LLNL-PRES-501654

FASTMath: Frameworks, Algorithms and Scalable Technologies for Mathematics

FASTMath Team

Lori Diachin, Institute Director

FASTMath SciDAC Institute

The FASTMath SciDAC project focuses on the development and use of mathematics software libraries

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

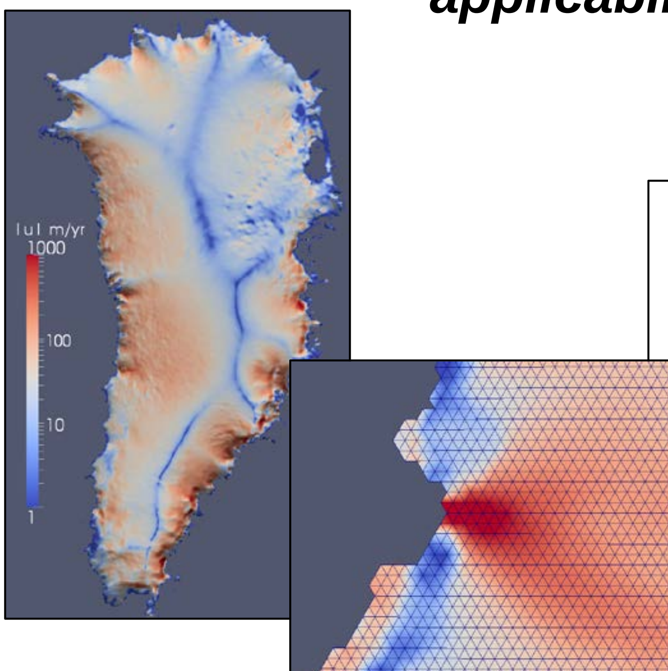

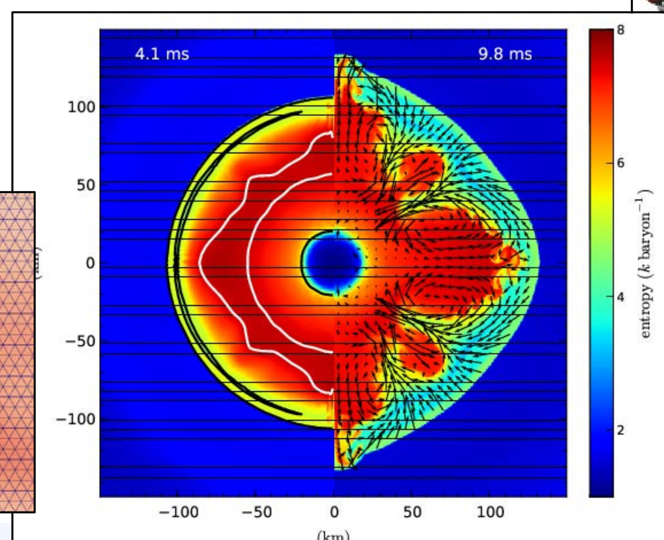

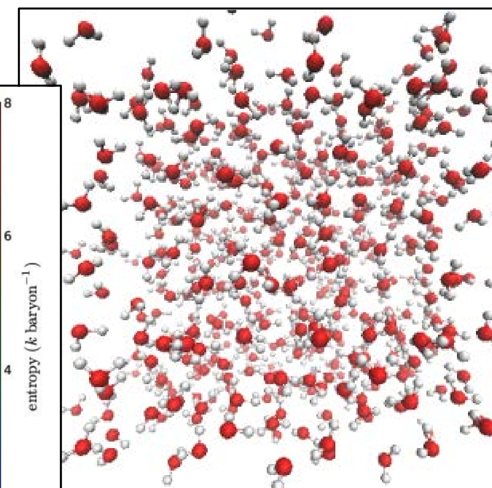

Argonne NATIONAL LABORATORY
BERKELEY LAB
Sandia National Laboratories
Berkeley University of California
Rensselaer
SMU
UBC THE UNIVERSITY OF BRITISH COLUMBIA

FASTMath encompasses three broad topical areas

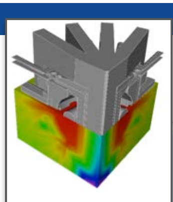

Tools for Problem Discretization

- Structured grid technologies
- Unstructured grid technologies
- Adaptive mesh refinement
- Complex geometry
- High-order discretizations
- Particle methods
- Time integration

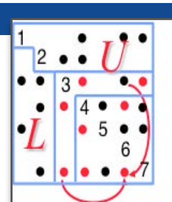

Solution of Algebraic Systems

- Iterative solution of linear systems
- Direct solution of linear systems
- Nonlinear systems
- Eigensystems
- Differential variational inequalities

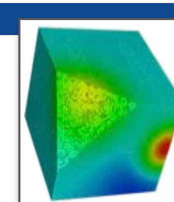

High Level Integrated Capabilities

- Adaptivity through the software stack
- Management of field data
- Coupling different physics domains
- Mesh/particle coupling methods

Argonne National Laboratory · Berkeley Lab · Sandia National Laboratories · Berkeley University of California · Rensselaer · SMU · UBC The University of British Columbia

FASTMath encompasses our algorithm development in widely used software

FASTMath Software

Structured Mesh Tools
- BoxLib (Ann Almgren)
- Chombo (Phil Colella)

