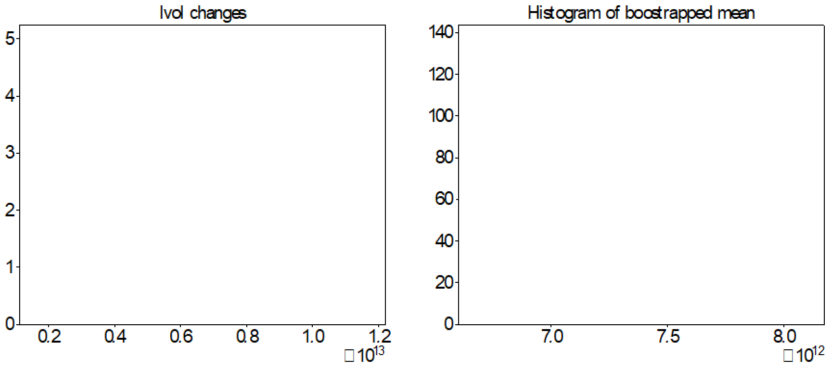

Data informed estimate of mean change in ice-sheet volume. Assumption of lognormal random field leads to bias towards ice-sheet mass gain. Assumptions must be revisited.

FASTMath UQ is also being used in a partnership with the Nuclear Energy project

- **Project Title**: Simulation of Fission Gas in Uranium Oxide Nuclear Fuel
- **Project objective: advance mechanistic understanding of fission gas behavior and release in UO2 nuclear fuel**
  - Develop and use a mesoscale simulator that takes advantage of leadership class computers.
- **UQ challenges for FASTMath**
  - Understand impact of uncertainty in model parameters on predictions
  - Address the coupling of uncertainty in the multiscale model hierarchy
  - Identify parameters with dominant impact on predictive uncertainty
  - Provide estimates of uncertainty in predictive quantities of interest
  - Enable model selection and model validation

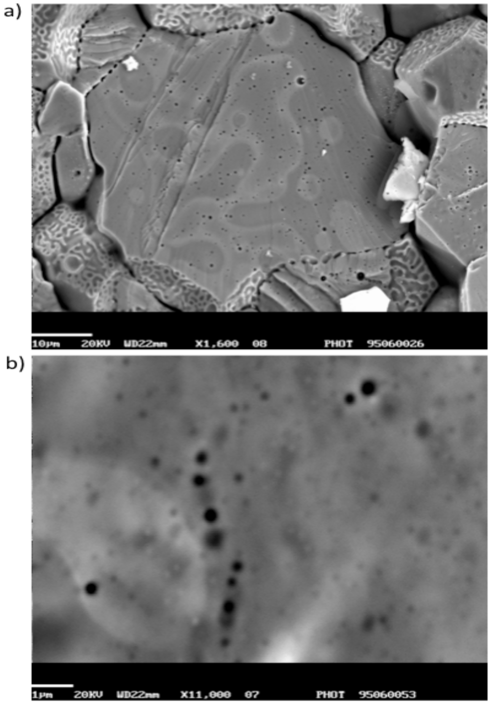

Uncertainty Quantification Approach

- **Probabilistic modeling of uncertainty**
- **Specify uncertain inputs based on expert opinion or employing Bayesian inference based on experimental data**
- **Identify important parameters using global sensitivity analysis (GSA)**
  - Employ Polynomial Chaos (PC) smoothing, Compressive Sensing (CS), and multilevel multifidelity methods as necessary to minimize need for high-fidelity computational samples
- **Propagate uncertainty forward through the coupled model hierarchy using generalized sparse quadrature methods**
- **Provide means for model selection using Bayesian evidence estimation**
- **Model validation**
  - Statistical comparisons between uncertain predictions and experimental data

NE Partnership: Work in progress

- **Collect expert opinion for uncertainty assessments of model input parameters (i.e. species reaction parameters)**
- **Perform global sensitivity analysis to identify important input parameters w.r.t. quantities of interest (QoIs)**
- **Construct sparse polynomial chaos surrogates to represent mapping from uncertain inputs to uncertain QoIs for performing efficient uncertainty propagation**

UO_2:

