SAND2008-5010P

Introduction to Multigrid Methods

Chris Siefert

Computational and Shock Multiphysics Group (1431)

Sandia National Laboratories

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.

Outline

- Background.
- Solving Linear Systems with Iterative Methods.
- Introduction to Multilevel Methods.
- Introduction to Algebraic Multigrid.
- Open Questions in Multilevel Methods.

What is Computational Science?

- What do we think of when we think of computational science?
  - Usually “big” things. . .
  - Airplanes, cars, rockets, etc.

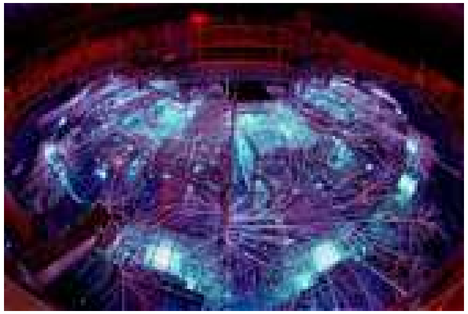

Sandia National Laboratories

What is Computational Science?

- What do we think of when we think of computational science?
  - Usually “big” things. . .
  - Airplanes, cars, rockets, etc.

  - BUT computational science touches everyday things as well!

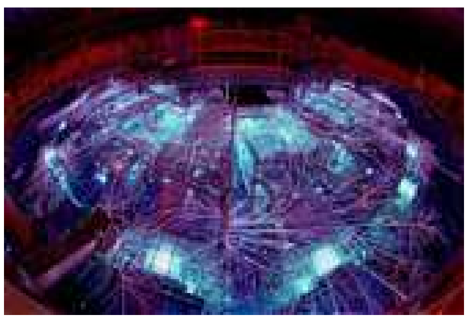

Sandia National Laboratories

Key Computational Kernels

- What are the important kernels for computational science?

Key Computational Kernels

- What are the important kernels for computational science?
  - Time integration (time-dependent problems).
  - Force calculations (particle methods).
  - Solving linear systems (implicit methods, static problems).
  - Interface tracking (shock problems).
  - Load balancing, graph algorithms (parallel problems, direct solvers).
  - Optimization (inverse problems).
  - Eigenvalues (structures problems).
  - . . . and more.

Sandia National Laboratories

Key Computational Kernels

- What are the important kernels for computational science?
  - Time integration (time-dependent problems).
  - Force calculations (particle methods).
  - Solving linear systems (implicit methods, static problems).
  - Interface tracking (shock problems).
  - Load balancing, graph algorithms (parallel problems, direct solvers).
  - Optimization (inverse problems).
  - Eigenvalues (structures problems).
  - . . . and more.

Sandia National Laboratories

Outline

- Background.
- Solving Linear Systems with Iterative Methods.
- Introduction to Multilevel Methods.
- Introduction to Algebraic Multigrid.
- Open Questions in Multilevel Methods.

Importance of Linear Algebra

- Solving linear systems was critical to the example
  $\Rightarrow$ One linear solve per time step!
- This is true of many simulations.
- We can do this w/ Gaussian elimination (GE).
- But is it fast enough?

Sandia National Laboratories

Is GE Good Enough?

A sparse matrix is "any matrix with enough zeros that it pays to take advantage of them." — J. Wilkinson

- For dense problems (almost all entries non-zero), yes.
- But what about sparse problems?
- Example: 1D Heat equation has 3 non-zeros per row.

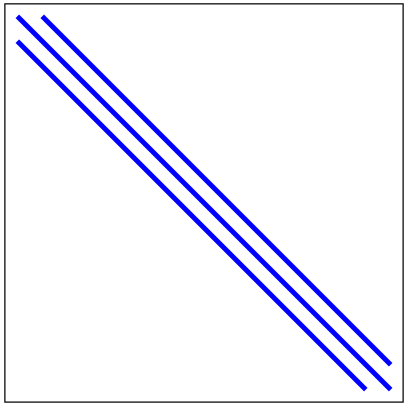

1D Heat Equation Sparsity

Introducing Iterative Methods

$$Ax = b$$

- Idea: Sparse matrix-vector products are cheap
  cost = # non-zeros.
- Let $D = \text{diag}(A)$ contain "a lot" of the matrix. Then,

$$\begin{aligned}(D + (A - D))x &= b \\ Dx &= b - (A - D)x \\ x &= D^{-1}(b - (A - D)x)\end{aligned}$$

- Jacobi's method:

$$x_{i+1} = x_i + D^{-1}(b - Ax_i)$$

- Total Operations $\approx$ nnz.

Speed of Various Methods

Consider a model Laplace problem of size: $n = k^d$, where $d = 2, 3$.

Method	2D	3D
Dense GE	k^6	k^9
Sparse GE	k^3	k^6
Jacobi	$k^4 \log k$	$k^5 \log k$

Table from:

Scientific Computing: An Introductory Survey, 2nd ed. by M.T. Heath

Sandia National Laboratories

Speed of Various Methods

Consider a model Laplace problem of size: $n = k^d$, where $d = 2, 3$.

Method	2D	3D
Dense GE	k^6	k^9
Sparse GE	k^3	k^6
Jacobi	$k^4 \log k$	$k^5 \log k$
Multigrid	k^2	k^3

Table from:

Scientific Computing: An Introductory Survey, 2nd ed. by M.T. Heath

Sandia National Laboratories

Outline

- Background.
- Solving Linear Systems with Iterative Methods.
- Introduction Multilevel Methods.
- Introduction to Algebraic Multigrid.
- Open Questions in Multilevel Methods.

Introducing Multilevel Methods

- Goal: Solve problem with specified mesh spacing, h.
- Idea: Approximate problem w/ coarse mesh H.

