

## ALGEBRAIC DISTANCE FOR ANISOTROPIC DIFFUSION PROBLEMS: MULTILEVEL RESULTS\*

ACHI BRANDT<sup>†</sup>, JAMES BRANNICK<sup>‡</sup>, KARSTEN KAHL<sup>§</sup>, AND IRENE LIVSHITS<sup>¶</sup>

**Abstract.** In this paper, we motivate, discuss the implementation, and present the resulting numerics for a new definition of strength of connection which uses the notion of *algebraic distance* as defined originally in the bootstrap algebraic multigrid framework (BAMG). We use this algebraic distance measure together with compatible relaxation and least-squares interpolation to derive an algorithm for choosing suitable coarse grids and accurate interpolation operators for algebraic multigrid algorithms. The main tool of the proposed strength measure is the least-squares functional defined by using a set of test vectors that in general is computed using the bootstrap process. The motivating application is the anisotropic diffusion problem, in particular, with non-grid aligned anisotropy. We demonstrate numerically that the measure yields a robust technique for determining strength of connectivity among variables for both two-grid and multigrid bootstrap algebraic multigrid methods. The proposed algebraic distance measure can also be used in any other coarsening procedure assuming that a rich enough set of near-kernel components of the matrix for the targeted system is known or is computed as in the bootstrap process.

**Key words.** bootstrap algebraic multigrid, least-squares interpolation, algebraic distances, strength of connection

**AMS subject classifications.** 65N55, 65N22, 65F10

**1. Introduction.** Multigrid (MG) is a methodology for designing numerical iterative methods for solving sparse matrix equations. One of the key issues in developing an efficient multigrid method is the selection of the coarse grids, which must consist of a smaller number of degrees of freedom than the fine grid, yet still be rich enough to allow for the accurate representation of a smooth error. In the Algebraic Multigrid (AMG) framework considered here, the coarse grids, the fine-to-coarse residual transfer (restriction) operators, the coarse-to-fine correction transfer (prolongation) operators, and the coarse-grid operators are all computed in a two-level setup algorithm that proceeds recursively. A main tool used in the AMG setup algorithm is strength of connectivity between variables. For many problems, e.g., those arising from a discretization of a partial differential equation (PDE), the connectivity among variables and thereby a proper choice of these multigrid components are known and well understood theoretically. This is, however, not the case for the anisotropic diffusion problems that are in the focus of this paper.

Generally speaking, multigrid methods for solving systems of sparse linear equations  $Ax = b$  are all based on the smoothing property of relaxation. An error vector  $e$  is called  $\tau$ -smooth if its residual is smaller than  $\tau\|e\|$ . The basic observation [1] is that the convergence of a proper relaxation process slows down only when the current error is  $\tau$ -smooth with  $\tau \ll 1$ ; the smaller the  $\tau$  the slower the convergence. In other words, the convergence rate of relaxation deteriorates as the error becomes dominated by eigenvectors with eigenvalues

---

\*Received December 3, 2012. Accepted April 10, 2015. Published online on September 29, 2015. Recommended by T. Manteuffel. Brannick's work was supported by the National Science Foundation under grants DMS-1217142 and DMS-1320608. Kahl's work was supported by the Deutsche Forschungsgemeinschaft through the Collaborative Research Centre SFB-TR 55 "Hadron physics from Lattice QCD". Livshits' work was supported by DOE subcontract B591416.

<sup>†</sup>Department of Applied Mathematics & Computer Science, The Weizmann Institute of Science, 76100 Rehovot, Israel ([achi.brandtweizmann.ac.il](mailto:achi.brandtweizmann.ac.il)).

<sup>‡</sup>Department of Mathematics, Pennsylvania State University, University Park, PA 16802, USA ([brannick@psu.edu](mailto:brannick@psu.edu)).

<sup>§</sup>Fachbereich Mathematik und Naturwissenschaften, Bergische Universität Wuppertal, D-42097 Wuppertal, Germany, ([kkahl@math.uni-wuppertal.de](mailto:kkahl@math.uni-wuppertal.de)).

<sup>¶</sup>Department of Mathematical Sciences, Ball State University, Muncie IN, 47306, USA ([ilivshits@bsu.edu](mailto:ilivshits@bsu.edu)).

small in magnitude. This observation and the assumption that when relaxation slows down, the error vector  $e$  can be approximated in a much lower-dimensional subspace, are the main ideas behind the multigrid methodology. Very efficient geometric multigrid solvers have been developed for the case that this lower-dimensional subspace of smooth errors corresponds to functions defined on a well-structured grid (the *coarse* level) and can be approximated by easy-to-derive equations based, for example, on discretizing the same continuum operator that has given rise to the fine-level equations  $Ax = b$ . The coarse-level equations are solved recursively using a similar combination of relaxation and still-coarser-level approximations to the resulting smooth error.

To deal with more general problems, e.g., the ones where the fine-level system may not be defined on a well-structured grid nor perhaps arise from any continuum problem, algebraic multigrid methods have been developed. These methods require techniques for deriving the set of coarse-level *variables* and the coarse-level *equations* based solely on the (fine-level) matrix  $A$ . The basic approach, developed in [5, 6, 25] and nowadays called classical AMG or RS-AMG, involves the following two steps:

- (1) Choosing the coarse-level variables as a subset  $C$  of the set  $\Omega$  of fine variables such that each variable in  $\Omega$  is *strongly connected* to variables in  $C$ .
- (2) Approximating the fine-level residual problem  $Ae = r$  by the coarse-level equations  $A_c e_c = r_c$  using the Galerkin prescription  $A_c = RAP$  and  $r_c = Rr$  yielding an approximation  $Pe_c$  to  $e$ .

The *interpolation matrix*  $P$  and the *restriction matrix*  $R$  are both defined directly in terms of the entries of the matrix  $A$ . Their construction relies on the notion of *strong connections*, developed originally for  $M$ -matrices, to provide a measure of the coupling of the variables used explicitly to coarsen the problem.

In the past two decades, numerous extensions of the classical AMG algorithm have been introduced, including modifications to the coarse-grid selection algorithms [17, 18, 20] and the definition of interpolation [12, 13, 14, 16, 27]. These works are motivated by the fact that the applicability of the original AMG algorithm is limited by the  $M$ -matrix heuristics upon which its strength of connectivity measure among variables is based. In particular, the measure and, hence, the classical AMG approach yields an efficient solver for problems where the matrix  $A$  has a dominant diagonal and, with small possible exceptions, all its off-diagonal elements have the same sign. Even then, the produced interpolation can have limited accuracy, insufficient for full multigrid efficiency. An example where the performance of AMG methods can deteriorate, and the one we consider here, is given by non-grid aligned scalar anisotropic diffusion problems. We refer to the paper [23], in which a smoothed aggregation variant of AMG is applied to this same test problem, and its performance is shown to deteriorate for non-grid aligned constant coefficient cases. Table VI in [23] gives results for two of the non-grid aligned tests we consider, namely, for the rotation angles  $\pi/4$  and  $\pi/8$ . As an example, we note that for  $\theta = \pi/8$ , the number of PCG iterations of the SA method doubles (grows from 9 to 18) as the problem size increases from  $n = 16^2$  to  $n = 128^2$ .

Advances in the design of AMG methods for anisotropic problems include the development of new notions of strength of connection that are used to choose coarse variables and the sparsity pattern and coefficients of interpolation [8, 9, 10, 23]. In [8], an energy-based measure of strength of connection was developed and tested for anisotropic diffusion problems. The approach uses a measure of strength of connectivity among variables based on approximations of the columns of the inverse of the system matrix that are computed using local relaxation. In [23], a related approach based on an evolution measure is developed and then applied to similar anisotropic model problems such as we consider here. The basic approach considers the connection between weighted-Jacobi-relaxation and the time integration of ordinary differ-

ential equations for the specific case of evolving  $\delta$ -functions to form the proposed strength measure. The works in [9, 10] develop a strength measure based on local energy estimates for interpolation to form aggregates for non-grid aligned anisotropic problems.

Closely related to these works is the approach of compatible relaxation [2], which uses a modified relaxation scheme to expose the character of the slow-to-converge (i.e., algebraically smooth) error. Coarse-grid points are then selected where this error is the largest, thus avoiding explicit use of a measure of strength of connection in choosing the set of coarse variables. A theoretical framework for CR-based coarsening is introduced in [19], and extensive tests for anisotropic diffusion problems on structured and unstructured meshes are found in [7, 11, 12, 21]. A related smoothed aggregation-based aggressive coarsening algorithm that uses the evolution measure from [23] is integrated with an energy minimizing form of interpolation in [26], yielding an effective solver for two-dimensional non-grid aligned anisotropic problems. In this work, the CR serves as an analysis tool to develop the aggregation scheme.

While these developments have resulted in marked improvements in certain cases, generally speaking, all existing methods require a substantial overlap of the coarse-grid basis functions (columns of  $P$ ) to obtain fast multigrid convergence for general non-grid aligned anisotropic problems. In the BAMG context, this amounts to using a high-caliber interpolation (i.e., with large interpolatory sets), which leads to a rapid fill-in of the coarse-grid operators.

In the present article, we focus on developing an alternative compatible relaxation coarsening algorithm that uses an *algebraic distance* measure to select coarse variables and interpolatory sets. Our aim is to investigate the suitability of such an approach together with a low-caliber (i.e., one with small interpolatory sets) interpolation constructed using the BAMG framework for solving anisotropic diffusion problems.

The algebraic distance measure we propose is based on a notion of strength of connectivity among variables that is derived from the local least-squares (LS) formulation for computing caliber-one interpolation in the BAMG process [3, 4, 24]. The approach first constructs a caliber-one LS interpolation operator for a given set of test vectors (representatives of algebraically smooth error) and then defines the algebraic distance between a fine point and its neighboring points in terms of the values of the local least-squares functionals resulting from the so-defined interpolation. The algebraic distance measure thus aims to address the issue of strength of connections in a general context—the goal being to determine explicitly those degrees of freedom from which a high quality least-squares interpolation for some given set of test vectors can be constructed. We note that the idea of basing strength of connection on the suitability of interpolation for a set of smooth (near kernel) test vectors has also been considered in [23].

The resulting measure of distance (connectivity) among variables is used to derive a strength graph which is then passed to a coloring algorithm [15, Chap. 8] to coarsen the unknowns at each stage of a compatible relaxation coarsening algorithm. Given the set of coarse variables, a similar measure is used in defining interpolation. The idea is to approximately minimize the values of the LS functional locally [3, 4]. This is accomplished by forming the LS-based interpolation for a given fine point and various candidate interpolatory sets (neighboring pairs (or sets) of coarse points) and then by choosing the set with smallest value of the associated LS functional to define the row of interpolation.

The remainder of the paper is organized as follows. Section 2 contains an introduction to the bootstrap algebraic multigrid components with an emphasis on the least-squares interpolation approach and compatible relaxation coarsening algorithm. Then, in Section 3, a general definition of strength of connection and the notion of algebraic distance as well as its connection to compatible relaxation are discussed. The diffusion equation with anisotropic

coefficients and its discretizations are introduced in Section 4 together with results of numerical experiments of our approach applied to these systems. Conclusions are presented in Section 5.

**2. Bootstrap AMG.** The bootstrap AMG framework, introduced in [3, §17.2], was proposed to extend AMG to general (non  $M$ -matrix) problems. The framework combines the following two general devices to inexpensively construct high-quality interpolation.

- (A) The interpolation is derived to provide the best least-squares fit to a set of  $\tau$ -smooth test vectors (TVs) obtained by a process described below.

Denote by  $C_i$  the set of coarse variables used in interpolating to the fine-grid variable  $i$ . It follows from the satisfaction of the compatible relaxation criterion (see the next section) that, with proper choices of  $C_i$  for all  $i$ , there exists an interpolation operator that approximates *any* vector  $x$  which is  $\tau$ -smooth with an error proportional to  $\tau$ , i.e., there exists an interpolation operator that satisfies the so-called weak approximation property. The proof of this result for the special case of the so-called ideal interpolation operator is given in [19]. The size  $|C_i|$  of this set should in principle increase as  $\tau$  decreases (and smaller  $\tau$  means overall better multigrid convergence), but in practice a pre-chosen and sufficiently small interpolation *caliber*

$$c := \max_{i \in \Omega} |C_i|$$

often yields small enough errors as demonstrated by the extensive numerical tests presented in [11].

The set  $C_i$  can often be adequately chosen by natural considerations such as the set of geometrical neighbors with  $i$  in their convex hull. If the chosen set is inadequate, the least-squares procedure will show a *poor fitness* (the interpolation errors are large compared with  $\tau$ ), and the set can then be extended. The least-squares procedure can also detect variables in  $C_i$  that can be discarded without a significant accuracy loss. Thus, this approach allows creating interpolation with whatever needed accuracy which is *as sparse as possible*.

