SIAM J. SCI. COMPUT.
Vol. 39, No. 5, pp. S88–S111

ALGEBRAIC MULTIGRID FOR STOKES EQUATIONS*

YVAN NOTAY†

Abstract. A method is investigated for solving stationary or time-dependent discrete Stokes equations. It uses one of the standard flavors of algebraic multigrid for coupled partial differential equations, which, however, is not applied directly to the linear system stemming from discretization, but to an equivalent system obtained with a simple algebraic transformation (which may be seen as a form of preconditioning in the literal sense). A two-grid analysis is provided, showing that the eigenvalues of the preconditioned matrix are within a region of the complex plane that is both bounded and away from the origin, independently of the mesh or grid size, as well as of other main problem parameters. On the other hand, whereas the approach can in principle be combined with any type of algebraic multigrid scheme, an investigation of the properties of the coarse grid matrices reveals that plain aggregation has to be preferred to maintain nice two-grid convergence at coarser levels. Eventually, numerical experiments are reported showing that the resulting method is both robust and cost effective, being significantly faster than a state-of-the-art competitor which combines MINRES with optimal block diagonal preconditioning.

Key words. algebraic multigrid, AMG, Stokes problem, saddle point, multigrid, convergence analysis, linear systems, preconditioning

AMS subject classifications. 65F08, 65F10, 65F50, 65N22

DOI. 10.1137/16M1071419

1. Introduction.

We consider the iterative solution of the stationary or time-dependent discrete Stokes equations (see next section for details and references). When they are applicable, *monolithic* or *all-at-once* multigrid methods are often efficient [39]. Because standard smoothing methods (such as Gauss–Seidel or damped Jacobi) are either not well defined or not convergent, relevant approaches are characterized by the use of a specific smoother [35]. This includes *collective* or *coupled* smoothing, in which the primary unknowns, pressure and the velocities in a grid cell, are updated simultaneously [40]. Another approach is called *distributive* smoothing, in which one first transforms the discrete system in such a way that standard decoupled smoothing (such as Gauss–Seidel) performs well on the transformed system [8, 47].

These approaches have been developed for *geometric* multigrid methods, and rarely considered in combination with *algebraic* multigrid (AMG) schemes (see, however, [20, 42, 43]). The development of AMG variants is in fact more difficult than for scalar partial differential equations (PDEs), and, so far, the most promising approaches [32, 44] do not apply AMG schemes directly to the discrete system, but to an equivalent system obtained through a simple algebraic transformation. This latter may be seen as a form of preconditioning in the literal sense. Roughly speaking, considering a partitioning of the system matrix that segregates the different types of discrete unknowns (velocity components and pressure), the transformation aims at increasing the weight of the diagonal blocks, so as to facilitate the use of *unknown-based*

*Received by the editors April 19, 2016; accepted for publication (in revised form) December 14, 2016; published electronically October 26, 2017.
http://www.siam.org/journals/sisc/39-5/M107141.html
Funding: This work was supported by the Energy oriented Centre of Excellence (EoCoE), grant agreement 676629, funded within the Horizon2020 framework of the European Union.
†Université Libre de Bruxelles, Service de Métrologie Nucléaire (C.P. 165-84), 50 Av. F.D. Roosevelt, B-1050 Brussels, Belgium (ynotay@ulb.ac.be, homepages.ulb.ac.be/~ynotay). This author is Research Director of the Fonds de la Recherche Scientifique – FNRS.

AMG schemes [12, 36], in which the prolongation operator is set up by considering separately the different types of unknowns (see below for details). The approach is similar to distributive smoothing in that it uses a transformed matrix. However, with distributive smoothing, the transformation is only used to obtain an efficient smoother, and the coarse grid correction is still based on the original matrix. Opposite to this, with the methods in [32, 44], the whole multigrid scheme is applied to the transformed system; differences and similarities with distributive smoothing are further commented on in section 3.1 below.

This AMG approach for the Stokes problem has been initiated by Webster in [44], where extensive numerical results are reported that highlight the potentialities of the method by showing some nice two-grid convergence rates at the fine grid level. However, the convergence seems to deteriorate at coarser levels, revealing a possible limitation, referred to as "loss of stability"; see also [45].

In the present paper, we aim at filling the many gaps left. We first develop a two-grid analysis, proving eigenvalue bounds for the simplest case of one single damped Jacobi smoothing step. Only mild assumptions are needed, and the analysis is compatible with virtually any sensible AMG method combined with the unknown-based coarsening approach.

However, the analysis holds only at the fine grid level, where the algebraic transformation is applied. As usual with AMG methods, coarse level matrices are obtained with Galerkin projection applied to the fine grid matrix—in the present case, the transformed matrix. We investigate the properties of these matrices to account for the stability issues observed in [44]. We develop a heuristic reasoning that relates the phenomenon to the type of coarsening, that is, to the type of AMG method that is used. Schemes that reproduce standard $h - 2h$ coarsening with bilinear interpolation lead to coarse level matrices with increasingly weaker diagonal blocks, explaining the observed deterioration of the convergence. On the other hand, nothing particular happens when the coarsening is based on plain aggregation, as also mentioned in [44].

These results complement our previous study in [32], where we propose and theoretically analyze a slightly different approach, that uses a transformation similar in spirit but more complex (*two-sided* instead of *one-sided*; see below for details). The stability issue mentioned in [44] is not examined, but a reported numerical experiment suggests that the method can be cost effective with many levels, at least when using (plain) aggregation-based AMG (along the lines of [22, 27, 29]).

Finally, we address the practical effectiveness of both variants from [44] and [32]. To this aim, we present the results of numerical experiments that include two-dimensional (2D) and three-dimensional (3D) problems, either stationary or time dependent, on both structured and unstructured grids. Because of the better properties of the associated coarse level matrices, plain aggregation is preferred, and we more precisely use the method from [29] with just the slight modifications needed to implement the unknown-based coarsening. It turns out that (near) grid independent convergence is observed in all cases, and that the number of iterations needed with both types of transformation are comparable. This, in fact, gives advantage to the simpler (one-sided) approach we focus on here, because the transformed matrices are then sparser and hence each iteration requires less computing time.

For comparison purposes, we also tested a state-of-the-art method combining MINRES with block diagonal preconditioning [16] based, in the time-dependent cases, on the Cahouet–Chabard technique [10]; here, AMG is also used, but in a more standard fashion, for the needed (approximate) inversion of certain matrix blocks corresponding to scalar PDEs. It turns out that the monolithic AMG methods discussed

in this work are significantly faster, the best variant outperforming block diagonal preconditioning by a factor of about two in most cases.

The remainder of this paper is organized as follows. In section 2, we present the class of Stokes problems considered in this work, as well as some of the test problems used in numerical experiments. In section 3, we present the transformation and develop the two-grid analysis. Coarse level matrices are discussed in section 4, the selected multigrid strategy is presented in section 5, large scale experiments are reported in section 6, and, eventually, conclusions are drawn in section 7.

2. Stokes equations and their discretization. We consider the following problem: find the velocity vector $\mathbf{u}$ and the pressure field p satisfying

$$\begin{aligned} \xi\mathbf{u} - \nu\Delta\mathbf{u} + \nabla p &= \mathbf{f} \quad \text{in } \Omega\,, \\ \nabla\cdot\mathbf{u} &= 0 \quad \text{in } \Omega\,, \end{aligned} \tag{2.1}$$

and appropriate boundary conditions on $\partial\Omega$. In (2.1), Ω is a bounded domain of $\mathbb{R}^2$ or $\mathbb{R}^3$, $\mathbf{f}$ represents a prescribed force, and the parameters $\nu > 0$ (viscosity) and $\xi \geq 0$ are given. The latter is often a quantity proportional to the inverse of the time step in an implicit time integration method applied to a nonstationary Stokes problem; $\xi = 0$ corresponds to the classical stationary Stokes problem.

In this work, we focus on standard finite difference and nodal finite elements discretizations,[1] which lead to a linear system of the form

$$\begin{pmatrix} A & B^T \\ B & -C \end{pmatrix} \begin{pmatrix} \mathbf{u} \\ \mathbf{p} \end{pmatrix} = \begin{pmatrix} \mathbf{b_u} \\ \mathbf{b_p} \end{pmatrix}. \tag{2.2}$$

In this system matrix, A is the discrete representation of the operator $\xi - \nu\Delta$; more precisely, A is block diagonal with one diagonal block per spatial dimension, being the discrete operator acting on one of the velocity components. It further follows that A is symmetric and positive definite (SPD). The matrix block B^T is the discrete gradient and $(-B)$ the discrete divergence; C is a stabilization term which is needed by some discretization schemes to avoid spurious solutions. Such spurious solutions arise when the discrete gradient admits more than the constant vector in its null space or near null space, i.e., when the discrete gradient is zero or near zero for some spurious pressure modes. The existence of such modes depends on which discretization scheme is used for velocities and pressure. We refer to, e.g., [46] and [16] for more details on, respectively, finite difference and finite element discretizations. Note a required property of the stabilization operator: if B is not full rank, C has to be positive definite on the null space of B^T, which further entails that the system matrix is nonsingular [3].

An important exception to this latter rule is when the boundary conditions are such that the physical pressure is only determined up to a constant. In such cases some additional condition is needed to make the problem well posed. Often, one imposes that the mean pressure is equal to zero. Then, the discrete system is singular but compatible. Regarding iterative solvers, this kind of singularity is, in general, harmless, and convergence to a particular solution is obtained without specific treatment; see, e.g., [11, 17, 31] for theoretical results and [14] for a practical example with a multigrid method. In the present work, all numerical examples are singular Stokes problems, and the theoretical results are formulated in a way that covers both singular and regular systems, taking advantage of the recent extension to singular systems of the algebraic convergence theory of two-grid methods [31].

[1] Other schemes lie outside of the scope, like discontinuous Galerkin methods; see [1] and the references therein for geometric multigrid methods in this context.

Test problems. Throughout this paper, results are illustrated with experiments based on the following two particular discretizations of (2.1), the second of which being considered in both two- and three-dimensional versions. (In the numerical results section (section 6), we additionally consider some examples of finite element discretizations with unstructured mesh.)

MAC scheme (2 dimensions). *In this problem,* Ω *is the unit square,* $\nu = 1$ *, and one imposes Dirichlet boundary conditions for all velocity components. One uses the marker and cell (MAC) finite difference discretization on a uniform staggered grid with mesh size* h *. As this scheme is naturally stable,* $C = 0$ *in this example.*

Collocated grid (2 dimensions/3 dimensions). *In this problem,* Ω *is the unit square/ cube,* $\nu = 1$ *, and one imposes Dirichlet boundary conditions for all velocity components. One uses the standard finite difference discretization on a collocated uniform grid with mesh size* h *. With this scheme, all unknowns are located at the vertices of grid cells, which makes the discretization somewhat easier but induces the presence of spurious pressure modes with zero discrete divergence* [46]*. Hence a form of stabilization is required and, according to the discussion in* [19]*,* C *is set equal to the five/seven point discretization of* $(16\,\nu)^{-1}h^2\Delta$ *(with Neumann boundary conditions).*

3. Algebraic transformations and two-grid analysis. In this section, we first present (sections 3.1 and 3.2) the proposed algebraic transformations and give (section 3.3) a basic description of associated AMG methods with unknown-based coarsening. We next (section 3.4) discuss the peculiarities stemming from the possible singularity of the system matrix. Finally (section 3.5), we develop the two-grid analysis for right-hand transformed matrices.

