

Recent Advances in System Solvers

Tom Manteuffel

University of Colorado at Boulder

Office of Advanced Scientific Computing Research

2007 Applied Mathematics Principle Investigators Meeting

May 22-24, 2007



- Introductory Remarks
- Linear Solvers
 - Context
 - MG/AMG
 - Adaptive AMG
- Nonlinear Systems
 - Context
 - Nested Iteration/Newton/Krylov/AMG
- Summary



- Introductory Remarks
- Linear Solvers
 - Context
 - MG/AMG
 - Adaptive AMG
- Summary



Theme song:

Everything is linear,



Theme song:

Everything is linear,

...in its own way



Theme song:

Everything is linear,

...in its own way

Take a simple Newton step,



Theme song:

Everything is linear,

...in its own way

Take a simple Newton step,

and iterate from 1 to k .

Sung to the tune of "Everything is Beautiful" by Ray Stevens



Krylov Methods \Leftrightarrow Polynomial Methods

$A\underline{x} = \underline{b}$	\underline{x}_0	initial guess
	\underline{x}_i	iterate
	$\underline{e}_i = \underline{x} - \underline{x}_i$	error
	$\underline{r}_i = \underline{b} - A\underline{x}_i$	residual



Krylov Methods \Leftrightarrow Polynomial Methods

$A\underline{x} = \underline{b}$	\underline{x}_0	initial guess
	\underline{x}_i	iterate
	$\underline{e}_i = \underline{x} - \underline{x}_i$	error
	$\underline{r}_i = \underline{b} - A\underline{x}_i$	residual

Error Equation: $p_i(0) = 1.0$

$$\underline{e}_i = p_i(A)e_0$$

$$\underline{r}_i = p_i(A)r_0$$



Jordan Decomposition

$$A = SJS^{-1}$$

Error Bound

$$\|e_i\| \leq \|S\| \|S^{-1}\| \|p_i(J)\|$$




Polynomial Methods: Error Bounds

Jordan Decomposition

$$A = SJS^{-1}$$

Error Bound

$$\|e_i\| \leq \|S\| \|S^{-1}\| \|p_i(J)\|$$




Polynomial Methods: Error Bounds

Jordan Decomposition

$$A = SJS^{-1}$$

Error Bound

$$\|e_i\| \leq \|S\| \|S^{-1}\| \|p_i(J)\|$$



If condition of A is large, it is hard to make a polynomial small on all of the eigenvalues and still have $p_i(0) = 1$.



$$CA\underline{x} = C\underline{b}$$

C – Any linear process

Choose C so that system with CA is easier to solve in some sense

For example, condition of CA is much smaller than that of A



$$CA\underline{x} = C\underline{b}$$

C – Any linear process

Examples: $A = L + D + U$

- $C = D^{-1}$,

Jacobi Preconditioning



$$CA\underline{x} = C\underline{b}$$

C – Any linear process

Examples: $A = L + D + U$

- $C = D^{-1}$,
- $C = (D + L)^{-1}$

Jacobi Preconditioning

Gauss/Seidel



$$CA\underline{x} = C\underline{b}$$

C – Any linear process

Examples: $A = L + D + U$

- $C = D^{-1}$,
- $C = (D + L)^{-1}$
- $C = ((D + L)D^{-1}(D + U))^{-1}$

Jacobi Preconditioning

Gauss/Seidel

Symmetric Gauss/Seidel



$$CA\underline{x} = C\underline{b}$$

C – Any linear process

Examples: $A = L + D + U$

- $C = D^{-1}$,
- $C = (D + L)^{-1}$
- $C = ((D + L)D^{-1}(D + U))^{-1}$
- $C = ((\hat{D} + L)\hat{D}^{-1}(\hat{D} + U))^{-1}$

Jacobi Preconditioning

Gauss/Seidel

Symmetric Gauss/Seidel

Incomplete Factorization



$$CA\underline{x} = C\underline{b}$$

C – Any linear process

Examples: $A = L + D + U$

- $C = D^{-1}$,
- $C = (D + L)^{-1}$
- $C = ((D + L)D^{-1}(D + U))^{-1}$
- $C = ((\hat{D} + L)\hat{D}^{-1}(\hat{D} + U))^{-1}$
- $C = A^*$

Jacobi Preconditioning

Gauss/Seidel

Symmetric Gauss/Seidel

Incomplete Factorization

Normal Equations



$$CA\underline{x} = C\underline{b}$$

C – Any linear process

Examples: $A = L + D + U$

- $C = D^{-1}$,
- $C = (D + L)^{-1}$
- $C = ((D + L)D^{-1}(D + U))^{-1}$
- $C = ((\hat{D} + L)\hat{D}^{-1}(\hat{D} + U))^{-1}$
- $C = A^*$
- $C =$ Multigrid V-cycle

Jacobi Preconditioning

Gauss/Seidel

Symmetric Gauss/Seidel

Incomplete Factorization

Normal Equations

PCG-MG



Preconditioning \Leftrightarrow Matrix Splitting



Preconditioning \Leftrightarrow Matrix Splitting

- Any matrix splitting can be used as a preconditioning



Preconditioning \Leftrightarrow Matrix Splitting

- Any matrix splitting can be used as a preconditioning
- Any linear process, C , can be used as a preconditioning



Preconditioning \Leftrightarrow Matrix Splitting

- Any matrix splitting can be used as a preconditioning
- Any linear process, C , can be used as a preconditioning
- Any preconditioning can accelerated by a polynomial method



- In general, if A comes from a PDE, optimal preconditioning requires a Multilevel algorithm
 - Optimal \Rightarrow condition of CA is independent of the mesh



- In general, if A comes from a PDE, optimal preconditioning requires a Multilevel algorithm
 - Optimal \Rightarrow condition of CA is independent of the mesh
 - Optimal \Rightarrow work grows linearly with the problem size



- In general, if A comes from a PDE, optimal preconditioning requires a Multilevel algorithm
 - Optimal \Rightarrow condition of CA is independent of the mesh
 - Optimal \Rightarrow work grows linearly with the problem size

If you want to solve a problem with billions of unknowns on 128,000 processors, you will need a multilevel algorithm somewhere.



Recent Developments

- A lot of recent activity in multilevel algorithms



Recent Developments

- A lot of recent activity in multilevel algorithms
- Especially in Algebraic Multigrid (AMG)




Recent Developments

- A lot of recent activity in multilevel algorithms
- Especially in Algebraic Multigrid (AMG)
 - More robust
 - More effective



Recent Developments

- A lot of recent activity in multilevel algorithms
- Especially in Algebraic Multigrid (AMG)
 - More robust
 - More effective



DYNA3D



Basic Components

- Simple relaxation or smoothing
 - Usually a matrix splitting or simple preconditioned one-step like damped Jacobi, Gauss/Sedel or block Gauss/Seidel
 - Resolves error in direction of eigenvectors with large eigenvalues
- Coarse-grid correction
 - Lower dimensional or simpler problem
 - Resolves error left by relaxation
- Recursion
 - Coarse-grid problem is solved by multigrid



Example: discrete forms of second-order elliptic operators


$$-\nabla \cdot A \nabla u + cu = f$$

- Large eigenvalues are associated with high frequency eigenvectors
- Simple iterative methods leave error geometrically smooth
- Coarse grid problem is a version of fine grid problem