- (B) Generally, the test vectors are constructed in a bootstrap manner, in which several tentative AMG levels are generated by interpolation fitted to only moderately smooth TVs; this tentative (multilevel) structure is then used to produce *improved* (much smoother) TVs, yielding a more accurate interpolation operator. The process continues if needed until fully efficient AMG levels have been generated.

The first test vectors are each produced by relaxing the homogeneous system

$$(2.1) \quad Av = 0$$

using different initial vectors, leading to an initial  $\tau$ -smooth test set. A mixture of random vectors and/or diverse geometrically smooth vectors can generally be used as initial approximations. The latter test vectors may not require relaxation at all or, perhaps, only near boundaries or other regions where singularities or discontinuities are present. In the case of discretized isotropic PDEs, if geometrically smooth vectors that satisfy the homogeneous boundary conditions are used, relaxation may not be needed at all. For the anisotropic problem considered, we use a constant vector and bootstrapped test vectors generated from initially random initial guesses to define an algebraic distance notion of strength of connection and the least-squares interpolation operator.

**2.1. Least-squares interpolation.** The basic idea of the least-squares interpolation approach is to approximate a set of test vectors,  $\{v^{(1)}, \dots, v^{(k)}\} \subset \mathbb{R}^n$ , minimizing the interpolation error for these vectors in a least-squares sense. In the context considered here, namely, applying the least-squares process to construct a classical AMG form of interpolation, each row of  $P$ , denoted by  $p_i$ , is defined as the minimizer of a local least-squares functional. Given

a splitting of variables  $C$  and  $F = \Omega \setminus C$ , for each  $i \in F$  and for various  $C_i \subset C$ , find  $p_i \in \mathbb{R}^{n_c}$ ,  $n_c = |C|$ , that minimize

$$(2.2) \quad \mathcal{LS}(p_i) = \sum_{\kappa=1}^k \omega_{\kappa} \left( v_i^{(\kappa)} - \sum_{j \in C_i} (p_i)_j v_j^{(\kappa)} \right)^2 \mapsto \min.$$

To guarantee the uniqueness of the solution of (2.2), the number of test vectors  $k$  should be greater than or equal to the caliber  $c$ . Further, in practice, we have observed that the accuracy of the least-squares interpolation operator and, hence, the performance of the resulting solver generally improves with increasing  $k$  [4] up to some value proportional to the caliber  $c$ . The weights  $\omega_{\kappa} > 0$  are chosen to reflect the global algebraic smoothness of the test vectors. We give their specific choice in the numerical experiments section.

The quality of the LS interpolation can be further improved by applying a sweep of local Jacobi relaxations to the test vectors prior to computing the LS-functional. Equivalently, the action of such pre-smoothing can be directly built in by considering a residual-based LS process, see [22]:

$$(2.3) \quad \mathcal{LS}(p_i) = \sum_{\kappa=1}^k \omega_{\kappa} \left( v_i^{(\kappa)} + \frac{1}{a_{ii}} r_i^{(\kappa)} - \sum_{j \in C_i} (p_i)_j v_j^{(\kappa)} \right)^2 \mapsto \min,$$

where  $r^{(\kappa)}$  is the residual of  $v^{(\kappa)}$  in (2.1).

**2.2. Compatible relaxation.** A general criterion for choosing an adequate set of coarse variables is the fast convergence of *compatible relaxation* (CR) as introduced in [2] and further developed in [11, 12, 19, 21].

Compatible relaxation can be generally employed in two capacities. First, given a relaxation scheme and a coarse-grid set  $C$ , it can predict multigrid convergence; habituated CR in [21] gives the best estimates. Second, given a relaxation scheme and a desired convergence rate, it can be used to construct an adequate coarse grid. A brief introduction to the details of CR relevant to the content of this paper is presented next.

Given a matrix  $A \in \mathbb{R}^{n \times n}$  and a suitable relaxation process with error propagation matrix  $E = I - M^{-1}A$ , assume that a classical AMG coarse-and-fine level splitting  $\Omega = C \cup F$  has been selected. Then, one choice of compatible relaxation is given by  $F$ -relaxation for the homogeneous system, i.e., relaxation applied only to the set of  $F$ -variables. In other words, ordering the unknowns according to the partitioning of  $\Omega$  into  $F$  and  $C$ :

$$u = \begin{bmatrix} u_f \\ u_c \end{bmatrix}, \quad A = \begin{bmatrix} A_{ff} & A_{fc} \\ A_{cf} & A_{cc} \end{bmatrix}, \quad \text{and} \quad M = \begin{bmatrix} M_{ff} & M_{fc} \\ M_{cf} & M_{cc} \end{bmatrix}.$$

The  $F$ -relaxation form of compatible relaxation is then defined by

$$(2.4) \quad u_f^{\nu+1} = E_f u_f^{\nu} \quad \text{with} \quad E_f = I - M_{ff}^{-1} A_{ff}.$$

If  $A$  is symmetric and positive definite and  $M$  is symmetric, then the asymptotic convergence rate of compatible relaxation,

$$\rho_f = \rho(E_f),$$

where  $\rho$  denotes the spectral radius, provides a measure of the quality of the coarse space, that is, a measure of the ability of the set of coarse variables to represent errors not eliminated by

the given fine-level relaxation. This measure can be approximated using  $F$ -relaxation for the homogeneous system with a random initial guess  $u_f^0$ . Since  $\lim_{\nu \rightarrow \infty} \|E_f^\nu\|^{1/\nu} = \rho(E_f)$  for any norm  $\|\cdot\|$ , the measure

$$(2.5) \quad \varrho_f = (\|u_f^\nu\|/\|u_f^0\|)^{1/\nu}$$

estimates  $\rho_f$  and provides a measure of the quality of the coarse-variable set.

As a tool for choosing  $C$ , CR can be used in the following way. Starting with an initial set of coarse variables  $C_0$ , a few sweeps of compatible relaxation will detect slowness if the current set is inadequate and expose viable candidates to be added to  $C$ , those that are slow to converge to zero in the CR process. A subset of these variables is then added to  $C_0$  to form a new coarse set  $C_1$ . The process repeats until fast convergence of CR is obtained. In practice,  $C_0$  may be empty or for structured problems, standard coarsening can be used. The process continues until the desired  $\rho_f$  is achieved.

Current CR algorithms do not use a strength of connection measure in constructing the coarse-variable set. Instead, they rely on the error produced by CR to form candidate sets of potential  $C$ -points. Our aim is to develop a more general notion of strength of connections based on algebraic distances and to explore its use in a compatible relaxation coarsening scheme and in defining the interpolatory set for each row of interpolation.

**3. Selecting coarse variables and interpolation via algebraic distances.** The classical definition of strength of connection is intended for the case of diagonally dominant M-matrices; it can break down when applied to problems involving more general classes of matrices such as the anisotropic problems considered in this paper. The near null space of a diagonally dominant M-matrix is typically well approximated locally by the constant vector, and the AMG strength of connections measure succeeds when this assumption is reflected in the coefficients of the system matrix  $A$ . (Similar observations motivated the work in [23].) If either the near null space cannot be accurately characterized as locally constant or this is not reflected in the matrix coefficients, then performance of classical AMG suffers.

As a more general measure of strength of connection, we consider a variable's ability to interpolate  $\tau$ -smooth errors for small  $\tau$  to its neighbors. Specifically, for a given fine-grid point  $i \in F$ , we construct a row of LS interpolation from each of its neighbors  $j \in N_i$ ,  $N_i$  being defined in terms of the graph of powers of  $A$ , and monitor the values of the LS functional. If for some  $j \in N_i$ , the LS functional is small relative to its size for other neighbors, then  $j$  is determined to be strongly connected to  $i$ . This process, repeated for each  $i \in F$ , allows to identify suitable coarse-grid points as those from which it is possible to build a high-quality LS interpolation to its fine-grid neighbors. Next, we describe in detail the idea of measuring strength between neighbors using algebraic distance and then discuss how this measure can be incorporated in a CR coarsening algorithm and in computing the nonzero sparsity pattern of interpolation.

**3.1. Strength of connection by algebraic distances.** In the simplest form, the definition of algebraic distances is straightforward. For any given pair of fine variables  $i, j \in \Omega$ , compute

$$(3.1) \quad \mu_{ij} = \frac{1}{\sum_{\kappa=1}^k \omega_\kappa \left( v_i^{(\kappa)} + \frac{1}{a_{ii}} r_i^{(\kappa)} - p_{ij} v_j^{(\kappa)} \right)^2},$$

where  $p_{ij}$  is the minimizer of the least-squares problem for a single variable  $j$ . We note that there are many possible choices for defining a measure of strength of connection based on  $\mathcal{LS}$ . We choose to take the reciprocal value of  $\mathcal{LS}$  to define strength since this coincides with

the usual notion of strength between unknowns in that two unknowns are classified as being strongly coupled by a large value of the given measure. Further, our experience suggests that this measure is robust for the targeted anisotropic problems. We note that the definition of  $\mu_{ij}$  is not symmetric since  $\mu_{ji}$  involves different entries of the test vectors used in computing LS interpolation than  $\mu_{ij}$  does.

For a given SPD matrix  $A$ , we define its connectivity graph as  $\mathcal{G} = (\mathcal{V}, \mathcal{E})$ , where  $\mathcal{V}$  and  $\mathcal{E}$  are the sets of vertices and edges. Here, an edge  $(i, j) \in E \iff (A)_{ij} \neq 0$ . Let  $\mathcal{G}_d(\mathcal{V}_d, \mathcal{E}_d)$  denote the graph of the matrix  $A^d$  and define  $\mathcal{G}_{d,i}(\mathcal{V}_{d,i}, \mathcal{E}_{d,i})$  as the subgraph associated with the  $i$ th vertex and its algebraic neighbors:

$$(3.2) \quad \mathcal{V}_{d,i} := \{j \mid (A^d)_{ij} \neq 0\} \quad \text{and} \quad \mathcal{E}_{d,i} := \{(i, j) \mid (A^d)_{ij} \neq 0\}.$$

Then, given a search depth  $d$  and a fine variable  $i \in \Omega$ , we compute  $\mu_{ij}$  for all  $j \in \mathcal{V}_{d,i}$ . This simplification, combined with the idea of deriving strength according to an algebraic distance measure based on simple caliber-one interpolation, allows us to control the complexity of the algorithm. More generally, the algebraic distance measure can be computed for sets of neighboring coarse points and, again, the LS functional can be used as an a posteriori measure of the quality of the interpolatory set  $C_i$ . We use these observations in the design of our algorithm for computing the interpolation coefficients discussed in Section 3.3.

**3.2. Compatible relaxation coarsening and algebraic distances.** In selecting  $C$ , we integrate the simplified variant of the algebraic distance notion of strength of connection based on caliber-one interpolation (3.1) into the CR-based coarsening algorithm developed in [11]. The notion of algebraic distances is used to form a subgraph of the graph of the matrix  $A^d$ ,  $d = 1, 2, \dots$ . Specifically, the algebraic distance between any two adjacent vertices in the graph  $\mathcal{G}_d$  of  $A^d$  is computed using (3.1). Then, for each vertex  $i \in F$ , we remove edges adjacent to  $i$  with small weights relative to the largest weight of all edges connected to  $i$ :

$$(3.3) \quad \mathcal{V}_{\mathcal{M}} = F, \quad \mathcal{E}_{\mathcal{M}} = \{(i, j) \mid i, j \in F \quad \text{and} \quad \mu_{ij} > \theta_{ad} \max_k \mu_{ik}\},$$

with  $\theta_{ad} \in (0, 1)$ . This, in turn, produces the graph  $\mathcal{M}^d(\mathcal{V}_{\mathcal{M}}, \mathcal{E}_{\mathcal{M}})$  of strongly connected vertices. Note that by definition, the strength graph is restricted to vertices  $i \in F$  so that it can be used in the subsequent CR coarsening stages. (Although the measure (3.1) is able to determine the coupling between any given two points, in order to make the idea practical we restrict its use to local neighborhoods; in cases where  $\mu_{ij} = \infty$ , variables  $i$  and  $j$  are defined as strong neighbors. In such cases, we do not consider these links in the later stages of the algorithm, that is, they are not used in computing the maximum in (3.3).) The strength graph  $\mathcal{M}^d$  is then passed to a coloring algorithm [15, Chap. 8] as in the classical AMG approach, in which coarse points are selected based on their number of strongly connected neighbors.

A general description of the overall CR coarsening approach is given by Algorithm 1. For additional specific details of the CR algorithm we refer the reader to [11], in which this scheme was developed for diffusion problems similar to those we consider here.

Here,  $\delta \in (0, 1)$  is the tolerance for the approximate CR convergence rate  $\rho_f$  computed using (2.5), and  $u^0$  is an initial guess for the compatible relaxation iterations. Further,  $d$  is fixed in advance, but more generally it can be adapted at each stage of the CR algorithm. We note also that the matrix  $A^d$ ,  $d > 1$ , is not implicitly formed except for its binary adjacency matrix, as this is all that is needed to construct  $\mathcal{M}^d$ . Our choices for these and other parameters of the algorithm used in our numerical experiments are given in Section 4.2.

