An Unstructured Multigrid Method Based on Geometric Smoothness*

Edmond Chow†

Abstract

For non-M-matrices, this paper proposes an unstructured multigrid method that only attempts to interpolate in the directions of geometrical smoothness. These directions are determined by analyzing samples of algebraically smooth error, e. Neighboring grid points i and j are called smoothly coupled if e_i and e_j are consistently nearby in value. In addition, these differences may be used to define interpolation weights. These new ideas may be incorporated into the algebraic multigrid method. Test results show that the new method can have much lower grid and operator complexities compared to AMG, leading to lower solve timings.

1 Introduction

In the algebraic multigrid (AMG) method, algebraically smooth error is the error that remains after the smoother has been applied and that must be reduced at the next level. AMG defines the strong couplings between a grid point and its neighboring grid points, and depends on being able to interpolate the algebraically smooth error along these strong couplings. Strong couplings are generally defined to be the large negative matrix entries, i.e., the coupling between points i and j is strong when the matrix entry a_{ij} is large and negative. For symmetric M-matrices and matrices where the positive off-diagonal entries are small ("essentially positive-type") [2], the algebraically smooth error *varies slowly* along strong couplings and is relatively easy to interpolate.

When the positive off-diagonal entries are large but the matrix is (weakly) diagonally dominant, algebraically smooth error still varies slowly along strong couplings [9]. Along large positive couplings, the algebraically smooth error is oscillatory [11] such that the absolute value of the smooth error is slowly-varying. This makes it possible to interpolate algebraically smooth but geometrically oscillating error for these problems.

When the matrix is not diagonally dominant and there exist large positive off-diagonal entries, algebraically smooth error no longer varies slowly along strong couplings in general. Further, it is not clear that it is possible to interpolate geometrically non-smooth error for these problems. Matrices of this sort arise commonly, for example, in linear quadrilateral or hexahedral finite element discretizations for anisotropic elliptic problems, but traditional AMG methods have difficulty with them. This paper proposes a multigrid method for these problems by directly identifying the directions of slowly-varying algebraically smooth error, and only interpolating in these directions.

*Submitted to *Numerical Linear Algebra with Applications*. This work was performed under the auspices of the U.S. Department of Energy by University of California Lawrence Livermore National Laboratory under contract No. W-7405-Eng-48.

†Center for Applied Scientific Computing, Lawrence Livermore National Laboratory, L-560, Box 808, Livermore, CA 94551 (echow@llnl.gov).

-1.0	1.9	-1.0
-3.9	8.0	-3.9
-1.0	1.9	-1.0

Figure 1: Stencil for Poisson's equation using elements with 1/10 aspect ratio.

For reference, we state AMG's definition of a strong coupling, also called *strong connection* or *strong dependence* [6], for interpolating grid point i. Point i is strongly coupled to j if

$$-a_{ij} \geq \theta \max_{k \neq i} \{-a_{ik}\} \tag{1}$$

where $0 < \theta \leq 1$ is called the *strength threshold*. We additionally say that for $\theta = 0$, point i is strongly coupled to j if $a_{ij} < 0$.

As a motivational example, consider solving Poisson's equation in a rectangular domain $(0,1) \times (0,10)$ with Dirichlet boundary conditions using linear quadrilateral elements on a 20×20 element mesh. Each element has aspect ratio 1/10. The AMG method by Ruge and Stüben [10] with 2-levels (2 symmetric Gauss-Seidel pre-smoothing steps, direct solve on the coarse level) using a strength threshold of 0.25 (the value suggested in the literature) requires 32 cycles to reduce the residual by six orders of magnitude. The stencil, shown in Figure 1, shows that this choice of strength threshold defines all connections except the positive connections as strong. However, the smooth error varies slowly only in the east-west direction in this problem; the smooth errors in the grid lines above and below the given point are unrelated, and interpolation using points in these grid lines is erroneous. A strength threshold of 0.26 or higher correctly classifies the couplings and AMG requires only 7 cycles for convergence with this choice. (When the anisotropy is stronger, the detrimental effect of the wrong strength threshold is even stronger.) For unstructured problems, the size of moderately negative entries does not clearly indicate whether smooth error is slowly-varying. To rectify this problem, it is possible to increase the strength threshold so that only the most negative connections are considered strong connections, erring on the side of misclassifying some strong directions. However, this strategy entails using large coarse grids (i.e., requiring more levels) and may degrade interpolation.

This paper suggests the use of *geometrically smooth couplings* instead of strong couplings for problems that are not (weakly) diagonally dominant such as the above. The proposed method explicitly identifies geometrically smooth couplings by using differences in samples of algebraically smooth error. Interpolation only uses these smooth couplings. Thus we do not attempt to interpolate across jumps in PDE coefficients, for example, since we assume that we do not know how to interpolate across these jumps. The method is identical to AMG except that it redefines strong couplings and that it uses geometrically-based interpolation. The coarsening algorithm, constructing the coarse grid matrix, and smoothers remain the same.

Other multigrid techniques have been proposed for problems with matrices with large positive off-diagonal entries. Stüben [12] suggests eliminating the off-diagonal positive entries in a stencil (or row) by substituting stencils corresponding to the positive entries. The resulting stencil can be used to determine which couplings are strong. It may be difficult, however, to interpret the resulting stencil, especially if not all positive entries can be eliminated. Another multigrid method is called AMGe, an algebraic multigrid method for finite element problems [5]. AMGe agglomerates elements to form coarse grid matrices, and utilizes a measure of interpolation quality to define interpolation for these problems.

This paper is organized as follows. A method for identifying geometrically smooth couplings is described in Section 2. Interpolation along couplings that are geometrically smooth is described in Section 3. Section 4 shows the results of numerical investigations and Section 5 concludes this paper.

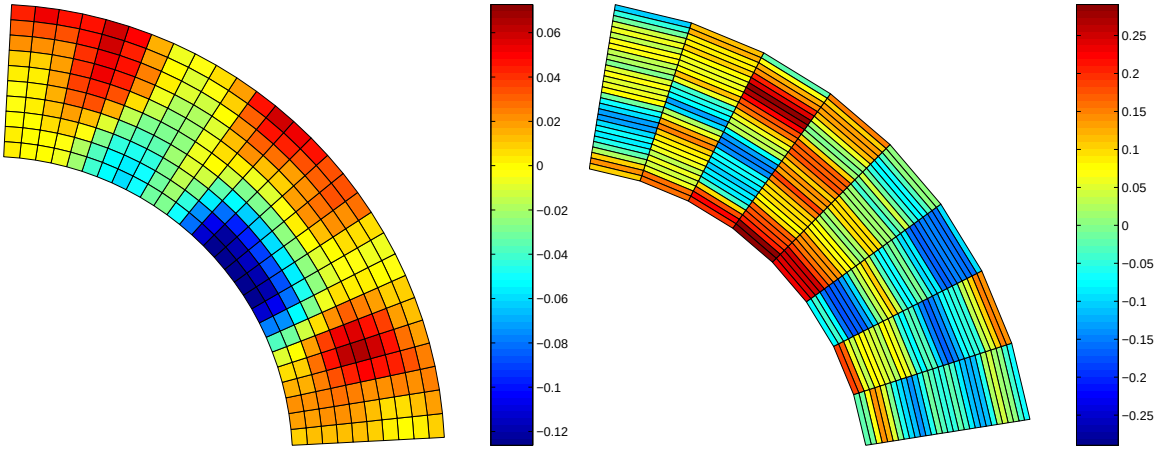

(a) Well-shaped elements

(b) Stretched elements

Figure 2: Samples of algebraically smooth error for two finite element discretizations.