Multigrid: example

Given Error




Multigrid: example

Given Error




Relax




Multigrid: example


Given Error



Relax

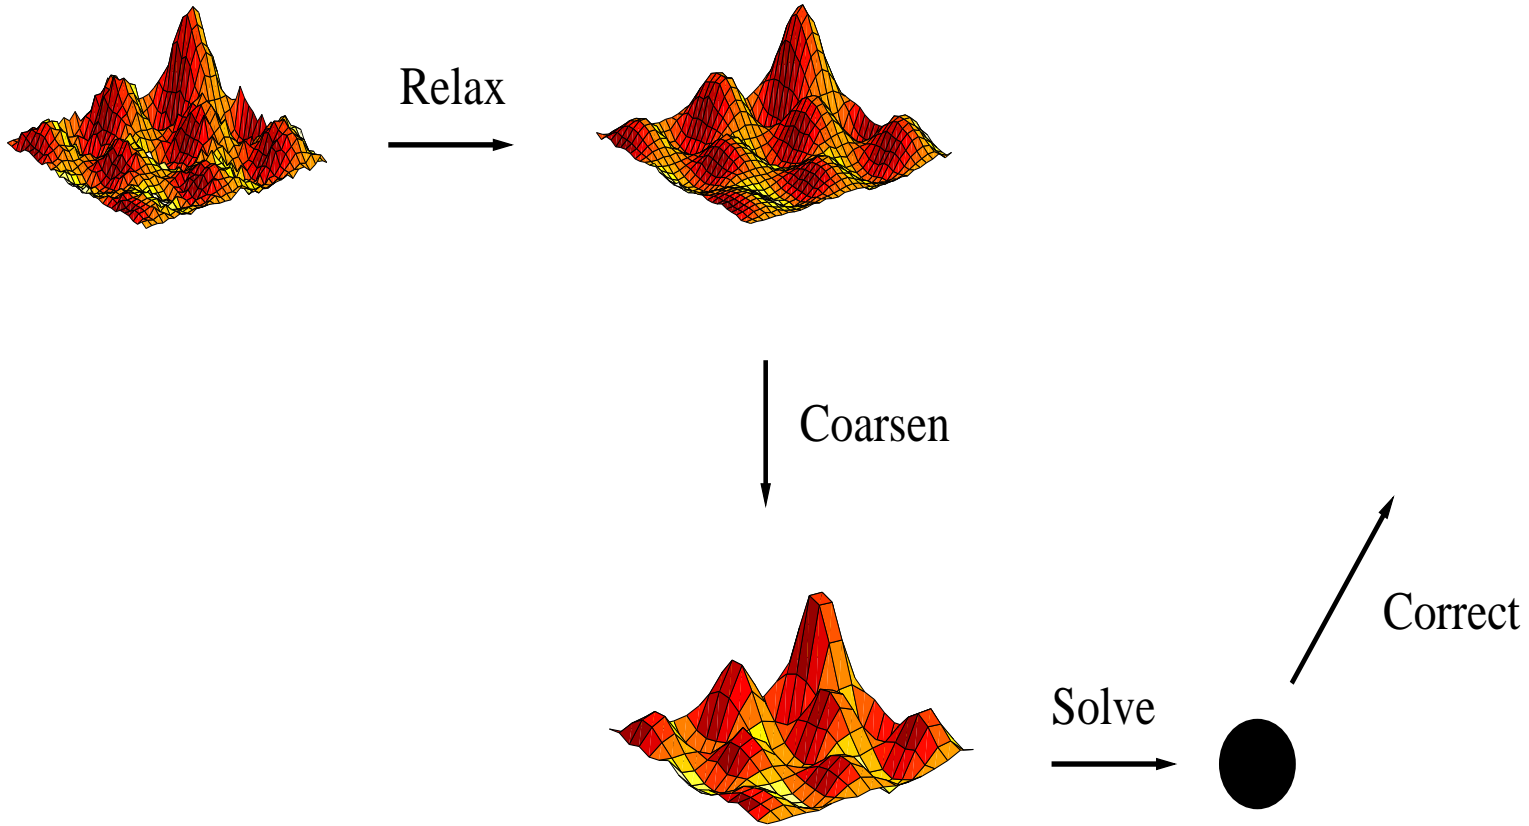


Coarsen

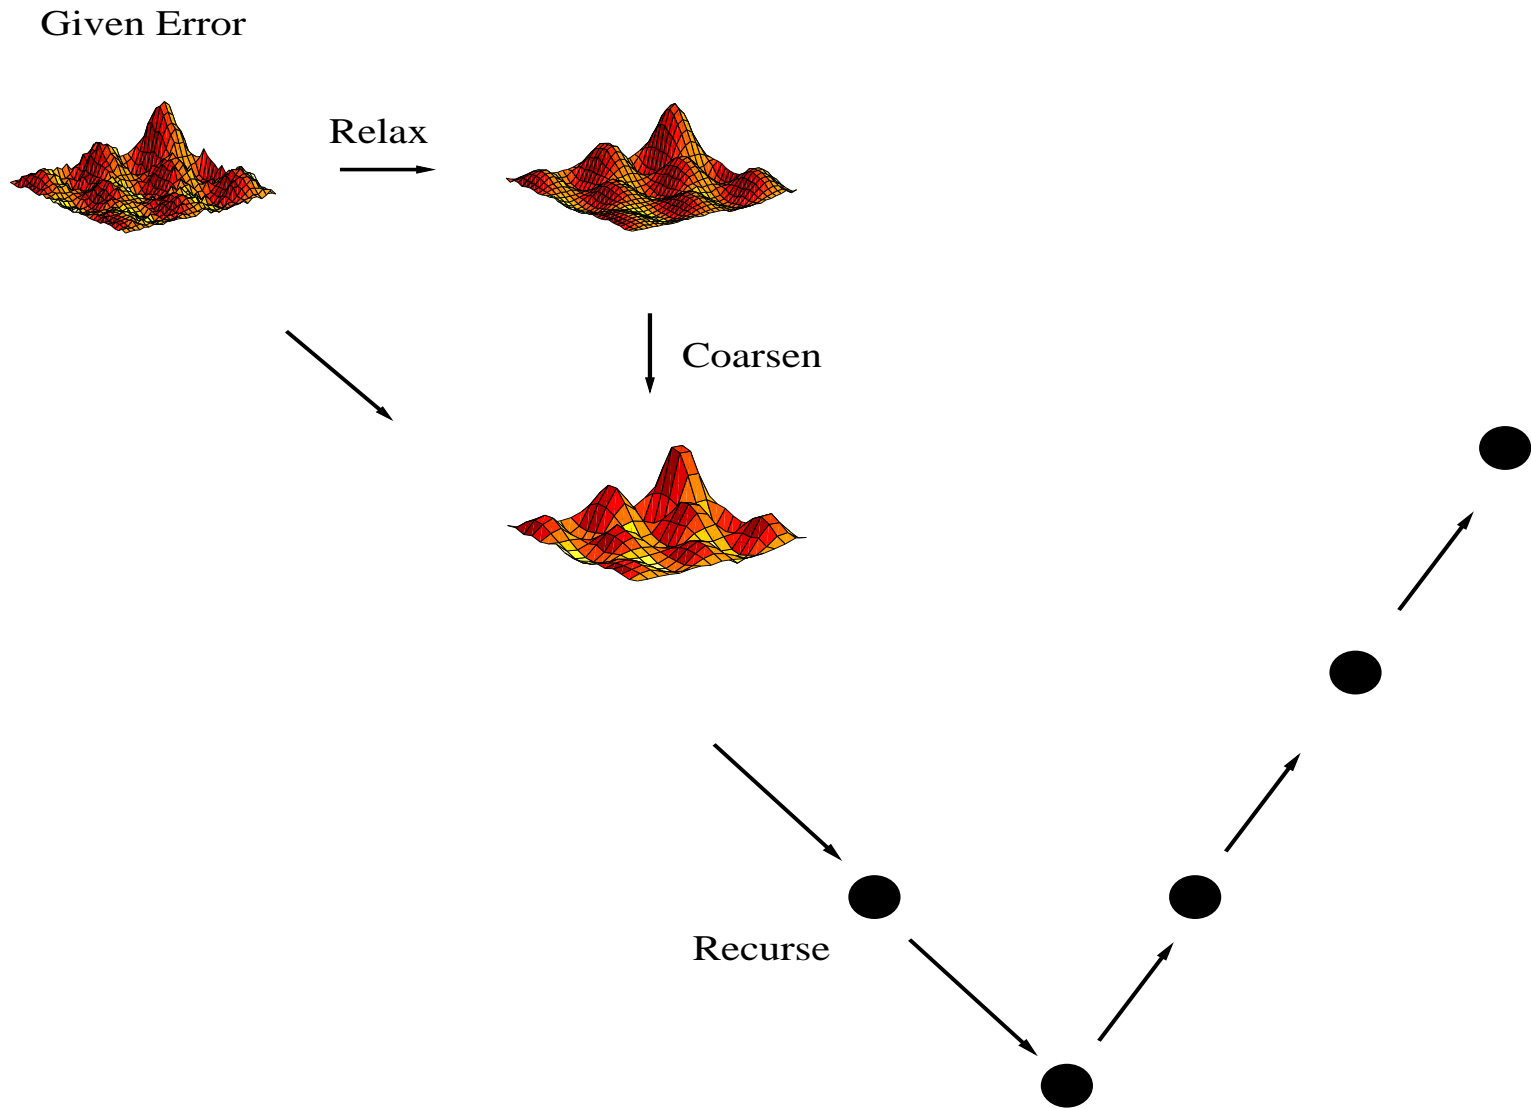


Multigrid: example

Given Error



Multigrid: example



Basic Components (again)

Multigrid algorithm is determined by

- Relaxation
- Interpolation from coarse grid to fine grid (P)
- Restriction from fine grid to coarse grid (R)
- Coarse-grid operator (A_c)



Basic Components (again)

Multigrid algorithm is determined by

- Relaxation
- Interpolation from coarse grid to fine grid (P)
- Restriction from fine grid to coarse grid (R)
- Coarse-grid operator (A_c)

In variational MG,

$$A_c = RA_fP$$



Multigrid Flavors

- Geometric Multigrid (GMG)
- Algebraic Multigrid (AMG)



Geometric multigrid (GMG)

- Coarse-grid problem is geometrically determined
- It is usually a smaller version of the fine grid problem
- Interpolation and restriction usually determined by the operator



Algebraic Multigrid (AMG)

- Directly address the matrix A without presumed knowledge of
 - Geometry
 - Operator



Algebraic Multigrid (AMG)

- Directly address the matrix A without presumed knowledge of
- Assume simple relaxation
 - For example, Damped Jacobi, Gauss/Seidel



Algebraic Multigrid (AMG)

- Directly address the matrix A without presumed knowledge of
- Assume simple relaxation
- Coarse-grid problem is chosen to resolve the “Algebraically smooth” error
 - Defined to be the error that relaxation does not resolve



Algebraic Multigrid (AMG)

- Directly address the matrix A without presumed knowledge of
- Assume simple relaxation
- Coarse-grid problem is chosen to resolve the “Algebraically smooth” error
- Work focuses on selection of a coarse grid and the intergrid transfer operators (R and P)



Algebraic Multigrid (AMG)

- Directly address the matrix A without presumed knowledge of
- Assume simple relaxation
- Coarse-grid problem is chosen to resolve the “Algebraically smooth” error
- Work focuses on selection of a coarse grid and the intergrid transfer operators (R and P)
- The coarse-grid operator is formed variationally ($A_c = RA_fP$)



Multigrid Flavors

- Geometric Multigrid (GMG)
- Algebraic Multigrid (AMG)



Multigrid Flavors

- Geometric Multigrid (GMG)
- Algebraic Multigrid (AMG)
 - AMG
 - Smoothed Aggregation (SA)