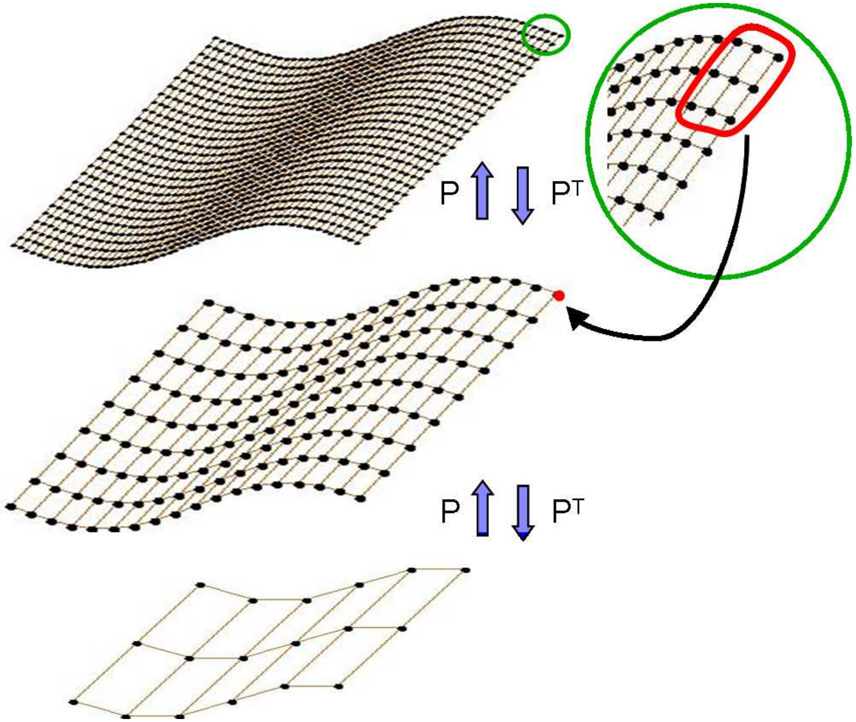

- Big Question: Will this work?

Fourier Series

- Consider a (real) Fourier series

$$f(x) = \frac{a_0}{2} + \sum_{i=1}^{\infty} \alpha_i \cos(2\pi x i)$$

- What do these functions look like?

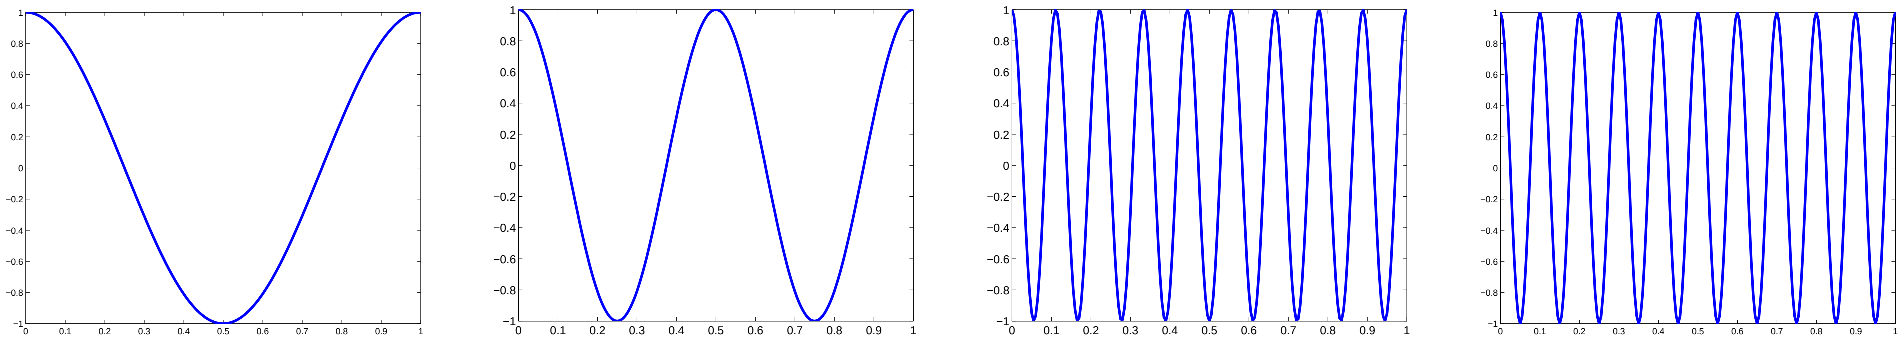

Smooth Oscillatory

Sampling Fourier Modes

- What modes can a discretization sample?

