SAND2007-1249C

High Performance Parallel Multigrid for Large-Scale Electromagnetics Simulations

Jonathan Hu, Stefan Domino, Allen Robinson, Ray Tuminaro

SIAM CS&E Conference

Monday, February 19th, 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-AC04-94AL85000.

Sandia National Laboratories

- **Goal: improve scalability of parallel smoothed aggregation (SA) in two Sandia applications**
  - **Fuego and ALEGRA**

- **Strong, i.e, fixed global problem size**
- **Weak, i.e., fixed work per processor**

- **Target platforms: Purple (LLNL) and Red Storm (SNL)**

Sandia National Laboratories

Algebraic Multigrid Background

Solve A $u = f$

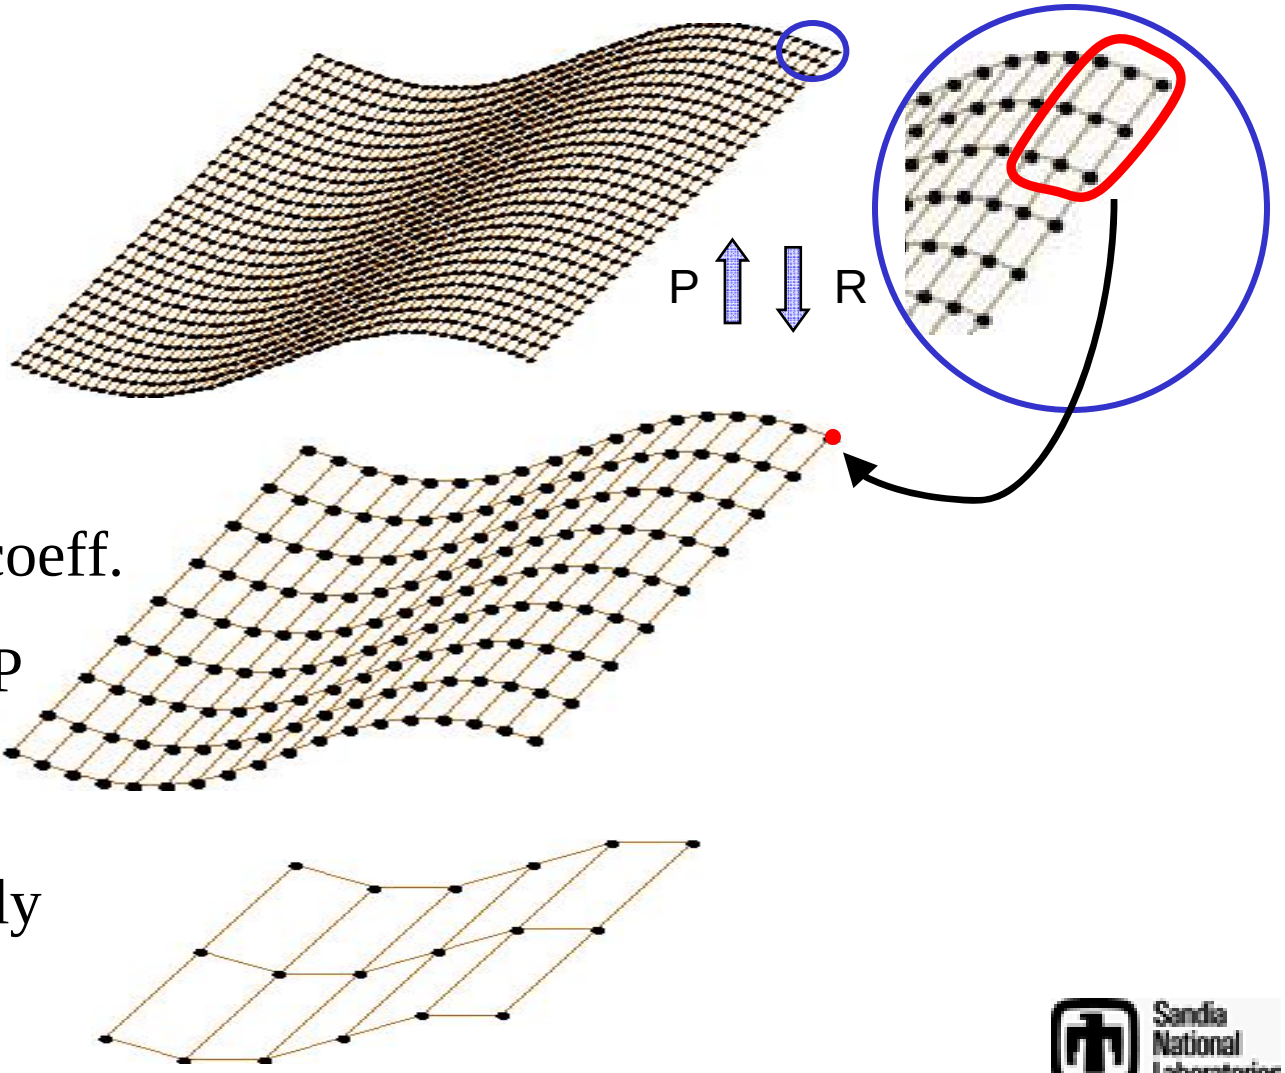

- Coarsen
- Determine interp. coeff.
- Project: $A_H = R\ A\ P$
- Define smoothers
- Continue recursively