3.1. Right-hand transformation. As preliminary step, we first change the sign of the last block of rows in (2.2), yielding

$$\begin{pmatrix} A & B^T \\ -B & C \end{pmatrix} \begin{pmatrix} \mathbf{u} \\ \mathbf{p} \end{pmatrix} = \begin{pmatrix} \mathbf{b_u} \\ -\mathbf{b_p} \end{pmatrix} . \tag{3.1}$$

Next, let D_A be some SPD matrix of the same size as A and whose inverse is sparse. In this work, we focus on $D_A = \mathrm{diag}(A)$, but the theoretical results are formulate for general D_A. We perform the change of variables

$$\begin{pmatrix} \mathbf{u} \\ \mathbf{p} \end{pmatrix} = \begin{pmatrix} I & -\alpha D_A^{-1} B^T \\ & I \end{pmatrix} \begin{pmatrix} \widehat{\mathbf{u}} \\ \widehat{\mathbf{p}} \end{pmatrix} , \tag{3.2}$$

where α is a positive parameter. This leads to the transformed system

$$\widehat{\mathcal{A}} \begin{pmatrix} \widehat{\mathbf{u}} \\ \widehat{\mathbf{p}} \end{pmatrix} = \begin{pmatrix} \mathbf{b_u} \\ -\mathbf{b_p} \end{pmatrix} , \tag{3.3}$$

where

$$\begin{aligned} \widehat{\mathcal{A}} = \begin{pmatrix} A & \widehat{B}^T \\ -B & \widehat{C} \end{pmatrix} &= \begin{pmatrix} A & B^T \\ -B & C \end{pmatrix} \begin{pmatrix} I & -\alpha D_A^{-1} B^T \\ & I \end{pmatrix} \\ &= \begin{pmatrix} A & (I - \alpha A\, D_A^{-1}) B^T \\ -B & C + \alpha\, B\, D_A^{-1} B^T \end{pmatrix} . \end{aligned} \tag{3.4}$$

The approach then solves the above system with a monolithic multigrid method using an unknown-based-type coarsening [12, 36], in which the prolongation operator is set

up by considering separately the different types of unknowns (in the present case, velocity components and pressure). More details on this topic are given in section 3.3 below.

It is worth discussing the connection with distributive smoothing [8, 35, 39, 47]. This latter approach also uses right preconditioning with an upper block triangular matrix. However, the transformation concerns only the smoothing stage: once relaxation has been applied to the transformed variable $\widehat{\mathbf{u}}$, $\widehat{\mathbf{p}}$, the original variables $\mathbf{u}$, $\mathbf{p}$ are restored and the coarse grid correction is performed in the usual way. Another difference lies in the upper triangular matrix used for transformation. Typically, distributive smoothing methods use $\alpha\, D_A = I$ and a discrete negative Laplace operator $(-\Delta_h)$ as bottom right block (instead of the identity). The main goal is indeed to make the transformed matrix close to lower block triangular, i.e., to obtain that the top right block $\widehat{B}^T$ is small in some sense. With the approach presented here, the main goal is different: it is crucial to ensure that $\widehat{C}$ has a suitable structure for the used AMG strategy, as discussed in section 3.3.

3.2. Other transformations. Instead of the changes of variables (3.2), one may multiply both sides of (3.1) to the left by

$$\begin{pmatrix} I & \\ \alpha B\, D_A^{-1} & I \end{pmatrix} ,$$

yielding a transformed system with matrix

$$\bar{\mathcal{A}} \;=\; \begin{pmatrix} A & B^T \\ -B(I-\alpha A\, D_A^{-1}) & C+\alpha\, B\, D_A^{-1} B^T \end{pmatrix} \;=\; \begin{pmatrix} A & B^T \\ -\widehat{B} & \widehat{C} \end{pmatrix} . \tag{3.5}$$

The method in [44] amounts to such left-hand transformation with $\alpha = 1$. As will be seen in section 3.5, the right- and left-hand transformations are in fact equivalent with respect to our theoretical analysis. We further checked that the practical performances are very similar with both approaches. Hence they should not be seen as essentially different, but, rather, as slight variations of the same method. When, as usual, the convergence is controlled via the residual norm, we prefer the right-hand transformation because then the original and transformed systems have the same residual up to a change of sign for the second block of equations. Accordingly, below we only report numerical results for right-hand transformations.

Right-hand and left-hand transformations can also be combined, leading to the two-sided transformation proposed in [32], with matrix

$$\begin{pmatrix} I & \\ \alpha B\, D_A^{-1} & I \end{pmatrix} \begin{pmatrix} A & B^T \\ -B & C \end{pmatrix} \begin{pmatrix} I & -\alpha D_A^{-1} B^T \\ & I \end{pmatrix}$$
$$= \begin{pmatrix} A & (I_n - \alpha A\, D_A^{-1}) B^T \\ -B(I_n - \alpha D_A^{-1}\, A) & C + B(2\alpha D_A^{-1} - \alpha^2 D_A^{-1} A D_A^{-1}) B^T \end{pmatrix} .$$

Similarly to the matrix in (3.1), the offdiagonal blocks are the opposite of the transpose of each other, which facilitates the theoretical analysis. We omit the details of this approach here, since a complete analysis is already available in [32], whereas the numerical results reported below give the advantage to one-sided transformations. As will be seen, the convergence is in fact about the same with both transformations. However, iterations are cheaper with the one-sided transformations because the transformed matrices have then significantly fewer nonzero entries than with the two-sided approach.

3.3. AMG methods with unknown-based coarsening. We first briefly describe multigrid methods for a general linear system

$$\mathcal{A}\,\mathbf{x} \;=\; \mathbf{b} \tag{3.6}$$

(without special structure). These methods are based on the recursive use of a two-grid method, which combines smoothing iterations with a coarse grid correction. Smoothing iterations are simple stationary iterations with a standard preconditioner. The coarse grid correction is based on solving a coarse representation of the problem with a reduced number of unknowns.

With the AMG methods, this correction is entirely determined by the prolongation matrix $\mathcal{P}$, of dimension $n \times n_c$, where n_c is the number of coarse unknowns. It extends to the fine grid a vector defined on the coarse space. The reverse operation is performed with the transpose of $\mathcal{P}$. If $\widetilde{\mathbf{x}}$ is the current approximation of the solution, the corresponding coarse grid correction is

$$\mathcal{P}\,\mathcal{A}_c^{-1}\mathcal{P}^T\,(\mathbf{b} - \mathcal{A}\,\widetilde{\mathbf{x}})\,,$$

where

$$\mathcal{A}_c \;=\; \mathcal{P}^T\mathcal{A}\,\mathcal{P} \tag{3.7}$$

is the coarse grid matrix. Letting $\widehat{\mathbf{x}}$ be the exact solution of (3.6), the error associated with $\widetilde{\mathbf{x}}$ is $\widehat{\mathbf{x}} - \widetilde{\mathbf{x}}$, and adding the coarse grid correction to $\widetilde{\mathbf{x}}$ amounts to multiplying this error by the iteration matrix

$$T_c \;=\; I - \mathcal{P}\,\mathcal{A}_c^{-1}\mathcal{P}^T\mathcal{A}\;.$$

The global iteration matrix associated with a stationary two-grid method is obtained by multiplying T_c by the iteration matrix(ces) associated with smoothing. In this work, we mainly use the symmetrized Gauss–Seidel smoothing, which combines presmoothing with forward Gauss–Seidel and postsmoothing with backward Gauss–Seidel. The two-grid iteration matrix is then

$$T_{\mathrm{TG}} \;=\; \big(I - \mathcal{U}^{-1}\mathcal{A}\big)\big(I - \mathcal{P}\,\mathcal{A}_c^{-1}\mathcal{P}^T\mathcal{A}\big)\big(I - \mathcal{L}^{-1}\mathcal{A}\big)\;,$$

where $\mathcal{L} = \mathrm{tril}(\mathcal{A})$ and $\mathcal{U} = \mathrm{triu}(\mathcal{A})$. Our theoretical results are, however, limited to a simpler variant with a single postsmoothing iteration based on the Jacobi method, for which

$$T_{\mathrm{TG}} \;=\; \big(I - \omega\,\mathcal{D}^{-1}\mathcal{A}\big)\big(I - \mathcal{P}\,\mathcal{A}_c^{-1}\mathcal{P}^T\mathcal{A}\big)\;,$$

where $\mathcal{D} = \mathrm{diag}(\mathcal{A})$, and where ω is a relaxation parameter.

In practice, the AMG methods are more effective when used as a preconditioner for a Krylov subspace method such as CG (in the symmetric case), GMRES, or GCR. The two-grid preconditioner $\mathcal{B}_{\mathrm{TG}}$ is related to the iteration matrix via

$$I - \mathcal{B}_{\mathrm{TG}}\,\mathcal{A} \;=\; T_{\mathrm{TG}}\;.$$

Thus, for the variant with a single Jacobi postsmoothing step, one finds

$$\mathcal{B}_{\mathrm{TG}} \;=\; \omega\,\mathcal{D}^{-1} + \;\big(I - \omega\,\mathcal{D}^{-1}\mathcal{A}\big)\mathcal{P}\,\mathcal{A}_c^{-1}\mathcal{P}^T\mathcal{A}\;.$$

As indicated above, the two-grid method is rarely used as such, but rather constitutes a building block for developing a multigrid scheme. In fact, the application of the two-grid method requires solving systems

$$\mathcal{A}_c\,\mathbf{x}_c \;=\; \mathbf{r}_c \tag{3.8}$$

with the coarse grid matrix, as indicated by the presence of $\mathcal{A}_c^{-1}$ in the definitions of $T_{\rm TG}$ and $\mathcal{B}_{\rm TG}$. Within a multigrid algorithm, instead of the exact solution, one uses the approximation obtained by performing 1 or 2 iterations with the same two-grid method, but applied this time at the coarse level. This thus brings us to a coarser level, and the process is then repeated until the coarse system is sufficiently small so that an exact solution can be obtained at low cost.

The chosen iterative scheme to solve the coarse systems defines the multigrid cycle: the V-cycle is obtained with one stationary iteration, the W-cycle with two stationary iterations, and the K-cycle [33] with two iterations accelerated by a Krylov subspace method.

With AMG schemes, the prolongation $\mathcal{P}$ is not fixed by the geometry, but obtained by applying appropriate algorithms to the system matrix. The corresponding solvers can then be used in black box mode. Note that these algorithms must also be used recursively: once $\mathcal{P}$ has been obtained for the fine grid on the basis of $\mathcal{A}$, the coarse grid matrix $\mathcal{A}_c$ is computed via (3.7), and then one has to apply again the algorithm to $\mathcal{A}_c$ to obtain the prolongation at this coarse level, and so on.

These algorithms have been mainly developed for matrices corresponding to the discretization of scalar PDEs. For systems of PDEs, the unknown-based coarsening approach deals separately with the different types of unknowns, defining a prolongation for each type based on the corresponding diagonal block in the system matrix. Thus, for the transformed matrix $\widehat{A}$, we will let

$$\mathcal{P} \;=\; \begin{pmatrix} P_A & \\ & P_{\widehat{C}} \end{pmatrix}, \tag{3.9}$$

and define P_A based on the A block and $P_{\widehat{C}}$ based on the $\widehat{C}$ block. Note that the approach remains of black-box-type, but it is necessary to provide the solver with a properly ordered matrix and indicate the size of the different blocks.

As A is block diagonal with each diagonal block corresponding to a discrete negative Laplacian, the standard coarsening algorithms will work well. This is less clear for $\widehat{C} = C + \alpha\, B\, D_A^{-1} B^T$, since the structure of this term depends on the discretization scheme. However, the dominant term will often be $\alpha\, B\, D_A^{-1} B^T$ since C, when it is nonzero, is just a stabilization term, so in principle small. Furthermore, since $(-B)$ is a discrete representation of the divergence, and B^T a discrete representation of the gradient, their product is close to a negative Laplacian $(-\Delta_h)$. Therefore, usually, $\widehat{C}$ will also have a favorable structure for the application of AMG methods. This is the first expected benefit of the transformation: without it, the unknown-based coarsening must compute P_C on the basis of C. It is therefore not usable when $C = 0$ and hazardous otherwise, as C does not necessarily have a favorable structure.

To make the above discussion more concrete, consider, for example, the MAC scheme. Then $\widehat{C} = B\, D_A^{-1} B^T$ corresponds, up to a scaling factor, to the standard finite difference formula for the discrete negative Laplacian (see below for more details). With finite element discretizations, it is harder to directly connect $\widehat{C}$ with a discrete Laplacian. However, for (stabilized) Q_1/P_0 and (stable) Q_2/Q_1 mixed finite elements[2] on regular 2D grids, we checked that $\widehat{C}$ has nonnegative row-sum and nonpositive offdiagonal entries at least away from domain boundaries—such properties often suffice to ensure the proper functioning of AMG methods. On the other hand, more positive offdiagonal entries appear with (stabilized) P_1/P_1 and (stable) Crouzeix–Raviart elements [13] (as described in [16], where they are denoted P_{2^*}/P_{-1}). However, negative

[2]We follow the taxonomy in [16].

offdiagonal entries still dominate whereas in section 6 we show with two examples that the proposed approach may work well with these discretizations.

