



LAWRENCE
LIVERMORE
NATIONAL
LABORATORY

aAMG based on Weighted Matching for Systems of Elliptic PDEs Arising From Displacement and Mixed Methods

P. Dambra, P. S. Vassilevski

June 26, 2014

The 18th European Conference on Mathematics for Industry
Taormina, Italy
June 9, 2014 through June 13, 2014

Disclaimer

This document was prepared as an account of work sponsored by an agency of the United States government. Neither the United States government nor Lawrence Livermore National Security, LLC, nor any of their employees makes any warranty, expressed or implied, or assumes any legal liability or responsibility for the accuracy, completeness, or usefulness of any information, apparatus, product, or process disclosed, or represents that its use would not infringe privately owned rights. Reference herein to any specific commercial product, process, or service by trade name, trademark, manufacturer, or otherwise does not necessarily constitute or imply its endorsement, recommendation, or favoring by the United States government or Lawrence Livermore National Security, LLC. The views and opinions of authors expressed herein do not necessarily state or reflect those of the United States government or Lawrence Livermore National Security, LLC, and shall not be used for advertising or product endorsement purposes.

α AMG based on Weighted Matching for Systems of Elliptic PDEs arising from Displacement and Mixed Methods

Pasqua D’Ambra and Panayot S. Vassilevski

Abstract Adaptive Algebraic Multigrid (or Multilevel) Methods (α AMG) are introduced to improve robustness and efficiency of classical algebraic multigrid methods in dealing with problems where no a-priori knowledge or assumptions on the near-null kernel of the underlined matrix are available. Recently we proposed an adaptive (bootstrap) AMG method, α AMG, aimed to obtain a composite solver with a desired convergence rate. Each new multigrid component relies on a current (general) *smooth vector* and exploits pairwise aggregation based on weighted matching in a matrix graph to define a new automatic, general-purpose coarsening process, which we refer to as “*the compatible weighted matching*”. In this work, we present results that broaden the applicability of our method to different finite element discretizations of elliptic PDEs. In particular, we consider systems arising from displacement methods in linear elasticity problems and saddle-point systems that appear in the application of the mixed method to Darcy problems.

1 Introduction

Algebraic Multigrid Methods (AMG) were introduced in the mid-1980s as *plug-in* solvers for large and sparse linear systems of equations $A\mathbf{x} = \mathbf{b}$, with the final aim to define automatic coarsening process only depending on the coefficient matrix [4, 15, 16]. These methods are particularly efficient for systems arising from scalar second-order elliptic partial differential equations (PDEs), where a characterization of the *algebraically smooth error*, which is the error component not reduced

P. D’Ambra

ICAR-CNR, Via P. Castellino, 111 80131 Napoli, Italy. e-mail: pasqua.dambra@cnr.it

P. S. Vassilevski

CASC-LLNL, P.O. Box 808, L-561, Livermore, CA 94551, USA. e-mail: panayot@llnl.gov.

The work of the 2nd author was performed under the auspices of the U.S. Department of Energy by Lawrence Livermore National Laboratory under Contract DE-AC52-07NA27344.

by a simple relaxation scheme (such as Gauss-Seidel relaxation), is available [9]. This error, corresponding to the eigenvectors of A with small associated eigenvalues (*near-null kernel of A*), must be nearly exactly represented in the coarse-space in order to be eliminated by the coarse-grid correction process. Therefore, a main focus in the current state-of-the-art AMG methods is to define strategies for building coarse variables and intergrid operators which are able to adapt themselves to the properties of the near-null kernel of the problem at hand in order to preserve efficiency and robustness for dealing with more general classes of problems than the traditional scalar elliptic PDEs, including systems of elliptic PDEs, convection-diffusion equations and also more general non-PDE problems. In this direction, Adaptive Algebraic Multigrids (α AMG) have been proposed [6, 8], where main idea is to use appropriate adaptive steps aimed to “identify” smooth error components which the current solver is not able to efficiently handle so that they can be used to improve the solver by modifying the coarsening scheme without using any specific a priori knowledge about these error components. In [11] we proposed a new α AMG method which relies on a bootstrap strategy aimed to compute a composite solver with a desired convergence rate. We demonstrated its effectiveness when applied to symmetric positive definite (s.p.d.) systems arising from finite element discretization of highly anisotropic scalar elliptic PDEs on structured and unstructured meshes. Here, we extend the application of the method to systems of elliptic PDEs coming from linear elasticity and Darcy flow in porous media in mixed setting. In Section 2 we outline the α AMG based on the compatible weighted matching; in Section 3 we describe the model problems and introduce the Bramble-Pasciak transformation used for extending our method in dealing with symmetric indefinite systems stemming from the mixed finite element discretization of Darcy problems; finally, in Section 4 we present results obtained by a prototype Matlab version of our α AMG solver.