---

**Algorithm 1** compatible\_relaxation      {Computes  $C$  using Compatible Relaxation}.

---

*Input:*  $u^0$ ,  $C_0$  { $C_0 = \emptyset$  is allowed}.  
*Output:*  $C$   
 Initialize  $C = C_0$   
 Initialize  $F = \Omega \setminus C$   
 Perform  $\nu$  CR iterations starting with an initial guess  $u^0$   
**while**  $\rho_f > \delta$  **do**  
      $C = C \cup \{ \text{maximum independent set of } \mathcal{F} \}$  based on  $\mathcal{M}^d$   
      $F = \{i \in \Omega \setminus C : \sigma_i > \text{tol}\}$   
     Perform  $\nu$  CR iterations of Eq. (2.4) with initial guess  $u^0$   
**end while**

---

Various choices of the candidate set measure,  $\sigma_i$ , used in determining potential  $C$ -points have been studied in the literature [2, 11, 21]; we follow the definition in [11]:

$$\sigma_i := \frac{|u_i^\nu|}{\|u^\nu\|_\infty}.$$

In practice, this measure gives the best overall results for smaller values of  $\nu$ , say  $\nu = 5$ .

**3.3. Defining interpolation by algebraic distances.** Given a set of coarse variables  $C$  and a set of  $\tau$ -smooth test vectors  $\{v^{(1)}, \dots, v^{(k)}\}$ , we use algebraic distance to find the coarse interpolatory sets  $C_i$  with a cardinality bounded by a given caliber  $c$ . More specifically, we consider all possible sets of  $C$ -points  $W$  of cardinality up to  $c$  in the  $d_{LS}$ -ring coarse-point neighborhood of a given  $F$ -point  $i$  defined as

$$(3.4) \quad \mathcal{N}_{d_{LS},i} := C \cap \mathcal{V}_{d_{LS},i},$$

where  $d_{LS}$  is the search depth for constructing least-squares interpolation. That is,  $\mathcal{V}_{d_{LS}}$ ,  $\mathcal{E}_{d_{LS}}$ , and  $\mathcal{G}_{d_{LS},i}$  are defined as in (3.2), where  $d_{LS}$  is a fixed positive integer. Thus, the sets of possible interpolatory points can be written as

$$\mathcal{W} := \{W \mid W \subseteq \mathcal{N}_{d_{LS},i} \text{ and } |W| \leq c\}.$$

Using an exhaustive search of this set, we find the minimizer of the least-squares functional: for each  $i \in F$

$$C_i = \arg \min_{W \in \mathcal{W}} \mathcal{LS}(p_i(W)),$$

where

$$\mathcal{LS}(p_i(W)) = \sum_{\kappa=1}^k \omega_\kappa \left( v_i^{(\kappa)} + \frac{1}{a_{ii}} r_i^{(\kappa)} - \sum_{j \in W} (p_i)_j v_j^{(\kappa)} \right)^2.$$

Here  $p_i$  denotes the minimizer of the least-squares problem (2.3) for the given set  $W$ . Thus, we must compute  $p_i$  and evaluate  $\mathcal{LS}(p_i(W))$  for all possible choices of  $W \in \mathcal{W}$ . The row of interpolation  $p_i$  is then chosen as the one that minimizes  $\mathcal{LS}(p_i(W))$ .

An additional detail of the approach is the penalization of large interpolatory sets. It is easily shown that for two sets  $W' \subset W''$ , their corresponding minimal least-squares functional values fulfill  $\mathcal{LS}' \geq \mathcal{LS}''$ . In order to keep the interpolation operator as sparse as possible, we require that the least-squares functional is reduced by a certain factor when increasing the

cardinality of  $W$ . That is, a new set of points  $W''$  is preferred over a set of points  $W'$  with  $|W''| > |W'|$  if

$$\mathcal{L}S'' < (\mathcal{L}S')^{\gamma(|W''|-|W'|)}.$$

Based on numerical experience, we choose  $\gamma = 1.5$  which tends to produce accurate and sparse interpolation operators for a large class of problems.

The above exhaustive search of all possible combinations of interpolatory sets with cardinality up to some given caliber is one of many possible strategies for selecting  $C_i$ . In our experience, this is often not necessary, and scanning a small number of possibilities based on the algebraic distance strength measure and caliber-one interpolation is sufficient for many problems. That is, the exact minimization of the LS functional is often not required to obtain sufficiently accurate interpolation. However, for the anisotropic diffusion problem we consider here, in the case of strong non-grid aligned anisotropies, an exhaustive search leads to the best overall results for the low-caliber interpolation, and therefore we use this strategy in our numerical experiments.

**4. Numerical results.** In this section, we present numerical tests that demonstrate the effectiveness of the algebraic distance measure of strength of connection when combined with CR coarsening and LS interpolation. The tests of the approach consist of a variety of 2D anisotropic diffusion problems discretized using finite differences and finite elements on an  $(N + 1) \times (N + 1)$  uniform grid.

**4.1. Model problem and discretizations.** We consider finite difference and bilinear finite element discretizations of the two-dimensional diffusion operator

$$(4.1) \quad \mathcal{L}u = \nabla \cdot \mathcal{K} \nabla u,$$

with an anisotropic diffusion coefficient

$$(4.2) \quad \mathcal{K} = \begin{bmatrix} \cos \alpha & -\sin \alpha \\ \sin \alpha & \cos \alpha \end{bmatrix} \begin{bmatrix} 1 & 0 \\ 0 & \epsilon \end{bmatrix} \begin{bmatrix} \cos \alpha & \sin \alpha \\ -\sin \alpha & \cos \alpha \end{bmatrix},$$

where  $0 < \epsilon < 1$  and  $0 \leq \alpha < 2\pi$ . Changing variables

$$\begin{bmatrix} \xi \\ \eta \end{bmatrix} = \begin{bmatrix} \cos \alpha & \sin \alpha \\ -\sin \alpha & \cos \alpha \end{bmatrix} \begin{bmatrix} x \\ y \end{bmatrix}$$

yields strong connections aligned with the direction  $\xi$ :

$$\mathcal{L}u(\xi, \eta) = u_{\xi\xi} + \epsilon u_{\eta\eta}.$$

An equivalent formulation in  $(x, y)$ -coordinates is given by

$$(4.3) \quad \mathcal{L}u(x, y) = au_{xx} + 2bu_{xy} + cu_{yy},$$

where  $a = \cos^2 \alpha + \epsilon \sin^2 \alpha$ ,  $b = \frac{(1 - \epsilon)}{2} \sin 2\alpha$ , and  $c = \sin^2 \alpha + \epsilon \cos^2 \alpha$ .

In formulating a finite difference discretization of (4.3), we consider a standard five-point discretization for the Laplacian term and then define the discretization of the  $u_{xy}$  term using lower-left and upper-right neighbors and not the lower-right and upper-left ones for each fine-grid point  $i \in \Omega$ , yielding an overall stencil that includes seven nonzero values. This gives a suitable discretization for  $\alpha = \pi/4$  with the stencil for  $u_{xy}$  given by

$$(4.4) \quad \tilde{S}_{xy} = \frac{1}{2h^2} \begin{bmatrix} 0 & -1 & 1 \\ 0 & 1 & -1 \\ 0 & 0 & 0 \end{bmatrix} + \frac{1}{2h^2} \begin{bmatrix} 0 & 0 & 0 \\ -1 & 1 & 0 \\ 1 & -1 & 0 \end{bmatrix}.$$

In contrast, for the  $\alpha = -\pi/4$  case, the direction of the anisotropy can not be captured by our chosen discretization and yields a matrix that is far from an M-matrix. For example, taking  $\epsilon = 0.1$ , the resulting system matrix for this choice of  $\alpha$  has a stencil

$$S_A = \frac{1}{2h^2} \begin{bmatrix} & -1 & 0.45 \\ -1 & 3.1 & -1 \\ 0.45 & -1 & \end{bmatrix},$$

and thus it is not even approximately an  $M$ -matrix, which makes an MG solution of this problem challenging. (We note that the  $\alpha = \pi/8$  case we consider in our tests also yields non-M-matrices for strong anisotropies.) Here, some of the off-diagonal entries in  $A$  are positive and, hence, the heuristics motivating the classical definition of strength of connection are not applicable. We consider this extreme case, although it is unlikely to arise in practice, to demonstrate the robustness of the BAMG approach as a coarsening strategy for the targeted anisotropic problems. We mention that on a structured grid, our chosen seven-point discretization is equivalent to the finite element discretization of the same elliptic boundary value problem using triangular finite elements.

In addition, we consider the bilinear finite element discretization of the same model problem again on an  $(N + 1) \times (N + 1)$  uniform grid. Letting  $\phi_j(x, y)$  denote a standard bilinear basis function that is one at node  $j$  and zero at all other nodes and writing the solution as  $\sum_{j=1}^n u_j \phi_j$ , the weak form of (4.1) is given by

$$-\sum_{j=1}^n u_j \int_0^1 \int_0^1 (\nabla \phi_i) \cdot \begin{bmatrix} a & b \\ b & c \end{bmatrix} \nabla \phi_j dx dy = \int_0^1 \int_0^1 f \phi_i dx dy, \quad i = 1, \dots, n.$$

And, the corresponding nine point stencil for the global stiffness matrix  $A$  is given by

$$S_A = \begin{bmatrix} -a + 3b - c & 2(a - 2c) & -a - 3b - c \\ 2(-2a + c) & 8(a + c) & 2(-2a + c) \\ -a - 3b - c & 2(a - 2c) & -a + 3b - c \end{bmatrix}.$$

As an example, when  $\theta = \frac{\pi}{4}$ , we have

$$S_A = \begin{bmatrix} \frac{1}{2} - \frac{5}{2}\epsilon & -1 - \epsilon & -\frac{5}{2} + \frac{1}{2}\epsilon \\ -1 - \epsilon & 8 + 8\epsilon & -1 - \epsilon \\ -\frac{5}{2} + \frac{1}{2}\epsilon & -1 - \epsilon & \frac{1}{2} - \frac{5}{2}\epsilon \end{bmatrix} \longrightarrow \begin{bmatrix} \frac{1}{2} & -1 & -\frac{5}{2} \\ -1 & 8 & -1 \\ -\frac{5}{2} & -1 & \frac{1}{2} \end{bmatrix} \quad \text{as } \epsilon \longrightarrow 0,$$

which again is not an  $M$  matrix.

**4.2. Formulating a two-level coarsening algorithm.** Our aim is to study the robustness of the algebraic distance notion of strength of connection for grid and non-grid aligned anisotropy. We focus our tests on the seven-point finite difference and nine-point finite element discretizations introduced earlier in this section for various values of the anisotropy angle  $\alpha$  and the anisotropy coefficient  $\epsilon$ .

In all tests, coarse grids are chosen using the compatible relaxation approach discussed in Section 3.2, and interpolation operators are then computed using the LS approach given in Section 3.3 with the caliber set to  $c = 1$  or  $c = 2$ . The stopping tolerance for CR steps is set as  $\rho_f < \delta = 0.7$ , the number of CR sweeps is chosen as  $\nu = 5$ , and the tolerance for the candidate set measure is  $\text{tol} = 1 - \rho_f$ . A larger choice of the stopping tolerance generally leads to more aggressive coarsening whereas our choice of  $\text{tol}$  ensures that points are added to the candidate set more sparingly at later stages of the CR coarsening process. We fix the

strength of connection parameter used in forming the strength graph at  $\theta_{ad} = 0.5$  and vary the graph distance  $d$  used to define the graph  $\mathcal{G}_d$  from which  $\mathcal{M}^d$  defined as in (3.3) is constructed. We note that generally the overall quality of the grids the algorithm produces depends only mildly on the choice of  $\theta_{ad}$ .

The search depth, used in defining the greedy approach for choosing LS interpolation as in (3.4) is fixed as  $d_{LS} = d + 2$ . Taking the value of the search depth larger than the coarsening depth allows the approach to scan a larger number of possible interpolatory sets in forming long-range interpolation whenever the problem requires it. In this way, the LS scheme for constructing interpolation is able to better follow a wider range of values of the anisotropy angle  $\alpha$ . We mention, in addition, that for the anisotropic problems we consider, it is important that the set of test vectors used in guiding the coarsening consists of the *characteristic components*, i.e., eigenvectors with small eigenvalues, induced by the anisotropy angle  $\alpha$ , especially for angles for which longer-range interpolation is needed in order to follow the anisotropy. In our two-grid tests, where only relaxation is used to compute the test vectors, we thus apply 40 iterations of Gauss-Seidel to seven distinct random initial guesses plus the constant vector in order to compute the eight test vectors used to construct LS interpolation in (2.3). The weights in (2.3) are defined as

$$\omega_\kappa = \frac{\langle v^{(\kappa)}, v^{(\kappa)} \rangle}{\langle Av^{(\kappa)}, v^{(\kappa)} \rangle}.$$

The number of test vectors and especially the amount of relaxation required to compute them can be reduced by replacing a single-grid relaxation procedure by a multilevel bootstrap cycling scheme [4] as we show in the next section where we reduce the number of relaxation steps from 40 to eight in the multilevel setup.

