

Adaptive reduction-based AMG[‡]Scott MacLachlan^{1,2,*}, Tom Manteuffel¹, Steve McCormick¹¹ *Department of Applied Mathematics, University of Colorado at Boulder, UCB 526, Boulder, CO, 80309-0526*² *Now at: Department of Computer Science and Engineering, University of Minnesota, 4-192 EE/CS Building, 200 Union Street SE, Minneapolis, MN, 55455*

SUMMARY

With the ubiquity of large-scale computing resources has come significant attention to practical details of fast algorithms for the numerical solution of partial differential equations. Included in this group are the class of multigrid and algebraic multigrid algorithms that are effective solvers for many of the large matrix problems arising from the discretization of elliptic operators. Algebraic multigrid (AMG) is especially effective for many problems with discontinuous coefficients, discretized on unstructured grids, or over complex geometries. While much effort has been invested in improving the practical performance of AMG, little theoretical understanding of this performance has emerged. This paper presents a two-level convergence theory for a reduction-based variant of AMG, called AMGr, which is particularly appropriate for linear systems that have M-matrix-like properties. For situations where less is known about the problem matrix, an adaptive version of AMGr that automatically determines the form of the reduction needed by the AMGr process is proposed. The adaptive cycle is shown, in both theory and practice, to yield an effective AMGr algorithm.

Copyright © 2006 John Wiley & Sons, Ltd.

KEY WORDS: iterative methods, multigrid, algebraic multigrid, adaptive multigrid, multigrid theory

1. Introduction

While modern scientific computing has benefited greatly from improvements in computational technology, our ability to simulate complex physical systems owes as much to improved algorithms, particularly in the area of linear systems solvers. The matrix equations that arise from discretizing partial differential equations (PDEs) can often be solved most efficiently using a multiscale solver, such as multigrid. Algebraic multigrid (AMG) [1, 2] offers many improvements in robustness and ease of use over classical geometric multigrid, especially for discretizations on complex geometries or irregular meshes, or when the coefficients of the PDE vary widely or are discontinuous.

[‡]This is a preprint of an article published in *Numerical Linear Algebra with Applications*, Volume 13, Pages 599-620. Copyright ©2006, John Wiley & Sons, Ltd.

*Correspondence to: maclach@cs.umn.edu

The performance of AMG has been assessed (c.f., [3]), and modifications proposed in many contexts, for a wide variety of problems on a wide variety of computer architectures. Ruge and Stüben give many examples and describe system variants of AMG in [2]. Parallel implementations of the AMG algorithm are discussed in [4, 5]. Within a parallel AMG algorithm, many details of the multigrid algorithm may also have a significant effect, such as the parameters of the relaxation scheme chosen [6] or the choice of coarse grids and interpolation operators [7].

While these advances play an important role in the applicability of AMG to many linear systems, there is less theoretical understanding of the performance of AMG. Brandt analyzes the evolution of errors and residuals through relaxation and coarse-grid correction, based on the entries of the matrix [8]. More recently, the element-based AMG (AMGe) framework [9, 10] provides insight into the choice of multigrid interpolation operators when individual element stiffness matrices are known in addition to the global linear system. Falgout et al. [11, 12] present a sharp theoretical analysis of the performance of an AMG algorithm, yet their measures are, in general, difficult both to compute exactly and to approximate. Here, we present a two-level variant of AMG, called AMGr, that is based on approximating the Schur complement by an equivalent operator. The performance bound for the resulting two-grid cycle depends only on the equivalence constants.

While classical AMG is robust in several respects, it can break down for problems that are not properly elliptic or, more precisely, for problems whose near-null-space components are not well understood. AMG relies on an algebraic sense of smoothness, which refers loosely to the character of error components that give relatively small residuals. Classical AMG uses knowledge of the local nature of this smoothness to “collapse” entries in certain rows of the matrix to obtain a suitable interpolation operator. Nominally, it assumes that smooth errors are constant along strong connections. Classical smoothed aggregation (SA) [13] uses representative smooth vectors (e.g., rigid body modes in elasticity) to define columns of the interpolation operator locally. Unfortunately, when representative smooth components are neither available nor well understood, these algebraic schemes cannot be effectively applied.

It is this limitation that motivates the development of adaptive algebraic multigrid methods (α AMG and α SA [14, 15, 16]) that attempt to determine the sense of smoothness automatically. The basic idea is to apply relaxation to the homogeneous problem ($Ax = 0$) to determine which components are slow to converge. Coarsening is then constructed based on the resulting algebraically smooth component(s). Because these components may not be fully representative, the resulting algorithm itself is applied to the homogeneous problem and coarsening is adjusted to match any slow-to-converge error components that are found. This process is continued until good efficiency is, hopefully, achieved.

To improve our understanding of adaptive AMG, here we develop a two-level, reduction-based version of it, called adaptive AMGr (α AMGr). We analyze the adaptive setup cycle to recover the lowest-energy mode of the linear system and, using this, find an equivalent diagonal operator to use in the AMGr algorithm. Global convergence is possible in special cases; however, because of the nonlinearity of this iteration, only a local result is given for the most general case considered.

This paper is arranged as follows. In Section 2, the AMGr algorithm is explained and a convergence proof is given. The adaptive extension of this algorithm, α AMGr, is discussed in Section 3. Numerical experiments are given in Section 4, and concluding remarks are made in Section 5.

2. Reduction-Based AMG (AMGr)

Consider the following representation for the symmetric $n \times n$ matrix A :

$$A = \begin{bmatrix} A_{ff} & -A_{fc} \\ -A_{fc}^T & A_{cc} \end{bmatrix}.$$

Here, we assume the existence of a partition of \mathbb{R}^n into F and C points. The development of AMGr is motivated by two-level theory from a reduction point of view; for a multigrid scheme based on F -point relaxation, it is easy to see that the ideal interpolation operator from an approximation point of view is $P = \begin{bmatrix} A_{ff}^{-1} A_{fc} \\ I \end{bmatrix}$. Because this is usually not practical from a computational point of view (A_{ff}^{-1} is generally dense), we instead ask the question of what sort of approximation to A_{ff} by a more-easily inverted matrix, D , is necessary to achieve good multigrid performance.

Assume, then, that the original degrees of freedom (points or variables in \mathbb{R}^n) are partitioned into those that are associated with the coarse level (set C) and those that are not (set F): $\mathbb{R}^n = F \cup C$. Further, suppose that this is done in such a way that the F -to- F connections are subdominant (i.e., the submatrix associated with the fine-level-only points is significantly better conditioned than the whole matrix). Such a splitting can be achieved by using compatible relaxation [17, 11, 18] in the multigrid coarsening process. Here, we show that, given such a splitting, the multigrid convergence factors, measured in the A -norm, can be bounded by a constant less than one that depends only on the spectral equivalence bounds between A_{ff} and D , indicating uniform multigrid convergence in that norm.

In what follows, lower case Greek letters are used to denote scalars, while lower case Roman letters denote vectors. The matrix, A , is assumed to be real, symmetric, and positive definite. Additional assumptions are imposed on A that specify the F -to- C dominance we need. We show that this algorithm obtains uniform convergence, for any given scale of dominance of the F - F block.

Write $A_{ff} = D + \mathcal{E}$, where D is positive definite and \mathcal{E} is such that $0 \leq \mathcal{E} \leq \epsilon D$ for some $\epsilon > 0$. The notation $A \leq B$ is taken to mean that $x^T A x \leq x^T B x$, for all x . Assume further that A and D satisfy

$$A_D = \begin{bmatrix} D & -A_{fc} \\ -A_{fc}^T & A_{cc} \end{bmatrix} \geq 0. \quad (1)$$

If A is diagonally dominant, then such a partition of A_{ff} is always possible with a diagonal matrix, D [19, Lemma 8]. It is, however, possible to satisfy these conditions in much more general circumstances than diagonally dominant matrices or, even, M-matrices. For any given non-negative definite matrices A_{ff} and D for which $\mathcal{E} = A_{ff} - D$ is non-negative definite, ϵ is simply the largest eigenvalue of the generalized eigenvalue problem, $\mathcal{E}x = \lambda D x$; since we assume D to be positive definite, then ϵ is guaranteed to be finite. Thus, given any symmetric non-negative definite matrix of form A_D and any symmetric and non-negative definite matrix \mathcal{E} , matrix A with $A_{ff} = D + \mathcal{E}$ satisfies Equation (1). Thus, the theory presented here is more general than the usual M-matrix based AMG theory, as in [8, 2]. Of particular interest is the case where A and its partitioning come from some fixed process for a given problem size (such as discretization of a PDE) and D is computed by a given formula (e.g., as part of the diagonal of A_{ff} , computed so that the remaining part, $\mathcal{E} = A_{ff} - D$, is non-negative definite). In this

case, our interest and expectation is not just that finite ϵ exists, but that the bound achieved is independent of the size of A , indicating the efficacy of the AMG algorithm.

The two-level AMGr scheme is defined by its basic components, relaxation and coarse-level correction. For relaxation, we choose a weighted F -point-only relaxation whose error-propagation matrix is given by

$$REL = \left(I - \sigma \begin{bmatrix} D^{-1} & 0 \\ 0 & 0 \end{bmatrix} A \right), \quad (2)$$

where $\sigma = \frac{2}{2 + \epsilon}$. To complement this, a variational coarse-level correction using the interpolation operator,

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

is used. Note that the error-propagation matrix for the coarse-level correction is then given by

$$CLC = I - P(P^TAP)^{-1}P^TA.$$

The error propagation matrix for our two-grid scheme, which consists of one relaxation step followed by coarse-level correction, is given by

$$MG_2 = CLC \cdot REL.$$

The analysis is simplified by noticing that any $e \in \mathbb{R}^n$ can be written as the A -orthogonal sum

$$e = \alpha \begin{bmatrix} A_{ff}^{-1}A_{fc} \\ I \end{bmatrix} v + \beta \begin{bmatrix} I \\ 0 \end{bmatrix} w, \quad (3)$$

where $\|v\|_{\hat{A}_{cc}} = \|w\|_{A_{ff}} = 1$ and $\hat{A}_{cc} = A_{cc} - A_{fc}^T A_{ff}^{-1} A_{fc}$ is the Schur complement of A_{ff} in A . Here, we use the energy-norm notation, $\|x\|_B = \langle Bx, x \rangle^{\frac{1}{2}}$, where B is any symmetric, positive definite matrix.