Linear Solvers
- Hypre (Rob Falgout)
- PETSc (Barry Smith)
- SuperLU (Sherry Li)
- ML/Trilinos (Jonathan Hu)

Unstructured Mesh Tools
- PUMI (Seegyoung Seol)
- MeshAdapt (Mark Shephard)
- MOAB (Vijay Mahadevan)
- Mesquite (Lori Diachin)
- PHASTA (Ken Jansen)
- APF (Cameron Smith)

Nonlinear Solvers/Differential Variational Inequalities
- SUNDIALS (Carol Woodward
- PETSc (Barry Smith)
- NOX/Trilinos (Andy Salinger)

Partitioning Tools
- Zoltan (Karen Devine)
- ParMA (Mark Shephard)

Time Integrators
- SUNDIALS (Carol Woodward)
- PETSc (Barry Smith)

Eigensolvers
- PARPACK (Chao Yang)

Structured grid capabilities focus on high order, mapped grids, embedded boundaries, AMR and particles

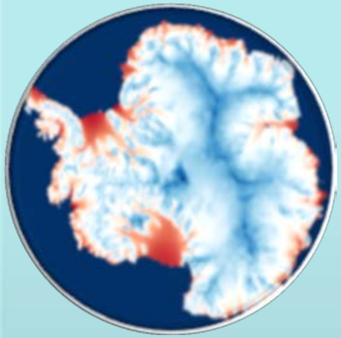

Structured AMR

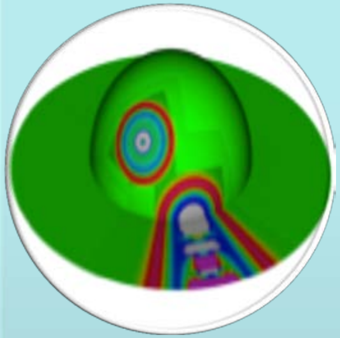

Mapped-multiblock grids

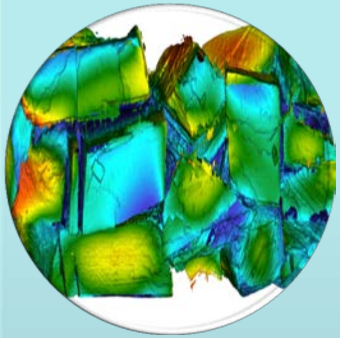

Embedded boundary methods

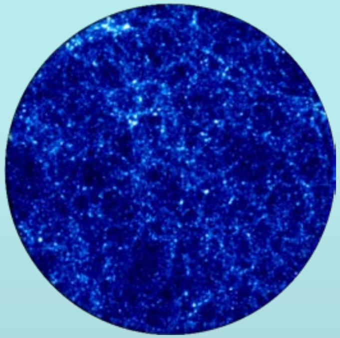

Particle-based methods

Application to cosmology, astrophysics, accelerator modeling, fusion, climate, subsurface reacting flows, low mach number combustion, etc.

Argonne National Laboratory · Berkeley Lab · Sandia National Laboratories · Berkeley University of California · Rensselaer · SMU · The University of British Columbia

Accomplishments in structured grids and particles

New AMR methods

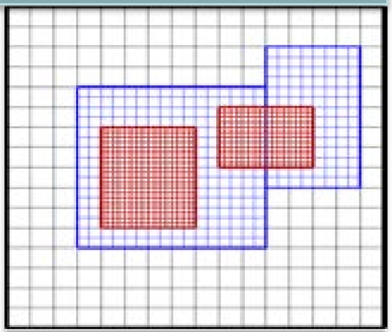

- Region based AMR for time stepping
- Fourth-order AMR for hyperbolic equations
- Allow mixed dimensionality

Mapped-multiblock FV Methods

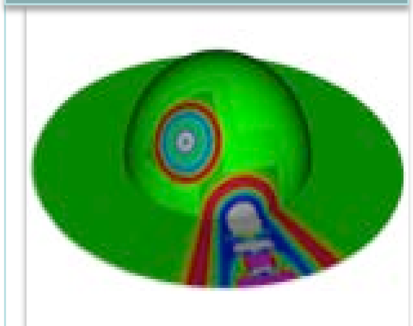

- Developing 4th order schemes
- Anisotropic refinement
- Mixed dimensions
- Extensions to PDEs on manifolds

Embedded Boundaries

- Grid generation from light source image data
- Interface to PETSc solvers
- Discretization for tensor operators

Particle methods

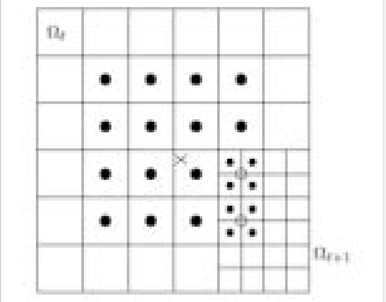

- Two-grid methods to improve scaling
- Space filling curve to assign particles to processors
- Load balancing schemes that separate particle and field calculations

Our unstructured grid capabilities focus on adaptivity, high order, and the tools needed for extreme scaling

