LLNL-PRES-674972

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

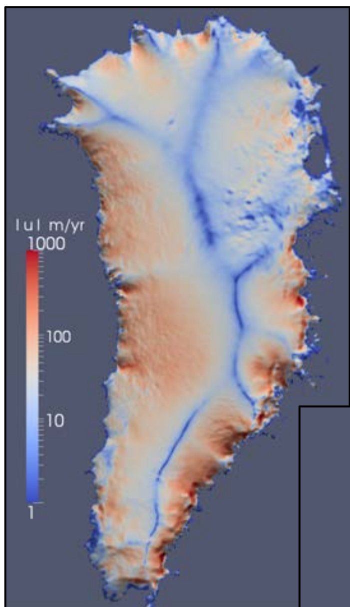

- ***Structured Meshes***
- ***Unstructured Meshes***
- ***Solution of Algebraic Systems***
- ***Integrated Technologies***

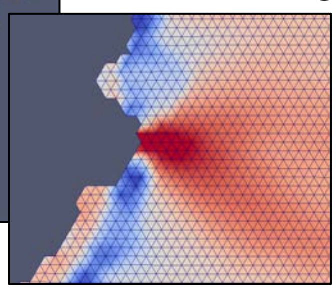

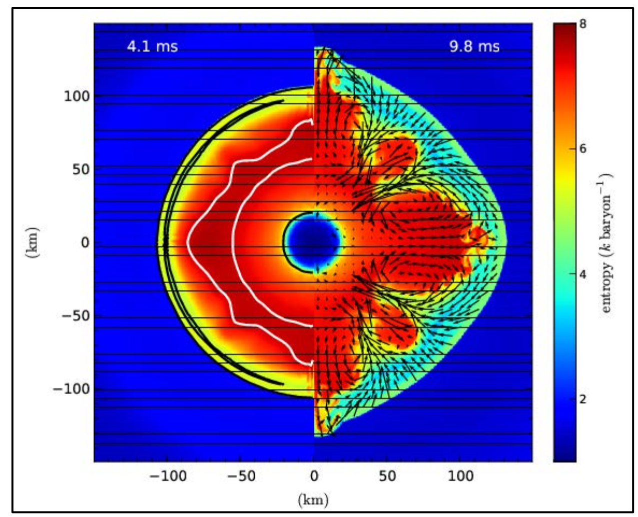

This presentation provides an overview of select FASTMath accomplishments

- Vignettes in each of the broad areas FASTMath encompasses
  - Motivation
  - State of the technology at the beginning of SciDAC-3
  - Key gaps
  - Approach
  - Main Accomplishments and Impact
- Mix of algorithmic advances, architecture-aware advances, and application advances
- Represents only a portion of FASTMath work

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

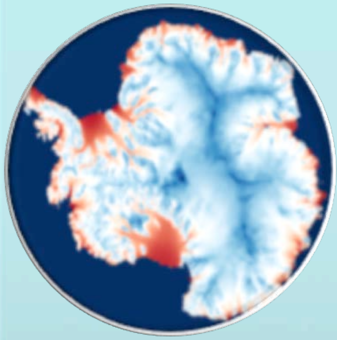

Structured AMR

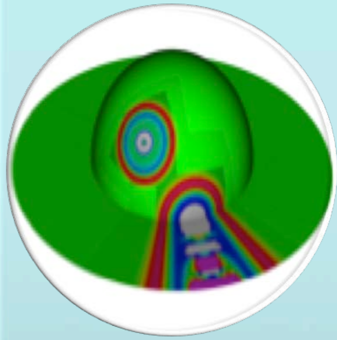

Mapped-multiblock grids

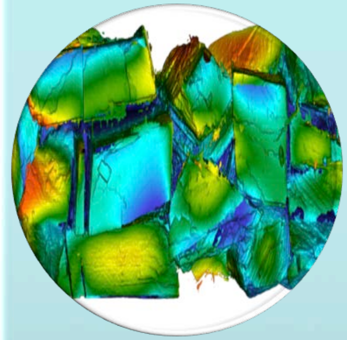

Embedded boundary methods

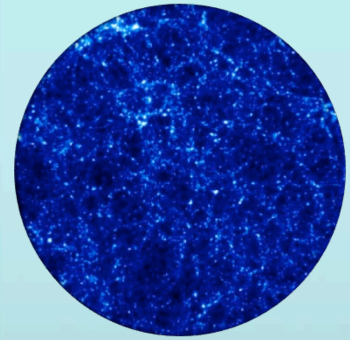

Particle-based methods

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

A new tiling framework for structured AMR improves suitability for modern, many-core NUMA architectures

- **Goals in 2011:** Improve loop-level threading which was ineffective for many-core systems
- **Approach**: Incorporate tiling constructs into core BoxLib framework to decrease the number of working sets, increase data locality and parallelism
  - Enables tiling w/ minimal re-coding of complex kernels
  - Demonstrated scalability with100's of threads/ processor
- **Impact:** Immediate impact across broad range of BoxLib-based multiphysics PDE applications
  - Combustion
  - Cosmology
  - Astrophysics
  - Incompressible flows

