

ABSOLUTE VALUE PRECONDITIONING FOR SYMMETRIC INDEFINITE LINEAR SYSTEMS*

EUGENE VECHARYNSKI[†] AND ANDREW V. KNYAZEVA[‡]

Abstract. We introduce a novel strategy for constructing symmetric positive definite (SPD) preconditioners for linear systems with symmetric indefinite matrices. The strategy, called absolute value preconditioning, is motivated by the observation that the preconditioned minimal residual method with the inverse of the absolute value of the matrix as a preconditioner converges to the exact solution of the system in at most two steps. Neither the exact absolute value of the matrix nor its exact inverse are computationally feasible to construct in general. However, we provide a practical example of an SPD preconditioner that is based on the suggested approach. In this example we consider a model problem with a shifted discrete negative Laplacian and suggest a geometric multigrid (MG) preconditioner, where the inverse of the matrix absolute value appears only on the coarse grid, while operations on finer grids are based on the Laplacian. Our numerical tests demonstrate practical effectiveness of the new MG preconditioner, which leads to a robust iterative scheme with minimalist memory requirements.

Key words. preconditioning, linear system, preconditioned minimal residual method, polar decomposition, matrix absolute value, multigrid, polynomial filtering

AMS subject classifications. 15A06, 65F08, 65F10, 65N22, 65N55

DOI. 10.1137/120886686

1. Introduction. Large sparse symmetric and indefinite systems arise in a variety of applications. For example, in the form of saddle point problems, such systems result from mixed finite element discretizations of underlying differential equations of fluid and solid mechanics; see, e.g., [3] and references therein. In acoustics, large sparse symmetric indefinite systems are obtained after discretizing the Helmholtz equation for certain media types and boundary conditions. Often the need to solve symmetric indefinite problems comes as an auxiliary task within other computational routines, such as the inner step in interior point methods in linear and nonlinear optimization [3, 26] or solution of the correction equation in the Jacobi–Davidson method [36] for a symmetric eigenvalue problem.

We consider an iterative solution of a linear system $Ax = b$, where the matrix A is real nonsingular and symmetric indefinite, i.e., the spectrum of A contains both positive and negative eigenvalues. In order to improve the convergence, we introduce a *preconditioner* T and formally replace $Ax = b$ by the *preconditioned system* $TAx = Tb$. If T is properly chosen, an iterative method for this system can exhibit a better convergence behavior compared to a scheme applied to $Ax = b$. Neither the preconditioner T nor the *preconditioned matrix* TA is normally explicitly computed.

If T is *not* symmetric positive definite (SPD), then TA , in general, is not symmetric with respect to any inner product [29, Theorem 15.2.1]. Thus, the introduction

*Submitted to the journal's Methods and Algorithms for Scientific Computing section August 1, 2012; accepted for publication (in revised form) January 2, 2013; published electronically March 14, 2013. This work is partially supported by National Science Foundation grant 1115734.

<http://www.siam.org/journals/sisc/35-2/88668.html>

[†]Computational Research Division, Lawrence Berkeley National Laboratory, Berkeley, CA 94720 (eugene.vecharynski@gmail.com).

[‡]Department of Mathematical and Statistical Sciences, University of Colorado, Denver, CO 80217-3364, and Mitsubishi Electric Research Laboratories, Cambridge, MA 02139 (andrew.knyazev@ucdenver.edu, <http://www.merl.com/people/?user=knyazev>, <http://math.ucdenver.edu/~aknyazevs/>).

of a non-SPD preconditioner replaces the original *symmetric* problem $Ax = b$ by a generally *nonsymmetric* $TAx = Tb$. Specialized methods for *symmetric* linear systems are no longer applicable to the preconditioned problem and must be replaced by iterative schemes for *nonsymmetric* linear systems, e.g., GMRES or GMRES(m) [35], Bi-CGSTAB [41], and QMR [17].

The approach based on the choice of a non-SPD preconditioner, which leads to solving a nonsymmetric problem, has several disadvantages. First, no short-term recurrent scheme that delivers an *optimal* Krylov subspace method is typically available for a nonsymmetric linear system [15]. In practice, this means that implementations of the optimal methods (e.g., GMRES) require an increasing amount of work and storage at every new step and hence are often computationally expensive.

Second, the convergence behavior of iterative methods for nonsymmetric linear systems is not completely understood. In particular, the convergence may not be characterized in terms of reasonably accessible quantities, such as the spectrum of the preconditioned matrix; see the corresponding results for GMRES and GMRES(m) in [20, 43]. This makes it difficult to predict computational costs.

If T is chosen to be SPD, i.e., $T = T^* > 0$, then the matrix TA of the preconditioned linear system is symmetric with respect to the T^{-1} -inner product defined by $(u, v)_{T^{-1}} = (u, T^{-1}v)$ for any pair of vectors u and v . Here (\cdot, \cdot) denotes the Euclidean inner product $(u, v) = v^*u$, in which the matrices A and T are symmetric. Due to this symmetry preservation, system $TAx = Tb$ can be solved using an *optimal* Krylov subspace method that admits a *short-term recurrent* implementation, such as preconditioned MINRES (PMINRES) [14, 28]. Moreover, the convergence of the method can be fully estimated in terms of the spectrum of TA .

In light of the above discussion, the choice of an SPD preconditioner for a symmetric indefinite linear system can be regarded as natural and favorable, especially if corresponding non-SPD preconditioning strategies fail to provide convergence in a small number of iterations. We advocate the use of SPD preconditioning.

The question of constructing SPD preconditioners for symmetric indefinite systems has been widely studied in many applications. For saddle point problems, the block-diagonal SPD preconditioning has been addressed, e.g., in [16, 37, 44]. In [2], it was proposed to use an inverse of the negative Laplacian as an SPD preconditioner for indefinite Helmholtz problems. This approach was further extended in [25] by introducing a shift into the preconditioner. Another strategy was suggested in [18], primarily in the context of linear systems arising in optimization. It is based on the so-called Bunch–Parlett factorization [9].

We introduce here a different idea (originating from the PhD thesis [42]) of constructing SPD preconditioners that resemble the inverse of the absolute value of the coefficient matrix. Throughout, the absolute value of A is defined as a matrix function $|A| = V|\Lambda|V^*$, where $A = V\Lambda V^*$ is the eigenvalue decomposition of A . We are motivated by the observation that PMINRES with $|A|^{-1}$ as a preconditioner converges to the exact solution in at most two steps. We refer to the new approach as the *absolute value* (AV) preconditioning and call the corresponding preconditioners the AV preconditioners.

The direct approach for constructing an AV preconditioner is to approximately solve $|A|z = r$. However, $|A|$ is generally not available, which makes the application of standard techniques, such as incomplete factorizations and approximate inverses, problematic. The vector $|A|^{-1}r$ can also be found using matrix function computations, normally fulfilled by a Krylov subspace method [19, 23] or a polynomial

approximation [30, 31]. Our numerical experience shows that the convergence, with respect to the outer iterations, of a linear solver can be significantly improved with this approach, but the computational costs of approximating $f(A)r = |A|^{-1}r$ may be too high, i.e., much higher than the cost of matrix-vector multiplication with A .

Introduction of the *general concept* of the AV preconditioning is the main theoretical contribution of the present work. As a proof of concept example of the AV preconditioning, we use a geometric multigrid (MG) framework. To investigate applicability and practical effectiveness of the proposed idea, we choose a model problem resulting from discretization of a shifted Laplacian (Helmholtz operator) on a unit square with Dirichlet boundary conditions. The obtained linear system is real symmetric indefinite. We construct an MG AV preconditioner that, used in the PMINRES iteration, delivers an efficient computational scheme.

Let us remark that the same model problem has been considered in [4], where the authors utilize the coarse grid approximation to reduce the indefinite problem to the SPD system. Satisfactory results have been reported for small shifts, i.e., for slightly indefinite systems. However, the limitation of the approach lies in the requirement on the size of the coarse space, which should be chosen sufficiently large. As we show below, the MG AV preconditioner presented in this paper allows keeping the coarsest problem reasonably small, even if the shift is large.

The numerical solution of Helmholtz problems is an object of active research; see, e.g., [1, 6, 12, 13, 21, 27, 40]. A typical Helmholtz problem is approximated by a complex symmetric (non-Hermitian) system. The real symmetric case of the Helmholtz equation, considered in this paper, is less common. However, methods for complex problems are evidently applicable to our particular real case, which allows us to make numerical comparisons with known Helmholtz solvers.

We test several solvers, based on the inverted Laplacian and the standard MG preconditioning, to compare with the proposed AV preconditioning. In fact, the inverted (shifted) Laplacian preconditioning [2, 25] for real Helmholtz problems can be viewed as a special case of our AV preconditioning. In contrast to preconditioners in [18] relying on the Bunch–Parlett factorization, we show that the AV preconditioners can be constructed without any decompositions of the matrix, which is crucial for very large or matrix-free problems.

This paper is organized as follows. In section 2, we present and justify the general notion of an AV preconditioner. The rest of the paper deals with the question of whether AV preconditioners can be efficiently constructed in practice. In section 3, we give a positive answer by constructing an example of a geometric MG AV preconditioner for the model problem. The efficiency of this preconditioner is demonstrated in our numerical tests in section 4. We conclude in section 5.