3.4. Singular matrices. Here we discuss the peculiarities of multigrid methods for singular systems. For the sake of simplicity, we restrict ourselves to the type of singularity commonly arising with Stokes problems, that is, we assume that the linear system (2.2) has a unique singular mode $(\mathbf{0}^T\ \mathbf{e}_p^T)^T$, with $\mathbf{e}_p$ such that $C\,\mathbf{e}_p = \mathbf{0}$ and $B^T\,\mathbf{e}_p = \mathbf{0}$. Regarding the transformed matrix (3.4), this implies $\widehat{C}\,\mathbf{e}_p = \mathbf{0}$ and $\mathbf{e}_p^T\widehat{C} = \mathbf{0}^T$, and, hence,

$$\mathcal{N}(\widehat{\mathcal{A}}) \;=\; \mathcal{N}(\widehat{\mathcal{A}}^T) \;=\; \mathrm{span}\left(\begin{pmatrix}\mathbf{0}\\ \mathbf{e}_p\end{pmatrix}\right)\ .$$

In practice, $\mathbf{e}_p$ is the constant vector, and $\widehat{C}$ has, therefore, zero row- and column-sum. Taking advantage that the constant vector is in the range of the prolongation with virtually any AMG scheme when the system matrix has zero row-sum, we then restrict ourselves to the cases where $\mathbf{e}_p \in \mathcal{R}(P_{\widehat{C}})$. Regarding a general linear system (3.6), the corresponding assumption is

$$\mathcal{N}(\mathcal{A}) \;=\; \mathcal{N}(\mathcal{A}^T) \;\subset\; \mathcal{R}(\mathcal{P})\ .$$

Clearly, the coarse grid matrix $\mathcal{A}_c = \mathcal{P}^T\mathcal{A}\,\mathcal{P}$ is then singular as well. However, according to the analysis in, e.g., [31], such a singularity is harmless because coarse systems (3.8) to be solved to implement the two-grid method are always compatible; moreover, which particular solution is picked up does not actually influence the convergence (even with Krylov subspace acceleration). It also follows that the recursive use of the method within a multigrid cycle does not raise any particular difficulty.

In the theoretical results below, we need to avoid referencing the inverse of $\mathcal{A}_c$ in the definition of the two-grid preconditioner. This is achieved using instead a pseudoinverse $\mathcal{A}_c^{\,g}$, which may be defined as any matrix such a that

$$\mathcal{A}_c\,\mathcal{A}_c^{\,g}\,\mathcal{A}_c \;=\; \mathcal{A}_c\ .$$

Observe that such a matrix is equal to $\mathcal{A}_c^{-1}$ when $\mathcal{A}_c$ is nonsingular, whereas, otherwise, $\mathcal{A}_c^{\,g}\,\mathbf{r}_c$ is a particular solution to any compatible system $\mathcal{A}_c\mathbf{x}_c = \mathbf{r}_c$ [2].

As regards the iterative methods for compatible singular systems, it is known that they are well posed if and only if the (preconditioned) system matrix is such that its zero eigenvalue has the same algebraic and geometric multiplicities [11, 17, 31]; then, only nonzero eigenmodes play a role in the convergence process. This motivates the formulation of the theorems in the next section, where we first issue a statement about the multiplicity of the zero eigenvalue, and next state bounds for nonzero eigenvalues.

3.5. Two-grid analysis. We start with the definition of the approximation property constant $K(G\,,P\,,D)$ associated with a triplet of matrices G, P, and D, where G is the (symmetric) matrix to which the two-grid scheme is applied, P the prolongation matrix of the two-grid scheme, and D an SPD matrix related to the smoother (in practice, often $D = \mathrm{diag}(G)$). The usage of this approximation property constant traces back to [6]. Here the definition is extended to positive semidefinite matrices, assuming then $\mathcal{N}(G) \subset \mathcal{R}(P)$, in agreement with the theory in [31] and the above considerations.

DEFINITION 3.1. *Let G and D be $n \times n$ matrices such that G is symmetric and nonnegative definite and D is SPD. Let P be an $n \times n_c$ matrix of rank $n_c < n$ such that $\mathcal{N}(G) \subset \mathcal{R}(P)$. The associated approximation property constant is*

$$K(G\,,P\,,D) \;=\; \sup_{\mathbf{v}\notin\mathcal{N}(G)} \frac{\mathbf{v}^T\,D\big(I - P\,\big(P^T D\,P\big)^{-1}\,P^T D\big)\,\mathbf{v}}{\mathbf{v}^T\,G\,\mathbf{v}}\;. \tag{3.10}$$

The transformed matrix (3.4) is nonsymmetric. For a general (nonsymmetric) matrix $\mathcal{A}$ with nonnegative definite symmetric part $\mathcal{A}_S = \frac{1}{2}(\mathcal{A}+\mathcal{A}^T)$, the two-grid convergence theory developed in [28] highlights the role of $K\big(\mathcal{A}_S\,,\mathcal{P}\,,\mathcal{D}\big)$. Note that $\mathcal{A}_S$ is nonnegative definite if and only if $\mathcal{A}$ is nonnegative definite in $\mathbb{R}^n$, that is, if and only if $\mathbf{v}^T\mathcal{A}\,\mathbf{v} \geq 0$ for all $\mathbf{v} \in \mathbb{R}^n$. The earlier results needed by our subsequent analysis are gathered in the following theorem. We give a short proof because the extension to singular matrices is only partly covered by the developments in [31].

THEOREM 3.2. *Let $\mathcal{A}$ and $\mathcal{D}$ be $n\times n$ matrices such that $\mathcal{A}$ is nonnegative definite in $\mathbb{R}^n$ and $\mathcal{D}$ is SPD. Let $\mathcal{P}$ be an $n \times n_c$ matrix of rank $n_c < n$ such that $\mathcal{N}(\mathcal{A}) = \mathcal{N}(\mathcal{A}^T) \subset \mathcal{R}(\mathcal{P})$. Setting $\mathcal{A}_c = \mathcal{P}^T\mathcal{A}\,\mathcal{P}$, let $\mathcal{A}_c^{\,g}$ be any matrix such that $\mathcal{A}_c\mathcal{A}_c^{\,g}\mathcal{A}_c = \mathcal{A}_c$, and let*

$$\mathcal{B}_{\rm TG} \;=\; \omega\,\mathcal{D}^{-1} + \;\big(I - \omega\,\mathcal{D}^{-1}\mathcal{A}\big)\mathcal{P}\,\mathcal{A}_c^{\,g}\mathcal{P}^T\;,$$

where ω is a positive number.

Then the algebraic multiplicity of the eigenvalue 0 of $\mathcal{B}_{\rm TG}\mathcal{A}$ is equal to $\dim(\mathcal{N}(\mathcal{A}))$, and there holds

$$\min_{\lambda\in\sigma(\mathcal{B}_{\rm TG}\mathcal{A})\setminus\{0,1\}} \Re e(\lambda) \;\geq\; \omega\,\left(\,K\big(\mathcal{A}_S\,,\mathcal{P}\,,\mathcal{D}\big)\,\right)^{-1}\;, \tag{3.11}$$

$$\max_{\lambda\in\sigma(\mathcal{B}_{\rm TG}\mathcal{A})\setminus\{0,1\}} \Big|\lambda - \frac{\omega\,\beta}{2}\Big| \;\leq\; \frac{\omega\,\beta}{2}\;, \tag{3.12}$$

where $\mathcal{A}_S = \frac{1}{2}(\mathcal{A}+\mathcal{A}^T)$, and

$$\beta \;=\; \sup_{\mathbf{v}\in\mathbb{C}^n\setminus\mathcal{N}(\mathcal{A})} \frac{\mathbf{v}^T\,\mathcal{A}^T\,\mathcal{D}^{-1}\mathcal{A}\,\mathbf{v}}{\mathbf{v}^T\,\mathcal{A}_S\,\mathbf{v}} \;=\; \sup_{\mathbf{v}\in\mathbb{C}^n\setminus\mathcal{N}(\mathcal{A})} \frac{\mathbf{v}^T\,\mathcal{A}\,\mathcal{D}^{-1}\mathcal{A}^T\,\mathbf{v}}{\mathbf{v}^T\,\mathcal{A}_S\,\mathbf{v}}\;. \tag{3.13}$$

In addition, if $\mathcal{A}$ is symmetric, then the eigenvalues of $\mathcal{B}_{\rm TG}\mathcal{A}$ are real, the inequality (3.11) is sharp, and $\beta = \lambda_{\max}(\mathcal{D}^{-1}\mathcal{A})$.

Proof. The assumptions used are more restrictive than those of Theorem 3.3 in [31], whose application proves the statement about the algebraic multiplicity of the eigenvalue 0, and, further, that the other eigenvalues are either equal to 1 or are the inverses of the nonzero eigenvalues of the generalized eigenvalue problem

$$\omega^{-1}\,\mathcal{D}\,\Big(I - \mathcal{P}\big(\mathcal{P}^T\,\mathcal{D}\,\mathcal{P}\big)^{-1}\mathcal{P}^T\,\mathcal{D}\Big)\;\mathbf{v} \;=\; \mu\;\mathcal{A}\,\mathbf{v}\;,\qquad \mathbf{v}\in\mathcal{V}\;, \tag{3.14}$$

where $\mathcal{V}$ is any subspace complementary to $\mathcal{N}(\mathcal{A})$.

Next, any $\mathbf{v}$ satisfying (3.14) for $\mu \neq 0$ also satisfies $\mathcal{P}^T\mathcal{A}\,\mathbf{v} = \mathbf{0}$ (and hence $\mathbf{v}^T\mathcal{A}^T\mathcal{P} = \mathbf{0}^T$), as seen by multiplying both sides to the left by $\mathcal{P}^T$. Then, multiplying both sides to the left by $\mathbf{v}^T\mathcal{A}^T\mathcal{D}^{-1}$ yields, with $\mu = \lambda^{-1}$,

$$\omega^{-1}\;\mathbf{v}^T\mathcal{A}^T\,\mathbf{v} \;=\; \lambda^{-1}\;\mathbf{v}^T\mathcal{A}^T\mathcal{D}^{-1}\mathcal{A}\,\mathbf{v}\;.$$

Thus, since $\mathbf{v} \in \mathcal{V}$ implies $\mathbf{v} \notin \mathcal{N}(\mathcal{A})$,

$$\Re\mathrm{e}(\lambda^{-1}) \;=\; \frac{\omega^{-1}\;\Re\mathrm{e}(\mathbf{v}^T\mathcal{A}^T\,\mathbf{v})}{\mathbf{v}^T\mathcal{A}^T\mathcal{D}^{-1}\mathcal{A}\,\mathbf{v}} \;=\; \frac{\omega^{-1}\;\mathbf{v}^T\mathcal{A}_S\,\mathbf{v}}{\mathbf{v}^T\mathcal{A}^T\mathcal{D}^{-1}\mathcal{A}\,\mathbf{v}} \;\geq\; (\omega\,\beta)^{-1}\;,$$

where we take β as given by the first equality (3.13). The above inequality is in fact equivalent to (3.12) because

$$\Re\mathrm{e}\left((\lambda_R + i\,\lambda_I)^{-1}\right) = \frac{\lambda_R}{\lambda_R^2+\lambda_I^2} \;\geq\; \frac{1}{\omega\,\beta} \quad\Longleftrightarrow\quad \left(\lambda_R - \frac{\omega\,\beta}{2}\right)^2 + \lambda_I^2 \;\leq\; \frac{(\omega\,\beta)^2}{4}\;.$$

Moreover, (3.14) admits the same eigenvalues if $\mathcal{A}$ is exchanged for its transpose $\mathcal{A}^T$ (since the matrix in the left-hand side is symmetric). Hence, we may permute the role of $\mathcal{A}$ and $\mathcal{A}^T$, proving the second equality (3.13).

On the other hand, exploiting (3.14), the proof of (3.11) given in Corollary 2.2 of [28] for regular matrices applies verbatim to the singular case as well.