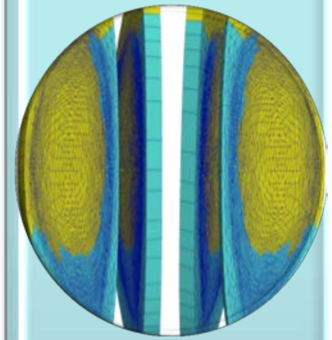

Parallel mesh infrastruct-ures

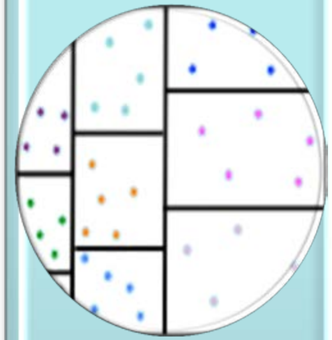

Dynamic load balancing

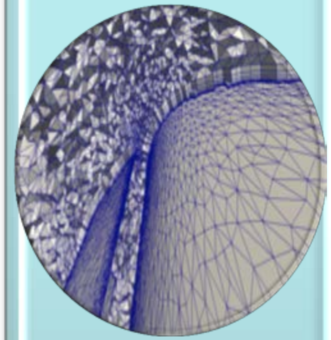

Mesh adaptation and quality control

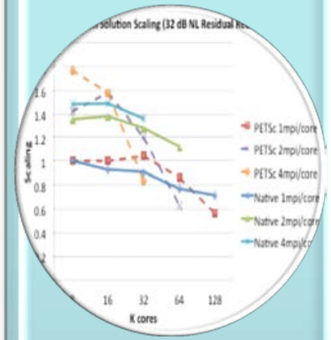

Parallel performance on unstructured meshes

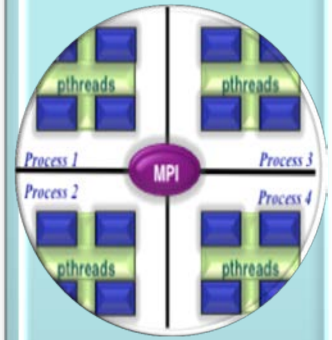

Architecture aware implementa-tions

Application to fusion, climate, accelerator modeling, NNSA applications, nuclear energy, manufacturing processes, etc.

Argonne National Laboratory · Berkeley Lab · Sandia National Laboratories · Berkeley · Rensselaer · SMU · The University of British Columbia

Accomplishments in unstructured grids

Parallel Unstructured Mesh Infrastructures

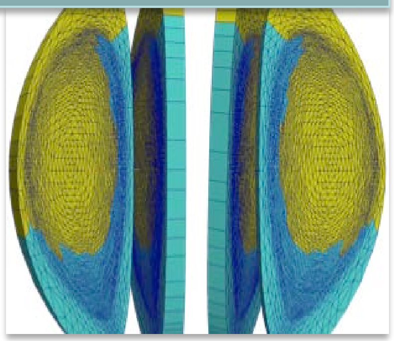

- PUMI scales to 750K cores and 92B elements
- MOAB scaling to 32K cores with improved memory access and Parallel I/O

Dynamic Load Balancing

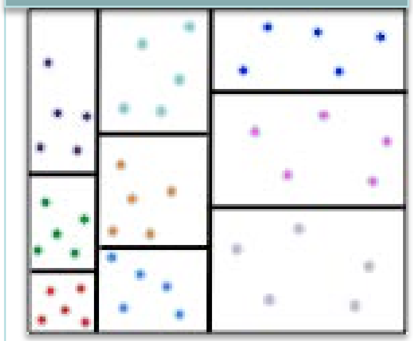

- Partitioning tools for hybrid programming
- Predictive load balancing to avoid memory problems
- Partitioning for multiple entity types
- Multilevel partitioning for large core counts (750K cores, 3.1M parts)

Adaptive Mesh Refinement

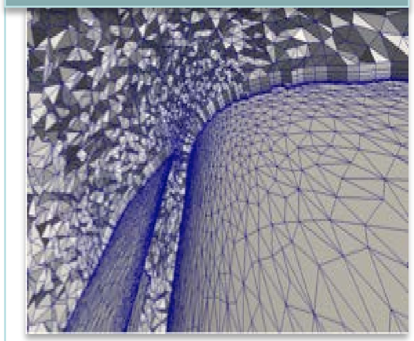

- Boundary layer adaptation (mixed meshes and boundary layer thickness)
- Support for curved geometries
- In-memory integration of MeshAdapt

Architecture-aware Implementations

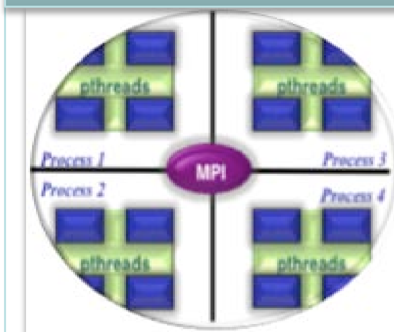

- Arch-aware task placement to reduce communication costs results in 34% reduction in stencil-based computational cost
- Parallel control utility developed

