Consequences for scalability arising from multi-material modeling

Allen C. Robinson

Jay Mosso, Chris Siefert, Jonathan Hu

Sandia National Laboratories

Tom Gardiner, Cray Inc.

Joe Crepeau, Applied Research Associates, Inc.

Numerical methods for multi-material fluid flows

Czech Technical University, Prague, Czech Republic

September 10 - 14, 2007

Sandia is a multiprogram laboratory operated by Sandia Corporation, a Lockheed Martin Company, for the United States Department of Energy's National Nuclear Security Administration under Contract DE-AC04-94AL84000.

The ALEGRA-HEDP mission: predictive design and analysis capability to effectively use the *Z* machine

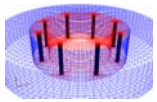

Meshing

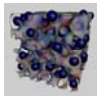

Material models

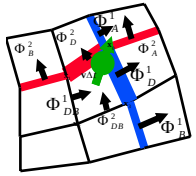

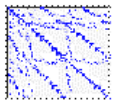

Algorithms

IMC (Joint LLNL/SNL development)

Optimization/UQ

Platforms

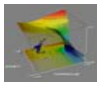

Analysis

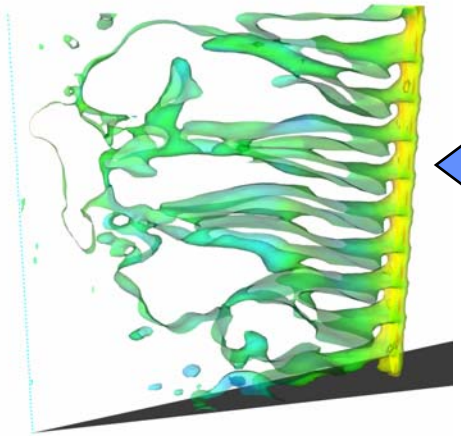

ALEGRA-HEDP

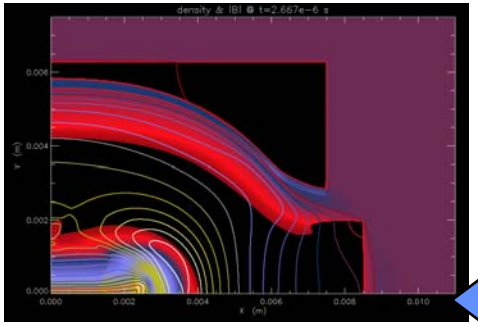

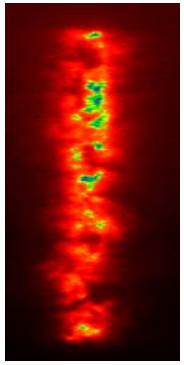

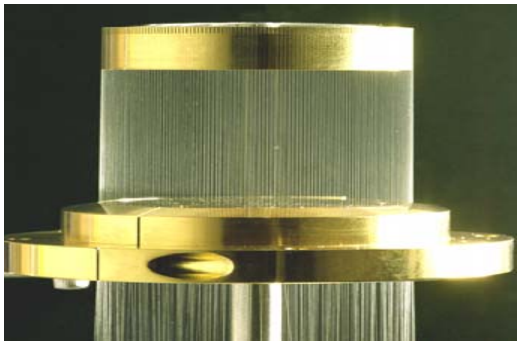

Z-pinches as x-ray sources

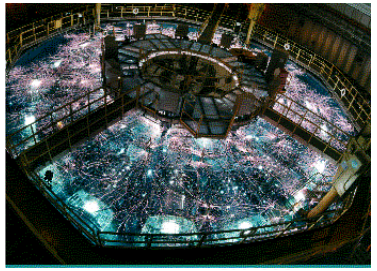

Magnetic flyers for EOS

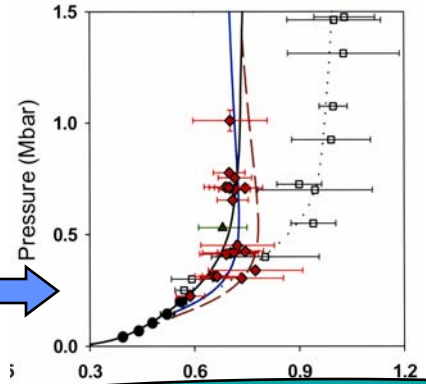

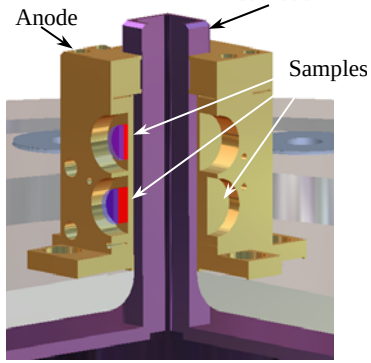

Computer and Information Sciences

Pulsed Power Sciences

Multimaterial and Multi-physics Modeling in arbitrary mesh ALE codes

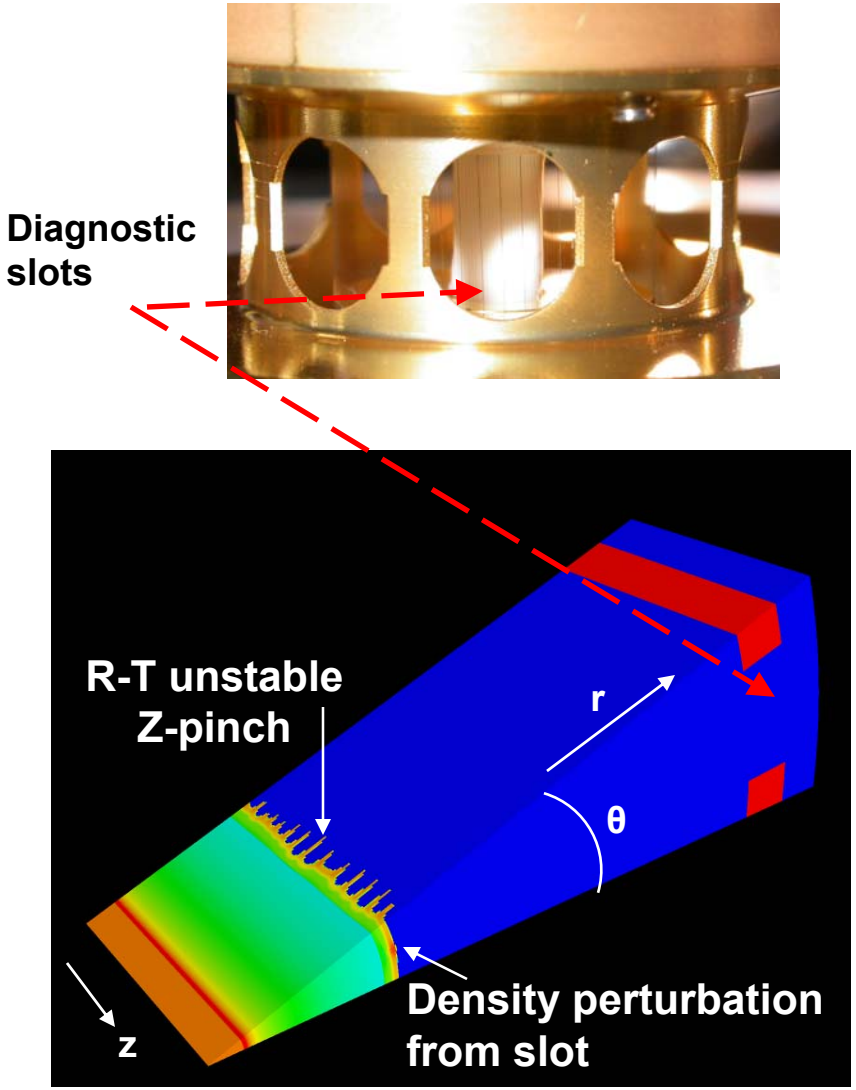

- Complicated geometries and many materials are a fact of life for realistic simulations.
- Future machines may be less tolerant of load imbalances.
- Multimaterial issues play a key role with respect to algorithmic performance. For example,
  - Interface Reconstruction
  - Implicit solver performance.
  - Material models
- What processes are required to confront and solve performance and load balancing issues in a timely manner?

What do current/future machines look like?

Representative largest platforms:

Purple:

Compute Nodes: 1336 nodes * 8 sockets/node = 12888

CPU (core) speed IBM Power5 (GHz) 1.9

Theoretical system peak performance = 93.39 TFlop/s

Red Storm:

Compute Nodes: 12960 sockets * 2 cores/socket = 25920

CPU speed AMD Opteron (GHz) 2.4

Theoretical system peak performance = 124 Tflop/s

What do future machines look like?

Representative largest platform in 5 years: (likely)

- 10+ Petaflops
- 40,000 sockets * 25 cores/socket = 1 Million cores
- .5 Gbyte/core

Representative largest platforms in 10 years: (crystal ball)

- Exaflops
- 100 Million cores.

Sounds great: but

- Memory bandwidth is clearly at serious risk
- Can latency and cross-sectional bandwidth keep up?

Minor software/algorithmic/process flaws today may be near fatal weaknesses tomorrow from both a scalability and robustness point of view.

ALEGRA Scalability Testing Process

Define sequences of gradually more complicated problems in a software environment that easily generates large scale scalability tests. (python/xml)

Budget/assign personnel and computer time to exercise these tests on a regular basis.