2. AV preconditioning for symmetric indefinite systems. Given an SPD preconditioner T , we consider solving a linear system with the *preconditioned minimal residual method*, implemented in the form of the PMINRES algorithm [14, 28]. In the absence of round-off errors, at step i , the method constructs an approximation $x^{(i)}$ to the solution of $Ax = b$ of the form

$$(2.1) \quad x^{(i)} \in x^{(0)} + \mathcal{K}_i(TA, Tr^{(0)})$$

such that the residual vector $r^{(i)} = b - Ax^{(i)}$ satisfies the optimality condition

$$(2.2) \quad \|r^{(i)}\|_T = \min_{u \in \mathcal{A}\mathcal{K}_i(TA, Tr^{(0)})} \|r^{(0)} - u\|_T.$$

Here, $\mathcal{K}_i(TA, Tr^{(0)}) = \text{span} \{Tr^{(0)}, (TA)Tr^{(0)}, \dots, (TA)^{i-1}Tr^{(0)}\}$ is the Krylov subspace generated by the matrix TA and the vector $Tr^{(0)}$, the T -norm is defined by $\|v\|_T^2 = (v, v)_T$ for any v , and $x^{(0)}$ is the initial guess. Scheme (2.1)–(2.2) represents an *optimal* Krylov subspace method and the PMINRES implementation is based on a *short-term recurrence*. The conventional convergence rate bound for (2.1)–(2.2) can be found, e.g., in [14] and relies solely on the distribution of eigenvalues of TA .

The following trivial, but important, theorem regards $|A|^{-1}$ as an SPD preconditioner for a symmetric indefinite system.

THEOREM 2.1. *The preconditioned minimal residual method (2.1)–(2.2) with preconditioner $T = |A|^{-1}$ converges to the solution of $Ax = b$ in at most two steps.*

Theorem 2.1 implies that $T = |A|^{-1}$ is an *ideal SPD preconditioner*. Note that the theorem holds not only for the preconditioned minimal residual method (2.1)–(2.2), but for all methods where convergence is determined by the degree of the minimal polynomial of TA .

In practical situations, the computation of an *ideal* SPD preconditioner $T = |A|^{-1}$ is prohibitively costly. However, we show that it is possible to construct inexpensive SPD preconditioners that resemble $|A|^{-1}$ and can significantly accelerate the convergence of an iterative method.

DEFINITION 2.2. *We call an SPD preconditioner T for a symmetric indefinite linear system $Ax = b$ an AV preconditioner if it satisfies*

$$(2.3) \quad \delta_0(v, T^{-1}v) \leq (v, |A|v) \leq \delta_1(v, T^{-1}v) \quad \forall v$$

with constants $\delta_1 \geq \delta_0 > 0$ such that the ratio $\delta_1/\delta_0 \geq 1$ is reasonably small.

Let us remark that Definition 2.2 of the AV preconditioner is informal because no precise assumption is made about how small the ratio δ_1/δ_0 should be. It is clear from (2.3) that δ_1/δ_0 measures how well the preconditioner T approximates $|A|^{-1}$, up to a positive scaling. If A represents a hierarchy of mesh problems, then it is desirable that δ_1/δ_0 is independent of the problem size. In this case, if A is SPD, Definition 2.2 of the AV preconditioner is consistent with the well-known concept of spectrally equivalent preconditioning for SPD systems; see [10].

The following theorem provides bounds for eigenvalues of the preconditioned matrix TA in terms of the spectrum of $T|A|$. We note that T and A , and thus TA and $T|A|$, do not in general commute. Therefore, our spectral analysis cannot be based on a traditional matrix analysis tool, a basis of eigenvectors.

THEOREM 2.3. *Given a nonsingular symmetric indefinite $A \in \mathbb{R}^{n \times n}$ and an SPD $T \in \mathbb{R}^{n \times n}$, let $\mu_1 \leq \mu_2 \leq \dots \leq \mu_n$ be the eigenvalues of $T|A|$. Then eigenvalues $\lambda_1 \leq \dots \leq \lambda_p < 0 < \lambda_{p+1} \leq \dots \leq \lambda_n$ of TA are located in intervals*

$$(2.4) \quad \begin{array}{l} -\mu_{n-j+1} \leq \lambda_j \leq -\mu_{p-j+1}, \quad j = 1, \dots, p; \\ \mu_{j-p} \leq \lambda_j \leq \mu_j, \quad j = p+1, \dots, n. \end{array}$$

Proof. We start by observing that the absolute value of the Rayleigh quotient of the generalized eigenvalue problem $Av = \lambda|A|v$ is bounded by 1, i.e.,

$$(2.5) \quad |(v, Av)| \leq (v, |A|v) \quad \forall v \in \mathbb{R}^n.$$

Now, we recall that the spectra of matrices $T|A|$ and TA are given by the generalized eigenvalue problems $|A|v = \mu T^{-1}v$ and $Av = \lambda T^{-1}v$, respectively, and introduce the corresponding Rayleigh quotients

$$(2.6) \quad \psi(v) \equiv \frac{(v, |A|v)}{(v, T^{-1}v)}, \quad \phi(v) \equiv \frac{(v, Av)}{(v, T^{-1}v)}, \quad v \in \mathbb{R}^n.$$

Let us fix any index $j \in \{1, 2, \dots, n\}$ and denote by S an arbitrary subspace of \mathbb{R}^n such that $\dim(S) = j$. Since inequality (2.5) also holds on S , using (2.6) we write

$$(2.7) \quad -\psi(v) \leq \phi(v) \leq \psi(v), \quad v \in S.$$

Moreover, taking the maxima in vectors $v \in S$, and after that the minima in subspaces $S \in S^j = \{S \subseteq \mathbb{R}^n : \dim(S) = j\}$, of all parts of (2.7) preserves the inequalities, so

$$(2.8) \quad \min_{S \in S^j} \max_{v \in S} (-\psi(v)) \leq \min_{S \in S^j} \max_{v \in S} \phi(v) \leq \min_{S \in S^j} \max_{v \in S} \psi(v).$$

By the Courant–Fischer theorem (see, e.g., [24, 29]) for the Rayleigh quotients $\pm\psi(v)$ and $\phi(v)$ defined in (2.6), we conclude from (2.8) that

$$-\mu_{n-j+1} \leq \lambda_j \leq \mu_j.$$

Recalling that j has been arbitrarily chosen, we obtain the following bounds on the eigenvalues of TA :

$$(2.9) \quad \begin{array}{l} -\mu_{n-j+1} \leq \lambda_j < 0, & j = 1, \dots, p; \\ 0 < \lambda_j \leq \mu_j, & j = p + 1, \dots, n. \end{array}$$

Next, in order to derive nontrivial upper and lower bounds for the p negative and $n - p$ positive eigenvalues λ_j in (2.9), we use the fact that eigenvalues ξ_j and ζ_j of the generalized eigenvalue problems $|A|^{-1}v = \xi T v$ and $A^{-1}v = \zeta T v$ are the reciprocals of the eigenvalues of the problems $|A|v = \mu T^{-1}v$ and $Av = \lambda T^{-1}v$, respectively, i.e.,

$$(2.10) \quad 0 < \xi_1 = \frac{1}{\mu_n} \leq \xi_2 = \frac{1}{\mu_{n-1}} \leq \dots \leq \xi_n = \frac{1}{\mu_1}$$

and

$$(2.11) \quad \zeta_1 = \frac{1}{\lambda_p} \leq \dots \leq \zeta_p = \frac{1}{\lambda_1} < 0 < \zeta_{p+1} = \frac{1}{\lambda_n} \leq \dots \leq \zeta_n = \frac{1}{\lambda_{p+1}}.$$

Similar to (2.5),

$$|(v, A^{-1}v)| \leq (v, |A|^{-1}v) \quad \forall v \in \mathbb{R}^n.$$

Thus, we can use the same arguments as those following (2.5) to show that relations (2.7) and (2.8), with a fixed $j \in \{1, 2, \dots, n\}$, also hold for

$$(2.12) \quad \psi(v) \equiv \frac{(v, |A|^{-1}v)}{(v, Tv)}, \quad \phi(v) \equiv \frac{(v, A^{-1}v)}{(v, Tv)}, \quad v \in \mathbb{R}^n,$$

where $\psi(v)$ and $\phi(v)$ are now the Rayleigh quotients of the generalized eigenvalue problems $|A|^{-1}v = \xi T v$ and $A^{-1}v = \zeta T v$, respectively. The Courant–Fischer theorem for $\pm\psi(v)$ and $\phi(v)$ in (2.12) allows us to conclude from (2.8) that

$$-\xi_{n-j+1} \leq \zeta_j \leq \xi_j.$$

Given the arbitrary choice of j in the above inequality, by (2.10)–(2.11) we get the following bounds on the eigenvalues of TA :

$$(2.13) \quad \begin{aligned} -1/\mu_{p-j+1} &\leq 1/\lambda_j < 0, & j = 1, \dots, p; \\ 0 &< 1/\lambda_j \leq 1/\mu_{j-p}, & j = p+1, \dots, n. \end{aligned}$$

Combining (2.9) and (2.13), we obtain (2.4). \square

Theorem 2.3 suggests two useful implications given by the corresponding corollaries below. In particular, the following result describes $\Lambda(TA)$, i.e., the spectrum of the preconditioned matrix TA , in terms of δ_0 and δ_1 in (2.3).

COROLLARY 2.4. *Given a nonsingular symmetric indefinite $A \in \mathbb{R}^{n \times n}$, an SPD $T \in \mathbb{R}^{n \times n}$, and constants $\delta_1 \geq \delta_0 > 0$ satisfying (2.3), we have*

$$(2.14) \quad \Lambda(TA) \subset [-\delta_1, -\delta_0] \cup [\delta_0, \delta_1],$$

where $\Lambda(TA)$ is the spectrum of TA .