Finally, for symmetric $\mathcal{A}$, the statement about β straightforwardly follows from $\mathcal{A} = \mathcal{A}^T = \mathcal{A}_S$. Further, since both right- and left-hand sides matrices in (3.14) are then symmetric and nonnegative definite, the largest eigenvalue μ is equal to $\omega^{-1}K\big(\mathcal{A},\mathcal{P},\mathcal{D}\big)$. This proves the sharpness of (3.11) in this case. ∎

The sharpness result in the SPD case gives a further interpretation of the approximation property constant: if ω is selected in such a way that the eigenvalue 1 is between the lower bound (3.11) and the upper bound (3.12), then the condition number of the two-grid method with a single Jacobi postsmoothing step is bounded by $\lambda_{\max}(\mathcal{D}^{-1}\mathcal{A})\cdot K\big(\mathcal{A},\mathcal{P},\mathcal{D}\big)$. Moreover, equality is reached with $\omega = \big(\lambda_{\max}(\mathcal{D}^{-1}\mathcal{A})\big)^{-1}$ (so that the upper bound from (3.12) is equal to 1); then,

$$K\big(\mathcal{A},\mathcal{P},\mathcal{D}\big) \;=\; \frac{\kappa_{\mathrm{TG}}}{\lambda_{\max}(\mathcal{D}^{-1}\mathcal{A})}\;, \tag{3.15}$$

where $\kappa_{\mathrm{TG}} = \frac{max_{\lambda\in\sigma(\mathcal{B}_{\mathrm{TG}}\mathcal{A})}\lambda}{min_{\lambda\in\sigma(\mathcal{B}_{\mathrm{TG}}\mathcal{A})\setminus\{0\}}\lambda}$ is the two-grid condition number.

We are now ready to prove our eigenvalue bounds for transformed matrices $\widehat{\mathcal{A}}$ preconditioned by a two grid scheme with one Jacobi post-smoothing step and an unknown-based type prolongation.

THEOREM 3.3. *Let $\widehat{\mathcal{A}}$ be an $n\times n$ matrix satisfying either* (3.4) *or* (3.5)*, where A and D_A are SPD, where C is symmetric nonnegative definite, and where α is a positive number. Let*

$$\mathcal{D} \;=\; \begin{pmatrix} D_A & \\ & D_{\widehat{C}} \end{pmatrix}\;,$$

where $D_{\widehat{C}}$ is SPD, and let

$$\mathcal{P} \;=\; \begin{pmatrix} P_A & \\ & P_{\widehat{C}} \end{pmatrix}$$

be an $n\times n_c$ matrix of rank $n_c < n$.

Assume that either $\widehat{\mathcal{A}}$ is nonsingular or $\mathcal{N}\big(\widehat{\mathcal{A}}\big) = \mathrm{span}\big\{\big(\mathbf{0}^T\;\mathbf{e}_p^T\big)^T\big\}$, where $\mathbf{e}_p \in \mathcal{R}(P_{\widehat{C}})$ is such that $C\,\mathbf{e}_p = \mathbf{0}$ and $B^T\,\mathbf{e}_p = \mathbf{0}$. Setting $\widehat{\mathcal{A}}_c = \mathcal{P}^T\widehat{\mathcal{A}}\,\mathcal{P}$, let $\widehat{\mathcal{A}}_c{}^g$ be any matrix such that $\widehat{\mathcal{A}}_c\,\widehat{\mathcal{A}}_c{}^g\,\widehat{\mathcal{A}}_c = \widehat{\mathcal{A}}_c$, and let

$$\mathcal{B}_{\mathrm{TG}} \;=\; \omega\,\mathcal{D}^{-1} + \;\big(I - \omega\,\mathcal{D}^{-1}\widehat{\mathcal{A}}\big)\mathcal{P}\;\widehat{\mathcal{A}}_c{}^g\;\mathcal{P}^T\;,$$

where ω is a positive number.

If $\alpha < 4\left(\lambda_{\max}(D_A^{-1}A)\right)^{-1}$*, then* $\widehat{\mathcal{A}}$ *is nonnegative definite in* $\mathbb{R}^n$*. Furthermore,* $\mathcal{B}_{\rm TG}\,\widehat{\mathcal{A}}$ *is regular if* $\widehat{\mathcal{A}}$ *is regular, and otherwise, the algebraic multiplicity of the eigenvalue* 0 *is equal to* 1*.*

In addition,

$$\min_{\lambda\in\sigma\left(\mathcal{B}_{\rm TG}\widehat{\mathcal{A}}\right)\setminus\{0,1\}} \Re e(\lambda) \;\geq\; \frac{\omega\;\left(1-\frac{\sqrt{\alpha\;\lambda_{\max}(D_A^{-1}A)}}{2}\right)}{\max\Big(\;K\big(A\,,P_A\,,D_A\big)\,,\,K\big(\widehat{C}\,,P_C\,,D_{\widehat{C}}\big)\;\Big)}\;, \tag{3.16}$$

$$\max_{\lambda\in\sigma\left(\mathcal{B}_{\rm TG}\widehat{\mathcal{A}}\right)\setminus\{0,1\}} \left|\lambda-\frac{\omega\,\beta}{2}\right| \;\leq\; \frac{\omega\,\beta}{2}\;, \tag{3.17}$$

where

$$\beta \;=\; \frac{\gamma+1+\sqrt{(\gamma-1)^2+\alpha\,\gamma\,\lambda_{\max}(D_A^{-1}A)}}{2\;\alpha\;\left(1-\frac{\alpha\;\lambda_{\max}(D_A^{-1}A)}{4}\right)} \tag{3.18}$$

with

$$\gamma \;=\; \alpha\;\lambda_{\max}\big(D_{\widehat{C}}^{-1}(C+B\,A^{-1}B^T)\big)\;. \tag{3.19}$$

Proof. We first consider a matrix $\widehat{\mathcal{A}}$ satisfying (3.4), and the results will be extended to matrices satisfying (3.5) thereafter. We start with the proof that the symmetric part

$$\widehat{\mathcal{A}}_S \;=\; \frac{1}{2}\left(\widehat{\mathcal{A}}+\widehat{\mathcal{A}}^T\right) \;=\; \begin{pmatrix} A & -\frac{\alpha}{2}A\,D_A^{-1}B^T \\ -\frac{\alpha}{2}B\,D_A^{-1}A & C+\alpha\,B\,D_A^{-1}B^T \end{pmatrix}$$

of $\widehat{\mathcal{A}}$ is such that

$$\widehat{\mathcal{A}}_S-\left(1-\frac{\sqrt{\alpha\;\lambda_{\max}(D_A^{-1}A)}}{2}\right)\widehat{\mathcal{A}}_D \qquad \text{is nonnegative definite}\;, \tag{3.20}$$

where

$$\widehat{\mathcal{A}}_D \;=\; \begin{pmatrix} A & \\ & C+\alpha\,B\,D_A^{-1}B^T \end{pmatrix}\;.$$

This will prove that $\widehat{\mathcal{A}}$ is nonnegative definite in $\mathbb{R}^n$. Then all assumptions of Theorem 3.2 are satisfied, which will immediately prove the statement about the multiplicity of the eigenvalue 0, whereas (3.16) and (3.17) will be deduced from (3.11) and (3.12).

We thus start with (3.20). Letting $\lambda_M=\lambda_{\max}(D_A^{-1}A)$, consider

$$\widehat{\mathcal{A}}_S-\left(1-\frac{\sqrt{\alpha\;\lambda_{\max}(D_A^{-1}A)}}{2}\right)\widehat{\mathcal{A}}_D=\frac{\sqrt{\alpha}}{2}\begin{pmatrix}\sqrt{\lambda_M}\,A & -\sqrt{\alpha}\,A\,D_A^{-1}B^T \\ -\sqrt{\alpha}\,B\,D_A^{-1}A & \sqrt{\lambda_M}\left(C+\alpha\,B\,D_A^{-1}B^T\right)\end{pmatrix}.$$

The top left block is SPD, whereas the Schur complement

$$\begin{aligned}&\frac{\sqrt{\alpha}}{2}\left(\lambda_M^{1/2}\left(C+\alpha\,B\,D_A^{-1}B^T\right)-\lambda_M^{-1/2}\alpha\,B\,D_A^{-1}A\,D_A^{-1}B^T\right)\\&=\frac{\sqrt{\alpha\,\lambda_M}}{2}\big(C+\alpha\,B\,D_A^{-1}\left(D_A-\lambda_M^{-1}A\right)D_A^{-1}B^T\big)\end{aligned}$$

is nonnegative definite, proving (3.20).

Further, considering Definition 3.1, (3.20) also implies that

$$K\big(\widehat{\mathcal{A}}_S\,,\mathcal{P}\,,\mathcal{D}\big)\;\le\;\left(1-\frac{\sqrt{\alpha\;\lambda_{\max}(D_A^{-1}A)}}{2}\right)^{-1}\;K\big(\widehat{\mathcal{A}}_D\,,\mathcal{P}\,,\mathcal{D}\big)\;.$$

Then, (3.16) follows from (3.11), since, in view of the block diagonal structure of both $\widehat{\mathcal{A}}_D$ and $\mathcal{D}\big(I-\mathcal{P}(\mathcal{P}^T\mathcal{D}\,\mathcal{P})\,\mathcal{P}^T\mathcal{D}\big)$, one has

$$K\big(\widehat{\mathcal{A}}_D\,,\mathcal{P}\,,\mathcal{D}\big)\;=\;\max\Big(\,K\big(A\,,P_A\,,D_A\big)\,,K\big(\widehat{C}\,,P_C\,,D_{\widehat{C}}\big)\Big)\;.$$

We now prove (3.17) by showing that, letting

$$\zeta\;=\;\frac{\overline{\gamma}+1+\sqrt{(\overline{\gamma}-1)^2+\alpha\,\overline{\gamma}\,\lambda_{\max}(D_A^{-1}A)}}{2\;\alpha\;\left(1-\frac{\alpha\;\lambda_{\max}(D_A^{-1}A)}{4}\right)}\;,$$

the matrix

$$\zeta\,\left(\widehat{\mathcal{A}}+\widehat{\mathcal{A}}^T\right)-2\,\widehat{\mathcal{A}}^T\mathcal{D}^{-1}\widehat{\mathcal{A}}\qquad\text{is nonnegative definite}\tag{3.21}$$

for any $\overline{\gamma}>\gamma$. With (3.12) this indeed yields the required result by taking the limit $\overline{\gamma}\to\gamma$ (and thus $\zeta\to\beta$). This trick is needed because the proof requires that $I-\overline{\gamma}^{-1}\,\alpha\,\widetilde{C}$ is SPD, where $\widetilde{C}=D_{\widehat{C}}^{-1/2}C\,D_{\widehat{C}}^{-1/2}$.

Let

$$\widetilde{\mathcal{A}}\;=\;\begin{pmatrix}\widetilde{A}&\widetilde{B}^T\\-\widetilde{B}&\widetilde{C}\end{pmatrix}\;=\;\mathcal{D}^{-1/2}\begin{pmatrix}A&B^T\\-B&C\end{pmatrix}\;\mathcal{D}^{-1/2}$$

with, thus, $\widetilde{A}=D_A^{-1/2}A\,D_A^{-1/2}$ and $\widetilde{B}=D_{\widehat{C}}^{-1/2}B\,D_A^{-1/2}$. Letting

$$\mathcal{U}\;=\;\begin{pmatrix}I&-\alpha D_A^{-1}B^T\\&I\end{pmatrix}\qquad\text{and}\qquad\widetilde{\mathcal{U}}\;=\;\mathcal{D}^{1/2}\mathcal{U}\;\mathcal{D}^{-1/2}\;=\;\begin{pmatrix}I&-\alpha\,\widetilde{B}^T\\&I\end{pmatrix}\;,$$

one has $\widehat{\mathcal{A}}=\mathcal{D}^{1/2}\widetilde{\mathcal{A}}\widetilde{\mathcal{U}}\,\mathcal{D}^{1/2}$ and, therefore,

$$\begin{aligned}&\widetilde{\mathcal{U}}^{-T}\mathcal{D}^{-1/2}\Big(\zeta\,(\widehat{\mathcal{A}}+\widehat{\mathcal{A}}^{\,T})-2\,\widehat{\mathcal{A}}^{\,T}\mathcal{D}^{-1}\widehat{\mathcal{A}}\Big)\mathcal{D}^{-1/2}\widetilde{\mathcal{U}}^{-1}\\&\quad=\;\zeta\,\left(\widetilde{\mathcal{U}}^{-T}\widetilde{\mathcal{A}}+\widetilde{\mathcal{A}}^T\widetilde{\mathcal{U}}^{-1}\right)-2\,\widetilde{\mathcal{A}}^T\widetilde{\mathcal{A}}\\&\quad=\;\zeta\,\begin{pmatrix}I&\\\alpha\,\widetilde{B}&I\end{pmatrix}\begin{pmatrix}\widetilde{A}&\widetilde{B}^T\\-\widetilde{B}&\widetilde{C}\end{pmatrix}+\zeta\,\begin{pmatrix}\widetilde{A}&-\widetilde{B}^T\\\widetilde{B}&\widetilde{C}\end{pmatrix}\begin{pmatrix}I&\alpha\,\widetilde{B}^T\\&I\end{pmatrix}\\&\qquad-2\,\begin{pmatrix}\widetilde{A}&-\widetilde{B}^T\\\widetilde{B}&\widetilde{C}\end{pmatrix}\begin{pmatrix}\widetilde{A}&\widetilde{B}^T\\-\widetilde{B}&\widetilde{C}\end{pmatrix}\\&\quad=\;\zeta\,\begin{pmatrix}2\,\widetilde{A}&\alpha\,\widetilde{A}\,\widetilde{B}^T\\\alpha\,\widetilde{B}\,\widetilde{A}&2\,\widetilde{C}+2\,\alpha\,\widetilde{B}\,\widetilde{B}^T\end{pmatrix}-2\,\begin{pmatrix}\widetilde{A}^2+\widetilde{B}^T\widetilde{B}&\widetilde{A}\,\widetilde{B}^T-\widetilde{B}^T\widetilde{C}\\\widetilde{B}\,\widetilde{A}-\widetilde{C}\,\widetilde{B}&\widetilde{C}^2+\widetilde{B}\,\widetilde{B}^T\end{pmatrix}.\end{aligned}\tag{3.22}$$