Take action as required to minimize impact of problematic results achieved on large scale systems.

Available Interface Reconstruction Options in ALEGRA

SLIC – Single Line Interface Reconstruction

SMYRA – Sandia Modified Young's Reconstruction
works with a unit cube description

New Smyra – Alternate version of SMYRA algorithm

PIR – Patterned Interface Reconstruction

- **Works with physical element description (not unit cubes)**
- **Additional smoothing steps yields second order accuracy**
- **Strict ordering and polygonal removal by material guarantees self-consistent geometry.**
- **More expensive**

Interface reconstruction is not needed for single material.

Problem Description

AdvectBlock – Single material simple advection

InterfaceTrack – 6 material advection problem in a periodic box with spheres and hemispheres

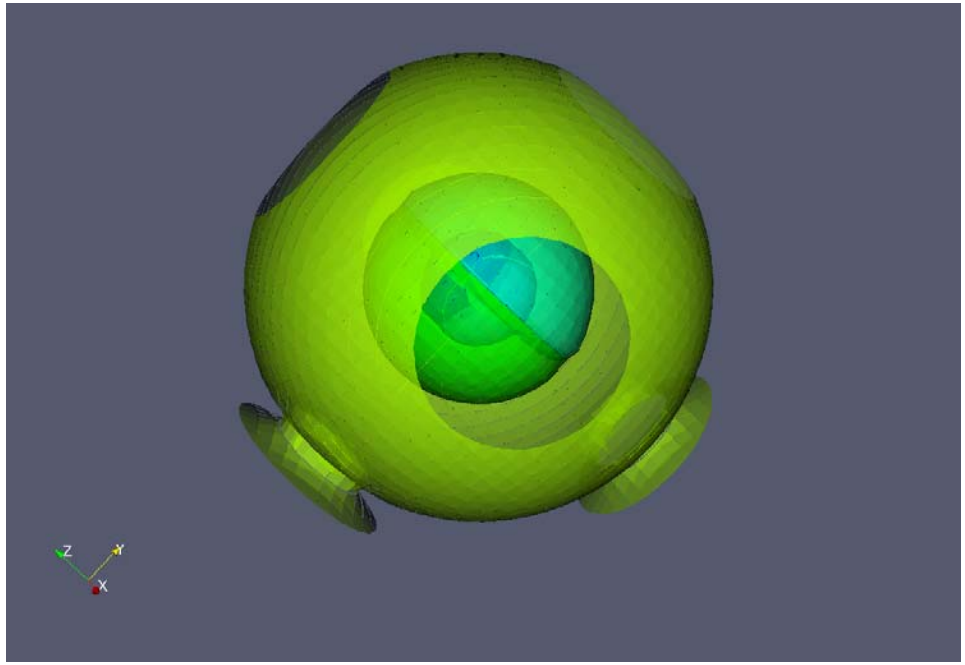

Large scale testing smokes out error (1/12/2007)

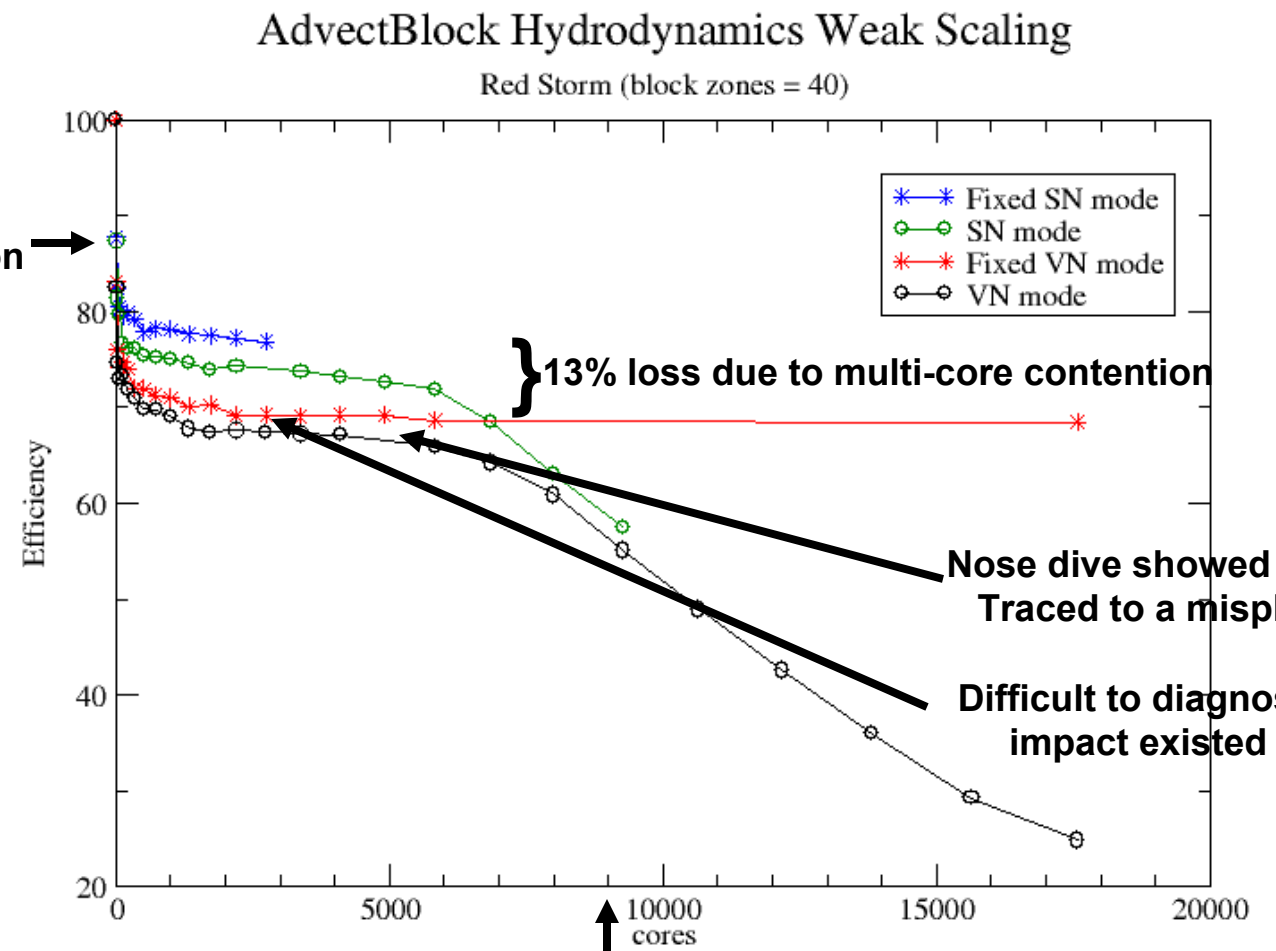

Parallel communication overhead

13% loss due to multi-core contention

SN = 1 core/node

VN= 2 cores/node

Nose dive showed up at 6000 cores
Traced to a misplaced all to one

Difficult to diagnosis performance impact existed at small scale

Before this fix was found Purple results showed similar results then suddenly dropped to 5% at this point.

Interface Track (6/22/2007)