Proof. The proof follows directly from (2.3) and (2.4) with $j = 1, p, p+1, n$. \square

The next corollary shows that the presence of reasonably populated clusters of eigenvalues in the spectrum of $T|A|$ guarantees the occurrence of corresponding clusters in the spectrum of the preconditioned matrix TA .

COROLLARY 2.5. *Given a nonsingular symmetric indefinite $A \in \mathbb{R}^{n \times n}$ and an SPD $T \in \mathbb{R}^{n \times n}$, let $\mu_l \leq \mu_{l+1} \leq \dots \leq \mu_{l+k-1}$ be a sequence of k eigenvalues of $T|A|$, where $1 \leq l < l+k-1 \leq n$ and $\tau = |\mu_l - \mu_{l+k-1}|$. Then, if $k \geq p+2$, the $k-p$ positive eigenvalues $\lambda_{l+p} \leq \lambda_{l+p+1} \leq \dots \leq \lambda_{l+k-1}$ of TA are such that $|\lambda_{l+p} - \lambda_{l+k-1}| \leq \tau$. Also, if $k \geq (n-p)+2$, the $k-(n-p)$ negative eigenvalues $\lambda_{n-k-l+2} \leq \dots \leq \lambda_{p-l} \leq \lambda_{p-l+1}$ of TA are such that $|\lambda_{n-k-l+2} - \lambda_{p-l+1}| \leq \tau$.*

Proof. The proof follows directly from bounds (2.4). \square

Corollary 2.4 implies that the ratio $\delta_1/\delta_0 \geq 1$ of the constants from (2.3) measures the quality of the AV preconditioner T . Indeed, the convergence speed of the preconditioned minimal residual method is determined by the spectrum of TA , primarily by the intervals of the right-hand side of inclusion (2.14). Additionally, Corollary 2.5 prompts that a “good” AV preconditioner should ensure clusters of eigenvalues in the spectrum of $T|A|$. This implies the clustering of eigenvalues of the preconditioned matrix TA , which has a favorable effect on the convergence behavior of a polynomial iterative method, such as PMINRES.

In the next section, we construct an example of the AV preconditioner for a particular model problem. We apply the MG techniques.

3. MG AV preconditioning for a model problem. Let us consider the real boundary value problem

$$(3.1) \quad -\Delta u(x, y) - c^2 u(x, y) = f(x, y), \quad (x, y) \in \Omega = (0, 1) \times (0, 1), \quad u|_\Gamma = 0,$$

where $\Delta = \partial^2/\partial x^2 + \partial^2/\partial y^2$ is the Laplace operator and Γ denotes the boundary of Ω . Problem (3.1) is a particular instance of the Helmholtz equation with Dirichlet boundary conditions, where $c > 0$ is a wave number.

After introducing a uniform grid of size h in both directions and using the standard five-point finite-difference stencil to discretize continuous problem (3.1), one obtains the corresponding discrete problem

$$(3.2) \quad (L - c^2 I)x = b,$$

where $A \equiv L - c^2 I$ represents a discrete negative Laplacian L (later called ‘‘Laplacian’’), satisfying the Dirichlet boundary condition shifted by a scalar c^2 .

The common rule of thumb (see, e.g., [13, 22]) for discretizing (3.1) is

$$(3.3) \quad ch \leq \pi/5.$$

Below, we call (3.2) the *model problem*. We assume that the shift c^2 is different from any eigenvalue of the Laplacian and is greater than the smallest but less than the largest eigenvalue. Thus, the matrix $L - c^2 I$ is nonsingular symmetric indefinite. In the following subsection, we apply the idea of the AV preconditioning to construct an MG AV preconditioner for system (3.2).

While our main focus throughout the paper is on the two-dimensional (2D) problem (3.1), in order to simplify presentation of theoretical analysis, we also refer to the one-dimensional (1D) analogue

$$(3.4) \quad -u''(x) - c^2 u(x) = f(x), \quad u(0) = u(1) = 0.$$

The conclusions drawn from (3.4), however, remain qualitatively the same for the 2D problem of interest, which we test numerically.

3.1. Two-grid AV preconditioner. Along with the fine grid of mesh size h underlying problem (3.2), let us consider a coarse grid of mesh size $H > h$. We denote the discretization of the Laplacian on this grid by L_H , and I_H represents the identity operator of the corresponding dimension. We assume that the exact fine-level absolute value $|L - c^2 I|$ and its inverse are not computable, whereas the inverse of the coarse-level operator $|L_H - c^2 I_H|$ can be efficiently constructed. In the two-grid framework, we use the subscript H to refer to the quantities defined on the coarse grid. No subscript is used for denoting the fine grid quantities.

While $|L - c^2 I|$ is not available, let us assume that we have its SPD approximation B , i.e., $B \approx |L - c^2 I|$ and $B = B^* > 0$. The operator B can be given in the explicit matrix form or through the action on a vector. We suggest the following general scheme as a two-grid AV preconditioner for model problem (3.2).

ALGORITHM 3.1 (the two-grid AV preconditioner). *Input:* r , $B \approx |L - c^2 I|$. *Output:* w .

1. Presmoothing. *Apply ν smoothing steps, $\nu \geq 1$:*

$$(3.5) \quad w^{(i+1)} = w^{(i)} + M^{-1}(r - Bw^{(i)}), \quad i = 0, \dots, \nu - 1, \quad w^{(0)} = 0,$$

where M defines a smoother. Set $w^{pre} = w^{(\nu)}$.

2. Coarse grid correction. *Restrict (R) $r - Bw^{pre}$ to the coarse grid, apply $|L_H - c^2 I_H|^{-1}$, and prolongate (P) to the fine grid. This delivers the coarse grid correction, which is added to w^{pre} :*

$$(3.6) \quad w_H = |L_H - c^2 I_H|^{-1} R(r - Bw^{pre}),$$

$$(3.7) \quad w^{cgc} = w^{pre} + Pw_H.$$

3. Postsmoothing. *Apply ν smoothing steps:*

$$(3.8) \quad w^{(i+1)} = w^{(i)} + M^{-*}(r - Bw^{(i)}), \quad i = 0, \dots, \nu - 1, \quad w^{(0)} = w^{cgc},$$

where M and ν are the same as in step 1. Return $w = w^{post} = w^{(\nu)}$.

In (3.6) we assume that $|L_H - c^2 I_H|$ is nonsingular, i.e., c^2 is different from any eigenvalue of L_H . The presmoothing is defined by the nonsingular M , while the postsmoothing is delivered by M^* . Note that the (inverted) absolute value appears only on the coarse grid, while the fine grid computations are based on the approximation B .

It is immediately seen that if $B = |L - c^2 I|$, Algorithm 3.1 represents a formal two-grid cycle [8, 39] for system

$$(3.9) \quad |L - c^2 I| z = r.$$

Note that the introduced scheme is rather general in that different choices of approximations B and smoothers M lead to different preconditioners. We address these choices in more detail in the following subsections.

It can be verified that the AV preconditioner given by Algorithm 3.1 implicitly constructs a mapping $r \mapsto w = T_{tg} r$, where the operator is

$$(3.10) \quad T_{tg} = (I - M^{-*} B)^\nu P |L_H - c^2 I_H|^{-1} R (I - B M^{-1})^\nu + F$$

with $F = B^{-1} - (I - M^{-*} B)^\nu B^{-1} (I - B M^{-1})^\nu$. The fact that the constructed preconditioner $T = T_{tg}$ is SPD follows directly from the observation that the first term in (3.10) is SPD provided that $P = \alpha R^*$ for some nonzero scalar α , while the second term F is SPD if the spectral radii of $I - M^{-1} B$ and $I - M^{-*} B$ are less than 1. The latter condition requires the pre- and postsmoothing iterations (3.5) and (3.8) to represent convergent methods for $By = r$. Note that the above argument essentially repeats the one used to justify symmetry and positive definiteness of a preconditioner based on the standard two-grid cycle for an SPD system; see, e.g., [5, 38].

In this paper we consider two different choices of the approximation B . The first choice is given by $B = L$, i.e., it is suggested to approximate the absolute value $|L - c^2 I|$ by the Laplacian L . The second choice is delivered by $B = p_m(L - c^2 I)$, where p_m is a polynomial of degree at most m such that $p_m(L - c^2 I) \approx |L - c^2 I|$.

3.2. Algorithm 3.1 with $B = L$. If $B = L$, Algorithm 3.1 can be regarded as a step of a standard two-grid method [8, 39] applied to the Poisson equation

$$(3.11) \quad Ly = r,$$

modified by replacing the operator L_H by $|L_H - c^2 I_H|$ on the coarse grid. The question remains of whether the algorithm delivers a form of an approximate solve for absolute value problem (3.9) and hence is suitable for AV preconditioning of (3.2). To be able to answer this question, we analyze the propagation of the initial error $e_0^{AV} = |L - c^2 I|^{-1} r$ of (3.9) under the action of the algorithm.

We start by relating errors of (3.9) and (3.11).

LEMMA 3.1. *Given a vector w , consider errors $e^{AV}(w) = |L - c^2 I|^{-1} r - w$ and $e^P(w) = L^{-1} r - w$ for (3.9) and (3.11), respectively. Then*

$$(3.12) \quad e^{AV}(w) = e^P(w) + (c^2 I - W_p) L^{-1} |L - c^2 I|^{-1} r,$$

where $W_p = 2V_p |\Lambda_p| V_p^*$, V_p is the matrix of eigenvectors of $L - c^2 I$ corresponding to the p negative eigenvalues $\lambda_1 \leq \dots \leq \lambda_p < 0$, and $|\Lambda_p| = \text{diag}\{|\lambda_1|, \dots, |\lambda_p|\}$.