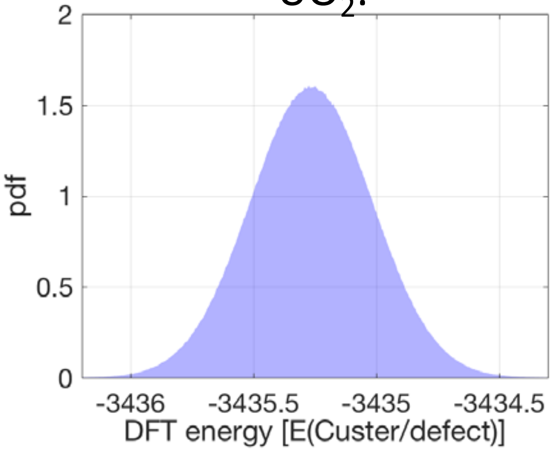

VU_2O_2:

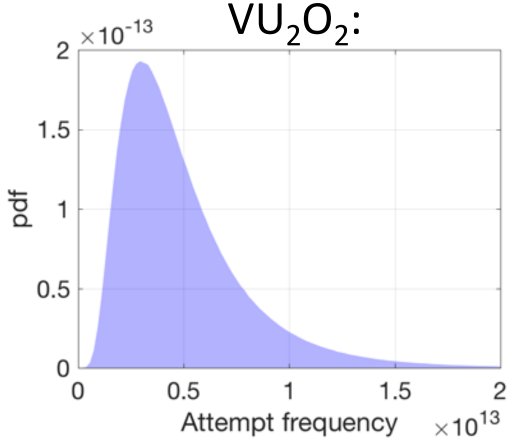

Eight core technology areas: Data Analytics

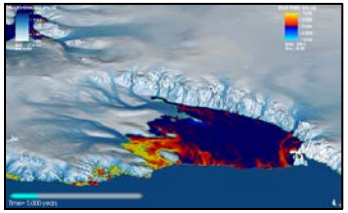

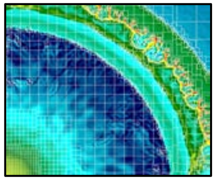

Structured Mesh Spatial Discretization

Data Analytics

Unstructured Mesh Spatial Discretization

Uncertainty Quantification

Synergistic activities link the eight areas together

Time Integrators

Numerical Optimization

Solution of Linear Systems

Solution of Eigenvalue Problems

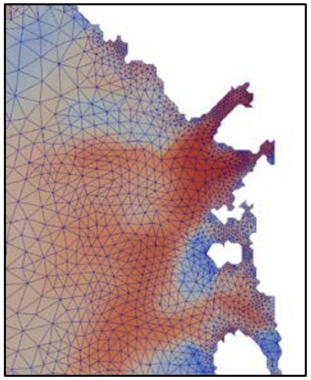

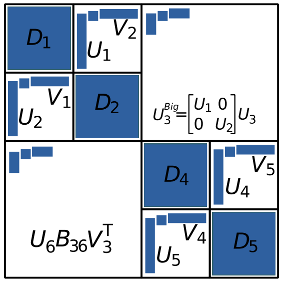

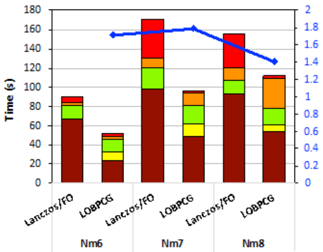

For more information contact: Rich Archibald (ORNL), archibaldrk@ornl.gov

Data Analytics

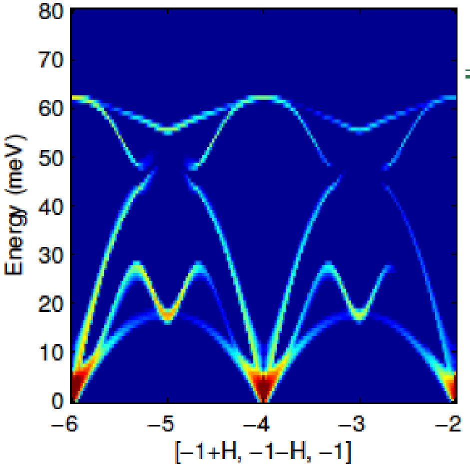

- **Goal:** Sparse functional representation of data, to enable faster IO and analysis of big datasets
- **Software tools:** Tasmanian, PUMI, TAO
- **FASTMath Tasks**

Sparse IO of Big Data

- Advanced sampling techniques for distributed data
- Research adaptive methods for sparse representation
- Build accurate uncertainty estimates for sparse representations of data