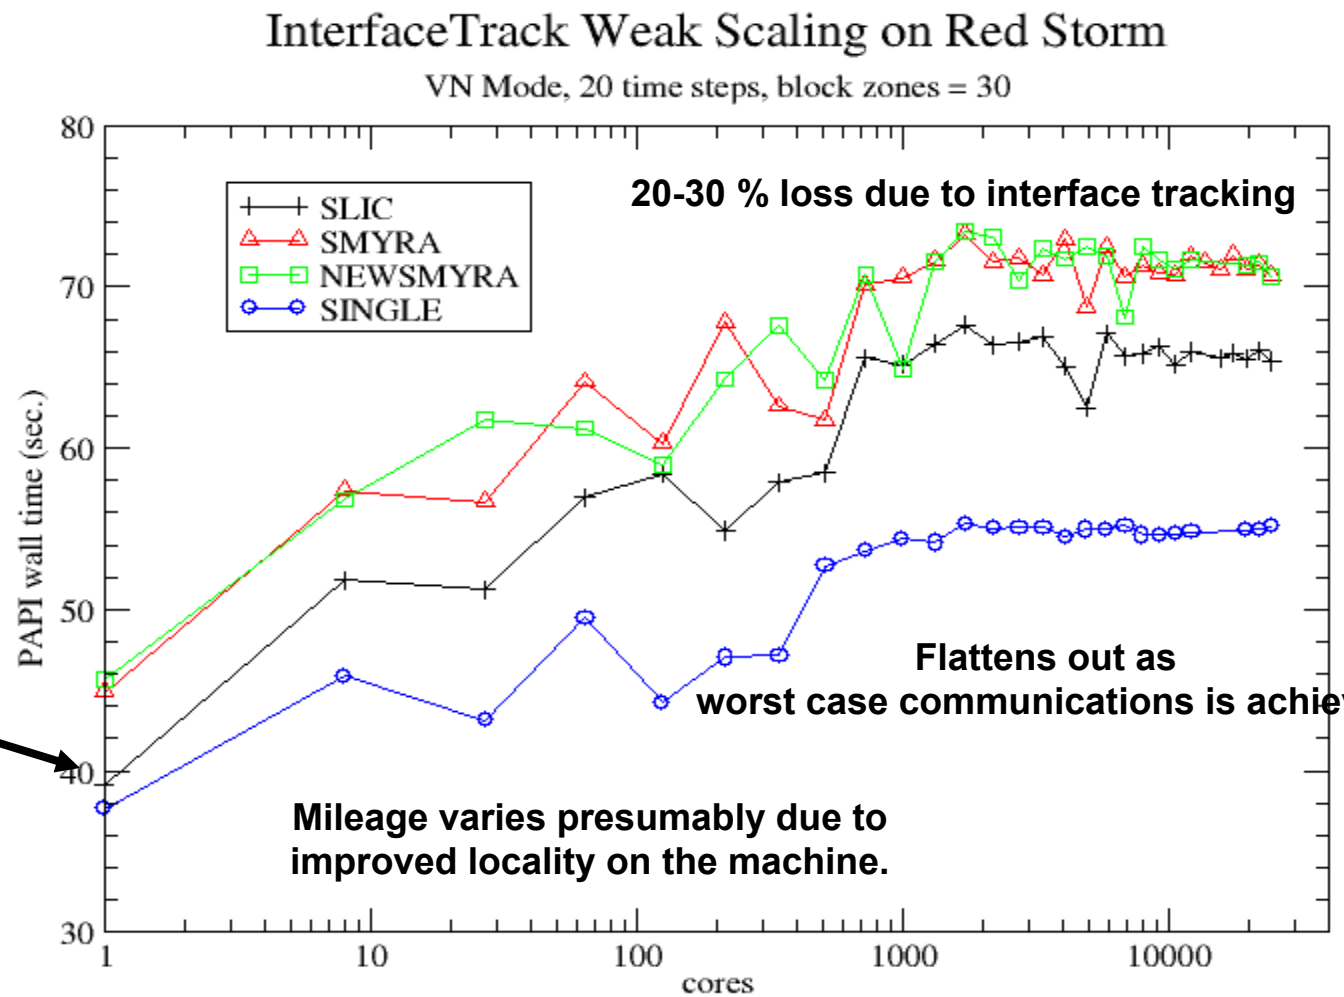

20-30 % loss due to interface tracking

Periodic bc is always "parallel" but no real communication occurs

Flattens out as worst case communications is achieved

Mileage varies presumably due to improved locality on the machine.

Next generation
Pattern Interface Reconstruction (PIR) Algorithm

Basic PIR is an extension of the Youngs' 3D algorithm:

- DL Youngs, *"An Interface Tracking Method for a Three-Dimensional Hydrodynamics Code"*, Technical Report 44/92/35, (AWRE 1984)
- Approximate interface normal by –Grad(Vf) $\hat{n}_0^G$
- Position planar interface (polygon) in element to conserve volume exactly for arbitrary shaped elements.
- not spatially second-order accurate

Smoothed PIR

- Planar algorithm generates a trial normal.
- Spherical algorithm generates an alternative trial normal.
- "roughness" measure determines which trial normal agrees best with the local neighborhood.

PIR Utility

- more accurately move materials through the computational mesh
- visualization

PIR Smoothing Algorithms

Smoothing uses Swartz Stability Points

- SJ Mosso, BK Swartz, DB Kothe, RC Ferrell, *"A Parallel Volume Tracking Algorithm for Unstructured Meshes"*, Parallel CFD Conference Proceedings, Capri, Italy, 1996.
- The centroid of each interface is a 'stable' position

Algorithm

- Compute the centroid of each interface
- Fit surface(s) to the neighboring centroids
- Compute the normal(s) of the fit(s)
- Choose the best normal
- Re-adjust positions to conserve volume
- Iterate to convergence.

$$Global\ Min(\hat{n}_{previous} \circ \hat{n}_{current}) > (1.0 - \varepsilon)$$

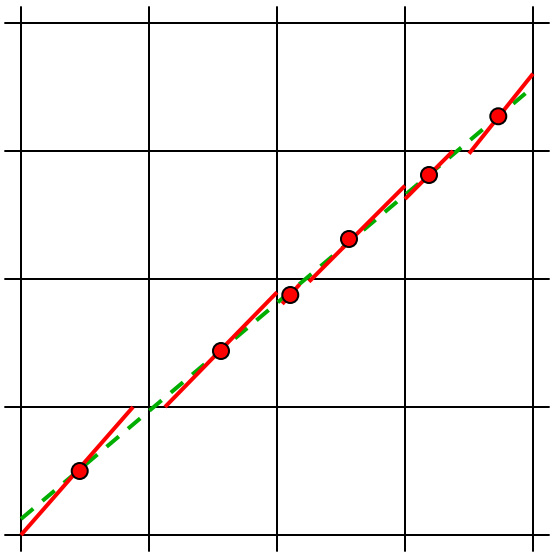

Planar Normal Algorithm

Least-Squares fit of a plane to the immediate 3D neighborhood

$$\vec{s}_i = \{x_i, y_i, z_i\}$$

$$Obj = \sum_{neighbors(i)} \left(\vec{n} \bullet \vec{s}_i\right)^2 - \lambda(\vec{n} \bullet \vec{n} - 1)$$

$$\left[\sum_{neighbors(i)} \vec{s}_i \otimes \vec{s}_i - \lambda \mathbf{I}\right] \vec{n} = 0$$

$$\begin{bmatrix} \sum x_i^2 & \sum x_i y_i & \sum x_i z_i \\ \sum x_i y_i & \sum y_i^2 & \sum y_i z_i \\ \sum x_i z_i & \sum y_i z_i & \sum z_i^2 \end{bmatrix} - \lambda I = 0$$

2 evecs ~ in plane
1 evec ~ out of plane (minimal eigenvalue)

$$\hat{n} \circ (\vec{S}_i - \vec{S}_0) = d_i$$

$$\vec{s}_i = \vec{S}_i - \vec{S}_0$$

$$\hat{n} \circ \vec{s}_i = d_i$$

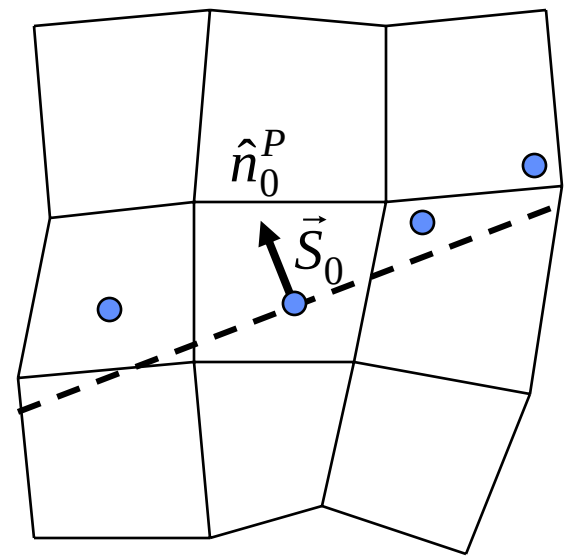

$$\hat{n}_0^P = \hat{n}(\lambda_{\min})$$

Spherical Normal Algorithm

Construct plane at midpoint of chord joining home S_0 and each neighboring S_i

Compute V closest to all midchord planes

$$\hat{c}_i = (\vec{S}_i - \vec{S}_0) / |\vec{S}_i - \vec{S}_0| = \vec{s}_i / |\vec{s}_i|$$

$$Obj = \sum_{neighbors(i)} (\hat{c}_i \circ (\vec{V} - \frac{1}{2}\vec{s}_i))^2$$

$$\frac{\partial Obj}{\partial V_x} = \frac{\partial Obj}{\partial V_y} = \frac{\partial Obj}{\partial V_z} = 0$$

$$\begin{bmatrix} \sum c_{x_i}{}^2 & \sum c_{x_i} c_{y_i} & \sum c_{x_i} c_{z_i} \\ \sum c_{x_i} c_{y_i} & \sum c_{y_i}{}^2 & \sum c_{y_i} c_{z_i} \\ \sum c_{x_i} c_{z_i} & \sum c_{y_i} c_{z_i} & \sum c_{z_i}{}^2 \end{bmatrix} \begin{bmatrix} V_x \\ V_y \\ V_z \end{bmatrix} = \begin{bmatrix} \frac{1}{2}\sum c_{x_i}{}^2 s_{x_i} \\ \frac{1}{2}\sum c_{y_i}{}^2 s_{y_i} \\ \frac{1}{2}\sum c_{z_i}{}^2 s_{z_i} \end{bmatrix}$$

$$\hat{n}_0^S = \frac{\vec{S}_0 - \vec{V}}{|\vec{S}_0 - \vec{V}|}$$

Planar fit is non-optimal for curved interfaces

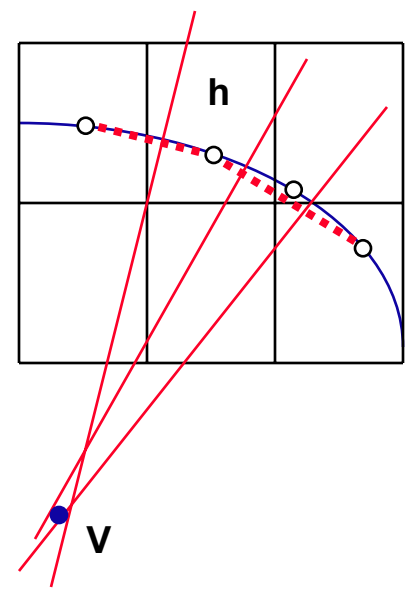

Roughness measure

Roughness is sum of a displacement volume and a relative orientation volume

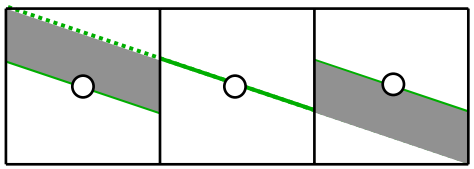

$$A_i \hat{n}_0 \circ \vec{s}_i$$

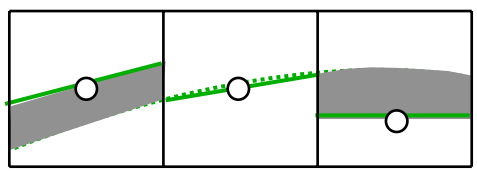

$$A_i abs(R_i - R) = A_i abs(\left|\vec{s}_i - \vec{V}\right| - \left|\vec{V}\right|)$$