Clearly, (3.21) holds if and only if the right-hand side of (3.22) is nonnegative definite. Further,

$$\lambda_{\max}\big(D_{\widehat{C}}^{-1}(C + B\,A^{-1}B^T)\big) \;=\; \lambda_{\max}\big(\widetilde{C} + \widetilde{B}\,\widetilde{A}^{-1}\widetilde{B}^T\big) \;<\; \alpha^{-1}\,\overline{\gamma}$$

implies that all eigenvalues of $\big(I-\overline{\gamma}^{-1}\,\alpha\,\widetilde{C}\big)^{-1}\,\widetilde{B}\,\widetilde{A}^{-1}\widetilde{B}^T$ are in the interval $[0,\alpha^{-1}\,\overline{\gamma})$. Hence, the eigenvalues of $\widetilde{A}^{-1}\widetilde{B}^T\big(I-\overline{\gamma}^{-1}\,\alpha\,\widetilde{C}\big)^{-1}\,\widetilde{B}$ are also in the interval $[0,\alpha^{-1}\,\overline{\gamma})$ [18, Theorem 1.3.22], that is, since $\big(I-\overline{\gamma}^{-1}\,\alpha\,\widetilde{C}\big)^{-1} = I+\overline{\gamma}^{-1}\,\alpha\,\widetilde{C}(I-\overline{\gamma}^{-1}\,\alpha\,\widetilde{C})^{-1}$, the matrix

$$\Big(\alpha^{-1}\,\overline{\gamma}\,\widetilde{A}-\overline{\gamma}^{-1}\,\alpha\,\widetilde{B}^T\,\widetilde{C}\big(I-\overline{\gamma}^{-1}\,\alpha\,\widetilde{C}\big)^{-1}\widetilde{B}\Big)-\widetilde{B}^T\,\widetilde{B} \qquad \text{is nonnegative definite.}$$

Therefore, the right-hand side of (3.22) is a fortiori nonnegative definite if it remains nonnegative definite when exchanging for the first term in the above expression the term $\widetilde{B}^T\,\widetilde{B}$ that appears in the top left block of the second matrix. With this substitution and some rearranging of the terms, the right-hand side of (3.22) becomes

$$\begin{aligned} &2\begin{pmatrix} \zeta\,\widetilde{A}-\widetilde{A}^2-\alpha^{-1}\,\overline{\gamma}\,\widetilde{A} & \Big(\dfrac{\alpha\,\zeta}{2}-1\Big)\,\widetilde{A}\,\widetilde{B}^T \\ \Big(\dfrac{\alpha\,\zeta}{2}-1\Big)\,\widetilde{B}\,\widetilde{A} & (\alpha\,\zeta-1)\,\widetilde{B}\,\widetilde{B}^T \end{pmatrix} \\ &+2\begin{pmatrix} \overline{\gamma}^{-1}\,\alpha\,\widetilde{B}^T\,\widetilde{C}(I-\overline{\gamma}^{-1}\,\alpha\,\widetilde{C})^{-1}\widetilde{B} & \widetilde{B}^T\widetilde{C} \\ \widetilde{C}\,\widetilde{B} & \zeta\,\widetilde{C}-\widetilde{C}^2 \end{pmatrix}, \end{aligned} \tag{3.23}$$

and we are thus left with the proof that this matrix is nonnegative definite.

We now show that both terms are actually nonnegative definite. For the second one we obtain

$$\begin{aligned} &\begin{pmatrix} \overline{\gamma}^{-1}\,\alpha\,\widetilde{B}^T\,\widetilde{C}(I-\overline{\gamma}^{-1}\,\alpha\,\widetilde{C})^{-1}\widetilde{B} & \widetilde{B}^T\widetilde{C} \\ \widetilde{C}\,\widetilde{B} & \zeta\,\widetilde{C}-\widetilde{C}^2 \end{pmatrix} \\ &\quad= \begin{pmatrix} \widetilde{B}^T\,\widetilde{C}^{1/2} & \\ & \widetilde{C}^{1/2} \end{pmatrix}\begin{pmatrix} \overline{\gamma}^{-1}\,\alpha\,(I-\overline{\gamma}^{-1}\,\alpha\,\widetilde{C})^{-1} & I \\ I & \zeta\,I-\widetilde{C} \end{pmatrix}\begin{pmatrix} \widetilde{C}^{1/2}\widetilde{B} & \\ & \widetilde{C}^{1/2} \end{pmatrix}, \end{aligned}$$

which is nonnegative definite if and only if the Schur complement of the middle term of the right-hand side is nonnegative definite, that is, if and only if

$$\zeta\,I-\widetilde{C}-\alpha^{-1}\,\overline{\gamma}\,(I-\overline{\gamma}^{-1}\,\alpha\,\widetilde{C}) \;=\; (\zeta-\alpha^{-1}\,\overline{\gamma})\,I \quad \text{is nonnegative definite.}$$

This is true because $\zeta\,\alpha$ is increasing with α for $0\le\alpha<4\,\big(\lambda_{\max}(D_A^{-1}A)\big)^{-1}$, and equal to $\max(\overline{\gamma},1)$ when $\alpha=0$, i.e., $\zeta>\alpha^{-1}\,\max(\overline{\gamma},1)$ for any $0<\alpha<4\,\big(\lambda_{\max}(D_A^{-1}A)\big)^{-1}$.

We eventually consider the first term in (3.23):

$$\begin{aligned} &\begin{pmatrix} \zeta\,\widetilde{A}-\widetilde{A}^2-\alpha^{-1}\,\overline{\gamma}\,\widetilde{A} & \Big(\dfrac{\alpha\,\zeta}{2}-1\Big)\,\widetilde{A}\,\widetilde{B}^T \\ \Big(\dfrac{\alpha\,\zeta}{2}-1\Big)\,\widetilde{B}\,\widetilde{A} & (\alpha\,\zeta-1)\,\widetilde{B}\,\widetilde{B}^T \end{pmatrix} \\ &\quad= \begin{pmatrix} \widetilde{A}^{1/2} & \\ & \widetilde{B} \end{pmatrix}\begin{pmatrix} (\zeta-\alpha^{-1}\,\overline{\gamma})\,I-\widetilde{A} & \Big(\dfrac{\alpha\,\zeta}{2}-1\Big)\,\widetilde{A}^{1/2} \\ \Big(\dfrac{\alpha\,\zeta}{2}-1\Big)\,\widetilde{A}^{1/2} & (\alpha\,\zeta-1)\,I \end{pmatrix}\begin{pmatrix} \widetilde{A}^{1/2} & \\ & \widetilde{B}^T \end{pmatrix}. \end{aligned}$$

We have just seen that $\alpha\,\zeta > \max(\overline{\gamma}, 1) \geq 1$. Hence the above matrix is nonnegative definite if and only the Schur complement corresponding to the elimination of the bottom right block in the middle term of the right hand is nonnegative definite, that is, if and only if

$$(\zeta - \alpha^{-1}\,\overline{\gamma})\,I - \left(1 + \frac{\left(\frac{\alpha\,\zeta}{2} - 1\right)^2}{\alpha\,\zeta - 1}\right)\,\widetilde{A} \;=\; (\zeta - \alpha^{-1}\,\overline{\gamma})\,I - (\alpha\,\zeta - 1)^{-1}\left(\frac{\alpha\,\zeta}{2}\right)^2\,\widetilde{A}$$

is nonnegative definite. This holds if, setting $\lambda_M = \lambda_{\max}(D_A^{-1}A) = \lambda_{\max}(\widetilde{A})$,

$$(\alpha\,\zeta - 1)(\zeta - \alpha^{-1}\,\overline{\gamma}) - \left(\frac{\alpha\,\zeta}{2}\right)^2\,\lambda_M \;=\; \alpha\left(1 - \frac{\alpha\,\lambda_M}{4}\right)\zeta^2 - (\overline{\gamma} + 1)\zeta + \alpha^{-1}\overline{\gamma} \;\geq\; 0\,.$$

Regarding matrices satisfying (3.4), the proof is then concluded by checking that this expression is in fact equal to 0 for the given value of ζ.

Consider now a matrix $\bar{\mathcal{A}}$ satisfying (3.5). Letting

$$\widehat{\mathcal{A}}_- \;=\; \begin{pmatrix} A & -\widehat{B}^T \\ B & \widehat{C} \end{pmatrix} \;=\; \begin{pmatrix} A & -(I - \alpha A\,D_A^{-1})B^T \\ B & C + \alpha\,B\,D_A^{-1}B^T \end{pmatrix}\,,$$

one sees that $\bar{\mathcal{A}} = \widehat{\mathcal{A}}_-^{\,T}$, whereas $\widehat{\mathcal{A}}_-$ is a matrix satisfying (3.4) with B changed for $-B$. However, a global change of sign of the entries in B does not affect either the assumptions or the results proved above (observe that none of the parameters in (3.16) or (3.17) is affected). Hence these results apply as well to this matrix $\widehat{\mathcal{A}}_-$. Moreover, their proof is based upon Theorem 3.2, which yields exactly the same conclusions whenever applied to a matrix $\mathcal{A}$ or its transpose $\mathcal{A}^T$. It then follows that the results proved above carry over matrices $\bar{\mathcal{A}}$ satisfying (3.5). ∎

It is worth illustrating these bounds with an example. We consider the MAC finite difference discretization of a problem (2.1) with constant coefficients $\nu > 0$ and $\xi \geq 0$, using a uniform staggered grid (as in the MAC scheme (2D) test problem).

The matrix A has two diagonal blocks and, away from boundaries, each of them corresponds to the classical five point stencil for the operator $\xi - \nu\Delta$:

$$\frac{\nu}{h^2}\begin{bmatrix} & -1 & \\ -1 & 4\left(1 + \dfrac{h^2\,\xi}{4\,\nu}\right) & -1 \\ & -1 & \end{bmatrix}\,.$$

Further, away from the boundaries, $\widehat{C} = B\,D_A^{-1}B^T$ corresponds to the stencil

$$\frac{1}{4\,\nu + \xi\,h^2}\begin{bmatrix} & -1 & \\ -1 & 4 & -1 \\ & -1 & \end{bmatrix}\,,$$

which is the standard five point approximation for $-\frac{h^2}{4\,\nu + \xi\,h^2}\,\Delta$. On the other hand, it is shown in [32, p. 73] that, for an infinite grid, there holds $\lambda_{\max}(B\,A^{-1}B^T) \leq \frac{8}{8\,\nu + \xi\,h^2}$.

