Multigrid & Optimal Solvers

Ray Tuminaro

Bochev, Elman, Gee, Howle, Hu, Lin, Sala, Schroder, Shuttleworth, Siefert, Shadid

GeoMath08

Sept, 2008

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.

What is Multigrid ?

Solve $A_3 u_3 = f_3$

Basic idea:

- Develop coarse approximations
- Accelerate convergence via coarse iterations to efficiently propagate information

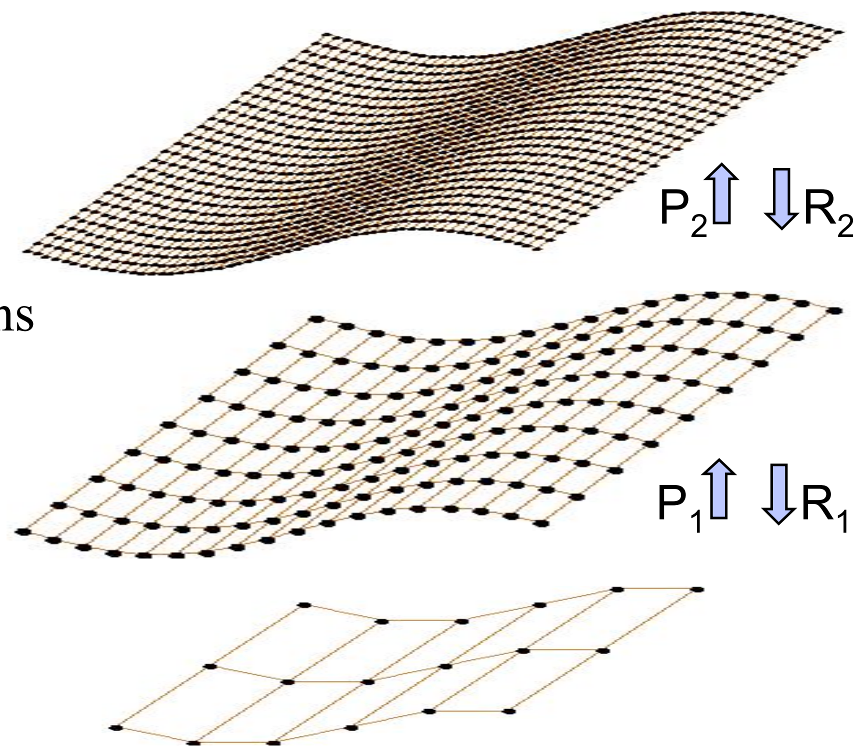

Smooth $A_3 u_3 = f_3$. Set $f_2 = R_2 r_3$.

Smooth $A_2 u_2 = f_2$. Set $f_1 = R_1 r_2$.

Solve $A_1 u_1 = f_1$ directly.

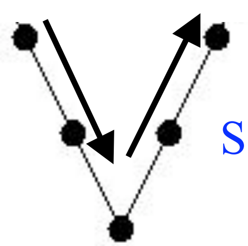

Set $u_3 = u_3 + P_2 u_2$. Smooth $A_3 u_3 = f_3$.

Set $u_2 = u_2 + P_1 u_1$. Smooth $A_2 u_2 = f_2$.

What is Algebraic Multigrid ?

Solve $A_3 u_3 = f_3$

- Construct Graph & Coarsen
- Determine P_i & R_i sparsity pattern
- Determine P_i & R_i's coefs
- Project: $A_i = R_i A_{i+1} P_i$