When presenting results of the solver constructed by the resulting BAMG setup algorithm, we use two pre- and post-Gauss-Seidel relaxation sweeps on all grids except the coarsest, where a direct solver is used. Here, the use of two pre- and post-smoothing steps is motivated by the fact that the setup algorithm generally produces aggressive coarsening. The estimates of the asymptotic convergence rates are computed as

$$\rho = \frac{\|e^\eta\|_A}{\|e^{\eta-1}\|_A},$$

where  $e^\eta$  is the error after  $\eta = 100$  MG iteration applied to the homogeneous system, starting with random initial approximations. We also report the operator complexity ratio

$$(4.5) \quad \gamma_o = \frac{\sum_l \text{mz}(A_l)}{\text{mz}(A)},$$

and the grid complexity ratio

$$(4.6) \quad \gamma_g = \frac{\sum_l |C_l|}{|\Omega|},$$

where  $C_0 = \Omega$ .

**4.3. Bootstrap coarsening.** In this section, we illustrate the choice of coarse grids and interpolation patterns that the BAMG setup algorithm constructs when applied to the finite difference and bilinear finite element approximations of the anisotropic test problem for  $\epsilon = 10^{-4}$  with various choices of the anisotropy angle  $\alpha$  in (4.2). The tests of the BAMG setup algorithm we consider are for interpolation calibers  $c = 1, 2$  and varying values of the

search depth  $d$ . In the plots in Figures 4.1–4.8, the black lines depict the interpolation pattern for each  $F$ -point (denoted by the smaller circles) from its neighboring  $C$ -points denoted by the larger circles. Generally, we observe that the coarsening and interpolation pattern follow the anisotropy when the choice of the discretization allows it. For example, in Figures 4.1 and 4.2 the coarsening perfectly follows the direction of the anisotropy for both discretizations when using a search depth  $d = 1$  and setting  $\alpha = 0$  and  $\alpha = \pi/4$ . On the other hand, we see that this is not true for the case for  $\alpha = -\pi/4$  and the seven-point finite difference approximation and for  $\alpha = \pi/8$  for both discretizations. For these angles, the direction of strong coupling is not able to capture the direction of the anisotropy for the given discretization of the problem when setting the search depth to  $d = 1$  since the matrix entries in these directions are zero. Hence, for these problems it is not possible to form a strength matrix from immediate algebraic neighbors which produces a coarse grid that allows the interpolation pattern to exactly follow the anisotropy. These observations in fact lead naturally to the use of the graph  $\mathcal{G}_d$  of  $A^d$ ,  $d > 1$ , to form the strength matrix. Overall, we see that the algebraic distance strength measure leads to least-squares interpolation that follows the general direction of anisotropy to the extent that the value of the search depth  $d$  allows it.

Another interesting observation here is that for the  $d = 2$  tests reported in Figures 4.5–4.8, we see that by using the graph of  $A^2$  to form the algebraic-distance-based strength graph, the coloring algorithm is now able to coarsen in the direction of anisotropy for the finite difference discretization and  $\alpha = -\pi/4$ . We note that the use of a larger search depth in defining the strength matrix has the additional benefit that it allows more aggressive coarsening while at the same time maintaining the characteristic directions induced by more general anisotropic directions. The ability of the setup to produce aggressive coarsening is seen in the tests for grid-aligned anisotropy when comparing the grids obtained for the search depths  $d = 1$  to  $d = 2$ .

Another difficult choice of  $\alpha$  occurs for the anisotropy direction  $\pi/8$  for which a longer-range interpolation and an extended search depth for coarse-grid candidates is required to properly capture the direction of the anisotropy. For this example, the reference angles made by coarse-fine connections are much closer to  $\pi/8$  when using  $d = 2$  (see Figures 4.5–4.8) than they are when using  $d = 1$  (see Figures 4.1–4.4) in the setup for both the finite difference and finite element discretizations.

To conclude the section, we consider the performance of the algorithm for two different choices of the interpolation caliber  $c = 1$  and 2. Here, the algorithm chooses the same coarse grids independent of the choice of caliber, but the sparsity structure and also the values of the interpolation operators change. Generally, we have seen that using a larger value of  $c$  improves the convergence rates of the two-level methods  $\rho$  in all cases, however, the complexity of solving the coarse-level system also increases.

**4.4. The coarse-level operator.** Another interesting deliverable of the proposed BAMG setup algorithm, in particular of its implementation of the compatible relaxation and the algebraic distances, is the pattern of the resulting coarse-grid stencil. Discretizations involving the finite difference discretization (4.4) favor  $\alpha = \pi/4$  and with the same argument result in the worst possible discretization for  $\alpha = -\pi/4$  for which the use of the upper-left and lower-right grid-point neighbors in the discretization of  $\partial_{xy}$  would be appropriate. We consider both cases  $\alpha = \pi/4$  and  $\alpha = -\pi/4$  for the seven-point finite difference discretization given in (4.4). For both cases, we assume  $\epsilon = 10^{-10}$ ,  $d = 2$ , and  $d_{LS} = 4$ .

We first confirm that for  $\alpha = \pi/4$ , the coarse-grid operator  $A_c = P^T A P$  preserves the intrinsic strength of connections inherited from the fine-grid operator  $A$ . A typical example of the stencil of  $A_c$  is given in Figure 4.9(a). Here, the details of the configuration of each stencil depend on the coarsening pattern in the neighborhood of the considered coarse-grid equation.

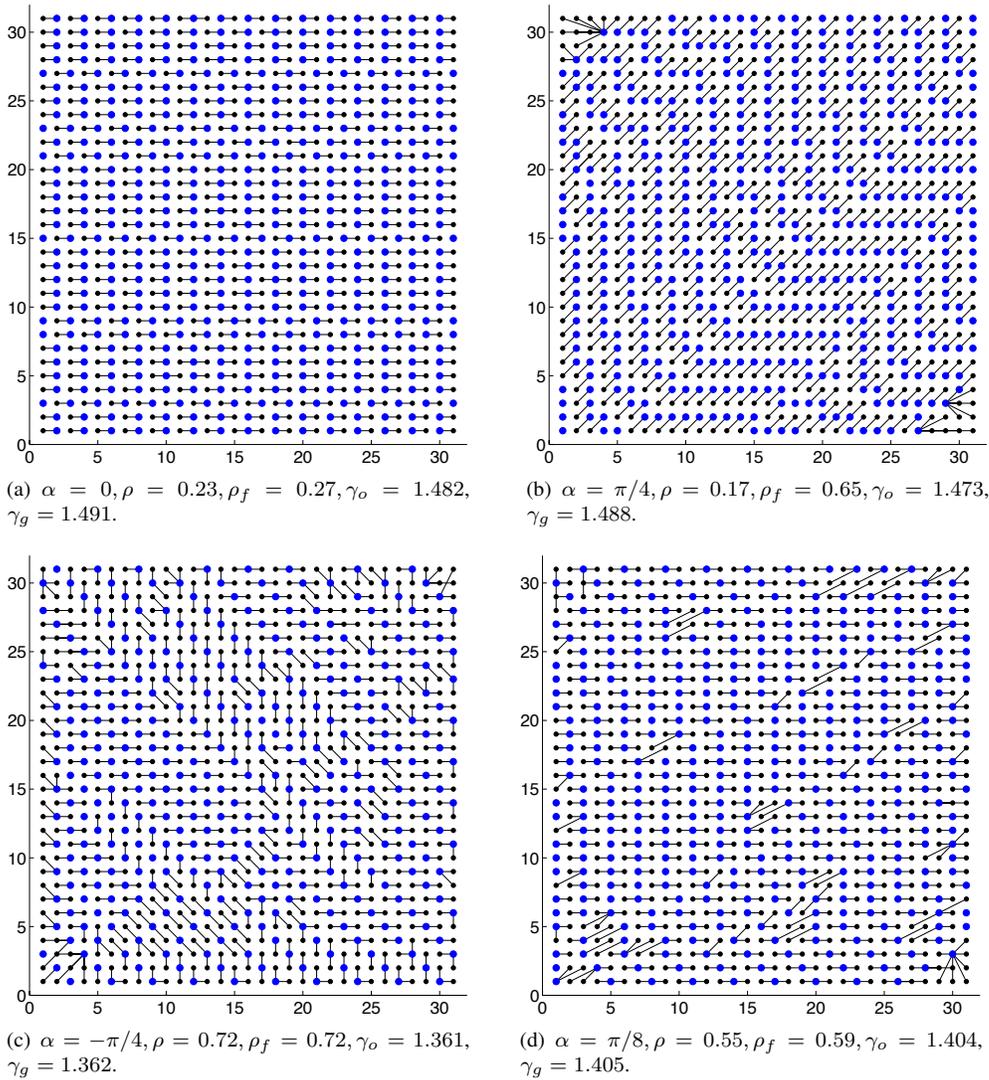


FIG. 4.1. Coarse grids and caliber  $c = 1$  interpolation patterns for the finite difference discretization with  $h = 1/32$  for various choices of  $\alpha$  using the graph of  $A$ , i.e.,  $d = 1$  and  $d_{LS} = 3$ , to define the strength matrix. Here, the smaller circles are  $F$ -points and the larger circles are  $C$ -points.

The results for the more challenging  $\alpha = -\pi/4$  case are provided in Figure 4.9(b). Here, we observe that although the discretization on the fine grid does not follow the anisotropy whatsoever, the nonzero pattern of the coarse-grid operator correctly aligns with the direction of anisotropy. This result demonstrates the ability of the algorithm to overcome, if needed, the disadvantage of a poorly chosen fine-grid discretization and regain a more favorable discretization on the first coarse grid. Further, the results for  $\alpha = \pi/4$  indicate, that all consecutive coarse grids (though not employed in the two-level algorithm) are likely to maintain a similar favorable discretization that accurately reflects the anisotropy, too.

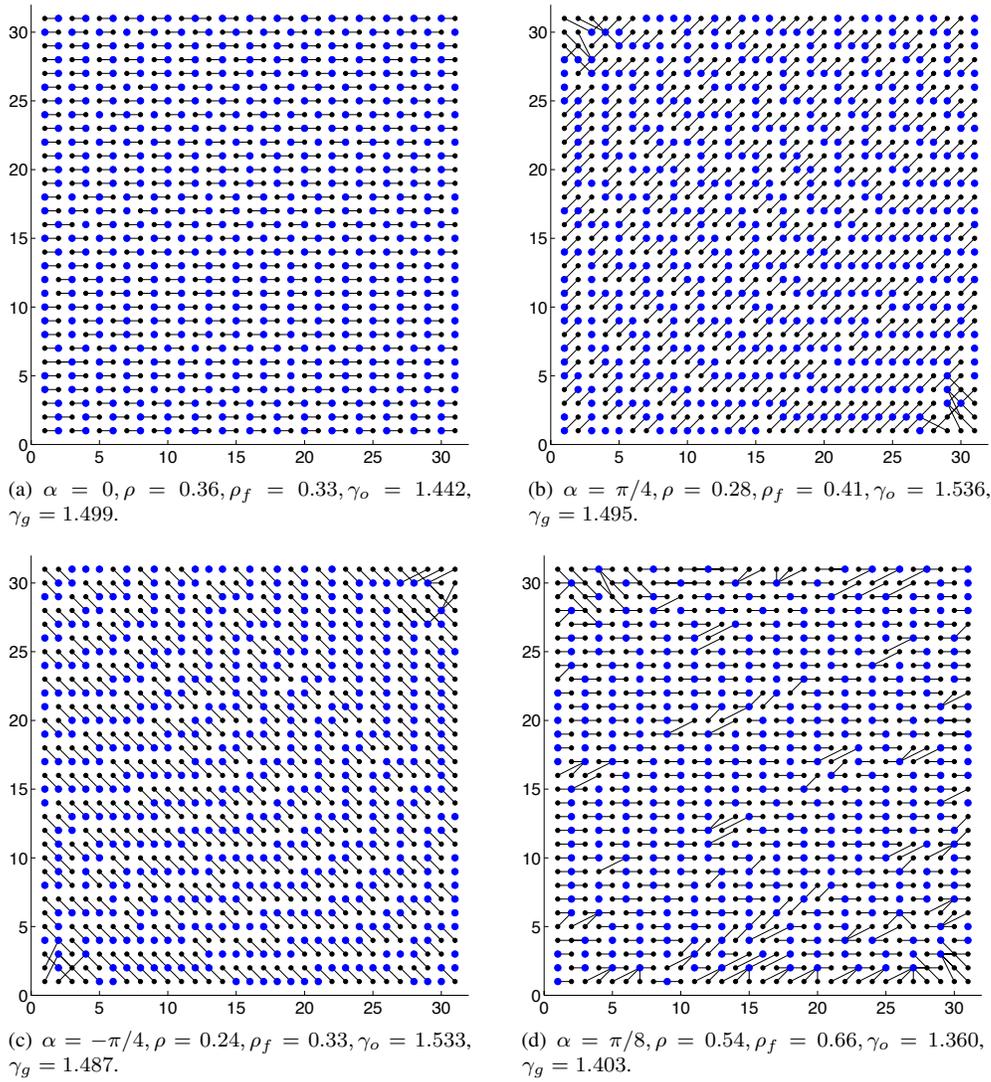


FIG. 4.2. Coarse grids and caliber  $c = 1$  interpolation patterns for the bilinear finite element discretization with  $h = 1/32$  for various choices of  $\alpha$  using the graph of  $A$ , i.e.,  $d = 1$  and  $d_{LS} = 3$ , to define the strength matrix. Here, the smaller circles are  $F$ -points and the larger circles are  $C$ -points.