○ **V**

Displacement roughness

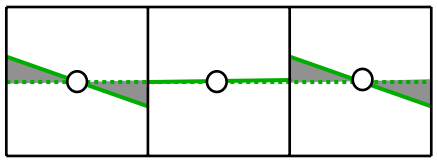

$$\frac{A_i}{2}\sqrt{A_i\left(1-(\hat{n}_0 \circ \hat{n}_i)^2\right)}$$

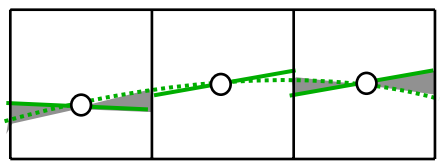

$$\frac{A_i}{2}\sqrt{A_i\left(1-\left(\frac{\vec{s}_i - \vec{V}}{\left|\vec{s}_i - \vec{V}\right|} \circ \frac{\vec{V}}{\left|\vec{V}\right|}\right)^2\right)}$$

○ **V**

Orientation roughness

Selection of 'best' normal

Three candidate normals: gradient, planar, spherical

Extrapolate shape and compute spatial s_i agreement and normal agreement: 'roughness'

Method with lowest roughness is selected

$$R_{gradient} = \sum_{neighbors(i)} A_i \left[\hat{n}_0^G \circ \vec{s}_i + \frac{1}{2}\sqrt{A_i\left(1-(\hat{n}_0^G \circ \hat{n}_i)^2\right)} \right]$$

$$R_{planar} = \sum_{neighbors(i)} A_i \left[\hat{n}_0^P \circ \vec{s}_i + \frac{1}{2}\sqrt{A_i\left(1-(\hat{n}_0^P \circ \hat{n}_i)^2\right)} \right]$$

$$R_{spherical} = \sum_{neighbors(i)} A_i \left[abs(\left|\vec{s}_i - \vec{V}\right| - \left|\vec{V}\right|) + \frac{1}{2}\sqrt{A_i\left[1-\left(\frac{\vec{s}_i - \vec{V}}{\left|\vec{s}_i - \vec{V}\right|}\right)^2\right]} \right]$$

InterfaceTrack Test Problem (modified – not periodic)

