Abstract

Substantial effort has recently focused on developing methods capable of solving very large linear systems that arise from discretizing partial differential equations, especially on unstructured grids. Algebraic multigrid (AMG) is of particular interest because of its promise of optimal performance without the need for explicit knowledge of the problem's origin. We introduce an extension of AMG based on an adaptive process that achieves good convergence on a broader class of problems than the original algorithm.

Model Problem

We consider, as a model problem, the steady-state diffusion equation

$$-\nabla \cdot D(x)\nabla u(x) = f(x)$$

plus boundary conditions, discretized via (bilinear) Finite Elements. We will consider the special case $D(x) \equiv 1$, the Poisson Equation. The matrices resulting from these problems are symmetric and positive definite (when $D(x)$ is).

Multigrid Basics

The efficiency of multigrid for large linear systems comes from identifying and exploiting the weaknesses of simple solution techniques. The coupling of simple stationary iterative methods, such as Jacobi or Gauss-Seidel, with corrections interpolated from a coarse-scale problem yields an optimal-order recursive solution technique.

Smoothing and Coarse-Grid Correction

The Jacobi iteration for solving the discretized Poisson equation takes the form of an averaging operation. Thus, errors which vary slowly between neighboring gridpoints are slow to be resolved. These errors, however, can be represented using fewer degrees of freedom and can thus be transferred to a coarser grid for resolution. Applying this idea recursively leads to a fast solution technique.

Geometric Coarsening

When we know the problem geometry, we can choose a coarser grid by eliminating points in a geometrically-regular pattern.

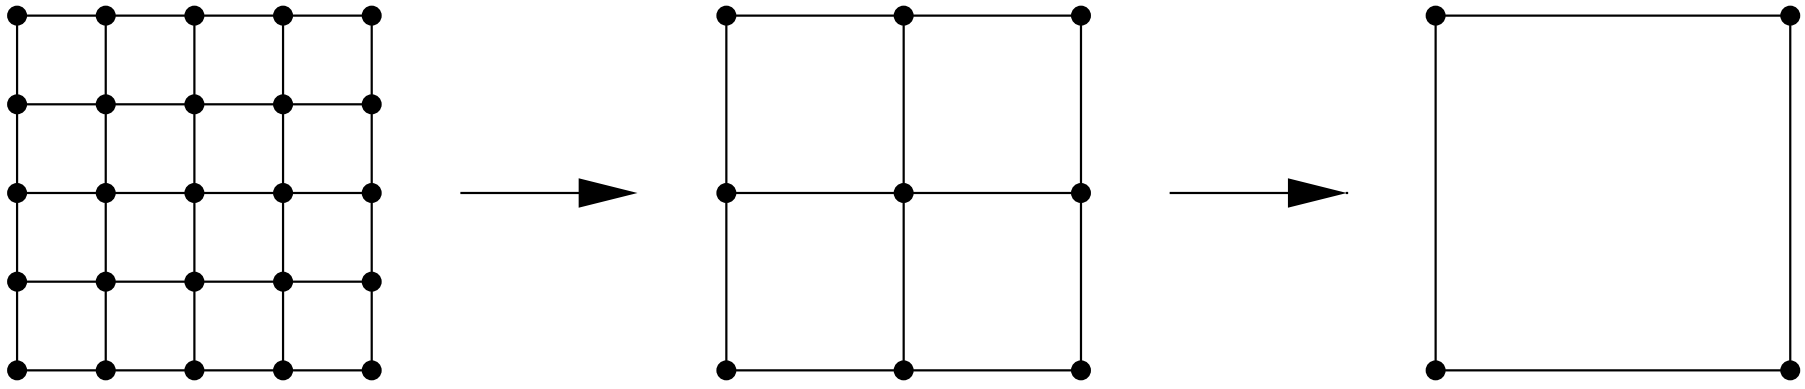

Given a correction computed on a coarse-grid, we use it to refine our fine-grid solution by interpolating the correction with a nearest-neighbor averaging (full weighting).

Algebraic Multigrid

In the absence of geometric information, choices must be based on algebraic information. Smoothing still occurs, but it cannot necessarily be interpreted geometrically. A vector is said to be *algebraically smooth* if it is slow to be resolved by relaxation. Algebraic multigrid methods look to choose coarse grids and intergrid transfer operators based on this sense of smoothness.