Proof. Observe that for any w ,

$$\begin{aligned} e^{AV}(w) &= |L - c^2 I|^{-1} r - w = |L - c^2 I|^{-1} r + (e^P(w) - L^{-1} r) \\ &= e^P(w) + |L - c^2 I|^{-1} L^{-1} (L - |L - c^2 I|) r. \end{aligned}$$

Denoting $A = L - c^2 I$, we use the expression $|A| = A - 2V_p \Lambda_p V_p^*$ to get (3.12). \square

Algorithm 3.1 transforms the initial error $e_0^P = L^{-1}r$ of (3.11) into

$$(3.13) \quad e^P = S_2^\nu K S_1^\nu e_0^P,$$

where $S_1 = I - M^{-1}L$ and $S_2 = I - M^{-*}L$ are pre- and postsmoothing operators, $K = I - P|L_H - c^2 I_H|^{-1}RL$ corresponds to the coarse grid correction step, and $e^P = L^{-1}r - w^{post}$. Denoting the error of absolute value system (3.9) after applying Algorithm 3.1 by $e^{Av} = |L - c^2 I|^{-1}r - w^{post}$ and observing that $e_0^P = |L - c^2 I|L^{-1}e_0^{Av}$, by (3.12)–(3.13) we obtain

$$(3.14) \quad e^{Av} = (S_2^\nu K S_1^\nu |L - c^2 I| + c^2 I - W_p) L^{-1} e_0^{Av}.$$

The last expression gives an explicit form of the desired error propagation operator, which we denote by G :

$$(3.15) \quad G = (S_2^\nu K S_1^\nu |L - c^2 I| + c^2 I - W_p) L^{-1}.$$

Below, as a smoother, we use a simple Richardson's iteration, i.e., $S_1 = S_2 = I - \tau L$, where τ is an iteration parameter. The restriction R is given by the full weighting and the prolongation P by the standard piecewise linear interpolation; see [8, 39].

At this point, in order to simplify further presentation, let us refer to the 1D analogue (3.4) of model problem (3.1). In this case, the matrix L is tridiagonal: $L = \text{tridiag}\{-1/h^2, 2/h^2, -1/h^2\}$. We assume that n , the number of interior grid nodes, is odd: $h = 1/(n+1)$. The coarse grid is then obtained by dropping the odd-numbered nodes. We denote the size of the coarse grid problem by $N = (n+1)/2 - 1$; $H = 1/(N+1) = 2h$. The tridiagonal matrix L_H denotes the discretization of the 1D Laplacian on the coarse level.

Recall that the eigenvalues of L are $\theta_j = \frac{4}{h^2} \sin^2 \frac{j\pi h}{2}$ with corresponding eigenvectors $v_j = \sqrt{2h} [\sin lj\pi h]_{l=1}^n$. Similarly, the eigenvalues of L_H are $\theta_j^H = \frac{4}{H^2} \sin^2 \frac{j\pi H}{2}$, and the coarse grid eigenvectors are denoted by $v_j^H = \sqrt{2H} [\sin lj\pi H]_{l=1}^N$. It is clear that operators $L - c^2 I$ and $L_H - c^2 I_H$ have the same sets of eigenvectors as L and L_H with eigenvalues $t_j = \theta_j - c^2$ and $t_j^H = \theta_j^H - c^2$, respectively.

Let $e_0^{Av} = \sum_{j=1}^n \alpha_j v_j$ be the expansion of the initial error in the eigenbasis of L . Since $e^{Av} = G e_0^{Av} = \sum_{j=1}^n \alpha_j (G v_j)$, we are interested in the action of the error propagation operator (3.15) on the eigenmodes v_j .

The action of the operators R and P on v_j and v_j^H , respectively, is well known; see, e.g., [8, pp. 80–81]. Thus, it is easy to obtain the following expression for $K v_j$:

$$(3.16) \quad K v_j = \begin{cases} \left(1 - c_j^4 \frac{\theta_j}{|t_j^H|}\right) v_j + s_j^2 c_j^2 \frac{\theta_j}{|t_j^H|} v_{n+1-j}, & j = 1, \dots, N, \\ v_j, & j = N+1, \\ \left(1 - c_j^4 \frac{\theta_j}{|t_{n+1-j}^H|}\right) v_j + s_j^2 c_j^2 \frac{\theta_j}{|t_{n+1-j}^H|} v_{n+1-j}^H, & j = N+2, \dots, n. \end{cases}$$

Here, $c_j = \cos \frac{j\pi h}{2}$ and $s_j = \sin \frac{j\pi h}{2}$. Since v_j are the eigenvectors of $S_1 = S_2 = I - \tau L$, $L - c^2 I$, L^{-1} , and W_p , (3.15) leads to explicit expressions for $G v_j$.

THEOREM 3.2. Let $c^2 < \theta_{N+1} = 2/h^2$. Then the error propagation operator G in (3.15) acts on the eigenvectors v_j of the 1D Laplacian as follows:

$$(3.17) \quad Gv_j = \begin{cases} g_j^{(11)}v_j + g_j^{(12)}v_{n+1-j}, & j = 1, \dots, N, \\ g_j v_j, & j = N + 1, \\ g_j^{(21)}v_j + g_j^{(22)}v_{n+1-j}, & j = N + 2, \dots, n, \end{cases}$$

where

$$(3.18) \quad g_j^{(11)} = (1 - \tau\theta_j)^{2\nu} \left(1 - c_j^A \frac{\theta_j}{|t_j^H|} \right) \frac{|t_j|}{\theta_j} + \frac{c^2}{\theta_j} - \frac{\beta_j}{\theta_j},$$

$$(3.19) \quad g_j^{(12)} = (1 - \tau\theta_j)^\nu s_j^2 c_j^2 \frac{|t_j|}{|t_j^H|} (1 - \tau\theta_{n+1-j})^\nu,$$

$$(3.20) \quad g_j = (1 - \tau\theta_j)^{2\nu} \frac{|t_j|}{\theta_j} + \frac{c^2}{\theta_j},$$

$$(3.21) \quad g_j^{(21)} = (1 - \tau\theta_j)^{2\nu} \left(1 - c_j^A \frac{\theta_j}{|t_{n+1-j}^H|} \right) \frac{|t_j|}{\theta_j} + \frac{c^2}{\theta_j},$$

$$(3.22) \quad g_j^{(22)} = (1 - \tau\theta_j)^\nu s_j^2 c_j^2 \frac{|t_j|}{|t_{n+1-j}^H|} (1 - \tau\theta_{n+1-j})^\nu;$$

and

$$\beta_j = \begin{cases} 2(c^2 - \theta_j), & \theta_j < c^2, \\ 0, & \theta_j > c^2. \end{cases}$$

Theorem 3.2 implies that for relatively small shifts, Algorithm 3.1 with $B = L$ and a proper choice of τ and ν reduces the error of (3.9) in the directions of almost all eigenvectors v_j . In a few directions, however, the error may be amplified. These directions are given by the smooth eigenmodes associated with θ_j that are close to c^2 on the right, as well as with θ_j that are distant from c^2 on the left. The number of the latter, if any, is small if ch is sufficiently small and becomes larger as ch increases.

Indeed, let $\tau = h^2/3$ so that $|1 - \tau\theta_j| < 1$ for all j and $|1 - \tau\theta_j| < 1/3$ for $j > N$. This choice of the parameter provides the least uniform bound for $|1 - \tau\theta_j|$ that correspond to the oscillatory eigenmodes [34, p. 415]. It is then readily seen that (3.19) and (3.22) can be made arbitrarily small within a reasonably small number ν of smoothing steps. Similarly, (3.20) and (3.21) can be made arbitrarily close to $c^2/\theta_j < 1$. If $c^2 \ll \theta_{N+1}$, then c^2/θ_j in (3.20) and (3.21) is close to zero. Thus, Theorem 3.2 shows that for relatively small shifts, smoothing provides small values of (3.19)–(3.22) and hence damps the oscillatory part of the error. Note that the damping occurs even though the smoothing is performed with respect to (3.11), not (3.9).

Now let us consider (3.18). Theorem 3.2 shows that if c^2 is close to an eigenvalue θ_j^H of the coarse-level Laplacian, i.e., if $t_j^H \approx 0$, then the corresponding reduction coefficient (3.18) can be large. This means that Algorithm 3.1 with $B = L$ has a potential difficulty of amplifying the error in the directions of a few smooth eigenvectors.

A similar effect is known to appear for standard MG methods applied to Helmholtz type problems; see [7, 13]. Below, we analyze (3.18) in more detail.

Let $\theta_j > c^2$. Then, using the relation $\theta_j^H = c_j^2\theta_j$, we can write (3.18) as

$$g_j^{(11)} = (1 - \tau\theta_j)^{2\nu} \left(1 - c_j^2 \frac{1}{|1 - c^2/\theta_j^H|} \right) \left(1 - \frac{c^2}{\theta_j} \right) + \frac{c^2}{\theta_j}.$$

Here, it is easy to see that as $c^2/\theta_j^H \rightarrow 0$, $g_j^{(11)} \rightarrow (1 - \tau\theta_j)^{2\nu} s_j^2 < 1/2$, meaning that the smooth eigenmodes corresponding to θ_j away from c^2 on the right are well damped.