Sandia National Laboratories

Coarse Grid Selection (Aggregation)

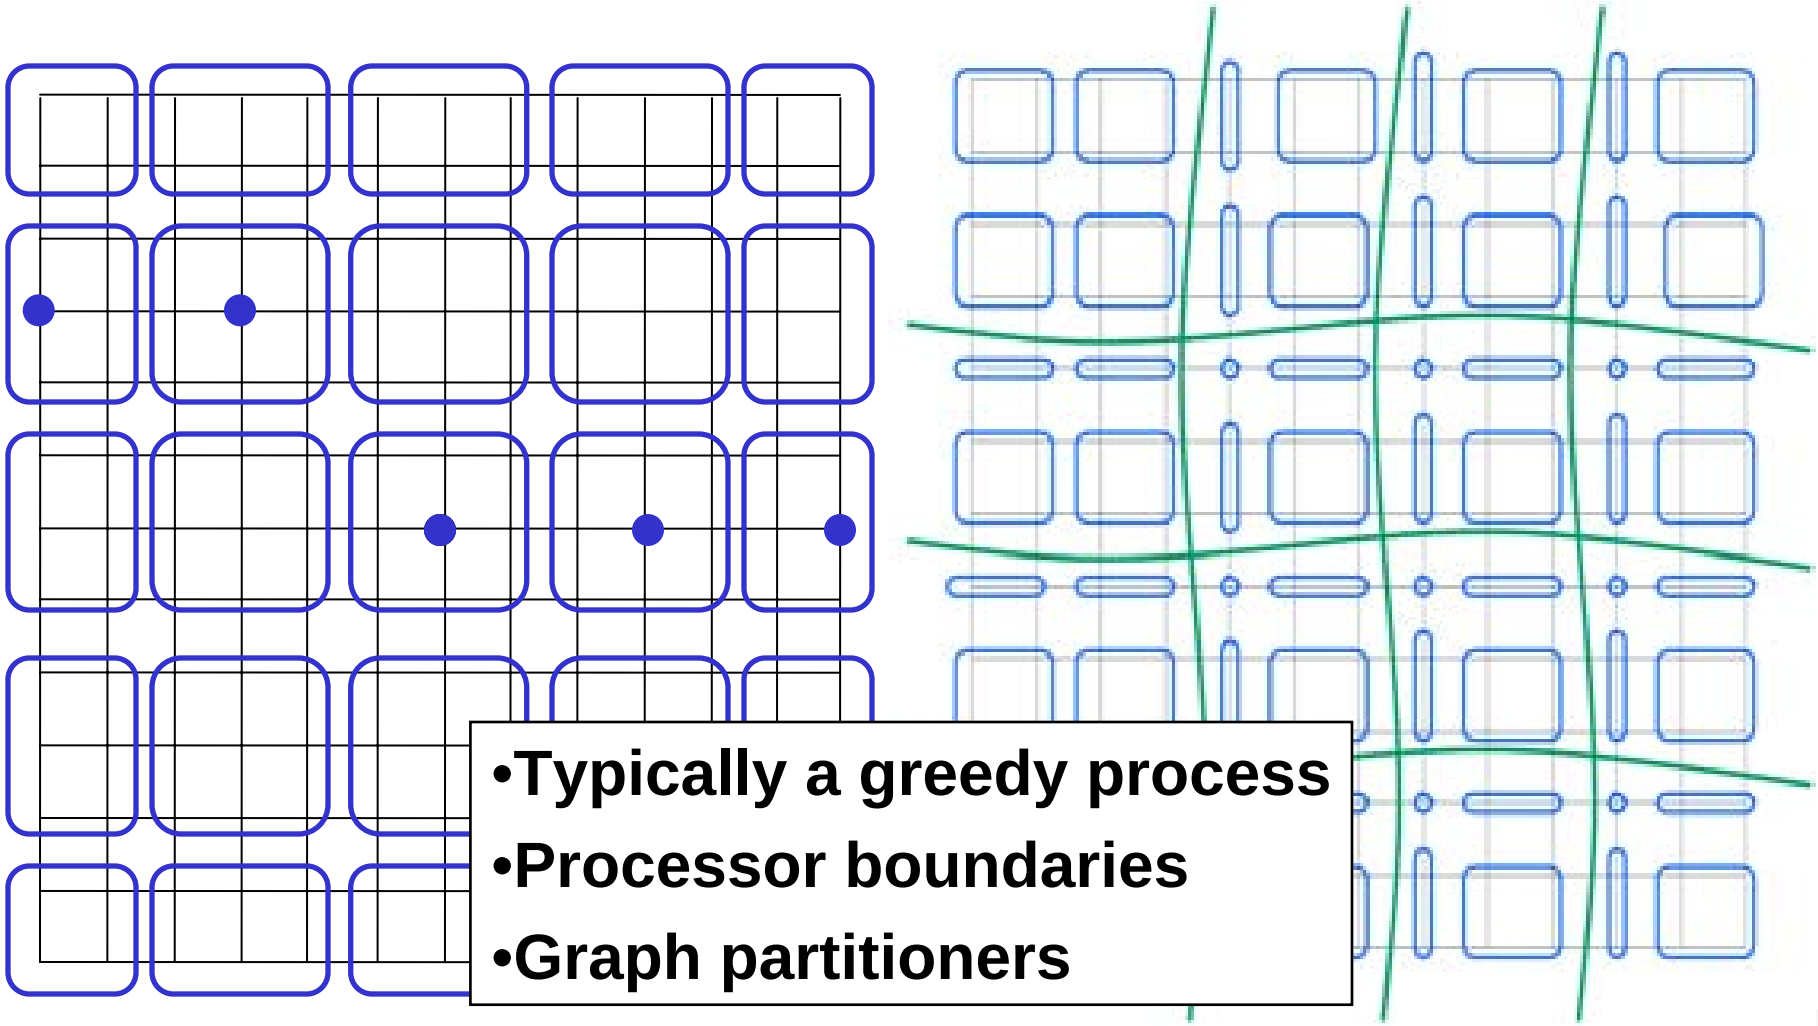

- Typically a greedy process
- Processor boundaries
- Graph partitioners

Smoothed Aggregation: Prolongator Construction

- Near null space $B_{(h)}$ given as input
- Rewrite $B_{(h)}$ to have local support
- ***Tentative prolongator*** $P^{(t)}$

$$B_{(h)} = QR \qquad P^{(t)}B_{(H)}$$

- **Final "smoothed" prolongator** P

$$P = (I - \{diag(A)\}^{(-1)} A)\, P^{(t)}$$

$PB_{(H)} = B_{(h)}$ still holds

- Can repeated recursively for $B_{(H)}$ …

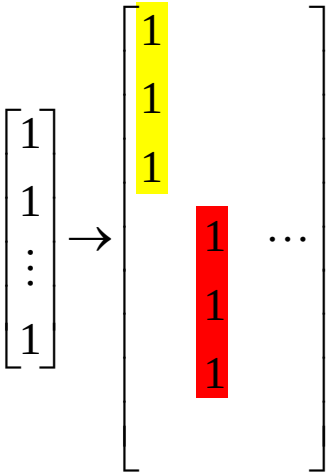

Sandia National Laboratories

Detecting Anisotropies

- **SA coarsens based on relative size of matrix entries:**

$$|a_{ij}| > \varepsilon |a_{ii}||a_{jj}|$$

  - **not always reliable …**

Bilinear FE, isotropic		
-0.0625	-0.0625	-0.0625
-0.0625	0.5	-0.0625
-0.0625	-0.0625	-0.0625
0.0625 / 0.5 = 0.1250		

Bilinear FE, 10:1 aspect		
-0.0420	-0.1658	-0.0420
0.0816	0.3366	0.0816
-0.0420	-0.1658	-0.0420
0.0816 / 0.3366 = 0.2424		

Sandia National Laboratories

Detecting Stretching Using Auxiliary Coordinate Matrix

- **Base dropping on auxiliary matrix, B**
- **Values of B based on coordinates**

$$B_{KJ} = -\frac{1}{dist\left(x_K - x_J\right)^2} \qquad B_{KK} = -\sum_{J \neq K} B_{KJ}$$

- **Nonzero pattern of B same as user's matrix**
- **B is constructed only once, during AMG setup**
- **Semi-coarsen using B along strong connections**