2 Main Features of α AMG based on Compatible Weighted Matching

In [11], we proposed a bootstrap process aimed to build a composite solver of the following form:

$$\mathbf{x}_k = \prod_{r=0}^{2m+1} (I - B_r^{-1}A)\mathbf{x}_{k-1}, \quad k = 1, 2, \dots, \quad (1)$$

with $B_{m+r} = B_{m+1-r}$, $r = 1, \dots, m+1$. Each B_r is an AMG-cycle built with its own aggregation procedure of unknowns driven by a weighted matching for the original matrix graph with weights depending on the most recently computed algebraically smooth vector \mathbf{x}_k with respect to the current composite solver. In more details, starting from a general (random) given vector, we build an initial AMG-cycle represented by the operator B_0 and apply it to the homogeneous system

$A\mathbf{x} = 0$, starting with a nonzero random initial iterate \mathbf{x}_0 , and successively computing $\mathbf{x}_k := (I - B^{-1}A)\mathbf{x}_{k-1}$ for a fixed number of iterations. The iterative process provides an approximation to the eigenvector of $B^{-1}A$ corresponding to the minimal eigenvalue of $B^{-1}A$, i.e., of the algebraically smooth vector corresponding to the current solver. This last vector is then used to build a new AMG-cycle represented by the operator B_1 to be composed as in (1) and tested on the homogeneous system. The bootstrap process is stopped when the process represented by (1) reaches a desired convergence rate.

Each new AMG operator B_r is built by using pairwise aggregation of unknowns driven by weighted matching algorithms for the matrix graph. Such matching algorithms are widely exploited in sparse matrix computations to enhance matrix diagonal dominance [12]. More aggressive coarsening (than pairwise aggregation) can be obtained by combining multiple sweeps of the pairwise aggregation. Our main idea was to exploit to the concept of compatible relaxation introduced in [5] for selecting the coarse-vector space. Since for the coarse space, we choose piecewise constant interpolant (that interpolates exactly the current smooth vector), we choose a complementary space such that on each aggregate (of pair of vertices) it is spanned by a vector orthogonal to the restriction of the smooth vector to that (pairwise) aggregate. To actually choose the aggregates, we use weights based on these orthogonal vectors so that the resulting A_{ff} matrix corresponding to the space complementary to the coarse space have maximal product of its diagonal entries, or maximal sum of its diagonal entries, etc., all of these strategies are implemented e.g., in [12]. For the actual details on the respective algorithms and results on scalar PDEs, we refer to [11]. Here, we investigate the use of more accurate interpolation operators obtained by weighted-Jacobi smoothing of the piecewise constant interpolation operators coupled with aggressive coarsening. This leads to smoothed aggregation type adaptive AMG method [7], which exhibits improved convergence and scalability properties with general reduction of setup costs.

Our coarsening process, which we referred to as *compatible weighted matching*, has the advantage to be independent of user-defined parameters; furthermore, it overcomes the limitations of the characterization of strength of connectivity between pairs of unknowns, well motivated only for algebraic systems with M-matrices. The latter concept is generally used in both the coarse space selection and in the interpolation scheme for classical AMG schemes. We stress that computing optimal matching has a super-linear computational complexity, whereas we are interested in (optimal) AMG with linear complexity, that is why we apply an approximate algorithm to find sub-optimal weighted matchings in a graph [13]; this approach was demonstrated to be effective in computing suitable compatible weighted matchings in the difficult case of highly non-grid aligned anisotropic scalar elliptic PDEs.

3 Case Studies: Linear Elasticity and Darcy Problems

We focus on two types of elliptic PDEs particularly relevant for many engineering applications, such as Lamé equations for linear elasticity and Darcy equations for flow in porous media in mixed system setting. Of main interest is to demonstrate the feasibility of our method on general s.p.d linear systems, where the coefficient matrix is not an M-matrix, as well as, on some symmetric but indefinite systems of saddle-point form.