Theorem 1. *Suppose we are given $A = \begin{bmatrix} A_{ff} & -A_{fc} \\ -A_{fc}^T & A_{cc} \end{bmatrix} \geq 0$ such that $A_{ff} = D + \mathcal{E}$, with D symmetric, $0 \leq \mathcal{E} \leq \epsilon D$, and $\begin{bmatrix} D & -A_{fc} \\ -A_{fc}^T & A_{cc} \end{bmatrix} \geq 0$. Define relaxation with error propagation operator $REL = (I - \sigma \begin{bmatrix} D^{-1} & 0 \\ 0 & 0 \end{bmatrix} A)$ for $\sigma = \frac{2}{2 + \epsilon}$, interpolation $P = \begin{bmatrix} D^{-1}A_{fc} \\ I \end{bmatrix}$, coarse-level correction with error propagation operator $CLC = I - P(P^TAP)^{-1}P^TA$, and the two-grid AMGr algorithm whose error propagation operator is given by $MG_2 = CLC \cdot REL$. Then,*

$$\|MG_2\|_A \leq \left(\frac{\epsilon}{1 + \epsilon} \left(1 + \left(\frac{\epsilon}{(2 + \epsilon)^2} \right) \right) \right)^{\frac{1}{2}} < 1. \quad (4)$$

Thus, the two-grid AMGr scheme converges uniformly for any fixed $\epsilon \geq 0$.

Proof: Let e be a unit vector in the A -norm, so that (3) implies that $\alpha^2 + \beta^2 = 1$. Considering the effect of relaxation on e , the A -orthogonal decomposition then gives us

$$\begin{aligned}
RELe &= e - \sigma \begin{bmatrix} D^{-1} & 0 \\ 0 & 0 \end{bmatrix} \begin{pmatrix} \beta A_{ff} w \\ \alpha \hat{A}_{cc} v - \beta A_{fc}^T w \end{pmatrix} \\
&= e - \beta \sigma \begin{pmatrix} D^{-1} A_{ff} w \\ 0 \end{pmatrix} \\
&= \alpha \begin{bmatrix} A_{ff}^{-1} A_{fc} \\ I \end{bmatrix} v + \beta \begin{bmatrix} I - \sigma D^{-1} A_{ff} \\ 0 \end{bmatrix} w.
\end{aligned}$$

However,

$$\begin{aligned}
\left\| \begin{bmatrix} I - \sigma D^{-1} A_{ff} \\ 0 \end{bmatrix} w \right\|_A &= \|(I - \sigma D^{-1} A_{ff}) w\|_{A_{ff}} \\
&\leq \rho \left(I - \sigma A_{ff}^{1/2} D^{-1} A_{ff}^{1/2} \right) \\
&\leq \max \left(\left| 1 - \frac{2}{2 + \epsilon} \right|, \left| \frac{2 + 2\epsilon}{2 + \epsilon} - 1 \right| \right) = \frac{\epsilon}{2 + \epsilon},
\end{aligned}$$

because $D \leq A_{ff} \leq (1 + \epsilon)D$ implies that $A_{ff}^{-1} \leq D^{-1} \leq (1 + \epsilon)A_{ff}^{-1}$. It is therefore easy to see that

$$RELe = \alpha \begin{bmatrix} A_{ff}^{-1} A_{fc} \\ I \end{bmatrix} v + \hat{\beta} \begin{bmatrix} I \\ 0 \end{bmatrix} \hat{w},$$

where $\|\hat{w}\|_{A_{ff}} = 1$ and

$$|\hat{\beta}| \leq |\beta| \left(\frac{\epsilon}{2 + \epsilon} \right). \quad (5)$$

To analyze the coarse-level correction step, let $\hat{e} = RELe$. We first bound the error after the coarse-level correction by the error modified by the interpolant, $\begin{bmatrix} D^{-1} A_{fc} \\ I \end{bmatrix} v$, that is optimally scaled (writing this scale as $\alpha\theta$, for convenience):

$$\begin{aligned}
\|CLC\hat{e}\|_A &= \min_u \left\| \hat{e} - \begin{bmatrix} D^{-1} A_{fc} \\ I \end{bmatrix} u \right\|_A \\
&\leq \min_{\theta} \left\| \alpha \begin{bmatrix} A_{ff}^{-1} A_{fc} \\ I \end{bmatrix} v + \hat{\beta} \begin{bmatrix} I \\ 0 \end{bmatrix} \hat{w} - \alpha\theta \begin{bmatrix} D^{-1} A_{fc} \\ I \end{bmatrix} v \right\|_A.
\end{aligned}$$

The key now is to write the correction as an A -orthogonal sum:

$$\begin{bmatrix} D^{-1} A_{fc} \\ I \end{bmatrix} v = \begin{bmatrix} A_{ff}^{-1} A_{fc} \\ I \end{bmatrix} v + \begin{bmatrix} (D^{-1} - A_{ff}^{-1}) A_{fc} \\ 0 \end{bmatrix} v,$$

which, together with the definitions of v and \hat{w} and the triangle inequality, yields

$$\begin{aligned} & \|CLC\hat{e}\|_A^2 \\ & \leq \min_{\theta} \left\| \alpha(1-\theta) \begin{bmatrix} A_{ff}^{-1}A_{fc} \\ I \end{bmatrix} v + \hat{\beta} \begin{bmatrix} I \\ 0 \end{bmatrix} \hat{w} - \alpha\theta \begin{bmatrix} (D^{-1} - A_{ff}^{-1})A_{fc} \\ 0 \end{bmatrix} v \right\|_A^2 \\ & = \min_{\theta} \left[\alpha^2(1-\theta)^2 \left\| \begin{bmatrix} A_{ff}^{-1}A_{fc} \\ I \end{bmatrix} v \right\|_A^2 + \left\| \hat{\beta}\hat{w} - \alpha\theta \begin{bmatrix} (D^{-1} - A_{ff}^{-1})A_{fc} \\ 0 \end{bmatrix} v \right\|_A^2 \right] \\ & \leq \min_{\theta} \left[\alpha^2(1-\theta)^2 + \left(\hat{\beta} + \alpha\theta \left\| \begin{bmatrix} (D^{-1} - A_{ff}^{-1})A_{fc} \\ 0 \end{bmatrix} v \right\|_A \right)^2 \right]. \end{aligned}$$

Note now that equation (1) implies that $A_{cc} \geq A_{fc}^T D^{-1} A_{fc}$, which, with the assumed bound on D , implies that

$$\begin{aligned} \left\| \begin{bmatrix} (D^{-1} - A_{ff}^{-1})A_{fc} \\ 0 \end{bmatrix} v \right\|_A^2 &= v^T A_{fc}^T A_{ff}^{-1/2} \left(A_{ff}^{1/2} D^{-1} A_{ff}^{1/2} - I \right)^2 A_{ff}^{-1/2} A_{fc} v \\ &\leq \epsilon v^T A_{fc}^T A_{ff}^{-1/2} \left(A_{ff}^{1/2} D^{-1} A_{ff}^{1/2} - I \right)^2 A_{ff}^{-1/2} A_{fc} v \\ &= \epsilon v^T A_{fc}^T \left(D^{-1} - A_{ff}^{-1} \right) A_{fc} v \\ &\leq \epsilon v^T \left(A_{cc} - A_{fc}^T A_{ff}^{-1} A_{fc} \right) v \\ &= \epsilon \|v\|_A^2 = \epsilon. \end{aligned}$$

Combining these last two bounds and using (5) yields

$$\|CLC\hat{e}\|_A^2 \leq \min_{\theta} \left[\alpha^2(1-\theta)^2 + \left(\left(\frac{\epsilon}{2+\epsilon} \right) \beta + \alpha\theta\sqrt{\epsilon} \right)^2 \right].$$

A little calculus shows that the optimal θ is

$$\theta = \frac{\alpha - \left(\frac{\epsilon}{2+\epsilon} \right) \sqrt{\epsilon} \beta}{(1+\epsilon)\alpha},$$

which, with a little algebra, simplifies the bound to

$$\|MG_2e\|_A^2 \equiv \|CLC\hat{e}\|_A^2 \leq \frac{\epsilon}{1+\epsilon} \left(\alpha + \left(\frac{\sqrt{\epsilon}}{2+\epsilon} \right) \beta \right)^2.$$

A little more calculus shows that the maximum of $\alpha + \left(\frac{\sqrt{\epsilon}}{2+\epsilon} \right) \beta$ over $\alpha \in [0, 1]$ (remembering that $\beta = \sqrt{1-\alpha^2}$) yields

$$\|MG_2e\|_A^2 \leq \frac{\epsilon}{1+\epsilon} \left(1 + \left(\frac{\sqrt{\epsilon}}{2+\epsilon} \right)^2 \right) = \frac{\epsilon \left((2+\epsilon)^2 + \epsilon \right)}{(1+\epsilon)(2+\epsilon)^2}.$$

It is easy to see that this expression is less than 1: $\frac{x}{1+x}$ is an increasing function of $x \geq 0$ and $\epsilon < \frac{(2+\epsilon)^2}{\epsilon}$ for $\epsilon > 0$, so

$$\frac{\epsilon}{1+\epsilon} < \frac{\frac{(2+\epsilon)^2}{\epsilon}}{1 + \frac{(2+\epsilon)^2}{\epsilon}} = \frac{(2+\epsilon)^2}{(2+\epsilon)^2 + \epsilon}.$$

□

Corollary 1. *Under the assumptions of Theorem 1, consider the two-level scheme with $\nu \geq 1$ relaxation sweeps followed by coarsening. The error propagation matrix for this scheme is $MG_2 = CLC \cdot REL^\nu$, and*

$$\|MG_2\|_A \leq \left(\frac{\epsilon}{1+\epsilon} \left(1 + \left(\frac{\epsilon}{2+\epsilon} \right)^{2(\nu-1)} \left(\frac{\epsilon}{(2+\epsilon)^2} \right) \right) \right)^{\frac{1}{2}}.$$

Note that $\frac{\epsilon}{2+\epsilon} < 1$. Hence, this factor monotonically decreases as ν increases, with limit $\frac{\epsilon}{1+\epsilon}$, the convergence factor for the two-level scheme based on exact F -point-only relaxation.