Our work on algebraic systems provides key solution technologies to applications

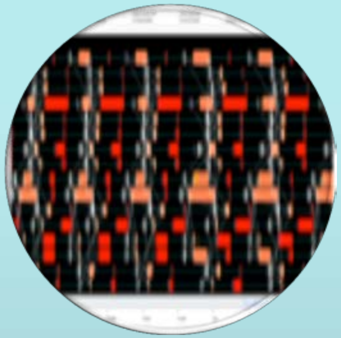

Linear system solution using direct and iterative solvers

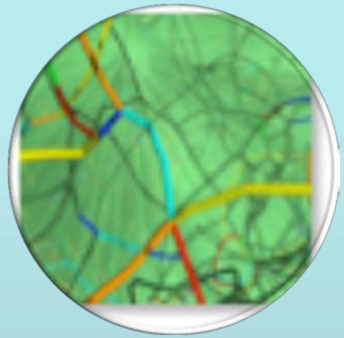

Nonlinear system solution using acceleration techniques and globalized Newton methods

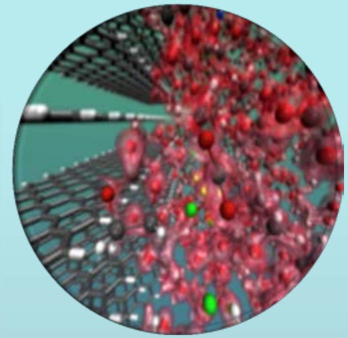

Eigensolvers using iterative techniques and optimization

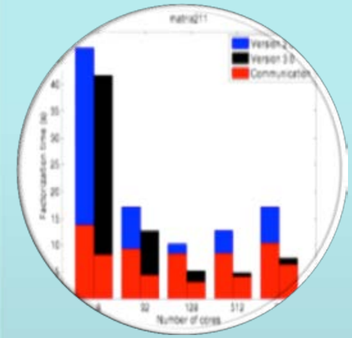

Architecture aware implementations

Application to fusion, nuclear structure calculation, quantum chemistry, accelerator modeling, climate, dislocation dynamics etc,

Accomplishments in algebraic solvers

Linear System Solution

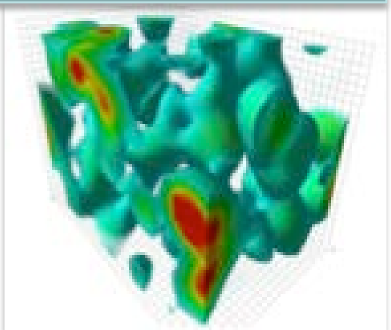

- Constrained energy minimization multigrid implemented
- HSS schemes for solving structured linear systems is 2X faster and uses 1/5 the memory of original

Nonlinear System Solution

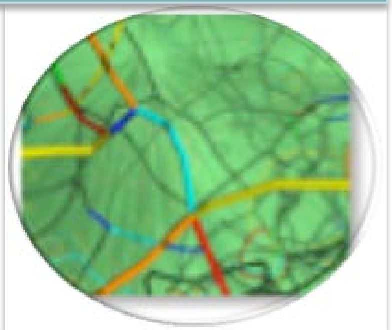

- Developed ARKode solver in SUNDIALS (IMEX method)
- SUNDIALS interfaced to SuperLU_MT
- Developed fixed point and Anderson accelerated solvers for SUNDIALS and NOX

Eigensolvers

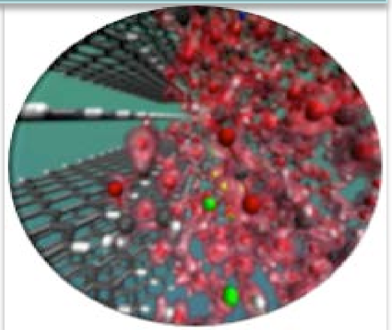

- Many eigenpairs of Hermitian matrices by reducing number of Rayleigh-Ritz calculations
- Topology aware data distributions for scalable eigensolver for nuclear configuration calcs
- Interior eigenvalues of non-Hermitian matrices

Architecture-aware Implementations

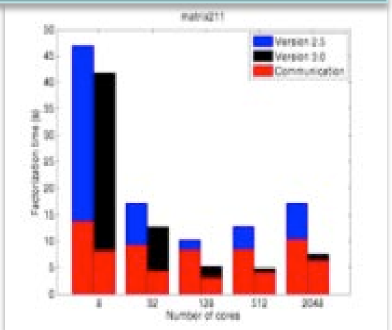

- Communication reduction in multigrid
- Communication bounds in iterative and direct methods
- Pipelined Krylov solvers to reduce global reductions
- New algorithms for DAG execution

Integrating technologies is a key value added by the FASTMath Institute

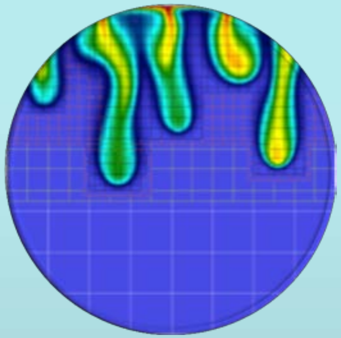

Mesh/solver interactions

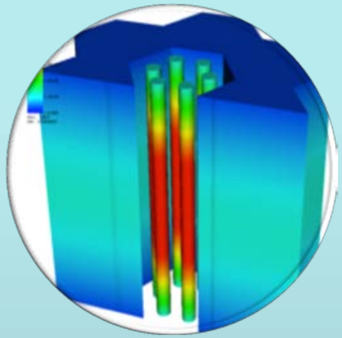

Mesh-to-mesh coupling methods

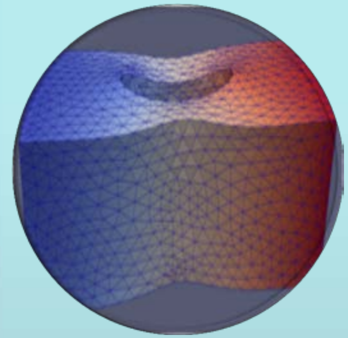

Unstructured mesh technologies into simulation workflows

Software unification strategies

Application to climate, plasma surface interactions, structural mechanics, nuclear energy, cosmology, fluid flow, etc.