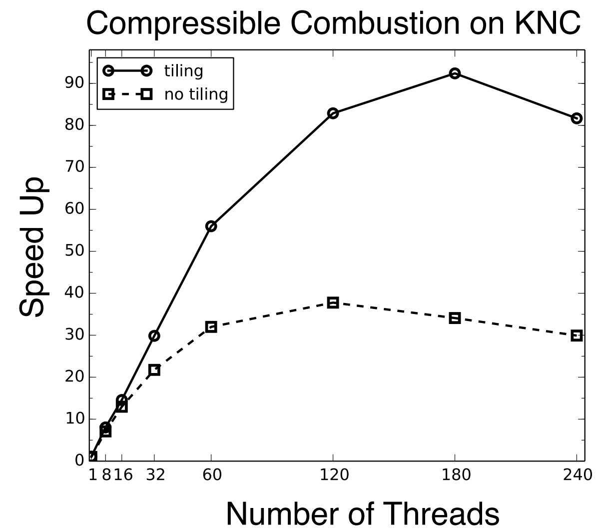

An Architecture Aware Algorithmic Advance

Contributors: A. Almgren, W. Zhang, T. Nguyen, J. Shalf, M. Day (LBNL), D. Unat (Koç, Turkey)

Robust software for embedded boundary grids enables new science in pore scale subsurface flow

- **Goal in 2011**: Significantly improve software needed to model flow in highly complex pore-scale geometries derived from image data
  - Prototype available
  - Optimization and productization required
- **Accomplishments**
  - Robust EB grid generation from image data
  - Large improvements in scalability, performance, memory usage
  - Integration of CrunchFlow reaction package, PETSc AMG solver
- **Impact**
  - Published science results in subsurface flow literature.
  - Ongoing engagement with subsurface flow team (Nanoscale CO_2 sequestration EFRC).

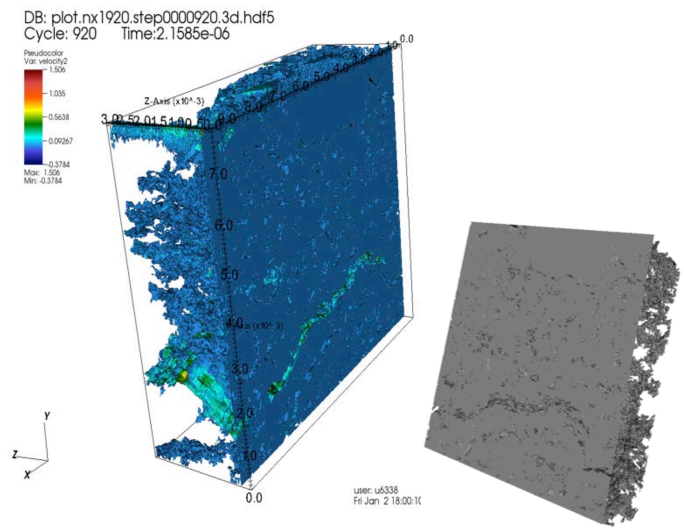

An Application Aware Algorithmic & SW Advance

Contributors: M. Adams, A. Dubey, D. Graves, T. Ligocki, D. Trebotich, B. Van Straalen (LBNL)

New particle-in-cell AMR algorithms have been developed for Vlasov equations

- **Goal in 2011**:
  - Develop this key numerical methodology for cosmology and plasma applications
  - Improve memory and load balance by separating particles from field grids
- **Accomplishments**
  - New scalable two-grid scheme using fast particle sorting package in Charm++
  - Improved performance, robustness of the existing infrastructure by changes to particle data structure.
- **Impact:** Enabled research in PIC algorithms: methods with well-defined convergence properties, increased arithmetic intensity and that eliminate noise and numerically-induced clumping (ExReDi RXSolver project).

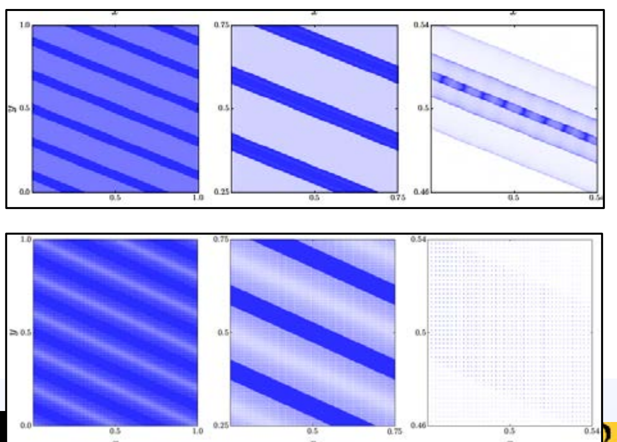

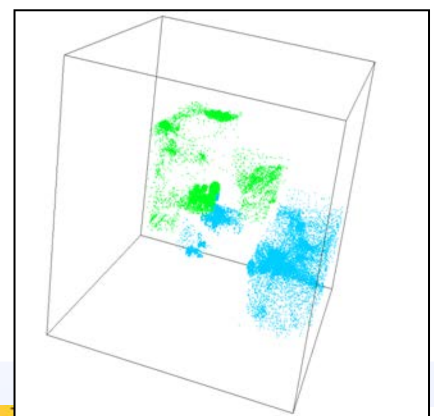

An Application Aware Algorithmic Advance

Contributors: P. Colella, A. Dubey, A. Myers, B. Van Straalen (LBNL)