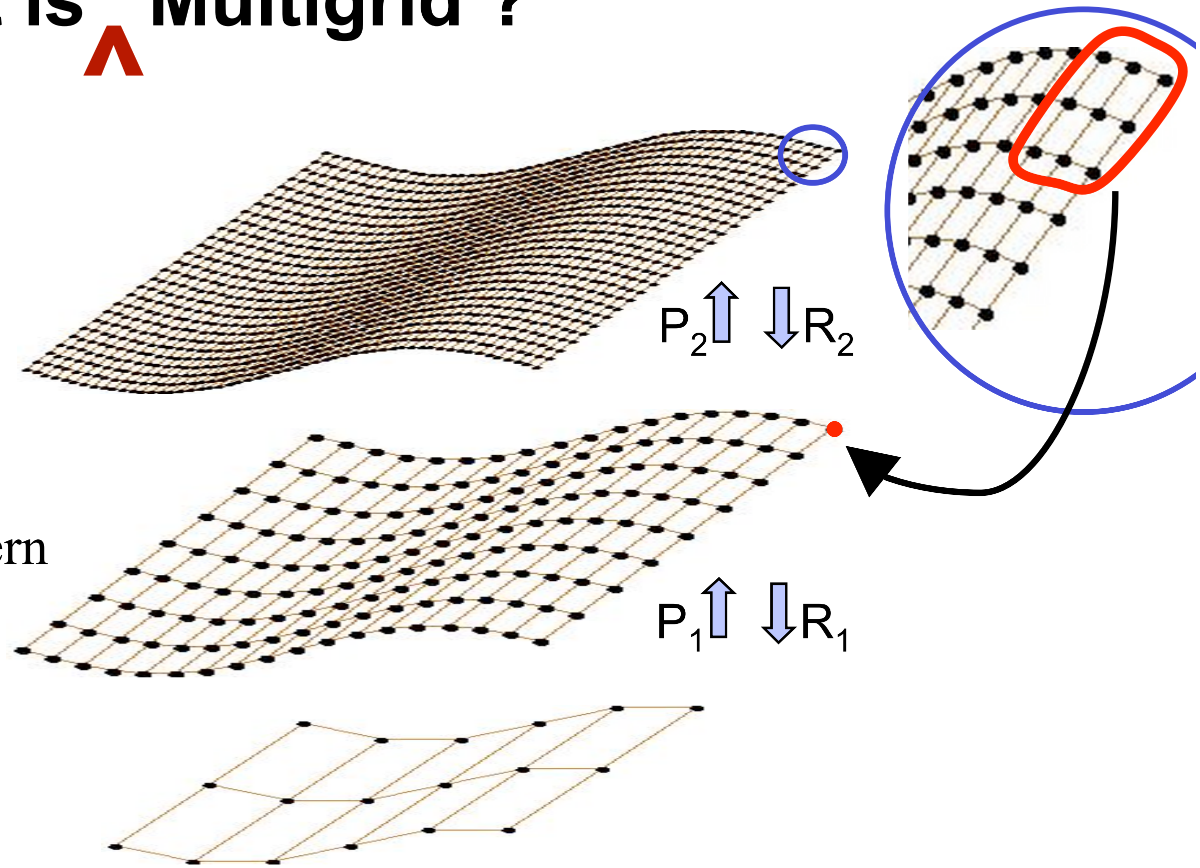

Smooth $A_3 u_3 = f_3$. Set $f_2 = R_2 r_3$.

Smooth $A_2 u_2 = f_2$. Set $f_1 = R_1 r_2$.

Solve $A_1 u_1 = f_1$ directly.

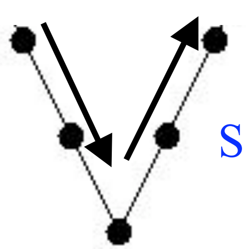

Set $u_3 = u_3 + P_2 u_2$. Smooth $A_3 u_3 = f_3$.

Set $u_2 = u_2 + P_1 u_1$. Smooth $A_2 u_2 = f_2$.

Sandia National Laboratories

How does it work?

Smoothing and coarse correction complementary!

- Error with increasing Jacobi iterations

- Can be represented on coarse mesh

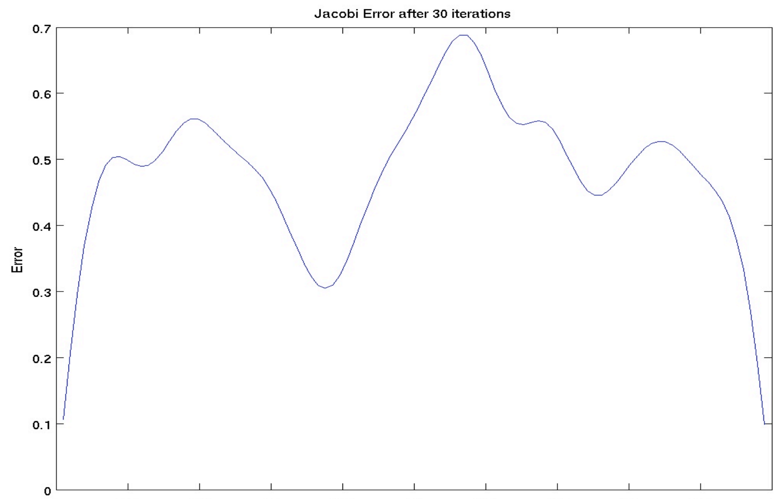

A Massively Parallel Algebraic Multigrid Solver

Smoothed Aggregation Capabilities

- Scalar & PDE systems (elliptic)
  - Symmetric, non-symmetric
  - variable dofs/block support (limited)
- Aggregation with arbitrary coarsening, load balancing, …

Smoothers

- Gauss-Seidel, polynomial, block methods, ILU, domain methods, ...

Package Leveraging

- Trilinos (Epetra, Ifpack, etc.)
- External: PETSc, SuperLU, Arpack, Parasails, kLU, ParMETIS, Zoltan, …
  - PETSc applications can construct and apply essentially any Trilinos preconditioner/iterative method, KSP solvers as smoothers