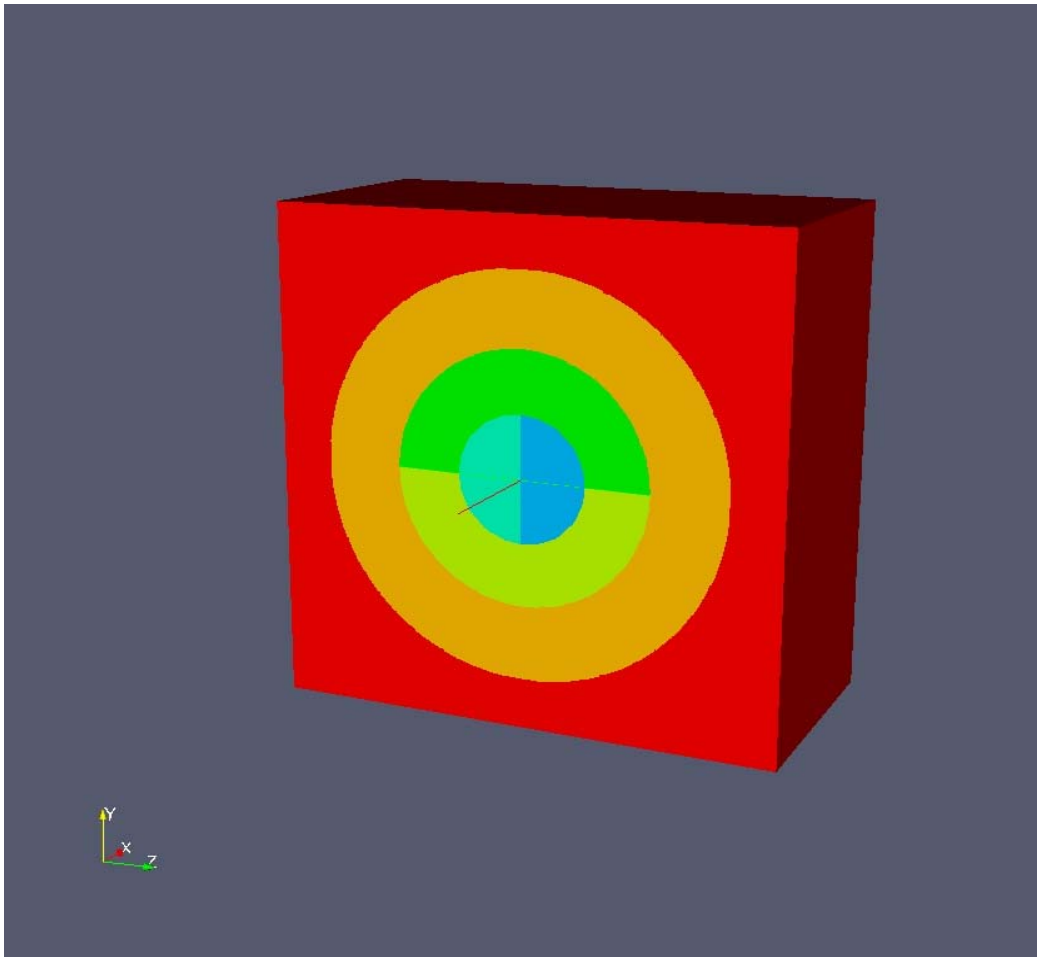

PIR Smoothing Algorithm Illustration

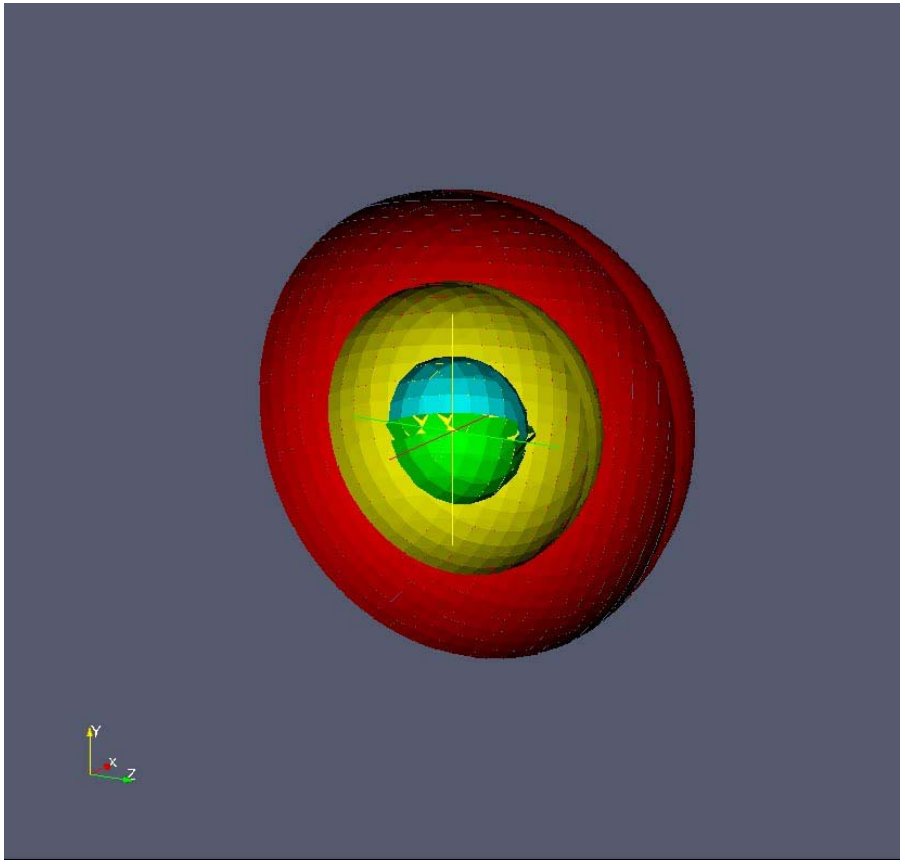

Unsmoothed

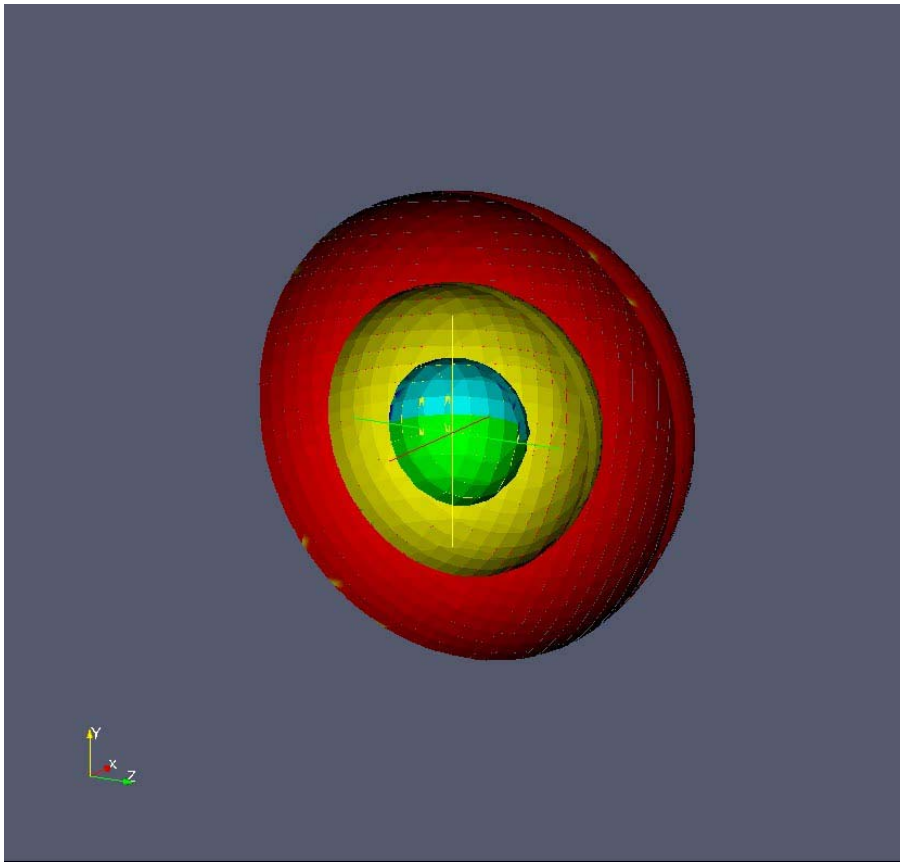

Smoothed

PIR Status

Smoothed PIR is nearing completion in both 2D and 3D as a fully functional feature in ALEGRA.

The method significantly reduces the numeric distortion of the shape of the body, as it moves through the mesh

Increased fidelity comes at cost. ~50% more floating point operations but ~10x cost.

Comparison of non-smoothed PIR with other options

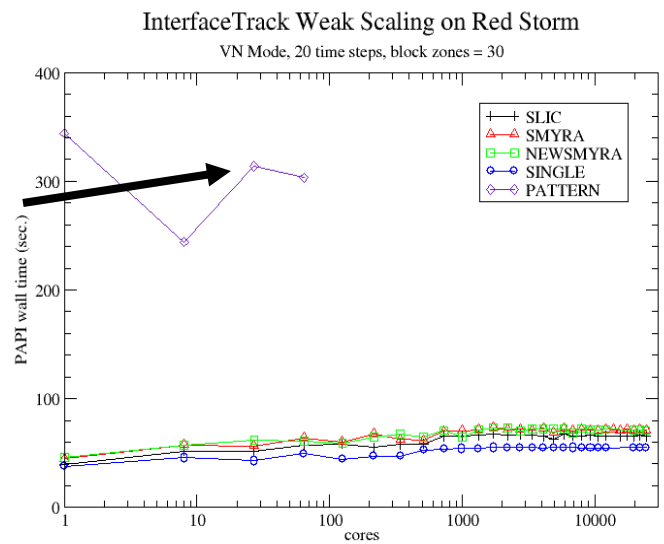

Why? Non-optimized code. Using tools such as valgrind with cachegrind we expect rapid improvements. Example: one line modification to STL::vector usage already resulted in 32% improvement in this algorithm!

Eddy Current Equations

$$\nabla \times \frac{1}{\mu}\nabla \times \mathbf{E} + \sigma \mathbf{E} = 0 \quad \text{in } \Omega$$
$$\mathbf{n} \times \mathbf{E} = 0 \quad \text{on } \Gamma$$
$$\mathbf{n} \times \frac{1}{\mu}\nabla \times \mathbf{E} = 0 \quad \text{on } \Gamma^*$$

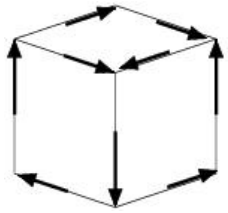

Model for many E&M phenomena.

Sandia interest: Z-Pinch. 3D magnetic diffusion step in Lagrangian operator split.

Challenge: Large null space of curl.

Solution: Compatible (edge) discretization.

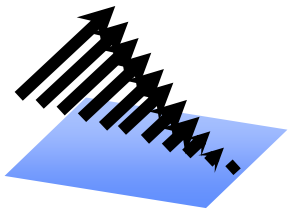

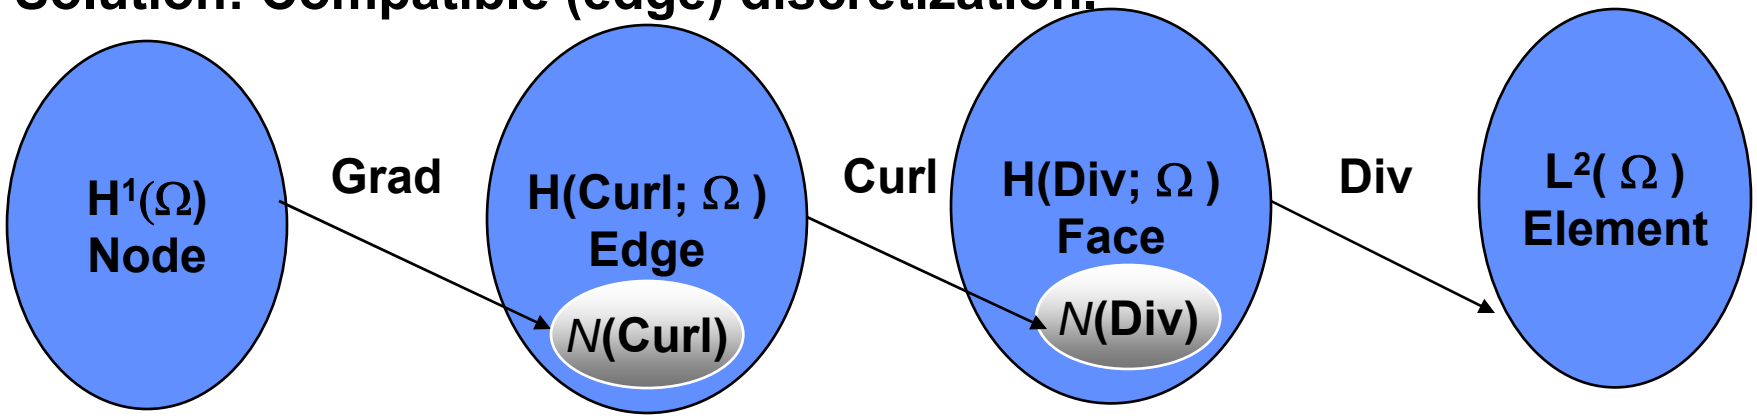

Algebraic Multigrid Solvers

Setup
- **Coarsen**
- **Project**
- **Recurse**

Each grid solves "smooth" modes on that grid.