Repartitioning for Parallel Performance

- Load-balancing often well suited for application, not linear solver
  - Coarsening can magnify imbalance
- Load-balancing can increase parallel efficiency on coarse levels
- **Main idea – load balance individual multigrid operators**
  - recursive coordinate bisection (Zoltan)
    - Balance both #rows, #nonzeros
  - Load balance nodal coarse grid after $A_H = RAP$
  - Build permutation matrix Π and do

$$\mathcal{R} = \Pi\, R\ ,\ \mathcal{P} = P\, \Pi^T\ ,\ \mathcal{A}_H = \Pi\, A_H\, \Pi^T$$

- Use processor subset on coarse meshes

Sandia National Laboratories

Eddy current experiments

$$\nabla \times \nabla \times E \quad + \quad \sigma E = rhs$$

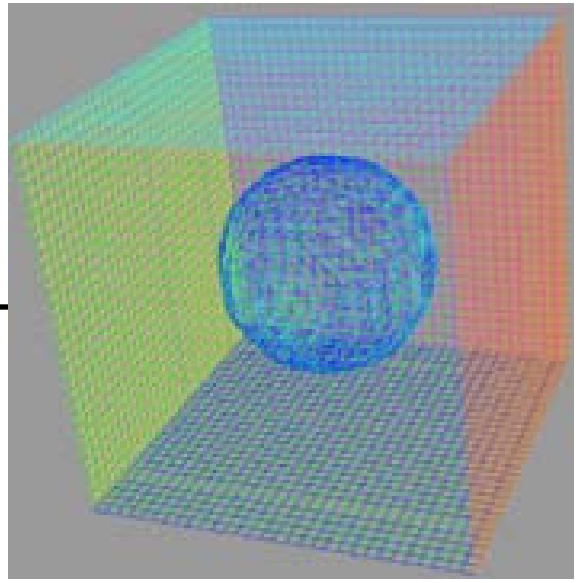

- Two problems on cubic domain
  - Constant conductivity
    - 81K dof (30^3 elements) per processor
  - Cubic domain, conductive central sphere, void outer region
    - Conductivity variation = 1e6
    - 86K dof per processor
- CG + Multigrid preconditioner
  - Two-stage hybrid smoother, degree 2 Chebyshev
  - Specialized interpolation operator
  - Direct coarse grid solve
- Purple and Red Storm platforms