With these results, we can estimate the constants involved in the theorem, up to boundary effects that are not expected to have a dramatic influence [7, 37]. From the first stencil we deduce

$$\lambda_{\max}(D_A^{-1}A) \;\leq\; 2\;,$$

whereas, combining the above bound on $\lambda_{\max}(B\,A^{-1}B^T)$ with $C=0$ and $\mathrm{diag}(\widehat{C}) = \left(\nu+\frac{\xi\,h^2}{4}\right)I$, we obtain

$$\lambda_{\max}\big(D_{\widehat{C}}^{-1}(C+B\,A^{-1}B^T)\big) \;\leq\; \frac{8\,\nu+2\,\xi\,h^2}{8\,\nu+\xi\,h^2} \;\leq\; 2\;.$$

Finally, the approximation property constants $K\big(A\,,P_A\,,D_A\big)$ and $K\big(\widehat{C}\,,P_C\,,D_{\widehat{C}}\big)$ depend on the method used to determine P_A and $P_{\widehat{C}}$, and cannot be estimated rigorously with most AMG methods. However, the sharpness result in Theorem 3.2 offers us an important clue: with (3.15), one sees that if the selected AMG method works well for the subproblems associated with A and $\widehat{C}$, then the approximation property constants will be independent of the grid size and around 1 or only moderately larger. Moreover, because both the diagonal blocks of A and $\widehat{C}$ correspond to five point stencils, any sensible AMG method will actually work well.

Thus, choosing appropriately the parameters ω and α, the inequalities (3.16) and (3.17) prove that all (nonzero) eigenvalues of the preconditioned matrix are within a region of the complex plane that is both bounded and away from the origin; moreover, it is independent of main problem parameters ν and ξ, and also of the mesh or grid size.

As regards α, one sees that the bounds are meaningful if it is significantly both larger than zero and smaller than $4\left(\lambda_{\max}(D_A^{-1}A)\right)^{-1} \approx 2$. In practice, we observed that the results are near optimal with $\alpha=1$, which also corresponds to the method in [44]. In addition, the transformed matrix is then slightly sparser, because using $\alpha=1$ sets to zero the diagonal entries in the term $(I-\alpha A\,D_A^{-1})$ appearing in the top right block of $\widehat{\mathcal{A}}$ (see (3.4) or (3.5)). Accordingly, we uniformly use $\alpha=1$ in all numerical experiments reported below.

Having said that, the clustering of the eigenvalues proved by Theorem 3.3 seems somehow less effective than those proved in [32] for the two-sided transformation. One may then wonder whether the one-sided approach is indeed less efficient, or if this results from a shortcoming in the analysis.

This motivates the following experiment, where we also investigate the convergence obtained with better smoothing schemes. In Table 1, we report the number of GMRES iterations needed to reduce the residual norm by 10^{-6} when using both types of transformation combined with two-grid preconditioners. Three different smoothing schemes are considered: one single Jacobi postsmoothing step ("Jac Post"), Jacobi smoothing with both one pre- and one post-smoothing step ("Jac Pre+Post"), and symmetrized Gauss–Seidel smoothing as described in section 3.3 ("SGS"). In all cases, the prolongations P_A and $P_{\widehat{C}}$ are based on geometric boxwise plain aggregation, a choice that is motivated in the next sections.

With a single Jacobi smoothing step, the two-sided transformation leads to faster convergence for the MAC scheme, whereas the performances are comparable for the collocated grid. Note that in this case the convergence with GMRES acceleration is logically practically independent of ω: indeed, with one single smoothing step, changing this parameter only rescales the eigenvalues that are not equal to 1 (see the discussion of (3.14) in the proof of Theorem 3.2). This is not true any more with

TABLE 1
Number of GMRES iterations to reduce the residual error by 10^{-6} when using two-grid preconditioning with different smoothing schemes.

	Right-hand transformation						Two-sided transformation					
	Jac Post		Jac Pre+Post			SGS	Jac Post		Jac Pre+Post			SGS
ω	0.5	0.8	0.5	0.65	0.8		0.5	0.8	0.5	0.65	0.8	
h^{-1}:												
MAC scheme (2D), $\xi = 0$												
128	24	24	15	200	200	12	19	19	15	15	14	11
256	24	24	15	200	200	12	19	19	15	14	14	11
MAC scheme (2D), $\xi = 100$												
128	24	24	14	200	200	12	19	19	15	14	14	11
256	24	24	15	200	200	12	19	18	15	14	14	11
Collocated grid (2D), $\xi = 0$												
128	18	18	13	12	11	9	19	19	14	13	12	10
256	18	18	13	12	11	9	19	19	14	13	12	10
Collocated grid (2D), $\xi = 100$												
128	18	18	13	12	11	9	19	18	14	12	12	9
256	18	18	13	12	11	9	19	19	14	13	12	10

more smoothing steps: then, ω should be small enough to ensure that no error mode is actually amplified by (stationary) smoothing iterations. Clearly, this leads to a stricter requirement for the right-hand transformation. However, when this latter is met, and also with symmetrized Gauss–Seidel smoothing, the convergence is overall the same with both transformation types; i.e., the two-sided transformation seems to lose the advantage it had for the MAC scheme when using a single smoothing step.

4. Two-grid method at coarser levels. Here we discuss the application of the two-grid method at coarser levels; that is, to solve systems with the coarse grid matrix obtained after one or several coarsening steps.

After one step, the coarse grid matrix is $\widehat{\mathcal{A}}_c = \mathcal{P}^T \widehat{\mathcal{A}} \mathcal{P}$, i.e., its structure depends upon the components P_A and $P_{\widehat{C}}$ of $\mathcal{P}$. When the prolongation is set up with an AMG algorithm, it is never fully structured even if the matrix stems from a constant coefficient PDE discretized on a uniform grid. Accordingly, a rigorous analysis of coarse grid matrices is generally untractable. However, in the present context, significant insight can be gained by considering the model geometric prolongations that AMG methods aim at imitating.

Thus, for constant coefficient isotropic problems, classical AMG methods along the lines of [9, 38] tend to reproduce the standard $h - 2h$ coarsening and the associated geometric bilinear interpolation.[3] Using this latter on our test problems, we observed that, throughout coarsening steps, the nature of the different matrix blocks is preserved: away from the boundaries, one obtains regular stencils corresponding to Laplace operators for the diagonal blocks, and to first order derivatives for the offdiagonal blocks. However, the scaling of the blocks is not as in the original transformed matrix $\widehat{\mathcal{A}}$: the weight of the entries in the offdiagonal blocks relative to the entries in the diagonal blocks (e.g., the main diagonal) is increased by a factor of about two with each coarsening step. This stems from the type of discrete operators involved.

[3]For a 5 point stencil in 2 dimensions, standard Ruge–Stüben AMG will produce a red-black coarsening of the fine mesh, and $h-2h$ coarsening only at subsequent levels; however, for such small stencils, one often prefers aggressive coarsening, which imitates the $h-2h$ coarsening right away; see, e.g., the discussion in [38, Section A.7.1].

The diagonal blocks correspond to scaled discrete Laplace operators. With finite differences, the entries are $\mathcal{O}(h^{-2})$, hence they are reduced by a factor of 4 with $h - 2h$ coarsening. On the other hand, offdiagonal blocks correspond to discrete first order differential operators, with $\mathcal{O}(h^{-1})$ entries that are reduced by a factor of 2 only. The picture is the same with finite element discretizations, where the entries are $\mathcal{O}(h^{d-2})$ for second order differential operators and $\mathcal{O}(h^{d-1})$ for first order ones, where d is the spatial dimension of the problem.

This relative increase of the weight of the offdiagonal blocks has a dramatic impact on the potentialities of the AMG method at coarser levels. To see this, observe that applying AMG to $\widehat{\mathcal{A}}$ would be just trivial if the offdiagonal blocks vanished. Hence, in some sense, what is proved in Theorem 3.3 is that the diagonal blocks are strong enough to dominate the effect of the offdiagonal ones. Clearly, the bounds would deteriorate if these latter were increased without a corresponding increase of the term $B\,D_A B^T$ in the bottom right block. In fact, for fixed ω, damped Jacobi smoothing iterations will be divergent once the amplification of the offdiagonal blocks passes some threshold. Of course, one may decrease ω at coarser levels, but this would imply eigenvalues closer to the origin and hence increasingly slower convergence as one goes deeper in the hierarchy.

This phenomenon was observed in [44] in combination with classical AMG coarsening. There, more sophisticated smoothing schemes are used like ILU(0) with or without Krylov acceleration. While this avoids the divergence of stationary two-grid iterations, the reported results indicate that the convergence becomes unacceptably slow beyond some coarsening threshold. Moreover, similarly disappointing results are also obtained in combination with smoothed aggregation AMG.

However, comments in [44] are more positive regarding plain aggregation (referred to as zero order interpolation). To investigate this other coarsening type, we applied geometric boxwise aggregation to our test problems, and we indeed observed that the above phenomenon does not happen anymore. The relative weights of the blocks remain as in the original matrix, which we explain by the well-known overweighting of coarse grid matrices associated with plain aggregation for second order differential operators see, e.g., the discussion in [21], and the proposition in [5] to overrelax coarse grid correction terms by a factor close to two, to compensate for this phenomenon.

Whereas this overweighting is sometimes seen as a weak point of aggregation-based AMG methods, here it comes as good news. Thanks to it, the structure of the matrix remains roughly similar at the successive coarse levels and, hence, one may expect that the recursive use of the two-grid method does not raise particular difficulty in this context.

These considerations are illustrated in Figure 1, where, considering the collocated grid (2D) problem with $\xi = 0$, we depict the convex hull of the iteration matrices associated with damped Jacobi smoothing at several levels. Clearly, with (geometric) bilinear interpolation, there is a severe divergence from the second coarse level, whereas nothing particular happens with (geometric, boxwise) plain aggregation coarsening.

5. Multigrid strategy. For scalar PDEs, there are many valuable methods based either on classical AMG methods along the lines of [9, 38], or on smoothed aggregation AMG [41]. However, the developments in the preceding section show that their use is somehow uneasy as unknown-based coarsening strategy for the transformed matrices considered in this work. To compensate for the loss of stability, one may, for instance, also apply the transformation at selected coarse levels [45], but this

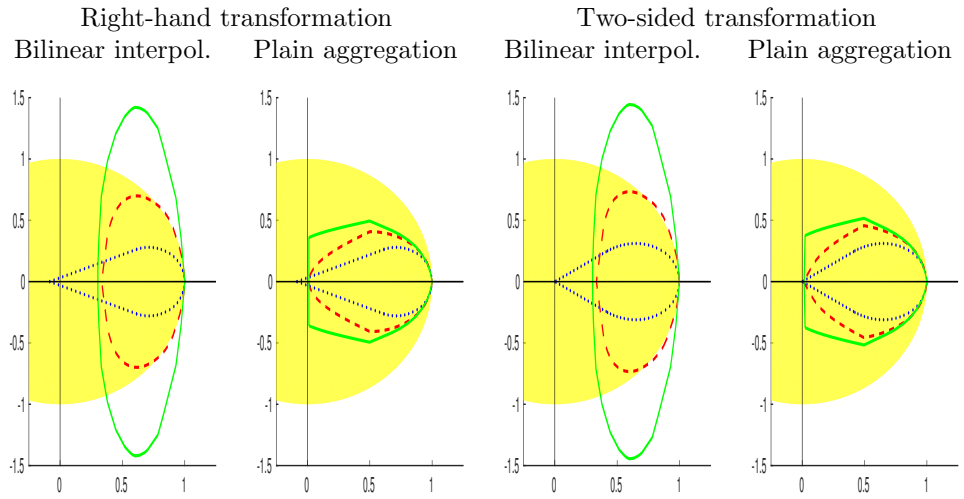

FIG. 1. *Convex hull of the eigenvalues of iteration matrices associated with damped Jacobi smoothing ($\omega = 0.5$) for the collocated grid (2D) problem with $\xi = 0$; ····· : fine grid level; - - -: first coarse level; ——: next coarse level; the yellow (shaded) region is a portion of the unit disk centered at the origin.*

increases the complexity of the scheme, making the coarse grid matrices significantly less sparse.

It seems then sensible to prefer the alternative offered by plain-aggregation-based AMG. Moreover, recent results about this approach show that it is already competitive for scalar, second order, elliptic PDEs, at least when carefully used [23, 27], that is, when applying an aggregation algorithm that takes into account the matrix entries so as to build only high quality aggregates [22, 29]. It is also important to use the K-cycle [33] to ensure that the optimal two-grid convergence [21, 22, 29] carries over the full multigrid scheme, i.e., to ensure that the convergence is independent of the number of levels.