Coefficients of the coarse-grid stencils presented in Figure 4.9 are given next. Here  $S_{cg}^+$  corresponds to  $\alpha = \pi/4$  (entries denoted by  $*$  are negligible with absolute values below  $10^{-11}$ ):

$$S_{cg}^+ = \begin{bmatrix} & & & * & -0.17 \\ & & & & * \\ & * & 0.33 & & \\ * & -0.17 & & * & \\ & & * & & \end{bmatrix},$$



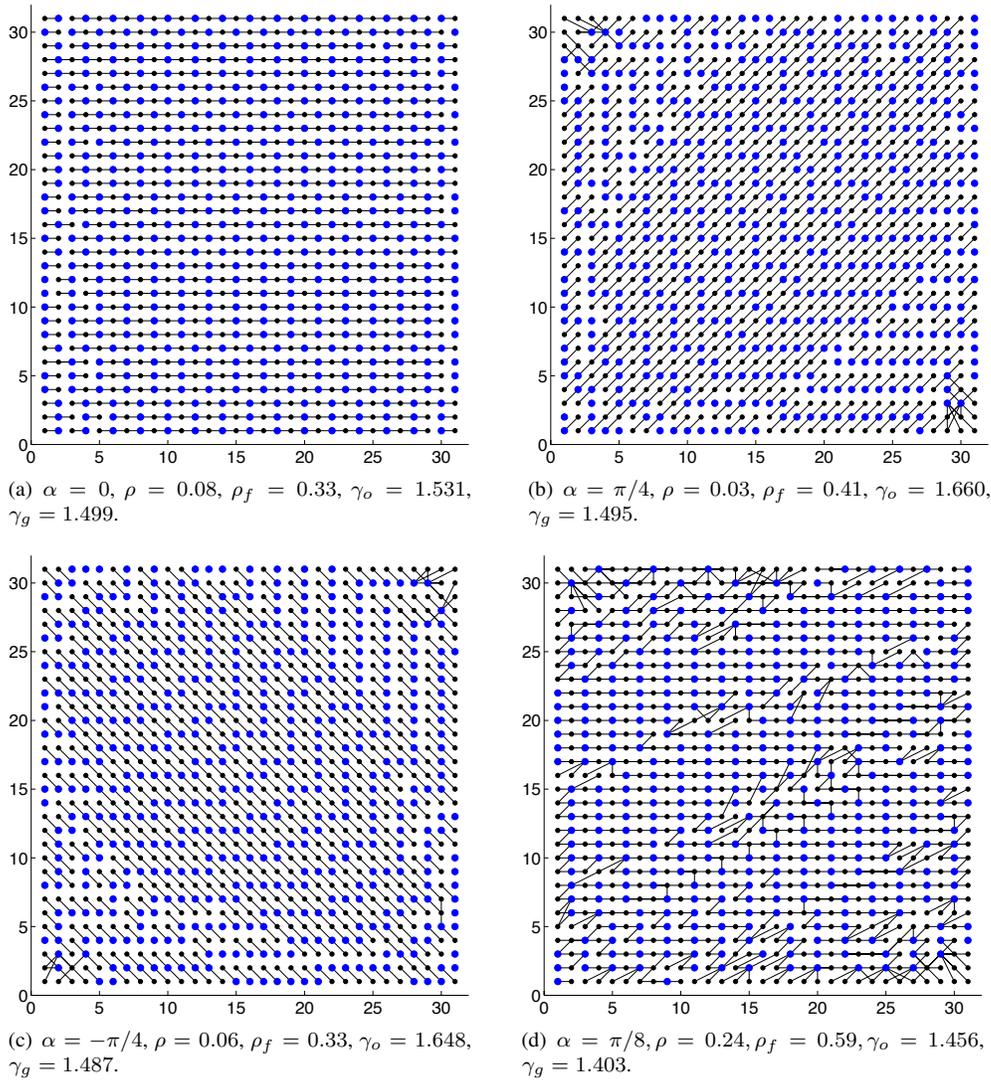


FIG. 4.4. Coarse grids and caliber  $c = 2$  interpolation patterns for the bilinear finite element discretization with  $h = 1/32$  for various choices of  $\alpha$ , using the graph of  $A$ , i.e.,  $d = 1$  and  $d_{LS} = 3$ , to define the strength matrix. Here, the smaller circles are  $F$ -points and the larger circles are  $C$ -points.

The distribution of the stencils' coefficients further illustrates the ability of the algebraic distance based strength of connections measure to choose the correct coarse-grid points for problems with anisotropic coefficients. In both cases, the nonzero sparsity pattern and dominant coefficients of the resulting coarse-grid operators follow the direction of anisotropy. Clearly, larger search depths, both for the CR and the LS procedures, yield an additional fill-in of the resulting coarse-level operator, with this effect becoming more profound in a multilevel setting. In such cases, it becomes especially important to choose accurate coarse grids based on a rich set of test functions. Strategies for computing the latter are discussed in the following sections.

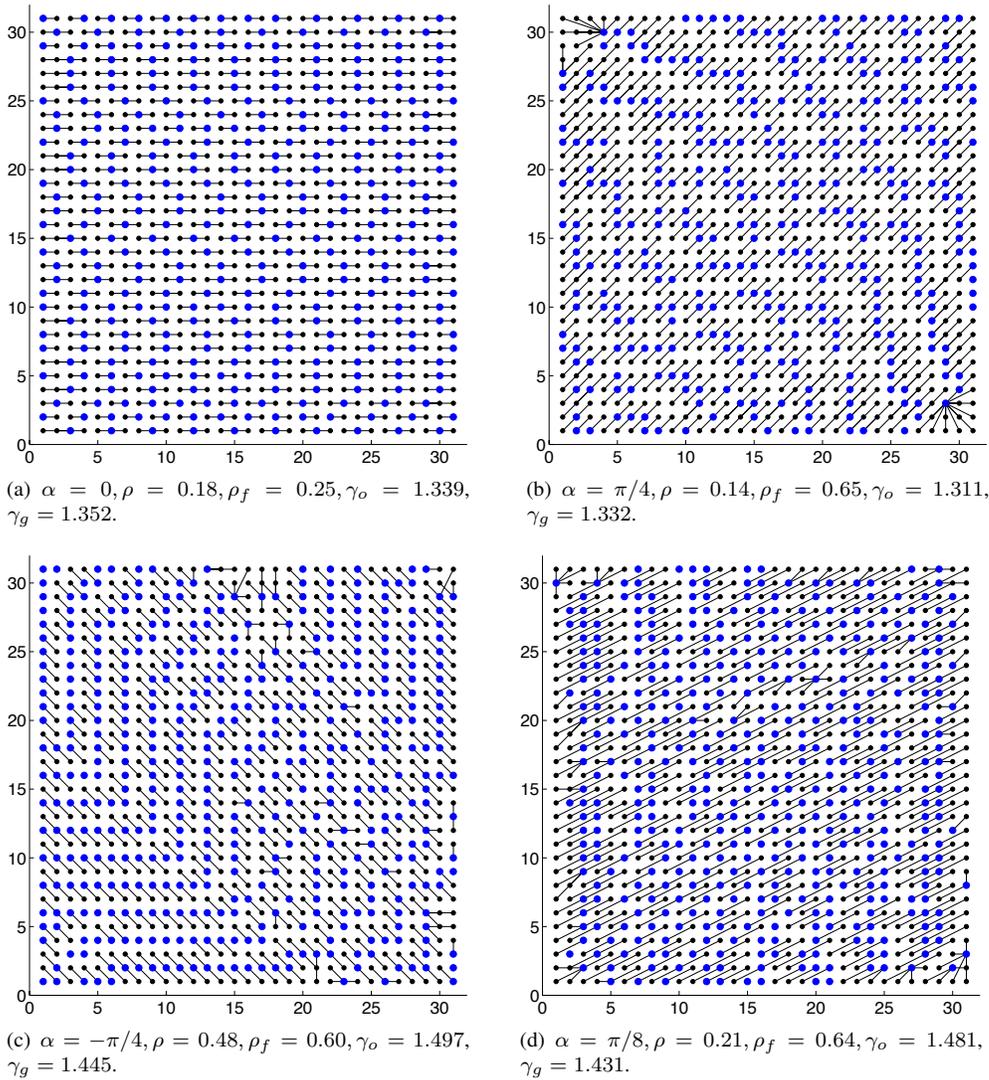


FIG. 4.5. Coarse grids and caliber  $c = 1$  interpolation patterns for the finite difference discretization with  $h = 1/32$  for various choices of  $\alpha$  using the graph of  $A^2$ , i.e.,  $d = 2$ , and  $d_{LS} = 4$ , to define the strength matrix. Here, the smaller circles are  $F$ -points and the larger circles are  $C$ -points.

**4.5. Two-level convergence.** Here, we present experiments with tests of the proposed AMG setup algorithm applied to (4.1) for various choices of the anisotropy angle  $\alpha$ , the anisotropy coefficient  $\epsilon$ , and the mesh spacing  $h = 1/N$ . In Tables 4.1–4.2 we present test results for the bootstrap setup with calibers  $c = 1$  and  $c = 2$  interpolation for the finite difference and finite element discretizations. We note that multilevel results are given in the next section. In the tables, the asymptotic convergence rates  $\rho$  of the two-grid solver produced by the BAMG setup algorithm are reported, along with the corresponding coarsening factors  $\gamma_g$  and operator complexity ratios  $\gamma_o$ . Here, for  $c = 1$ , we observe a strong dependence of the computed convergence rates, grid, and operator complexities on the problem parameters  $\epsilon$ ,  $\alpha$ , and  $h$ . For  $c = 2$ , this dependence is less pronounced with only a slight dependence on  $h$

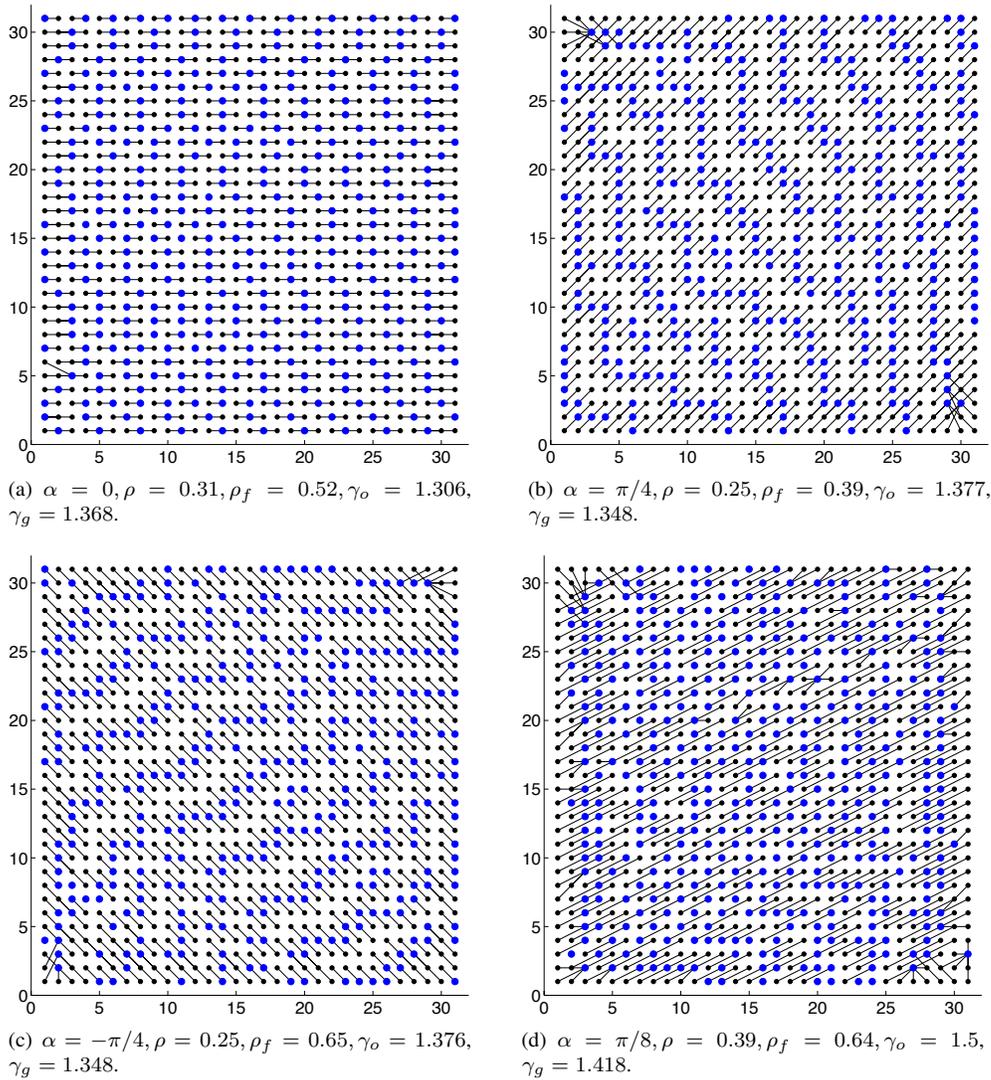


FIG. 4.6. Coarse grids and caliber  $c = 1$  interpolation patterns for the bilinear finite element discretization with  $h = 1/32$  for various choices of  $\alpha$ , using the graph of  $A^2$ , i.e.,  $d = 2$  and  $d_{LS} = 4$ , to define the strength matrix.