P=prolongator

P^T=restriction

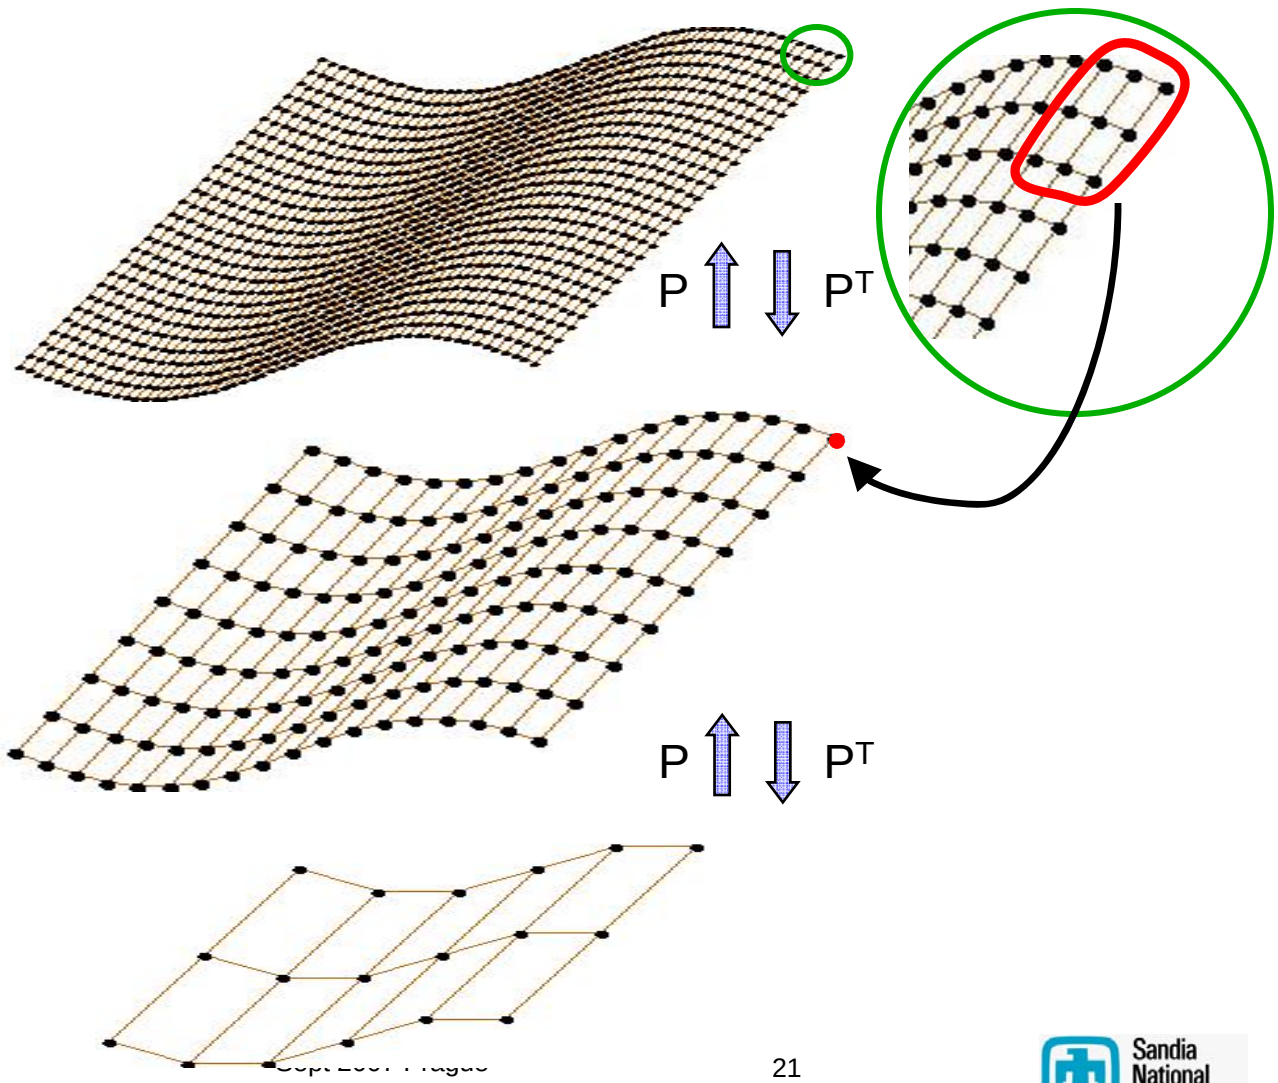

H(curl) Multigrid

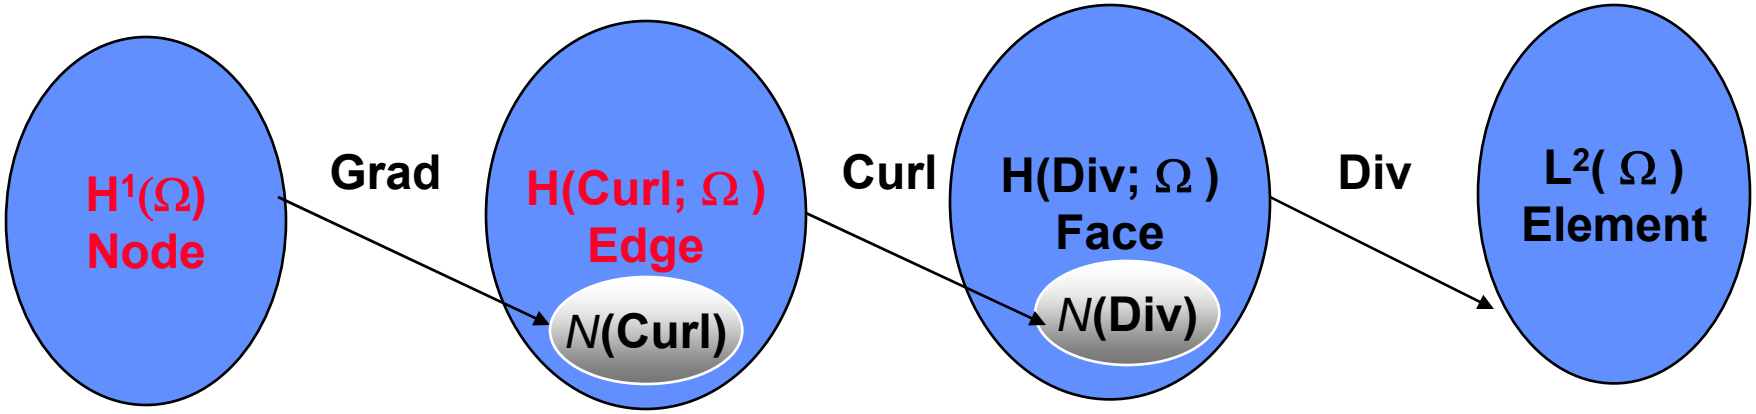

Operates on two grids: nodes and edges.

We have developed two H(curl) AMG solvers

- **Special (commuting) prolongator (Hu, et al., 2006)**
- **Discrete Hodge Laplacian reformulation (Bochev, et al., 2007, in review).**

New AMG: Laplace Reformulation

Idea: Reformulate to Hodge Laplacian:

$$\nabla \times \nabla \times - \nabla \nabla \cdot = -\Delta$$

Use a discrete Hodge decomposition:

$$e = a + \nabla \phi, \text{ where } \nabla \cdot a = 0$$

Resulting preconditioner looks like:

$$\begin{bmatrix} \nabla \times \nabla \times + \nabla \nabla \cdot + \sigma \mathbf{I} & \sigma \nabla \\ \nabla \cdot \sigma & \nabla \cdot \sigma \nabla \end{bmatrix} \begin{bmatrix} a \\ \phi \end{bmatrix} = \begin{bmatrix} f \\ \nabla \cdot f \end{bmatrix}$$

Hodge part interpolated to vector nodal Laplacian

Then apply standard AMG algorithms to each diagonal block

"Multigrid was designed for Laplacians."

Theory: Multigrid & Multimaterial

Recent work by Xu and Zhu (2007) for Laplace is encouraging.

Idea: Material jumps have limited effect on AMG.

- **Only a small number of eigenvalues get perturbed.**
- **The reduced condition number is $O(|\log h|^2)$ without these eigenvalues.**

Caveats:

- **Theory is only for Laplace (not Maxwell).**
- **Assumes number of materials is small.**
- **If we really have varying properties "which we do in real problems" then more bad EVs…**

Test Problems (10^6 jump in conductivity)

Sphere: ball in a box.

- Half-filled elements near surface.

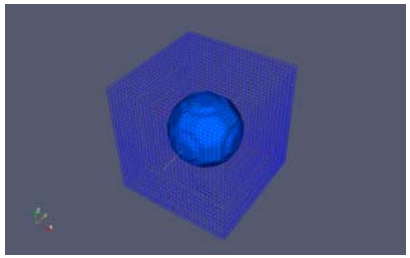

Liner: cylindrical liner.

- Non-orthogonal mesh, slight stretching.

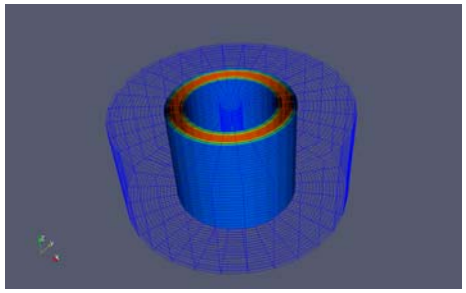

LinerF: fingered cylindrical liner.

- Non-orthogonal, slight mesh stretching.
- Material fingering.

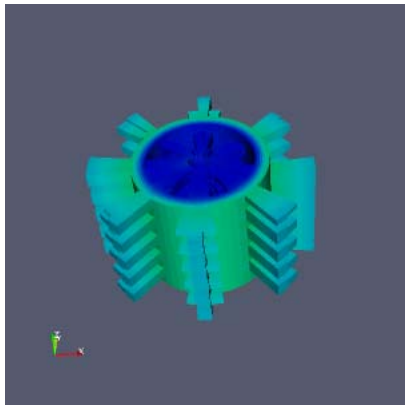

Weak scaling tests

Multimaterial Issues & Scalability

Basic Issue: coefficient (σ) changes.

Physics & discretization issues.

- **Multimaterial & mesh stretching.**
- **Material fingering.**
- **Half-filled elements at material boundaries.**

Multigrid issues.