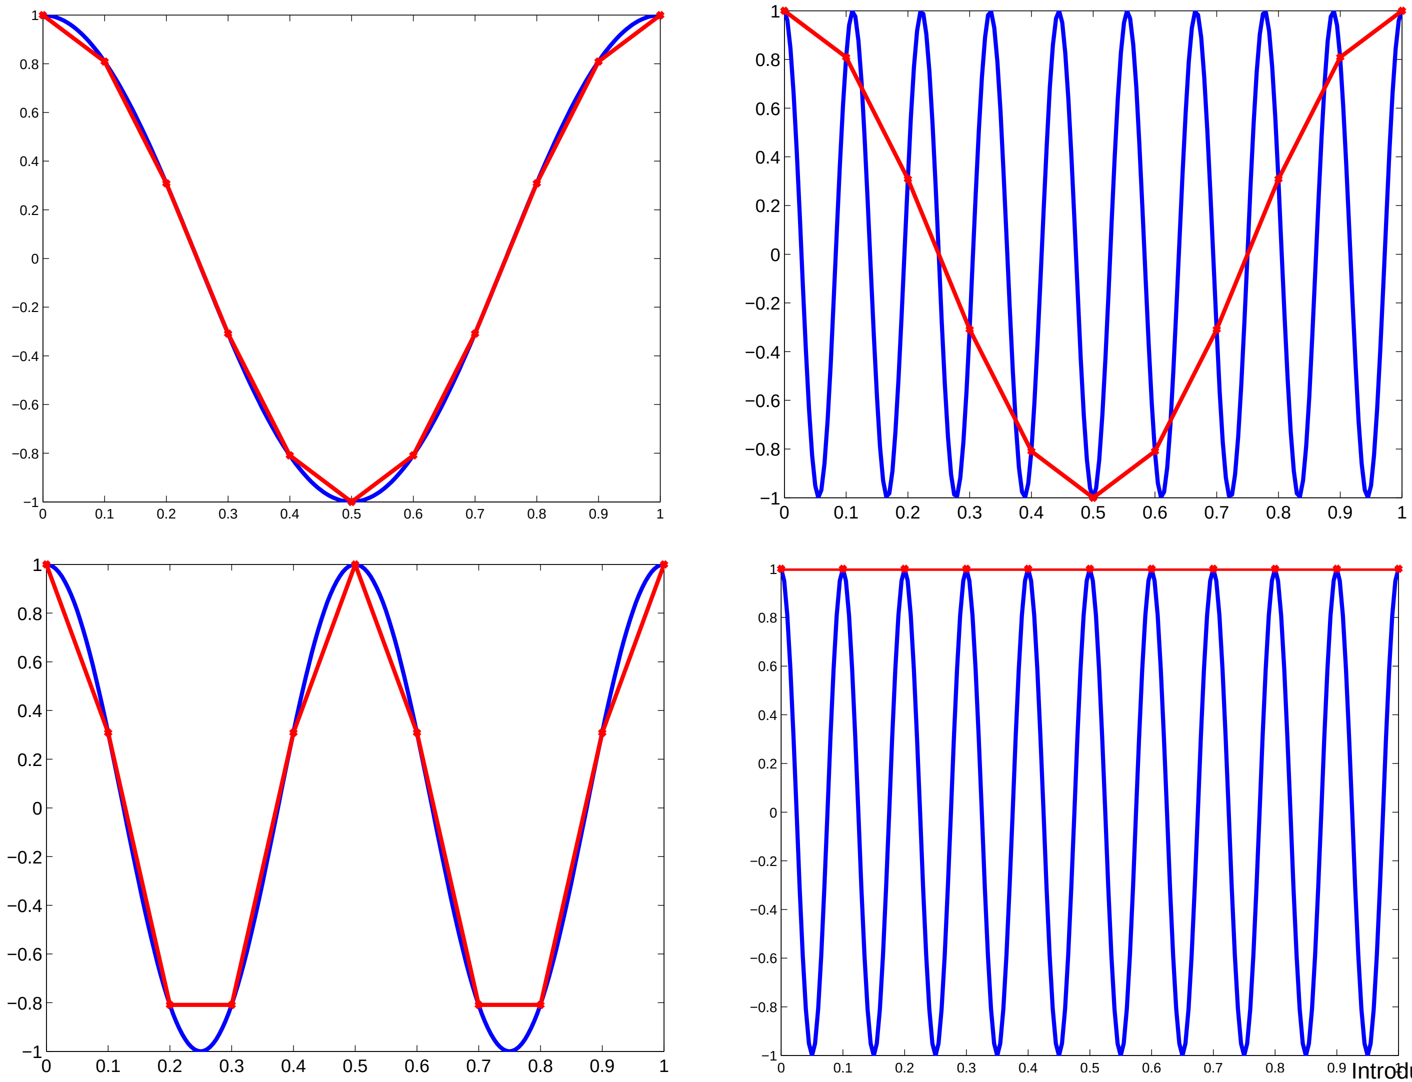

Multigrid & Fourier Modes

- Question: What does this have to do with multigrid?
- Coarse grids can only resolve smooth modes.
- Coarse grids cannot resolve oscillatory modes (aliasing).
- Next question: What about oscillatory modes?

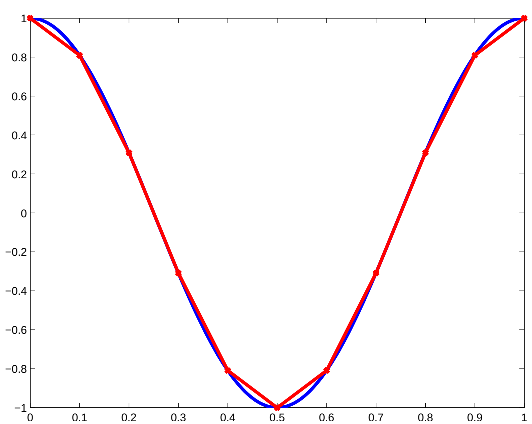

Coarse Grid OK.

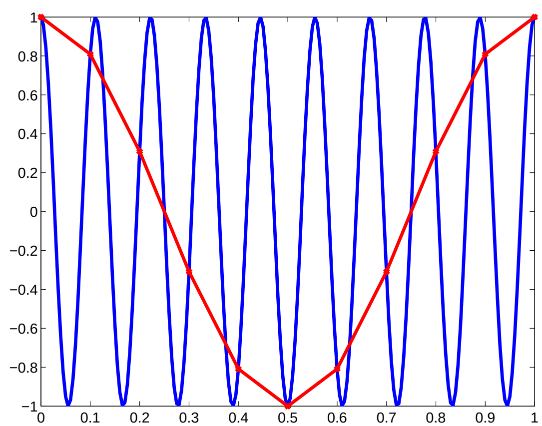

Coarse Grid no help.