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

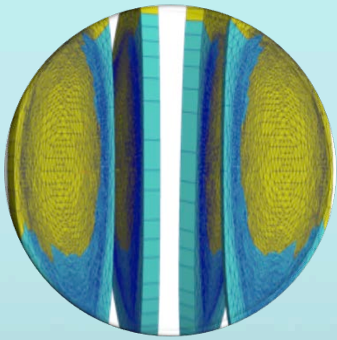

Parallel mesh infrastructures

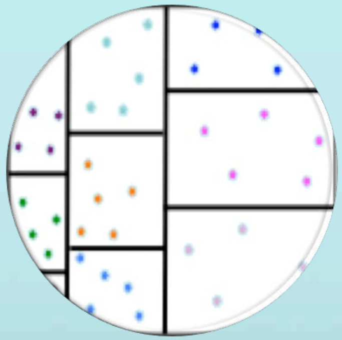

Dynamic load balancing

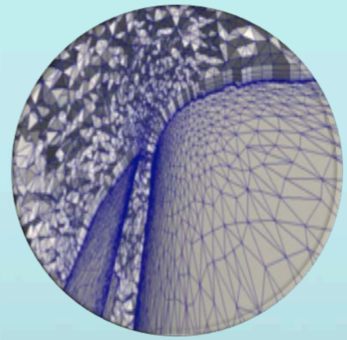

Mesh adaptation and quality control

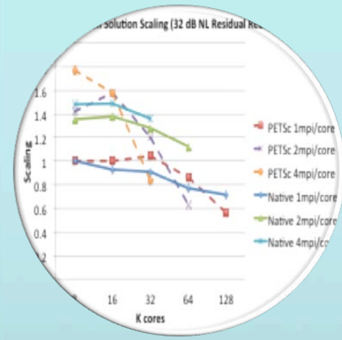

Parallel performance on unstructured meshes

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

Unstructured mesh methods are effective for discretizing fields in fusion plasma simulations

- **Goal in 2011**: Effectively use unstructured meshes in fusion application codes
  - Reduce mesh generation from hours to minutes
  - Incorporate AMR, parallel mesh infrastructure, and combined particle mesh methods
- **Approach:**
  - Develop field following techniques for mesh generation, combine meshing components and provide application specific control
  - Use FASTMath MeshAdapt driven by error indicators
  - Use FASTMath PUMI and Zoltan and employ directly in-memory integrations with mesh generation, mesh adaptation and solvers
- **Impact:** Support EPSI and CEMM fusion centers
  - Generation of better controlled meshes in minutes
  - Mesh adaptation effectively captures edge modes
  - M3D-C1 executes 2D to 3D simulations

Coarse mesh matched to the geometry and physics

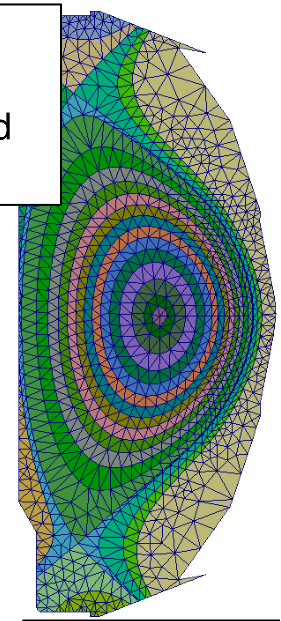

Adaptive mesh to capture edge modes

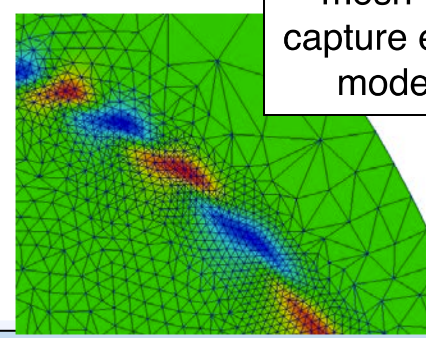

An Application Aware Algorithmic Advance

Contributors: E.S. Seol, F. Zhang, M.S. Shephard (RPI), R. Hager, S.-H. Ku, C.-S. Chang, S.C. Jardin (PPPL), N.M Ferraro (GA)

New algorithms improve application scalability through strategic placement of tasks on nodes

- **Goal in 2011:** Given a (possibly non-contiguous) allocation of nodes, assign interdependent MPI tasks to "nearby" nodes in the network
  - Develop metric to optimize (max hops, avg hops, congestion, etc.)
  - Create software for general, inexpensive task placement
- **Approach:**
  - For tasks, use geometric proximity as a proxy for interdependence
  - For nodes, use nodes' geometric coords in torus/mesh network
  - Apply inexpensive geometric partitioning algorithms to both application tasks and nodes, giving consistent ordering of both
- **Accomplishments:**
  - Determined *Average Number of Hops* is good proxy for communication costs
  - Reduced MiniGhost time on 64K by 34% relative to default and by 24% relative to node-only grouping
  - Reduced overall MueLu multigrid solver time by >10%