- **Aggregates crossing material boundaries.**
- **What is an appropriate semi-coarsening?**
- **H(grad) theory not directly applicable.**

Old H(curl) Iterations (7/9/2007)

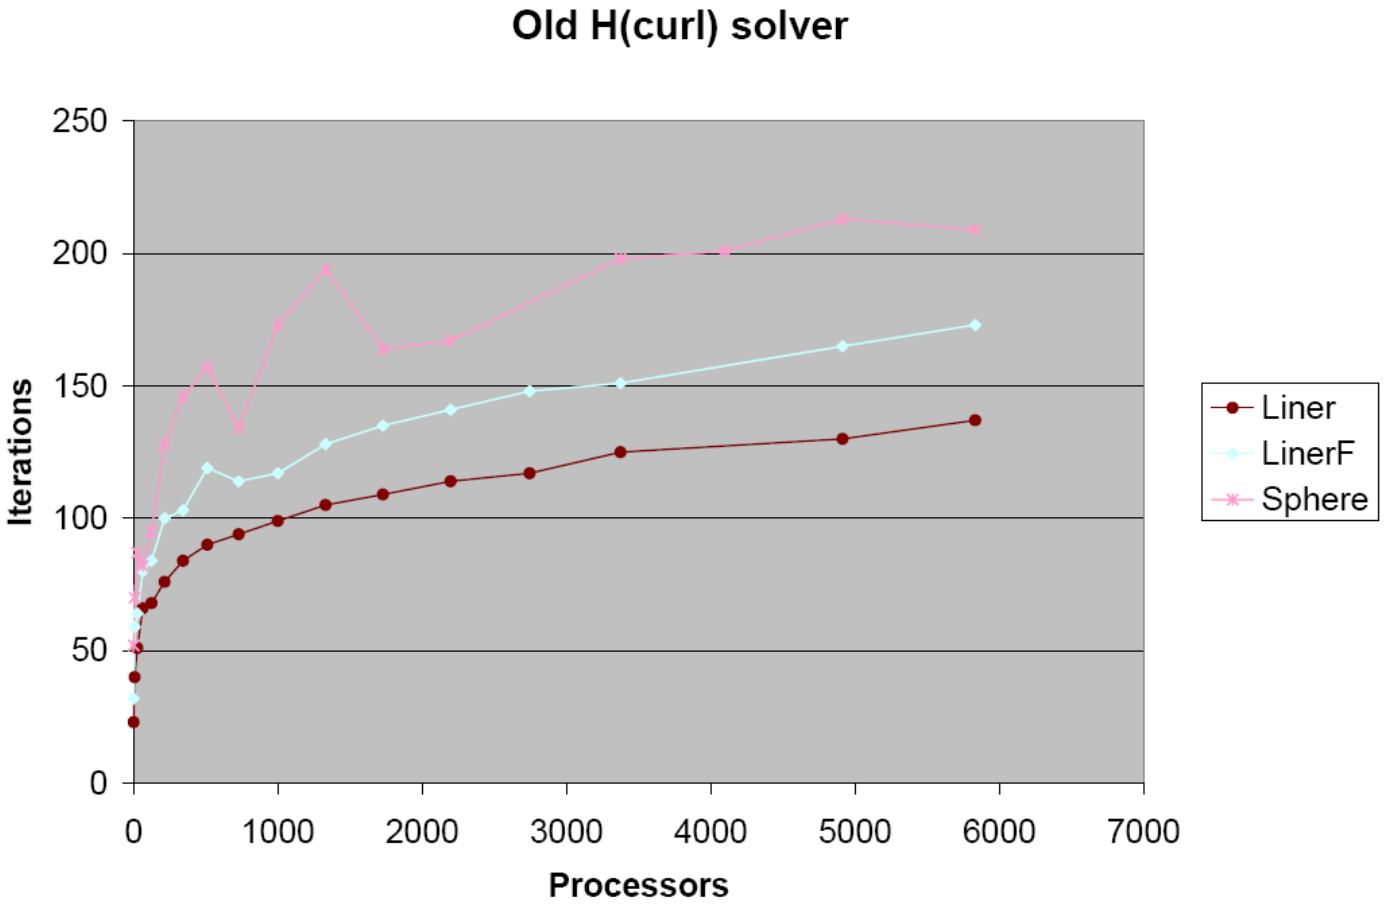

Liner and Liner F – 1 Hiptmair fine smooth, LU coarse grid, smooth prolongator

Sphere -2 Hiptmair fine sweeps, 6 coarse Hiptmair, smooth prolongator off

Performance sensitive to solver settings and problem.

Old H(curl) Run Time (7/9/2007)

Liner and Liner F – 1 Hiptmair fine smooth, LU coarse grid, smooth prolongator
Note degradation due to fingering
Sphere -2 Hiptmair fine sweeps, 6 coarse Hiptmair, smooth prolongator off

Performance sensitive to solver settings and problem.