Algebraic Connection

Classical AMG chooses the coarse grid as a subset of the fine grid. Considering each point, i, the coarse grid is chosen so that i is connected to at least one coarse grid point which can be used to accurately interpolate a value to i. An algebraic connection (a_{ij}) is said to be *strong* if $a_{ij} \geq \theta \max_k |a_{ik}|$. For our model problem we can see that algebraically smooth error varies slowly along strong connections, and thus we choose the coarse grid so that each point has at least one strong connection to a coarse-grid point.

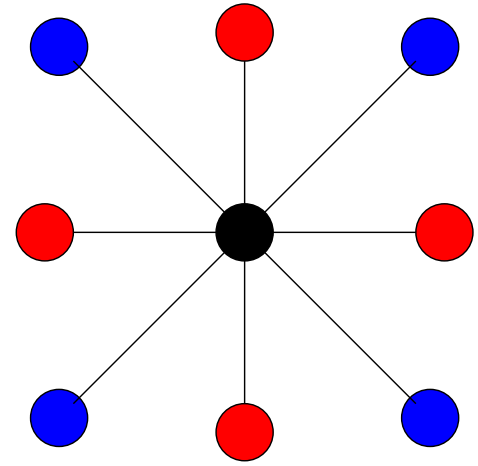

Fine Grid Points

Coarse Grid Points

Variations

Smoothed Aggregation uses a collection of aggregates of fine-grid points as coarse-grids. Columns of interpolation are chosen by smoothing known near-nullspace vectors on each aggregate.
Element-based AMG (AMGe) chooses interpolation based on access to the local finite element stiffness matrices.

Black Box Multigrid (BoxMG) chooses coarse grids geometrically, but uses an algebraic interpolation definition that can be shown to preserve the continuity of normal flux.

Weaknesses of Classical AMG

Classical AMG methods are based on the assumption that smooth components (both algebraic and geometric) are near-constant. This assumption is appropriate for matrices that comes from simple discretizations of our model problems, but is not robust enough to handle all situations of interest. For scalar problems, some discretization methods do not yield matrices for which the discrete nullspace is near-constant. For systems of PDEs, inter-variable coupling often results in components which relaxation is slow to resolve, but which are not near-constant (e.g. elasticity).

Adaptive AMG (α AMG)

Multigrid methods achieve their efficiency through the complementarity of smoothing and coarse-grid correction. We consider the framework where relaxation is fixed and seek to construct coarsening to quickly eliminate algebraically smooth components. This means that if smoothing is inefficient on a component, we must determine interpolation so that coarse-grid correction can provide an accurate correction for that component.

Main Ideas

- Smoothing on $Ax = 0$ with a random initial guess quickly exposes errors that are slow to be resolved.
- The coarse-grid and interpolation operator are then chosen so that this error can be accurately corrected from the coarse-grid.
- The coarse-grid operator is formed using the Galerkin condition ($A_c = P^T A P$).
- The full algorithm is then specified by recursion.

Systems of PDEs

Systems of PDEs are somewhat more difficult as their near-nullspaces have dimension higher than one, so we must base interpolation on multiple, distinct, slow-to-converge components. In the context of smoothed aggregation this is natural - we simply add columns to the interpolation operator. In the context of classical AMG, this can be achieved by determining interpolation to fit multiple vectors.

Numerical Results

We have implemented two versions of this method. Adaptive Smoothed Aggregation (αSA) *(submitted to SISC)* employs these ideas in a Smoothed Aggregation context and is applicable to two- and three-dimensional problems for both scalar PDEs and systems of PDEs. Adaptive AMG (αAMG) is in an earlier stage of development and is applicable to two-dimensional, scalar PDEs. The current implementation of αAMG determines interpolation in an algebraic fashion, but relies on a geometric choice of coarse-grids.

2D Laplace

Consider the Poisson equation on the unit square in two-dimensions, with pure Dirichlet boundary conditions. If we impose the full-coarsening used by αAMG on classical AMG, we achieve convergence factors bounded above by 0.07 as h varies from $\frac{1}{32}$ to $\frac{1}{1024}$. Allocating 12 work units to initial relaxations, we were able to recover these rates for αAMG.