Accomplishments in integrated technologies

Mesh Solver Interactions

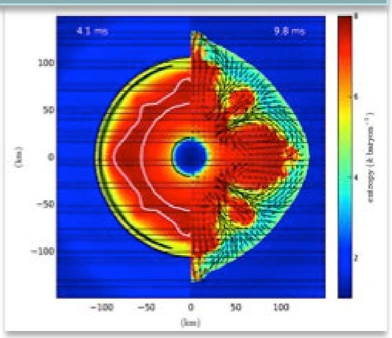

- BoxLib/PETSc, hypre, Trilinos interfaces for porous media and radiation
- Chombo/PETSc for climate and pore scale modeling
- PHASTA-PUMI-PETSc for compressible flow
- MOAB/PETSc for plasma surface interactions

Mesh-to-mesh coupling

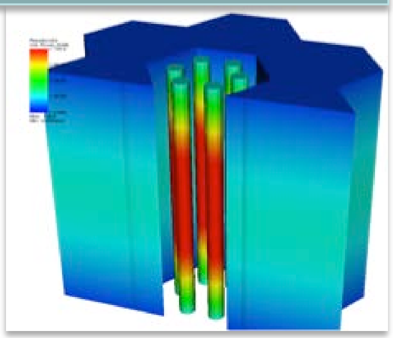

- Developed CouPE for tight multi-physics coupling with loose coupling interfaces
- Strong scalability up to 512K cores for coupling

Simulation Workflows

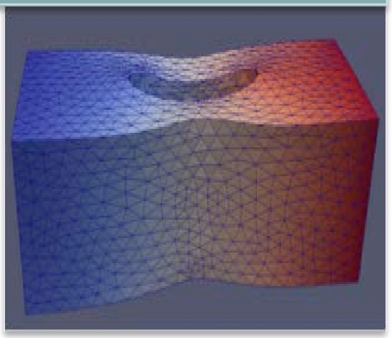

- Eliminate file I/O in parallel adaptive simulations
- Attached parallel fields interface to support data and information transfer
- Integrated PUMI, Albany, MeshAdapt, Zoltan

Software unification

- Tutorials lower barriers to SW adoption
- Common build/configure practices
- One stop shopping web site
- Reference installations at ANL and LBNL

We have helped the application teams significantly reduce time to solution in their simulations

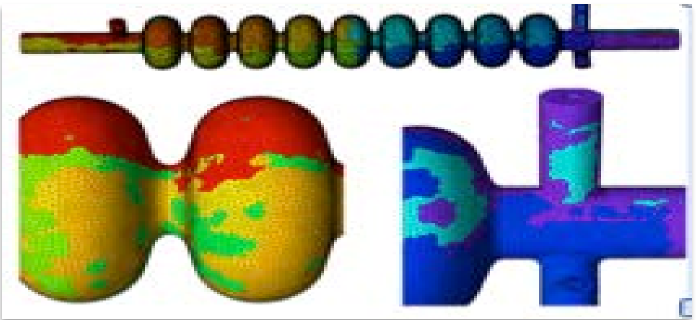

Sparse direct solves improve time to solution 20X for accelerators allowing 8 cavity simulation (Spentzouris)

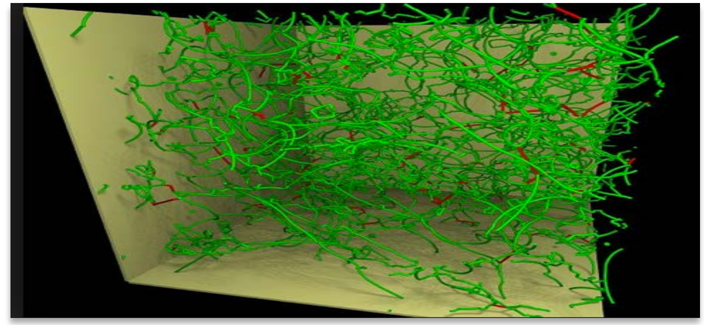

Acceleration-based nonlinear solvers speed up dislocation dynamics 35-50%; multistage Runge-Kutta methods reduce time steps by 94% (Arsenlis)