 - Adaptive AMG (αAMG , αSA)



AMG is characterized by choice of the Coarse Grid, Interpolation, P , and Restriction, R .

For simplification, assume A symmetric and $R = P^t$



Divide degrees of freedom into the Coarse DOF and Fine DOF

$$A = \begin{bmatrix} A_{ff} & A_{fc} \\ A_{cf} & A_{cc} \end{bmatrix}$$



Divide degrees of freedom into the Coarse DOF and Fine DOF

$$A = \begin{bmatrix} A_{ff} & A_{fc} \\ A_{cf} & A_{cc} \end{bmatrix}$$

After relaxation error is algebraically smooth

$$\|A\underline{e}\| \ll \|\underline{e}\|$$

error in direction of large eigenvalues has been reduced



Divide degrees of freedom into the Coarse DOF and Fine DOF

$$A = \begin{bmatrix} A_{ff} & A_{fc} \\ A_{cf} & A_{cc} \end{bmatrix}$$

After relaxation error is algebraically smooth

$$\|A\underline{e}\| \ll \|\underline{e}\|$$

error in direction of large eigenvalues has been reduced

$$\begin{bmatrix} A_{ff} & A_{fc} \\ A_{cf} & A_{cc} \end{bmatrix} \begin{pmatrix} \underline{e}_f \\ \underline{e}_c \end{pmatrix} \simeq \begin{pmatrix} \underline{0} \\ \underline{0} \end{pmatrix}$$



$$\begin{bmatrix} A_{ff} & A_{fc} \\ A_{cf} & A_{cc} \end{bmatrix} \begin{pmatrix} \underline{e}_f \\ \underline{e}_c \end{pmatrix} \approx \begin{pmatrix} \underline{0} \\ \underline{0} \end{pmatrix}$$



$$\begin{bmatrix} A_{ff} & A_{fc} \\ A_{cf} & A_{cc} \end{bmatrix} \begin{pmatrix} \underline{e}_f \\ \underline{e}_c \end{pmatrix} \approx \begin{pmatrix} \underline{0} \\ \underline{0} \end{pmatrix}$$

$$A_{ff}\underline{e}_f + A_{fc}\underline{e}_c = 0$$



$$\begin{bmatrix} A_{ff} & A_{fc} \\ A_{cf} & A_{cc} \end{bmatrix} \begin{pmatrix} \underline{e}_f \\ \underline{e}_c \end{pmatrix} \approx \begin{pmatrix} \underline{0} \\ \underline{0} \end{pmatrix}$$

$$A_{ff}\underline{e}_f + A_{fc}\underline{e}_c = 0$$

$$\underline{e}_f = -A_{ff}^{-1}A_{fc}\underline{e}_c$$



$$\begin{bmatrix} A_{ff} & A_{fc} \\ A_{cf} & A_{cc} \end{bmatrix} \begin{pmatrix} \underline{e}_f \\ \underline{e}_c \end{pmatrix} \simeq \begin{pmatrix} \underline{0} \\ \underline{0} \end{pmatrix}$$

$$A_{ff}\underline{e}_f + A_{fc}\underline{e}_c = 0$$

$$\underline{e}_f = -A_{ff}^{-1}A_{fc}\underline{e}_c$$

Perfect Interpolation

$$\begin{pmatrix} \underline{e}_f \\ \underline{e}_c \end{pmatrix} = \begin{bmatrix} -A_{ff}^{-1}A_{fc} \\ I \end{bmatrix} \underline{e}_c = P\underline{e}_c$$