On the other hand, the good results obtained above (see Table 1) with symmetrized Gauss–Seidel suggest selecting this smoothing scheme.

This leads us to the following multigrid strategy: two-grid scheme obtained from the combination of symmetrized Gauss–Seidel smoothing with unknown-based coarsening by plain aggregation, using more specifically the algorithm in [26]; multigrid scheme as in [27], obtained with the standard K-cycle for nonsymmetric matrices, in which all coarse systems are solved with two GCR iterations [15] using the two-grid preconditioner at the considered level (except at the coarsest level where a sparse direct solver is used).

In fact, this method requires only slight modifications of the method in [26], needed to implement the unknown-based coarsening. Furthermore, the corresponding code is available as the "block" version of the AGMG software [24], which we therefore used for the numerical experiments reported in the next section.

6. Numerical results. The main test problems and the used multigrid method have already been described at the end of section 2 and in section 5 (respectively). Let us further remember that the method is not applied directly to the original linear system (2.2), but to an equivalent, transformed, system. When we refer to the right-hand transformation, we mean the system (3.3) where system matrix $\widehat{\mathcal{A}}$ is as in

TABLE 2

Ratio between the numbers of nonzero entries in the transformed and original matrices (rat_{Tr}), algorithmic complexity ($\mathcal{C}_A$), and global complexity ($\mathcal{C}_G$).

	Right-hand transformation						Two-sided transformation					
	Size 1			Size 2			Size 1			Size 2		
ξ	rat_{Tr}	$\mathcal{C}_A$	$\mathcal{C}_G$	rat_{Tr}	$\mathcal{C}_A$	$\mathcal{C}_G$	rat_{Tr}	$\mathcal{C}_A$	$\mathcal{C}_G$	rat_{Tr}	$\mathcal{C}_A$	$\mathcal{C}_G$
MAC scheme (2D)												
0	1.9	1.3	2.5	1.9	1.3	2.6	3.1	1.3	3.9	3.1	1.3	3.9
10	1.9	1.3	2.5	1.9	1.3	2.6	3.1	1.3	3.9	3.1	1.3	3.9
100	1.9	1.3	2.5	1.9	1.3	2.6	3.1	1.3	3.9	3.1	1.3	3.9
1000	1.9	1.3	2.5	1.9	1.3	2.6	3.1	1.3	3.9	3.1	1.3	3.9
Collocated grid (2D)												
0	1.5	1.4	2.1	1.5	1.4	2.1	2.7	1.3	3.7	2.7	1.3	3.7
10	1.5	1.4	2.1	1.5	1.4	2.1	2.7	1.3	3.7	2.7	1.3	3.7
100	1.5	1.4	2.1	1.5	1.4	2.1	2.7	1.3	3.7	2.7	1.3	3.7
1000	1.5	1.4	2.1	1.5	1.4	2.1	2.7	1.3	3.7	2.7	1.3	3.7
Collocated grid (3D)												
0	1.7	1.8	3.1	1.7	1.8	3.1	3.3	1.9	6.3	3.4	2.0	6.6
10	1.7	1.7	2.9	1.7	1.7	2.9	3.3	1.9	6.3	3.4	2.0	6.6
100	1.7	1.8	3.0	1.7	1.7	2.9	3.3	1.9	6.2	3.4	2.0	6.6
1000	1.7	1.6	2.8	1.7	1.7	2.9	3.3	1.8	6.1	3.4	2.0	6.6

(3.4) with $\alpha = 1$; the two-sided transformation refers to the more complex transformation sketched in section 3.2 and fully described in [32].

In all cases, the multigrid method is used as a preconditioner for the GCR method restarted each 10 or each 30 iterations [15]. The right-hand side is a vector with random velocity components and zero pressure components, the initial approximation is the zero vector, and iterations are stopped when the relative residual norm is below 10^{-6}. All results are reported for two different grid sizes: in the 2D examples, *Size* 1 and *Size* 2 refer to, respectively, $h^{-1} = 256$ and $h^{-1} = 1024$, whereas, in the 3D example, they refer to $h^{-1} = 48$ and $h^{-1} = 96$. Hence the number of unknowns ranges approximately from 2×10^5 to 3×10^6 in 2 dimensions, and from 4×10^5 to 3.5×10^6 in 3 dimensions (i.e., there is about one order of magnitude between Size 1 and Size 2).

In Table 2, we report on *complexities*, that is, on the memory usage involved in the solution method. Two factors have to be taken into account. First, the transformed matrices have more nonzero entries than the original system matrix; in the table, the ratio between the numbers of nonzero entries in the transformed and original matrices is reported in the columns labeled "rat_{Tr}." Next, the multigrid preconditioner involves some overhead, which is characterized by the algorithm complexity $\mathcal{C}_A$, defined as the sum of all nonzero entries in the matrices at all levels divided by the number of nonzero entries in the fine grid matrix (here, the transformed matrix). Finally, in the present context, we further define the global complexity $\mathcal{C}_G$ as the sum of all nonzero entries in the matrices at all levels divided by the number of nonzero entries in the *original* fine grid matrix; in other words: $\mathcal{C}_G = \mathcal{C}_A \cdot \text{rat}_{\text{Tr}}$. Despite these cumulative effects, the complexities are acceptable in all cases, except perhaps in the 3D example with the two-sided transformation. The one-sided approach has, in fact, a clear advantage here, the transformed matrices having between 1.5 and 2 times fewer nonzero entries than with the two-sided transformation.

The iteration counts are reported in Table 3. Here, for comparison purpose, we include a standard method ("Block Diag. Prec."), which solves the original system

TABLE 3
Number of iterations to reduce the residual norm by a factor of 10^{-6} .

	Right-hand transf.				Two-sided transf.				Block Diag. Prec.			
	GCR(10)		GCR(30)		GCR(10)		GCR(30)		1 inner it.		2 inner it.	
Size	1	2	1	2	1	2	1	2	1	2	1	2
ξ:												
						MAC scheme (2D)						
0	17	17	16	17	16	17	16	17	57	64	41	44
10	17	17	16	17	16	17	16	17	55	62	37	41
100	15	17	15	16	15	16	15	16	49	57	33	37
1000	13	16	13	15	13	15	13	15	41	49	26	33
						Collocated grid (2D)						
0	20	27	20	26	15	17	15	16	61	66	46	50
10	20	27	19	26	15	16	15	16	58	66	44	48
100	18	24	17	23	15	16	14	15	53	61	38	42
1000	14	19	14	19	13	15	13	14	43	51	33	37
						Collocated grid (3D)						
0	15	17	15	17	15	15	14	15	68	71	48	51
10	15	17	15	17	15	15	14	15	63	69	46	49
100	14	16	14	16	13	14	13	14	53	59	38	43
1000	15	14	15	13	12	12	12	11	44	46	32	34

(2.2) by MINRES with state-of-the-art block diagonal preconditioner

$$\begin{pmatrix} \widetilde{A} & \\ & \widetilde{S} \end{pmatrix} ,$$

where $\widetilde{A}$ is an approximation of A , and $\widetilde{S}$ is an approximation of the Schur complement $C + B\,A^{-1}B^T$ [16]. The needed inverse of $\widetilde{A}$ is obtained by applying the same AMG method as for the transformed matrices; we consider either one single application of the multigrid preconditioner ("1 inner it."), or two FCG iterations [25] using this preconditioner to solve a system with matrix A ("2 inner it."). The latter option allows us to investigate whether it can be cost effective to use a more costly but more accurate approximation to A , as is the case when the the inverse of $\widetilde{A}$ obtained by applying a more standard AMG method, that provides a better approximation than the plain aggregation method considered here, but at a significantly higher cost (see the discussion in [30]).

Regarding the Schur complement approximation, in the stationary case ($\xi = 0$), it is standard to use $\widetilde{S} = \nu^{-1} I$ with finite difference discretizations. Time-dependent cases require more care, but the Cahouet–Chabard method [10] is both effective and optimal with respect to the mesh size and other problem parameters [34]. It uses $\widetilde{S}^{-1} = \nu\, I + \xi(-\widetilde{\Delta}_h)^{-1}$, where $\widetilde{\Delta}_h$ is an approximation to a discrete Laplace operator with Neumann boundary conditions for the pressure variables; again, for $(-\widetilde{\Delta}_h)^{-1}$, we consider either one application of the AMG preconditioner or two inner iterations, in each case applied to solve a system with an exact discrete Laplace operator $-\Delta_h$. Note that the latter has to be supplied to the solver, hence this preconditioner is not fully algebraic.

One sees that the block diagonal preconditioner requires significantly more iterations than monolithic AMG preconditioners, even when using enhanced approximations $\widetilde{A}$ and $\widetilde{S}$ with two inner iterations. This explains the timing results reported in

TABLE 4
Total solution time in microseconds per unknown.

	Right-hand transf.				Two-sided transf.				Block Diag. Prec.			
	GCR(10)		GCR(30)		GCR(10)		GCR(30)		1 inner it.		2 inner it.	
Size	1	2	1	2	1	2	1	2	1	2	1	2
ξ:												
MAC scheme (2D)												
0	2.8	2.7	2.5	2.8	3.7	4.1	3.7	4.3	6.2	5.5	6.8	7.0
10	2.5	2.7	2.5	2.8	3.6	4.2	3.7	4.4	6.5	6.9	8.8	9.1
100	2.3	2.7	2.4	2.7	3.5	4.1	3.6	4.1	5.7	6.4	7.9	8.3
1000	2.1	2.6	2.1	2.6	3.3	3.9	3.3	4.4	4.8	5.5	6.3	7.4
Collocated grid (2D)												
0	2.9	3.9	3.0	4.2	4.1	4.9	4.2	4.9	5.5	5.8	7.9	7.9
10	2.9	3.9	2.8	4.2	4.2	4.8	4.2	4.9	6.6	7.2	10.6	10.5
100	2.7	3.6	2.6	3.8	4.2	4.8	4.1	4.7	6.1	6.8	9.3	9.3
1000	2.3	3.1	2.4	3.3	3.9	4.7	4.0	4.6	5.1	5.8	7.9	8.3
Collocated grid (3D)												
0	9.4	5.1	8.5	4.8	8.3	9.5	8.1	9.7	6.6	7.2	8.8	9.7
10	7.4	4.4	7.4	4.6	8.1	9.6	8.0	9.7	7.4	8.4	11.0	11.9
100	5.2	4.3	5.2	4.4	7.7	9.4	7.7	9.3	6.2	7.2	9.0	10.5
1000	6.5	4.0	6.5	3.9	7.1	8.7	7.1	8.5	4.7	5.7	6.9	8.2

Table 4,[4] where one sees that the methods presented in this paper are significantly faster despite a higher cost per iteration. Thus, for the largest tested size, AMG applied to the right-hand transformed matrix is between 1.4 and 2.5 times faster than MINRES with block diagonal preconditioning (using the most cost effective variant which is finally the one with a single AMG application for the diagonal blocks).

Finally, we also tested two mixed finite element discretizations of stationary Stokes problems on unstructured grids. The first one is P_{2^*}/P_{-1} based on Crouzeix–Raviart elements [13], as described in [16]; the domain is the unit square, and a zoom on the central part of the mesh is displayed on Figure 2 (left). The second discretization is P_1/P_1 stabilized according to the method in [4]; the domain is the unit cube, and Figure 2 (right) offers a view on a cut of the mesh displaying one eighth of it.

For these matrices, there more are nonzero entries per row in B, making the two-sided transformation memory consuming and even less competitive. On the other hand, the block diagonal preconditioning method requires the pressure mass matrix, which was not provided. Hence, only right-hand transformation is tested here. Results are reported in Table 5, where we also reproduce the results obtained with finite difference discretizations (largest sizes). One sees that the method reaches similar efficiency for the mixed finite elements, except that the ratio between the numbers of nonzero entries in the transformed and original matrices ($\mathrm{rat}_{\mathrm{Tr}}$) is here somehow larger. As a consequence, the time needed to solve the system per nonzero entry *in the original matrix* is slightly less than twice the time needed for finite difference discretizations (while the time per unknown is significantly increased, on account of the much larger number of nonzero entries per row).