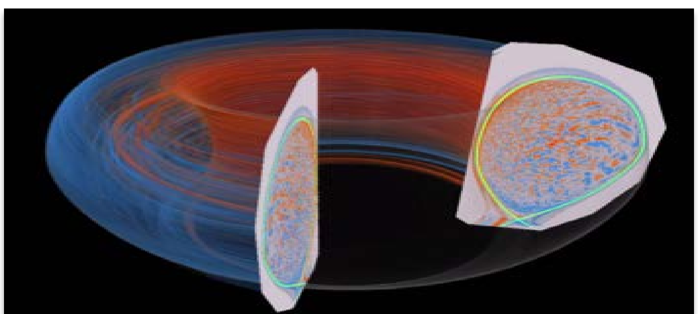

Sped up flux surface creation to improve 2D mesh generation in fusion application from 11.5 hours to 1 minute (Chang)

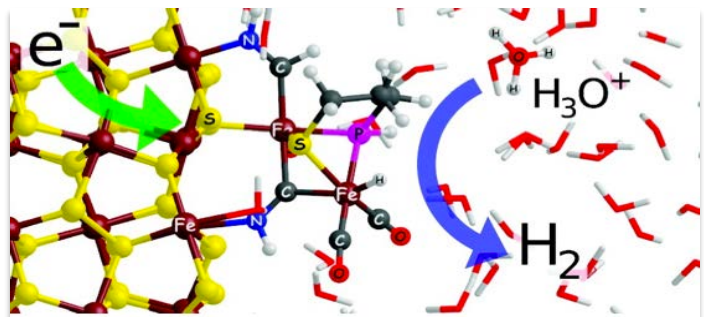

Sophisticated eigensolvers significantly improve materials calculations in many domains including ions in solution (Car), excited state phenomenon (Chelikowsky, Head-Gordon)

We have helped the application teams achieve unprecedented resolution and increased reliability

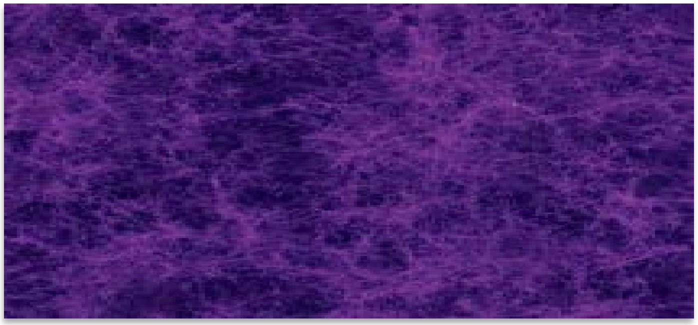

Astrophysics Lyman-α forest simulation at 4096^3 in an 80Mpc/h box; produced statistics at 1% accuracy for first time (Habib)

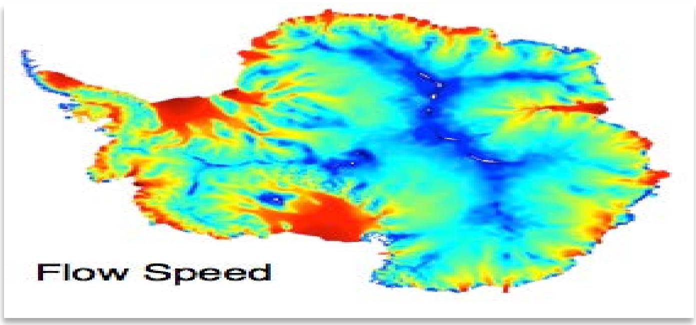

Predictions of grounding line match experiment for first time in ice sheet modeling due to AMR (Price)

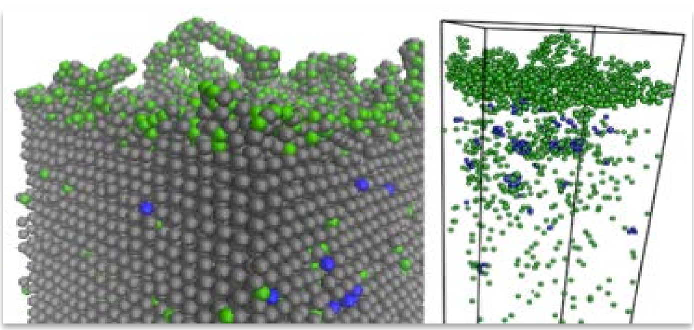

Implicit ODE integrators combined with AMG linear solvers enables solution of 4D reaction-diffusion eqns for plasma surface interactions (Wirth)

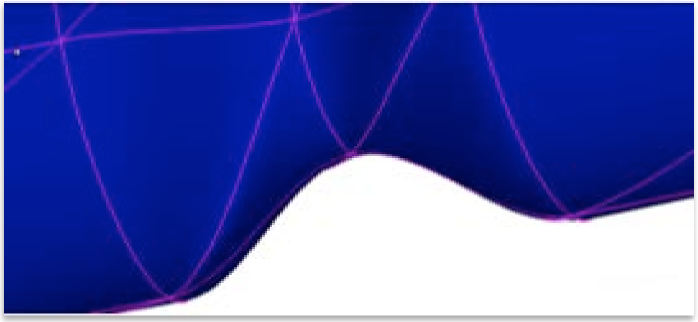

High-order unstructured meshes for particle accelerators overcome mesh generation/ adaptation bottlenecks (Spentzouris)

Argonne NATIONAL LABORATORY · BERKELEY LAB · Sandia National Laboratories · Berkeley · Rensselaer · SMU · UBC THE UNIVERSITY OF BRITISH COLUMBIA

Our strategy toward architecture awareness focuses on both inter- and intra-node issues