Fast Estimation & Evaluation

- Develop and design high order regularizers that optimize functional representations of data
- Surrogate models that accelerate estimation and evaluation of sparse data approximation

Streaming Data Ranking

- Algorithms that maximize information transfer
- Ordered sparse functional representations of data
- Parallel methods for streaming distributed datasets

- **Applications:** Fusion, workflows, experimental facilities
- **For more information:** Rick Archibald (ArchibaldRK@ornl.gov)

FASTMath provides in-situ analysis of streaming climate data

Scientific Achievement

Developing effective online data monitoring and archiving strategy over temporal and spatial domains while respecting practical storage and memory capacity constraints

Significance and Impact

Streaming methods improve the quality and quantity of information available to scientists for analysis

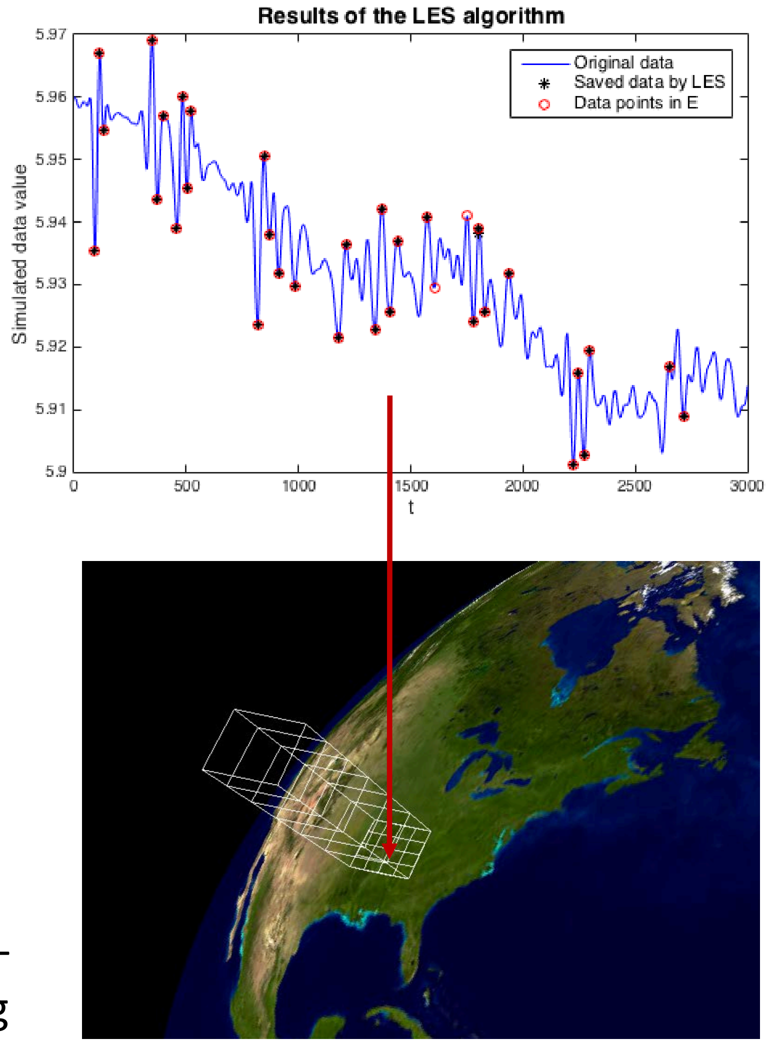

Research Details

- Typically averaged values or restricted temporal snap-shots of data can be saved in climate simulation
- Intelligently select and record the most informative extreme values in the raw data generated from real-time simulations in the context of better monitoring climate changes

X. Xian, R. Archibald, B. Mayer, K. Liu & J. Li, "An effective online data monitoring and saving strategy for large-scale climate simulations", Quality Technology & Quantitative Management, 10.1080/16843703.2017.1414112, 2018

U.S. DEPARTMENT OF ENERGY | Office of Science

FASTMATH

Work was performed at ORNL
For more information contact Rick Archibald:
archibaldrk@ornl.gov