Interfaces

- C++, matlab, matlab-like, web-based

http://trilinos.sandia.gov/packages/ml

- Developer's Guide, User's Guide, MLAPI

A Massively Parallel Algebraic Multigrid Solver

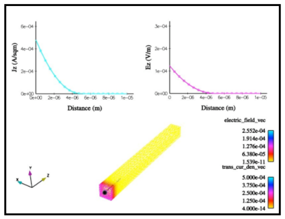

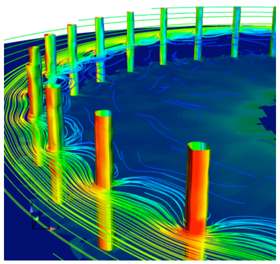

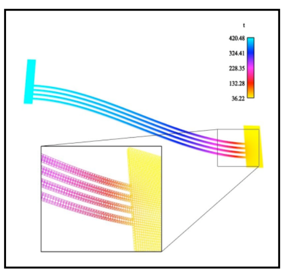

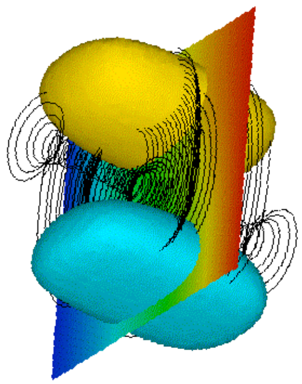

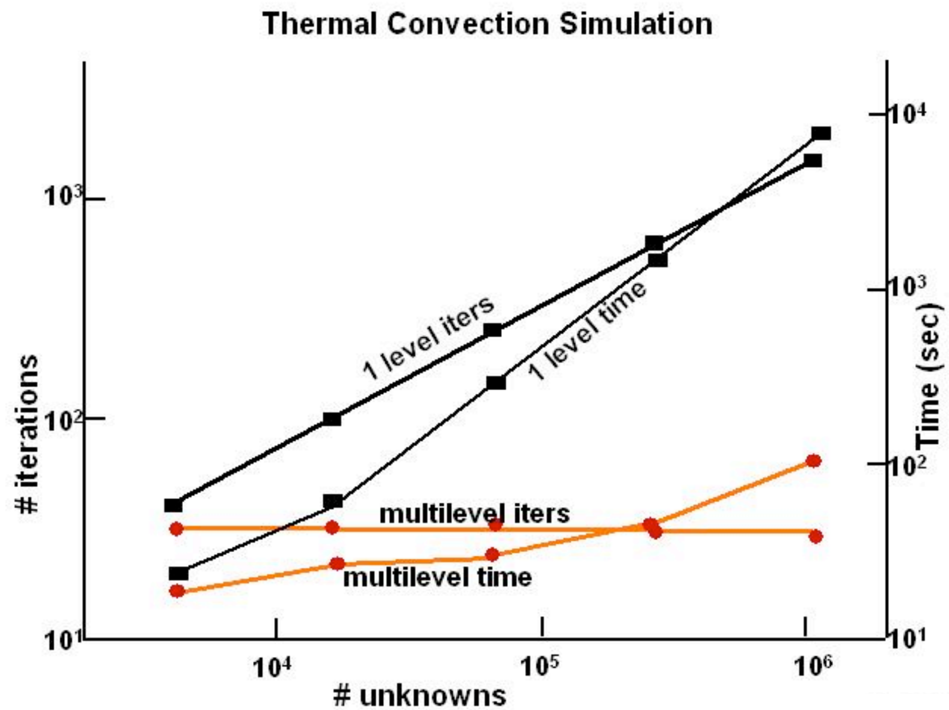

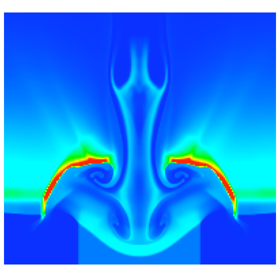

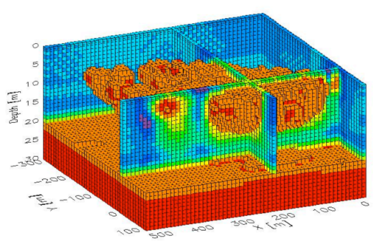

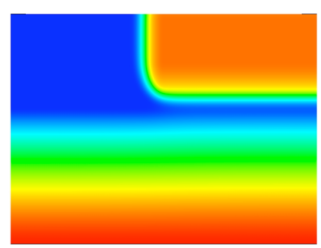

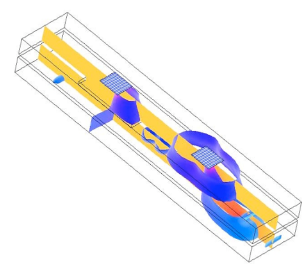

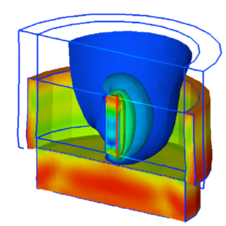

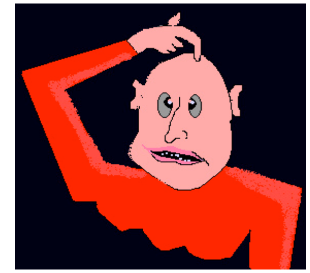

What could possibly go wrong?

- Smoothing doesn't damp error in expected ways
  - Anisotropic problems, bad aspect ratios
  - Non-elliptic operators
  - …
- Coarse operator does not
  - Large PDE coefficient
  - Stability
  - M-matrix
  - Null space
- PDE Systems
  - Constraints
- Nonsymmetry

$\varepsilon u_{xx} + u_{yy} = f$ & ***Jacobi smoothing***

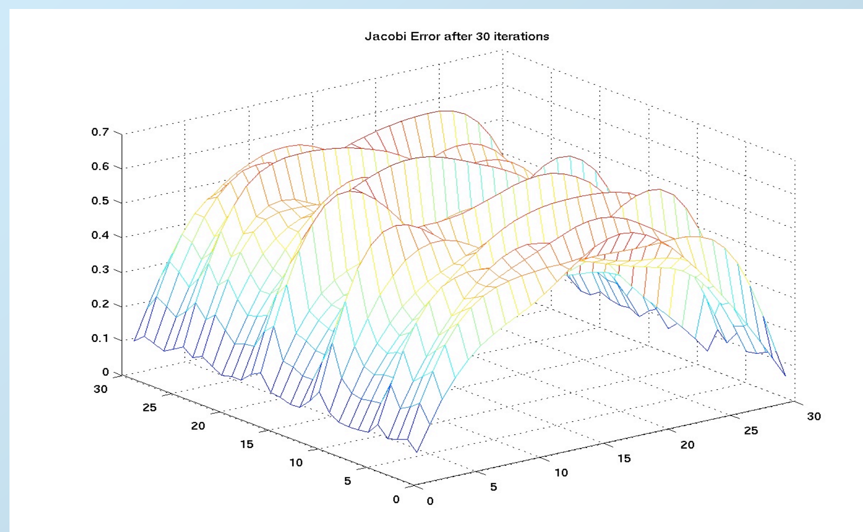

What could possibly go wrong?