Inter-node: Massive Concurrency

- Reduce communication
- Increase concurrency
- Reduce synchronization
- Address memory footprint
- Enable large communication/computation overlap

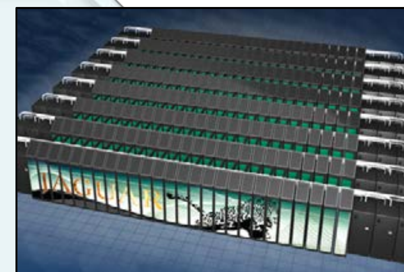

Intra-node: Deep NUMA

- MPI + threads for many packages
- Compare task and data parallelism
- Thread communicator to allow passing of thread information among libraries
- Low-level kernels for vector operations that support hybrid programming models

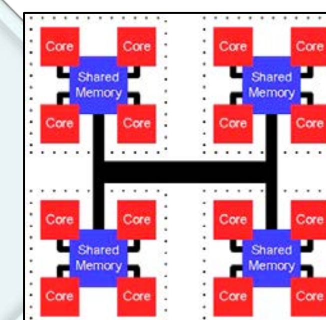

We are developing new algorithms that address architecture awareness

Reduce communication

- AMG: develop non-Galerkin approaches, use redundancy or agglomeration on coarse grids, develop additive AMG variants (hypre) (2X improvement)
- Hierarchical partitioning optimizes communication at each level (Zoltan) (27% improvement in matrix-vector multiply)
- Relaxation and bottom solve in AMR multigrid (Chombo) (2.5X improvement in solver, 40% overall)
- HSS methods

Increase concurrency

- New spectrum slicing eigensolver in PARPACK (Computes 10s of thousands of eigenvalues in small amounts of time)
- New pole expansion and selected inversion schemes (PEXSI) (now scales to over 100K cores)
- Utilize BG/Q architecture for extreme scaling demonstrations (PHASTA) (3.1M processes on 768K cores unstructured mesh calculation)

Reduce synchronization points

- Implemented pipelined versions of CG and conjugate residual methods; 4X improvement in speed (PETSc) (30% speed up on 32K cores)

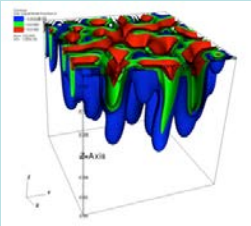

Used in PFLOTRAN applications

Address memory footprint issues

- Predictive load balancing schemes for AMR (Zoltan) (Allows AMR runs to complete by maintaining memory footprint)
- Hybrid programming models

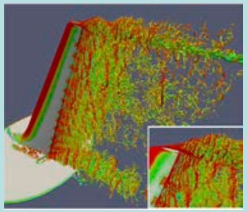

Used in PHASTA extreme scale applications

Increase communication and computation overlap

- Improved and stabilized look-ahead algorithms (SuperLU) (3X run time improvement)

Used in Omega3P accelerator simulations

We are refactoring our software to support hybrid programming models

Block structured AMR in Chombo (MPI + OpenMP)	• Adding course grained thread loop over blocks and micro-blocking reduced communication costs and memory footprint; performance improvements limited
Unstructured grids using Parallel Control Utility (MPI and threads)	• Utility layer that allows support of both MPI and threads; showed 30% efficiency improvement and 10% memory reduction on BG/Q
Partitioning with Zoltan2 (MPI + OpenMP)	• Implementation in multi-dimensional jagged geometric partitioning scaled to 8B elements on 64K nodes
Direct linear solvers in SuperLU_DIST (MPI +OpenMP + CUDA)	• Aggregation of small BLAS operations into larger ones to hide long-latency operations resulted in 2.7X faster performance and 5X reduction in memory costs on 100-node GPU cluster
Linear solver kernels (Pthreads and OpenMP)	• Introduction of a thread communicator allows passing this information among libraries for portable performance with maintainable kernels
Time integrators in SUNDIALS (MPI + OpenMP)	• New threaded kernels and integration with SuperLU_MT provide speed up and flexibility

FASTMath heavily leverages and helps drive research in the base research programs at ASCR

FASTMath technologies leverage Applied Math base

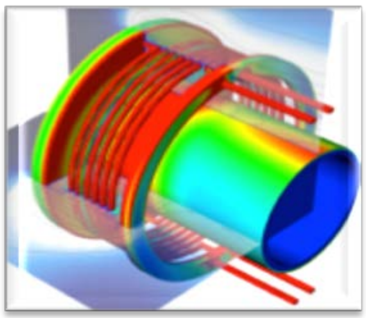

Examples: Algorithms in Chombo, BoxLib, PETSc, SUNDIALS, hypre, ML and Mesquite

Image courtesy of hypre Project

Driving requirements in the Applied Math base program

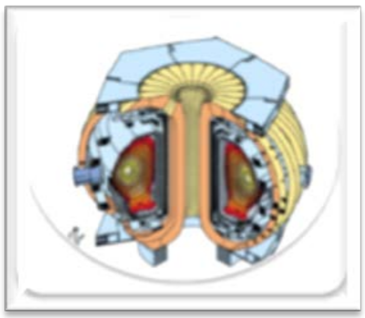

Example: Fusion collaboration led to a new base math project in high-order discretization methods