Jacobi to the Rescue

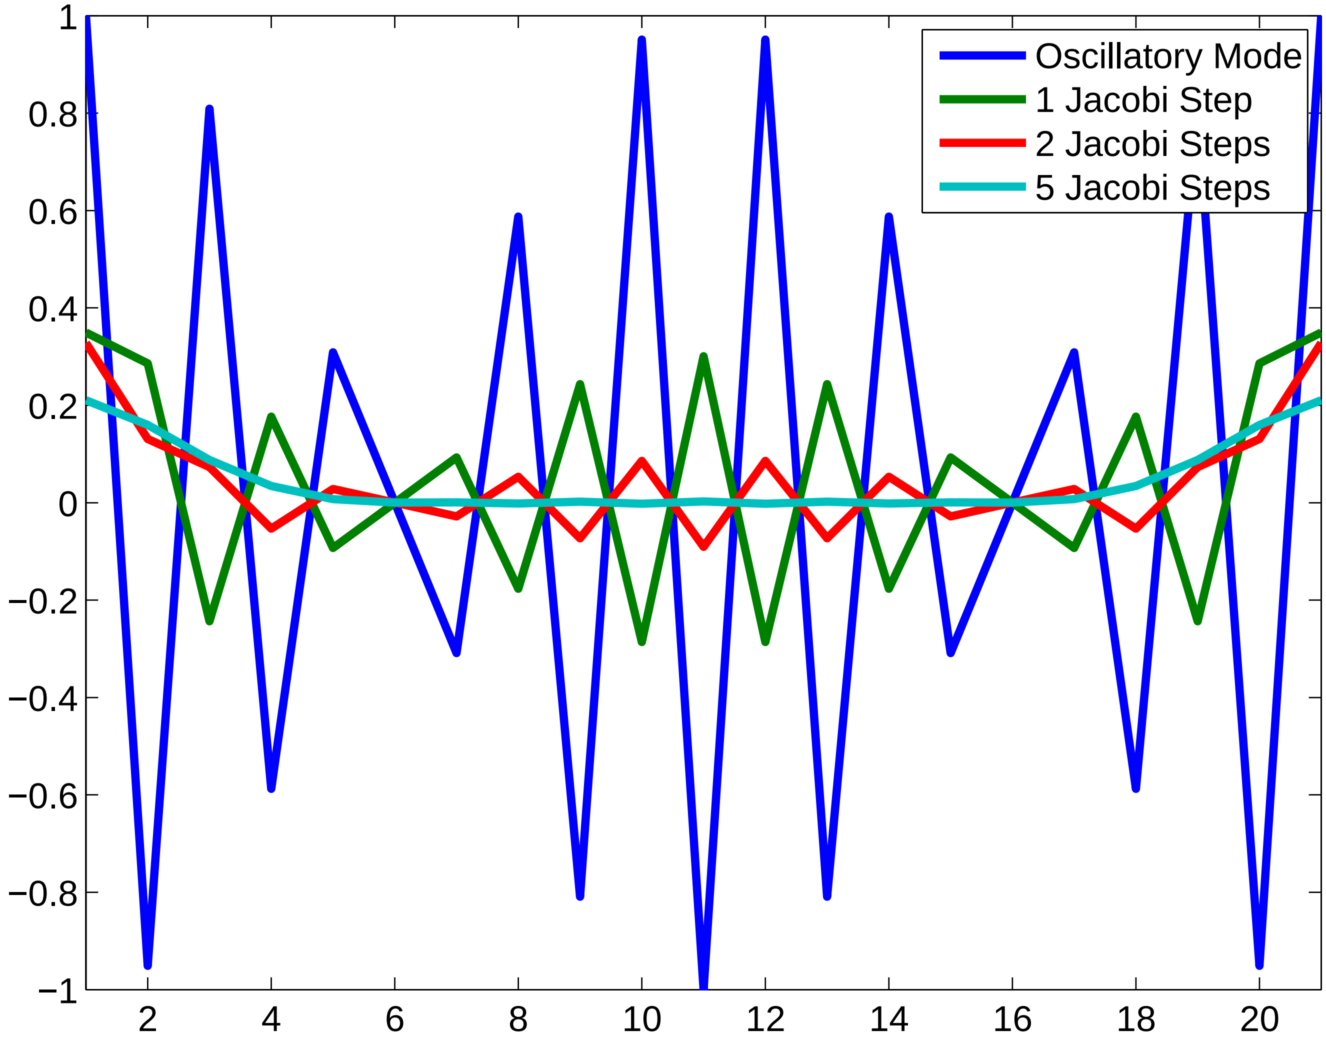

Multigrid by Picture

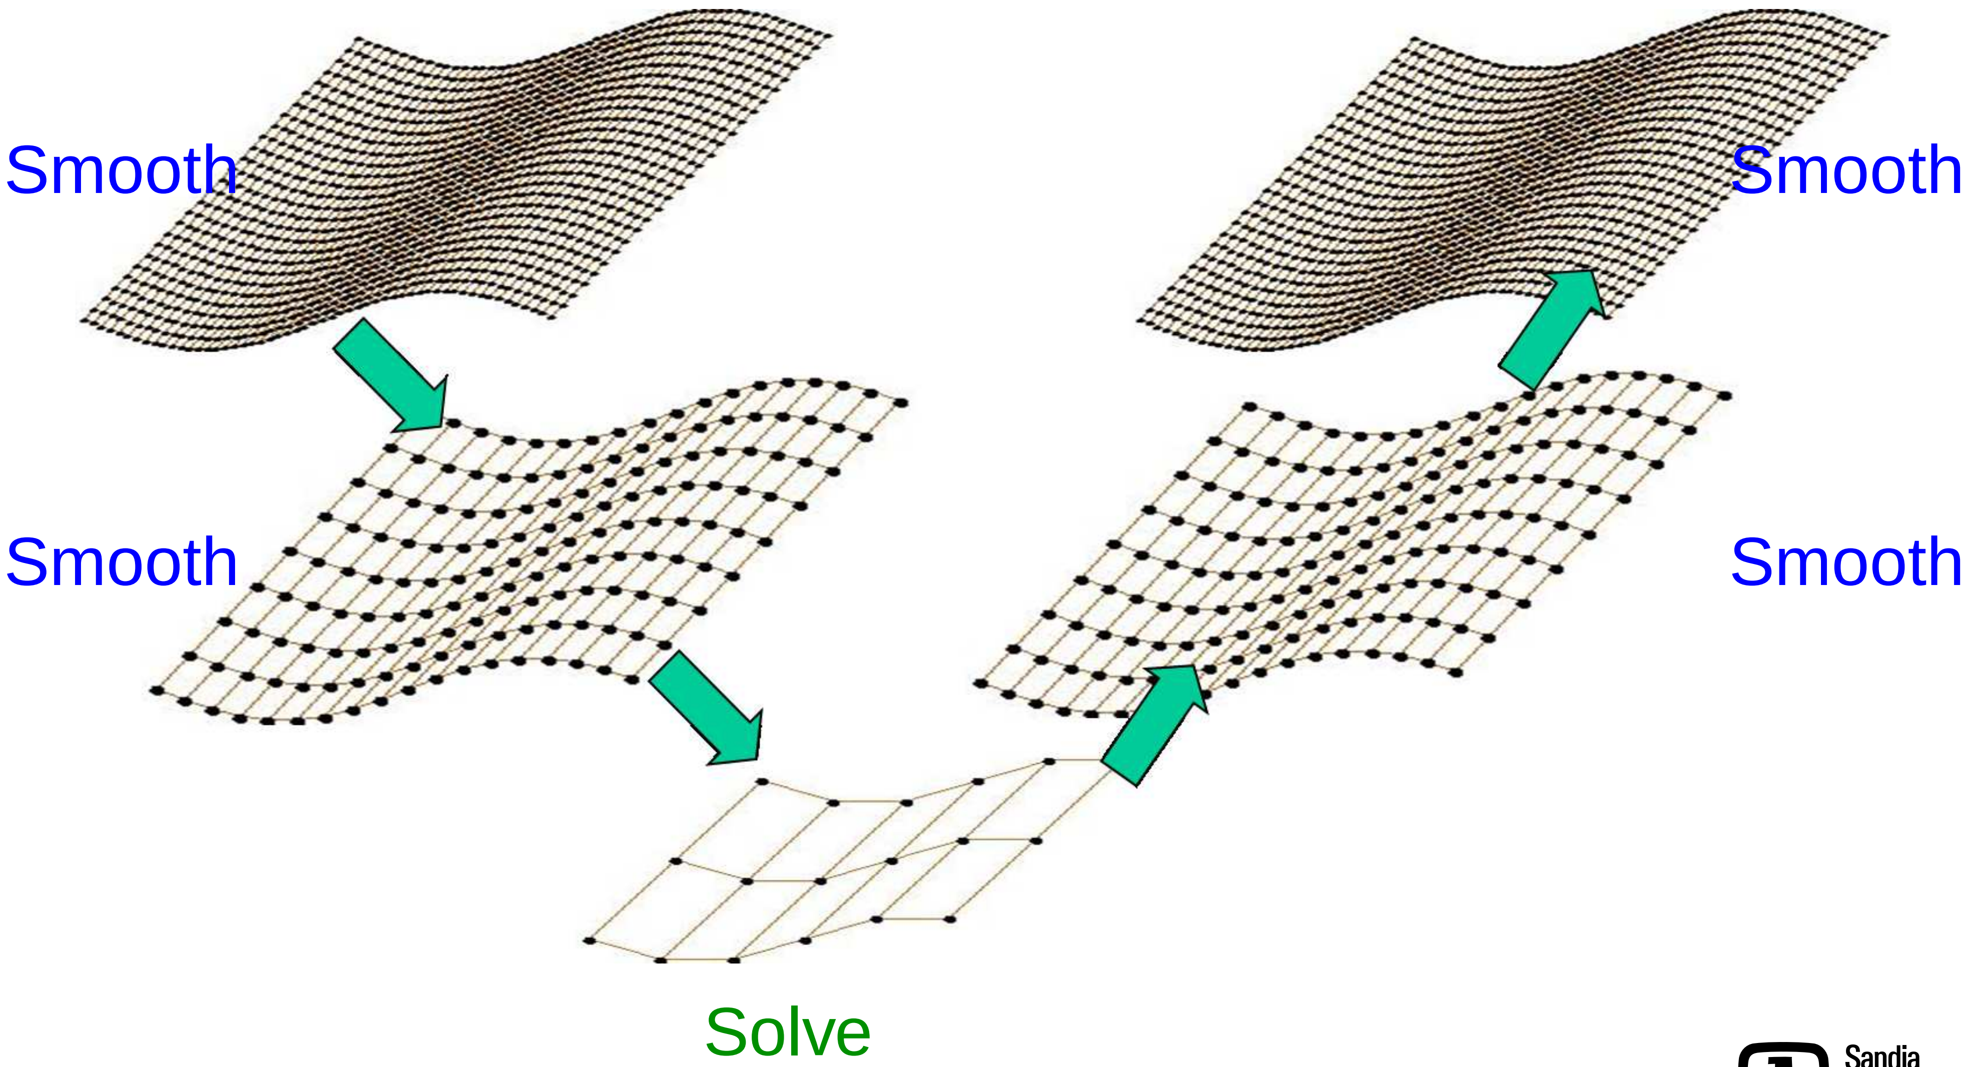

Multigrid Method for $A_h x = b$

Loop until convergence...

1. Smooth on fine grid.
   jacobi(A_h, x, b).
2. Transfer residual $(b - A_h x)$ to coarse grid (restriction).
   $r_c = P^T(b - A_h x)$.
3. Solve on coarse grid.
   $x_c = A_H^{-1} r_c$.
4. Transfer solution to fine grid (prolongation).
   $x = x + P x_c$
5. Smooth on fine grid.
   jacobi$(A_, x, b)$.

Outline

- Background.
- Solving Linear Systems with Iterative Methods.
- Introduction to Multilevel Methods.
- Introduction to Algebraic Multigrid.
- Open Questions in Multilevel Methods.

Multigrid without a Grid

- Multigrid requires a hierarchy of grids
  $\Rightarrow$ inconvenient for the user.
- Can we automatically build the hierarchy?
- Yes! This is called algebraic multigrid (AMG).
- Question: Does the smooth/oscillatory distinction make sense?

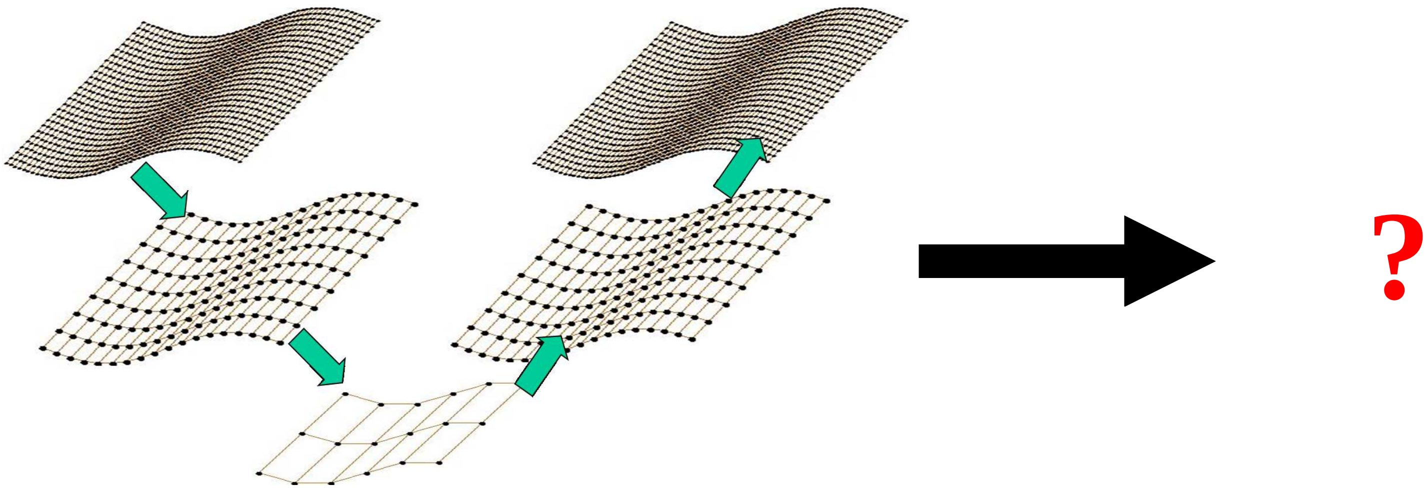

The Logic of AMG