Constant conductivity – weak scaling

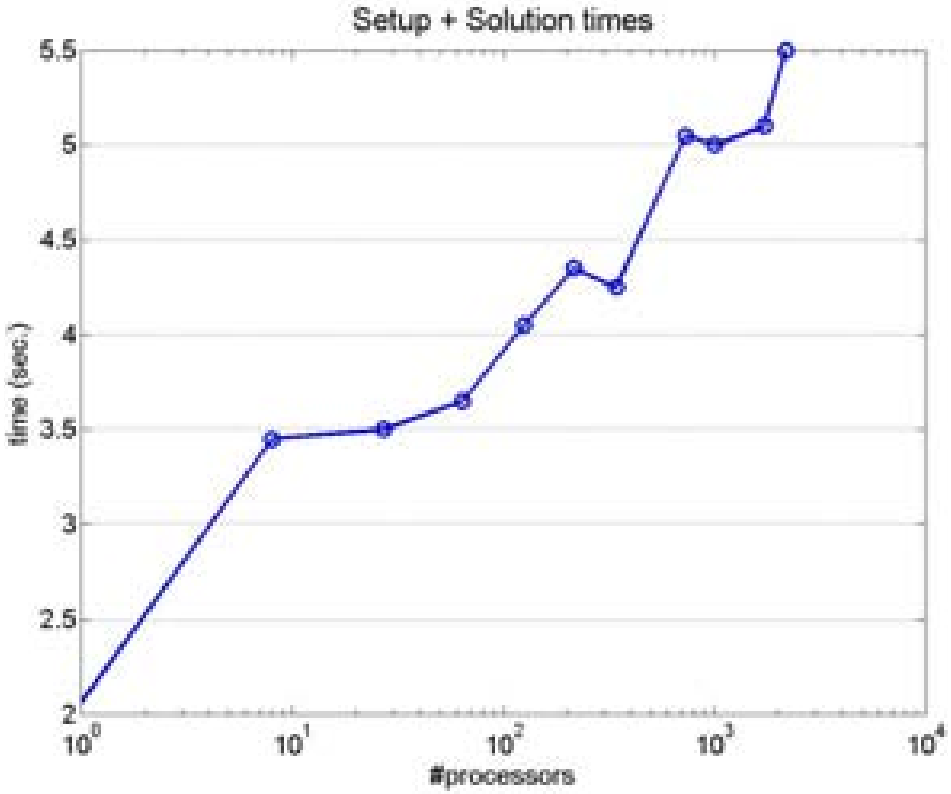

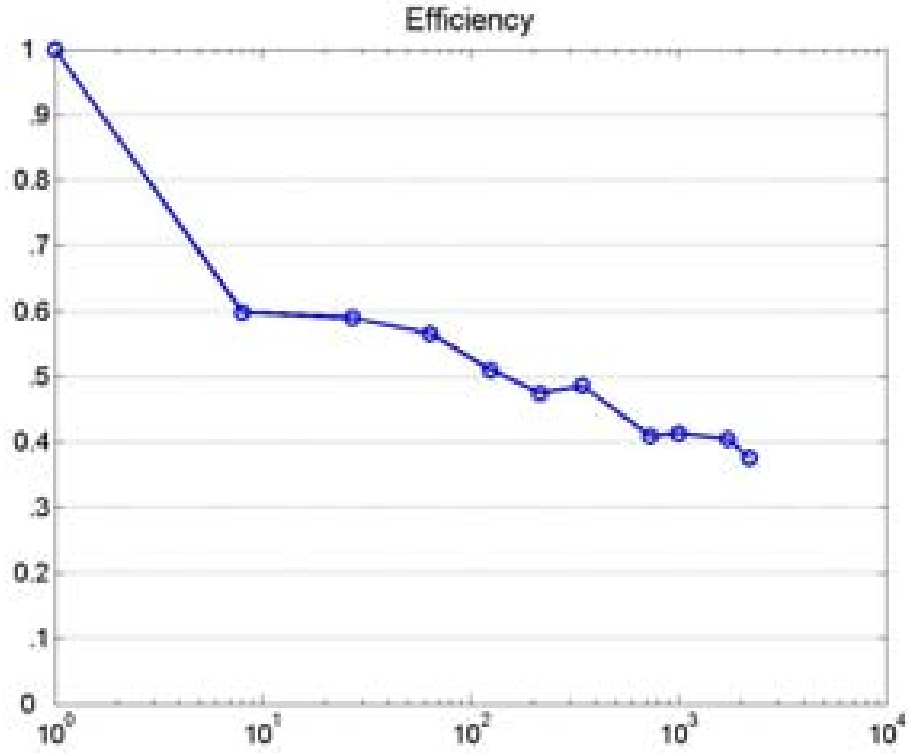

- **Average over 5 time steps**

Sandia National Laboratories

Conductivity jump – weak scaling

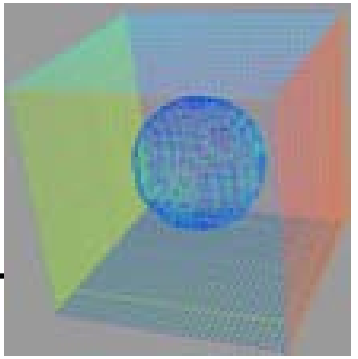