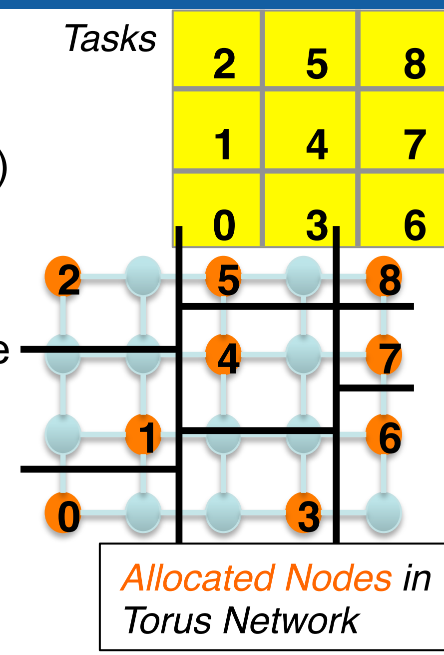

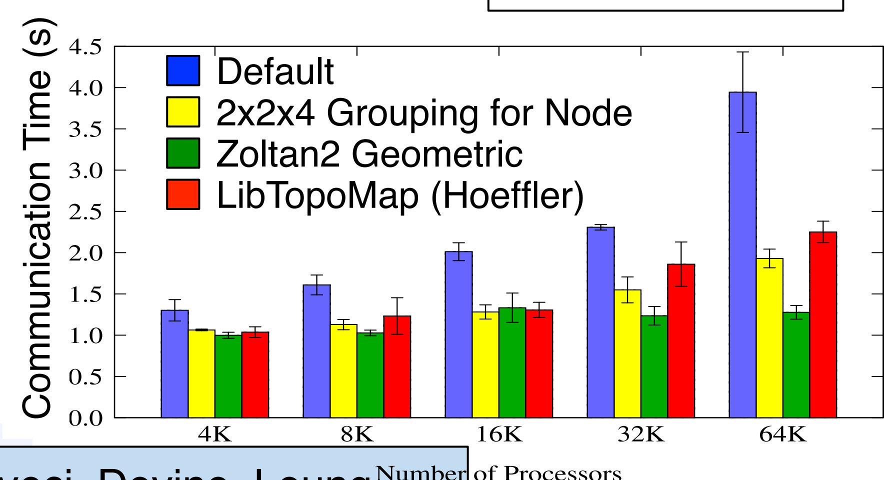

An Architecture Aware Algorithmic Advance

Contriibutors: Deveci, Devine, Leung, Prokopenko, Rajamanickam (SNL)

Our work on algebraic systems provides key solution technologies to applications

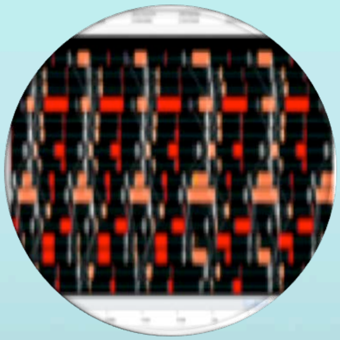

Linear system solution using direct and iterative solvers

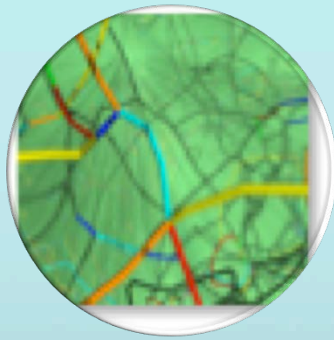

Nonlinear system solution using acceleration techniques and globalized Newton methods

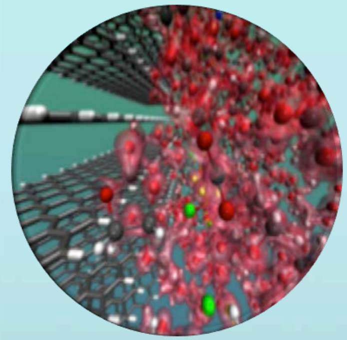

Eigensolvers using iterative techniques and optimization

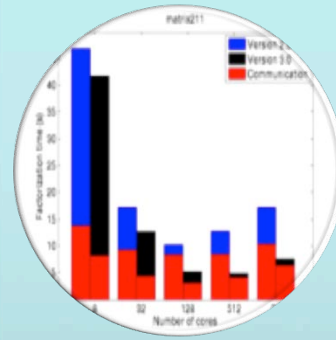

Architecture aware implementations

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

We are re-evaluating 'old' algorithms made new again for multigrid methods on extreme scale computers

- **Goal in 2011**: Shift to data centric complexity – reevaluate segmental refinement (SR) algorithms
  - Data locality: loop fusion increases arithmetic intensity
  - Inherently asynchronous: fewer data dependencies
  - Each application of SR removes communication completely at some level of memory hierarchy
- **Approach**:
  - Investigate algorithmic tradeoffs (accuracy/cost) of avoiding "far communication"
  - Investigate asymptotic complexity
  - Compose multiple SR methods
  - Apply to new operators and equations
- **Accomplishment:** First published multilevel SR result (SISC)
  - Observed accuracy dependence on parameter
  - Demonstrated modest reduction in communication costs

Segmental refinement tree communication patterns

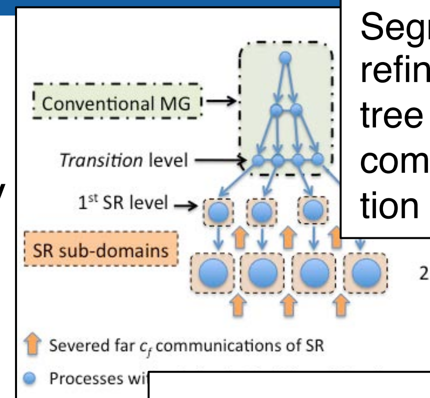

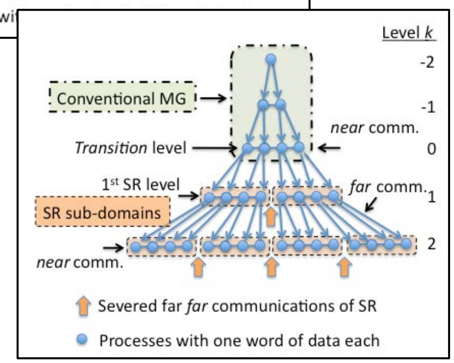

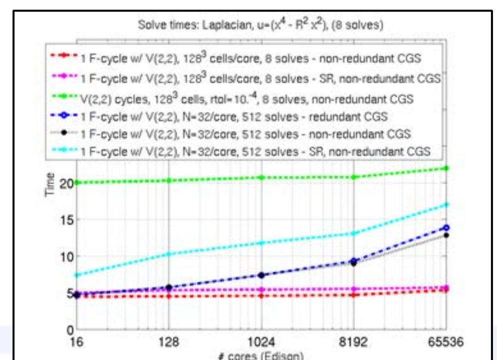

An Architecture Aware Algorithmic Advance	Contributors: M. Adams (LBNL), J. Brown, B. Smith (ANL)