that is restricted mostly to the non-grid aligned cases. The exception is  $\alpha = 0$  and  $\epsilon = 0.1$ , where we see a slight increase in  $\rho$  as the problem size is increased from  $N = 64$  to  $N = 128$ . Moreover in all cases, the convergence rates and complexities are uniformly bounded with respect to  $\epsilon$  and  $\alpha$  for fixed  $h$ . These results are promising when considering that all tests were performed with the same strength parameter  $\theta_{ad} = 0.5$ . In fact, all parameters in the setup algorithm were fixed illustrating that the individual components of the BAMG setup are robust for the targeted anisotropic problems even those leading to non- $M$ -matrix systems as in the  $\alpha = -\pi/4$  and  $\alpha = \pi/8$  cases. Further, we note that the bootstrap setup with caliber  $c = 2$  interpolation handles the isotropic case when  $\alpha = 0$  and  $\epsilon = 1$  with similar efficiency, producing a two-grid method with convergence rate  $\rho = 0.28$  and complexities  $\gamma_g = 0.25$  and  $\gamma_o = 1.6$  for  $h = 1/32, 1/64, 1/128$ .

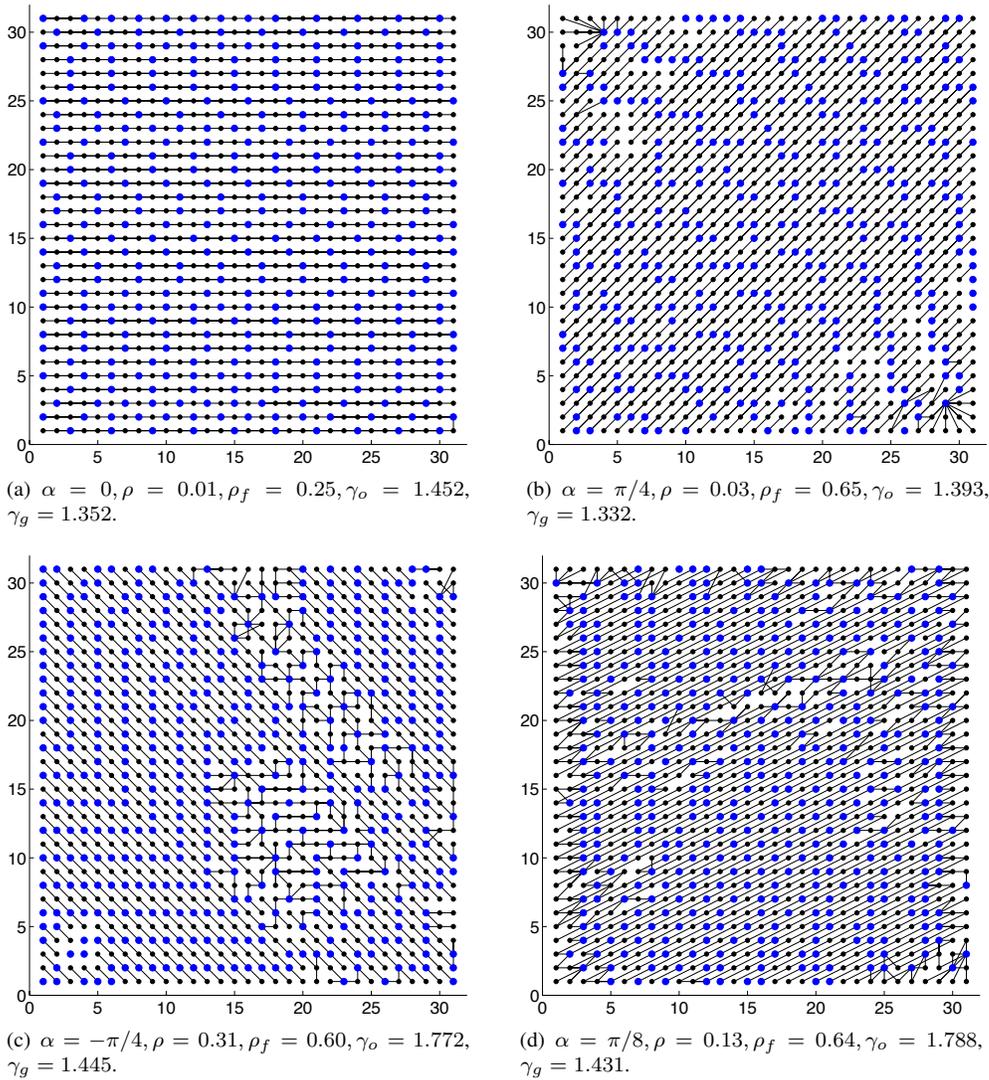


FIG. 4.7. Coarse grids and caliber  $c = 2$  interpolation patterns for the finite difference discretization with  $h = 1/32$  for various choices of  $\alpha$  using the graph of  $A^2$ , i.e.,  $d = 2$  and  $d_{LS} = 4$ , to define the strength matrix.

**4.6. Multilevel convergence.** Next, we consider the performance of a multilevel BAMG algorithm. We report results of a nonlinear Algebraic Multilevel Iteration (AMLI) W-cycle preconditioner constructed by using recursively the proposed bootstrap setup algorithm applied to the same anisotropic test problems discretized using finite differences and bilinear finite elements. Though numerical results for the stand-alone V-cycle solver and preconditioner are not reported, we note that the convergence rates of both approaches deteriorate for increasing problem sizes and strength of anisotropy in the non-aligned cases. This observation in fact motivated our use of the AMLI W-cycle preconditioner. For details of the nonlinear AMLI-cycle we refer to [28]. Here, we limit the numerical tests to caliber  $c = 2$  interpolation since this choice produced the best results in the tests of the two-grid method given the previous section. Figure 4.10 provides a schematic outline of the bootstrap V- and W-cycle setup

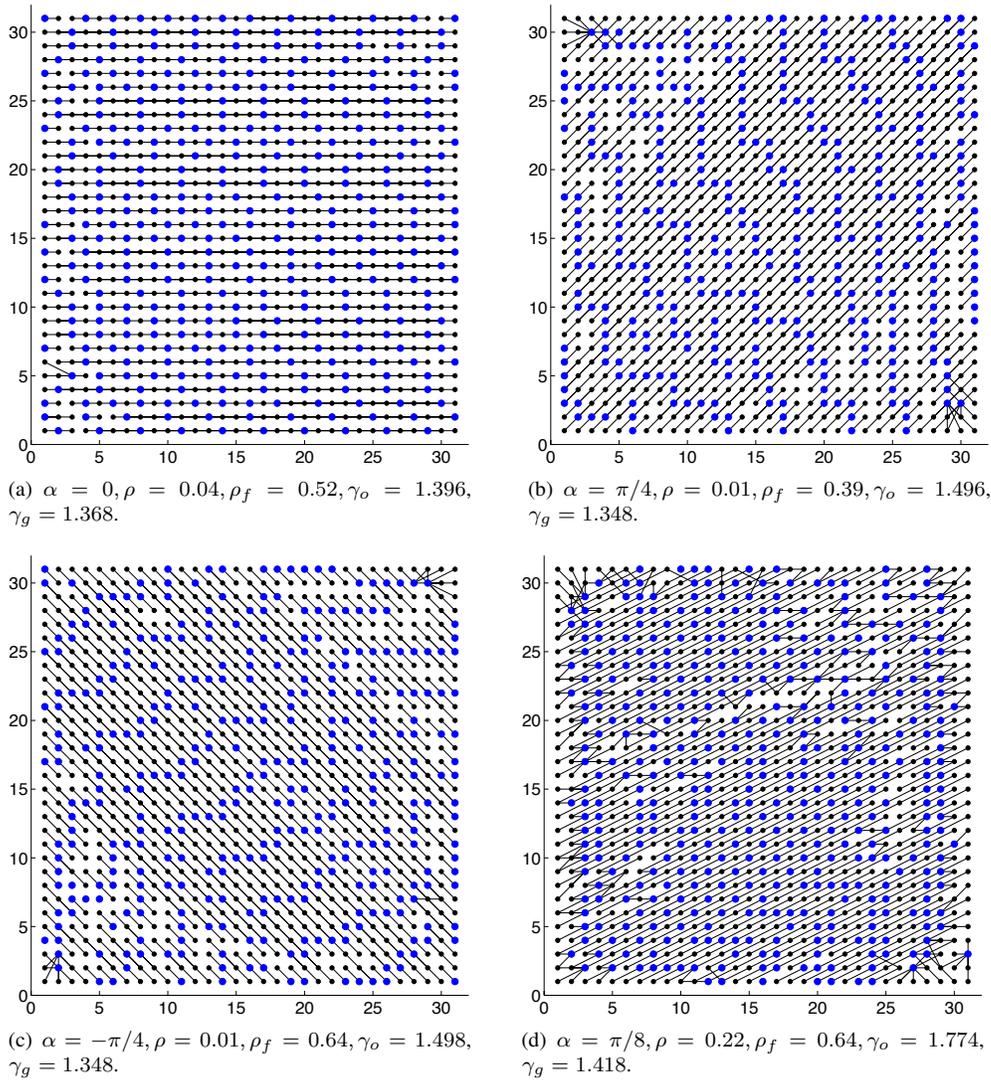


FIG. 4.8. Coarse grids and caliber  $c = 2$  interpolation patterns for the bilinear finite element discretization with  $h = 1/32$  for various choices of  $\alpha$ , using the graph of  $A$ , i.e.,  $d = 2$  and  $d_{LS} = 4$ , to define the strength matrix.

algorithms. A main ingredient of the bootstrap setup is its use of a multilevel generalized eigensolver to compute the bootstrapped test vectors once an initial multigrid hierarchy has been constructed. The goal is to enrich the set of TVs by using approximations of the lowest eigenmodes of the finest-level matrix  $A_0 = A$  obtained by computing eigenvectors on the coarsest level  $L$  and then transferring them, with some additional local processing, to the finest grid. The main ideas of the multilevel eigensolver are as follows.

Given the initial Galerkin operators  $A_0, A_1, \dots, A_L$  on each level and the corresponding interpolation operators  $P_{l+1}^l, l = 0, \dots, L - 1$ , define the composite interpolation operators as  $P_l = P_1^0 \cdot \dots \cdot P_l^{l-1}, l = 1, \dots, L$ . Then, for any given vector  $x_l \in \mathbb{C}^{n_l}$  we have

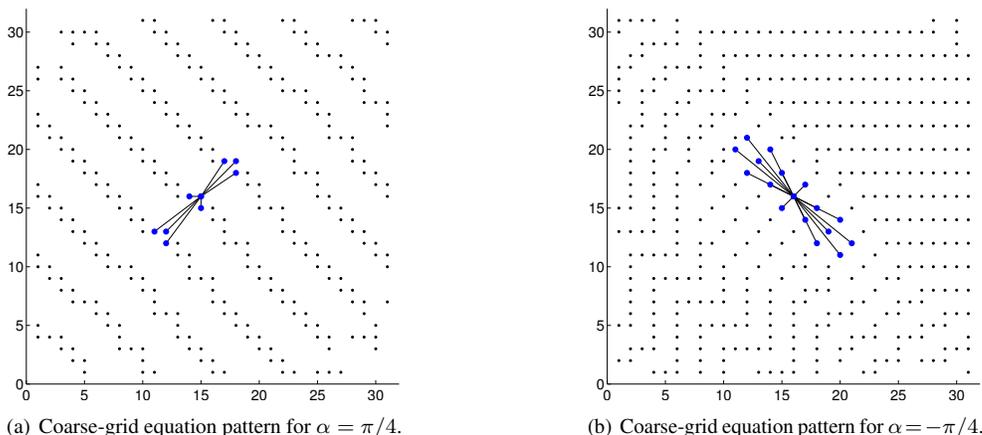


FIG. 4.9. Nonzero pattern of one of the stencils of the coarse-grid operator  $A_c$ , centered at  $i \in C$  and connected with  $j \in C$  such that  $(A_c)_{ij} \neq 0$ , for  $h = 1/32$  and  $\alpha = \pi/4$ . The smaller dots in the graph are all other coarse-grid variables.