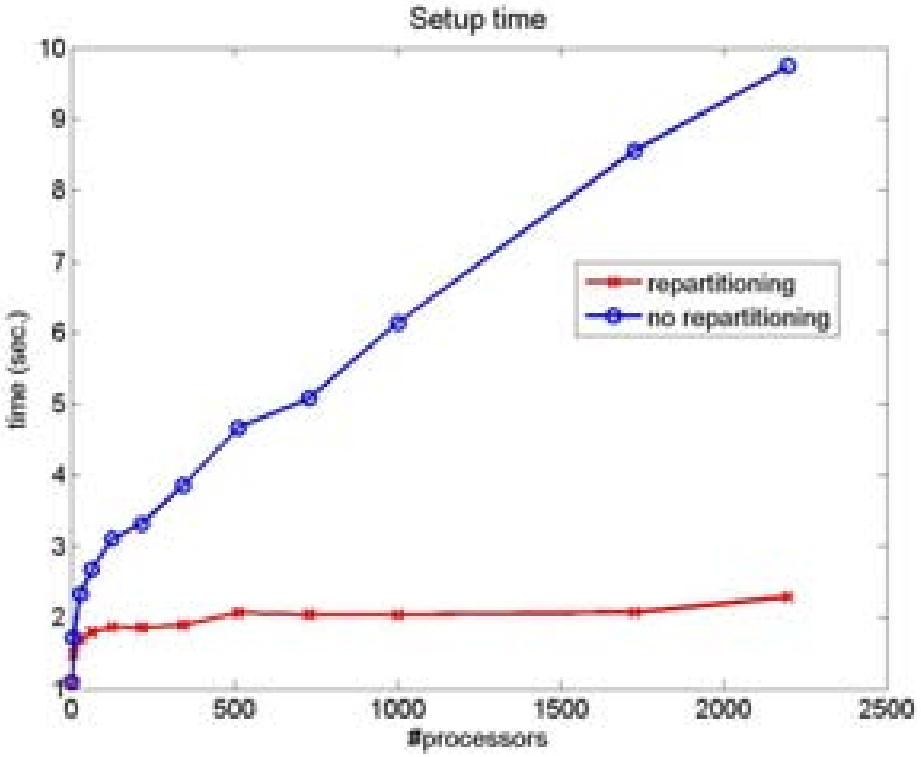

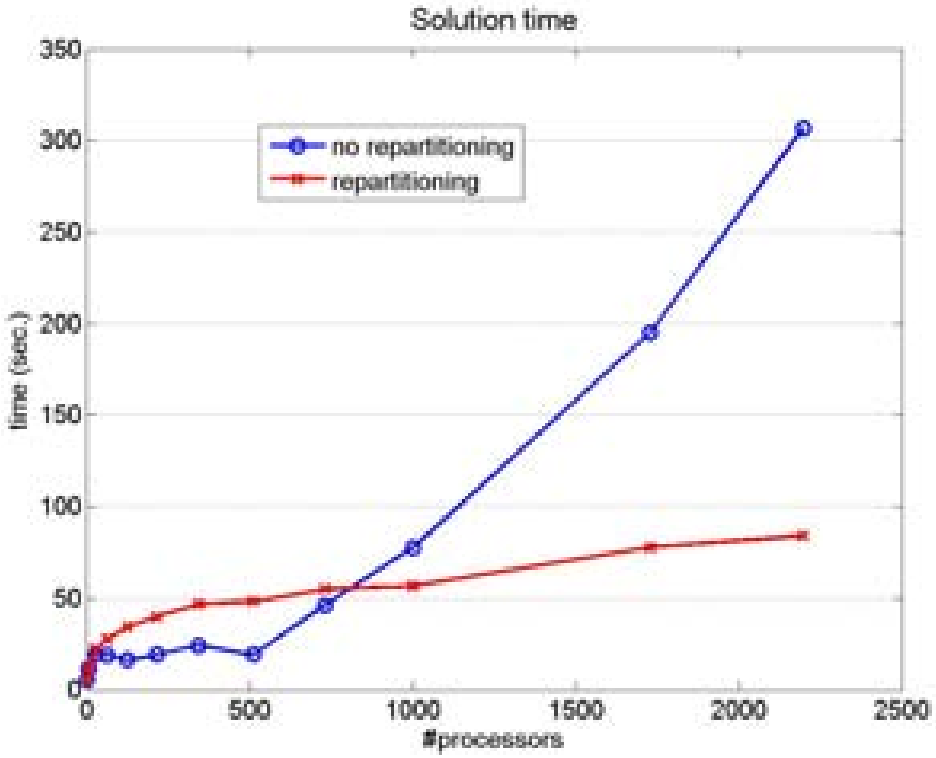

Sandia National Laboratories

Conductivity jump – weak scaling

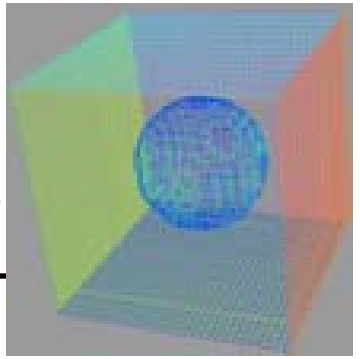

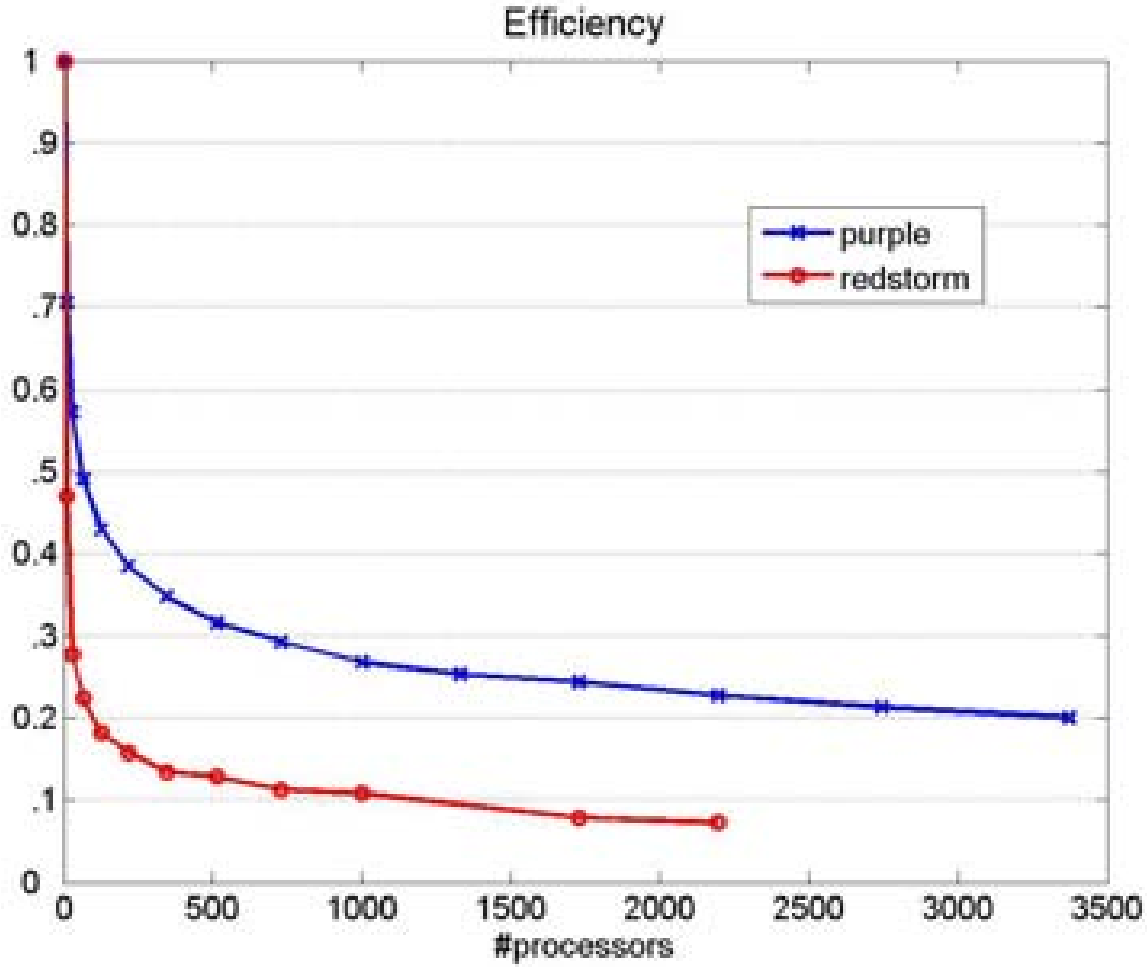

Sandia National Laboratories

Fuego Helium Plume

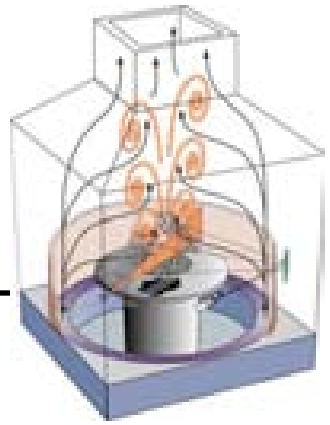

- Uniformly refined meshes from 506K to 32M elements
- Simplified cylindrical geometry
- Pressure solve

- Preconditioning
  - nonsymmetric SA multigrid
  - 2nd degree polynomial smoothing
  - aggregation via auxiliary coordinate matrix