2 Identifying geometric smoothness

2.1 Smoothness matrix

Given a coefficient matrix A at a given grid level, a *sample* of algebraically smooth error may be generated as the result of applying the smoother (to be used in the multigrid solution process) to the homogeneous equations

$$Ae = 0 \tag{2}$$

with a random initial guess for e. This relaxed vector e is representative of the actual errors remaining after smoothing that are generated in the multigrid solution process. As we will see, however, many samples of algebraically smooth error will be needed to fully characterize this error. Relaxed vectors produced this way have been used in other algorithms, for example, Bootstrap AMG [4].

Figure 2 shows samples of algebraically smooth error for two linear quadrilateral finite element discretizations of Poisson's equation on a segment of an annulus. The smoother was two steps of symmetric Gauss-Seidel (SGS). The figure shows that when the grid is stretched, the smooth error is only geometrically smooth in the radial direction; there is no relation in the error components that can be discerned in the angular direction.

To try to quantitatively characterize geometrically smooth error, we proceed as follows. Let n be the dimension of a sample of algebraically smooth error e. Define an n-by-n sparse matrix M that specifies a sparsity pattern that does not include the diagonal, and define $\mathrm{dist}(i,j)$ as the graph distance between grid points i and j. We then define the "difference matrix" as a function of e

$$D_{ij}(e) = \begin{cases} \frac{1}{\|e\|_2}\frac{|e_i - e_j|}{\mathrm{dist}(i,j)} & \text{if } M_{ij} \neq 0 \\ \text{undefined} & \text{otherwise} \end{cases} \tag{3}$$

whose entries can be interpreted as the directional derivatives of e. Large values of D_{ij} indicate geometric non-smoothness in the grid; small values of D_{ij} do not imply smoothness unless the same entries are small for many samples of algebraically smooth error. We thus define the "smoothness matrix" using the average of R difference matrices (using R samples of algebraically

smooth error)

$$S_{ij} = \begin{cases} \left[\frac{1}{R}\sum_{k=1}^{R} D_{ij}(e^{(k)})\right]^{-1} & \text{if } M_{ij} \neq 0 \\ 0 & \text{otherwise} \end{cases} . \tag{4}$$

Note the reciprocal used in (4); now large values of S_{ij} indicate geometric smoothness and small values of S_{ij} indicate geometric non-smoothness.

Finally, we define a "smooth coupling" as a coupling between points i and j such that $S_{ij} \geq \tau$, where τ is a threshold to be discussed in Section 2.2. A thresholded version of S may be passed to AMG to use as its "strength matrix" which is used to define the coarse grid and interpolation.

A few notes about these definitions are in order.

- D is an approximation to a derivative in a graph sense rather than a geometric sense. This is reasonable because coarsening is defined in terms of graph neighbors rather than points in a geometric neighborhood.
- The matrix M is used to make sure that differences are only computed between nearby grid points. Most often, the pattern of M will be taken to be the pattern of the matrix A less the diagonal. This definition of M implies that a grid point will at most be interpolated using its nearest graph neighbors. An expanded pattern for M (e.g., the pattern of A^2 less the diagonal) may be useful for non-grid-aligned problems or to generate very coarse grids, but interpolation using farther graph neighbors is required. In AMG, strength of connection is only defined for the nearest graph neighbors of a node, but the above technique can be used to identify smooth couplings that are not these nearest graph neighbors.
- The scaling by $\|e\|_2$ in (3) makes the difference matrix independent of how much the error was reduced by the smoother. Localized scalings may be appropriate in some cases, and will be discussed in Section (2.3).
- Instead of defining the smoothness matrix using an average of difference matrices, it is possible to define it using a component-wise maximum of difference matrices. However, we did not find the latter to give significantly different results.
- The matrix S is symmetric, i.e., $S_{ij} = S_{ji}$, which is suited for symmetric problems. We plan to investigate how to extend the above ideas to nonsymmetric problems.

The calculation of the smoothness matrix is adapted to the smoother being used in the multigrid solution process. In a parallel processing example, if a non-overlapping block Jacobi smoother is used, then the smoothness matrix will automatically prevent interpolation across processor boundaries. Another technique, called compatible relaxation [3] uses the smoother to help select the coarse grid, and is similar in that it is also adapted to the smoother chosen for the solution process.

2.2 Smoothness threshold

A smoothness threshold τ discriminates between geometrically smooth and non-smooth couplings. It is not clear that a fixed value of τ is appropriate for all problems. Further, different smoothness thresholds may be required on different coarse grid levels.

The following procedure for automatically choosing the smoothness threshold at each level is surprisingly effective. We assume that the smoother generates algebraically smooth error such that every grid point should have at least one smooth coupling. Then, given a smoothness matrix S, we choose

$$\tau = \min_i \max_j S_{ij}.$$

Now, points i and j are called smoothly coupled if S_{ij} is greater than or equal to τ.

2.3 Local scaling of the smoothness matrix

The scaling by $\|e\|_2$ in (3) is required to make the difference matrix independent of how much the error was reduced by the smoother and independent of the scaling of the initial guess for (2). For some problems, it is possible that the error is reduced at different rates in different parts of the grid, and *local scalings* of the relaxed vectors may be appropriate. For a relaxed vector e, an obvious local scaling is to scale e_i by $\max_{M_{ij} \neq 0}(e_j)$.

We use a simpler approach, which is to scale the smoothness matrix such that S_{ij} is scaled by $\max_k S_{ik}$. When only one relaxed vector is used, this is identical to scaling this relaxed vector as described above. Our local scaling is also similar to to the way strong connections are defined in AMG; see (1).

A smoothness threshold, $0 \leq \tau \leq 1$, is also used to classify smooth and non-smooth couplings when local scaling is used. However, we do not have an automatic procedure for choosing the smoothness threshold in this case.

2.4 Grid and operator complexity

The following definitions, from [6], are useful to quantify the storage and work required by a multigrid V-cycle. *Grid complexity* is the total number of grid points, on all grids, divided by the number of grid points on the finest grid. *Operator complexity* is the total number of nonzero entries, in all coarse and fine grid matrices, divided by the number of nonzero entries in the fine grid matrix. Operator complexity is an indication of the work per V-cycle of the multigrid process.

It turns out (see Section 4) that the coarsening procedure based on smooth couplings chooses fewer coarse grid points (C-points) than coarsening based on approximately the same number of AMG strong couplings. This leads to smaller-dimension coarse grids, fewer levels, and smaller operator and grid complexities. To try to understand why these coarsenings behave differently, we examine the number of smooth or strong couplings at each grid point. Figure 3 plots the histograms of the number of strong or smooth couplings at each grid point on the finest grid for three methods. The figure shows that the AMG rule for strong couplings limits the maximum number of strong couplings at a grid point. This is due to its use of local scaling, which assumes that not all couplings at a grid point can be strong. Local scaling of the smoothness matrix has the same difficulty (although it is not as evident from this figure) but it is not as severe.