Perfect Interpolation

$$\begin{pmatrix} \underline{e}_f \\ \underline{e}_c \end{pmatrix} = \begin{bmatrix} -A_{ff}^{-1}A_{fc} \\ I \end{bmatrix} \underline{e}_c = P\underline{e}_c$$



Perfect Interpolation

$$\begin{pmatrix} \underline{e}_f \\ \underline{e}_c \end{pmatrix} = \begin{bmatrix} -A_{ff}^{-1}A_{fc} \\ I \end{bmatrix} \underline{e}_c = P\underline{e}_c$$

After relaxation

$$\begin{aligned} AP\underline{e}_c &= \underline{r} \\ P^t AP\underline{e}_c &= P^t \underline{r} \\ A_c \underline{e}_c &= \underline{r}_c \end{aligned}$$



Perfect Interpolation

$$\begin{pmatrix} \underline{e}_f \\ \underline{e}_c \end{pmatrix} = \begin{bmatrix} -A_{ff}^{-1}A_{fc} \\ I \end{bmatrix} \underline{e}_c = P\underline{e}_c$$

After relaxation

$$\begin{aligned} AP\underline{e}_c &= \underline{r} \\ P^t AP\underline{e}_c &= P^t \underline{r} \\ A_c \underline{e}_c &= \underline{r}_c \end{aligned}$$

A_c is the Schur Complement

$$A_c = A_{cc} - A_{cf}A_{ff}^{-1}A_{fc}$$

A_c is Dense



Problem: A_{ff}^{-1} is Dense $\Rightarrow A_c$ is Dense



Problem: A_{ff}^{-1} is Dense $\Rightarrow A_c$ is Dense

Solution: Sparse Approximation of A_{ff}^{-1}

- $A_{ff}^{-1} \rightarrow D_{ff}^{-1}$ Diagonal of A_{ff}
- $A_{ff}^{-1} \rightarrow \hat{D}_{ff}^{-1}$ Lumped Diagonal of A_{ff}
- $A_{ff}^{-1} \rightarrow C_{ff}$ Sparse approximate inverse of A_{ff}



Problem: A_{ff}^{-1} is Dense $\Rightarrow A_c$ is Dense

Solution: Sparse Approximation of A_{ff}^{-1}

- $A_{ff}^{-1} \rightarrow D_{ff}^{-1}$ Diagonal of A_{ff}
- $A_{ff}^{-1} \rightarrow \hat{D}_{ff}^{-1}$ Lumped Diagonal of A_{ff}
- $A_{ff}^{-1} \rightarrow C_{ff}$ Sparse approximate inverse of A_{ff}

For example: simple iteration on $A_{ff} = D_{ff} - B_{ff}$

$$A_{ff}^{-1} \rightarrow (I + D_{ff}^{-1} B_{ff}) D_{ff}^{-1}$$

Iterated Interpolation, Long Range Interpolation, Compatible Relaxation



Problem: A_{ff}^{-1} is Dense $\Rightarrow A_c$ is Dense

Solution: Sparse Approximation of A_{ff}^{-1}

- $A_{ff}^{-1} \rightarrow D_{ff}^{-1}$ Diagonal of A_{ff}
- $A_{ff}^{-1} \rightarrow \hat{D}_{ff}^{-1}$ Lumped Diagonal of A_{ff}
- $A_{ff}^{-1} \rightarrow C_{ff}$ Sparse approximate inverse of A_{ff}

Are any of these any good?



Weak Approximation Property

Interpolation must approximate an eigenvector up to the same accuracy as the size of the corresponding eigenvalue



Weak Approximation Property

Interpolation must approximate an eigenvector up to the same accuracy as the size of the corresponding eigenvalue

Weak approximation property: there exists constant C

$$M(P, \underline{u}) := \min_{\underline{v}} \frac{\|\underline{u} - P\underline{v}\|^2}{\langle A\underline{u}, \underline{u} \rangle} \leq \frac{C}{\|A\|}$$



Weak Approximation Property

Interpolation must approximate an eigenvector up to the same accuracy as the size of the corresponding eigenvalue

Weak approximation property: there exists constant C

$$M(P, \underline{u}) := \min_{\underline{v}} \frac{\|\underline{u} - P\underline{v}\|^2}{\langle A\underline{u}, \underline{u} \rangle} \leq \frac{C}{\|A\|}$$

Two-grid Convergence Factor

$$\rho \leq 1 - O\left(\frac{1}{C}\right)$$

Measure can be enforced locally



Attempt to identify connections between unknowns that are important



Attempt to identify connections between unknowns that are important

Strength of Connection: Original definition: i is strongly depends on the set

$$S_i := \{j : |a_{ij}| \geq \theta \max_{k \neq i} |a_{ik}|\}$$

for some parameter θ . (e.g. $\theta = .25$)



Attempt to identify connections between unknowns that are important

Strength of Connection: Original definition: i is strongly depends on the set

$$S_i := \{j : |a_{ij}| \geq \theta \max_{k \neq i} |a_{ik}|\}$$

for some parameter θ . (e.g. $\theta = .25$)

New, more general, definitions of strength derived from local approximation of A^{-1}



Attempt to identify connections between unknowns that are important

Strength of Connection: Original definition: i is strongly depends on the set

$$S_i := \{j : |a_{ij}| \geq \theta \max_{k \neq i} |a_{ik}|\}$$

for some parameter θ . (e.g. $\theta = .25$)

Strength of connection fundamental in choosing the coarse grid



AMG Alphabet Soup

AMG

Classical AMG (84)

SA

Soothed Aggregation (96)

AMGe

finite element AMG (01)

AMG ℓ

element free AMGe (02)

ρ AMGe

spectral AMGe (03)



AMG Alphabet Soup

AMG	Classical AMG (84)
SA	Soothed Aggregation (96)
AMGe	finite element AMG (01)
AMG ϵ	element free AMGe (02)
ρ AMGe	spectral AMGe (03)

Adaptive Algorithms

AMG	adaptive AMG (84)
BAMG	Bootstrap AMG (01)
α SA	adaptive Soothed Aggregation (04)
CR	Compatible Relaxation (04)
α AMG	adaptive AMG (06)
α AMGr	adaptive AMGr (06)



- Classical or RS - AMG (84)



- Classical or RS - AMG (84)
 - Developed by Brandt/McCormick/Ruge (84)
 - Implemented by Ruge/Stuben (85)