Fusion SciDAC Research Workflow and Machine Learning

Summit

ITER

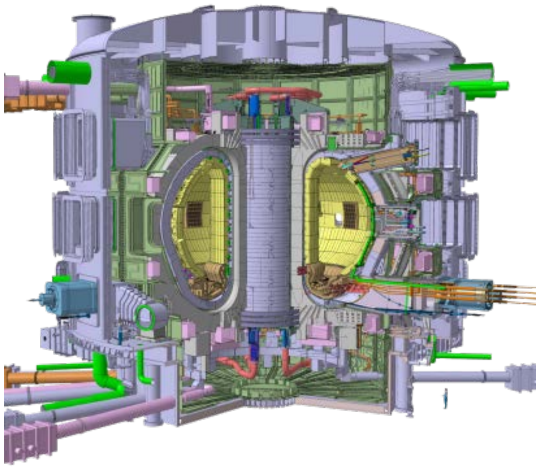

Preloading of Experimental data onto data nodes to support discovery, validation & ML

Large-scale computing

- Preconditioner, multi-scale, load balancing, fault-tolerance, meshing, etc algorithms need help from ML

- Data too big for filesystem
- In-memory on-the-fly analysis & filtering, and reduction: ML

Improvement, ML

Small scale computing

- Reduced physics models
- For quick experimental support
- Accuracy is an issue.

Experimental validation

- ITER data bigger than HEP data
- Possesses all 5 V properties
- Simulations could guide intra and inter experimental shots for faster progress: ML
- Two-way (near-) real-time data exchange between remote simulations & experiment, being supported by ML.

Prototyping of this workflow is being pursued among ADIOS, KSTAR and XGC.