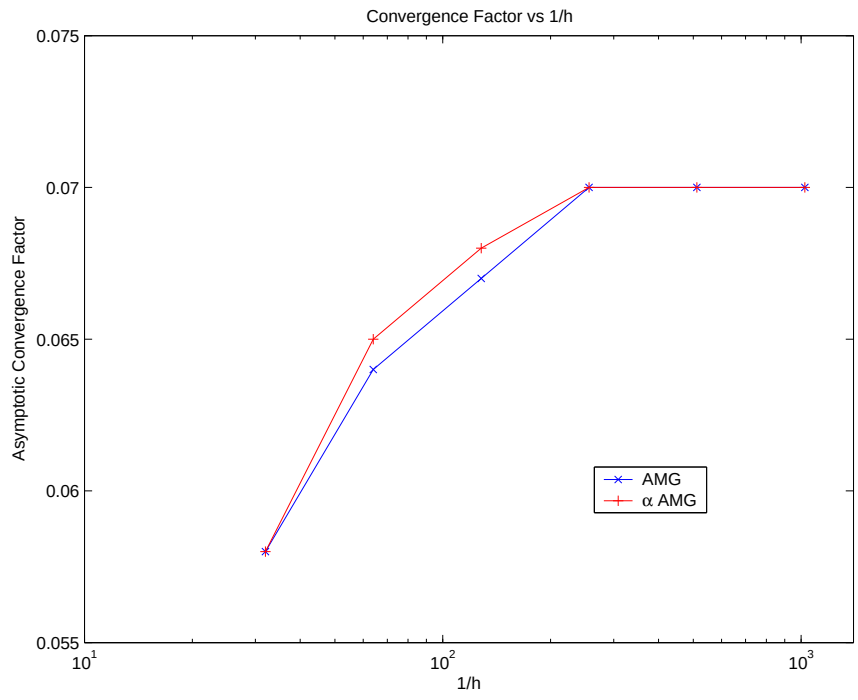

2D Diffusion

Consider the model problem with piecewise constant $D(x, y)$ chosen to be 100 in the square $[\frac{1}{3}, \frac{2}{3}]^2$, with Dirichlet boundary conditions on the left and right boundaries and Neumann boundary conditions on the top and bottom.

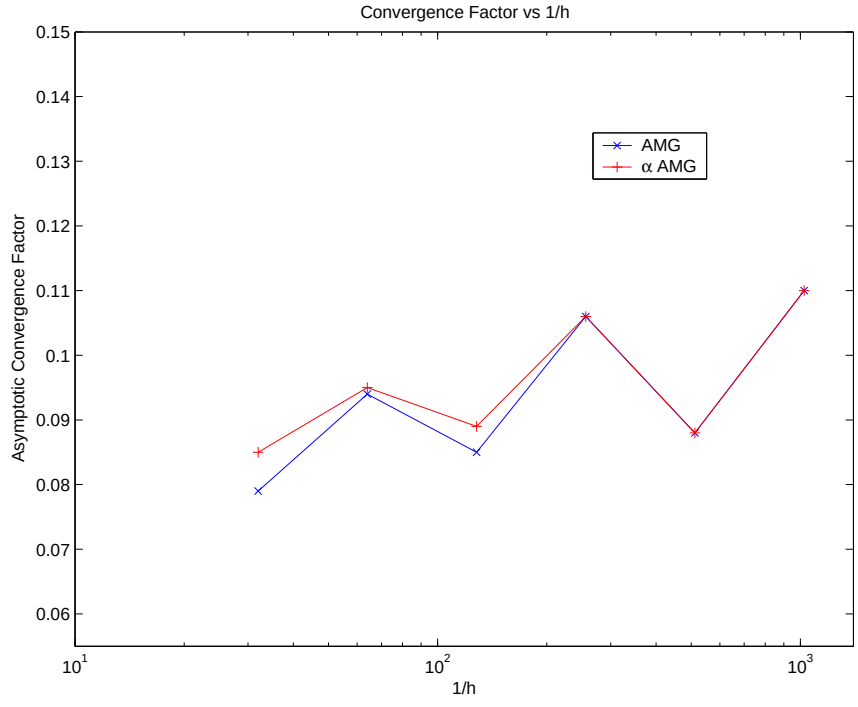

Imposing the same full-coarsening, we achieve convergence factors bounded above by 0.11 as h varies from $\frac{1}{32}$ to $\frac{1}{1024}$ for both AMG and αAMG with 18 work units allocated to initial relaxation.