TABLE 4.1

Approximate asymptotic convergence rates of the two-grid solver applied to the seven point FD Anisotropic Laplace problem with Dirichlet boundary conditions for various choices of  $\alpha$ ,  $\epsilon$  and  $h$ . Here, the proposed setup algorithm is applied with search depth for the CR coarsening algorithm set as  $d = 2$ , and with the search depth for caliber  $c = 1$  and  $c = 2$  least-squares interpolation set as  $d_{LS} = 4$ . The reported results correspond to convergence rates  $\rho$ , and, in parenthesis, coarsening factors  $\gamma_g$ , and operator complexity factors  $\gamma_0$  computed for  $h = 1/32$ ,  $h = 1/64$ , and  $h = 1/128$ .

$\alpha$	$c = 1$ and $\epsilon = 0.1$			$c = 2$ and $\epsilon = 0.1$		
0	.23(1.3, 1.3)/.36(1.4, 1.4)/.45(1.4, 1.4)	.04(1.3, 1.6)/.13(1.4, 1.5)/.20(1.4, 1.5)				
$\pi/4$	.14(1.4, 1.4)/.37(1.3, 1.4)/.49(1.4, 1.4)	.01(1.4, 1.5)/.04(1.3, 1.5)/.05(1.4, 1.5)				
$-\pi/4$	.36(1.4, 1.4)/.60(1.3, 1.4)/.77(1.3, 1.4)	.07(1.4, 1.6)/.27(1.3, 1.5)/.31(1.3, 1.5)				
$\pi/8$	.14(1.3, 1.3)/.42(1.3, 1.3)/.67(1.3, 1.3)	.01(1.3, 1.4)/.12(1.3, 1.4)/.15(1.3, 1.4)				
$\alpha$	$c = 1$ and $\epsilon = 0.0001$			$c = 2$ and $\epsilon = 0.0001$		
0	.18(1.4, 1.3)/.41(1.4, 1.4)/.61(1.4, 1.4)	.01(1.4, 1.5)/.05(1.4, 1.6)/.05(1.4, 1.5)				
$\pi/4$	.14(1.3, 1.3)/.33(1.4, 1.4)/.49(1.4, 1.4)	.03(1.3, 1.4)/.05(1.4, 1.5)/.06(1.4, 1.6)				
$-\pi/4$	.48(1.5, 1.5)/.69(1.4, 1.5)/.84(1.4, 1.5)	.31(1.5, 1.8)/.38(1.4, 1.7)/.42(1.4, 1.7)				
$\pi/8$	.21(1.4, 1.5)/.68(1.4, 1.5)/.89(1.4, 1.5)	.13(1.4, 1.8)/.35(1.4, 1.7)/.43(1.4, 1.7)				
$\alpha$	$c = 1$ and $\epsilon = 0$			$c = 2$ and $\epsilon = 0$		
0	.18(1.4, 1.3)/.62(1.4, 1.4)/.66(1.4, 1.4)	.01(1.4, 1.3)/.08(1.4, 1.5)/.09(1.4, 1.5)				
$\pi/4$	.10(1.3, 1.3)/.33(1.4, 1.3)/.42(1.4, 1.4)	.04(1.3, 1.4)/.05(1.4, 1.5)/.06(1.4, 1.6)				
$-\pi/4$	.48(1.5, 1.5)/.70(1.4, 1.5)/.87(1.4, 1.6)	.31(1.5, 1.8)/.37(1.4, 1.7)/.41(1.4, 1.8)				
$\pi/8$	.38(1.4, 1.5)/.76(1.4, 1.6)/.91(1.4, 1.7)	.12(1.4, 1.8)/.35(1.4, 1.7)/.40(1.4, 1.8)				

$\langle x_l, x_l \rangle_{A_l} = \langle P_l x_l, P_l x_l \rangle_A$ . Furthermore, defining  $T_l = P_l^H P_l$ , we obtain

$$\frac{\langle x_l, x_l \rangle_{A_l}}{\langle x_l, x_l \rangle_{T_l}} = \frac{\langle P_l x_l, P_l x_l \rangle_A}{\langle P_l x_l, P_l x_l \rangle_2}.$$

This observation in turn implies that on any level  $l$ , given a vector  $v^{(l)} \in \mathbb{C}^{n_l}$  and  $\lambda^{(l)} \in \mathbb{C}$

TABLE 4.2

*Approximate asymptotic convergence rates of the two-grid solver applied to the nine point bilinear FE Anisotropic Laplace problem with Dirichlet boundary conditions for various choices of  $\alpha$ ,  $\epsilon$  and  $h$ . Here, the proposed setup algorithm is applied with search depth for the CR coarsening algorithm set as  $d = 2$ , and with the search depth for caliber  $c = 1$  and  $c = 2$  least-squares interpolation set as  $d_{LS} = 4$ . The reported results correspond to convergence rates  $\rho$ , and, in parenthesis, coarsening factors  $\gamma_g$ , and operator complexity factors  $\gamma_o$  computed for  $n = 32^2$ ,  $h = 1/32$ ,  $h = 1/64$ , and  $h = 1/128$ .*

$\alpha$	$c = 1$ and $\epsilon = 0.1$	$c = 2$ and $\epsilon = 0.1$
0	.15(1.4, 1.3)/.33(1.3, 1.3)/.45(1.3, 1.4)	.05(1.4, 1.4)/.18(1.3, 1.4)/.21(1.3, 1.5)
$\pi/4$	.18(1.4, 1.4)/.43(1.3, 1.4)/.49(1.4, 1.4)	.04(1.4, 1.5)/.09(1.3, 1.4)/.11(1.4, 1.5)
$-\pi/4$	.17(1.3, 1.4)/.42(1.3, 1.3)/.77(1.3, 1.4)	.02(1.3, 1.5)/.19(1.3, 1.4)/.24(1.3, 1.6)
$\pi/8$	.36(1.3, 1.3)/.49(1.3, 1.3)/.67(1.3, 1.3)	.22(1.3, 1.8)/.26(1.3, 1.3)/.33(1.3, 1.5)
$\alpha$	$c = 1$ and $\epsilon = 0.0001$	$c = 2$ and $\epsilon = 0.0001$
0	.31(1.4, 1.3)/.55(1.4, 1.3)/.61(1.4, 1.4)	.04(1.4, 1.4)/.05(1.4, 1.6)/.05(1.4, 1.5)
$\pi/4$	.25(1.4, 1.4)/.46(1.4, 1.4)/.58(1.4, 1.4)	.01(1.4, 1.5)/.19(1.4, 1.5)/.22(1.4, 1.5)
$-\pi/4$	.25(1.4, 1.4)/.52(1.4, 1.4)/.84(1.4, 1.5)	.01(1.4, 1.5)/.20(1.4, 1.6)/.25(1.4, 1.6)
$\pi/8$	.39(1.4, 1.5)/.68(1.4, 1.5)/.89(1.4, 1.5)	.22(1.4, 1.8)/.29(1.4, 1.6)/.36(1.4, 1.6)
$\alpha$	$c = 1$ and $\epsilon = 0$	$c = 2$ and $\epsilon = 0$
0	.32(1.4, 1.3)/.60(1.4, 1.3)/.66(1.4, 1.4)	.05(.4, 1.4)/.10(1.4, 1.4)/.13(1.4, 1.4)
$\pi/4$	.25(1.3, 1.4)/.46(1.4, 1.4)/.72(1.4, 1.4)	.04(1.3, 1.4)/.21(1.4, 1.6)/.23(1.4, 1.7)
$-\pi/4$	1.3(1.4, 1.5)/.51(1.4, 1.4)/.74(1.4, 1.4)	.01(1.4, 1.5)/.20(1.4, 1.6)/.26(1.4, 1.6)
$\pi/8$	.40(1.4, 1.5)/.68(1.4, 1.5)/.79(1.4, 1.5)	.22(1.4, 1.8)/.33(1.4, 1.6)/.45(1.4, 1.6)

such that  $A_l v^{(l)} = \lambda^{(l)} T_l v^{(l)}$ , we have the Raleigh quotient

$$(4.7) \quad \text{rq}(P_l v^{(l)}) := \frac{\langle P_l v^{(l)}, P_l v^{(l)} \rangle_A}{\langle P_l v^{(l)}, P_l v^{(l)} \rangle_2} = \lambda^{(l)}.$$

In this way, the eigenvectors and eigenvalues of the operators in the multigrid hierarchy on all levels are related directly to the eigenvectors and eigenvalues of the finest-grid operator  $A$ . Note that the eigenvalue approximations in (4.7) are continuously updated within the algorithm so that the overall approach resembles an inverse Rayleigh-Quotient iteration found in eigenvalue computations; cf. [29]. For additional details of the algorithm and its implementation we refer to the paper [4].

The cost per iteration of a single  $V(\mu_1, \mu_2)$ -setup cycle with  $\mu_1$  pre- and  $\mu_2$  post-smoothing steps can be roughly estimated in terms of the cost of a single fine-grid relaxation step, i.e., one work unit, which we define as the number of nonzero entries in the fine-level matrix denoted by  $\text{nnz}(A_0)$ . Letting, as before, the operator complexity ratio be the total number of nonzero entries of the matrices on all levels as in (4.5) of the multilevel hierarchy, the cost of a single  $V(\mu_1, \mu_2)$ -cycle is then roughly given by

$$(\mu_1 + \mu_2) \times \gamma_o \times k \quad \text{work units,}$$

where  $k$  denotes the number of test vectors computed in the BAMG setup. For example, the cost of a single  $V(4, 4)$ -cycle, which computes eight relaxed vectors  $v \in \mathcal{V}^r$  and eight eigenvector approximations  $v \in \mathcal{V}^e$ , requires roughly  $64\gamma_o$  work units. In a W-cycle setup one solves recursively each of the coarse problems twice so that if we assume that the coarsening ratio is  $1/3$  (as it is for most of the two-grid tests presented in the previous section), then the overall cost of a W-cycle setup increases by at most a factor of 3, when compared to the cost of the V-cycle setup. We note in addition that the estimate of the cost of the V-BAMG setup as approximated by (4.6) also estimates the cost of the regular adaptive smoothed aggregation setup phase with one candidate vector when we set  $\mu_1 = 1$  and  $\mu_2 = 0$ .

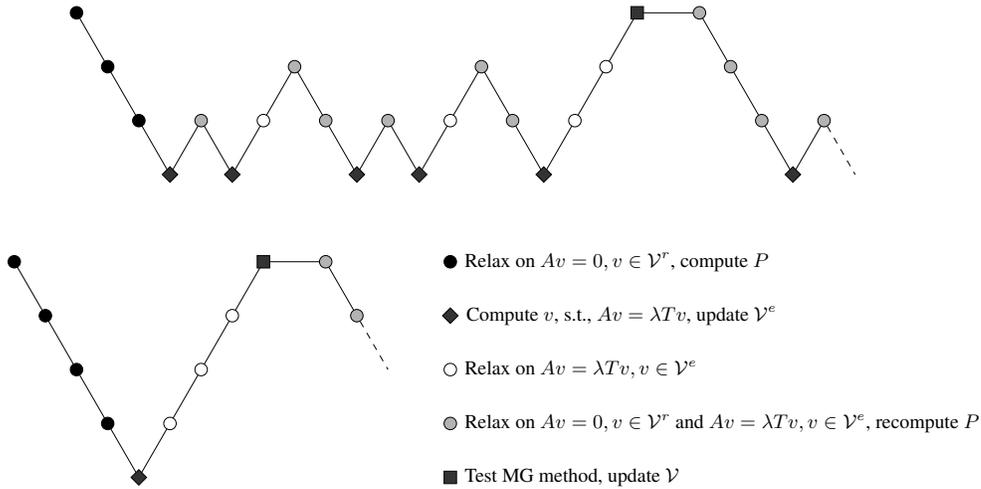


FIG. 4.10. Galerkin Bootstrap AMG W-cycle and V-cycle setup schemes.