Fusion Data Hierarchy

Fusion data possesses all 5V properties

- Volume
- Velocity
- Value
- Variety
- Veracity

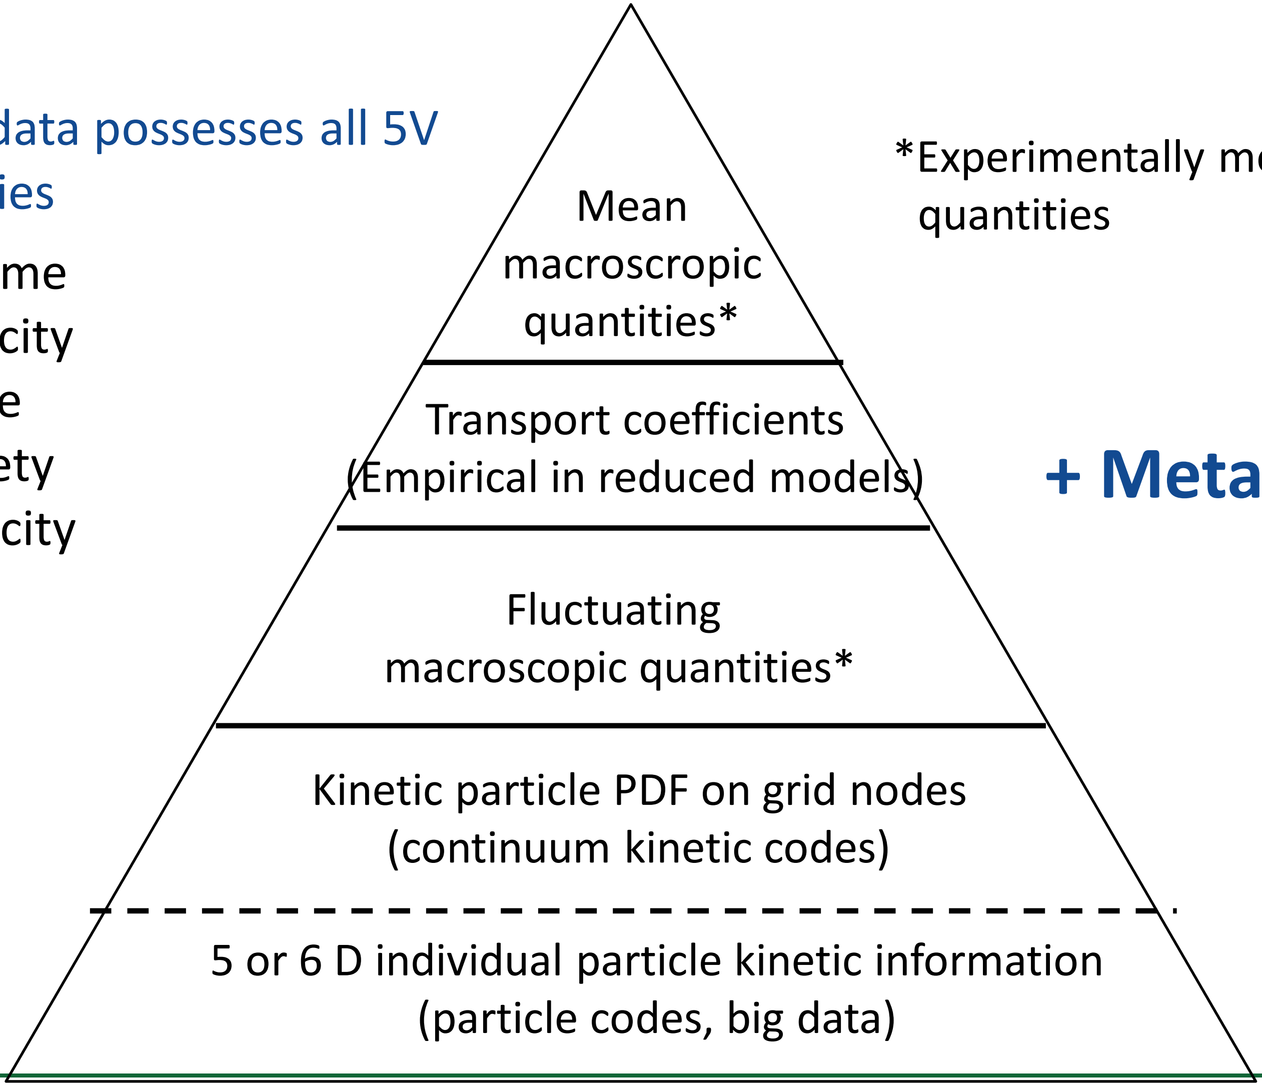

Mean macroscropic quantities*

Transport coefficients (Empirical in reduced models)

Fluctuating macroscopic quantities*

Kinetic particle PDF on grid nodes (continuum kinetic codes)

5 or 6 D individual particle kinetic information (particle codes, big data)

*Experimentally measurable quantities

+ Meta Data

Applied Math Recommendations and Challenges in Fusion SciDAC ML

- **Create a unified (or a few) salient and standardized data sets to create community target for ML development.**
  - Fusion is a multiscale-multiphysics science.
  - Each fusion SciDAC project is studying different physics or different space-time space, and looks for different feature
  - ML may need to target data sets from each SciDAC centers differently
- **Mathematical and UQ methods will need to have access to raw data for optimal ML results, while curation of them is essential to ensure key data bases are complete.**
  - Fusion SciDACs can allow the ML tools to be applied to the raw data after curation of them: in-memory for the in-situ simulation data, and raw diagnostics data before processing.
- **Given the nature of major fusion SciDA simulation data, in-memory in-situ ML is necessary. Adaptation of current tools is required in FASTMath in collaboration with the proper SciDAC teams.**
- **Given the rich hierarchical data and simulation optimization needs in Fusion SciDACs, three-way collaboration among FASTMath, RAPIDS and Fusion SciDAC appears necessary for timely progress.**

Math-CS collaborations led to significant advances in SciDAC-3 application sciences

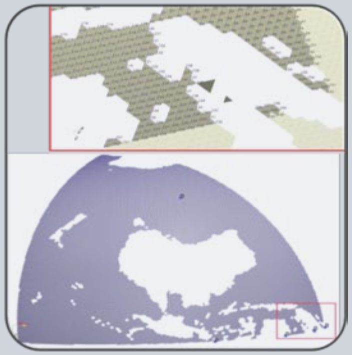

Increased accuracy of MPAS Ocean modeling system through improved scaling and improved partitioning algorithms

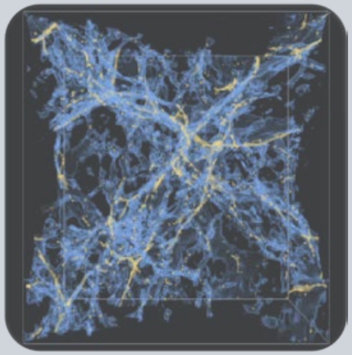

Development of highly-scalable, many core-aware multigrid solver increased ability of cosmologies to simulated evolution and fate of the universe (HPGMG in BoxLib)