Proof: It is easy to see that (5) becomes

$$|\hat{\beta}| \leq |\beta| \left(\frac{\epsilon}{2+\epsilon} \right)^\nu$$

and that the remainder of the proof of Theorem 1 holds with appropriate modification.
□

Our convergence bound depends only on the equivalence bound, $D \leq A_{ff} \leq (1+\epsilon)D$. Achieving good efficiency then requires a balance in the choice of D such that the parameter, ϵ , is small and relaxation and coarse-grid correction are easy to perform. At one extreme, choosing $D = A_{ff}$ shows that this is a direct method - the coarse-scale operator is, simply, the Schur complement, and relaxation is a direct solve of the fine-scale equations. At the other extreme, choosing D to be diagonal typically yields a nonzero convergence factor (albeit smaller than 1), but each iteration is significantly more cost effective, due to the Jacobi relaxation and sparsity of the coarse-grid correction calculation. Indeed, compatible relaxation yields diagonally dominant A_{ff} blocks for many operators, allowing an efficient choice of a diagonal matrix, D , that achieves a tight bound, ϵ .

Because REL is self-adjoint in the A -inner product, and CLC is a projection in the A -inner product, the multigrid convergence factors, measured in the A -norm, for symmetric multigrid cycles, $MG_2 = REL^\nu \cdot CLC \cdot REL^\nu$, may be easily related to the non-symmetric cycles analyzed in Theorem 1 and its corollary. Relative to the A -inner product, the symmetric cycle is the composition of the adjoint of the non-symmetric cycle analyzed above and itself. Thus, the multigrid convergence factor for symmetric cycles is simply the square of the cycles analyzed in Theorem 1 and its corollary.

3. Adaptive AMGr (α AMGr)

Just as classical AMG relies on a good representation of the errors that are slow to be reduced by relaxation (the algebraically smooth errors), the theoretical bounds on AMGr depend on the knowledge of a good splitting of A_{ff} into $D + \mathcal{E}$ with the assumed properties. In classical AMG, the choice of relaxation is taken to be fixed and, so, the goal is to define interpolation that effectively complements this choice. When a good representation of the algebraically smooth errors of relaxation is not available a priori, the adaptive multigrid methodology [14, 15, 16] may be used to probe the performance of relaxation and expose prototypes of such errors. Here, we describe a similar, adaptive version of AMGr; however, because both relaxation and interpolation in AMGr depend on the splitting of A_{ff} , our goal is not to expose the slow-to-converge errors of a fixed relaxation scheme, but rather to expose a good splitting of A_{ff} in the case that such a splitting is not known beforehand and, thus, improve relaxation and interpolation simultaneously.

Despite the difference in goals, our approach is similar to that of the adaptive multigrid methodology. The relaxation scheme chosen in (2) will be slow to resolve errors, e , that yield small residuals regardless of the chosen splitting, because Ae will always be small relative to e . Thus, fine-grid relaxation quickly exposes errors that are both algebraically smooth and in the near-null space of A . As discussed in [16], we use a Rayleigh quotient minimization on the coarse grid to better expose the near-null space of A in this two-level algorithm (for a multilevel algorithm, relaxation would be used on every level but the coarsest). An improved representative of the near-null space can then be used to determine an improved splitting. To understand this process more clearly, we now assume (without loss of generality) that A is singular, but nonnegative definite. It is easy to see that Theorem 1 still holds with the understanding that convergence is in $\mathcal{N}(A)^\perp$ (the orthogonal complement of the null space of A).

Although we no longer assume that D and \mathcal{E} are known, we continue to suppose that the partition $\mathbb{R}^n = F \cup C$ is given and that D and \mathcal{E} exist under the assumptions of the previous section. Thus, the matrix, A , can be written as

$$A = \begin{bmatrix} A_{ff} & -A_{fc} \\ -A_{fc}^T & A_{cc} \end{bmatrix} \geq 0.$$

Below, n_f and n_c denote the dimensions of the square matrices A_{ff} and A_{cc} , respectively, and subscripts f and c refer to the corresponding components of $v \in \mathbb{R}^n$. Thus, $v = (v_f, v_c)^T$ and $n = n_f + n_c$. Now, assume also that A is singular with a one-dimensional null space spanned by a unit vector, $\mathbf{r} \in \mathbb{R}^n: A\mathbf{r} = 0$, $\|\mathbf{r}\| = 1$. Continuing to assume that $A_D \geq 0$ and $\mathcal{E} \geq 0$, we must conclude that $A_D\mathbf{r} = 0$ and $\mathcal{E}\mathbf{r}_f = 0$. To see this, note that the splitting $A_{ff} = D + \mathcal{E}$ yields

$$0 = \langle A\mathbf{r}, \mathbf{r} \rangle = \langle A_D\mathbf{r}, \mathbf{r} \rangle + \langle \mathcal{E}\mathbf{r}_f, \mathbf{r}_f \rangle.$$

By assumption, the two terms on the right are nonnegative, so they must both be zero. Since A_D and \mathcal{E} are both nonnegative, it must be that $A_D\mathbf{r} = 0$ and $\mathcal{E}\mathbf{r}_f = 0$.

Finding a general matrix, D , with which to split A_{ff} allows many possible algorithmic choices, but here we consider the case that D is diagonal. This is motivated by thinking of the splitting of \mathbb{R}^n into F and C coming from an algorithm such as compatible relaxation and, so, the A_{ff} block exhibits some form of diagonal dominance. Now, without loss of generality

and for convenience, we further assume that $D = I$, implicitly rescaling A if necessary. Note that such a rescaling changes neither our assumptions on A nor the iteration in the following theorems (as viewed on the coarse scale), although the constants in the statement of Theorem 2 do change. Our aim is to construct a two-level method that yields grid-independent convergence factors for solving $Ax = b$, in general, and $Ax = 0$, in particular, without assuming knowledge of \mathbf{r} nor of the splitting of $A_{ff} = I + \mathcal{E}$.

Suppose now that the vector, u , is given as an approximation to the true null-space component, \mathbf{r} . We write $u = \mathbf{r} + e$, where e is taken to be somehow orthogonal to \mathbf{r} (in a manner to be defined). Note that our goal is to expose the eigenvector, \mathbf{r} , up to any scale factor. The implicit rescaling of u such that $\langle u, \mathbf{r} \rangle = \langle \mathbf{r}, \mathbf{r} \rangle$ is acceptable, because u is not used in what follows without being normalized against its scaling. Defining $RQ(v) = \frac{\langle Av, v \rangle}{\langle v, v \rangle}$, the Rayleigh quotient, our two-level adaptive AMGr algorithm then takes the following abstract form:

1. Relax on $Au = 0$.
2. Define P such that $Pu_c = u$, for some u_c .
3. Set $u^{\text{new}} = P \left(\underset{w \in \mathbb{R}^{n_c}}{\text{argmin}} RQ(Pw) \right)$.

In this form, the relationship between the adaptive AMGr setup algorithm and two-level eigensolvers is apparent. Indeed, the connection between good algebraic multigrid interpolation and eigenvectors has long been known (cf. [2, 20]). For stochastic matrices, aggregation-disaggregation methods take a form similar to the algorithm above. Local and global convergence theory for these methods, however, relies on the positivity of the stochastic operators [21], which does not directly hold for discretized PDEs. The RQMG methodology [22] could be applied directly to finding the eigenvector belonging to the smallest eigenvalue of A , but would, in general, require updates to the coarsening process used based on the evolving prototype. The eigensolver proposed here is much simpler and designed for exactly the task at hand.

To simplify the theoretical development, we replace the iterative block F -point solver of Equation (2) by exact F -point relaxation, $u_f = A_{ff}^{-1} A_{fc} u_c$. While impractical for general systems, the assumption on A_{ff} is that it is well approximated by some diagonal matrix, D , in the sense that $D \leq A_{ff} \leq (1 + \epsilon)D$, so the A_{ff} block of A is easy to approximately invert. In this light, it is easy to see that the following theory readily extends to the case that just a few relaxation steps are applied to $A_{ff} u_f = A_{fc} u_c$.

Interpolation, P , is then defined analogously to AMGr, by computing an approximate D based on the decomposition of $A_{ff} = D + \mathcal{E}$. Under the assumptions that $A_D \geq 0$ and $\mathcal{E} \geq 0$, we have $A_{ff} \mathbf{r}_f = D \mathbf{r}_f$ and, hence, $\mathbf{r}_f = A_{ff}^{-1} A_{fc} \mathbf{r}_c = D^{-1} A_{fc} \mathbf{r}_c$, so we define interpolation by approximating D based on our approximation, u , to \mathbf{r} . This is done by choosing the form of interpolation to be $\Lambda^{-1} A_{fc}$ for some diagonal matrix $\Lambda > 0$, and then determining Λ to match the fine-level values: $\Lambda^{-1} A_{fc} u_c = u_f = A_{ff}^{-1} A_{fc} u_c$. Note that the choice of diagonal Λ means that this single condition is enough to uniquely specify all of Λ . The interpolation operator, P , is then given by

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

Using u to update Λ is analogous to the updating of interpolation in adaptive AMG [16] based on the prototype of algebraically smooth error. An improved representation of the true null space of A yields a better splitting of A_{ff} , improving the convergence of the resulting AMGr cycle. For later use, it is convenient to define an “ideal” interpolation operator,

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

that results from taking $D = A_{ff}$. As remarked in Section 2, this yields $\epsilon = 0$, giving an exact two-level scheme.