Image courtesy of ESL Project

New joint math/cs projects in Exascale solvers initiative

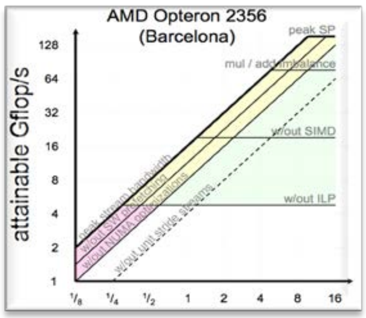

Example: Discussions between FASTMath and Super led to ExReDi Project

Image courtesy of Sam Williams, LBNL

Use of FASTMath capabilities in base CS projects

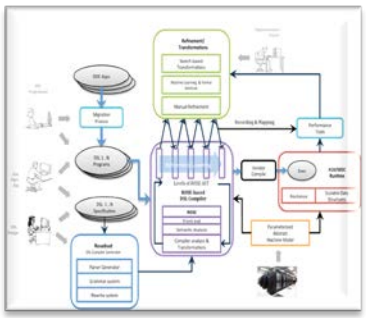

Example: DTEC X-Stack project is exploring use of domain specific languages in Chombo

Image courtesy of DTEC Project

The FASTMath team includes experts from four national laboratories and six universities

Lawrence Berkeley National Laboratory

Mark Adams
Ann Almgren
Phil Colella
Anshu Dubey
Dan Graves
Sherry Li
Lin Lin
Terry Ligocki
Mike Lijewski
Peter McCorquodale
Esmond Ng
Brian Van Straalen
Chao Yang
Subcontract: Jim Demmel (UC Berkeley)

Lawrence Livermore National Laboratory

Barna Bihari
Lori Diachin
Milo Dorr
Rob Falgout
Mark Miller
Jacob Schroder
Carol Woodward
Ulrike Yang
Subcontract: Carl Ollivier-Gooch (Univ of British Columbia)
Subcontract: Dan Reynolds (Southern Methodist)

Rensselear Polytechnic Inst.

E. Seegyoung Seol
Onkar Sahni
Mark Shephard
Cameron Smith
Subcontract: Ken Jansen (UC Boulder)

Argonne National Laboratory

Jed Brown
Lois Curfman McInnes
Todd Munson
Vijay Mahadevan
Barry Smith
Subcontract: Jim Jiao (SUNY Stony Brook)
Subcontract: Paul Wilson (Univ of Wisconsin)

Sandia National Laboratories

Karen Devine
Glen Hansen
Jonathan Hu
Vitus Leung
Siva Rajamanickam
Michel Wolf
Andrew Salinger

FASTMath posters provide more details

1. **FASTMath Structured Meshing Technologies**
   POC: Phil Colella, LBNL
2. **FASTMath Unstructured Mesh Technologies**
   POC: Mark Shephard, RPI
3. **FASTMath Partitioning and Task Placement**
   POC: Karen Devine, SNL
4. **FASTMath Iterative Solver Technologies**
   POC: Ulrike Yang, LLNL
5. **FASTMath Direct Solver Technologies**
   POC: Sherry Li, LBNL
6. **FASTMath Nonlinear and ODE Solver Technologies**
   POC: Carol Woodward, LLNL
7. **FASTMath Eigensolver Technologies**
   POC: Esmond Ng, LBNL
8. **FASTMath Unstructured Mesh/Solver Interactions**
   POC: Vijay Mahadevan, ANL
9. **Construction of Explicit (PETSc) sparse matrices from (Chombo) AMR Hierarchies**
   POC: Mark Adams, LBNL
10. **FASTMath Component-Based Unstructured Mesh Simulation Workflows**
    POC: Glen Hansen, SNL

For more information, please contact any of the following or visit our web site

- FASTMath Institute Director:
  - Lori Diachin,
    diachin2@llnl.gov, 925-422-7130

- FASTMath Executive Council
  - Phil Colella, Structured Mesh Tools
    pcolella@lbl.gov, 510-486-5412
  - Esmond Ng, Nonlinear/Eigensolvers
    egng@lbl.gov, 510-495-2851
  - Andy Salinger, Integrated Technologies
    agsalin@sandia.gov, 505-845-3523
  - Mark Shephard, Unstructured Mesh Tools
    shephard@scorec.rpi.edu, 518-276-8044
  - Barry Smith, Linear Solvers,
    bsmith@mcs.anl.gov, 630-252-9174

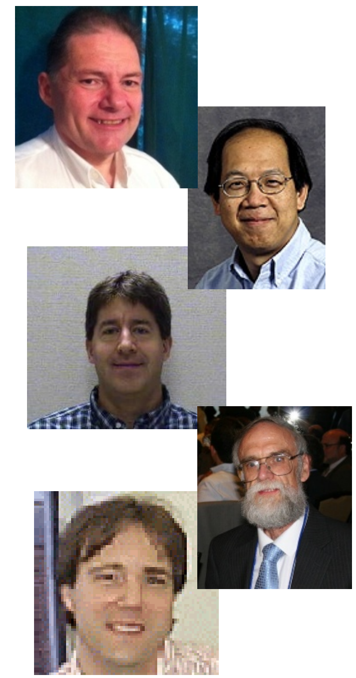

http://www.fastmath-scidac.org