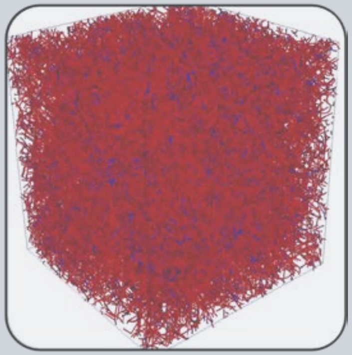

New nonlinear solver and evaluation of OpenMP performance enabled largest dislocation dynamics simulation using ParaDiS

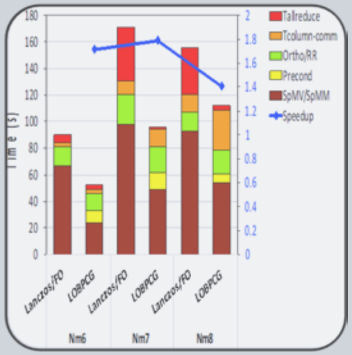

Improved efficiency of eigensystem solution reduced time to solution in large, sparse ab initio calculations in nuclear structure

FASTMath and RAPIDS will actively collaborate to continue to improve math libraries and application experience

- **Performance improvements to math library software**
  - Improved scaling (identify performance bottlenecks, find ‘performance bugs’, eliminate unnecessary communication)
  - Improved on-node performance (programming models, memory)
- **Using performance models to improve foundational understanding of algorithms**
- **Advanced visualization tools for FASTMath tools (e.g., AMR)**
- **In situ visualization tools used in unstructured mesh simulation workflow**
- **Use of CS abstractions to improve or accelerate application development**
  - Domain Specific Language compilers/tools
  - Leverage abstractions developed by RAPIDS for I/O to unify application experience