SMU THE UNIVERSITY OF BRITISH COLUMBIA

Significant work has focused on reducing communication costs in algebraic multigrid methods

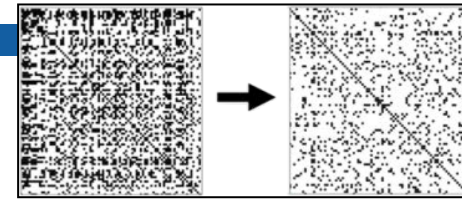

- **Goal in 2011:** Reduce communication in algebraic multigrid to improve overall scalability on next generation architectures
- **Approaches:** Use of redundancy and agglomeration for coarse grids
  - *Non-Galerkin AMG*: Replaces coarse grid operators with sparser ones
  - *Mult-Additive AMG*: Combines improved communication computation overlap of additive AMG with convergence properties of multiplicative AMG
  - *AMG-domain decomposition (AMG-DD*): Employs cheap global problems to speed up convergence

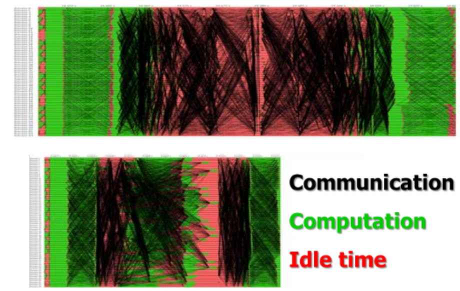

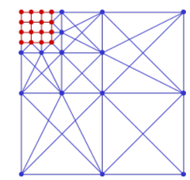

- **Accomplishment:**
  - Use of redundancy and agglomeration, non-Galerkin AMG and mult-additive AMG achieved speedups up to 2x (in some cases more) over existing AMG
  - AMG-DD has the potential for FMG convergence with only log N latency vs. (log N)²!

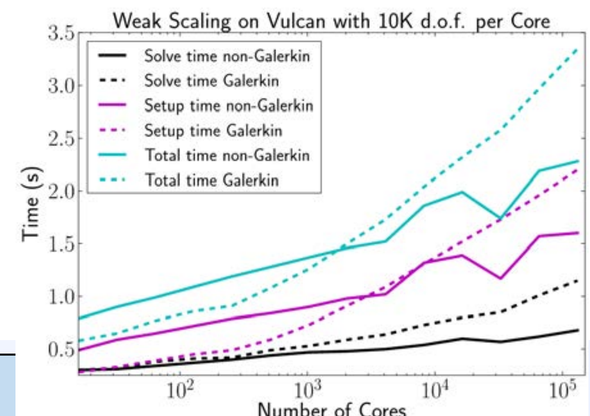

An Architecture Aware Algorithmic Advance

Contributors: R. Falgout, J. Schroder, U. Yang (LLNL)

Software and algorithmic improvements improved direct solver time many core architectures

- **Goal in 2011:** Restructure existing sparse direct solver codes to be architecture-friendly and develop new algorithms and codes that are nearly O(n) scaling
- **Approach:**
  - New DAG-based scheduling to shorten critical path
  - Offload fine-grained tasks to GPU or MIC accelerators using MPI + OpenMP + CUDA; pipeline execution of CPU and GPU tasks
  - Low-rank representation of well-separated blocks reduce flops and comm.
- **Accomplishments**:
  - SuperLU: Process idle time significantly reduced; 3x faster on multi-GPU or multi-Xeon Phi clusters; 2-5x reduction in memory usage.
  - New code STRUMPACK (STRUctured Matrix PACKage): O(n log n) flops for many 3D PDEs; Up to 5.4x faster than dense LU for BEM matrices; up to 7x faster than traditional sparse MF solver for model PDEs.

10 sparse matrices
3 solvers
12-core Intel Ivy Bridge
- tdr190k – ComPASS SciDAC (HEP)
- A22 – CEMM SciDAC (FES)

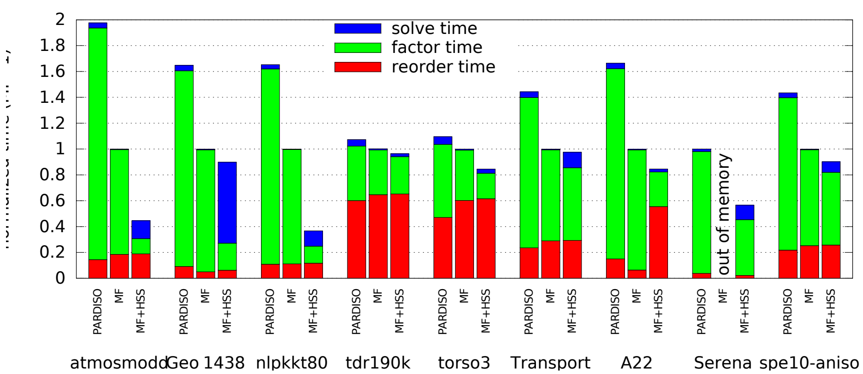

An Architecture Aware Algorithmic Advance

Contributors: X. Li (LBNL)