The most widely used mathematical model for studying deformation of materials due to the application of external forces are the following Lamé equations, which are equilibrium equations written in terms of the displacement field \mathbf{u} :

$$\mu \Delta \mathbf{u} + (\lambda + \mu) \operatorname{grad}(\operatorname{div} \mathbf{u}) = \mathbf{f} \quad \mathbf{x} \in \Omega \quad (2)$$

where $\mathbf{u} = \mathbf{u}(\mathbf{x})$ is the displacement vector, Ω is the 3D spatial domain, and λ and μ are the Lamé constants. A mix of Dirichlet boundary conditions and so-called traction conditions are usually applied to have a unique solution. Discretization of (2) by finite element method, if each scalar component of the displacement vector $\mathbf{u} = (u, v, w)$ is considered separately (*unknown-based* [14] discretization), leads to s.p.d. systems of equations whose coefficient matrix can be written in the following block form:

$$A = \begin{bmatrix} A_{uu} & A_{uv} & A_{uw} \\ A_{vu} & A_{vv} & A_{vw} \\ A_{wu} & A_{wv} & A_{ww} \end{bmatrix}$$

We note that if $\nu \gg \lambda$, the above matrix is spectrally equivalent to its block diagonal, corresponding to the matrix coming from discretization of Laplace equation per each unknown component. In this case, block-wise version of the classical AMG are efficient solver. In general, A is not strongly block-diagonally dominant and problem-dependent multigrid operators have to be considered to improve convergence of AMG [1]. In the present work we demonstrate that our α AMG is able to obtain a solver with a desired convergence rate for general elasticity problems, without any a-priori information on the problem neither on the discretization scheme.

The second type of systems of PDEs we considered in this work comes from the Darcy problem of flows in porous media. It is a boundary value problem associated to the following second order elliptic equation:

$$-\operatorname{div} k(\mathbf{x}) \operatorname{grad} p = f(\mathbf{x}) \quad \mathbf{x} \in \Omega, \quad (3)$$

where $p = p(\mathbf{x})$ is the flow pressure, Ω is the spatial domain, and $k(\mathbf{x})$ is the permeability coefficient. In a mixed finite-element formulation, the flow velocity field $\mathbf{u} = -k \nabla p$ is introduced and equation (3) becomes $\operatorname{div} \mathbf{u} = f$. The resulting problem is a system of two first order equations which can be discretized by using pair of finite element spaces leading to the following indefinite system of saddle-point form:

$$\begin{bmatrix} A & B^T \\ B & 0 \end{bmatrix} \begin{bmatrix} \mathbf{u} \\ p \end{bmatrix} = \begin{bmatrix} \mathbf{f} \\ f \end{bmatrix},$$

where A is an s.p.d. matrix. Such linear systems, especially for highly variable or discontinuous permeability coefficient, are very challenging for general iterative solvers, and more specifically for algebraic multigrid (see [2, 17]). Here we propose to use an approach based on the Bramble-Pasciak preconditioner [3] which transforms the saddle-point matrix into a s.p.d. matrix. They utilize a preconditioner matrix M for the A block, such that $A - M$ is s.p.d., and transform the saddle-point matrix into the following s.p.d. one:

$$\widehat{\mathcal{A}} = \begin{bmatrix} AM^{-1} - I & 0 \\ BM^{-1} & -I \end{bmatrix} \begin{bmatrix} A & B^T \\ B & 0 \end{bmatrix} = \begin{bmatrix} AM^{-1}A - A & (AM^{-1} - I)B^T \\ B(M^{-1}A - I) & BM^{-1}B^T \end{bmatrix}. \quad (4)$$

A good choice in practice is a diagonal matrix M assembled from the local element-based diagonal matrices $diag(M_{fem})$, where $M_{fem} = 1/2\lambda_{min}D_{fem}$ and $D_{fem} = diag(A_{fem})$. Here, A_{fem} is the local element mass matrix for each finite element and λ_{min} is the minimal eigenvalue of the generalized local eigenvalue problem $A_{fem}\mathbf{q} = \lambda D_{fem}\mathbf{q}$. In this case the transformed matrix (4) can be explicitly computed at a cost of a moderate increase in the total number of nonzero elements. In the following Section we report some numerical results related to the application of our adaptive AMG on Darcy problems discretized by the mixed finite-element method, in the above transformed s.p.d. form.