2.5 Number of samples of algebraically smooth error

Using more samples of algebraically smooth error gives more accurate discrimination of the smooth couplings, but obviously incurs greater cost. Figure 4 plots the sorted entries in the smoothness matrix for an anisotropic diffusion problem on a 20-by-20 grid, as a function of the number of samples of algebraically smooth error. In this example, the smooth couplings are known, and the optimal threshold is at the knee near 40. The knee develops after averaging about ten samples, but there is no knee in general.

The plot gives an idea of the values in the smoothness matrix, but not necessarily the number of couplings that are classified correctly. The number of sample vectors will be investigated further in Section 4.3 with respect to convergence rate and solution time.

2.6 Implementation details

Algorithm 1 shows the calculation of the entries of the smoothness matrix S (without local scaling). These entries will typically be stored using a sparse data structure, and it is more efficient to completely compute each entry or each column of S before beginning to compute

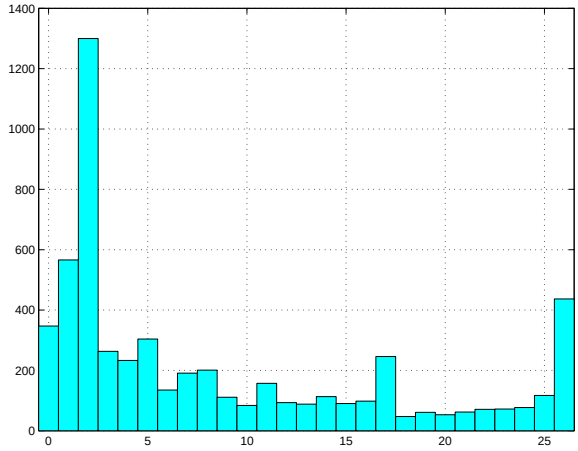

(a) Smooth couplings, 48390 couplings (1894 C-points, 6 levels).

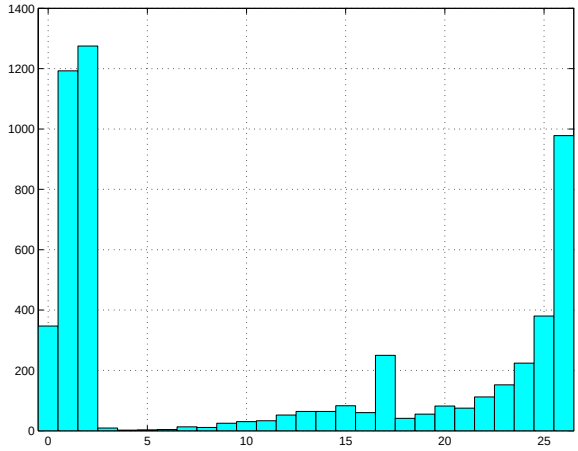

(b) Smooth couplings using local scaling, τ=0.25, 64953 couplings (1736 C-points, 6 levels).

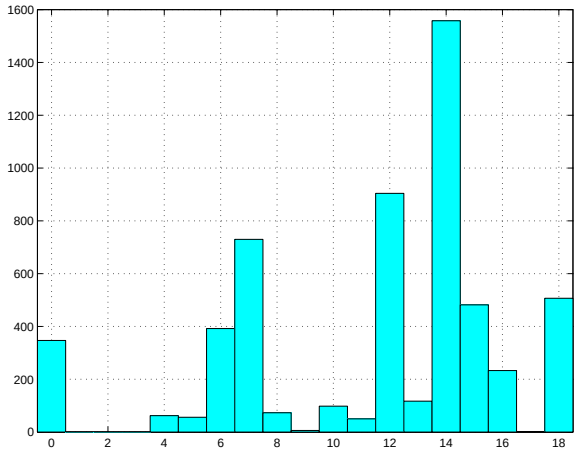

(c) AMG strong couplings, θ=0.25, 64457 couplings (2480 C-points, 11 levels).

Figure 3: Histograms of the number of strong or smooth couplings at each grid point. (The y-axis shows the number of grid points with the number of strong or smooth couplings on the x-axis.) These results are for the test matrix UU-2, described in Section 4.1. Each row of the test matrix typically has 27 nonzeros per row. The grid points with zero couplings are grid points specifying Dirichlet boundary conditions.

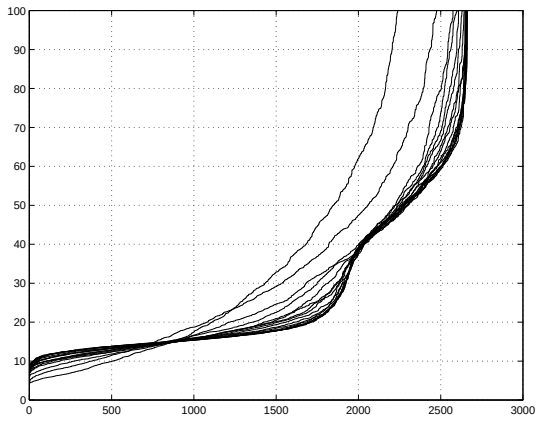

Figure 4: Plots of the entries in the smoothness matrices for 1 to 20 samples of algebraically smooth error (size of entries in smoothness matrix is sorted along the x-axis). The entries for 20 samples has the most developed knee.

the next entry or column. It is also more efficient to smooth a block a vectors than to smooth vectors individually (although this optimization is not used in the tests in this paper). If the computation is organized this way, then the number of samples R of algebraically smooth error must be known beforehand. In addition, since S is symmetric, only its upper or lower triangular portion needs to be computed.

2.7 Examples

Finite elements with stretched quadrilaterals

Figure 5 shows typical entries in the smoothness matrix for the stencil in Figure 1. The smoother was 2 steps of SGS. Larger entries indicate smoother couplings. In the figure, it is now clear that the couplings to the points in the grid lines above and below the given point are equally non-smooth.

22.7	20.8	21.8
51.7		54.8
15.5	15.4	15.9

Figure 5: Entries in the smoothness matrix for the stencil in Figure 1.

Figure 6 plots the weighted graph of the smoothness matrix for an anisotropic diffusion problem using stretched elements on a segment of an annulus. The edges of the graph are colored according to the entries in the smoothness matrix. As expected for this problem, geometrically smooth connections are in the radial direction. Some nearly horizontal connections can also be classified as smooth due to the anisotropy preferring the x direction. The boundary regions are shown to be less smooth since error at the boundaries is removed fastest.

Domain decomposition

As mentioned, the smoothness matrix is adapted to the smoother being used. Consider a non-overlapping block Jacobi smoother (using SGS for approximate solves within the blocks) for a diffusion problem partitioned for two processors. Figure 7 shows that the smoothness matrix detects the geometric non-smoothness across the processor boundary in the algebraically smooth error. The illustration is similar for problems with jumps in PDE coefficients.