If $\theta_j < c^2$, then (3.18) takes the form

$$g_j^{(11)} = (1 - \tau\theta_j)^{2\nu} \left(\frac{c^2/\theta_j - c_j^2 - c_j^4}{c^2/\theta_j - c_j^2} \right) \left(\frac{c^2}{\theta_j} - 1 \right) + \left(2 - \frac{c^2}{\theta_j} \right).$$

Since $c_j^2 \in (1/2, 1)$, for any $c^2/\theta_j > 1$, we can obtain the bound

$$\frac{c^2/\theta_j - 2}{c^2/\theta_j - 1} \leq \frac{c^2/\theta_j - c_j^2 - c_j^4}{c^2/\theta_j - c_j^2} \leq \frac{c^2/\theta_j - 3/4}{c^2/\theta_j - 1/2}.$$

Additionally, $3^{-2\nu} < (1 - \tau\theta_j)^{2\nu} < 1$. Thus,

$$l_j < g_j^{(11)} < \frac{3(c^2/\theta_j) - 1}{4(c^2/\theta_j) - 2},$$

where $l_j = 0$ if $1 < c^2/\theta_j \leq 2$ and $l_j = 2 - c^2/\theta_j$ if $c^2/\theta_j > 2$.

The inequality implies that $|g_j^{(11)}| < 1$ for $1 < c^2/\theta_j \leq 3$, i.e., the algorithm reduces the error in the directions of several smooth eigenvectors associated with θ_j to the left of c^2 . At the same time, we note that as $c^2/\theta_j \rightarrow \infty$, $g_j^{(11)} \rightarrow \infty$, i.e., the smooth eigenmodes corresponding to θ_j that are distant from c^2 on the left can be amplified. Clearly, if ch is sufficiently small, then the number of such error components is not large (or is none) and grows as ch increases.

The above analysis shows that Algorithm 3.1 with $B = L$ indeed represents a solve for (3.9), where the solution is approximated everywhere, possibly except for a subspace of a small dimension. In the context of preconditioning, this translates into the fact that the preconditioned matrix has a spectrum clustered around 1 and -1 with a few outliers generated by the amplification of the smooth eigenmodes. If the shift is sufficiently small, the number of such outliers is not large, which only slightly delays the convergence of the outer PMINRES iterations and does not significantly affect the efficiency of the overall scheme.

3.3. Algorithm 3.1 with $B = p_m(L - c^2I)$. The analysis of the previous subsection suggests that the quality of Algorithm 3.1 with $B = L$ may deteriorate as ch increases. This result is not surprising, since for larger ch the relation $L \approx |L - c^2I|$ becomes no longer meaningful. Below we introduce a different approach for approximating the fine grid absolute value. In particular, we consider constructing *polynomial approximations* $B = p_m(L - c^2I)$, where $p_m(\lambda)$ is a polynomial of degree at most $m > 0$ such that $p_m(L - c^2I) \approx |L - c^2I|$.

Let us first refer to the ideal particular case, where $p_m(L - c^2I) = |L - c^2I|$. This can happen, e.g., if $p_m(\lambda)$ is an interpolating polynomial of $f(\lambda) = |\lambda|$ on the spectrum of $L - c^2I$, $m = n - 1$. In such a situation, Algorithm 3.1 with $B = p_m(L - c^2I)$ results in the following transformation of the initial error:

$$(3.23) \quad e^{\text{AV}} = \bar{S}_2^\nu \bar{K} \bar{S}_1^\nu e_0^{\text{AV}},$$

where $\bar{S}_1 = I - M^{-1}|L - c^2I|$ and $\bar{S}_2 = I - M^{-*}|L - c^2I|$ are pre- and postsmoothing operators and $\bar{K} = I - P|L_H - c^2I_H|^{-1}R|L - c^2I|$ corresponds to the coarse grid correction step. The associated error propagation operator is further denoted by \bar{G} ,

$$(3.24) \quad \bar{G} = \bar{S}_2^\nu \bar{K} \bar{S}_1^\nu.$$

For clarity, we again consider the 1D counterpart (3.4) of the model problem. As a smoother, we choose Richardson’s iteration with respect to absolute value system (3.9), i.e., $\bar{S}_1 = \bar{S}_2 = I - \tau|L - c^2I|$. It is important to note here that the eigenvalues $|t_j|$ of the absolute value operator are, in general, no longer ascendingly ordered with respect to j as is the case for θ_j ’s and t_j ’s. Moreover, in contrast to L and $L - c^2I$, the top part of the spectrum of $|L - c^2I|$ may be associated with both smooth and oscillatory eigenmodes. In particular, this means that Richardson’s iteration may fail to properly eliminate the oscillatory components of the error, which is an undesirable outcome of the smoothing procedure. To avoid this, we require that $|t_1| < t_{N+1}$. It is easy to verify that the latter condition is fulfilled if

$$(3.25) \quad ch < 1.$$

Note that (3.25) automatically holds if discretization rule (3.3) is enforced. Repeating the above argument for the 2D case also leads to (3.25).

Let the restriction and prolongation operators R and P be the same as in the previous subsection. Similar to (3.16), we obtain an explicit expression for the action of the coarse grid correction operator \bar{K} on eigenvectors v_j :

$$(3.26) \quad \bar{K}v_j = \begin{cases} \left(1 - c_j^4 \frac{|t_j|}{|t_j^H|}\right) v_j + s_j^2 c_j^2 \frac{|t_j|}{|t_j^H|} v_{n+1-j}, & j = 1, \dots, N, \\ v_j, & j = N + 1, \\ \left(1 - c_j^4 \frac{|t_j|}{|t_{n+1-j}^H|}\right) v_j + s_j^2 c_j^2 \frac{|t_j|}{|t_{n+1-j}^H|} v_{n+1-j}^H, & j = N + 2, \dots, n. \end{cases}$$

The following theorem is the analogue of Theorem 3.2.

THEOREM 3.3. *The error propagation operator \bar{G} in (3.24) acts on the eigenvectors v_j of the 1D Laplacian as follows:*

$$(3.27) \quad \bar{G}v_j = \begin{cases} \bar{g}_j^{(11)} v_j + \bar{g}_j^{(12)} v_{n+1-j}, & j = 1, \dots, N, \\ \bar{g}_j v_j, & j = N + 1, \\ \bar{g}_j^{(21)} v_j + \bar{g}_j^{(22)} v_{n+1-j}, & j = N + 2, \dots, n, \end{cases}$$

where

$$(3.28) \quad \bar{g}_j^{(11)} = (1 - \tau|t_j|)^{2\nu} \left(1 - c_j^4 \frac{|t_j|}{|t_j^H|}\right),$$

$$(3.29) \quad \bar{g}_j^{(12)} = (1 - \tau|t_j|)^\nu s_j^2 c_j^2 \frac{|t_j|}{|t_j^H|} (1 - \tau|t_{n+1-j}|)^\nu,$$

$$(3.30) \quad \bar{g}_j = (1 - \tau|t_j|)^{2\nu},$$

$$(3.31) \quad \bar{g}_j^{(21)} = (1 - \tau|t_j|)^{2\nu} \left(1 - c_j^4 \frac{|t_j|}{|t_{n+1-j}^H|}\right),$$

$$(3.32) \quad \bar{g}_j^{(22)} = (1 - \tau|t_j|)^\nu s_j^2 c_j^2 \frac{|t_j|}{|t_{n+1-j}^H|} (1 - \tau|t_{n+1-j}|)^\nu.$$

We conclude from Theorem 3.3 that in the ideal case where $p_m(L - c^2I) = |L - c^2I|$, Algorithm 3.1 with $B = p_m(L - c^2I)$ and a proper choice of τ and ν reduces the error

of system (3.9) in the directions of all eigenvectors v_j , possibly except for a few that correspond to θ_j close to the shift c^2 . Unlike in the case of Algorithm 3.1 with $B = L$, as ch grows, no amplified error components appear in the directions of eigenvectors associated with θ_j distant from c^2 on the left. This suggests that Algorithm 3.1 with $B = p_m(L - c^2I) = |L - c^2I|$ provides a more accurate solve for (3.9) with larger ch .

To see this, let us first assume that $\tau = h^2/(3 - c^2h^2)$. Since (3.25) implies that $|t_j| = t_j$ for $j > N$, this choice is known to give the smallest uniform bound on $|1 - \tau|t_j||$ corresponding to the oscillatory eigenmodes v_j , which is $|1 - \tau|t_j|| < 1/(3 - ch) < 1/2$ with the last inequality resulting from (3.25). Hence, coefficients (3.29)–(3.32) can be reduced within a reasonably small number ν of smoothing steps.

Next, we note that (3.28), which is not substantially affected by smoothing, can be large if c^2 is close to θ_j^H , i.e., if $t_j^H \approx 0$. At the same time, we can write (3.28) as

$$g_j^{(11)} = (1 - \tau|t_j|)^{2\nu} \left(1 - c_j^2 \left| 1 + \frac{c^2 s_j^2}{t_j^H} \right| \right),$$

which shows that $|g_j^{(11)}|$ approaches $(1 - \tau|t_j|)^{2\nu} s_j^2 < 1/2$ as $|t_j^H|$ increases, i.e., smooth error components associated with θ_j away from c^2 are well damped.

Thus, if used as a preconditioner, Algorithm 3.1 with $B = p_m(L - c^2I) = |L - c^2I|$ aims at clustering the spectrum of the preconditioned matrix around 1 and -1 , with a few possible outliers that result from the amplification of the smooth eigenmodes associated with θ_j close to c^2 . Unlike in the case where $B = L$, the increase of ch does not additionally amplify the smooth error components distant from c^2 on the left. Therefore, Algorithm 3.1 with $B = p_m(L - c^2I) = |L - c^2I|$ can be expected to provide a more accurate preconditioner for larger shifts.