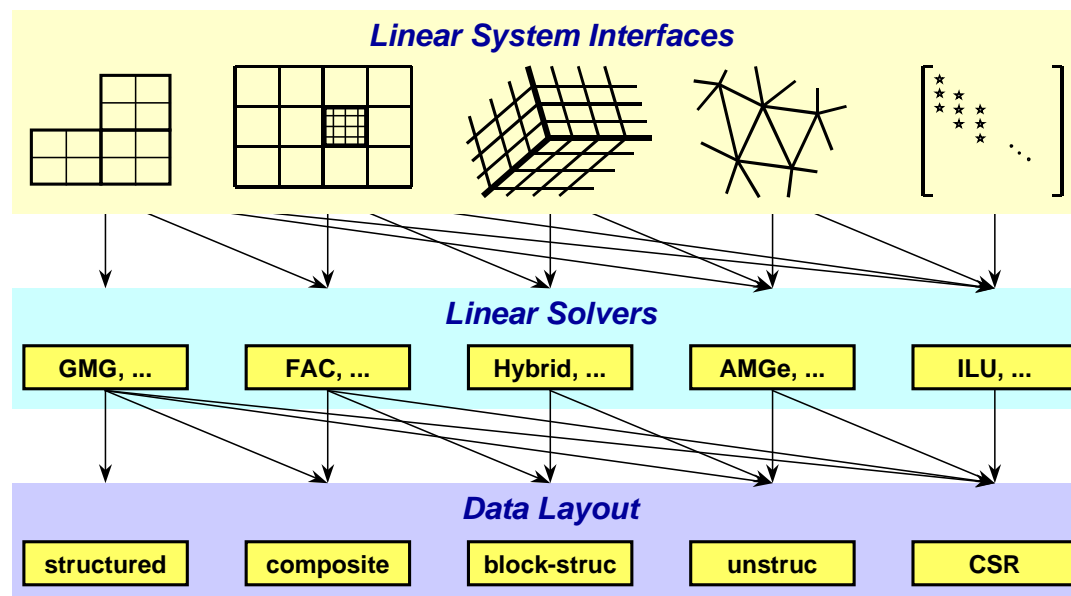
Classical AMG

- Classical or RS - AMG (84)
- Workhorse in many applications



Classical AMG

- Classical or RS - AMG (84)
- Workhorse in many applications
- Implemented as BoomerAMG in HYPRE



- Classical or RS - AMG (84)
- Workhorse in many applications
- Implemented as BoomerAMG in HYPRE
- Weaknesses:
 - Systems of PDEs
 - Singularities
 - Operator complexity in 3D



- Classical or RS - AMG (84)
- Workhorse in many applications
- Implemented as BoomerAMG in HYPRE
- Weaknesses:
 - Systems of PDEs
 - Singularities
 - Operator complexity in 3D
- Based on M-matrix principles



AMG Alphabet Soup

AMG	Classical AMG (84)
SA	Soothed Aggregation (96)
AMGe	finite element AMG (01)
AMG ϵ	element free AMGe (02)
ρ AMGe	spectral AMGe (03)

Adaptive Algorithms

AMG	adaptive AMG (84)
BAMG	Bootstrap AMG (01)
α SA	adaptive Soothed Aggregation (04)
CR	Compatible Relaxation (04)
α AMG	adaptive AMG (06)
α AMGr	adaptive AMGr (06)



Smoothed Aggregation (SA)

Brezina, Mandel, Vanek (96)



Smoothed Aggregation (SA)

Brezina, Mandel, Vanek (96)

- Requires knowledge of one (or more) global (near) null-space vector(s), \underline{v}_j .



Smoothed Aggregation (SA)

Brezina, Mandel, Vanek (96)

- Requires knowledge of one (or more) global (near) null-space vector(s), \underline{v}_j .
- Divide the Graph of A into disjoint aggregates, $\{\mathcal{A}_i\}$



Smoothed Aggregation (SA)

Brezina, Mandel, Vanek (96)

- Requires knowledge of one (or more) global (near) null-space vector(s), \underline{v}_j .
- Divide the Graph of A into disjoint aggregates, $\{\mathcal{A}_i\}$
- Associate one (or more) coarse-level DOF with each aggregate



Smoothed Aggregation (SA)

Brezina, Mandel, Vanek (96)


- Requires knowledge of one (or more) global (near) null-space vector(s), \underline{v}_j .
- Divide the Graph of A into disjoint aggregates, $\{\mathcal{A}_i\}$
- Associate one (or more) coarse-level DOF with each aggregate
- Construct a tentative interpolation matrix, \hat{P} , by chopping up the near null-space vector(s)



Smoothed Aggregation: Interpolation

Null-space vector: $\underline{v} = (v_1, v_2, \dots, v_n)^t$

$$\hat{P} = \begin{bmatrix} v_1 \\ \vdots \\ v_{n_{f1}} \\ \\ v_{n_2} \\ \vdots \\ v_{n_{f2}} \\ \\ v_{n_c} \\ \vdots \\ v_{n_{fc}} \end{bmatrix}$$



Note: \underline{v} is in $\text{Range}(P)$



Smoothed Aggregation: Interpolation

Normalize

$$\hat{P}^t \hat{P} = I$$



Smoothed Aggregation: Interpolation

Normalize

$$\hat{P}^t \hat{P} = I$$

Smooth \hat{P}

$$P = (I - \alpha A) \hat{P}$$



Smoothed Aggregation: Interpolation

Normalize

$$\hat{P}^t \hat{P} = I$$

Smooth \hat{P}

$$P = (I - \alpha A) \hat{P}$$

Construct coarse-grid operator

$$A_c = P^t A P = \hat{P}^t (I - \alpha A) A (I - \alpha A) \hat{P}$$



Smoothed Aggregation: Interpolation

Normalize

$$\hat{P}^t \hat{P} = I$$

Smooth \hat{P}

$$P = (I - \alpha A) \hat{P}$$

Construct coarse-grid operator

$$A_c = P^t A P = \hat{P}^t (I - \alpha A) A (I - \alpha A) \hat{P}$$

- Choose α to reduce the condition of A_c



Smoothed Aggregation: Interpolation

Normalize

$$\hat{P}^t \hat{P} = I$$

Smooth \hat{P}

$$P = (I - \alpha A) \hat{P}$$

Construct coarse-grid operator

$$A_c = P^t A P = \hat{P}^t (I - \alpha A) A (I - \alpha A) \hat{P}$$

- Choose α to reduce the condition of A_c

Recurse



$$A_c = P^t A P = \hat{P}^t (I - \alpha A) A (I - \alpha A) \hat{P}$$

- Reduces the condition of A_c
- Maintains good approximation of null-space vector, \underline{v}
 - Null-space, \underline{v} , still in the range of P
 - Other near null-space vectors still well approximated by P
- Yields aggressive coarsening



Multiple Null Space Vectors

Accommodate multiple (near) null-space vectors, $V = [\underline{v}_1, \dots, \underline{v}_k]$

$$V_j = \begin{bmatrix} v_{1n_j} & \cdot & v_{kn_j} \\ \vdots & & \vdots \\ v_{1n_{f_j}} & \cdot & v_{kn_{f_j}} \end{bmatrix}$$

$$\hat{P} = \begin{bmatrix} V_1 & & \\ & V_2 & \\ & & V_{n_c} \end{bmatrix}$$



Multiple Null Space Vectors

Accommodate multiple (near) null-space vectors, $V = [\underline{v}_1, \dots, \underline{v}_k]$

$$V_j = \begin{bmatrix} v_{1n_j} & \cdot & v_{kn_j} \\ \vdots & & \vdots \\ v_{1n_{f_j}} & \cdot & v_{kn_{f_j}} \end{bmatrix}$$

$$\hat{P} = \begin{bmatrix} V_1 & & \\ & V_2 & \\ & & V_{n_c} \end{bmatrix}$$

Normalize

$$\hat{P}^t \hat{P} = I$$

Smooth \hat{P}

$$P = (I - \alpha A) \hat{P}$$

Coarse-grid operator

$$A_c = P^t A P$$

Recurse



- *Very effective for systems, like linear Elasticity, where (near) null-space (rigid body motions) is known.*



- *Very effective for systems, like linear Elasticity, where (near) null-space (rigid body motions) is known.*
- *Effective in the context of irregular meshes*



- *Very effective for systems, like linear Elasticity, where (near) null-space (rigid body motions) is known.*
- *Effective in the context of irregular meshes*
- *Aggressive coarsening yields good complexity*