Algorithm 1 Calculate smoothness matrix S (without local scaling)

Define the number of samples R
Initialize the sparse data structure for S, given the chosen sparsity pattern of M
Construct R samples of algebraically smooth error $e^{(k)}$, $k = 1, \ldots, R$ by relaxing $Ae = 0$ with random initial guesses
Scale $\tilde{e}^{(k)} = e^{(k)} / \|e^{(k)}\|_2$
for all (i, j) such that $i < j$ and (i, j) is in the pattern of M **do**
  $d = \frac{1}{\text{dist}(i,j) \cdot R} \sum_{k=1}^{R} |\tilde{e}_i^{(k)} - \tilde{e}_j^{(k)}|$
  $\{d$ can be zero if there are explicit zeros in the matrix $A\}$
  if $d \neq 0$ **then**
    $S_{ij} = 1/d$
  else
    $S_{ij} = 0$
  end if
  $S_{ji} = S_{ij}$
end for

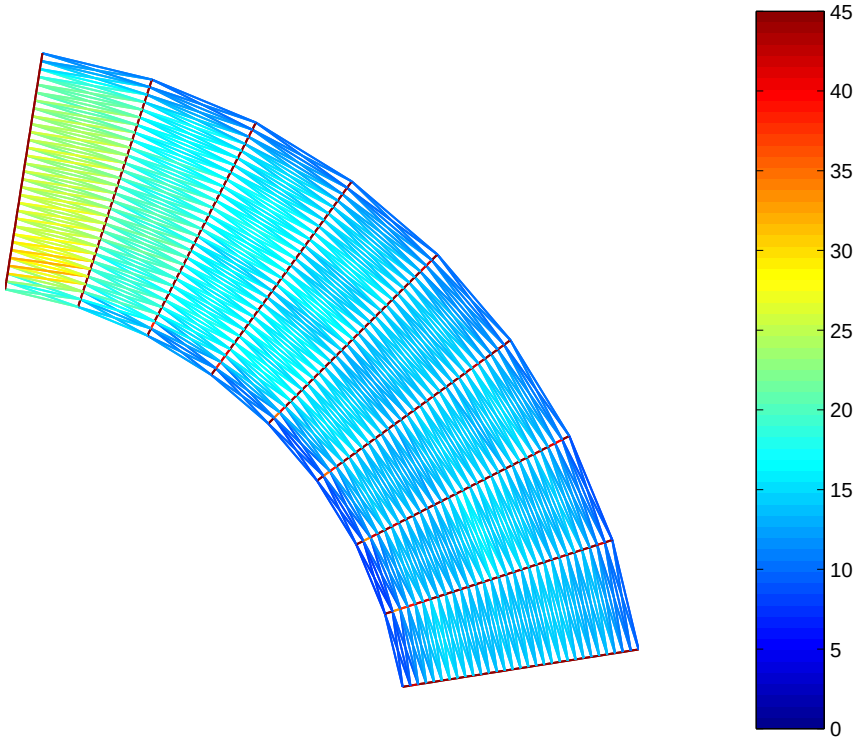

Figure 6: Weighted graph of the smoothness matrix for an anisotropic diffusion problem (coefficients $k_x = 10$, $k_y = 1$) using stretched elements on a segment of an annulus. The edges of the graph are colored according to the entries in the smoothness matrix.

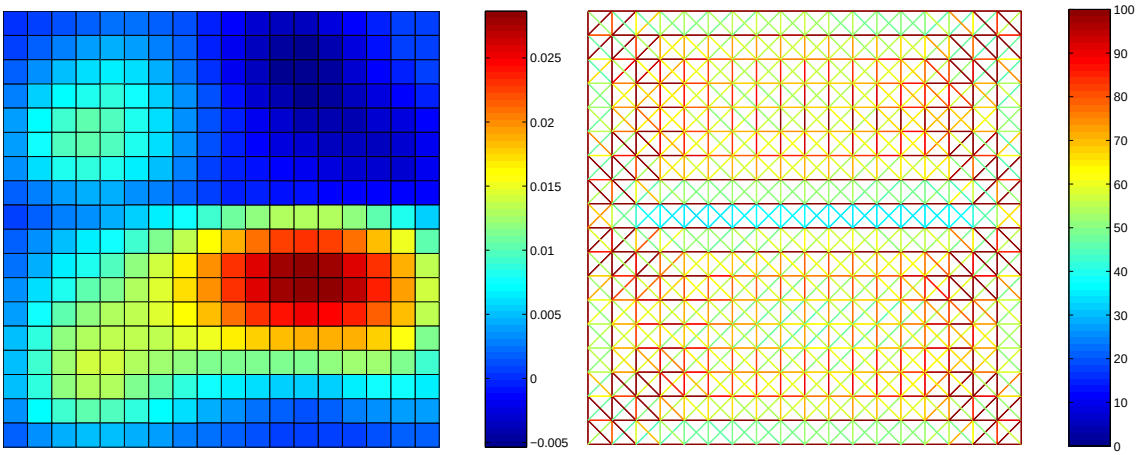

(a) Smooth error when two subdomains are used.

(b) Graph of smoothness matrix.

Figure 7: Smooth error and graph of smoothness matrix for a problem partitioned for parallel computing. (This figure should be viewed in color.)

Rotated anisotropic diffusion

The smoothness matrix can identify smooth couplings that are not immediate couplings in the graph of the matrix. For problems where strong anisotropy is not aligned with the grid points, these couplings may provide better interpolation. Figure 8 illustrates a rotated anisotropic diffusion problem, where the direction of anisotropy is in the horizontal direction. The smoothness matrix was computed using the sparsity pattern of A^2 less the diagonal. The graph corresponding to the largest entries in this smoothness matrix is plotted (the other entries complicate the graph). In passing, we note that geometric smoothness in the algebraically smooth error is not perfectly aligned with the anisotropy unless many smoothing steps are used.

3 Interpolation

In this very short section, we describe a few alternatives for choosing the interpolation matrix when the important couplings for a grid point are the geometrically smooth couplings defined above. These interpolation schemes are geometric. The common interpolations defined for AMG would not work, or would work poorly, because they would lead to many negative interpolation weights.

3.1 Linear interpolation

Given that points i and j are smoothly coupled, point j contributes to the value of point i proportionally to the reciprocal of its geometric distance from point i. If only couplings to coarse points are used, this is called linear interpolation using direct couplings only. The interpolation weights are scaled such that the constant vector is interpolated perfectly.

If j is not a coarse grid point, then point can j contribute to the value at point i *indirectly* using the coarse points that are smoothly coupled to point j. This is called linear interpolation with both direct and indirect couplings, and is similar to the way grid points are interpolated indirectly in AMG. Again, the weights are scaled to interpolate the constant vector perfectly.

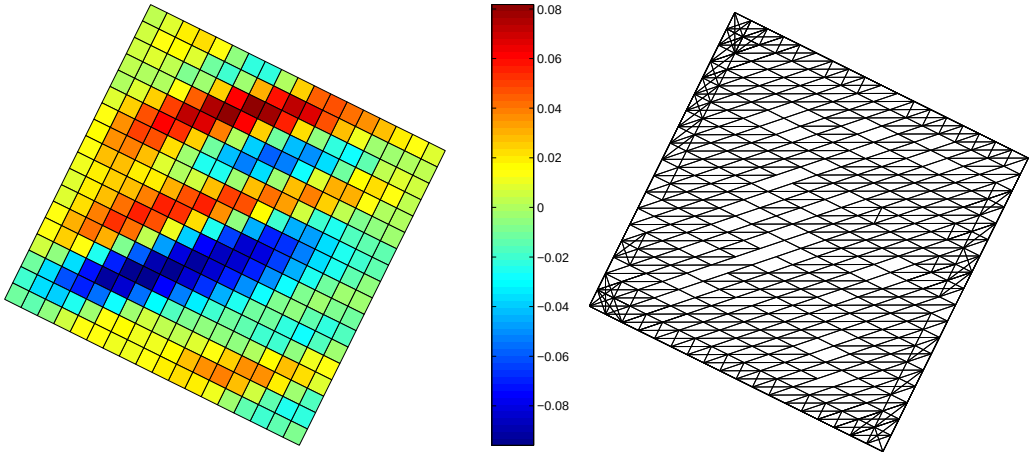

(a) Smooth error (coefficients $k_x = 10$, $k_y = 0.1$). (b) Smoothness graph.