FASTMath has developed new software for flexible time integrators that support advanced application use

- **Goal in 2011:** construct new time integration package based on additive Runge-Kutta (ARK) methods as a component of the SUNDIALS solver suite.
- **Approach:** Construct ARKode in SUNDIALS
  - Additive methods enable explicit, implicit or ImEx integration
  - Single-step (multi-stage) approach allows for spatial adaptivity between steps.
  - High-order embeddings enable robust error control and time step adaptivity.
  - Emulate the user interface for CVODE (one of the most widely-used and highly efficient time integrators available), to enable rapid testing and comparison between solver approaches.
- **Accomplishment:**
  - Beta release 2013; official SUNDIALS release March 2015
  - Incorporated into ParaDiS, a large-scale dislocation dynamics simulator requiring extreme spatial adaptivity
  - Currently incorporating into BOUT++ (fusion)
  - Rapidly-expanding user base.

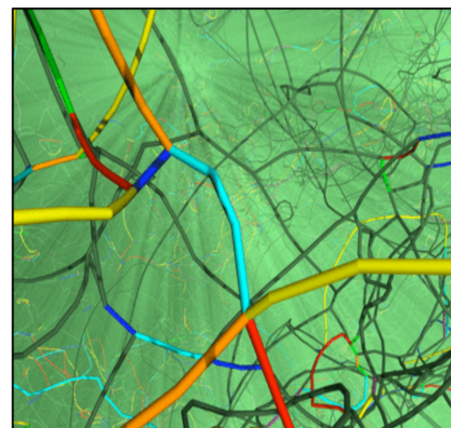

An Application Aware Algorithmic & SW Advance

Contributors: D. Reynolds (SMU), C. Woodward (LLNL)

Use of advanced nonlinear solvers and time integrators to significantly improve dislocation dynamics simulation

- **Goal:** Enable larger and more accurate time steps and solutions at reduced computational cost
  - Current approach: second order Trapezoid integrator and linear fixed point nonlinear solver
- **Approach:**
  - Applied Anderson acceleration to nonlinear solver
  - Tested third, fourth, and fifth order multistage integrators, found third order more efficient
- **Accomplishment:** New nonlinear solver and integrator results in
  - Faster solution times while maintaining accuracy
    - Accurate time steps orders of magnitude larger than existing methods
    - Efficiency gains: order of magnitude on simple tests to factors of 2.5 or more for challenging simulations

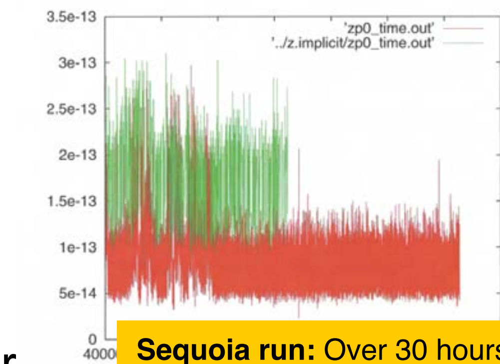

Sequoia run: Over 30 hours of wall clock time, KINSOL solver with accelerated fixed point (green) resulted in substantially larger time steps than original solver (red). Always showed a run time advantage

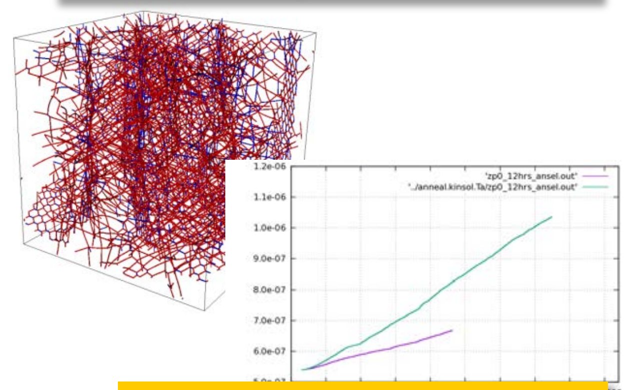

Annealing problem using KINSOL with accelerated solver results in ~50% increase in simulated time over 12 wall clock hours

An Application Advance	Contributors: C. Woodward, D. Gardner (LLNL), D. Reynolds (SMU)

We have developed efficient eigengsolvers for material science and chemistry applications

- **Goal:** Develop new algorithms that overcome the scalability issues, address convergence issues, and preserve and leverage special problem structure
- **Approach:**
  - Project precondtioned congugate gradient (PPCG) algorithm reduces number of Rayleigh Ritz calculations, solves *k* 3 x 3 eigenvalue problems instead one big 3k x 3k one each iteration
  - Generalized preconditioned locally harmonic Ritz (GPLHR) method optimizes subspace generation and projection
  - New eigensolver preserves eigenvalue pairings for Bethe-Salpeter (BSE) problem
  - New eigensolver for time-dependent DFT
- **Accomplishment:**
  - 2x speedup of Quantum Esspresso and Qbox eigensolver for large systems
  - More robust GPLHR eigensolver in Qchem
  - Faster BSE solver in BerkeleyGW and TDDFT eigensolver in NWChem