[4]Timings are reported for a standard Linux workstation equipped with Intel i5-4570 @ 3.20 GHz processor and 32 Gb DDR RAM memory; solvers were called from the Matlab environment, but all computationally intensive routines are written in Fortran and have been compiled with the GNU compiler (gfortran).

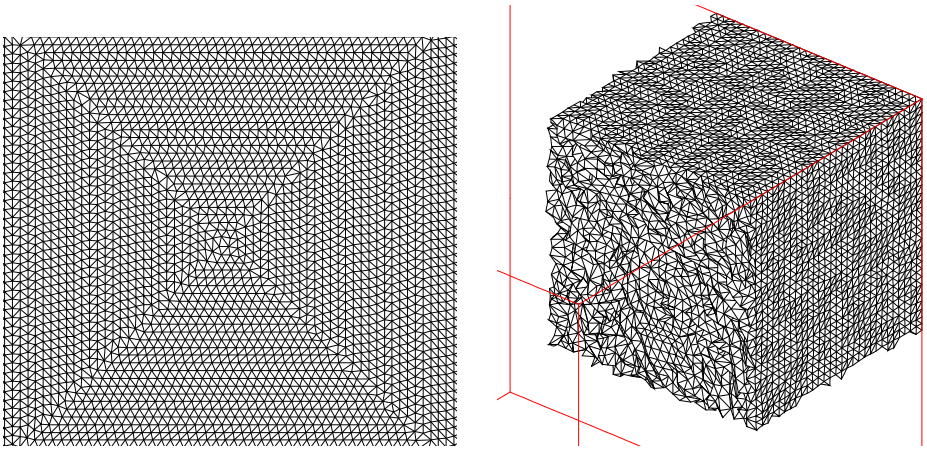

FIG. 2. *Zoom on the central part of the 2D unstructured mesh (left), and view of a cut of the unstructured 3D mesh, displaying one eighth of it.*

TABLE 5

Results with the right-hand transformation for the stationary problems on structured and unstructured meshes; "tm" refers to the total solution time and is reported in microsecond, either per unknown ($\frac{tm}{n}$) or per nonzero entry in the original matrix ($\frac{tm}{nnz}$).

						GCR(10)			GCR(30)		
	$\frac{n}{10^6}$	$\frac{n}{nnz}$	rat_{Tr}	$\mathcal{C}_A$	$\mathcal{C}_G$	it	$\frac{\text{tm}}{n}$	$\frac{\text{tm}}{nnz}$	it	$\frac{\text{tm}}{n}$	$\frac{\text{tm}}{nnz}$
MAC(2D)	3.1	6.0	1.9	2.3	2.5	17	2.7	0.44	17	2.8	0.47
Coll.(2D)	3.1	7.7	1.5	2.4	2.1	27	3.9	0.51	26	4.2	0.55
Coll.(3D)	3.5	9.9	1.7	1.8	3.1	17	5.1	0.51	17	4.8	0.49
P_{2*}/P_{-1}(2D)	1.3	17.3	3.4	1.5	4.9	22	14.3	0.83	22	14.5	0.84
P_1/P_1(3D)	0.8	36.3	2.5	1.8	4.6	23	32.3	0.89	22	32.1	0.88

7. Conclusions. We have shown that monolithic AMG methods can be successfully applied to solve discrete Stokes equations, using the standard unknown-based coarsening approach in which the prolongation operator is set up by considering separately the different types of unknowns. Two conditions, however, are to be satisfied. First, the AMG method should not be applied to the linear system stemming from the discretization, but to an equivalent system obtained through a simple algebraic transformation. Second, when more than two levels are needed, plain aggregation-based AMG has to be preferred, because the induced coarse level matrices are better suited to the recursive application of the method. When both these requirements are met, monolithic AMG appears both robust and cost effective with respect to state-of-the-art block preconditioning.

Regarding future research, one may observe that the algebraic transformations considered in this work can be also applied to linearized Navier–Stokes problems. Hence it is worth investigating if the approach can be successful in this context as well.

Acknowledgments. Stephane Gounand is gratefully acknowledged for his comments and having provided the problems with unstructured meshes. Suggestions by anonymous referees are also acknowledged.

REFERENCES

[1] J. H. Adler, T. R. Benson, and S. P. MacLachlan, *Preconditioning a mass-conserving discontinuous Galerkin discretization of the Stokes equations*, Numer. Linear Algebra Appl., (2016), http://dx.doi.org/10.1002/nla.2047.

[2] A. Ben Israel and T. N. E. Greville, *Generalized Inverses: Theory and Applications*, Wiley, New York, 1974.

[3] M. Benzi, G. H. Golub, and J. Liesen, *Numerical solution of saddle point problems*, Acta Numer., 14 (2005), pp. 1–137.

[4] P. B. Bochev, C. R. Dohrmann, and M. D. Gunzburger, *Stabilization of low-order mixed finite elements for the Stokes equations*, SIAM J. Numer. Anal., 44 (2006), pp. 82–101.

[5] D. Braess, *Towards algebraic multigrid for elliptic problems of second order*, Computing, 55 (1995), pp. 379–393.

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

[7] A. Brandt, *Rigorous quantitative analysis of multigrid,* I. *Constant coefficients two-level cycle with L_2-norm*, SIAM J. Numer. Anal., 31 (1994), pp. 1695–1730.

[8] A. Brandt and O. Livne, *Multigrid Techniques:* 1984 *Guide with Applications to Fluid Dynamics*, Classics Appl. Math. 67, rev. ed., SIAM, Philadelphia, 2011.

[9] A. Brandt, S. F. McCormick, and J. W. Ruge, *Algebraic multigrid (AMG) for sparse matrix equations*, in Sparsity and Its Application, D. J. Evans, ed., Cambridge University Press, Cambridge, 1984, pp. 257–284.

[10] J. Cahouet and J.-P. Chabard, *Some fast 3D finite element solvers for the generalized Stokes problem*, Internat. J. Numer. Methods Fluids, 8 (1988), pp. 869–895.

[11] Z.-H. Cao, *On the convergence of general stationary linear iterative methods for singular linear systems*, SIAM J. Matrix Anal. Appl., 29 (2008), pp. 1382–1388.

[12] T. Clees, *AMG Strategies for PDE Systems with Applications in Industrial Semiconductor Simulation*, dissertation, Mathematisch-Naturwissenschaftlichen Fakultät, Universität Köln, Köln, Germany, 2005.

[13] M. Crouzeix and P.-A. Raviart, *Conforming and nonconforming finite element methods for solving the stationary Stokes equations* I, ESAIM Math. Model. Numer. Anal., 7 (1973), pp. 33–75.

[14] J. Dendy, *Black box multigrid for periodic and singular problems*, Appl. Math. Comput., 25 (1988), pp. 1–10.

[15] S. C. Eisenstat, H. C. Elman, and M. H. Schultz, *Variational iterative methods for nonsymmetric systems of linear equations*, SIAM J. Numer. Anal., 20 (1983), pp. 345–357.

[16] H. Elman, D. Silvester, and A. Wathen, *Finite Elements and Fast Iterative Solvers*, Oxford University Press, Oxford, 2005.

[17] K. Hayami and M. Sugihara, *A geometric view of Krylov subspace methods on singular systems*, Numer. Linear Algebra Appl., 18 (2011), pp. 449–469.

[18] R. Horn and C. Johnson, *Matrix Analysis*, 2nd ed., Cambridge University Press, New York, 2013.

[19] J. Linden, B. Steckel, and K. Stüben, *Parallel multigrid solution of the Navier-Stokes equations on general 2D domains*, Parallel Comput., 7 (1988), pp. 461–475.

[20] B. Metsch, *Algebraic Multigrid (AMG) for Saddle Point Systems*, dissertation, Institut für Numerische Simulation, Universität Bonn, Bonn, Germany, 2013.

[21] A. C. Muresan and Y. Notay, *Analysis of aggregation-based multigrid*, SIAM J. Sci. Comput., 30 (2008), pp. 1082–1103.

[22] A. Napov and Y. Notay, *An algebraic multigrid method with guaranteed convergence rate*, SIAM J. Sci. Comput., 34 (2012), pp. A1079–A1109.

[23] A. Napov and Y. Notay, *Algebraic multigrid for moderate order finite elements*, SIAM J. Sci. Comput., 36 (2014), p. A1678–A1707.

[24] Y. Notay, *AGMG*: Interative Solution with AGgregation - Based algebraic MultiGrid, http://homepages.ulb.ac.be/~ynotay/AGMG.

[25] Y. Notay, *Flexible conjugate gradients*, SIAM J. Sci. Comput., 22 (2000), pp. 1444–1460.

[26] Y. Notay, *A robust algebraic multilevel preconditioner for non symmetric M-matrices*, Numer. Linear Algebra Appl., 7 (2000), pp. 243–267.

[27] Y. Notay, *An aggregation-based algebraic multigrid method*, Electron. Trans. Numer. Anal., 37 (2010), pp. 123–146.

[28] Y. Notay, *Algebraic analysis of two-grid methods: The nonsymmetric case*, Numer. Linear Algebra Appl., 17 (2010), pp. 73–96.

[29] Y. Notay, *Aggregation-based algebraic multigrid for convection-diffusion equations*, SIAM J. Sci. Comput., 34 (2012), pp. A2288–A2316.

[30] Y. Notay, *A new analysis of block preconditioners for saddle point problems*, SIAM J. Matrix Anal. Appl., 35 (2014), pp. 143–173.
[31] Y. Notay, *Algebraic two-level convergence theory for singular systems*, SIAM J. Matrix Anal. Appl., 37 (2016), pp. 1419–1439.
[32] Y. Notay, *A new multigrid approach for Stokes problems*, Numer. Math., 132 (2016), pp. 51–84.
[33] Y. Notay and P. S. Vassilevski, *Recursive Krylov-based multigrid cycles*, Numer. Linear Algebra Appl., 15 (2008), pp. 473–487.
[34] M. A. Olshanskii, J. Peters, and A. Reusken, *Uniform preconditioners for a parameter dependent saddle point problem with application to generalized Stokes interface equations*, Numer. Math., 105 (2006), pp. 159–191.
[35] C. Oosterlee and F. Gaspar, *Multigrid relaxation methods for systems of saddle point type*, Appl. Numer. Math., 58 (2008), pp. 1933–1950.
[36] J. W. Ruge and K. Stüben, *Algebraic multigrid*, in Multigrid Methods, S. F. McCormick, ed., Front. Appl. Math. 3, SIAM, Philadelphia, 1987, pp. 73–130.
[37] R. Stevenson, *On the Validity of Local Mode Analysis of Multi-grid Methods*, dissertation, Utrecht University, Utrecht, The Netherlands, 1990.
[38] K. Stüben, *An introduction to algebraic multigrid*, in U. Trottenberg, C. W. Oosterlee, and A. Schüller, Multigrid, Academic Press, London, 2001, pp. 413–532.
[39] U. Trottenberg, C. W. Oosterlee, and A. Schüller, *Multigrid*, Academic Press, London, 2001.
[40] S. Vanka, *Block-implicit multigrid solution of Navier-Stokes equations in primitive variables*, J. Comput. Phys., 65 (1986), pp. 138–158.
[41] P. Vaněk, J. Mandel, and M. Brezina, *Algebraic multigrid based on smoothed aggregation for second and fourth order elliptic problems*, Computing, 56 (1996), pp. 179–196.
[42] M. Wabro, *Algebraic Multigrid Methods for the Numerical Solution of the Incompressible Navier-Stokes Equations*, dissertation, Johannes Kepler Universität Linz, Linz, Austria, 2003.
[43] M. Wabro, *AMGe—coarsening strategies and application to the Oseen equations*, SIAM J. Sci. Comput., 27 (2006), pp. 2077–2097.
[44] R. Webster, *Stability of algebraic multigrid for Stokes problems*, Internat. J. Numer. Methods Fluids, 71 (2013), pp. 488–505.
[45] R. Webster, *Stabilisation of AMG Solvers for Saddle-Point Stokes Problems*, manuscript.
[46] P. Wesseling, *Principles of Computational Fluid Dynamics*, Springer Ser. Comput. Math. 29, Springer, Berlin, 1992.
[47] G. Wittum, *Multi-grid methods for Stokes and Navier-Stokes equations*, Numer. Math., 54 (1989), pp. 543–563.