4 Numerical Results

In this Section we report some preliminary results which illustrate the ability of our method to solve the systems of equations introduced in Section 3 both in 2D and in 3D domains. We investigate the convergence behavior and the setup cost for increasing mesh size of the discretization. The setup cost is measured in terms of AMG components $nstages$ built by the bootstrap process to reach a desired convergence rate set to 0.7. The obtained convergence rate ρ was estimated by applying the solver in (1) for 15 iterations at each new built. We also report, per each test case and per each mesh with n nodes, the average number of levels $nlev$ of all solver components and the average of their operator complexity $cmpx$, which gives information on the cost of the application of one cycle; $cmpx$ is defined as the ratio between the sum of nonzero entries of the matrices of all levels and the number of nonzero entries of the fine-grid matrix. Each AMG component, built on the base of the *compatible weighted matching* coarsening method, is a general ν -fold cycle [17], where 1-sweep is alternated with 3 sweeps in the next level; In this way, we ensure linear cost per cycle since our coarsening is based on pairwise aggregation. Symmetric Gauss-Seidel relaxation (1 iteration) is employed as pre/post smoothing while direct solver (based on LU factorization) is used at the coarsest level. In order to achieve aggressive coarsening we combine four steps of pairwise aggregation based on com-

patible matching, which allows us to define coarse matrices with a coarsening ratio of at most 16 at each level; the process is stopped when the size of the coarsest matrix is at most 100.

As test case for linear elasticity, we consider equation (2) on a beam characterized by $\mu = 0.42$ and $\lambda = 1.7$; one side of the beam is considered fixed and the opposite end is pushed downward. The problem is discretized using linear finite elements on triangular (2D) and tetrahedral meshes (3D) on different mesh sizes, obtained by uniform refinement, with the software package MFEM (<http://mfem.googlecode.com>). In Table 1, we summarize our results obtained both in the case of constant piecewise interpolation, i.e., unsmoothed aggregation, (on the left) and with smoothed aggregation (on the right). We observe that our method

<i>Composite αAMG Setup</i>					
	n	$nstages$	ρ	$nlev$	$cmpx$
<i>beam 2D</i>	4386	8	0.61	3	1.12
	16962	10	0.69	3	1.12
	66690	12	0.69	4	1.12
<i>beam 3D</i>	2475	9	0.61	3	1.20
	15795	11	0.67	3	1.20
	111843	16	0.67	4	1.23

<i>Composite αAMG Setup</i>					
	n	$nstages$	ρ	$nlev$	$cmpx$
<i>beam 2D</i>	4386	5	0.69	3	1.25
	16962	7	0.63	3	1.20
	66690	9	0.68	4	1.23
<i>beam 3D</i>	2475	8	0.53	3	1.53
	15795	10	0.60	3	1.78
	111843	12	0.64	4	2.40

Table 1 Linear Elasticity Problems: Setup cost when unsmoothed (on the left) and smoothed aggregation (on the right) are used.

is able to achieve convergence factors less than the desired one for all the cases, although no a-priori information on the spectral properties of the matrices neither on the particular features of the system of PDEs and of its discretization were used. We notice that the number of the necessary bootstrap steps generally increases with increasing the mesh size, especially for 3D problems; the largest size mesh requires 5 more bootstrap steps with respect to the medium size mesh. The total number of bootstrap steps, as expected, is reduced if the smoothed aggregation is applied; furthermore, smoothed aggregation coupled with our aggressive coarsening based on a combination of more steps of pairwise aggregation produces a moderate increase in the operator complexity, leading to a general reduction both in the setup and the application cost of the method.

For the Darcy problems, we consider saddle-point systems stemming from a realistic problem with highly variable permeability coefficients, describing a 3D petroleum reservoir obtained from the 10th Society of Petroleum Engineers (SPE) Comparative Solution Project [10]. We present results for Dirichlet problems (i.e. pressure given on the boundary) discretized by using MFEM with structured hexahedral meshes. For discretization, we used first-order Raviart-Thomas spaces [17] for velocity and piecewise-constant functions for pressure. We apply the Bramble-Pasciak transformation described in Section 3 to obtain the corresponding s.p.d. matrix (4). We observe that for the considered test case and the employed mesh sizes, the number of nonzeros in the transformed matrix has an increase of about

80% with respect to the original saddle-point matrix. In Table 2 we report results for different mesh sizes (note that here n is the size of the saddle-point matrix) for both unsmoothed and smoothed aggregation, when the algorithmic choices were the same as in the elasticity problems.

<i>Composite αAMG Setup</i>					
	n	$nstages$	ρ	$nlev$	$cmpx$
<i>SPE10</i>	1403	2	0.50	2	1.07
	10652	3	0.68	3	1.12
	33645	5	0.65	4	1.13
	88800	7	0.65	4	1.14

<i>Composite αAMG Setup</i>					
	n	$nstages$	ρ	$nlev$	$cmpx$
<i>SPE10</i>	1403	2	0.57	2	1.12
	10652	3	0.69	3	1.34
	33645	5	0.66	4	1.46
	88800	6	0.69	4	1.54

Table 2 Darcy Problems: Setup cost when unsmoothed (on the left) and smoothed aggregation (on the right) are used.