A global convergence result can be obtained for the special case of a two-dimensional coarse grid. We include this simplified result here because it motivates the general case. Writing $u_c = \mathbf{r}_c + e_c$, with $\mathbf{r}_c \perp e_c$, gives an orthogonal basis for the coarse-grid space as the span of $\{\mathbf{r}_c, e_c\}$. This reduction in complexity allows for a complete analysis of the α AMGr algorithm, with the coarse-grid eigenproblem solved by a single step of inverse iteration.

Theorem 2. *Let $n_c = 2$, and let approximation u , with $u_c = \mathbf{r}_c + e_c$, $\|\mathbf{r}_c\| = 1$, and $\mathbf{r}_c \perp e_c$, be given. Define the usual α AMGr interpolation operator, $P = \begin{bmatrix} \Lambda^{-1}A_{fc} \\ I \end{bmatrix}$, with $\Lambda^{-1}A_{fc}u_c = A_{ff}^{-1}A_{fc}u_c$, and coarse-grid matrix $A_P = P^TAP$. Assume also that $A\mathbf{r} = 0$, $A_{ff} = I + \mathcal{E}$, for $0 \leq \mathcal{E} \leq \epsilon I$, and $\begin{bmatrix} I & -A_{fc} \\ -A_{fc}^T & A_{cc} \end{bmatrix} \geq 0$. Finally, let $\bar{\eta} = \min_i(\mathbf{r}_f)_i$ and $\eta = \min(\min_i(\mathbf{r}_f + A_{fc}e_c)_i, \bar{\eta}) > 0$. Then the coarse-grid solution, v_c , to the inverse iteration, $v_c = A_P^{-1}u_c$, satisfies*

$$v_c = k(\zeta\mathbf{r}_c + e_c),$$

for constant k and $\zeta \geq 1 + \frac{1}{1+\|e_c\|^2} \frac{\eta^2}{\epsilon(1+\epsilon)\|\mathbf{r}_f\|^2}$.

Proof: Let $W = [\mathbf{r}_c, \frac{e_c}{\|e_c\|}]$ be the orthogonal transformation from Cartesian coordinates to the (\mathbf{r}_c, e_c) coordinate system. The inverse iteration, $v_c = A_P^{-1}u_c$, may then be rewritten as

$$v_c = W(W^T A_P W)^{-1} W^T (\mathbf{r}_c + e_c).$$

Note that $Pu_c = Qu_c$ and $Q\mathbf{r}_c = \mathbf{r}$, so that

$$\begin{aligned} P\mathbf{r}_c &= \mathbf{r} + \begin{pmatrix} (\Lambda^{-1}-I)\mathbf{r}_f \\ 0 \end{pmatrix}, \\ Pe_c &= Qe_c - \begin{pmatrix} (\Lambda^{-1}-I)\mathbf{r}_f \\ 0 \end{pmatrix}. \end{aligned}$$

Compute $\alpha \equiv \mathbf{r}_c^T A_P \mathbf{r}_c = \mathbf{r}_f^T (\Lambda^{-1} - I) A_{ff} (\Lambda^{-1} - I) \mathbf{r}_f$. Notice then that $e_c^T A_P \mathbf{r}_c = -\alpha$, because $Pe_c = Pu_c - P\mathbf{r}_c = Qu_c - P\mathbf{r}_c$ and Qu_c is A -orthogonal to any vector that is zero on all coarse points.

Thus, taking $\gamma = \|e_c\|$ and $\beta = e_c^T \hat{A}_{cc} e_c$, the inverse iteration matrix may be rewritten as

$$W^T P^T A_P W = \begin{bmatrix} \alpha & -\frac{\alpha}{\gamma} \\ -\frac{\alpha}{\gamma} & \frac{\alpha+\beta}{\gamma^2} \end{bmatrix}.$$

Inverse iteration then takes the form

$$\begin{aligned} v_c &= W \begin{bmatrix} \alpha & -\frac{\alpha}{\gamma} \\ -\frac{\alpha}{\gamma} & \frac{\alpha+\beta}{\gamma^2} \end{bmatrix}^{-1} W^T (\mathbf{r}_c + e_c) \\ &= W \frac{\gamma^2}{\alpha\beta} \begin{bmatrix} \frac{\alpha+\beta}{\gamma^2} & \frac{\alpha}{\gamma} \\ \frac{\alpha}{\gamma} & \alpha \end{bmatrix} \begin{pmatrix} 1 \\ \gamma \end{pmatrix} \\ &= \left(\frac{\alpha + \beta}{\alpha\beta} + \frac{\gamma^2}{\beta} \right) \mathbf{r}_c + \left(\frac{\gamma + \gamma^3}{\beta} \right) \frac{e_c}{\gamma} \\ &= \left(\frac{1 + \gamma^2}{\beta} \right) \left(\left(\frac{\alpha + \beta}{(1 + \gamma^2)\alpha} + \frac{\gamma^2}{1 + \gamma^2} \right) \mathbf{r}_c + e_c \right). \end{aligned}$$

Now consider $\beta = e_c^T \hat{A}_{cc} e_c$, and note that

$$\hat{A}_{cc} = A_{cc} - A_{fc}^T A_{fc} + A_{fc}^T (I - A_{ff}^{-1}) A_{fc} \geq A_{fc}^T (I - A_{ff}^{-1}) A_{fc},$$

because $A_{cc} - A_{fc}^T A_{fc} \geq 0$, by our assumptions as in Equation (1). Then, since

$$(I - A_{ff}^{-1}) = (I - (I + \mathcal{E})^{-1}) = \mathcal{E}(I + \mathcal{E})^{-1},$$

we see that

$$\hat{A}_{cc} \geq \frac{1}{\epsilon} A_{fc}^T \mathcal{E}^2 (I + \mathcal{E})^{-1} A_{fc},$$

because

$$\mathcal{E}^2 (I + \mathcal{E})^{-1} \leq \epsilon \mathcal{E} (I + \mathcal{E})^{-1}.$$

In particular, this implies that $\beta \geq \frac{1}{\epsilon} e_c^T A_{fc}^T \mathcal{E}^2 (I + \mathcal{E})^{-1} A_{fc} e_c$. We now turn our attention to α , which can be bounded above as follows:

$$\begin{aligned} \alpha &= \|(\Lambda^{-1} - I)\mathbf{r}_f\|_{A_{ff}}^2 \leq (1 + \epsilon) \|(\Lambda^{-1} - I)\mathbf{r}_f\|^2 \\ &= (1 + \epsilon) \sum_i \left(\frac{((A_{ff}^{-1} - I)A_{fc}e_c)_i}{(\mathbf{r}_f + A_{fc}e_c)_i} \right)^2 (\mathbf{r}_f)_i^2 \\ &\leq \frac{1 + \epsilon}{\eta^2} \sum_i ((A_{ff}^{-1} - I)A_{fc}e_c)_i^2 \sum_i (\mathbf{r}_f)_i^2 \\ &\leq \frac{(1 + \epsilon)\|\mathbf{r}_f\|^2}{\eta^2} \sum_i ((A_{ff}^{-1} - I)A_{fc}e_c)_i^2. \end{aligned}$$

Hence,

$$\begin{aligned} \alpha &= \frac{(1 + \epsilon)\|\mathbf{r}_f\|^2}{\eta^2} \|(A_{ff}^{-1} - I)A_{fc}e_c\|^2 \\ &\leq \frac{(1 + \epsilon)\|\mathbf{r}_f\|^2}{\eta^2} \|(A_{ff}^{-1} - I)A_{fc}e_c\|_{A_{ff}}^2 \\ &= \frac{(1 + \epsilon)\|\mathbf{r}_f\|^2}{\eta^2} e_c^T A_{fc}^T A_{ff}^{-1} (A_{ff} - I) A_{ff} (A_{ff} - I) A_{ff}^{-1} A_{fc} e_c \\ &= \frac{(1 + \epsilon)\|\mathbf{r}_f\|^2}{\eta^2} e_c^T A_{fc}^T \mathcal{E}^2 (I + \mathcal{E})^{-1} A_{fc} e_c \leq \frac{\epsilon(1 + \epsilon)\|\mathbf{r}_f\|^2}{\eta^2} \beta. \end{aligned}$$

Thus, $v_c = \left(\frac{1+\gamma^2}{\beta}\right) \left(\left(\frac{\alpha+\beta}{(1+\gamma^2)\alpha} + \frac{\gamma^2}{1+\gamma^2}\right) \mathbf{r}_c + e_c\right)$ can be seen to have the form, $v_c = k(\zeta \mathbf{r}_c + e_c)$, for $k = \left(\frac{1+\gamma^2}{\beta}\right)$, and

$$\begin{aligned} \zeta &= \left(\frac{\alpha+\beta}{(1+\gamma^2)\alpha} + \frac{\gamma^2}{1+\gamma^2}\right) \geq 1 + \frac{\beta}{(1+\gamma^2)\alpha} \\ &\geq 1 + \frac{1}{1+\gamma^2} \frac{\eta^2}{\epsilon(1+\epsilon)\|\mathbf{r}_f\|^2}. \end{aligned}$$

□

Corollary 2. *Under the hypotheses of Theorem 2, the inverse iteration process, $u_c \leftarrow A_P^{-1}u_c$, converges uniformly to $u_c = \mathbf{r}_c$.*

Proof: Let $\bar{\eta} = \min_i (\mathbf{r}_f)_i$. Renormalizing $v_c = \mathbf{r}_c + \frac{1}{\zeta}e_c$, we see that $\frac{1}{\zeta} < 1$, so

$$\eta^{\text{new}} = \min \left(\min_i \left(\mathbf{r}_f + \frac{1}{\zeta} A_{fc} e_c \right)_i, \bar{\eta} \right) \geq \eta^{\text{old}},$$

and that

$$\gamma^{\text{new}} = \left\| \frac{1}{\zeta} e_c \right\| < \gamma^{\text{old}}.$$

Thus, for any initial $\eta^{(0)}$ and $\gamma^{(0)}$, we have ζ from Theorem 2 satisfying

$$\zeta \geq 1 + \frac{1}{1+(\gamma^{(0)})^2} \frac{(\eta^{(0)})^2}{\epsilon(1+\epsilon)}.$$

□

This simplified setting allows a complete analysis of the α AMGr algorithm, under very natural assumptions. The splitting of $A_{ff} = I + \mathcal{E}$ does not reduce the generality, as discussed above. The requirement that $\eta > 0$ arises naturally from requiring well-posedness of the AMGr process. For any iterate, $u = \mathbf{r} + e$, Λ^{-1} is defined pointwise as $(\Lambda^{-1})_i = \frac{(A_{ff}^{-1} A_{fc} u_c)_i}{(A_{fc} u_c)_i}$. The requirement that $\eta > 0$ then ensures that, for any iterate, u , $(A_{fc} u_c)_i = (\mathbf{r}_f + A_{fc} e_c)_i \neq 0$ for all i . At the solution, $A_{fc} u_c = \mathbf{r}_f$ and, so, the additional condition on $\bar{\eta}$ simply enforces that α AMGr is well-defined at the solution as well.