- Smoothing doesn't damp error in ex
  - Anisotropic problems, bad aspe
  - Non-elliptic operators
  - …
- Coarse operator does not pre
  - Large PDE coefficient variations
  - Stability
  - M-matrix
  - Null space
- PDE Systems
  - Constraints
- Nonsymmetry

fine scale features & coarse levels

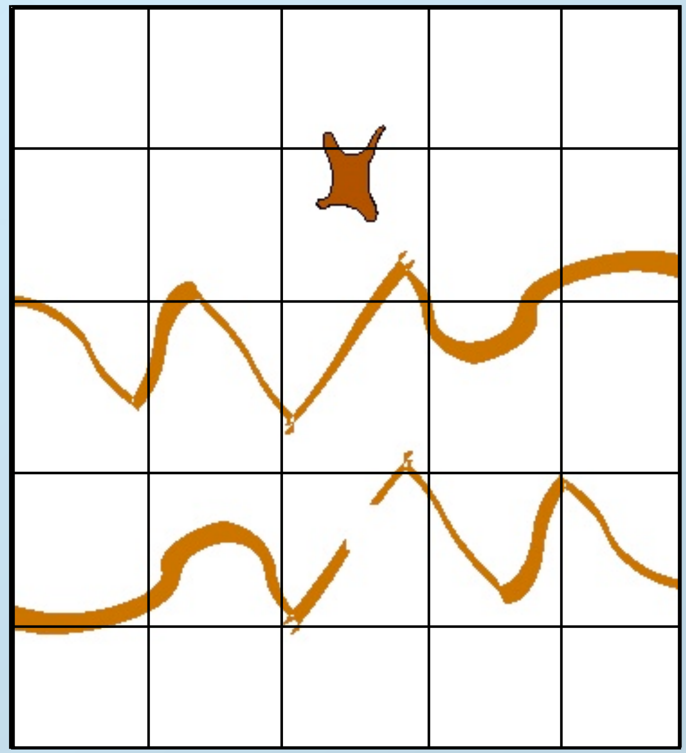

What could possibly go wrong?

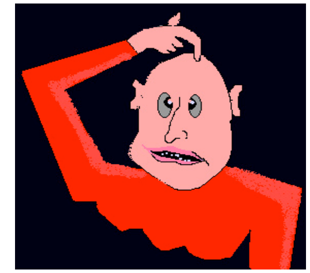

- Smoothing doesn't damp error in expected ways
  - Anisotropic problems, bad aspect ratios
  - Non-elliptic operators
  - …
- Coarse operator does not
  - Large PDE coefficient
  - Stability
  - M-matrix
  - Null space
- PDE Systems
  - Constraints
- Nonsymmetry

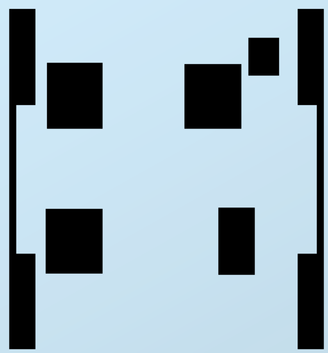

Jacobi & Gauss-Seidel smoothers divide by the diagonal???

Following Physics/ Irregular Coarsening

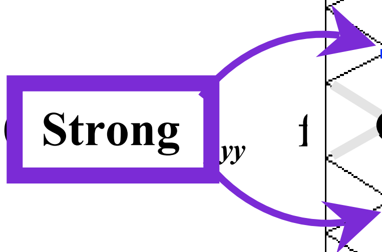

Q1 stencil f

Detection:

Special Prolongators

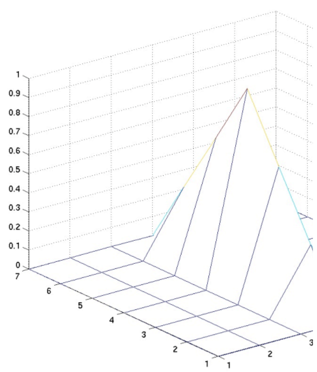

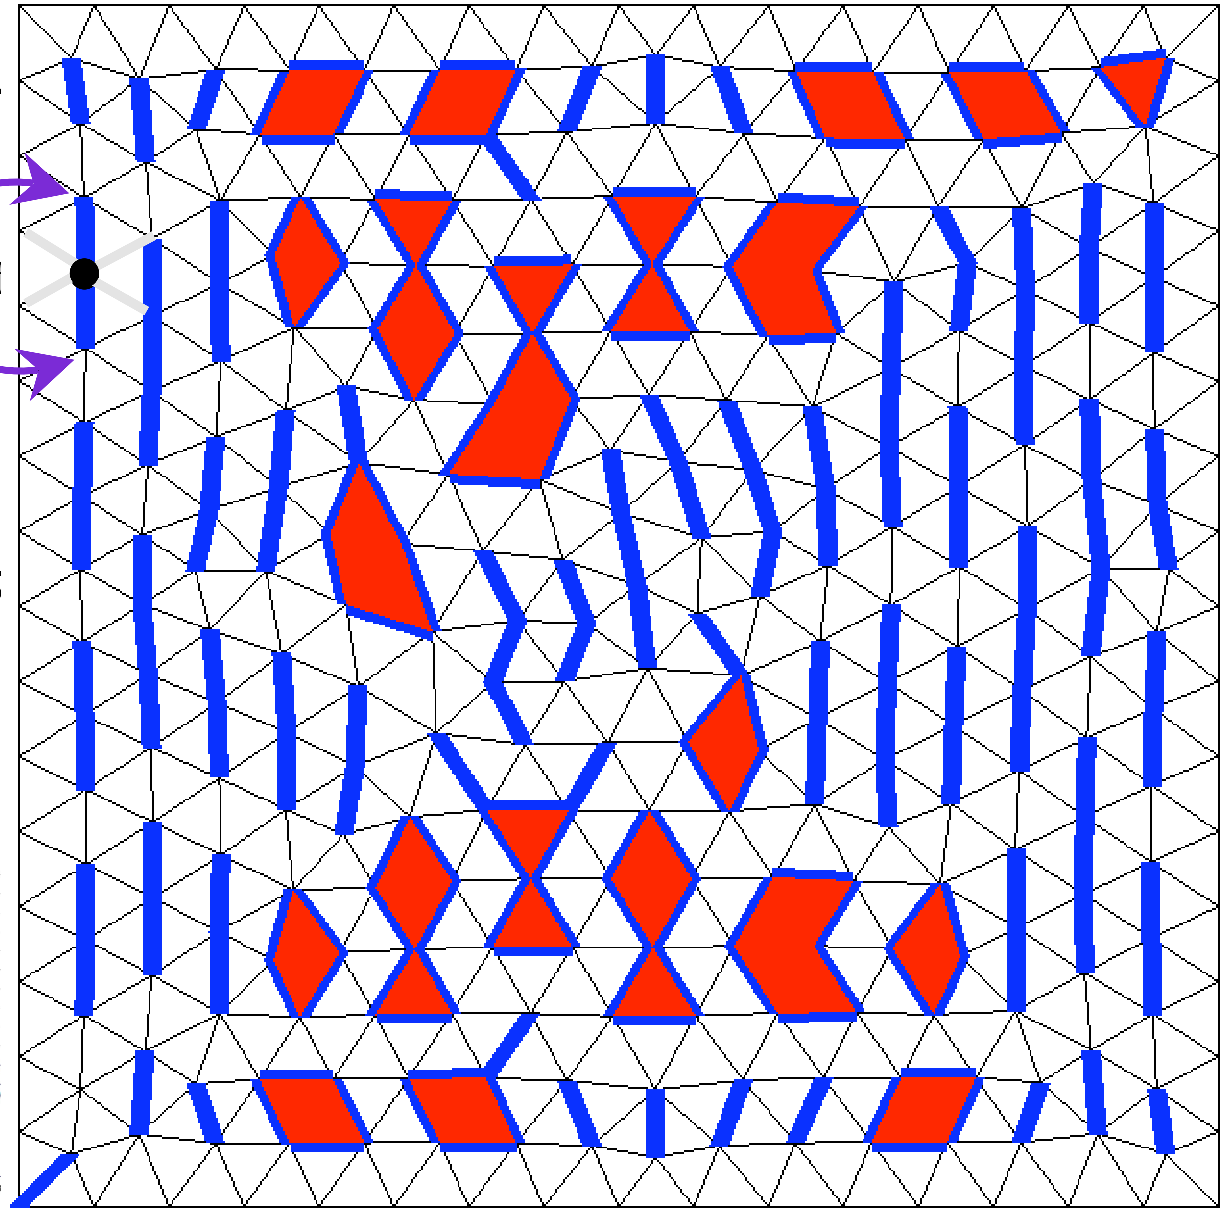

Strength of Connection

Too global !!

$$Au:\ (\varepsilon u_x)_x\ ,\ u_x(0) = 0,\ u(1) = 0$$

$$\varepsilon = \begin{cases} .001 & x \le 0.5 \\ 1.0 & x > 0.5 \end{cases}.$$

Column of

A^{-1} @ $x = .5$

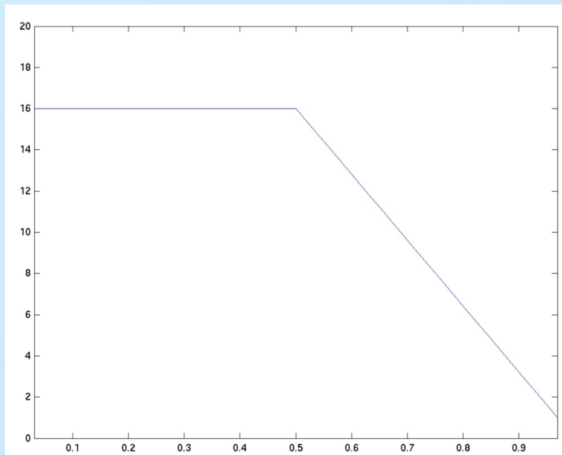

– $y_t = -A\,y$ with $y(0) = \delta_i$

– $z_t(\Delta t) \approx \Delta t\, A\, \delta_i + \delta_i$

$z(\infty) = A^{-1}\, \delta_i$

– **matrix coefs ≤ ODE ≤ matrix inverse coefs**

Local discretization errors/lack of mass matrix

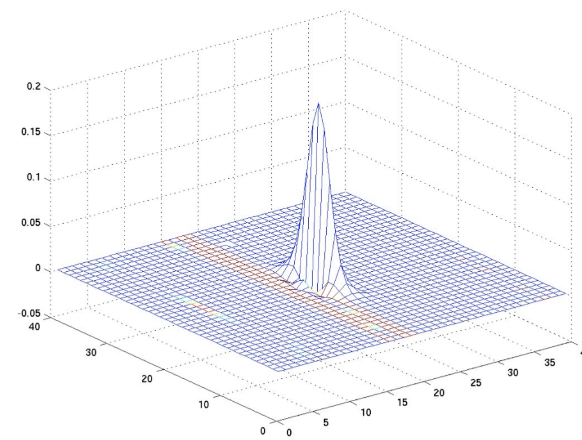

New Strength Measure

$$z = (I - \Delta t\, D^{-1} A)^k\, \delta_i$$

Δt chosen for stability & k chosen to reduce δ_i by ½

Strength ≡ How well can z be interpolated within neighborhood?