- Start by choosing a smoother.
- If the smoother damps the error. . .
  it is *algebraically* smooth.
- If the smoother doesn't damp the error . . .
  choose the grids so that it is smooth somewhere.
- Note: Contrast this with geometric MG where you pick the grid hierarchy first.

Sandia National Laboratories

The Two Schools of AMG

- Classical AMG (Ruge-Stüben)
  - Choose subset of nodes for the coarse mesh (C-points).
  - Fine-only nodes (F-points) interpolate off of neighboring C-points.
- Smoothed Aggregation (SA)
  - Group or “aggregate” unknowns together to form coarse unknowns.
  - Interpolate based on grouping plus smoothing.

Smoothed Aggregation AMG

• • • • • • • • •

- Define the near null space.
  - Often a null space for an unconstrained problem.
  - Example: Constant vector for heat equation.
  - Example: Rigid body modes for structural mechanics.

Smoothed Aggregation AMG

Near Nullspace

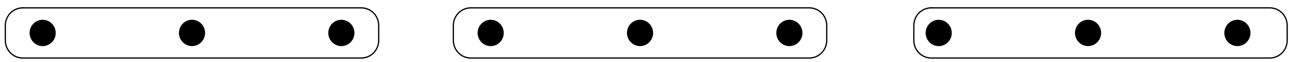

- Define the near null space.
- Aggregate unknowns.