Because we have been unable to extend this line of proof to the general case (see, however, [23, §5.3.3]), we turn instead to a more general approach that yields a local convergence result. In place of inverse iteration, we consider a two-grid cycle based on an exact solution to the coarse-grid eigenproblem. Thus, given an approximation, u , to \mathbf{r} , consider the following adaptive AMGr setup cycle:

1. Solve $A_{ff} u_f = A_{fc} u_c$.
2. Define $P = \begin{bmatrix} \Lambda^{-1} A_{fc} \\ I \end{bmatrix}$, for diagonal Λ such that $\Lambda^{-1} A_{fc} u_c = A_{ff}^{-1} A_{fc} u_c$.
3. Set $u^{\text{new}} = P \left(\underset{w \in \mathbb{R}^{n_c}}{\text{argmin}} RQ(Pw) \right)$.

We show that this defines a contraction mapping (measured in a particular norm) and, thus, converges to the solution, $u = \mathbf{r}$ and $\Lambda = I$, in some neighborhood of it.

In what follows, we use the notation

$$\mathcal{D}(w) = \text{Diag}(\dots, w_i, \dots),$$

for a given vector, w , to be the diagonal matrix with w_i as the diagonal terms. We also use the Jacobian notation,

$$J(G, x)|_{x=y},$$

to denote the Jacobian of map G with respect to x , evaluated at $x = y$, or $J(z, x)|_{x=y}$ to denote the Jacobian of the map from x to z , evaluated at $x = y$.

Theorem 3. *Suppose we are given an $A = \begin{bmatrix} A_{ff} & -A_{fc} \\ -A_{fc}^T & A_{cc} \end{bmatrix} \geq 0$ such that $A_{ff} = I + \mathcal{E}$, $0 \leq \mathcal{E} \leq \epsilon I$, and $\begin{bmatrix} I & -A_{fc} \\ -A_{fc}^T & A_{cc} \end{bmatrix} \geq 0$. Further, suppose that $A_{fc} \neq 0$ (convergence in the case that $A_{fc} = 0$ is trivial). Assume that there exists a vector, \mathbf{r} , $\mathbf{r}_i \neq 0$ for all i , such that $A\mathbf{r} = 0$, but that $v^T Av > 0$ for any other vector, $v \perp \mathbf{r}$.*

Given u_c , write $u_c = \mathbf{r}_c + e_c$, for $e_c \perp \mathbf{r}_c$, rescaling u_c if necessary. Define the diagonal matrix, Λ , such that $\Lambda^{-1}A_{fc}u_c = A_{fc}^{-1}A_{fc}u_c$, and let $P = \begin{bmatrix} \Lambda^{-1}A_{fc} \\ I \end{bmatrix}$ and $A_P = P^T A P$. Denote the new error by $e_c^{\text{new}} = v_c - \mathbf{r}_c$, where v_c results from solving the coarse-grid eigenvalue problem, $A_P v_c = \lambda_{\min} v_c$, $\langle v_c, \mathbf{r}_c \rangle = \|\mathbf{r}_c\|^2$. Let this iteration process be denoted by a map, $G : e_c \rightarrow e_c^{\text{new}}$.

Then G is a contraction map at the solution, $e_c = 0$, in the $\bar{A}_{cc} = (A_{cc} - 2A_{fc}^T A_{fc} + A_{fc}^T A_{ff} A_{fc})$ -norm. Further, the norm of the Jacobian of G with respect to e_c at $e_c = 0$ satisfies

$$\|J(G, e_c)|_{e_c=0}\|_{\bar{A}_{cc}} \leq \frac{\epsilon}{1 + \epsilon},$$

and so G remains a contraction in some neighborhood of $e_c = 0$.

Proof: The coarse-grid step is to find

$$e_c^{\text{new}} = \underset{\tilde{e}_c \perp \mathbf{r}_c}{\text{argmin}} RQ(A_P, \mathbf{r}_c + \tilde{e}_c),$$

where $RQ(B, v) = \frac{\langle Bv, v \rangle}{\langle v, v \rangle}$. This is always possible, under the assumption that e_c is close enough to 0 that the eigenvector associated with the minimum eigenvalue is guaranteed to have a non-zero projection onto \mathbf{r}_c , which is in the null space of the ideal coarse-grid operator, $A_{cc} - 2A_{fc}^T A_{fc} + A_{fc}^T A_{ff} A_{fc}$. Since we are concerned with properties of the mapping at the solution, $e_c = 0$, this assumption is easily satisfied.

The fine-grid error, e_f^{new} , is then defined by an exact fine-grid relaxation,

$$A_{ff}(\mathbf{r}_f + e_f^{\text{new}}) - A_{fc}(\mathbf{r}_c + e_c^{\text{new}}) = 0.$$

The iteration may be viewed as the mapping, $G : e_c \rightarrow e_c^{\text{new}}$. Defining $Z = \Lambda^{-1} - I$, and letting z be the vector with components $z_i = Z_{ii}$ (i.e., $Z = \mathcal{D}(z)$), the iteration proceeds as

$$e_c \rightarrow z \rightarrow e_c^{\text{new}}.$$

First, rewrite the α AMG interpolation equation,

$$(I + Z)A_{fc}(\mathbf{r}_c + e_c) = A_{ff}^{-1}A_{fc}(\mathbf{r}_c + e_c),$$

as

$$ZA_{fc}(\mathbf{r}_c + e_c) = (A_{ff}^{-1} - I)A_{fc}(\mathbf{r}_c + e_c).$$

Then, defining $U = \mathcal{D}(A_{fc}\mathbf{r}_c + A_{fc}e_c)$ and making use of the fact that $ZA_{fc}u_c = Uz$, we have

$$z = U^{-1}(A_{ff}^{-1} - I)A_{fc}(\mathbf{r}_c + e_c),$$

if U^{-1} exists (i.e., if $(A_{fc}u_c)_i \neq 0$ for any i). Note that A is the sum of the two positive semi-definite parts, $\begin{bmatrix} \mathcal{E} & 0 \\ 0 & 0 \end{bmatrix}$ and $\begin{bmatrix} I & -A_{fc} \\ -A_{fc}^T & A_{cc} \end{bmatrix}$, and that $A\mathbf{r} = 0$, so $\mathcal{E}\mathbf{r}_f = 0$, $A_{ff}\mathbf{r}_f = \mathbf{r}_f$, and $A_{fc}\mathbf{r}_c = \mathbf{r}_f$. Thus, U^{-1} exists at the solution, $e_c = 0$ (since $U = \mathcal{D}(\mathbf{r}_f)$, which, by assumption, is nonzero), and, by continuity, it must exist in some neighborhood of the solution, which is sufficient for our analysis.

Now, noting that $z = z(e_c)$, consider the Jacobian, $J(z, e_c) = \frac{\partial z}{\partial e_c}$, evaluated at e_c in direction v_c :

$$J(z, e_c)|_{e_c}[v_c] = U^{-2} \left(U(A_{ff}^{-1} - I)A_{fc}v_c - \mathcal{D}(A_{fc}v_c)(A_{ff}^{-1} - I)A_{fc}(\mathbf{r}_c + e_c) \right).$$

At the solution, $e_c = 0$, the second term in $J(z, e_c)$ vanishes because $A_{ff}^{-1}A_{fc}\mathbf{r}_c = A_{fc}\mathbf{r}_c$, and we have

$$J(z, e_c)|_{e_c=0} = \mathcal{D}(A_{fc}\mathbf{r}_c)^{-1}(A_{ff}^{-1} - I)A_{fc}. \tag{6}$$

For the second part of the map, we first find

$$u_c = \operatorname{argmin}_{v_c: \langle v_c, \mathbf{r}_c \rangle = \langle \mathbf{r}_c, \mathbf{r}_c \rangle} RQ(A_P, v_c), \tag{7}$$

and set $e_c^{\text{new}} = u_c - \mathbf{r}_c$. This yields an implicit relationship between z and e_c^{new} . Taking the gradient of (7), we have the set of equations

$$A_P u_c - \frac{\langle A_P u_c, u_c \rangle}{\langle u_c, u_c \rangle} u_c = 0$$

plus the constraint equation

$$\langle u_c, \mathbf{r}_c \rangle - \langle \mathbf{r}_c, \mathbf{r}_c \rangle = 0.$$

Note that $A_P = A_P(z)$ and, to belabor the point, this has the form of a set of equations

$$F(u_c, z) = A_P(z)u_c - \frac{\langle A_P(z)u_c, u_c \rangle}{\langle u_c, u_c \rangle} u_c = 0. \tag{8}$$

With the constraint, there are $n_c + 1$ equations, but they are consistent, as the first n_c do not determine the constant multiplier and are therefore singular. The constraint fixes the scaling.