- *Very effective for systems, like linear Elasticity, where (near) null-space (rigid body motions) is known.*
- *Effective in the context of irregular meshes*
- *Aggressive coarsening yields good complexity*
- *Amenable to parallel implementation*



- *Very effective for systems, like linear Elasticity, where (near) null-space (rigid body motions) is known.*
- *Effective in the context of irregular meshes*
- *Aggressive coarsening yields good complexity*
- *Amenable to parallel implementation*
- *Conceptionally straightforward*



Compare SA to AMG

- SA constructs P column by column
- AMG constructs P row by row
- Both attempt to accurately interpolate algebraically smooth vectors
- Both try to reduce the complexity (number of nonzeros) of the coarse-grid operator



AMG Alphabet Soup

AMG	Classical AMG (84)
SA	Soothed Aggregation (96)
AMGe	finite element AMG (01)
AMG ϵ	element free AMGe (02)
ρ AMGe	spectral AMGe (03)

Adaptive Algorithms

AMG	adaptive AMG (84)
BAMG	Bootstrap AMG (01)
α SA	adaptive Soothed Aggregation (04)
CR	Compatible Relaxation (04)
α AMG	adaptive AMG (06)
α AMGr	adaptive AMGr (06)



AMG Gang

Ball State

I. Livshits

Delft

S. MacLachlan

Boulder

M. Brezina

Davidson College

T. Chartier

T. Manteuffel

FIT

J. Jones

S. McCormick

Penn State

J. Brannick

J. Ruge

J. Xu

G. Sanders

L. Zikatanov

B. Sheehan

SNL

J. Hu

LLNL

A. Baker

R. Tuminaro

A. Cleary

Urbana-Champaign

D. Alber

R. Falgout

L. Olson

V. Henson

Utah

O. Livne

T. Kolev

Weizmann Institute

A. Brandt

B. Lee

P. Vassilevski

U. Yang



AMG Alphabet Soup

AMG	Classical AMG (84)
SA	Soothed Aggregation (96)
AMGe	finite element AMG (01)
AMG ϵ	element free AMGe (02)
ρ AMGe	spectral AMGe (03)

Adaptive Algorithms

AMG	adaptive AMG (84)
BAMG	Bootstrap AMG (01)
α SA	adaptive Soothed Aggregation (04)
CR	Compatible Relaxation (04)
α AMG	adaptive AMG (06)
α AMGr	adaptive AMGr (06)



AMGe


Finite element AMG (04)



AMGe

- Uses local stiffness matrices
- Aggregates elements like SA
- Uses local null-space to determine local interpolation properties


Finite element AMG (04)



AMGe


- Uses local stiffness matrices
- Aggregates elements like SA
- Uses local null-space to determine local interpolation properties
- Effective for
 - Anisotropic Problems
 - Systems PDEs

Finite element AMG (04)




AMGe

- Uses local stiffness matrices
- Aggregates elements like SA
- Uses local null-space to determine local interpolation properties



Finite element AMG (04)



	AMG	AMGe
ρ	.98	.26

AMG Alphabet Soup

AMG	Classical AMG (84)
SA	Soothed Aggregation (96)
AMGe	finite element AMG (01)
AMG ϵ	element free AMGe (02)
ρ AMGe	spectral AMGe (03)

Adaptive Algorithms

AMG	adaptive AMG (84)
BAMG	Bootstrap AMG (01)
α SA	adaptive Soothed Aggregation (04)
CR	Compatible Relaxation (04)
α AMG	adaptive AMG (06)
α AMGr	adaptive AMGr (06)



ρ AMGe


Spectral AMGe (02)



ρ AMGe

- Based on local stiffness matrices like AMGe
- Aggregates elements like SA
- Creates local columns in interpolation matrix based on local null-space


Spectral AMGe (02)



ρ AMGe

- Based on local stiffness matrices like AMGe
- Aggregates elements like SA
- Creates local columns in interpolation matrix based on local null-space
- Blends rather than smooths columns of P


Spectral AMGe (02)



ρ AMGe

- Based on local stiffness matrices like AMGe
- Aggregates elements like SA
- Creates local columns in interpolation matrix based on local null-space
- Blends rather than smooths columns of P
- Effective when global null-space vectors not available, but local stiffness matrices are available

Spectral AMGe (02)



AMG Alphabet Soup

AMG	Classical AMG (84)
SA	Soothed Aggregation (96)
AMGe	finite element AMG (01)
AMG ϵ	element free AMGe (02)
ρ AMGe	spectral AMGe (03)

Adaptive Algorithms

AMG	adaptive AMG (84)
BAMG	Bootstrap AMG (01)
α SA	adaptive Soothed Aggregation (04)
CR	Compatible Relaxation (04)
α AMG	adaptive AMG (06)
α AMGr	adaptive AMGr (06)



AMG ϵ


element free AMGe (02)



AMG $\dot{\epsilon}$

- Based on principles of AMGe
- Aggregates elements like SA
- Creates local stiffness matrices from neighboring elements


element free AMGe (02)



AMG \neq

- Based on principles of AMGe
- Aggregates elements like SA
- Creates local stiffness matrices from neighboring elements
- Effective when local stiffness matrices are not available

element free AMGe (02)



AMG Alphabet Soup

AMG

Classical AMG (84)

SA

Soothed Aggregation (96)

AMGe

finite element AMG (01)

AMG ϵ

element free AMGe (02)

ρ AMGe

spectral AMGe (03)

Adaptive Algorithms

AMG

adaptive AMG (84)

BAMG

Bootstrap AMG (01)

α SA

adaptive Soothed Aggregation (04)

CR

Compatible Relaxation (04)

α AMG

adaptive AMG (06)

α AMGr

adaptive AMGr (06)