Sphere - Old/New Comparison

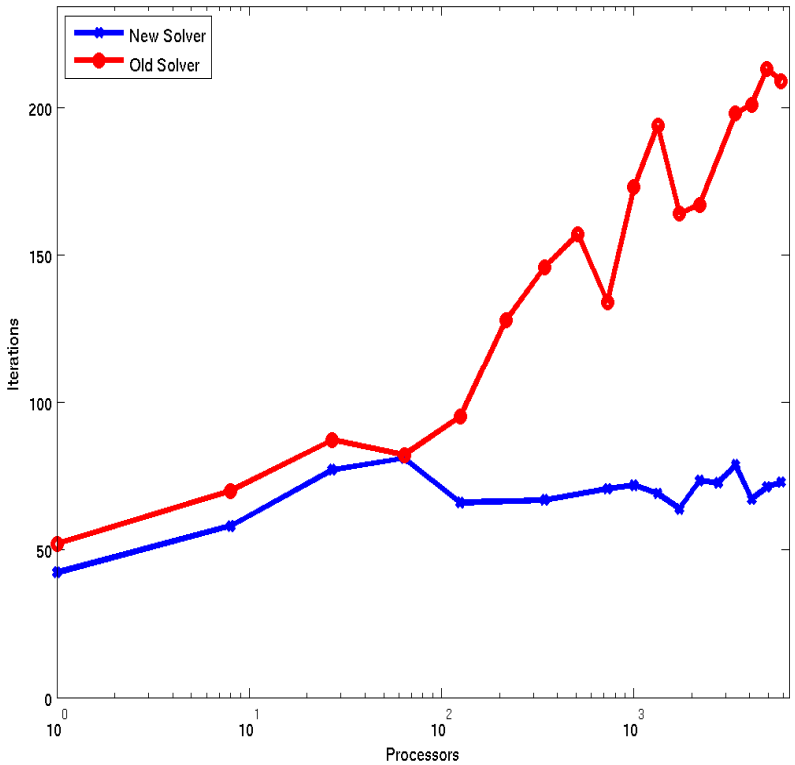

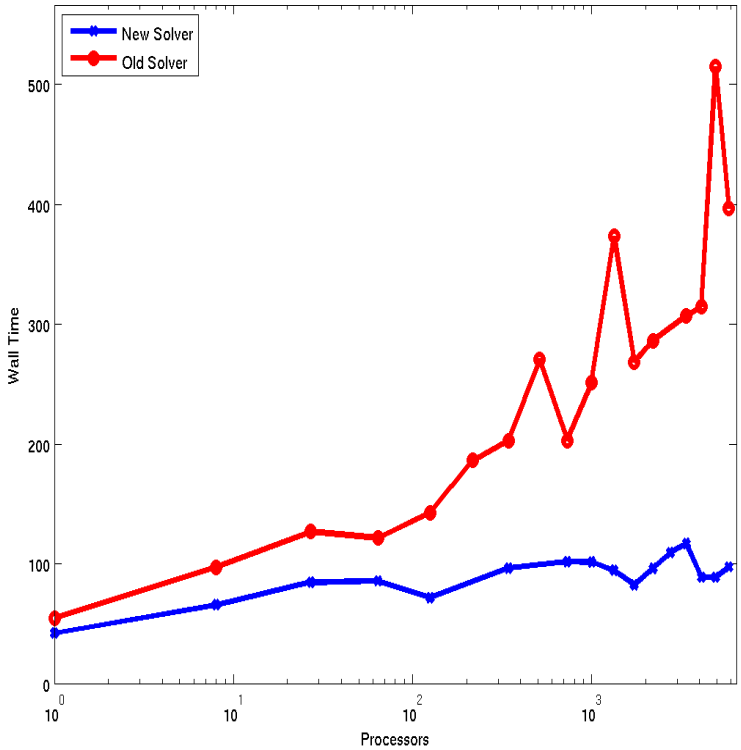

Liner - Old/New Comparison

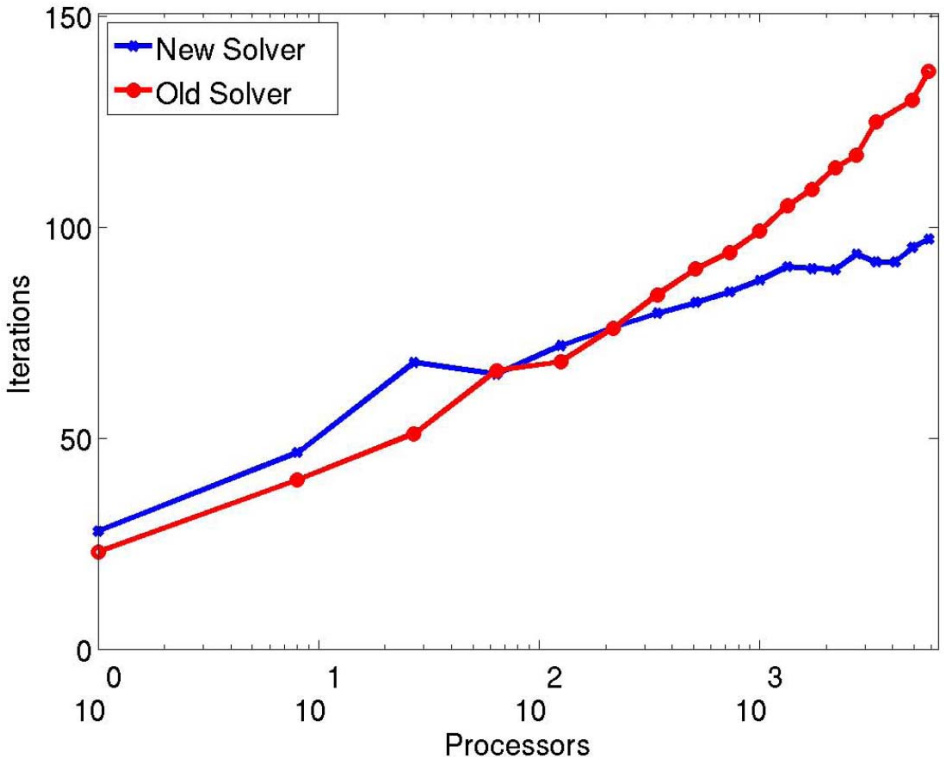

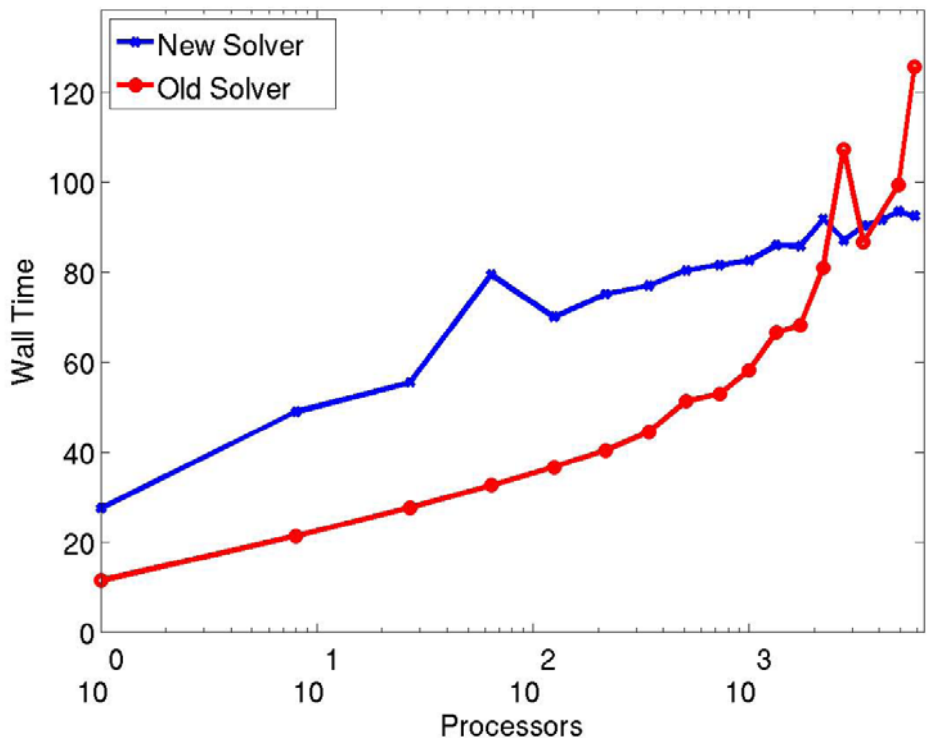

LinerF- Old/New Comparison

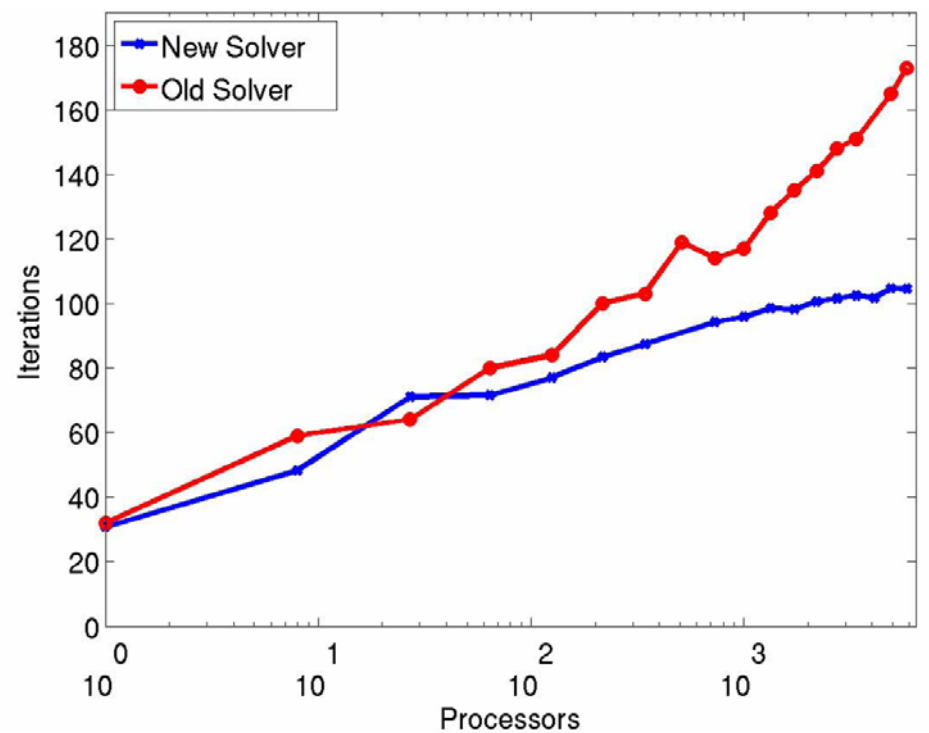

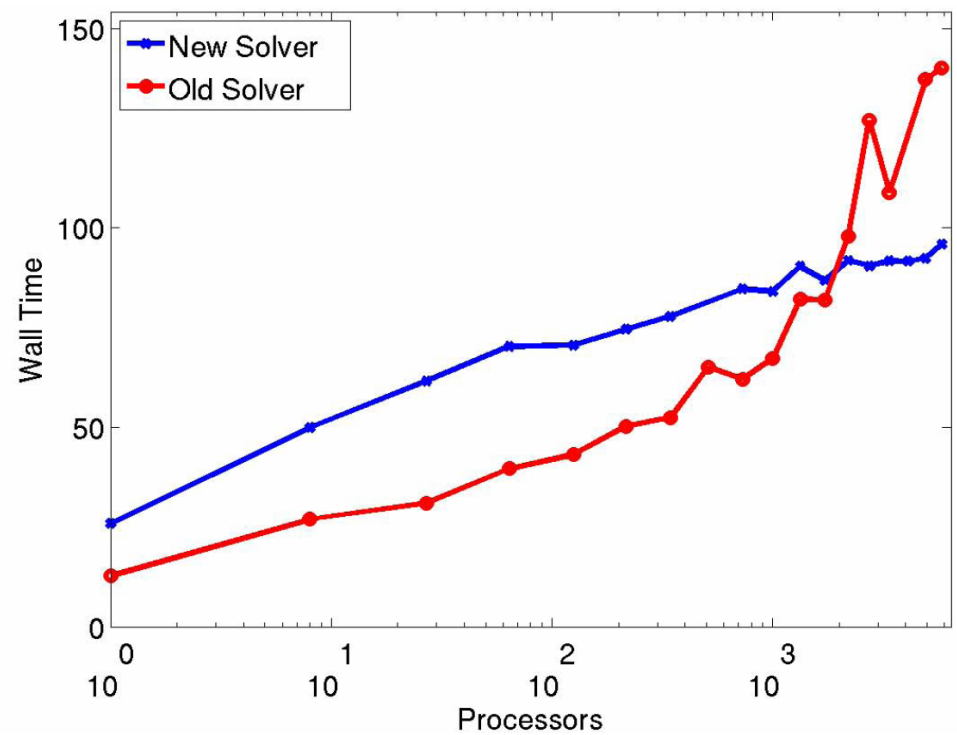

Observations

Multimaterial issue have a significant effect for AMG performance.

However, getting the right overall multigrid solver settings seems at least as important as the effect of multimaterial issues on the multigrid performance on a given problem.

We need to expand our test suite to include smoothly varying properties

Improve matrix of tests versus AMG option settings.

Investigate whether “optimal” default settings exist.

Expensive process.

Summary

Multimaterial modeling impacts scalable performance.

Interface reconstruction algorithms impact scalable performance to a significant degree. High quality reconstruction such as PIR is needed but comes at a cost. This justifies dedicated attention to performance issues related to high order interface reconstruction.

AMR multigrid performance can be strongly dependent on material discontinuities, details of the problem and solver settings. New H(curl) Hodge Laplacian multigrid show promise at large scale.

A continual testing and improvement process is required for large scale capacity computing success today and even more so in the future.

Continued emphasis on answering questions of optimal algorithmic choices appears to be key to achieving future requirements.