Although our analysis targets the ideal but barely feasible case where $p_m(L - c^2I) = |L - c^2I|$, it motivates the use of *polynomial approximations* $p_m(L - c^2I) \approx |L - c^2I|$ and provides a theoretical insight into the superior behavior of such an option for larger ch . In the rest of this subsection we describe a method for constructing such polynomial approximations. Our approach is based on the finding that the problem is easily reduced to constructing *polynomial filters*.

We start by introducing the step function

$$h_\alpha(\lambda) = \begin{cases} 1, & \lambda \geq \alpha, \\ 0, & \lambda < \alpha, \end{cases}$$

where α is a real number, and noting that $\text{sign}(\lambda) = 2h_0(\lambda) - 1$, so that

$$(3.33) \quad |L - c^2I| = (2h_0(L - c^2I) - I) (L - c^2I).$$

Here $h_0(L - c^2I) = Vh_0(\Lambda)V^*$, where V is the matrix of eigenvectors of $L - c^2I$ and $h_0(\Lambda) = \text{diag}\{0, \dots, 0, 1, \dots, 1\}$ is obtained by applying the step function $h_0(\lambda)$ to the diagonal entries of the matrix Λ of the associated eigenvalues. Clearly the number of zeros on the diagonal of $h_0(\Lambda)$ equals the number of negative eigenvalues of $L - c^2I$.

Let $q_{m-1}(\lambda)$ be a polynomial of degree at most $(m - 1)$ such that $q_{m-1}(\lambda)$ approximates $h_0(\lambda)$ on the interval $[a, b]$, where a and b are the lower and upper bounds on the spectrum of $L - c^2I$, respectively. In order to construct an approximation $p_m(L - c^2I)$ of $|L - c^2I|$, we replace the step function $h_0(L - c^2I)$ in (3.33) by the polynomial $q_{m-1}(L - c^2I)$. Thus,

$$(3.34) \quad |L - c^2I| \approx p_m(L - c^2I) = (2q_{m-1}(L - c^2I) - I) (L - c^2I).$$

The matrix $L - c^2I$ is readily available on the fine grid. Therefore, we have reduced the problem of evaluating the polynomial approximation p_m of the absolute value operator to constructing a polynomial q_{m-1} that approximates the step function h_0 . More specifically, since Algorithm 3.1 can be implemented without the explicit knowledge of the matrix B , i.e., B can be accessed only through its action on a vector, we need to construct approximations of the form $q_{m-1}(L - c^2I)v$ to $h_0(L - c^2I)v$, where v is a given vector.

The task of constructing $q_{m-1}(L - c^2I)v \approx h_0(L - c^2I)v$ represents an instance of *polynomial filtering*, which is well known; see, e.g., [11, 32, 45]. In this context, due to the property of filtering out certain undesirable eigencomponents, the step function h_0 is called a *filter function*. The approximating polynomial q_{m-1} is referred to as a *polynomial filter*.

State-of-the-art polynomial filtering techniques such as [32] would first replace the discontinuous step function $h_0(\lambda)$ by a smooth approximation on $[a, b]$ and then approximate the latter by a polynomial in the least-squares sense. In this paper, we follow a simpler approach based on the direct approximation of $h_0(\lambda)$ using *Chebyshev polynomials* [30, 31]. The constructed polynomial q_{m-1} allows defining $q_{m-1}(L - c^2I)v \approx h_0(L - c^2I)v$ and hence $p_m(L - c^2I)v \approx |L - c^2I|v$. Thus, the entire procedure provides means to replace a matrix-vector product with the unavailable $|L - c^2I|$ by, essentially, a few multiplications with $L - c^2I$. As we further show, the degree m of the approximating polynomial can be kept reasonably low. Moreover, in the MG framework discussed in the next subsection, the algorithm has to be invoked only on sufficiently coarse grids.

3.4. The MG AV preconditioner. Now let us consider a hierarchy of $s + 1$ grids numbered by $l = s, s - 1, \dots, 0$, with the corresponding mesh sizes $\{h_l\}$ in decreasing order ($h_s = h$ corresponds to the finest grid and h_0 to the coarsest). For each level l we define the discretization $L_l - c^2I_l$ of the differential operator in (3.1), where L_l is the Laplacian on grid l and I_l is the identity of the same size.

In order to extend the two-grid AV preconditioner given by Algorithm 3.1 to the *multigrid*, instead of inverting the absolute value $|L_H - c^2I_H|$ in (3.6), we recursively apply the algorithm to the restricted vector $R(r - Bw^{pre})$. This pattern is then followed in the V-cycle fashion on all levels with the inversion of the absolute value of the shifted Laplacian on the coarsest grid. The matrix B on level l is denoted by B_l . Each B_l is assumed to be SPD and is expected to approximate $|L_l - c^2I_l|$.

In the previous subsections we have considered two choices of B for the two-grid preconditioner in Algorithm 3.1. In the MG framework, these choices give $B_l = L_l$ and $B_l = p_{m_l}(L_l - c^2I_l)$, where p_{m_l} is a polynomial of degree at most m_l on level l .

The advantage of the first option, $B_l = L_l$, is that it can be easily constructed and the application of B_l to a vector is inexpensive even if the size of the operator is very large. According to our analysis for the 1D model problem in subsection 3.2, the approach is suitable for ch_l sufficiently small. Typically this is a case for l corresponding to finer grids. However, ch_l increases with every new level. This may result in the deterioration of accuracy of the overall MG preconditioning scheme, unless the size of the coarsest level is kept sufficiently large.

The situation is different for the second option $B_l = p_{m_l}(L_l - c^2I_l)$. In this case, applications of B_l may be expensive on finer grids because they require a sequence of matrix-vector multiplications with large shifted Laplacian operators. However, on coarser levels, i.e., for larger ch_l , this is not restrictive because the involved operators are significantly decreased in size compared to the finest level. Additionally,

as suggested by the analysis in subsection 3.3, if $p_{m_l}(L_l - c^2 I_l)$ represent reasonable approximations of $|L_l - c^2 I_l|$ on levels l , one can expect a higher accuracy of the whole preconditioning scheme compared to the choice $B_l = L_l$.

Our idea is to combine the two options. Let $\delta \in (0, 1)$ be a “switching” parameter, where for finer grids $ch_l < \delta$. We choose

$$(3.35) \quad B_l = \begin{cases} L_l, & ch_l < \delta, \\ p_{m_l}(L_l - c^2 I_l), & ch_l \geq \delta. \end{cases}$$

The polynomials $p_{m_l}(L_l - c^2 I_l)$ are accessed through their action on a vector.

Summarizing our discussion, if started from the finest grid $l = s$, the following scheme gives the multilevel extension of the two-grid AV preconditioner defined by Algorithm 3.1. The subscript l is introduced to match quantities to the corresponding grid. We assume that the parameters δ , m_l , ν_l and the smoothers M_l are prespecified.

ALGORITHM 3.2 (AV-MG(r_l): the MG AV preconditioner). *Input* r_l . *Output* w_l .

1. Set B_l by (3.35).
2. Presmoothing. Apply ν_l smoothing steps, $\nu_l \geq 1$:

$$(3.36) \quad w_l^{(i+1)} = w_l^{(i)} + M_l^{-1}(r_l - B_l w_l^{(i)}), \quad i = 0, \dots, \nu_l - 1, \quad w_l^{(0)} = 0,$$

where M_l defines a smoother on level l . Set $w_l^{pre} = w_l^{(\nu)}$.

3. Coarse grid correction. Restrict $(R_{l-1}) r_l - B_l w_l^{pre}$ to the grid $l-1$, recursively apply AV-MG, and prolongate (P_l) back to the fine grid. This delivers the coarse grid correction added to w_l^{pre} :

$$(3.37) \quad w_{l-1} = \begin{cases} |L_0 - c^2 I_0|^{-1} R_0 (r_l - B_l w_l^{pre}), & l = 1, \\ AV-MG(R_{l-1} (r_l - B_l w_l^{pre})), & l > 1; \end{cases}$$

$$(3.38) \quad w_l^{cgc} = w_l^{pre} + P_l w_{l-1}.$$

4. Postsmoothing. Apply ν_l smoothing steps:

$$(3.39) \quad w_l^{(i+1)} = w_l^{(i)} + M_l^{-*}(r_l - B_l w_l^{(i)}), \quad i = 0, \dots, \nu_l - 1, \quad w_l^{(0)} = w_l^{cgc},$$

where M_l and ν_l are the same as in step 2. Return $w_l = w_l^{post} = w_l^{(\nu)}$.

The described MG AV preconditioner implicitly constructs a mapping denoted by $r \mapsto w = T_{mg} r$, where the operator $T = T_{mg}$ has the structure

$$(3.40) \quad T_{mg} = (I - M^{-*} B)^\nu P T_{mg}^{(s-1)} R (I - B M^{-1})^\nu + F$$

with F as in (3.10) and $T_{mg}^{(s-1)}$ defined according to the recursion

$$(3.41) \quad \begin{aligned} T_{mg}^{(l)} &= (I - M_l^{-*} B_l)^{\nu_l} P_l T_{mg}^{(l-1)} R_{l-1} (I - B_l M_l^{-1})^{\nu_l} + F_l, \\ T_{mg}^{(0)} &= |L_0 - c^2 I_0|^{-1}, \quad l = 1, \dots, s-1, \end{aligned}$$

where $F_l = B_l^{-1} - (I_l - M_l^{-*} B_l)^{\nu_l} B_l^{-1} (I_l - B_l M_l^{-1})^{\nu_l}$.