Figure 8: Algebraically smooth error.

3.2 Interpolation using the geometric smoothness matrix

The entries in the smoothness matrix S may be used directly to define the interpolation weights. Large entries in S indicate geometric smoothness and provide better interpolation, while smaller entries in S should be less relied upon in interpolation. In our version, if a fine grid point is smoothly coupled to another fine grid point, indirect interpolation is used. This is called *GS-weighted* interpolation. Again, the weights are scaled to interpolate the constant vector perfectly. GS-weighted interpolation may be appropriate for problems that have multiple degrees of freedom at each grid point.

4 Numerical investigations

Geometrically smooth couplings may be used instead of AMG's strong couplings, and linear or GS-weighted interpolation may replace the interpolation in AMG to define a new unstructured multigrid method. We call this method GSMG (Geometric Smoothness Multigrid). When we need to refer to the specific variant of GSMG that uses local scaling of the smoothness matrix, we call it GSMG-L.

GSMG has been implemented by adding modules to BoomerAMG [8], an AMG code based on algorithms by Ruge and Stüben [10]. In this section, we first describe our test problems, then show results for AMG (using BoomerAMG), GSMG, GSMG-L, and GSMG using different types of interpolation. The tests were run on an 850 MHz Pentium III computer with 256 kbytes of cache memory and 256 Mbytes of main memory.

4.1 Test problems

Our test problems arise from preconditioning the iterative solution of the finite element equations $Ku = f$ for a 3-D elasticity problem. When K is ordered such that the x-, y-, and z-direction displacements are grouped among themselves, K has the block structure

$$K = \begin{bmatrix} K_{xx} & K_{xy} & K_{xz} \\ K_{yx} & K_{yy} & K_{yz} \\ K_{zx} & K_{zy} & K_{zz} \end{bmatrix}.$$

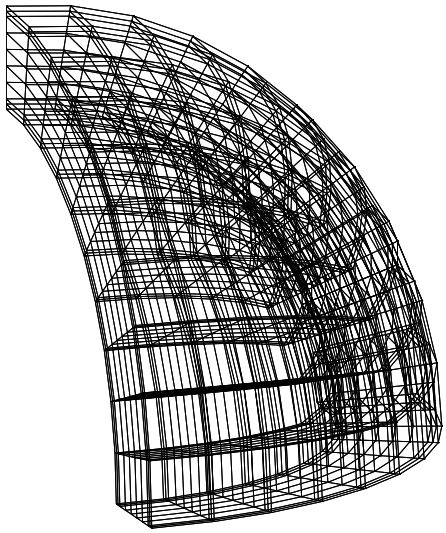

Figure 9: Gridding of an octant of three concentric spherical shells with (s_1, s_2, a_1) = (2, 5, 10).

The block diagonal matrix blockdiag(K_{xx}, K_{yy}, K_{zz}) is spectrally equivalent to K with respect to the meshsize parameter [1] and forms a good preconditioner if solves with the diagonal blocks are efficient. These blocks correspond to anisotropic second order operators, with the strength of the anisotropy depending on the material Poisson ratio. Our test problems are the diagonal blocks K_{xx} from different finite element discretizations of the same physical problem.

The physical problem of interest is three concentric spherical shells; two steel shells surround a third shell composed of lucite. Thus there are material coefficient discontinuities in this problem. An octant of the physical problem is discretized using linear hexahedral finite elements on a block-structured grid. The steel shells are 0.5 units thick, the lucite shell is 2.0 units thick, and the outer radius of the outer shell is 20.5 units. In the discretization, the steel shells are s_1 elements thick, the lucite shell is s_2 elements thick, and a_1 elements are used in the angular direction along the side of an octant. Figure 9 shows a gridding of the problem with (s_1, s_2, a_1) = (2, 5, 10). As is typical for these test problems, the elements have poor aspect ratios.

Table 1 describes the three test problems used in the following numerical investigations. The second and third problems are more difficult since they have poorer element aspect ratios. The third problem is a scaled-up version of the second problem, and allows a comparison of results on large and small problems.

	s_1	s_2	a_1	aspect ratio	n	nnz
UU-1	3	10	40	1/4.8	21437	538069
UU-2	3	10	20	1/9.7	5627	136759
UU-3	6	20	40	1/9.7	41613	1065157

Table 1: Test problems: the K_{xx} matrices for three griddings of the spherical shells problem, showing the gridding parameters, the worst element aspect ratio, and the number of equations n and number of nonzeros nnz in the matrices.

One-point integration for the finite elements is used, combined with hourglass damping [7] to eliminate spurious zero modes. We have found experimentally that this integration and damping procedure leads to matrices whose algebraically smooth error is very geometrically non-smooth unless very many smoothing steps are used. When only 1 or 2 smoothing steps are used, as is typical, AMG and GSMG converge very slowly. For this reason, in our tests below, we will use

up to 20 steps for the SGS smoother. Both pre- and post-smoothing are used.

The results tabulated in the sections below show the number of V-cycles for convergence (assumed when the initial residual (with a zero initial guess) is reduced by six orders of magnitude), the convergence rate on the final step (not the average convergence rate), the setup and solve timings, and the grid and operator complexities. Further, the tables show the number of smooth or strong couplings (S-couplings) found on the finest grid, the number of coarse points (C-points) found on the finest grid, and the number of levels.

4.2 AMG results

Tables 2–4 show sample results for AMG (many other strength thresholds θ were tested). The best convergence rates are achieved for θ from about 0.75 to 1, which is much higher than the nominal value of 0.25 suggested in the literature. Higher values of θ are required to ensure that the strong couplings are indeed good couplings to use for interpolation. The best timings are achieved for θ of 0.95 (not shown) for UU-1 and 0.99 (not shown) or 1.00 for the other two problems. This is because the grid and operator complexities can be extremely large, and the best timings are achieved when these complexities are small. AMG is not scalable for these large values of θ (based on comparing UU-2 and UU-3 results); AMG is scalable for θ where the best convergence rate is achieved. In summary, better convergence rates and complexities are limited for these problems due to poor choices of strong couplings.

20 step smoother									
θ	S-couplings (1st lev)	C-points (1st lev)	levels	cycles	conv. rate	time (s)		complexity	
						setup	solve	grid	operator
0.00	347954	12023	6	13	0.433	37.11	144.08	1.72	5.57
0.25	245480	10678	12	15	0.481	11.38	424.59	2.16	14.49
0.50	125917	9998	14	25	0.657	4.85	494.85	2.21	10.31
0.75	54094	10075	12	8	0.304	1.49	66.83	2.12	4.24
1.00	20782	9661	11	17	0.587	0.47	63.93	1.82	1.80

10 step smoother									
θ	S-couplings (1st lev)	C-points (1st lev)	levels	cycles	conv. rate	time (s)		complexity	
						setup	solve	grid	operator
0.00	347954	12023	6	18	0.542	37.14	100.13	1.72	5.57
0.25	245480	10678	12	20	0.581	11.46	278.44	2.16	14.49
0.50	125917	9998	14	30	0.709	4.90	302.76	2.21	10.31
0.75	54094	10075	12	9	0.354	1.49	37.73	2.12	4.24
1.00	20782	9661	11	22	0.664	0.50	40.43	1.82	1.80

5 step smoother									
θ	S-couplings (1st lev)	C-points (1st lev)	levels	cycles	conv. rate	time (s)		complexity	
						setup	solve	grid	operator
0.00	347954	12023	6	25	0.643	36.89	71.88	1.72	5.57
0.25	245480	10678	12	29	0.679	11.45	210.50	2.16	14.49
0.50	125917	9998	14	38	0.753	4.86	195.20	2.21	10.31
0.75	54094	10075	12	12	0.406	1.49	26.78	2.12	4.24
1.00	20782	9661	11	29	0.737	0.49	29.54	1.82	1.80

Table 2: AMG results for UU-1 as a function of the strength threshold θ.