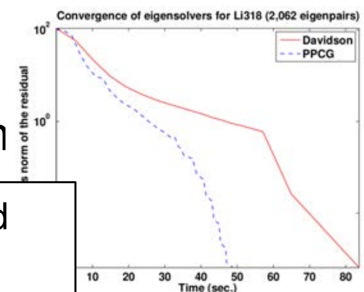

DFT band structure calculation

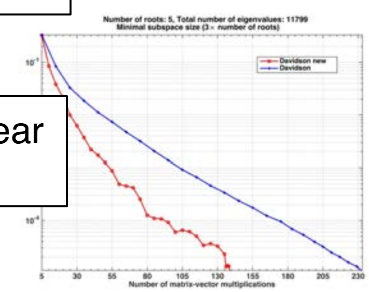

TDDFT linear response

An Application and Architecture Aware Algorithmic Advance

Contributors: E. Vecharynski, M. Shao, C. Yang, E. Ng (LBNL)

SMU

UBC THE UNIVERSITY OF BRITISH COLUMBIA

Integrating technologies is a key value added by the FASTMath Institute

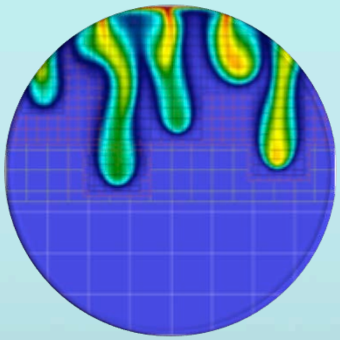

Mesh/solver interactions

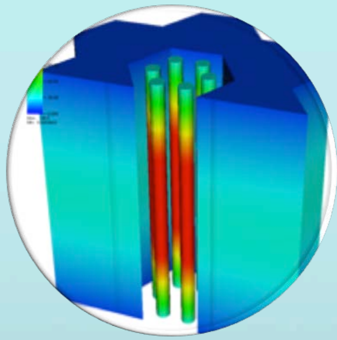

Mesh-to-mesh coupling methods

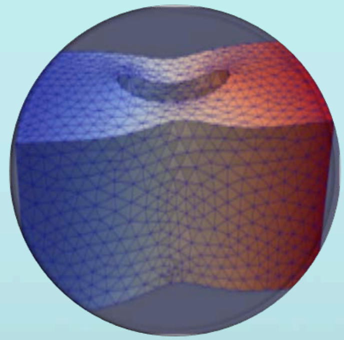

Unstructured mesh technologies into simulation workflows

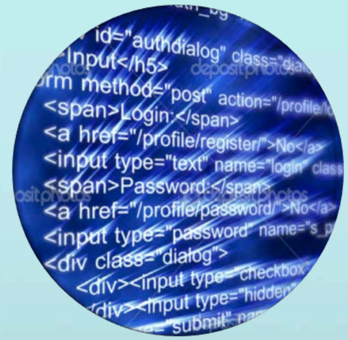

Software unification strategies

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

We have demonstrated extreme scale computations are possible with an integrated suite of FASTMath Tools

- **Goal in 2011:** Provide an end-to-end unstructured grid solver workflow that matches the complex needs of DOE PDE-based applications
  - Provide an in-team, beta driver/tester of integrated technologies
  - Provide basis for a new suite of DOE applications at extreme scale
- **Approach:**
  - Use PHASTA as baseline software
  - Integrate PETSc solver to expand preconditioner options
  - Integrate MeshAdapt to adapt anywhere, including anisotropic boundary layer
  - Integrate PUMI and ParMa for entire workflow scaling
- **Accomplishment:**
  - 100 fold improvement to PETSc global assembly time
  - PHASTA workflow scales to 92B elements on 3M processes

767 vertical tale simulation

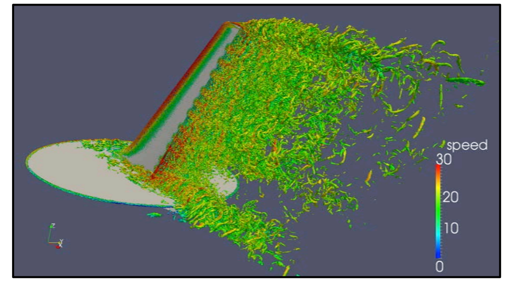

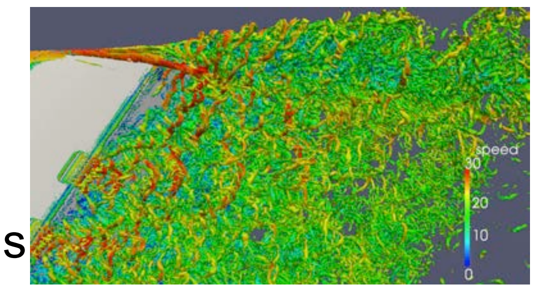

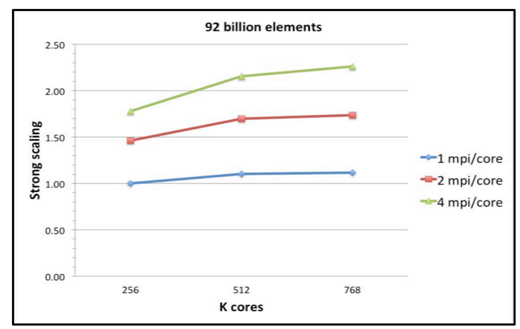

Achitecture Aware Integrated Technology Advance	Contributors: K.E. Jansen, J. Brown, M. Rasquin, B. Matthews (Colorado), M.S. Shephard, O Sahni, D. Ibanez, C.W. Smith (RPI)

Developed a new adjoint-based inversion capability in Albany for ice sheet initialization problems

Goal in 2011: Develop a new method for initializing model parameters to match observations for PISCEES Application

Approach: General-purpose implementation in Albany using Trilinos libraries:

- automatic differentiation, preconditioners, optimization

Impact:

→ Initialization is on critical path of ACME climate science runs.