The structure of the multilevel preconditioner $T = T_{mg}$ in (3.40) is the same as that of the two-grid preconditioner $T = T_{lg}$ in (3.10) with $|L_H - c^2 I_H|^{-1}$ replaced by the recursively defined operator $T_{mg}^{(m-1)}$ in (3.41). Thus, the symmetry and positive

definiteness of $T = T_{mg}$ follow from the same property of the two-grid operator through relations (3.41), provided that $P_l = \alpha_l R_{l-1}^*$ and the spectral radii of $I_l - M_l^{-1}B_l$ and $I_l - M_l^{-*}B_l$ are less than 1 throughout the coarser levels. We remark that preconditioner (3.40)–(3.41) is nonvariable, i.e., it preserves the global optimality of PMINRES.

The simplest possible approach for computing w_0 in (3.37) is to explicitly construct $|L_0 - c^2 I_0|^{-1}$ through the full eigendecomposition of the coarse-level Laplacian and then apply it to $R_0(r_1 - B_1 w_1^{pre})$. An alternative approach is to determine w_0 as a solution of the linear system $(L_0 - c^2 I_0 + 2V_0|\Lambda_0|V_0^*)w_0 = R_0(r_1 - B_1 w_1^{pre})$, where V_0 is the matrix of eigenvectors associated with the negative eigenvalues of $L_0 - c^2 I_0$ contained in the corresponding diagonal matrix Λ_0 . In the latter case, the full eigendecomposition of L_0 is replaced by the *partial* eigendecomposition targeting negative eigenpairs, followed by a linear solve.

Since we use Richardson's iteration with respect to $p_{m_l}(L_l - c^2 I_l)$ as a smoother on coarser grids, as motivated by the discussion in subsection 3.3, the guidance for the choice of the coarsest grid is given by condition (3.25). More specifically, in the context of the standard coarsening procedure ($h_{l-1} = 2h_l$), we select hierarchies of grids satisfying $ch_l < 1$ for $l = s, \dots, 1$ and $ch_0 > 1$. As shown in the next section, even for reasonably large c^2 , the coarsest-level problems are small.

The parameter δ in (3.35) should be chosen to ensure the balance between computational costs and the quality of the MG preconditioner. In particular, if δ is reasonably large, then the choices of B_l are dominated by the option $B_l = L_l$, which is inexpensive but may not be suitable for larger shifts on coarser levels. On the other extreme, if δ is close to zero, then the common choice corresponds to $B_l = p_{m_l}(L_l - c^2 I_l)$, which provides a better preconditioning accuracy for larger shifts but may be too computationally intense on finer levels. In our numerical experiments, we keep $\delta \in [1/3, 3/4]$.

As we demonstrate in the next section, the degrees m_l of the occurring polynomials p_{m_l} should not be large, i.e., only a few matrix-vector multiplications with $L_l - c^2 I_l$ are required to obtain satisfactory approximations of absolute value operators. For properly chosen δ , these additional multiplications need to be performed on grids that are significantly coarser than the finest grid, i.e., the involved matrices $L_l - c^2 I_l$ are orders of magnitude smaller than the original fine grid operator. As confirmed by our numerical experiments, the overhead caused by the polynomial approximations appears to be marginal and does not much affect the computational cost of the overall preconditioning scheme.

4. Numerical experiments. This section presents a numerical study of the MG preconditioner in Algorithm 3.2. Our goal here is twofold. On the one hand, the reported numerical experiments serve as a proof of concept of the AV preconditioning described in section 2. In particular, we show that the AV preconditioners can be constructed at essentially the same cost as the standard preconditioning methods (MG in our case). On the other hand, we demonstrate that the MG AV preconditioner in Algorithm 3.2 combined with the optimal PMINRES iteration, in fact, leads to an efficient and economical computational scheme, further called MINRES-AV-MG, which outperforms several known competitive approaches for the model problem.

Let us briefly describe the alternative preconditioners used for our comparisons. Throughout, we use MATLAB for our numerical examples.

The inverted Laplacian preconditioner. This strategy, introduced in [2], is a representative of an SPD preconditioning for model problem (3.2), where the

TABLE 4.1

The largest problem sizes satisfying $ch_l \geq \delta$ for different values of the shift c^2 , “switching” parameters δ , and the standard coarsening scheme $h_{l-1} = 2h_l$. The last row ($\delta = 1$) corresponds to the sizes of the coarsest problems for different c^2 .

	$c^2 = 300$	$c^2 = 400$	$c^2 = 1500$	$c^2 = 3000$	$c^2 = 4000$
$\delta = 1/3$	961	961	3969	16129	16129
$\delta = 1/2$	961	961	3969	3969	3969
$\delta = 3/4$	225	225	961	3969	3969
$\delta = 1$	225	225	961	961	961

preconditioner is applied through solving systems $Lw = r$, i.e., $T = L^{-1}$. As has been previously discussed, for relatively small shifts c^2 , the Laplacian L constitutes a good SPD approximation of $|L - c^2I|$. In this sense, the choice $T = L^{-1}$ perfectly fits, as a special case, into the general concept of the AV preconditioning presented in section 2. We refer to PMINRES with $T = L^{-1}$ as MINRES-Laplace.

Usually, one wants to solve the system $Lw = r$ only approximately, i.e., use $T \approx L^{-1}$. This can be efficiently done, e.g., by applying the V-cycle of a standard MG method [8, 39]. In our tests, however, we perform the exact solves using the MATLAB “backslash” so that the reported results reflect the best possible convergence with the inverted Laplacian type preconditioning.

The indefinite MG preconditioner. We consider a standard V-cycle for problem (3.2). Formally, it can be obtained from Algorithm 3.2 by setting $B_l = L_l - c^2I_l$ on all levels and replacing the first equality in (3.37) by the linear solve with $L_0 - c^2I_0$. The resulting MG scheme is used as a preconditioner for restarted GMRES and for Bi-CGSTAB. We refer to these methods as GMRES(k)-MG and Bi-CGSTAB-MG, respectively; k denotes the restart parameter. A thorough discussion of the indefinite MG preconditioning for Helmholtz problems can be found, e.g., in [13].

In our tests, we consider the 2D model problem (3.2) corresponding to (3.1) discretized on the grid of size $h = 2^{-8}$ (the fine problem size $n = 65025$). The exact solution x^* and the initial guess x_0 are randomly chosen. The right-hand-side $b = (L - c^2I)x^*$, which allows evaluating the actual errors along the steps of an iterative method. All the occurring MG preconditioners are built upon the standard coarsening scheme (i.e., $h_{l-1} = 2h_l$), restriction is based on the full weighting, and prolongation is based on piecewise multilinear interpolation [8, 39].

Let us recall that Algorithm 3.2 requires setting a parameter δ to switch between $B_l = L$ and $p_{m_l}(L - c^2I)$ on different levels; see (3.35). Assuming standard coarsening, Table 4.1 presents the largest problem sizes corresponding to the condition $ch_l \geq \delta$ for a few values of δ and c^2 . In other words, given δ and c^2 , each cell of Table 4.1 contains the largest problem size for which the polynomial approximation of $|L_l - c^2I_l|$ is constructed. Unless otherwise explicitly stated, we set $\delta = 1/3$. Note that according to the discussion in subsection 3.4 (condition (3.25)), the row of Table 4.1 corresponding to $\delta = 1$ delivers the sizes n_0 of the coarsest problems for different shift values.

Table 4.1 shows that the coarsest problems remain relatively small even for large shifts. The polynomial approximations are constructed for coarser problems of significantly reduced dimensions, which in practical applications are negligibly small compared to the original problem size.

As a smoother on all levels of Algorithm 3.2 we use Richardson’s iteration, i.e., $M_l^{-1} \equiv \tau_l I_l$. On the finer levels, where $B_l = L_l$, we choose $\tau_l = h_l^2/5$ and $\nu_l = 1$. On the coarser levels, where $B_l = p_{m_l}(L_l - c^2I_l)$, we set $\tau_l = h_l^2/(5 - c^2h_l^2)$ and $\nu_l = 5$.

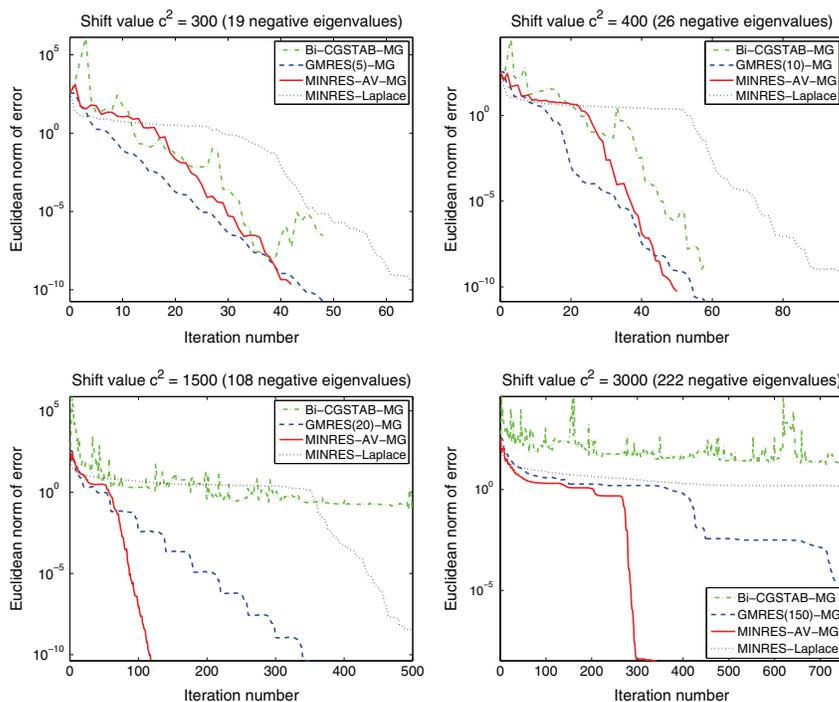


FIG. 4.1. Comparison of several preconditioned schemes; $n = 65025$.