$$min_x \;||\, B\,x - z\,||_Q \;\text{ with }\; \langle B\,x - z,\, e\rangle_Q = 0$$

- B is simple prolongator
- x & z : nearest neighbor subset of i
- Scalar/piecewise constant case ⇒ $S_{ij} = \left|\frac{z_j}{z_i - z_j}\right|$

AMG convergence & strength

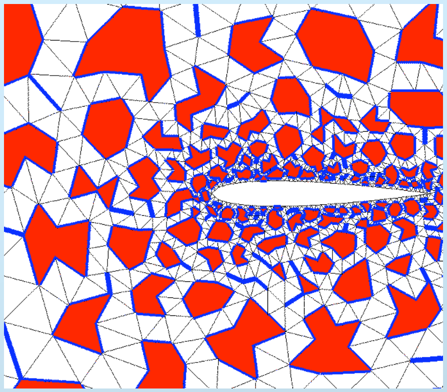

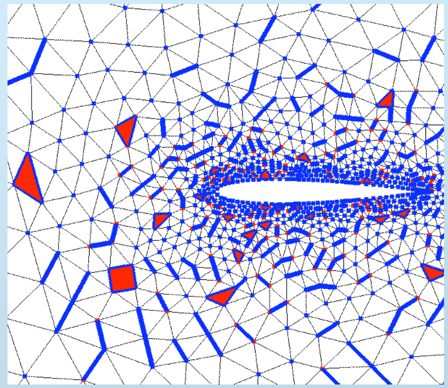

≈DOFs	coefs	ODE
20k	10	11
75k	13	12
300k	15	14

diffusion

≈DOFs	coefs	ODE
2k	33	13
14k	67	13
110k	129	18

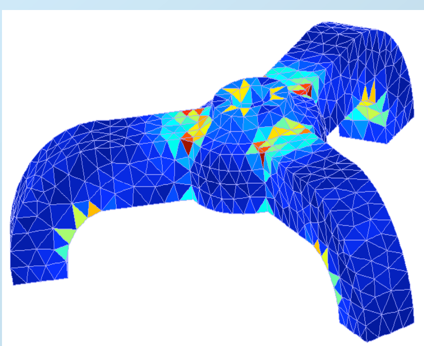

elasticity

mesh	coefs	ODE
32x32	19	16
64x64	29	22
128x128	132	23

recirculating flow

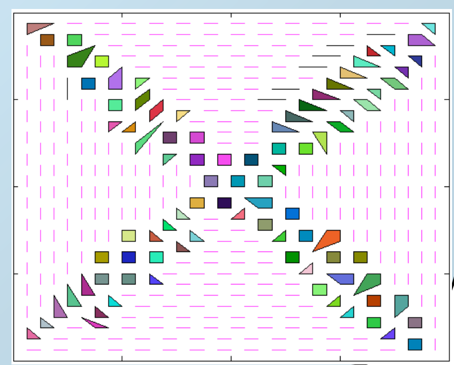

Incompressible Fluid Flow

$$\begin{pmatrix} F & B^T \\ B & 0 \end{pmatrix} \begin{pmatrix} u \\ p \end{pmatrix} = f$$

B: divergence , B^T = gradient

- Pressure Projection Methods
  - Explicit update for u (or well-conditioned solve)
  - Laplace solve for p
- Fully Implicit + AMG applied to 2x2 system
- Fully Implict + Block Preconditioners

Pressure Poisson solves

- **air/sand jet**

	Isotropic coarsen	Geometric coarsen
Its	196	22
Time (seconds)	42.95	4.24

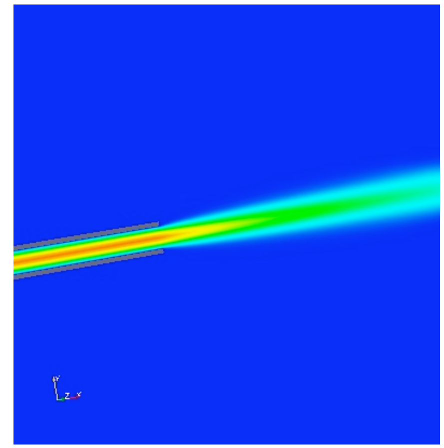

- **He Plume problem** (16384 procs)

	Isotropic coarsen	Geometric coarsen
Its	161	44
Time (seconds)	17.22	4.86

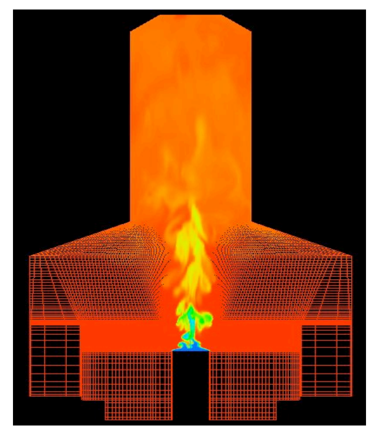

All at once …

- Group Dofs @ nodes
- Construct graph for block matrix
  - generalize strength
- Generalize P