Organizational Structure

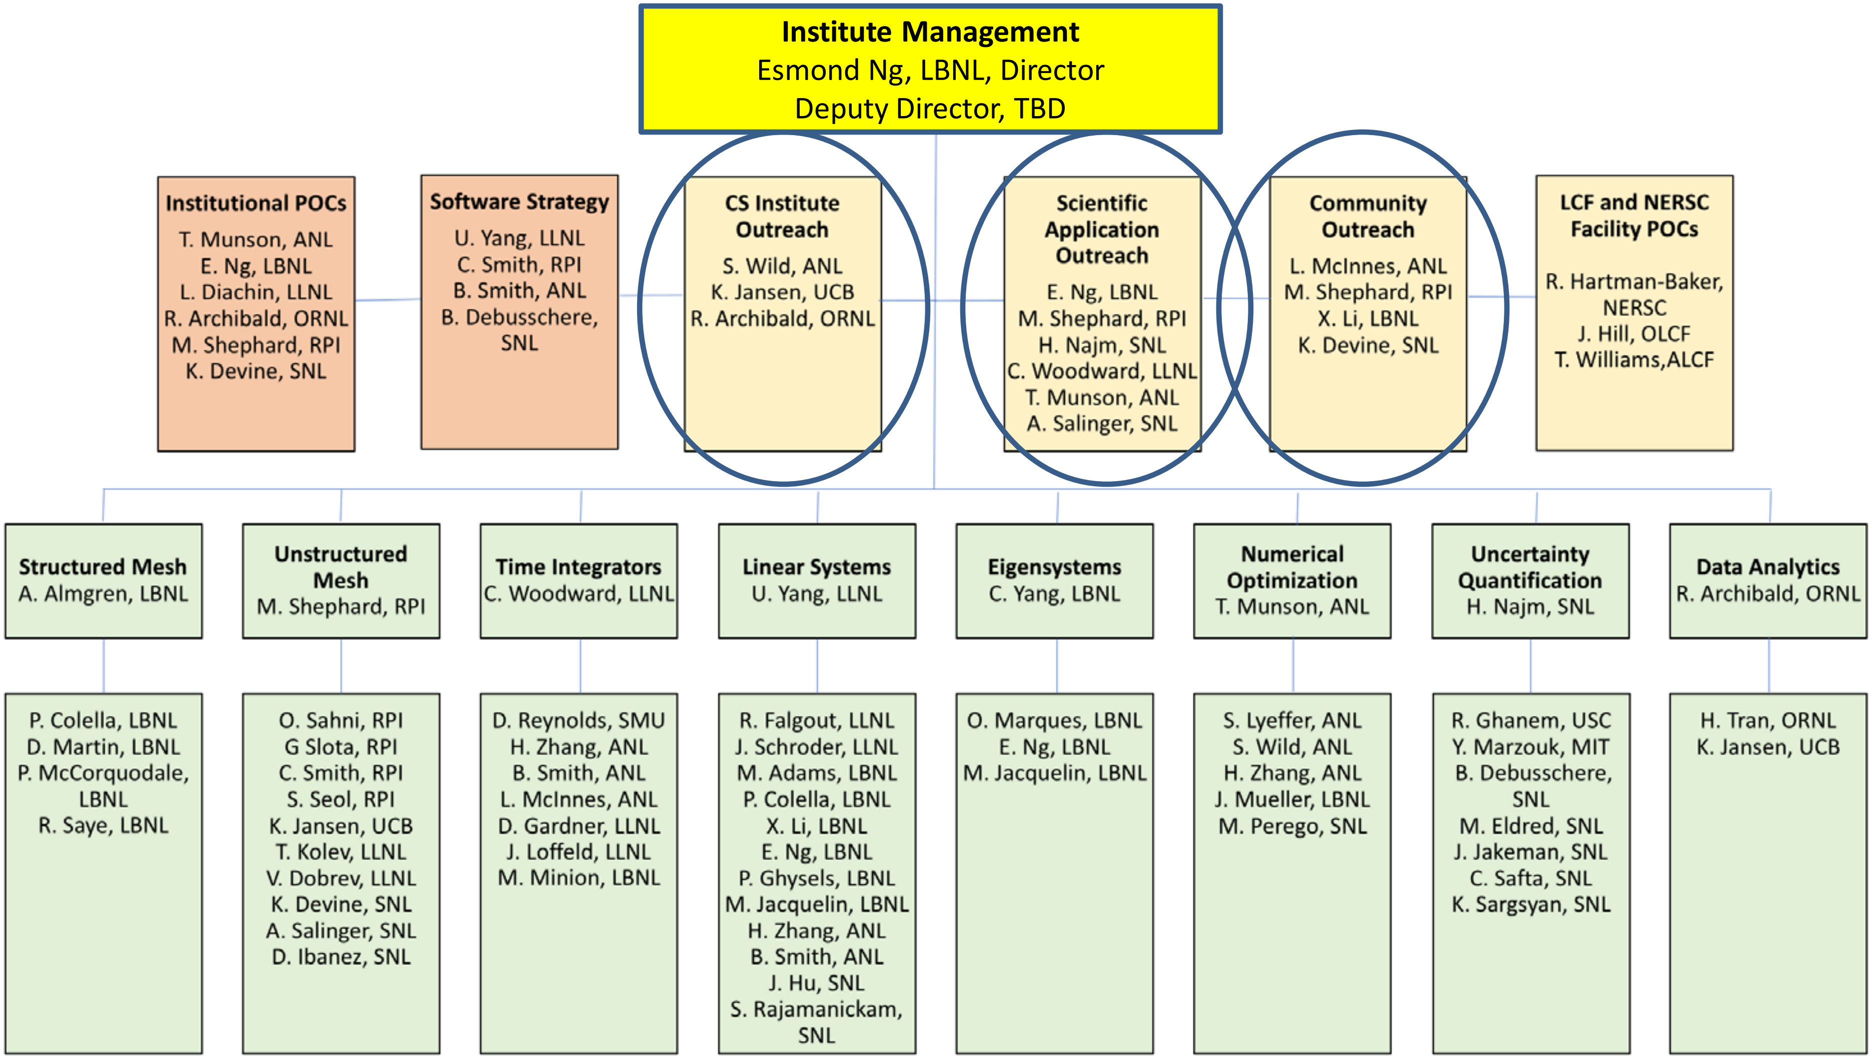

Institute Management
Esmond Ng, LBNL, Director
Deputy Director, TBD

Institutional POCs
T. Munson, ANL
E. Ng, LBNL
L. Diachin, LLNL
R. Archibald, ORNL
M. Shephard, RPI
K. Devine, SNL

Software Strategy
U. Yang, LLNL
C. Smith, RPI
B. Smith, ANL
B. Debusschere, SNL

CS Institute Outreach
S. Wild, ANL
K. Jansen, UCB
R. Archibald, ORNL

Scientific Application Outreach
E. Ng, LBNL
M. Shephard, RPI
H. Najm, SNL
C. Woodward, LLNL
T. Munson, ANL
A. Salinger, SNL