Smoothed Aggregation AMG

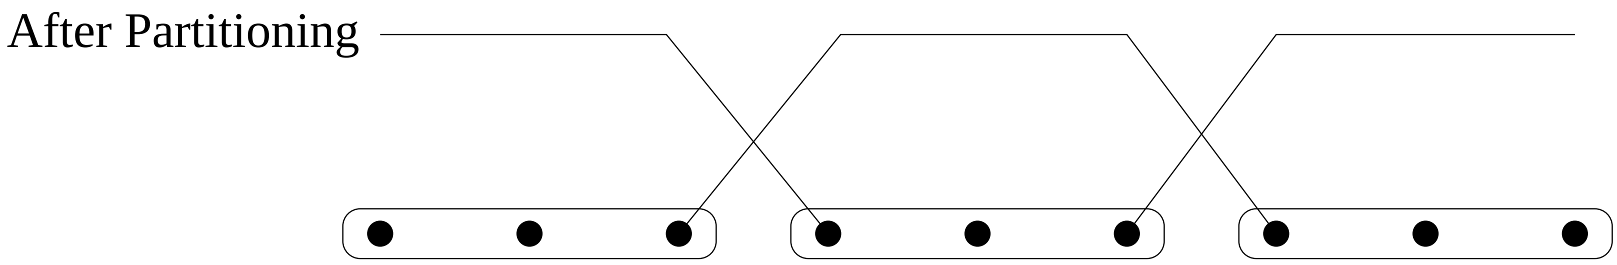

- Define the near null space.
- Aggregate unknowns.
- Partition near null space between aggregates.
  - Preserves near null space on coarse grids.

Smoothed Aggregation AMG

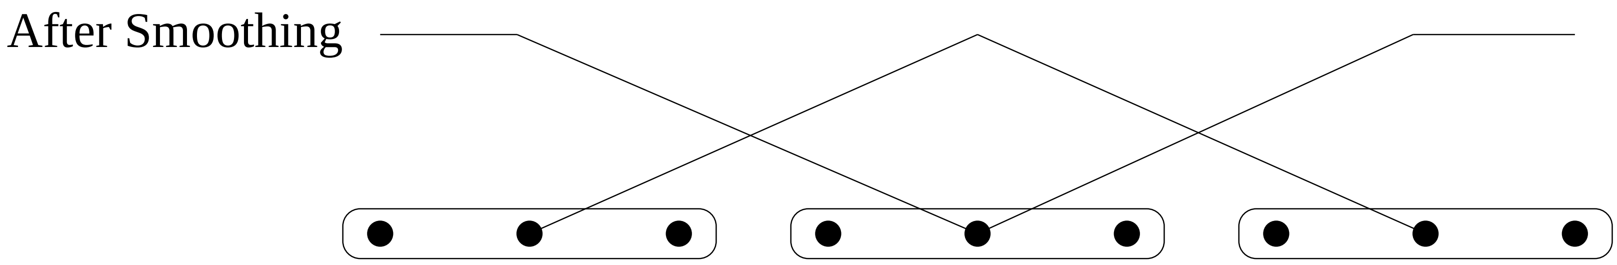

- Define the near null space.
- Aggregate unknowns.
- Partition near null space between aggregates.
- Smooth the prolongator using a step of Jacobi.
  - Preserves null space.
  - Improved interpolation.

Sandia National Laboratories

Multigrid at Sandia

- ML is Sandia's AMG package.
- It provides scalable multilevel/multigrid preconditioners.
- Method types
  - Smoothed Aggregation (SA) - symmetric or nearly symmetric problems.
  - Non-symmetric SA - non-symmetric problems.
  - MatrixFree - matrix-free SA.
  - DD / DD-ML - domain decomposition.
  - Maxwell / RefMaxwell - Maxwell's equations.