The cost of computing the algebraic distance measure for each fine-grid point is roughly given by the number of its neighbors as defined by  $A^d$  times the number of test vectors that is used. Thus, the overall cost is proportional to the problem size  $n$  on the finest level. The cost of computing the LS interpolation operators is computed similarly.

We apply two W-cycle bootstrap cycles using 4 pre- and post-smoothing steps to compute the set of relaxed vectors and set of bootstrap vectors with  $|\mathcal{V}^r| = |\mathcal{V}^e| = 8$ , which together are then used in computing the algebraic distance measure for defining strength of connection and the least-squares interpolation operator on each level. The compatible relaxation algorithm is applied with search depth  $d = 2$ , and the LS interpolation is formed using  $d_{LS} = d + 2$  as before.

In the solve phase, we consider a standard W-cycle as a preconditioner to the conjugate gradient iteration (PCG) and nonlinear AMLI W-cycles as a preconditioner to the flexible conjugate gradient iteration. Here, we report the number of iterations needed by both approaches to reduce the residual by  $10^8$ , whereas our previous two-level results were for the stand alone (unpreconditioned) solver. We note the dependence of the iteration counts of the PCG method on the problem size for the non-grid aligned cases when using standard W-cycles; in all tests, the number of iterations increases by 2–5 iterations as we half the mesh spacing, i.e., increase the problem size by a factor of 4 in our two-dimensional setting. We note in addition that this dependence of the iteration counts on the problem size is nearly eliminated when we use the AMLI W-cycle preconditioner. Overall, since the coarse grids and interpolation stencils follow the direction of the anisotropy on all levels of the multigrid hierarchy, these results suggest that a higher-caliber ( $c > 2$ ) interpolation is needed to obtain scalable multilevel results.

**5. Concluding remarks.** The LS functional gives a flexible and robust tool for measuring AMG strength of connectivity via algebraic distances. It is computed for pairs of points to define a strength graph used to choose coarse points and for sets of points to determine interpolatory sets. The proposed coarsening approach combines algebraic distances, compatible relaxation, and least-squares interpolation. It provides an effective mechanism for the non-grid aligned anisotropic diffusion problems considered. The approach chooses suitable coarse-grid variables and prolongation operators for a wide range of anisotropies, without the need for parameter tuning. Moreover, even when the initial fine-grid discretization is

TABLE 4.3

Nonlinear AMLI W-cycle preconditioned flexible CG/standard W-cycle preconditioned CG applied to the seven-point finite difference anisotropic Laplace problem with Dirichlet boundary conditions for  $\epsilon = 0.1, 0.0001$ ,  $n = 32^2, 64^2, 128^2$ , and various choices of  $\alpha$ . The reported results correspond to the number of iterations needed to reduce the residual by  $10^8$  and the grid and operator complexities for the method constructed using a W-cycle bootstrap setup algorithm again with the same parameters  $d = 2$  and  $d_{LS} = 4$  that were used for the tests of the two-grid method reported in Table 4.1.

$\alpha$	FD and $\epsilon = 0.1$		
0	4 / 5 (1.4,1.8)	4 / 6 (1.5,1.9)	4 / 6 (1.5,1.9)
$\pi/4$	4 / 6 (1.4,1.5)	4 / 6 (1.5,1.6)	4 / 6 (1.5,1.6)
$-\pi/4$	5 / 6 (1.4,1.9)	6 / 9 (1.5,2.0)	7 / 13 (1.5,2.0)
$\pi/8$	4 / 6 (1.3,1.5)	4 / 9 (1.3,1.6)	5 / 11 (1.4,1.7)
Levels	3	4	5

$\alpha$	FE and $\epsilon = 0.1$		
0	5 / 7 (1.4,1.4)	5 / 8 (1.4,1.4)	6 / 8 (1.5,1.4)
$\pi/4$	5 / 6 (1.4,1.5)	5 / 7 (1.5,1.5)	5 / 8 (1.5,1.6)
$-\pi/4$	5 / 7 (1.5,1.6)	5 / 10 (1.5,1.7)	6 / 12 (1.6,1.7)
$\pi/8$	6 / 8 (1.4,1.3)	6 / 12 (1.4,1.4)	7 / 14 (1.4,1.4)
Levels	3	4	5

$\alpha$	FD and $\epsilon = 0.0001$		
0	3 / 4 (1.4,1.7)	4 / 5 (1.4,1.9)	4 / 6 (1.4,2.0)
$\pi/4$	4 / 6 (1.4,1.5)	4 / 6 (1.5,1.6)	4 / 6 (1.5,1.7)
$-\pi/4$	6 / 9 (1.6,2.1)	8 / 13 (1.6,2.2)	9 / 16 (1.6,2.2)
$\pi/8$	6 / 8 (1.3,1.3)	6 / 12 (1.5,1.9)	6 / 17 (1.5,2.0)
Levels	3	4	5

$\alpha$	FE and $\epsilon = 0.0001$		
0	5 / 7 (1.4,1.7)	6 / 7 (1.4,1.7)	6 / 9 (1.4,1.8)
$\pi/4$	7 / 8 (1.4,1.8)	8 / 8 (1.4,1.9)	9 / 9 (1.4,2.0)
$-\pi/4$	6 / 9 (1.5,1.8)	7 / 12 (1.5,1.9)	8 / 15 (1.6,1.9)
$\pi/8$	6 / 8 (1.5,1.9)	6 / 9 (1.5,1.9)	6 / 11 (1.5,2.0)
Levels	3	4	5

unfavorable, i.e., chosen in the direction opposite to the one defined by the anisotropy (as in the  $\alpha = -\pi/4$  case), the method constructs a suitable interpolation operator and, further, produces a coarse-grid operator which better captures the anisotropy directions, correcting the deficiency of the fine-grid operator. Moreover, we have shown that using caliber  $c = 2$  LS interpolation leads to a nearly optimal multilevel method for the targeted constant coefficient anisotropic diffusion problems. As noted earlier, the main challenge faced in fully extending the approach to an optimal multilevel one for variable coefficient anisotropic problems is that of designing an algorithm capable of constructing long-range interpolation with caliber  $c > 2$ , as needed to accurately capture general anisotropies that at the same time maintains low-grid and operator complexities. This is a research topic that we are currently investigating.

## REFERENCES

- [1] A. BRANDT, *Algebraic multigrid theory: the symmetric case*, Appl. Math. Comput., 19 (1986), pp. 23–56.
- [2] ———, *General highly accurate algebraic coarsening*, Electron. Trans. Numer. Anal., 10 (2000), pp. 1–20. <http://etna.mcs.kent.edu/vol.10.2000/pp1-20.dir>
- [3] ———, *Multiscale scientific computation: review 2001*, in Multiscale and Multiresolution Methods, J. Barth, T. Chan, and R. Haimes, eds., Lecture Notes in Computational Science and Engineering, 20, Springer, Berlin, 2002, pp. 3–95.

- [4] A. BRANDT, J. BRANNICK, K. KAHL, AND I. LIVSHITS, *Bootstrap AMG*, SIAM J. Sci. Comput., 33 (2011), pp. 612–632.
- [5] A. BRANDT, S. MCCORMICK, AND J. RUGE, *Algebraic multigrid (AMG) for automatic multigrid solution with application to geodetic computations*, Tech. Report, Colorado State University, Fort Collins, 1983.
- [6] ———, *Algebraic multigrid (AMG) for sparse matrix equations*, in Sparsity and its Applications, D. J. Evans, ed., Cambridge University Press, Cambridge, 1985, pp. 257–284.
- [7] J. BRANNICK, *Adaptive Algebraic Multigrid Coarsening Strategies*, PhD. Thesis, Department of Applied Mathematics, University of Colorado at Boulder, Boulder, 2005
- [8] J. BRANNICK, M. BREZINA, S. MACLACHLAN, T. MANTEUFFEL, S. MCCORMICK, AND J. RUGE, *An energy-based AMG coarsening strategy*, Numer. Linear Algebra Appl., 13 (2006), pp. 133–148.
- [9] J. BRANNICK, Y. CHEN, J. KRAUSS, AND L. ZIKATANOV, *An algebraic multigrid method based on matching of graphs*, in Domain Decomposition Methods in Science and Engineering XX, R. Banks, M. Holst, O. B. Widlund, and J. Xu, eds., Lecture Notes in Computational Science and Engineering, 91, Springer, Heidelberg, 2013, pp. 143–150.
- [10] J. BRANNICK, Y. CHEN, AND L. ZIKATANOV, *An algebraic multilevel method for anisotropic elliptic equations based on subgraph matching*, Numer. Linear Algebra Appl., 19 (2012), pp. 279–295.
- [11] J. BRANNICK AND R. D. FALGOUT, *Compatible relaxation and coarsening in algebraic multigrid*, SIAM J. Sci. Comput., 32 (2010), pp. 1393–1416.
- [12] J. BRANNICK AND L. ZIKATANOV, *Algebraic multigrid methods based on compatible relaxation and energy minimization*, in Domain Decomposition Methods in Science and Engineering XVI, O. Widlund and D. E. Keyes, eds., Lecture Notes in Computational Science and Engineering, 55, 2007, pp. 15–26.
- [13] M. BREZINA, A. J. CLEARY, R. D. FALGOUT, V. E. HENSON, J. E. JONES, T. A. MANTEUFFEL, S. MCCORMICK, AND J. RUGE, *Algebraic multigrid based on element interpolation (AMGe)*, SIAM J. Sci. Comput., 22 (2000), pp. 1570–1592.
- [14] M. BREZINA, R. FALGOUT, S. MACLACHLAN, T. MANTEUFFEL, S. MCCORMICK, AND J. RUGE, *Adaptive amg (aAMG)*, SIAM J. Sci. Comput., 25 (2004), pp. 1896–1920.
- [15] W. L. BRIGGS, V. E. HENSON, AND S. MCCORMICK, *A Multigrid Tutorial*, 2nd ed., SIAM, Philadelphia, 2000.
- [16] T. CHARTIER, R. D. FALGOUT, V. E. HENSON, J. JONES, T. MANTEUFFEL, S. MCCORMICK, J. RUGE, AND P. S. VASSILEVSKI, *Spectral AMGe ( $\rho$ AMGe)*, SIAM J. Sci. Comput., 25 (2003), pp. 1–26.
- [17] A. J. CLEARY, R. D. FALGOUT, V. E. HENSON, AND J. E. JONES, *Coarse-grid selection for parallel algebraic multigrid*, in Solving Irregularly Structured Problems in Parallel, A. Ferreira, J. Rolim, H. Simon, and S.-H. Teng, eds., Lecture Notes in Computer Science, 1457, Springer, Berlin, 1998, pp. 104–115.
- [18] H. DE STERCK, U. M. YANG, AND J. J. HEYS, *Reducing complexity in parallel algebraic multigrid preconditioners*, SIAM J. Matrix Anal. Appl., 27 (2006), pp. 1019–1039.
- [19] R. FALGOUT AND P. VASSILEVSKI, *On generalizing the algebraic multigrid framework*, SIAM J. Numer. Anal., 42 (2004), pp. 1669–1693.
- [20] V. E. HENSON AND U. M. YANG, *BoomerAMG: a parallel algebraic multigrid solver and preconditioner*, Appl. Numer. Math., 41 (2002), pp. 155–177.
- [21] O. E. LIVNE, *Coarsening by compatible relaxation*, Numer. Linear Algebra Appl., 11 (2004), pp. 205–227.
- [22] T. MANTEUFFEL, S. MCCORMICK, M. PARK, AND J. RUGE, *Operator-based interpolation for bootstrap algebraic multigrid*, Numer. Linear Algebra Appl., 17 (2010), pp. 519–537.
- [23] L. N. OLSON, J. SCHRODER, AND R. S. TUMINARO, *A new perspective on strength measures in algebraic multigrid*, Numer. Linear Algebra Appl., 17 (2010), pp. 713–733.
- [24] D. RON, I. SAFRO, AND A. BRANDT, *Relaxation-based coarsening and multiscale graph organization*, Multiscale Model. Simul., 9 (2011), pp. 407–423.
- [25] J. RUGE AND K. STÜBEN, *Algebraic multigrid*, in Multigrid Methods, S. McCormick, ed., Frontiers in Applied Mathematics, 3, SIAM, Philadelphia, 1987, pp. 73–130.
- [26] J. B. SCHRODER, *Smoothed aggregation solvers for anisotropic diffusion*, Numer. Linear Algebra Appl., 19 (2012), pp. 296–312.
- [27] P. VANĚK, J. MANDEL, AND M. BREZINA, *Algebraic multigrid by smoothed aggregation for second and fourth order elliptic problems*, Computing, 56 (1996), pp. 179–196.
- [28] P. VASSILEVSKI, *Multilevel Block Factorization Preconditioners*, Springer, New York, 2008.
- [29] J. H. WILKINSON, *The Algebraic Eigenvalue Problem*, Clarendon Press, Oxford, 1965.