20 step smoother									
θ	S-couplings (1st lev)	C-points (1st lev)	levels	cycles	conv. rate	time (s) setup	time (s) solve	complexity grid	complexity operator
0.00	83456	2480	6	39	0.769	3.01	72.88	1.59	4.12
0.25	64457	2480	11	30	0.720	1.83	132.21	2.06	9.64
0.50	30350	2480	12	43	0.784	0.66	119.46	2.19	6.47
0.75	14441	2480	11	11	0.397	0.21	11.87	1.95	2.61
1.00	5271	2477	9	12	0.455	0.11	8.57	1.80	1.66

10 step smoother									
θ	S-couplings (1st lev)	C-points (1st lev)	levels	cycles	conv. rate	time (s) setup	time (s) solve	complexity grid	complexity operator
0.00	83456	2480	6	58	0.851	3.02	58.36	1.59	4.12
0.25	64457	2480	11	45	0.808	1.82	100.95	2.06	9.64
0.50	30350	2480	12	53	0.822	0.65	76.68	2.19	6.47
0.75	14441	2480	11	19	0.589	0.21	10.67	1.95	2.61
1.00	5271	2477	9	14	0.515	0.10	5.44	1.80	1.66

5 step smoother									
θ	S-couplings (1st lev)	C-points (1st lev)	levels	cycles	conv. rate	time (s) setup	time (s) solve	complexity grid	complexity operator
0.00	83456	2480	6	91	0.902	3.00	45.27	1.59	4.12
0.25	64457	2480	11	66	0.858	1.83	77.14	2.06	9.64
0.50	30350	2480	12	65	0.852	0.66	49.95	2.19	6.47
0.75	14441	2480	11	34	0.743	0.22	10.18	1.95	2.61
1.00	5271	2477	9	19	0.586	0.11	3.88	1.80	1.66

Table 3: AMG results for UU-2 as a function of the strength threshold θ.

20 step smoother									
θ	S-couplings (1st lev)	C-points (1st lev)	levels	cycles	conv. rate	time (s) setup	time (s) solve	complexity grid	complexity operator
0.00	679552	19520	7	38	0.777	188.09	872.36	1.62	5.69
0.25									
0.50									
0.75	115709	20776	14	11	0.407	2.45	150.53	2.07	3.42
1.00	40272	19332	11	22	0.671	0.93	163.85	1.83	1.76

10 step smoother									
θ	S-couplings (1st lev)	C-points (1st lev)	levels	cycles	conv. rate	time (s) setup	time (s) solve	complexity grid	complexity operator
0.00	679552	19520	7	55	0.842	184.26	602.29	1.62	5.69
0.25									
0.50									
0.75	115709	20776	14	18	0.586	2.46	125.43	2.07	3.42
1.00	40272	19332	11	27	0.717	0.94	99.97	1.83	1.76

5 step smoother									
θ	S-couplings (1st lev)	C-points (1st lev)	levels	cycles	conv. rate	time (s) setup	time (s) solve	complexity grid	complexity operator
0.00	679552	19520	7	81	0.889	180.65	490.04	1.62	5.69
0.25									
0.50									
0.75	115709	20776	14	32	0.740	2.44	118.83	2.07	3.42
1.00	40272	19332	11	37	0.775	0.94	73.55	1.83	1.76

Table 4: AMG results for UU-3 as a function of the strength threshold θ. Blanks in the table indicate that AMG failed due to excessive memory requirements.

4.3 GSMG results, and number of sample vectors

Tables 5–7 show results for GSMG as a function of the number of sample relaxed vectors used. The smoothness threshold was chosen automatically. It is immediately clear that the grid and operator complexities are much smaller than those for AMG, although more couplings are used for interpolation. A related result is that GSMG uses fewer coarse grid points and fewer levels. The GSMG convergence rate is generally better or comparable, except when only 5 steps are used in the smoother.

GSMG has a relatively large setup cost, but the lower operator complexity generally makes the solve cost much lower than AMG's solve cost. This can be advantageous in situations where multiple systems need to be solved with the same matrix. The total (setup and solve) time may be larger or smaller depending on the problem.

For these problems, GSMG requires about 5 to 10 sample relaxed vectors to give the best total time. However, the number of V-cycles and the solve time can be reduced further if more sample vectors are used, up to about 20 or 25. The best number of samples to use does not seem to depend on the number of smoother steps.

GSMG operates scalably (by comparing results for UU-2 and UU-3) for these problems when a large number of smoothing steps (e.g., 20) is used.

For UU-2 and UU-3, the number of smooth couplings seems to increase when the number of smoothing steps is decreased. This implies that 1) the selected smoothness threshold depends on the number of smoothing steps, and/or 2) the relaxed vectors appear geometrically smoother if less smoothing is used.

20 step smoother									
sample	S-couplings	C-points	levels	cycles	conv.	time (s)		complexity	
vectors	(1st lev)	(1st lev)			rate	setup	solve	grid	operator
5	413118	3982	5	9	0.291	9.87	26.39	1.22	1.40
10	395234	4282	7	7	0.209	17.77	22.29	1.27	1.54
15	401322	4282	7	7	0.203	24.77	21.52	1.26	1.50
20	404530	4138	7	7	0.210	32.19	21.35	1.26	1.47
25	389628	4562	6	7	0.217	40.23	21.61	1.27	1.52

10 step smoother									
sample	S-couplings	C-points	levels	cycles	conv.	time (s)		complexity	
vectors	(1st lev)	(1st lev)			rate	setup	solve	grid	operator
5	397814	4278	5	13	0.413	6.62	19.57	1.24	1.46
10	378938	4762	6	10	0.318	10.15	17.29	1.29	1.57
15	370308	4801	5	9	0.344	13.91	14.10	1.28	1.53
20	378622	4639	6	8	0.299	17.16	12.30	1.28	1.53
25	363602	5147	6	9	0.314	21.54	13.97	1.29	1.52

5 step smoother									
sample	S-couplings	C-points	levels	cycles	conv.	time (s)		complexity	
vectors	(1st lev)	(1st lev)			rate	setup	solve	grid	operator
5	460252	3573	5	20	0.607	4.52	15.63	1.20	1.36
10	470442	3238	6	18	0.541	5.79	13.36	1.18	1.29
15	471892	3233	5	18	0.533	7.71	13.43	1.18	1.29
20	480870	2978	5	17	0.508	9.02	12.19	1.17	1.24
25	488626	2828	6	19	0.604	10.35	13.30	1.16	1.20

Table 5: GSMG results for UU-1 as a function of the number of sample relaxed vectors.

