# Toward a More Robust Sparse Solver ...with some ideas on resilience and scalability

Luke Olson CS @ Illinois

#### The Point

- The state of algebraic solvers (AMG)
- Thinking about MG and resilience
- Collaboration in the Joint Lab to accelerate this

#### Sparse Problem

- complex, non-symmetric
- non-elliptic
- preconditioning necessary
- encompass a large solve time



#### Multilevel

attenuate high energy quickly with with relaxation

attenuate low energy error through coarse-grid correction





I.Determine sense of energy2.Find strength of edges3.Group edges4.Form interpolation

#### Interpolation



#### Interpolation

$$e_1 \leftarrow (I - P(P^T A P)^{-1} P^T A) Ge_0$$
  
coarse grid correction relax

$$Ge_0 \in \mathcal{R}(P) \quad \Rightarrow \quad e_1 = 0$$

interpolation should capture what relaxation misses

- P should have low energy (low A-norm or  $A^*A$ -norm)
  - 1. determine sparsity pattern
  - 2. minimize energy column-wise (parallel)

# Interpolation (basic)

• Set the sparsity pattern from aggregation





#### Toward General Interpolation

- Want P so that  $u_{low} \in \mathcal{R}(P)$
- Grow sparsity pattern along strong edges:

 $S^k P_{tent}$ 

• Minimize residual:

 $AP_j = 0$  for each column j

Constraint the minimization with

$$Pu_{low}^c = u_{low}$$

#### Toward General Interpolation

• Hermitian (and positive definite): use CG

$$AP_j = 0 \Leftrightarrow \min \|P_j\|_A$$
$$R = P^*$$

• Non-Hermitian: use GMRES

$$AP_j = 0 \Leftrightarrow \min \|P_j\|_{A^*A}$$
$$A^*R_j^* = 0 \Leftrightarrow \min \|R_j^*\|_{AA^*}$$

- Range of interpolation/restriction target "right"/"/" left" low-energy
- Cost is comparable to that of standard smoothing

### Example: recirculating flow



h	std.	opt.	
1/64	>150	24	
1/128	>150	28	
1/256	>150	33	
1/512	>150	33	

iterations

#### What interpolation gives us

- Sense of optimality (or a sense of violating optimality)  $u \in \mathcal{R}(P) \quad \Rightarrow \quad \|u\|_A \ll 1$
- Energetic stability  $\|P(P^T P)^{-1}P^T\|_A$

uniformly bounded gives convergence

- Galerkin coarsening  $A_{\text{coarse}} = P^T A_{\text{fine}} P$
- High growth in non-zeros per row:
  - high complexity
  - high data dependence between compute nodes



- Galerkin coarsening  $A_{\text{coarse}} = P^T A_{\text{fine}} P$
- High growth in non-zeros per row:
  - high complexity
  - high data dependence between compute nodes



- Basic idea\* is to take  $A_{lgebraic}$  \* Falgout, Sch Algebraic Mult  $A_{
m c}=P^TA_{
m f}P$ 

\* Falgout, Schroder. Non-Galerkin Coarse Grids for Algebraic Multigrid. LLNL 2013.

$$\|I - \hat{A}_c^{-1} A_c\| \to \min$$

- $\hat{A}_c$  has improved sparsity
- I. Determine a sparsity pattern

2. Enforce collapsing to maintain spectral equivalence



potential: lower corruption, few (potential) faults in MPI

# Why think about MG resilience

- Large amount of time spent in solving Ax = b
  - I. opportunity for silent detection
  - 2. opportunity for less aggressive recovery
  - 3. refocus on convergence
- A good solver will attenuate some errors, recognize others
- Natural hierarchy → micro-checkpointing

# Trying a fault



- Jon Calhoun's fault injection framework
  - LLVM IR bitcode
  - fault model to inject in arithmetic, pointers, control, etc

# Trying a fault



- Example: residual corruption in bit m
- 2D Poisson
- results in an arithmetic error
- other corruptions possible
  - pointer errors most common, but easily manifest as a SEG

### Trying a fault



bit injected into residual in the second iteration

### What happens

- Error propagates with the mat-vec Ax
- Coarse levels intensify fault (relative to amount of data)
  - less likely and cheaper to detect
- if a fault is detected and it is severe then recover, adjust
- Need accurate measure of **silent** errors
  - handle persistent errors
  - recover from transient ones

#### measures

- energy metric (expensive, accurate):  $\eta = \langle Ax, x \rangle - 2 \langle b, x \rangle$
- residual metric (cheap,, bound):  $\langle r,r \rangle$



#### measures

Metric	Probability of encounter		level	overhead
Convergence	94%	-	0	18%
segfault	24%	-	I	6%
energy check	21%			
residual check	3%		2	< %
loop assertion	6%		3	11%

- I 00 tests
- aniso diff.  $40k \times 40k$
- injection rate ~ I-2 per solve

#### measures

- How do we effectively measure stability locally?
- Do the measures capture the right silent errors (efficiently)?

# Some directions

- Fault model to match data
  - potential collaboration: Cappello (ANL/INRIA)





- Redundancy and recovery
  - potential collaboration: Brown et al (ANL)
- Algorithm recovery with micro checkpointing
  - potential collaboration: Cappello (ANL/INRIA)
- Extend interpolation and non-Galerkin methodology to DD-type
  - <u>potential collaboration</u>: Brown (ANL), Nataf (LJLL), Grigori (INRIA)

- this is work with
  - Jon Calhoun (fault injection, analysis)
  - Amanda Bienz (non-Galerkin scaling)
  - Andrew Reisner (redundancy models)
  - others @ Illinois, LLNL, ANL