We observe that the adaptive solver is able to obtain the required convergence rate with a moderate number of setup steps, demonstrating the potential of the coupling between Bramble-Pasciak transformation and the adaptive solver to handle well indefinite systems of saddle-point type coming from realistic flow problems. The increase in the number of bootstrap steps needed to obtain the desired convergence rate for increasing mesh size is moderate, showing good scalability properties also in the case of unsmoothed aggregation. We also observe that in this case the impact of smoothed aggregation based on a simple Jacobi smoother on the convergence behavior and scalability is not as significant.

References

1. Baker, A. H., Kolev, Tz. V., Yang, U. M.: Improving Algebraic Multigrid Interpolation Operators for Linear Elasticity. *Numerical Linear Algebra with Applications*. vol. 17, pp. 495–517 (2010).
2. Benzi, M., Golub G. H., Liesen J.: Numerical Solution of Saddle Point Problems. *Acta Numerica*. vol. 14, pp. 1–137 (2005).
3. Bramble, J. H., Pasciak, J. E.: A Preconditioning Technique for Indefinite Systems Resulting from Mixed Approximations of Elliptic Problems. *Mathematics of Computation*. vol. 50, pp. 1–17 (1988).
4. Brandt A., McCormick S.F., Ruge, J.W.: Algebraic multigrid (AMG) for Sparse Matrix Equations. In: D.J. Evans (ed.), *Sparsity and Its Applications*, pp. 257–284. Cambridge University Press, Cambridge, (1984).
5. Brandt, A.: General Highly Accurate Algebraic Coarsening. *Electronic Transactions on Numerical Analysis*, Vol. 10, pp. 1–20 (2000).
6. Brandt, A., Brannick, J., Kahl, K., Livshitz, I., Bootstrap AMG. *SIAM Journal on Scientific Computing*, vol. 33, pp. 612–632 (2011).
7. Brezina, M., Falgout, R. D., MacLachlan, S., Manteuffel, T., McCormick, S., Ruge, J.: Adaptive Smoothed Aggregation α SA Multigrid. *SIAM Review*, vol. 47, pp. 317–346 (2005).
8. Brezina, M., Falgout, R. D., MacLachlan, S., Manteuffel, T., McCormick, S. F., Ruge J.: Adaptive Algebraic Multigrid. *SIAM Journal on Scientific Computing*, vol. 27, pp. 1261–1286 (2006).

9. Briggs W. L., Henson V. E., McCormick S. F.: A Multigrid Tutorial, 2nd Ed. SIAM, Philadelphia (2000).
10. Christie, M.A., Blunt, M. J.: Tenth SPE Comparative Solution Project: a Comparison of Upscaling Techniques. SPE Reservoir Eng. Evaluat., vol. 4, pp. 308–317 (2001).
11. D'Ambra, P., Vassilevski, P.S.: Adaptive AMG with Coarsening based on Compatible Weighted Matching. LLNL-TR-613612, to appear in Computing and Visualization in Science.
12. Duff, I. S., Koster, J.: On Algorithms for Permuting Large Entries to the Diagonal of a Sparse Matrix. SIAM Journal on Matrix Analysis and Applications, vol. 22, pp. 973–996 (2001).
13. Preis, R.: Linear Time $1/2$ -Approximation Algorithm for Maximum Weighted Matching in General Graphs. In STACS'99, LNCS, Springer-Verlag, vol. 1563, pp. 259–269 (1999).
14. Ruge, J.W.: AMG for Problems of Elasticity. Applied Mathematics and Computation, Elsevier, New York (1986).
15. Ruge, J.W., Stüben, K.: An Introduction to Algebraic Multigrid (AMG). In: S. F. McCormick (ed.), *Multigrid Methods*, vol. 3, pp. 73–130. Frontiers in Applied Mathematics, SIAM, Philadelphia (1987).
16. Stüben K.: A Review of Algebraic Multigrid. Journal of Computational and Applied Mathematics, vol. 128, pp. 281–309 (2001).
17. Vassilevski, P. S.: Multilevel Block Factorization Preconditioners, Matrix-based Analysis and Algorithms for Solving Finite Element Equations. Springer, New York, (2008).