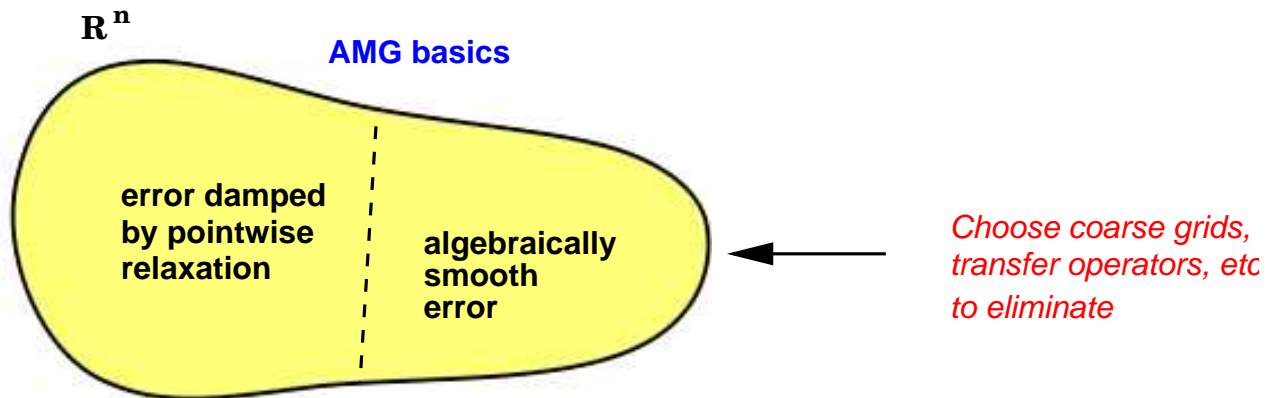


- AMG methods employ (relatively) simple relaxation



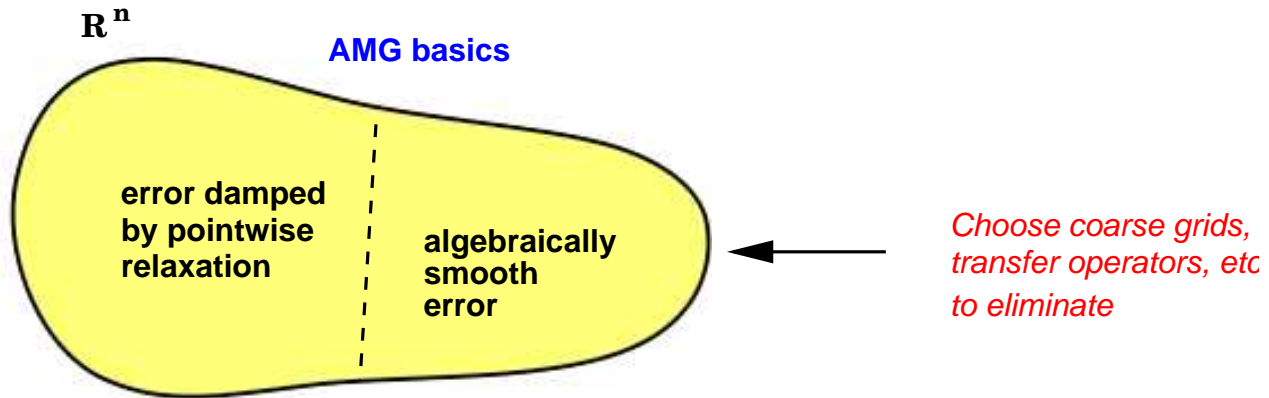
Adaptive AMG

- AMG methods employ (relatively) simple relaxation
- The coarse-grid problem must capture all modes not effectively reduced by relaxation



Adaptive AMG

- AMG methods employ (relatively) simple relaxation
- The coarse-grid problem must capture all modes not effectively reduced by relaxation






- Algebraically smooth vectors are not necessarily geometrically smooth



Algebraically smooth error can be oscillatory

- Error after seven Gauss/Seidel iterations on

$$-u_{xx} - \epsilon u_{yy} = f$$






$\epsilon=1$	$\epsilon = .001$
--------------	-------------------




Algebraically smooth error can be oscillatory

- Error after seven Gauss/Seidel iterations on

$$-u_{xx} - \epsilon u_{yy} = f$$



- Adaptive AMG can “follow physics”



- Let current method tell you what type of error is not being reduced effectively



- Let current method tell you what type of error is not being reduced effectively
- Adjust AMG components to capture this error



- Let current method tell you what type of error is not being reduced effectively
- Adjust AMG components to capture this error
- Do no harm: make sure change does not awaken previously reduced errors



- Let current method tell you what type of error is not being reduced effectively
- Adjust AMG components to capture this error
- Do no harm: make sure change does not awaken previously reduced errors
- Do as much of the work as possible on the coarser grids



- Let current method tell you what type of error is not being reduced effectively
- Adjust AMG components to capture this error
- Do no harm: make sure change does not awaken previously reduced errors
- Do as much of the work as possible on the coarser grids
- Test the current method and modify as necessary



Adaptive Smoothed Aggregation

- Given A , choose simple relaxation, call it the current method, C



Adaptive Smoothed Aggregation

- Given A , choose simple relaxation, call it the current method, C
- Iterate with the current method on $CA\underline{x} = \underline{0}$
 - If it is acceptable, stop



Adaptive Smoothed Aggregation

- Given A , choose simple relaxation, call it the current method, C
- Iterate with the current method on $CA\underline{x} = \underline{0}$
 - If it is acceptable, stop
- Approximate largest eigenvalue/vector of $(I - CA)$
 - Can be accomplished with a multilevel process



Adaptive Smoothed Aggregation

- Given A , choose simple relaxation, call it the current method, C
- Iterate with the current method on $CA\underline{x} = \underline{0}$
 - If it is acceptable, stop
- Approximate largest eigenvalue/vector of $(I - CA)$
 - Can be accomplished with a multilevel process
- Construct new coarse interpolation, P , and coarse-grid operator, A_c



Adaptive Smoothed Aggregation

Current approximation to the Null-space vector: $\underline{v} = (v_1, v_2, \dots, v_n)^t$

$$\hat{P} = \begin{bmatrix} v_1 \\ \vdots \\ v_{nf_1} \\ & v_{n_2} \\ & \vdots \\ & v_{nf_2} \\ & & v_{n_c} \\ & & \vdots \\ & & v_{nf_c} \end{bmatrix}$$

Normalize

$$\hat{P}^t \hat{P} = I$$

Smooth \hat{P}

$$P = (I - \alpha A) \hat{P}$$

Coarse-grid operator

$$A_c = P^t A P$$

Recurse



Adaptive Smoothed Aggregation

- Recursively construct V -cycle, call it the current method, C
 - Don't come back until your finished!



Adaptive Smoothed Aggregation

- Recursively construct V -cycle, call it the current method, C
 - Don't come back until your finished!
- Iterate with the current method on $CA\underline{x} = \underline{0}$
 - If acceptable, stop
 - Better approximation to null-space \underline{v}_1
 - Add new column to each aggregate $\underline{v}_1, \underline{v}_2$



Adaptive Smoothed Aggregation

- Recursively construct V -cycle, call it the current method, C
 - Don't come back until your finished!
- Iterate with the current method on $CA\underline{x} = \underline{0}$
 - If acceptable, stop
 - Better approximation to null-space \underline{v}_1
 - Add new column to each aggregate $\underline{v}_1, \underline{v}_2$
- Recurse



Adaptive Flavors

AMG

adaptive AMG (84)

BAMG

Bootstrap AMG (01)

α SA

adaptive Soothed Aggregation (04)

CR

Compatible Relaxation (04)

α AMG

adaptive AMG (06)

α AMGr

adaptive AMGr (06)



Adaptive Flavors

AMG	adaptive AMG (84)
BAMG	Bootstrap AMG (01)
α SA	adaptive Soothed Aggregation (04)
CR	Compatible Relaxation (04)
α AMG	adaptive AMG (06)
α AMGr	adaptive AMGr (06)

All depend on determining a local representation of algebraically smooth vectors



Perfect Interpolation

$$A = \begin{bmatrix} A_{ff} & A_{fc} \\ A_{cf} & A_{cc} \end{bmatrix}$$

$$P = \begin{bmatrix} -A_{ff}^{-1}A_{fc} \\ I \end{bmatrix}$$



Perfect Interpolation

$$A = \begin{bmatrix} A_{ff} & A_{fc} \\ A_{cf} & A_{cc} \end{bmatrix} \quad P = \begin{bmatrix} -A_{ff}^{-1}A_{fc} \\ I \end{bmatrix}$$

Choose diagonal matrix Δ_{ff}

$$\Delta_{ff}A_{fc}\underline{v}_1 = A_{ff}^{-1}A_{fc}\underline{v}_1$$



Perfect Interpolation

$$A = \begin{bmatrix} A_{ff} & A_{fc} \\ A_{cf} & A_{cc} \end{bmatrix} \quad P = \begin{bmatrix} -A_{ff}^{-1}A_{fc} \\ I \end{bmatrix}$$

Choose diagonal matrix Δ_{ff}

$$\Delta_{ff}A_{fc}\underline{v}_1 = A_{ff}^{-1}A_{fc}\underline{v}_1$$

- Adaptive approximation to smallest eigenvalue/vector(s), \underline{v}_1



CR

Livne(04), Brannick(05)

$$A = \begin{bmatrix} A_{ff} & A_{fc} \\ A_{cf} & A_{cc} \end{bmatrix} \quad P = \begin{bmatrix} -A_{ff}^{-1}A_{fc} \\ I \end{bmatrix}$$