  - Easiest : $(v^h)_{ij} = f((v^H)_{ij})$

    where v is one particular component

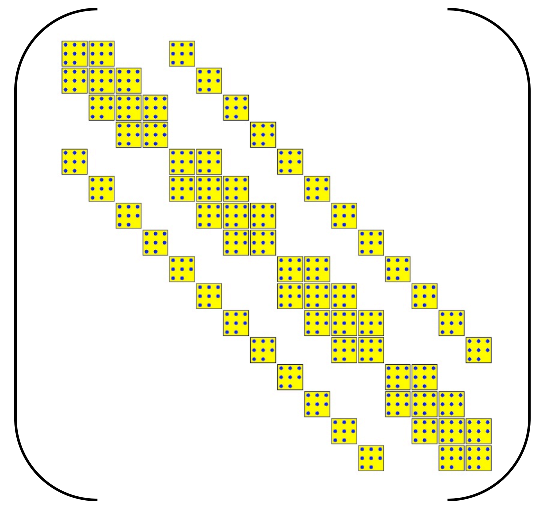

- Use heavy smoothers
  - Block Gauss-Seidel, ILU, subdomain solves

- Aggressive coarsening

Coarse level stability is a concern for Saddle Pt. Problems!!

Arbitrary Matrix Coarsening

- New: aggressive coarsening
  - graph partitioning to create larger aggregates
  - Only a few levels, cheaper RAP products

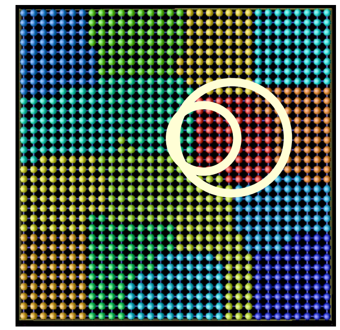

Transient LES-k

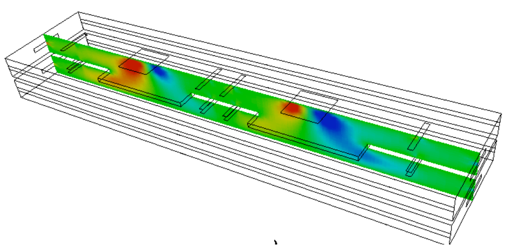

proc		DD/ILUT		3-level (GS/ILU/KLU)			
	DOFs	iter/N	sec/N	medium	coarse	iter/N	sec/N
32	2.13M	161	70	24930	290	89	47
256	16.4M	320	167	2E+05	2265	101	55
2048	129.1M	552	450	15.1M	17805	113	105

CVD Reactor

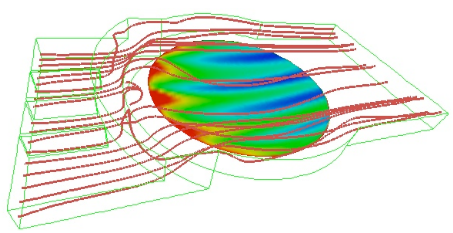

proc	1-level DD/ ILU			3-level (ILU/ILU/KLU)			
	DOFs	iter/N	sec/N	medium	coarse	iter/N	sec/N
16	636168	76	133	9888	152	38	132
128	4.84M	201	164	96416	1928	72	138
1024	37.8M	480	281	586360	9160	113	176

NPN BJT

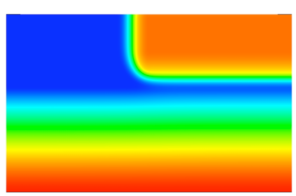

DOFs	1 level	3-level	3-level /new P
112 M	858.0	61.4	30.2

Block Preconditioners

$$\begin{bmatrix} \mathcal{F} & \mathcal{G} \\ \mathcal{G}^{T} & 0 \end{bmatrix} = \begin{bmatrix} I & 0 \\ \mathcal{G}^{T}\mathcal{F}^{-1} & I \end{bmatrix} \begin{bmatrix} \mathcal{F} & \mathcal{G} \\ 0 & \mathcal{S} \end{bmatrix} \qquad \mathcal{S} = -\mathcal{G}^{T}\mathcal{F}^{-1}\mathcal{G}$$

$$\begin{bmatrix} \mathcal{F} & \mathcal{G} \\ \mathcal{G}^{T} & 0 \end{bmatrix} \mathcal{M} = \begin{bmatrix} I & 0 \\ \mathcal{G}^{T}\mathcal{F}^{-1} & I \end{bmatrix} \quad \text{with} \quad \mathcal{M} = \begin{bmatrix} \mathcal{F} & \mathcal{G} \\ 0 & \mathcal{S} \end{bmatrix}$$

⇒ 2 iterations

$$\mathcal{M} = \begin{bmatrix} \mathcal{F} & \mathcal{G} \\ & Q \end{bmatrix}^{-1}$$

A couple of AMG cycles for $\mathcal{F}^{-1}$ & Q^{-1}

What to choose for Q^{-1}?

Commuting

Suppose

$$\mathcal{F}\,\mathcal{G} = \mathcal{G}\mathcal{F}_p$$

e.g., Stokes

$\mathcal{G} = \nabla,\ \ \mathcal{F} = \Delta,$

$\mathcal{F}_p = \Delta$

Then,

$$Q^{-1} = \mathcal{F}_p \underbrace{(\mathcal{G}^T\mathcal{G})^{-1}}_{\Delta^{-1}} \rightarrow \underbrace{(\mathcal{G}^T\mathcal{F}^{-1}\mathcal{G})}_{\mathcal{S}} Q^{-1} \approx I$$

Or $$\mathcal{F}_p = (\mathcal{G}\,\mathcal{G})^{-1}\ \mathcal{G}^T\mathcal{F}\,\mathcal{G} \ \Rightarrow\ Q^{-1} = \Delta^{-1}\mathcal{G}^T\mathcal{F}\mathcal{G}\Delta^{-1}$$

AMG cycles for Δ^{-1}

Issues:

- **Discrete vs. differential commuting**
- **Stabilization**
- **Boundary conditions**

Performance

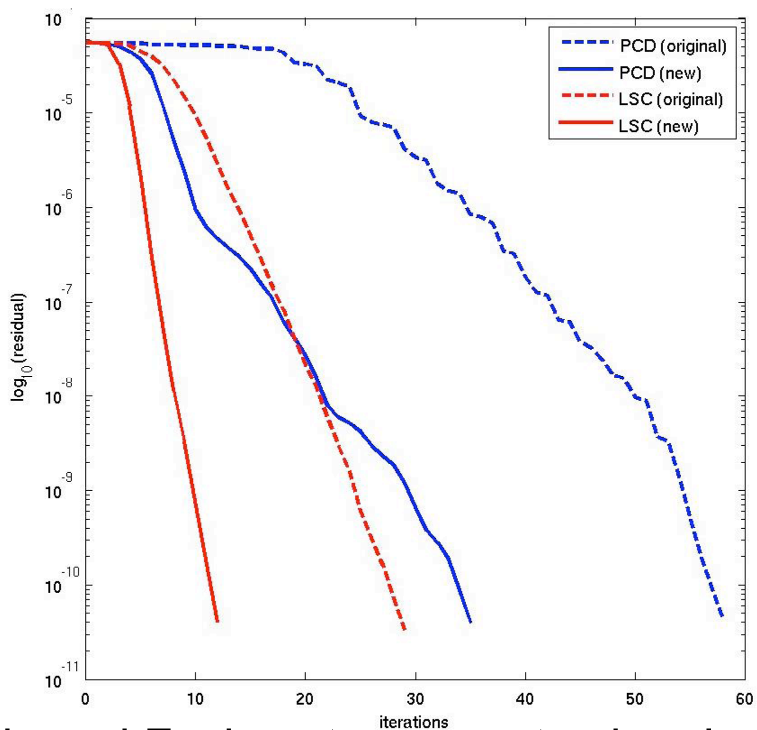

Backward Facing step , exact sub-solves

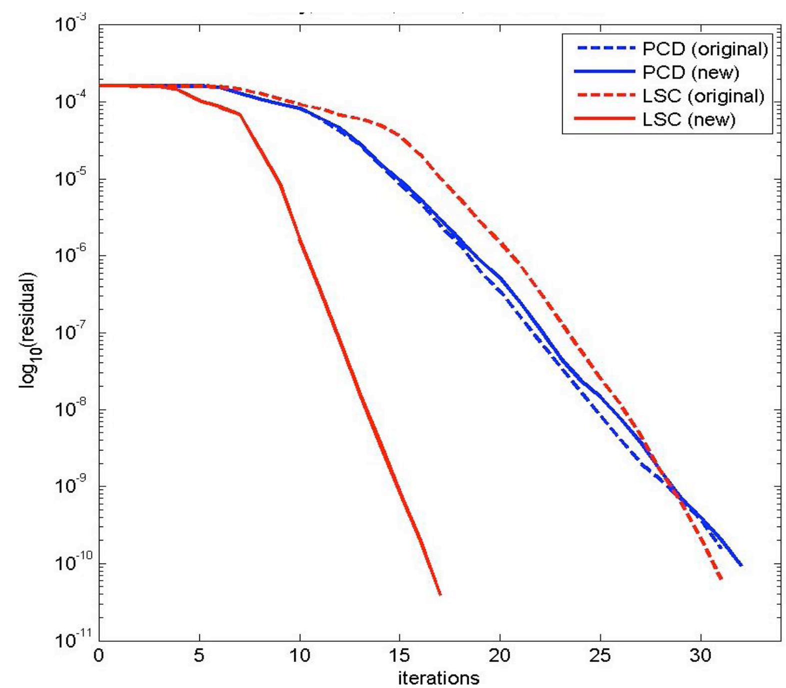

Lid driven cavity, exact sub-solves

Flow over Obstruction

AMG sweeps

Re	Mesh	1 level DD	PC-D	procs
10	2.1 M	151.2 (2004.1)	21.7 (1507.5)	8
	16.8 M	667.2 (20908.0)	24.7 (1997.7)	64
50	2.1 M	132.4 (2676.1)	38.7 (1797.2)	8
	16.8 M	637.2 (18646.0)	44.7 (2397.7)	64

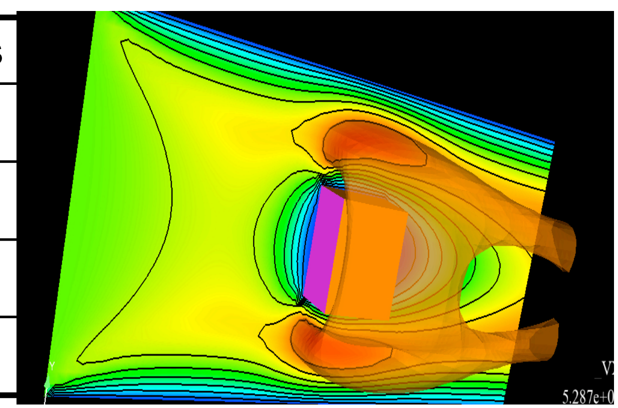

Concluding Remarks

- AMG methods continue to evolve according to application needs
- Care needs to be exercised, e.g.
  - Anisotropic
  - PDE constraints
  - Nonsymmetric, heterogeneous, etc.
- Strength of connection Important
- Incompressible flow
  - Pressure projection
  - All-at-once
  - Block preconditioners
- AMG is often effective even if not always optimal

I hope to learn more about GeoScience problems

Sandia National Laboratories