Community Outreach
L. McInnes, ANL
M. Shephard, RPI
X. Li, LBNL
K. Devine, SNL

LCF and NERSC Facility POCs
R. Hartman-Baker, NERSC
J. Hill, OLCF
T. Williams,ALCF

Structured Mesh
A. Almgren, LBNL

P. Colella, LBNL
D. Martin, LBNL
P. McCorquodale, LBNL
R. Saye, LBNL

Unstructured Mesh
M. Shephard, RPI

O. Sahni, RPI
G Slota, RPI
C. Smith, RPI
S. Seol, RPI
K. Jansen, UCB
T. Kolev, LLNL
V. Dobrev, LLNL
K. Devine, SNL
A. Salinger, SNL
D. Ibanez, SNL

Time Integrators
C. Woodward, LLNL

D. Reynolds, SMU
H. Zhang, ANL
B. Smith, ANL
L. McInnes, ANL
D. Gardner, LLNL
J. Loffeld, LLNL
M. Minion, LBNL

Linear Systems
U. Yang, LLNL

R. Falgout, LLNL
J. Schroder, LLNL
M. Adams, LBNL
P. Colella, LBNL
X. Li, LBNL
E. Ng, LBNL
P. Ghysels, LBNL
M. Jacquelin, LBNL
H. Zhang, ANL
B. Smith, ANL
J. Hu, SNL
S. Rajamanickam, SNL

Eigensystems
C. Yang, LBNL

O. Marques, LBNL
E. Ng, LBNL
M. Jacquelin, LBNL

Numerical Optimization
T. Munson, ANL

S. Lyeffer, ANL
S. Wild, ANL
H. Zhang, ANL
J. Mueller, LBNL
M. Perego, SNL

Uncertainty Quantification
H. Najm, SNL

R. Ghanem, USC
Y. Marzouk, MIT
B. Debusschere, SNL
M. Eldred, SNL
J. Jakeman, SNL
C. Safta, SNL
K. Sargsyan, SNL

Data Analytics
R. Archibald, ORNL

H. Tran, ORNL
K. Jansen, UCB

For more information…

- **Contacts:**
  - Esmond Ng, Institute Director (egng@lbl.gov)
  - Any of the numerical components leads
- **Web site:**
  - www.fastmath-scidac.org

See our posters today and tomorrow! All the area leads are here at the meeting and can provide more information.

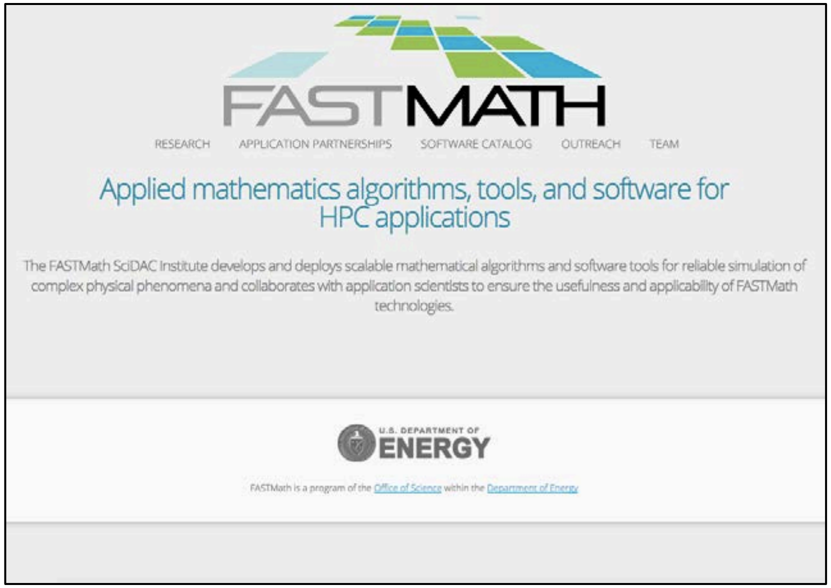

Applied mathematics algorithms, tools, and software for HPC applications

The FASTMath SciDAC Institute develops and deploys scalable mathematical algorithms and software tools for reliable simulation of complex physical phenomena and collaborates with application scientists to ensure the usefulness and applicability of FASTMath technologies.

FASTMath is a program of the Office of Science within the Department of Energy