→ Capability already used in an NNSA design application.

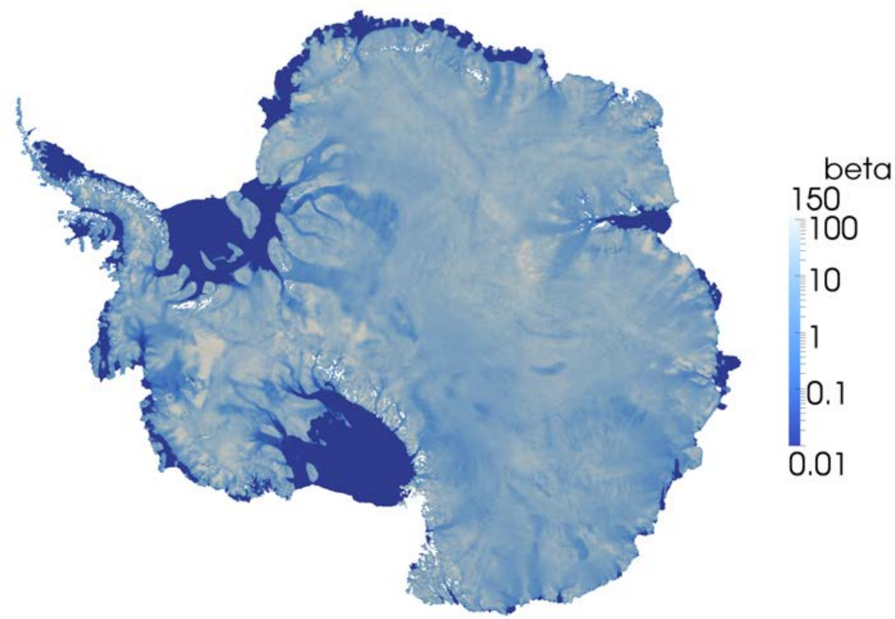

Inversion result for Basal friction coefficient (kPa yr/m) at each of 700K nodes of Antarctica surface mesh.

An Application Aware Integrated Technology Advance

Contributors: M. Perego, A. Salingler (SNL)

FASTMath posters provide more details

1. **hypre: High Performance Preconditioners**
   POC: U. Yang, LLNL
2. **MueLu: Multigrid Framework for Advanced Architectures**
   POC: Andrey Prokopenko, SNL
3. **Parallel Unstructured Mesh Infrastructure**
   POC: Mark Shephard, RPI
4. **Dynamic Partitioning using Mesh Adjacencies**
   POC: Mark Shephard, RPI
5. **Massively Parallel Adaptive Simulations using PETSc for Turbulent Boundary Layer Flows**
   POC: Mark Shephard, RPI
6. **Construction of Parallel Adaptive Simulation Loops**
   POC: Mark Shephard, RPI
7. **Albany: Integrating Algorithmic Components to Build Advanced Applications**
   POC: Andy Salinger, SNL
8. **The Zoltan2 Toolkit: Partitioning, Task Placement, Coloring, and Ordering**
   POC: Karen Devine, SNL
9. **FASTMath Structured Mesh and Particle Technologies**
   POC: Anshu Dubey, LBNL
10. **Tiling in BoxLib: Implementation and Performance of Logical Tiling in an AMR**
    POC: Ann Almgren, LBNL
11. **Sparse direct solvers on distributed CPU-GPU machines**
    POC: Sherry Li, LBNL
12. **SUNDIALS: Suite of Nonlinear and DIfferential/Algebraic Solvers**
    POC: Carol Woodward, LLNL
13. **SIGMA: Scalable Interfaces for Geometry and Mesh based Applications**
    POC: Vijay Mahadeven
14. **Large-scale Eigensolvers for SciDAC Applications**
    POC: Chao Yang, LBNL

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

The SciDAC Program has been critical for the support and development of key comp math libraries

- **New Algorithms**
  - Application-aware and Architecture-aware
  - Across all FASTMath topical areas

- **More Robust Software**
  - Enabling transition to many-/multi-core architectures
  - Incorporation of mixed programming models
  - Scalability
  - Tighter integration and unification of FASTMath packages

- **Application Partnerships**: FASTMath supports general infrastructure needed by SciDAC applications across SC and NNSA

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

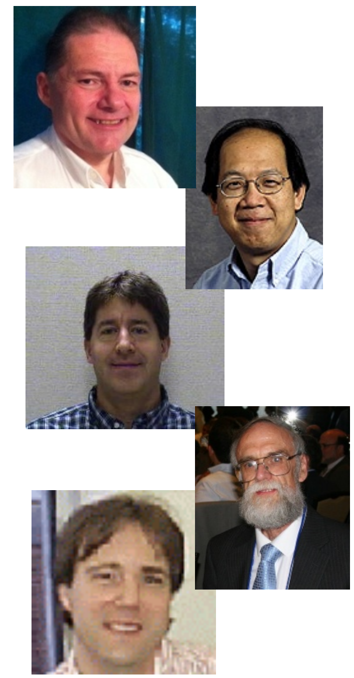

http://www.fastmath-scidac.org