20 step smoother									
sample	S-couplings	C-points	levels	cycles	conv.	time (s)		complexity	
vectors	(1st lev)	(1st lev)			rate	setup	solve	grid	operator
5	66340	1537	6	13	0.421	2.45	9.96	1.37	1.78
10	60420	1646	6	10	0.359	4.24	7.52	1.40	1.77
15	50024	1881	7	9	0.362	6.38	7.06	1.44	1.83
20	48390	1894	6	9	0.334	8.09	6.84	1.43	1.77
25	51474	1831	6	9	0.351	9.80	6.80	1.43	1.75

10 step smoother									
sample	S-couplings	C-points	levels	cycles	conv.	time (s)		complexity	
vectors	(1st lev)	(1st lev)			rate	setup	solve	grid	operator
5	93754	1131	5	34	0.754	1.39	12.01	1.27	1.53
10	83410	1336	6	16	0.503	2.31	6.05	1.34	1.64
15	77654	1429	6	17	0.539	3.37	6.47	1.34	1.62
20	70336	1553	6	13	0.443	4.21	4.99	1.37	1.65
25	65636	1633	6	15	0.495	5.44	6.01	1.39	1.74

5 step smoother									
sample	S-couplings	C-points	levels	cycles	conv.	time (s)		complexity	
vectors	(1st lev)	(1st lev)			rate	setup	solve	grid	operator
5	119578	747	5	68	0.871	0.86	10.55	1.18	1.25
10	118906	778	5	54	0.827	1.25	8.56	1.18	1.28
15	117028	851	5	53	0.827	1.66	8.55	1.19	1.31
20	114298	849	5	47	0.822	2.07	7.67	1.18	1.29
25	110540	957	5	43	0.808	2.51	7.19	1.20	1.32

Table 6: GSMG results for UU-2 as a function of the number of sample relaxed vectors.

20 step smoother									
sample	S-couplings	C-points	levels	cycles	conv.	time (s)		complexity	
vectors	(1st lev)	(1st lev)			rate	setup	solve	grid	operator
5	687806	9621	6	19	0.579	23.95	129.46	1.30	1.65
10	671350	10134	6	13	0.473	39.51	90.91	1.32	1.67
15	633350	10810	8	11	0.413	56.62	76.65	1.35	1.72
20	620908	11043	7	12	0.434	73.68	83.83	1.36	1.71
25	599488	11275	7	11	0.392	93.43	78.89	1.37	1.74

10 step smoother									
sample	S-couplings	C-points	levels	cycles	conv.	time (s)		complexity	
vectors	(1st lev)	(1st lev)			rate	setup	solve	grid	operator
5	834420	7656	6	35	0.757	13.03	106.01	1.22	1.47
10	859320	7265	6	28	0.693	19.30	80.89	1.21	1.38
15	887646	6763	6	37	0.765	25.22	102.31	1.19	1.34
20	898644	6672	6	28	0.678	31.55	77.86	1.19	1.34
25	896832	6757	7	23	0.637	39.67	67.34	1.22	1.41

5 step smoother									
sample	S-couplings	C-points	levels	cycles	conv.	time (s)		complexity	
vectors	(1st lev)	(1st lev)			rate	setup	solve	grid	operator
5	979128	5302	6	67	0.857	7.59	94.71	1.15	1.21
10	980378	5173	6	58	0.833	10.67	81.14	1.15	1.19
15	977984	5400	7	55	0.830	14.28	78.28	1.16	1.21
20	976900	5249	6	57	0.831	17.60	80.29	1.15	1.19
25	982190	5574	7	55	0.831	20.72	77.33	1.16	1.19

Table 7: GSMG results for UU-3 as a function of the number of sample relaxed vectors.

4.4 GSMG-L results

The GSMG-L results in Tables 8–10 show that the grid and operator complexities are larger than those of GSMG, but smaller than those of AMG. Further, GSMG-L is able to achieve slightly better convergence rates than both GSMG and AMG. However, the total time for GSMG-L is larger than the total time for GSMG due to its larger complexities, and the total time for GSMG-L may be larger or smaller than the total time for AMG depending on its setup cost. Twenty sample vectors were used in these tests.

Whereas AMG needs approximately the same number of C-points regardless of the strength threshold, GSMG-L needs far fewer C-points when the smoothness threshold is small. Thus GSMG-L seems to have a better distribution of its smooth couplings than AMG has of its strong couplings.

Like AMG, the GSMG-L results are somewhat problematic since the best timings are generally achieved when the grid and operator complexities are smallest rather than when the convergence rate is fastest. This similarity in the AMG and GSMG-L results can be explained by the similar way they use local scalings to select the relevant couplings.

20 step smoother									
τ	S-couplings (1st lev)	C-points (1st lev)	levels	cycles	conv. rate	time (s) setup	time (s) solve	complexity grid	complexity operator
0.00	499542	2400	5	8	0.228	23.81	17.92	1.13	1.12
0.25	277751	6648	8	6	0.158	69.77	37.84	1.47	3.17
0.50	147494	9103	10	6	0.193	84.33	48.75	1.74	4.04
0.75	53504	10820	12	7	0.248	61.33	41.88	2.03	3.03
1.00	20740	9976	11	12	0.485	35.32	41.71	1.83	1.65

10 step smoother									
τ	S-couplings (1st lev)	C-points (1st lev)	levels	cycles	conv. rate	time (s) setup	time (s) solve	complexity grid	complexity operator
0.00	499542	2400	5	13	0.433	13.91	16.16	1.13	1.12
0.25	287391	6558	7	8	0.279	36.25	23.88	1.46	3.01
0.50	145395	9248	9	6	0.209	43.14	23.95	1.75	4.06
0.75	51801	10845	12	8	0.306	30.32	23.22	2.03	2.98
1.00	20740	9961	11	18	0.619	17.62	30.74	1.83	1.65

5 step smoother									
τ	S-couplings (1st lev)	C-points (1st lev)	levels	cycles	conv. rate	time (s) setup	time (s) solve	complexity grid	complexity operator
0.00	499542	2400	5	25	0.648	8.01	16.18	1.13	1.12
0.25	406307	4894	7	23	0.643	17.74	28.42	1.33	2.30
0.50	211570	8771	9	11	0.365	25.96	24.26	1.71	4.18
0.75	67649	11165	12	9	0.378	19.55	16.04	2.06	3.33
1.00	20740	9607	10	27	0.721	10.01	25.24	1.80	1.65

Table 8: GSMG-L results for UU-1 as a function of the smoothness threshold τ.

4.5 Alternative interpolations

Table 11 shows the number of GSMG V-cycles required for convergence using linear interpolation, linear interpolation using direct connections only, and interpolation using weights from the smoothness matrix (GS-weighted). AMG interpolation did not lead to convergence when the relevant couplings were based on geometric smoothness.