Implicit differentiation of F yields

$$J(F, u_c)J(u_c, e_c)J(e_c, z) + J(F, z) = 0, \tag{9}$$

where, of course, $J(u_c, e_c) = I$. We deal with the constraint equation later.

We now need to construct $J(F, u_c)$ and $J(F, z)$. Notice that the coarse-grid operator satisfies

$$A_P = P^T A P = A_{cc} - 2A_{fc}^T \Lambda^{-1} A_{fc} + A_{fc}^T \Lambda^{-1} A_{ff} \Lambda^{-1} A_{fc}.$$

Referring to the definition of A_P and recalling that $A_{ff} = I + \mathcal{E}$ and $\Lambda^{-1} = I + Z$, then

$$A_P = (A_{cc} - 2A_{fc}^T A_{fc} + A_{fc}^T A_{ff} A_{fc}) + A_{fc}^T (Z^2 + Z\mathcal{E} + \mathcal{E}Z + Z\mathcal{E}Z) A_{fc}.$$

This yields

$$\begin{aligned} J(F, z)|_z = & A_{fc}^T (2Z\mathcal{D}(A_{fc}u_c) + \mathcal{D}(\mathcal{E}A_{fc}u_c) + \mathcal{E}\mathcal{D}(A_{fc}u_c) + \mathcal{D}(\mathcal{E}ZA_{fc}u_c) + Z\mathcal{E}\mathcal{D}(A_{fc}u_c)) \\ & - \frac{u_c u_c^T}{u_c^T u_c} A_{fc}^T (2Z\mathcal{D}(A_{fc}u_c) + \mathcal{D}(\mathcal{E}A_{fc}u_c) + \mathcal{E}\mathcal{D}(A_{fc}u_c) \\ & + \mathcal{D}(\mathcal{E}ZA_{fc}u_c) + Z\mathcal{E}\mathcal{D}(A_{fc}u_c)). \end{aligned}$$

The important thing to notice is that, at the solution, $e_c = 0$, $u_c = \mathbf{r}_c$, $Z = 0$, and $z = 0$, yielding

$$J(F, z)|_{z=0} = A_{fc}^T \mathcal{E}\mathcal{D}(A_{fc}\mathbf{r}_c) - \frac{\mathbf{r}_c \mathbf{r}_c^T}{\mathbf{r}_c^T \mathbf{r}_c} A_{fc}^T \mathcal{E}\mathcal{D}(A_{fc}\mathbf{r}_c).$$

However, $\mathbf{r}_c^T A_{fc}^T \mathcal{E} = (\mathcal{E}A_{fc}\mathbf{r}_c)^T = (\mathcal{E}\mathbf{r}_f)^T = 0^T$, so

$$J(F, z)|_{z=0} = A_{fc}^T \mathcal{E}\mathcal{D}(A_{fc}\mathbf{r}_c).$$

Now, to compute $J(F, u_c)$, we have

$$J(F, u_c) = A_P - RQ(A_P, u_c)I - u_c(\nabla_{u_c} RQ(A_P, u_c))^T,$$

where, as stated before,

$$\nabla_{u_c} RQ(A_P, u_c) = \frac{2}{\langle u_c, u_c \rangle} (A_P u_c - \frac{\langle A_P u_c, u_c \rangle}{\langle u_c, u_c \rangle} u_c).$$

At the solution, this yields

$$J(F, u_c)|_{u_c=\mathbf{r}_c} = A_{cc} - 2A_{fc}^T A_{fc} + A_{fc}^T A_{ff} A_{fc} = A_{cc} - A_{fc}^T A_{fc} + A_{fc}^T \mathcal{E}A_{fc}.$$

Now consider the constraint. If we think of it as the last in Equation (8), this adds one extra equation to the set in (9). The last row of $J(F, z)$ is 0^T because the constraint does not depend on z . The constraint equation also adds one row to $J(F, u_c)$, and that last row is \mathbf{r}_c^T . That is, the columns of $J(e_c, z)$ are all orthogonal to \mathbf{r}_c . We must check that this augmented system has a solution, because matrix $J(F, u_c)|_{u_c=\mathbf{r}_c} = A_{cc} - 2A_{fc}^T A_{fc} + A_{fc}^T A_{ff} A_{fc}$ has \mathbf{r}_c as its null space. By the assumption that \mathbf{r} is the unique null vector (up to scale) of A , we conclude that \mathbf{r}_c is the only null vector (up to scale) of $A_{cc} - A_{fc}^T A_{fc} + A_{fc}^T \mathcal{E}A_{fc}$, and, thus, that $J(F, u_c)|_{u_c=\mathbf{r}_c}$ has rank $n_c - 1$. Note that each column of the right side of $-J(F, z)|_{z=0} = -A_{fc}^T \mathcal{E}\mathcal{D}(A_{fc}\mathbf{r}_c)$ is orthogonal to \mathbf{r}_c , because $\mathbf{r}_c^T A_{fc}^T \mathcal{E} = 0^T$. Then, since $J(F, u_c)|_{u_c=\mathbf{r}_c}$ is symmetric, each column of $-J(F, z)|_{z=0}$

must be in its range. (This is Fredholm's theorem: $(\text{Range}(B))^\perp = \text{Null}(B^T)$.) So, we know that a solution, $J(e_c, z)$, exists. If we additionally ask for the solution to satisfy the constraint, $J(e_c, z)\mathbf{r}_c = 0$, we then get a unique value for $J(e_c, z)|_{z=0}$, because this constraint is orthogonal to $J(F, u_c)|_{u_c=\mathbf{r}_c}$, giving an $n_c + 1 \times n_c$ system with rank n_c and a consistent right side.

Using the results in Equations (6) and (9), we now have

$$(A_{cc} - A_{fc}^T A_{fc} + A_{fc}^T \mathcal{E} A_{fc}) J(e_c, z) J(z, e_c) = A_{fc}^T \mathcal{E} (I - A_{ff}^{-1}) A_{fc} = A_{fc}^T \mathcal{E}^2 (I + \mathcal{E})^{-1} A_{fc},$$

where we are interested in $J(e_c, z) J(z, e_c)$ over the space $(\mathbf{r}_c)^\perp$. We choose to measure the size of $J(e_c, z) J(z, e_c)$ in the A_P -norm that, at the solution, is the $\bar{A}_{cc} = (A_{cc} - A_{fc}^T A_{fc} + A_{fc}^T \mathcal{E} A_{fc})$ -norm. On this space, \bar{A}_{cc}^{-1} exists and is well-posed by assumption that the unique (up to scale) null-space component of A is \mathbf{r} , and so the unique (up to scale) null-space component of \bar{A}_{cc} is \mathbf{r}_c . Noting that, for any y_f , $\mathbf{r}_c^T A_{fc}^T \mathcal{E}^{\frac{1}{2}} y_f = 0$, as $A_{fc} \mathbf{r}_c = \mathbf{r}_f$, we can write

$$\|J(e_c, z) J(z, e_c)\|_{\bar{A}_{cc}}^2 = \max_{w_c: w_c \neq 0, w_c \perp \mathbf{r}_c} \left(\frac{w_c^T A_{fc}^T (I + \mathcal{E})^{-1} \mathcal{E}^2 A_{fc} \bar{A}_{cc}^{-1} A_{fc}^T \mathcal{E}^2 (I + \mathcal{E})^{-1} A_{fc} w_c}{w_c^T \bar{A}_{cc} w_c} \right),$$

because \bar{A}_{cc}^{-1} only applies to vectors in $(\mathbf{r}_c)^\perp$, where it is well-defined. By assumption, $A_{cc} \geq A_{fc}^T A_{fc}$, so $\bar{A}_{cc} \geq A_{fc}^T \mathcal{E} A_{fc}$. Since \bar{A}_{cc} and $A_{fc}^T \mathcal{E} A_{fc}$ have the same null space,

$$\|J(e_c, z) J(z, e_c)\|_{\bar{A}_{cc}}^2 \leq \max_{w_c: w_c \neq 0, w_c \perp \mathbf{r}_c} \left(\frac{w_c^T A_{fc}^T (I + \mathcal{E})^{-1} \mathcal{E}^2 A_{fc} (A_{fc}^T \mathcal{E} A_{fc})^{-1} A_{fc}^T \mathcal{E}^2 (I + \mathcal{E})^{-1} A_{fc} w_c}{w_c^T \bar{A}_{cc} w_c} \right),$$

but

$$\rho(\mathcal{E}^{\frac{1}{2}} A_{fc} (A_{fc}^T \mathcal{E} A_{fc})^{-1} A_{fc}^T \mathcal{E}^{\frac{1}{2}}) = \rho((A_{fc}^T \mathcal{E} A_{fc})^{-\frac{1}{2}} A_{fc}^T \mathcal{E}^{\frac{1}{2}} \mathcal{E}^{\frac{1}{2}} A_{fc} (A_{fc}^T \mathcal{E} A_{fc})^{-\frac{1}{2}}) = 1,$$

so $\mathcal{E}^{\frac{1}{2}} A_{fc} (A_{fc}^T \mathcal{E} A_{fc})^{-1} A_{fc}^T \mathcal{E}^{\frac{1}{2}} \leq I$, and

$$\|J(e_c, z) J(z, e_c)\|_{\bar{A}_{cc}}^2 \leq \max_{w_c: w_c \neq 0, w_c \perp \mathbf{r}_c} \left(\frac{w_c^T A_{fc}^T (I + \mathcal{E})^{-1} \mathcal{E}^3 (I + \mathcal{E})^{-1} A_{fc} w_c}{w_c^T \bar{A}_{cc} w_c} \right).$$