Sandia National Laboratories

Trilinos Summary

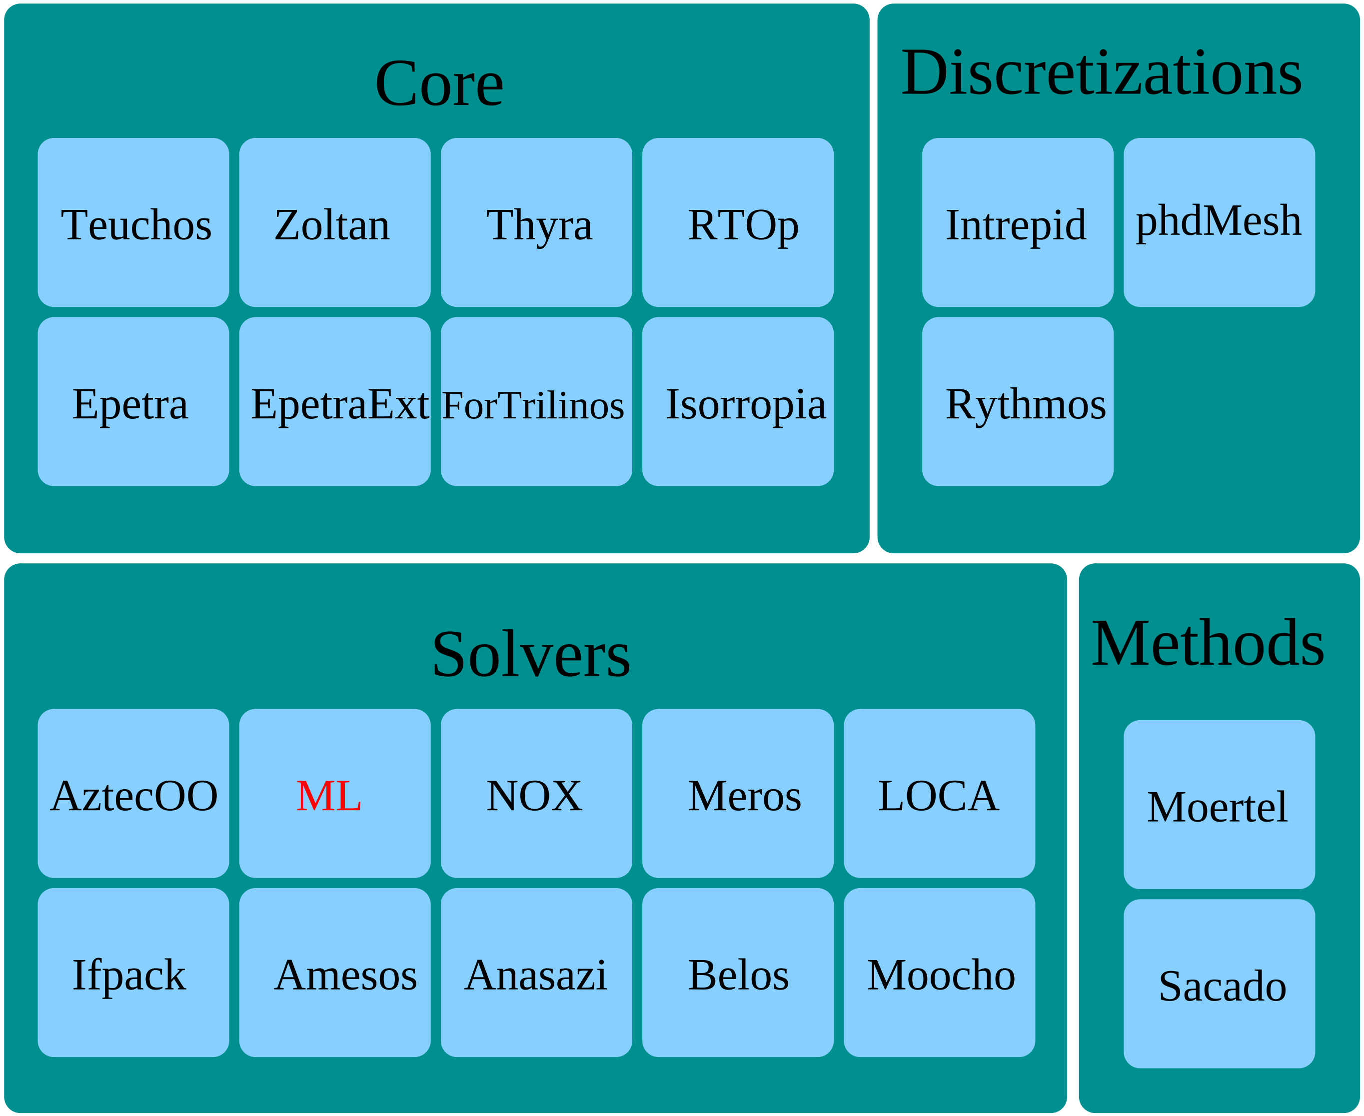

Sandia National Laboratories

Outline

- Background.
- Solving Linear Systems with Iterative Methods.
- Introduction to Multilevel Methods.
- Introduction to Algebraic Multigrid.
- Open Questions in Multilevel Methods.

Open Questions in Multigrid

- MG is designed for Laplace/Heat problems
- On other problems additional issues arise.
- Mathematical issues: anisotropy, systems, variable materials.
- Computer science issues: parallelism, scalability.

Math Issue #1: Anisotropy

$$\frac{\partial^2 u}{\partial x^2} + \epsilon \frac{\partial^2 u}{\partial y^2} = f$$

- Anisotropic operators have direction-dependent behavior.
- Example: Heat diffuses “faster” in y direction (ϵ small).
- Tests varying ϵ w/ $10,000$ unknowns.

	$\epsilon = 1$	$\epsilon = 10^{-1}$	$\epsilon = 10^{-2}$	$\epsilon = 10^{-3}$	$\epsilon = 10^{-4}$
Iterations	14	20	53	129	189

- This is BAD!

Reacting to Anisotropy

- Better meshes fix some problems.

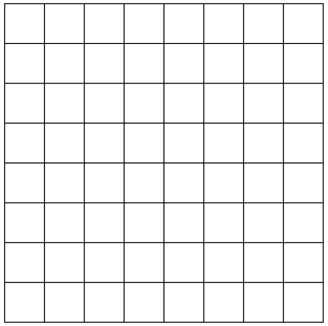

Isotropic Mesh

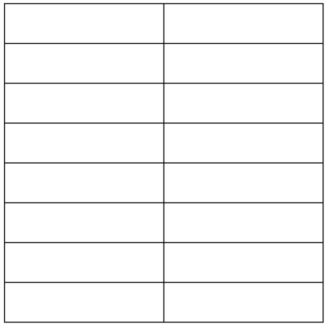

Anisotropic Mesh

- Meshes alone cannot solve hard problems.
- Research problem: Robust detection of anisotropy.
- Research problem: Non-axial anisotropy.

Math Issue #2: PDE Systems

- PDE systems multiple different types of variables (e.g. displacement, velocity, pressure, temperature, etc.).
- Example: Linear elasticity.
- One solution: Smoothed aggregation — explicitly preserve null space on coarse levels.
- Research problem: Fluid problems (e.g. Navier-Stokes).

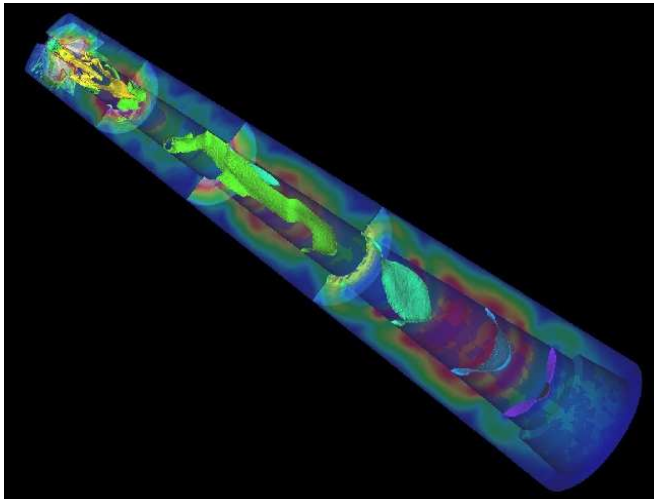

Image courtesy of the CSAR/UIUC
`http://www.csar.uiuc.edu`

Sandia National Laboratories

Math Issue #3: Multimaterial

- Material interfaces can be sites of discontinuities $\Rightarrow$ oscillatory modes at boundaries.
- Features can be hard to resolve on coarse grid.

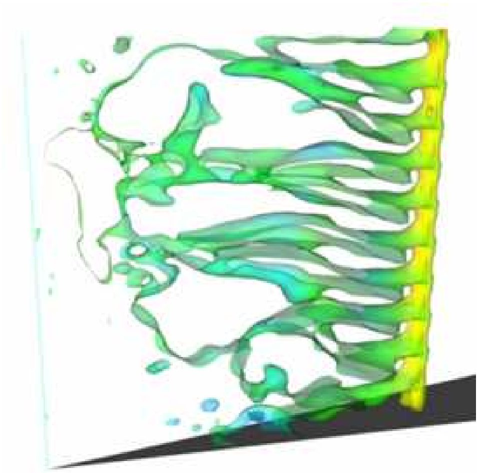

- Research problem: Detecting material interfaces.
- Research problem: Handling disappearing features.

Sandia National Laboratories

CS Issues: Parallelism

- More processors *should* lead to faster solutions.
- Strong scaling — fix work, increase processors.
- Example: 2,000 steps of Jacobi.