- Computing platform: Red Storm

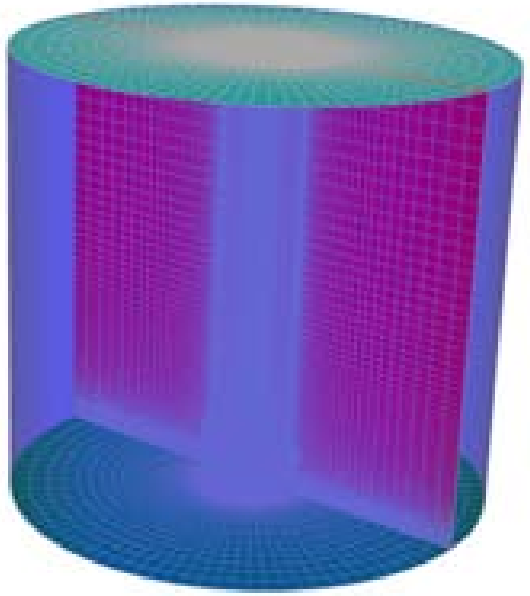

506K element mesh

Sandia National Laboratories

Fuego Helium Plume: Strong Scaling

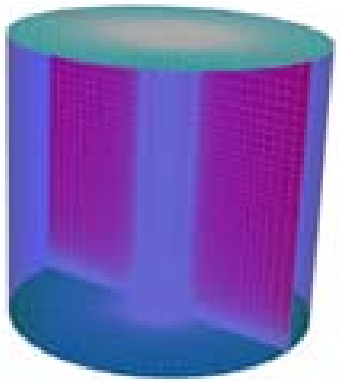

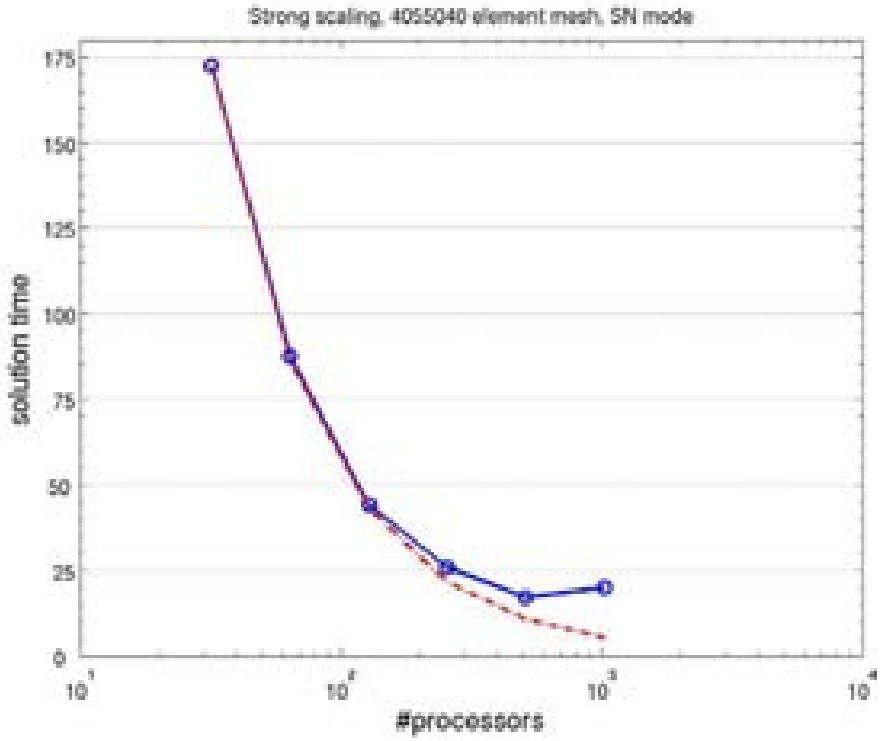

4M element mesh

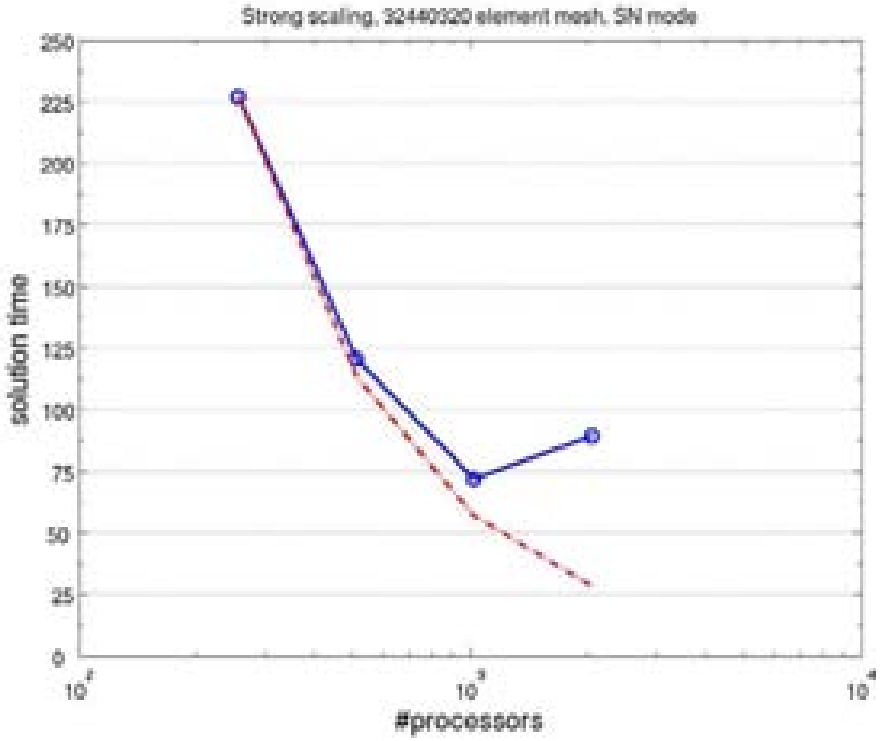

32M element mesh

Sandia National Laboratories

Concluding Remarks

- **Repartitioning can be important**
  - **Setup cost made up for in solve**
  - **Alleviates stress on aggregation**
- **Simple tests have helped us identify bottlenecks**
  - **There are probably more…**

- **Future: a reformulated Maxwell solver should give more direct control of aggregation, e.g.**
  - **C. Siefert, MS 82 (Thursday a.m.)**

Sandia National Laboratories