The results show essentially the same performance except when GS-weighted interpolation was used with a small number of smoothing steps. Thus the use of geometric coordinates (for

20 step smoother									
τ	S-couplings (1st lev)	C-points (1st lev)	levels	cycles	conv. rate	time (s) setup	time (s) solve	complexity grid	complexity operator
0.00	122522	600	5	22	0.638	5.74	11.87	1.12	1.11
0.25	64953	1736	6	8	0.252	13.12	9.50	1.48	2.81
0.50	23694	2443	8	7	0.230	14.85	10.07	1.75	3.30
0.75	10607	2592	10	8	0.242	9.24	7.27	1.90	2.25
1.00	5270	2552	9	10	0.355	7.10	6.99	1.81	1.64

10 step smoother									
τ	S-couplings (1st lev)	C-points (1st lev)	levels	cycles	conv. rate	time (s) setup	time (s) solve	complexity grid	complexity operator
0.00	122522	600	5	44	0.796	2.98	11.58	1.12	1.11
0.25	67244	1778	7	14	0.467	8.26	9.43	1.48	2.87
0.50	24567	2459	8	8	0.254	7.99	6.00	1.76	3.33
0.75	10009	2641	10	9	0.284	5.19	4.49	1.92	2.30
1.00	5270	2536	9	13	0.476	3.97	5.01	1.80	1.63

5 step smoother									
τ	S-couplings (1st lev)	C-points (1st lev)	levels	cycles	conv. rate	time (s) setup	time (s) solve	complexity grid	complexity operator
0.00	122522	600	5	86	0.891	1.86	12.99	1.12	1.11
0.25	95171	1317	6	54	0.839	3.54	13.42	1.35	2.11
0.50	45808	2357	8	23	0.646	5.60	11.08	1.73	3.92
0.75	14364	2787	10	12	0.408	3.29	3.62	1.95	2.63
1.00	5270	2449	9	19	0.591	2.11	3.79	1.77	1.61

Table 9: GSMG-L results for UU-2 as a function of the smoothness threshold τ.

20 step smoother									
τ	S-couplings (1st lev)	C-points (1st lev)	levels	cycles	conv. rate	time (s) setup	time (s) solve	complexity grid	complexity operator
0.00	989686	4800	5	20	0.599	48.39	90.90	1.13	1.12
0.25	517821	15008	8	7	0.225	253.85	150.08	1.59	5.45
0.50	170967	19798	11	7	0.241	169.43	114.48	1.87	4.11
0.75	86516	19297	13	8	0.305	107.46	84.24	1.96	2.55
1.00	40260	18689	11	17	0.591	71.09	119.27	1.76	1.64

10 step smoother									
τ	S-couplings (1st lev)	C-points (1st lev)	levels	cycles	conv. rate	time (s) setup	time (s) solve	complexity grid	complexity operator
0.00	989686	4800	5	38	0.769	28.17	95.87	1.13	1.12
0.25	577647	13944	9	24	0.653	165.16	252.02	1.56	5.42
0.50	190193	19778	11	11	0.389	90.17	92.10	1.88	4.23
0.75	81996	19825	13	10	0.371	54.97	52.55	1.98	2.57
1.00	40260	18164	11	26	0.703	34.18	86.01	1.74	1.60

5 step smoother									
τ	S-couplings (1st lev)	C-points (1st lev)	levels	cycles	conv. rate	time (s) setup	time (s) solve	complexity grid	complexity operator
0.00	989686	4800	5	74	0.875	16.19	97.16	1.13	1.12
0.25	846116	9238	8	57	0.854	70.70	190.16	1.35	3.21
0.50	404468	17774	10	36	0.740	64.96	197.95	1.80	5.29
0.75	124615	21112	13	19	0.611	38.54	66.83	2.03	3.20
1.00	40260	17338	11	40	0.798	19.24	71.99	1.70	1.58

Table 10: GSMG-L results for UU-3 as a function of the smoothness threshold τ.

linear interpolation) is not mandatory in GSMG.

	smoother steps	linear	linear (direct only)	GS-weighted
UU-1	20	7	7	7
	10	8	9	8
	5	17	18	20
UU-2	20	9	9	8
	10	13	16	13
	5	47	44	60
UU-3	20	12	13	11
	10	28	26	31
	5	57	54	75

Table 11: GSMG results showing the number of V-cycles required for convergence using linear interpolation, linear interpolation using direct connections only, and interpolation using weights from the smoothness matrix.

5 Conclusions

This paper has argued that geometric smoothness may more accurately determine the relevant couplings in AMG for matrices that are not diagonally dominant. Determining the geometrically smooth couplings comes at additional setup cost, but methods based on geometric smoothness, such as GSMG, may have much smaller solve timings. The smaller solve timings are primarily the result of smaller operator complexities. On the negative side, GSMG is much more sensitive to the geometric smoothness of the algebraically smooth error. GSMG may require many smoothing steps for some problems.

One of the main differences between AMG and GSMG is that AMG uses a local definition of strong couplings (i.e., the strength of the coupling is based on local information) whereas in GSMG, a global threshold is used to determine whether a coupling is smooth for the entire grid. The latter strategy seems to lead to a method that results in better complexities for the finite element problems tested here.

The geometric character of algebraically smooth error needs to be better understood. For example, can smooth error vary sharply in some regions of the grid, but can still be interpolated geometrically? Methods can be designed that test how well neighboring points provide interpolation in order to decide on the most relevant couplings for a grid point.

Acknowledgments

The author is grateful to Ulrike Meier Yang for advice on using the BoomerAMG code. The author also wishes to thank Robert D. Falgout, Van Emden Henson, Jim E. Jones, and Panayot S. Vassilevski for their helpful comments during this work.

References

[1] O. Axelsson. On iterative solvers in structural mechanics; separate displacement orderings and mixed variable methods. *Mathematics and Computers in Simulation*, 40:11–30, 1999.

[2] A. Brandt. Algebraic multigrid theory: The symmetric case. *Appl. Math. Comput.*, 19:23–56, 1986.

[3] A. Brandt. General highly accurate algebraic coarsening. *Elect. Trans. Numer. Anal.*, 10:1–20, 2000.

[4] A. Brandt. Multiscale scientific computation: Review 2000, 2000.

[5] M. Brezina, A. J. Cleary, R. D. Falgout, V. E. Henson, J. E. Jones, T. A. Manteuffel, S. F. McCormick, and J. W. Ruge. Algebraic multigrid based on element interpolation (AMGe). *SIAM J. Sci. Comput.*, 22:1570–1592, 2000.

[6] W. L. Briggs, V. E. Henson, and S. F. McCormick. *A Multigrid Tutorial.* SIAM Books, Philadelphia, 2000. Second edition.

[7] D. P. Flanagan and T. Belytschko. A uniform strain hexahedron and quadrilateral with orthogonal hourglass control. *Intl. J. Num. Meth. Engrg.*, 17:679–706, 1981.

[8] V. E. Henson and U. M. Yang. BoomerAMG: a parallel algebraic multigrid solver and preconditioner. *Applied Numerical Mathematics*, to appear.

[9] W. Z. Huang. Convergence of algebraic multigrid methods for symmetric positive definite matrices with weak diagonal dominance. *Appl. Math. Comput.*, 46:145–164, 1991.

[10] J. W. Ruge and K. Stüben. Algebraic multigrid (AMG). In S. F. McCormick, editor, *Multigrid Methods*, volume 3 of *Frontiers in Applied Mathematics*, pages 73–130. SIAM, Philadelphia, PA, 1987.

[11] K. Stüben. Algebraic multigrid (AMG): An introduction with applications. In U. Trottenberg, C. Oosterlee, and A. Schuller, editors, *Multigrid.* Academic Press, 2000.

[12] K. Stüben. Some studies on algebraic multigrid (AMG). In *International Workshop on Algebraic Multigrid Methods*, St. Wolfgang, 2000.