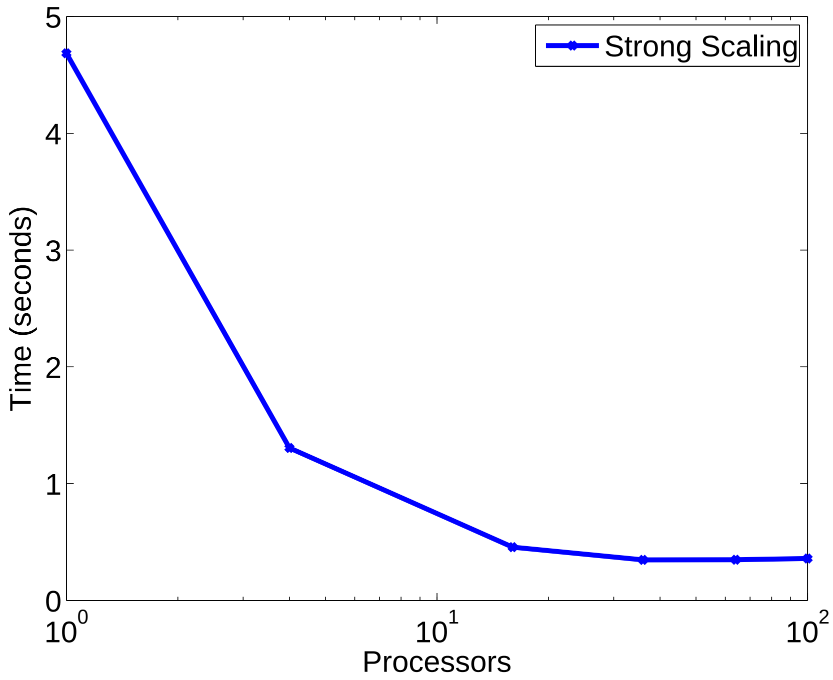

CS Issues: Parallelism

- More processors *should* lead to faster solutions.
- Strong scaling — fix work, increase processors.
- Example: 2,000 steps of Jacobi.

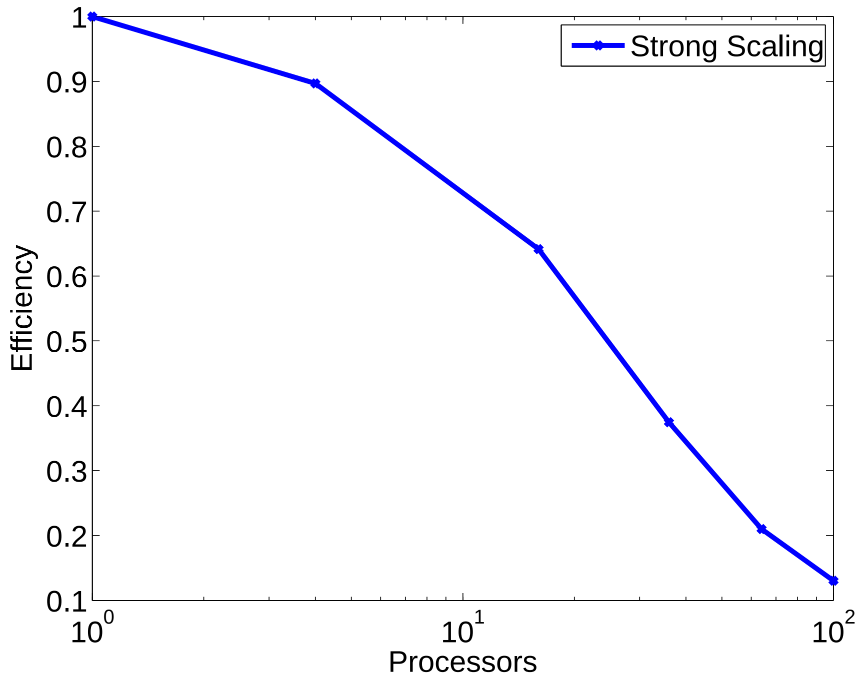

- Question: What causes the loss in efficiency?

Understanding Efficiency

- Answer: Computation to communication ratio.
- Weak scaling — fix work per processor.

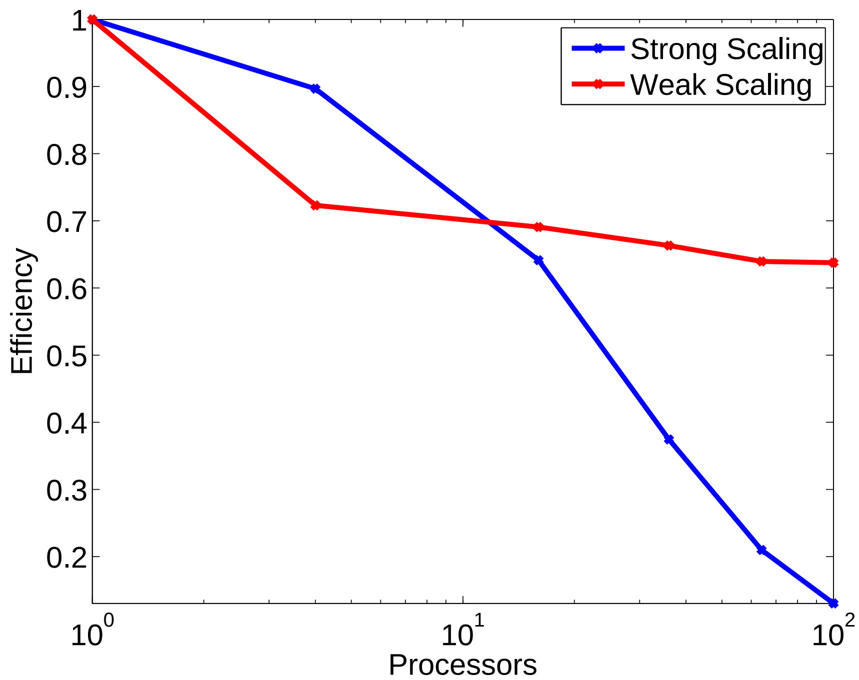

- Message: What works on a small # of procs, might not work on a large #.

CS Issue #1: Scalability

- Coarse grids $\Rightarrow$ less work per proc $\Rightarrow$ poor performance.
- One solution: Move data to leave some procs idle.
- Research problem: What is the best way to repartition?
- Research problem: How to address poor performance on really big (terascale) computers.

Red Storm(SNL) 26,569 procs

Jaguar(ORNL) 23,016 procs

Sandia National Laboratories

Take Home

"I would rather have today's algorithms on yesterday's computers than vice versa." - Reported by P. Toint

- Importance of good algorithms.
- Rationale behind multilevel algorithms.
- Nature of the "big questions" in multilevel algorithm research.
  - Math: Anisotropy, multimaterial, PDE systems.
  - CS: parallelism, scalability.
- My web site: `http://www.sandia.gov/~csiefer`
- Trilinos project: `http://trilinos.sandia.gov`

Sandia National Laboratories