Scaled Matrices

To test the robustness of αAMG, we multiplied the matrices from the previous two examples by a diagonal matrix with entries ranging from 10^{-7} to 2. This causes AMG performance to significantly degrade, however αAMG performance does not suffer.

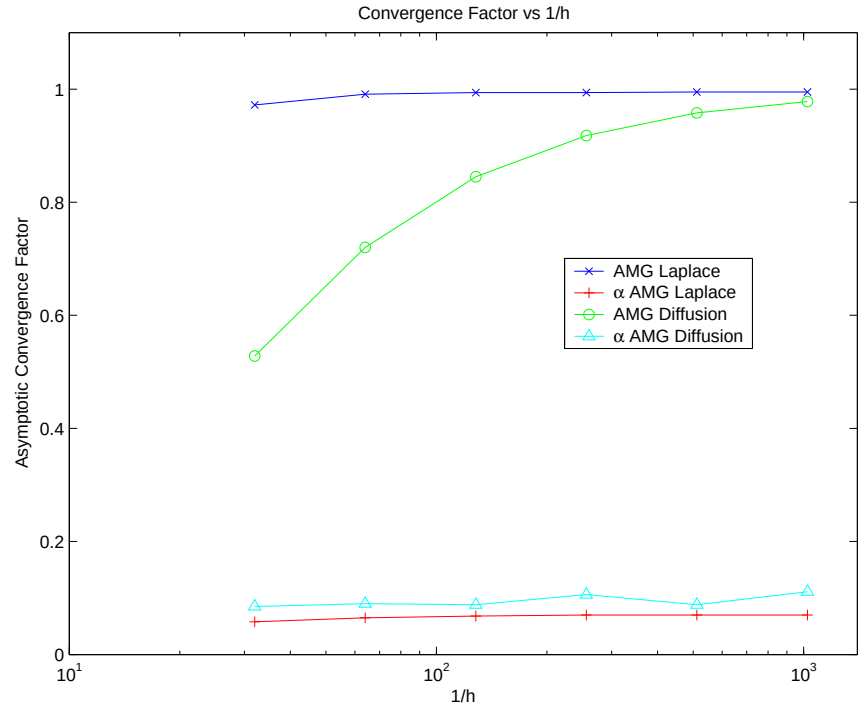

2D Elasticity

The αSA algorithm can be quite effective in solving elasticity problems. For a 2D elasticity problem with 181,202 degrees of freedom, standard Smoothed Aggregation required 23 iterations to reduce the residual by a factor of 10^{-12}. When the degrees of freedom were rotated by random angles, over 5,000 iterations were required. αSA required 18 iterations on the non-rotated problem (although took about 3 times more CPU time), and was able to solve the rotated problem in 18 iterations, taking only 2.66 times longer than standard SA on the non-rotated problem.

3D Elasticity

For a 3D problem with 201,720 degrees of freedom, standard Smoothed Aggregation required 16 iterations for the same residual reduction. After performing a nodal rotation in all 3 dimensions, 739 iterations were required. αSA required 17 iterations to solve the non-rotated problem (and took 6 times longer). To solve the rotated problem, αSA took only 15 iterations, requiring 6 times more CPU time than standard Smoothed Aggregation for the non-rotated problem, but over 6 times less CPU time than standard SA on the rotated problem.

Other Algorithms

Similar methods and similar algorithms are currently being studied in a number of research groups.

- Achi Brandt and Oren Livne have investigated adaptivity in the context of AMG and compatible relaxation.
- Tim Chartier is investigating adaptivity in the context of AMGe and spectral AMGe.

We have active collaborations with both of these projects.

Conclusions

The framework developed for Adaptive Multigrid Methods has been demonstrated to yield improved results for a number of “difficult” problems. While we cannot hope to improve upon the optimal performance exhibited by classical Multigrid, AMG, and Smoothed Aggregation for problems with known near-nullspaces, we have demonstrated that, when the near-nullspace is not known a priori, we can recover the desired performance.