We can again use the bound, $\bar{A}_{cc} \geq A_{fc}^T \mathcal{E} A_{fc}$, to simplify the denominator, but we must be careful not to allow the denominator to become zero. To do this, we take

$$\begin{aligned} \|J(e_c, z) J(z, e_c)\|_{\bar{A}_{cc}}^2 &\leq \max_{w_c: w_c \neq 0, w_c \perp \mathbf{r}_c} \left(\frac{w_c^T A_{fc}^T (I + \mathcal{E})^{-1} \mathcal{E}^3 (I + \mathcal{E})^{-1} A_{fc} w_c}{w_c^T (A_{cc} - A_{fc}^T A_{fc}) w_c + w_c^T A_{fc}^T \mathcal{E} A_{fc} w_c} \right), \\ &= \max_{w_c: \mathcal{E}^{\frac{1}{2}} A_{fc} w_c \neq 0, w_c \perp \mathbf{r}_c} \left(\frac{w_c^T A_{fc}^T (I + \mathcal{E})^{-1} \mathcal{E}^3 (I + \mathcal{E})^{-1} A_{fc} w_c}{w_c^T (A_{cc} - A_{fc}^T A_{fc}) w_c + w_c^T A_{fc}^T \mathcal{E} A_{fc} w_c} \right), \\ &\leq \max_{w_c: \mathcal{E}^{\frac{1}{2}} A_{fc} w_c \neq 0} \left(\frac{w_c^T A_{fc}^T (I + \mathcal{E})^{-1} \mathcal{E}^3 (I + \mathcal{E})^{-1} A_{fc} w_c}{w_c^T A_{fc}^T \mathcal{E} A_{fc} w_c} \right). \end{aligned}$$

Note that we first change the conditions on the set over which the maximum is taken. The generalization to $\mathcal{E}^{\frac{1}{2}}A_{fc}w_c \neq 0$ is possible as if $\mathcal{E}^{\frac{1}{2}}A_{fc}w_c = 0$ and $w_c \perp \mathbf{r}_c$, the expression to be maximized is zero (as the numerator is zero, but the denominator is not, since $(A_{cc} - A_{fc}^T A_{fc} + A_{fc}^T \mathcal{E} A_{fc})w_c$ is zero only when w_c is (up to scale) \mathbf{r}_c). Because the condition $\mathcal{E}^{\frac{1}{2}}A_{fc}w_c \neq 0$ already excludes the case that $w_c = 0$, the set considered in the maximum can be taken to be the set w_c such that $\mathcal{E}^{\frac{1}{2}}A_{fc}w_c \neq 0$ and $w_c \perp \mathbf{r}_c$. The term $\langle (A_{cc} - A_{fc}^T A_{fc})w_c, w_c \rangle$ may then be discarded from the denominator as it is positive. In doing so, we expand the set over which we maximize by dropping the requirement that $w_c \perp \mathbf{r}_c$, as the critical case when w_c is a scalar multiple of \mathbf{r}_c is already covered by the other condition.

Suppose now that V is the unitary matrix that diagonalizes \mathcal{E} , i.e.,

$$\Psi = \mathcal{D}(\psi) = V\mathcal{E}V^*.$$

Under this change of basis, we have

$$\|J(e_c, z)J(z, e_c)\|_{A_{cc}}^2 \leq \max_{x_f: \Psi^{\frac{1}{2}}x_f \neq 0} \frac{\langle \Psi^3(I + \Psi)^{-2}x_f, x_f \rangle}{\langle \Psi x_f, x_f \rangle} \leq \left(\frac{\epsilon}{1 + \epsilon} \right)^2,$$

by assumption on \mathcal{E} and, thus, Ψ .

Thus, the norm of the Jacobian of the mapping, $J(G, e_c)|_{e_c=0} = J(e_c, z)J(z, e_c)$, is bounded uniformly less than 1, yielding convergence in a neighborhood of the solution.

□

4. Numerical Results

We consider a set of test problems derived primarily from finite element discretizations of elliptic PDEs. Coarse grids are chosen both consistently with standard multigrid coarsening principles and to expose the role of the parameter, ϵ , in the convergence bound of Theorem 1. The adaptive AMGr setup procedure is also tested, both in the case of a singular matrix, A , as discussed in Section 3, and in the more general (and more important) case of a near-singular matrix.

Parameter ϵ , which is used in the relaxation step (see (2)), can be computed exactly from the decomposition $A_{ff} = D + \mathcal{E}$, $D \leq A_{ff}$, using the maximum eigenvalue of the generalized eigenvalue problem, $A_{ff}x = \lambda_{max}Dx$. A convenient upper bound for ϵ in the case of diagonally dominant A_{ff} comes from using Gerschgorin's theorem to bound the largest eigenvalue of A_{ff} , $\lambda_{max}(A_{ff})$, from above, and the smallest eigenvalue of D , $\lambda_{min}(D)$, from below. With these bounds, a convenient upper bound is $\epsilon \leq \frac{\lambda_{max}(A_{ff})}{\lambda_{min}(D)} - 1$. While this works well for homogeneous, isotropic operators when a good splitting is already known, it is less useful in the context of adaptive AMGr. Instead, an accurate estimate of ϵ may be obtained using a few steps of a Lanczos algorithm. In the tables below, we compute the exact ϵ for a given D , as well as that given by Gerschgorin's theorem, denoted ϵ_G , or the Lanczos algorithm, denoted ϵ_L , as appropriate.

We begin by considering A to be the matrix obtained by the discretization of Poisson's equation in two dimensions, $-\Delta u = f$, with homogeneous Dirichlet boundary conditions,

grid	D	ν	ϵ	$\rho(\epsilon)$	ϵ_G	$\rho(\epsilon_G)$
16×16	diagonal	1	4.90	0.71	6.00	0.75
16×16	diagonal	2	4.90	0.51	6.00	0.56
16×16	diagonal	3	4.90	0.36	6.00	0.42
16×16	diagonal	4	4.90	0.36	6.00	0.36
32×32	diagonal	1	4.98	0.71	6.00	0.75
32×32	diagonal	2	4.98	0.51	6.00	0.56
32×32	diagonal	3	4.98	0.37	6.00	0.42
32×32	diagonal	4	4.98	0.37	6.00	0.37
16×16	tridiagonal	1	4.13	0.67	6.00	0.75
16×16	tridiagonal	2	4.13	0.47	6.00	0.56
16×16	tridiagonal	3	4.13	0.34	6.00	0.42
16×16	tridiagonal	4	4.13	0.36	6.00	0.36

Table I. Asymptotic convergence factors, ρ , for two-level AMGr cycles with given choices of D and ν . Convergence factors are reported using both an exact ϵ calculated from the generalized eigenproblem, $\rho(\epsilon)$, and that which results from a Gerschgorin bound, $\rho(\epsilon_G)$.

using bilinear finite elements on a regular rectangular mesh. The partition of $\mathbb{R}^n = F \cup C$ is done in a multigrid full-coarsening fashion, removing every other row and column of degrees of freedom from the fine-scale mesh. Thus, the fine-scale block, A_{ff} , is diagonally dominant, because of the M-matrix structure of A and the fact that each row of A associated with F has diagonal $\frac{8}{3}$ and at least two entries in A_{fc} with size $-\frac{1}{3}$.

Table I shows the performance of AMGr for various choices of D and ν . One simple choice of D for this problem is the diagonal matrix that is the difference between A_{ff} and the matrix that has the same off-diagonal entries as A_{ff} , with the sum of these entries on the diagonal. This matrix can be easily computed as $\mathcal{D}(A_{ff}u_f)$, where $(u_f)_i = 1$ for all i . For this choice of D , and small values of ν , the relaxation step dominates the convergence factor, with $\rho \approx \left(\frac{\epsilon}{2+\epsilon}\right)^\nu$ for both the exact and approximate values of ϵ . As ν grows, the error reduction betters the theoretical bound. Thus, we see that the bound in Equation (4) is not a sharp bound on the overall AMGr process, although the bound on relaxation in Equation (5) is sharp and, thus, the performance of relaxation depends strongly on the chosen value of ϵ . Notice that the performance of AMGr is nearly indistinguishable between the two mesh sizes considered, indicating that ϵ , and not h , is an important factor in performance. To further test the influence of D and ϵ , we consider taking D to be not just the diagonally dominant part of A_{ff} , but to also include some off-diagonal structure. Knowing that choosing $D = A_{ff}$ leads to an exact solver, we choose D to be the tridiagonal matrix such that $\mathcal{E} = A_{ff} - D \geq 0$, but $\mathcal{E}u_f = 0$. The results presented in Table I show that this choice leads to a slightly improved ϵ when calculated exactly, but the same approximation as ϵ_G . The improved exact value is reflected in the slight improvement in performance when the exact value is used, although no real change is seen in the performance of the algorithm using the approximate ϵ_G . This again indicates the importance of ϵ in the performance of relaxation for this algorithm.