- Principle: Coarse grid is adequate if A_{ff} is well conditioned



CR

Livne(04), Brannick(05)

$$A = \begin{bmatrix} A_{ff} & A_{fc} \\ A_{cf} & A_{cc} \end{bmatrix} \quad P = \begin{bmatrix} -A_{ff}^{-1}A_{fc} \\ I \end{bmatrix}$$

- Principle: Coarse grid is adequate if A_{ff} is well conditioned
- Use simple relaxation on A_{ff} , together with a greedy independent set algorithm, to choose coarse grid



α AMG and α SA surprisingly effective on a wide range of problems


- Highly irregular meshes
- Strongly anisotropic
- Adaptively refined meshes
- Discontinuous coefficients (heterogeneous material)
- Singularities
- Hyperbolic problems
- QCD



Adaptive AMG for Lattice QCD


- Quantum Chromodynamics (QCD) calculations involve huge linear systems and large-scale (petascale) computing
- Requires solving the complex and non-hermitian discretized Dirac operator
- Each equation may be solved 1000s times

$$\begin{aligned}
 M(\mathcal{U}) &= D(\mathcal{U}) - m_0 I \\
 &= \begin{bmatrix} A - m_0 I & B \\ -B^* & A - m_0 I \end{bmatrix}
 \end{aligned}$$




QCD: 2D Schwinger Model

- The system becomes extremely ill-conditioned for typical choices of m_0





$$m_0 = 0$$

→



$$m_0 \approx m_{cr}$$

- Near null space is unknown and oscillatory



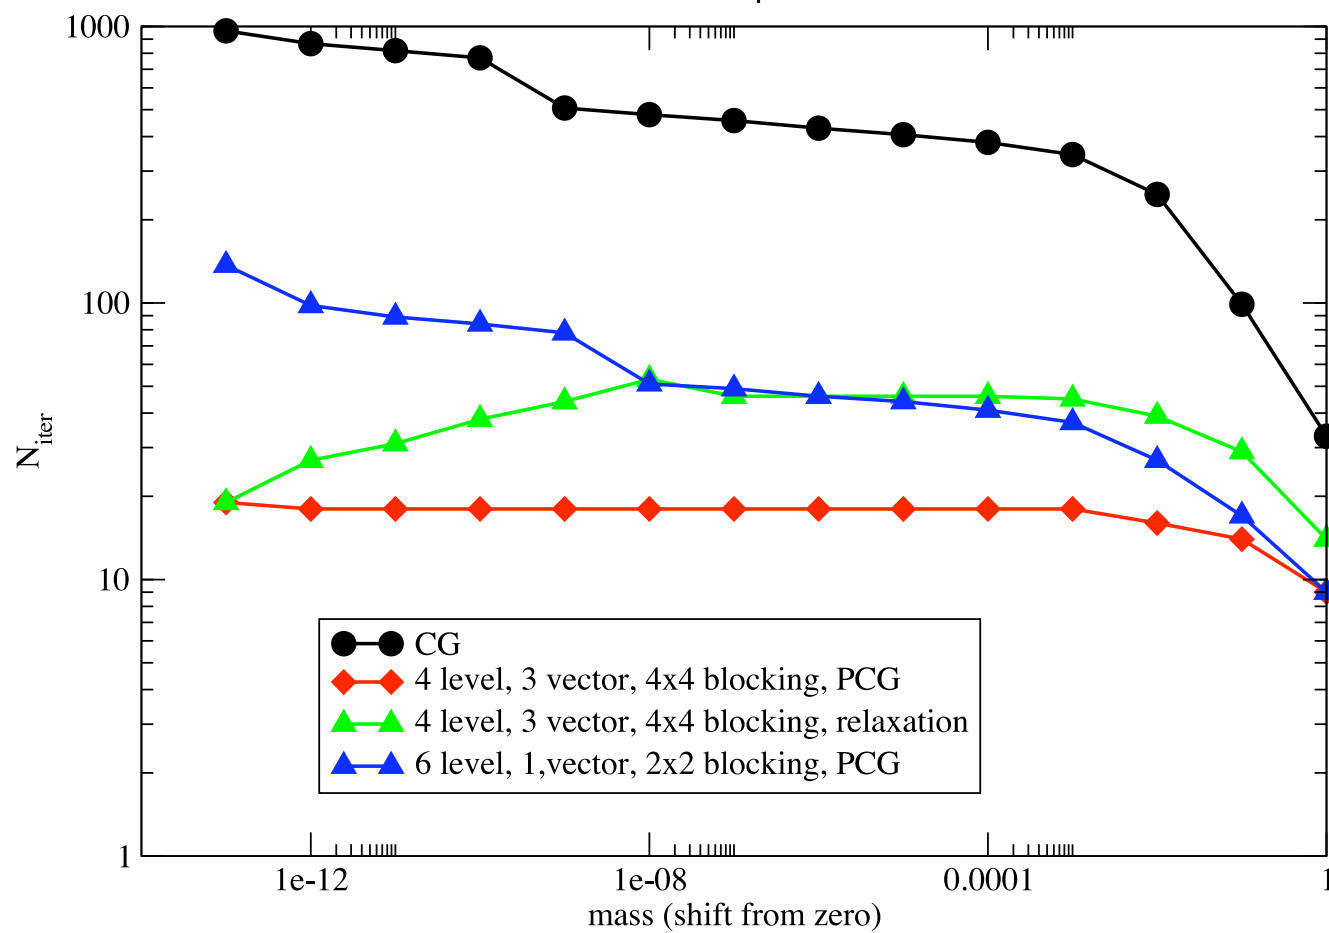
2D Schwinger Model

- Form the normal equations and apply αSA
- Set-up requires 100s of Work Units
 - (WU = matrix vector multiply)
- Interpolation requires 8 – 10 columns on each aggregate
- For small mass shift, faster than the current method (Diagonally scaled PCG) on even one right-hand side



Mass scaling of Gauged Laplacian

$$V=128^2, \beta=1.0$$



- Real Problem: 4D model – preliminary results promising



Real Problem

- Real Problem: 4D model – preliminary results promising
- Real Real Problem: Dirac Equations



Real Problem

- Real Problem: 4D model – preliminary results promising
- Real Real Problem: Dirac Equations

α SA allows the QCD community to do problems that they could not do before



- Linear systems from PDEs require multilevel algorithms
- GMG optimal for structured grids
- AMG/SA effective for unstructured grids, known (near) null-space
- α AMG/SA greatly expand the domain of applicability



- Linear systems from PDEs require multilevel algorithms
- GMG optimal for structured grids
- AMG/SA effective for unstructured grids, known (near) null-space
- α AMG/SA greatly expand the domain of applicability
- Adaptive AMG/SA a group effort



AMG Gang

Ball State

I. Livshits

Delft

S. MacLachlan

Boulder

M. Brezina

Davidson College

T. Chartier

T. Manteuffel

FIT

J. Jones

S. McCormick

Penn State

J. Brannick

J. Ruge

J. Xu

G. Sanders

L. Zikatanov

B. Sheehan

SNL

J. Hu

LLNL

A. Baker

R. Tuminaro

A. Cleary

Urbana-Champaign

D. Alber

R. Falgout

L. Olson

V. Henson

Utah

O. Livne

T. Kolev

Weizmann Institute

A. Brandt

B. Lee

P. Vassilevski

U. Yang