We first test the adaptive AMGr setup cycle in same theoretical setting as in Section 3, that of a singular matrix, A , whose A_{ff} block can be dominated by a diagonal matrix, D . As test

coefficient	μ	$RQ(u)$	ν	ϵ	$\rho(\epsilon)$	ϵ_L	$\rho(\epsilon_L)$
$K = 1$	1	3.35×10^{-2}	1	9.19	0.86	9.19	0.86
$K = 1$	2	1.31×10^{-2}	1	6.08	0.75	6.08	0.75
$K = 1$	3	3.72×10^{-3}	1	5.10	0.72	5.10	0.72
$K = 1$	3	3.72×10^{-3}	2	5.10	0.54	5.10	0.54
$K = 1$	3	3.72×10^{-3}	3	5.10	0.39	5.10	0.39
$K = 1$	3	3.72×10^{-3}	4	5.10	0.39	5.10	0.39
$K = 10^{-8} + 10x^2 + 10y^2$	1	1.29×10^{-2}	1	8.58	0.84	8.58	0.84
$K = 10^{-8} + 10x^2 + 10y^2$	2	1.21×10^{-2}	1	5.92	0.75	5.92	0.75
$K = 10^{-8} + 10x^2 + 10y^2$	3	1.20×10^{-2}	1	5.10	0.72	5.10	0.72
$K = 10^{-8} + 10x^2 + 10y^2$	3	1.20×10^{-2}	2	5.10	0.54	5.10	0.54
$K = 10^{-8} + 10x^2 + 10y^2$	3	1.20×10^{-2}	3	5.10	0.38	5.10	0.38
$K = 10^{-8} + 10x^2 + 10y^2$	3	1.20×10^{-2}	4	5.10	0.39	5.10	0.39

Table II. Performance of adaptive AMGr setup routine for various coefficients, K , and numbers of setup cycles, μ . The RQ of the resulting approximation, u , to \mathbf{r} is reported, as well as the performance of the resulting AMGr algorithms with both the exact ϵ and approximate ϵ_L , for given values of ν .

problems, we consider the discrete operators obtained by bilinear finite element discretization of second-order elliptic operators, $-\nabla \cdot K \nabla u = f$, subject to Neumann boundary conditions, for different choices of K . For these examples, we fix the grid to be a 16×16 element mesh, and choose Λ as in α AMGr. We then perform μ setup cycles with an initial guess of u such that $u_i \in (0, 1)$ is chosen from a uniform distribution, and measure the Rayleigh Quotient of the resulting approximation, u , to the true null-space component, \mathbf{r} . The exact ϵ and approximate ϵ_L are then computed and the resulting AMGr methods tested. Note that this approximate choice of Λ instead of the true D allows us to break the assumptions of AMGr, by trying a method such that $\Lambda \not\leq A_{ff}$.

Results for these tests are shown in Table II. For $K = 1$, we see a steady reduction in the Rayleigh Quotient of the approximate null-space component as we increase the number of setup stages. The optimal ϵ also decreases as a better u is exposed. For the variable K case, we see slower initial reduction in $RQ(u)$ (although this initial slowness does not persist - after 4 setup cycles, $RQ(u) = 3.59 \times 10^{-3}$, and after 5 cycles, $RQ(u) = 2.83 \times 10^{-4}$). As a better approximation of the null space is achieved, the parameter ϵ improves, and the resulting method performs better. In both of these cases, the importance of accurately estimating ϵ is easily demonstrated. Because of the significant variation in K in the second example, $\lambda_{min}(\Lambda)$ can be quite small, giving a very large value for ϵ_G . The poor performance of relaxation with weighting determined by ϵ_G then prevents good convergence of the AMGr algorithm. For example, for the variable-coefficient problem with $\mu = 3$, $\epsilon_G = 1573$, and the upper bound on the convergence factor of AMGr with four relaxation sweeps is 0.999996. Instead of using the simple Gerschgorin bound for these problems, we use the approximation, ϵ_L after at most five steps of the Lanczos process. While computation of ϵ_L is more expensive than that of ϵ_G , we see that even a few steps of the Lanczos process is enough to generate an approximation to the true ϵ that is accurate to 3 digits.

Finally, we experiment with the choice of coarse grids for AMGr and α AMGr for

coarse grid	μ	$RQ(u)$	ν	ϵ	$\rho(\epsilon)$	ϵ_G	$\rho(\epsilon_G)$
standard coarsening 8×8	3	8.55×10^{-2}	4	6.27	0.67	11.16	0.56
red-black coarsening	3	7.85×10^{-2}	1	1.85	0.48	1.89	0.49
red-black coarsening	3	7.85×10^{-2}	2	1.85	0.43	1.89	0.24
semi-coarsening 8×15	3	7.78×10^{-2}	1	0.65	0.25	0.66	0.25
semi-coarsening 8×15	3	7.78×10^{-2}	2	0.65	0.07	0.66	0.07

Table III. Performance of AMGr and adaptive AMGr for various choices of the coarse grid

the constant-coefficient diffusion problem, with Dirichlet boundary conditions. The theory developed in Section 3 does not strictly apply in this case, but the algorithm itself is well defined. Because these experiments are mainly aimed at examining the roles of coarse grids and the ϵ they induce, we do not present a complete study of parameters μ and ν in Table III. Choosing a standard multigrid coarse grid for this problem yields an exact ϵ of approximately 6.27, with performance similar to that observed above after just 3 setup cycles. Choosing a denser coarse grid, such as the red-black checkerboard coarse grid, yields much improvement in ϵ , with the matrix A_{ff} being more diagonally dominant. With this improvement in ϵ comes a similar improvement in the performance of the AMGr scheme. Choosing a coarse-grid of similar density, that of semi-coarsening (removing every other line of nodes in one direction), gives an even smaller ϵ , and better AMGr performance. Because these alternate coarsenings yield stronger diagonal dominance (and not just dominance by a more general D), the estimates ϵ_G are much better approximations to the true ϵ .

5. Conclusions

The reduction-based variants of the classical and adaptive AMG algorithms presented here allow theoretical analysis of their performance based on the decomposition of a well-conditioned A_{ff} block. The convergence bound for AMGr depends only on the number of pre-relaxations used and the parameter, ϵ , from the spectral bounds, $D \leq A_{ff} \leq (1+\epsilon)D$. The adaptive AMGr setup scheme is shown to converge globally in the special case that $n_c = 2$ and, in general, locally. Numerical examples given demonstrate the convergence behavior of the α AMGr setup phase and performance of AMGr.

ACKNOWLEDGEMENTS

This work was sponsored by the Department of Energy under grant numbers DE-FG02-03ER25574 and DE-FC02-01ER25479, Lawrence Livermore National Laboratory under contract numbers B541045 and B537461, Sandia National Laboratory under contract number 15268, and the National Science Foundation under grant number DMS-0410318.

REFERENCES

1. Brandt A, McCormick SF, and Ruge JW. Algebraic multigrid (AMG) for sparse matrix equations, in *Sparsity and Its Applications*, ed.. Evans DJ. Cambridge University Press. Cambridge. (1984).
2. Ruge JW and Stüben K. Algebraic multigrid (AMG), in *Multigrid Methods*, ed.. McCormick SF. volume 3 of *Frontiers in Applied Mathematics*. 73–130. SIAM. Philadelphia, PA. (1987).
3. Cleary AJ, Falgout RD, Henson VE, Jones JE, Manteuffel TA, McCormick SF, Miranda GN, and Ruge JW. Robustness and scalability of algebraic multigrid. *SIAM J. Sci. Comput.* **21**(5), 1886–1908. (2000).
4. Henson V and Yang U. BoomerAMG: a parallel algebraic multigrid solver and preconditioner. *Applied Numerical Mathematics*. **41**, 155–177. (2002).
5. Krechel A and Stüben K. Parallel algebraic multigrid based on subdomain blocking. *Parallel Computing*. **27**, 1009–1031. (2001).
6. Yang U. On the use of relaxation parameters in hybrid smoothers. *Num. Lin. Alg. Appl.* **11**, 155–172. (2004).
7. De Sterck H, Yang U, and Heys J. Reducing complexity in parallel algebraic multigrid preconditioners. *SIAM J. on Matrix Analysis and Applications*. **27**(4), 1019–1039. (2006). Also available as LLNL Technical Report UCRL-JRNL-206780.
8. Brandt A. Algebraic multigrid theory: The symmetric case. *Appl. Math. Comput.* **19**, 23–56. (1986).
9. Brezina M, Cleary A, Falgout R, Henson V, Jones J, Manteuffel T, McCormick S, and Ruge J. Algebraic multigrid based on element interpolation (AMGe). *SIAM J. Sci. Comp.* **22**, 1570–1592. (2000).
10. Chartier T, Falgout R, Henson V, Jones J, Manteuffel T, McCormick S, Ruge J, and Vassilevski P. Spectral AMGe (ρ AMGe). *SIAM J. Sci. Comp.* 1–26. (2003).
11. Falgout RD and Vassilevski PS. On generalizing the AMG framework. *SIAM J. Numer. Anal.* **42**(4), 1669–1693. (2004). Also available as LLNL technical report UCRL-JC-150807.
12. Falgout RD, Vassilevski PS, and Zikatanov LT. On two-grid convergence estimates. *Num. Lin. Alg. Appl.* **12**, 471–494. (2005). Also available as LLNL technical report UCRL-JC-150807.
13. Vaněk P, Mandel J, and Brezina M. Algebraic multigrid by smoothed aggregation for second and fourth order elliptic problems. *Computing*. **56**, 179–196. (1996).
14. Brezina M, Falgout R, MacLachlan S, Manteuffel T, McCormick S, and Ruge J. Adaptive smoothed aggregation (α SA). *SIAM J. Sci. Comp.* **25**(6), 1896–1920. (2004).
15. Brezina M, Falgout R, MacLachlan S, Manteuffel T, McCormick S, and Ruge J. Adaptive smoothed aggregation (α SA) multigrid. *SIAM Rev.* **47**(2), 317–346. (2005).
16. Brezina M, Falgout R, MacLachlan S, Manteuffel T, McCormick S, and Ruge J. Adaptive algebraic multigrid. *SIAM J. Sci. Comp.* **27**(4), 1261–1286. (2006).
17. Brandt A. General highly accurate algebraic coarsening. *Elect. Trans. Numer. Anal.* **10**, 1–20. (2000).
18. Livne O. Coarsening by compatible relaxation. *Num. Lin. Alg. Appl.* **11**(2), 205–227. (2004).
19. Notay Y. Algebraic multigrid and algebraic multilevel methods: a theoretical comparison. *Num. Lin. Alg. Appl.* **12**, 419–451. (2005).
20. McCormick SF and Ruge JW. Multigrid methods for variational problems. *SIAM J. Numer. Anal.* **19**, 924–929. (1982).
21. Marek I and Pultarová I. A note on local and global convergence analysis of iterative aggregation-disaggregation methods. *Lin. Alg. App.* **413**, 327–341. (2006).
22. Mandel J and McCormick SF. A multilevel variational method for $Au = \lambda Bu$ on composite grids. *J. Comput. Phys.* **80**, 442–452. (1989).
23. MacLachlan S. *Improving Robustness in Multiscale Methods*. Ph.D. dissertation. University of Colorado at Boulder. Boulder, Colorado. 2004.