Similar to the 1D case considered in subsections 3.2 and 3.3, both choices of τ_l correspond to the optimal smoothing of oscillatory eigenmodes with respect to the 2D operators L_l and $|L_l - c^2 I_l|$, respectively [8, 39]. Since $p_{m_l}(L_l - c^2 I_l)$ only approximates $|L_l - c^2 I_l|$ in practice, the choice of τ_l on the coarser grids is likely to be not optimal. Therefore, for $ch_l \geq \delta$, we have increased the number of smoothing steps to 5. In all tests, the degrees m_l of polynomials are set to 10. The intervals containing $\Lambda(L_l - c^2 I_l)$, required for evaluating p_{m_l} , are $[-c^2, 8h_l^{-2} - c^2]$. The inverted coarse grid absolute value $|L_0 - c^2 I_0|^{-1}$ is constructed by the full eigendecomposition.

In Figure 4.1, we compare MINRES-AV-MG with the above introduced alternative preconditioned schemes for the model problem. Each plot corresponds to a different shift value. The restart parameter k varies for all runs of GMRES(k)-MG, increasing (left to right and top to bottom) as c^2 grows from 300 to 3000. In our tests, the size n_0 of the coarsest problem in Algorithm 3.2 is 225 (Figure 4.1, top) and 961 (Figure 4.1, bottom); see Table 4.1 with $\delta = 1$. The same n_0 is used for the MG preconditioner in the corresponding runs of GMRES(k)-MG and Bi-CGSTAB-MG.

Figure 4.1 shows that MINRES-AV-MG noticeably outperforms PMINRES with the inverted Laplacian preconditioner. For smaller shifts ($c^2 = 300, 400$), MINRES-AV-MG is comparable, in terms of the iteration count, to GMRES(k)-MG and Bi-CGSTAB-MG; $k = 5, 10$. For larger shifts ($c^2 = 1500, 3000$); however, MINRES-AV-MG provides a superior convergence behavior. In particular, the scheme exhibits faster convergence than GMRES(k)-MG under less demanding storage requirements, while Bi-CGSTAB-MG fails to converge ($k = 50, 150$).

If the polynomial approximations in Algorithm 3.2 appear only on sufficiently coarse grids and the size n_0 of the coarsest problem is relatively small, then the additional costs introduced by the coarser grid computations of the AV MG preconditioner

TABLE 4.2
Time comparison of the MG preconditioned schemes in Figure 4.1.

	$c^2 = 300$	$c^2 = 400$	$c^2 = 1500$	$c^2 = 3000$
t_B/t_{AV}	1.1	1.1	—	—
t_G/t_{AV}	1.4	1.3	2.6	1.9

are negligible relative to the costs of operations on the finer grids, which are the same as in the standard V-cycle for the indefinite problem. This means that the complexity of Algorithm 3.2 is similar to that of the MG preconditioner in GMRES(k)-MG and Bi-CGSTAB-MG.

To be more precise, in the tests reported in Figure 4.1, bottom, a single application of Algorithm 3.2 has required 15–20% more time than the indefinite MG preconditioner, even though the polynomial approximations in Algorithm 3.2 have been constructed for problem sizes as large as 16129 if $c^2 = 3000$, which is of the same order as the original problem size n . For larger problem sizes, the time difference becomes negligible. For example, if $h = 2^{-9}$ ($n = 261121$), Algorithm 3.2 results in only 5% time increase, and the relative time difference becomes indistinguishable for smaller h . The application of all the MG preconditioners in Figure 4.1, top, required essentially the same time.

The above discussion suggests that our numerical comparison, based on the number of iterations, is representative. Additionally, in Table 4.2 we provide the time comparisons for the MG preconditioned schemes. In particular, we measure the actual time required by the runs of MINRES-AV-MG (t_{AV}), GMRES(k)-MG (t_G), and Bi-CGSTAB-MG (t_B) in Figure 4.1 and report the speed-ups.

We have observed that the performance of GMRES(k)-MG can be improved by increasing the restart parameter. In Figure 4.1, however, the values k have been chosen to balance storage and convergence behavior. In particular, we set k to be sufficiently small, so that the storage required for GMRES(k)-MG is as close as possible to that of MINRES-AV-MG, while the convergence of the method is not lost. Since Bi-CGSTAB-MG is based on a short-term recurrence, its storage expenses are similar to that of MINRES-AV-MG.

The unsatisfactory performance of GMRES(k)-MG in Figure 4.1 can in part be attributed to the observation that smoothing based on Richardson's (or, more generally, Jacobi) iteration becomes increasingly unstable as grids coarsen. In particular, as shown in [13], on the intermediate levels with $ch_l \geq 1/2$ this smoothing scheme strongly amplifies the smooth error eigenmodes. A straightforward remedy is to invoke the coarse grid solve on the largest grid that fails to satisfy $ch_l < 1/2$.

In Table 4.3, we compare MINRES-AV-MG and GMRES(k)-MG with different values of the restart parameter. We report the iteration counts required to reduce

TABLE 4.3
Number of iterations of MINRES-AV-MG (MINR) and GMRES(k)-MG (GMR(k)) required to reduce the initial error by 10^{-8} ; $n = 65025$. The preconditioner in GMRES(k)-MG uses Richardson's smoothing on levels $ch_l < 1/2$ and invokes the coarse grid solve on the level that follows. Numbers in parentheses correspond to the right preconditioned GMRES(k)-MG. Dash denotes that the method failed to converge within 1000 steps.

	MINR	GMR(5)	GMR(10)	GMR(20)	GMR(25)	GMR(35)
$c^2 = 1500$	89	29(44)	20(22)	16(18)	16(18)	16(18)
$c^2 = 3000$	282	—(—)	—(—)	—(—)	223(269)	69(69)
$c^2 = 4000$	310	—(—)	—(—)	—(—)	—(—)	395(471)

the initial error by 10^{-8} for systems with $c^2 = 1500, 3000,$ and 4000 . The indefinite MG preconditioner in GMRES(k)-MG is configured to run Richardson’s smoothing on grids $ch_l < 1/2$ and perform the coarse grid solve on the level that follows. We test the indefinite MG preconditioner in the left (used so far) and right preconditioned versions of GMRES(k).

The above described setting of the right preconditioned GMRES(k)-MG represents a special case of the Helmholtz solver introduced in [13]. In this paper, instead of the “early” coarse grid solve, the MG preconditioning scheme performs further coarsening on levels $ch_l \geq 1/2$ with GMRES used as a smoother. Since the resulting preconditioner is nonlinear, (the full) FGMRES [33] is used for outer iterations.

It is clear that the results in Table 4.3 provide an insight into the best possible outer iteration counts that can be expected from the *restarted* version of the method in [13]. In fact, the same observation is used by the authors of [13] for hand-tuning the smoothing schedule of their preconditioner to study its best-case performance. Table 4.3 then demonstrates that regardless of the smoothing schedule and the smoother’s stopping criterion, the (F)GMRES based solvers, such as the one in [13], require increasing storage to maintain the convergence as c^2 grows, whereas the robustness of MINRES-AV-MG is not lost under the minimalist memory costs.

Thus, if the shift is large and the amount of storage is limited, so that k is forced to be sufficiently small, the (F)GMRES(k) outer iterations may fail to converge within a reasonable number of steps, even if the coarse grid solve in the MG preconditioner is performed “early.” We note, however, that if storage is available or the shifts are not too large, the (F)GMRES based methods may represent a valid option.

We have also tested the Bunch–Parlett factorization [18] as a coarse grid solve in the MG framework. In particular, as a preconditioner for MINRES, we have used Algorithm 3.2 with $B_l = L_l$ on all levels and the coarsest grid absolute value in (3.37) replaced by the application of the “perfect” Bunch–Parlett factorization based preconditioner [18]. We have obtained results that are inferior to the schemes considered in this paper for shifts not too small, e.g., for $c^2 > 200$ if $h = 2^{-8}$. The unsatisfactory behavior may be related to the fact that the inverted Laplacian $T = L^{-1}$ and the ideal absolute value $T = |L - c^2I|^{-1}$ preconditioners share the same eigenvectors with $A = L - c^2I$, while the preconditioner from [18] does not.

The standard MG preconditioners, as in GMRES(k)-MG and Bi-CGSTAB-MG, are known to have optimal costs, linearly proportional to n . As discussed above, the same is true for the AV preconditioner in Algorithm 3.2. Therefore if in addition the number of iterations in the iterative solver preconditioned with Algorithm 3.2 does not depend on the problem size, the overall scheme is optimal.

We verify this optimality in Table 4.4, which shows the mesh-independence of the convergence of PMINRES with the MG AV preconditioner. The rows of the table

TABLE 4.4

Mesh-independent convergence of PMINRES with the MG AV preconditioner. The numbers in parentheses correspond to $\delta = 3/4$. The default value of δ is $1/3$.

	$h = 2^{-6}$	$h = 2^{-7}$	$h = 2^{-8}$	$h = 2^{-9}$	$h = 2^{-10}$	$h = 2^{-11}$
$c^2 = 300$	31(31)	31(31)	30(32)	30(32)	30(32)	30(30)
$c^2 = 400$	37(40)	38(40)	37(40)	37(40)	37(40)	37(39)
$c^2 = 1500$	67(97)	97(119)	89(109)	88(108)	89(106)	90(107)
$c^2 = 3000$	228(229)	222(284)	279(332)	256(298)	257(296)	256(298)

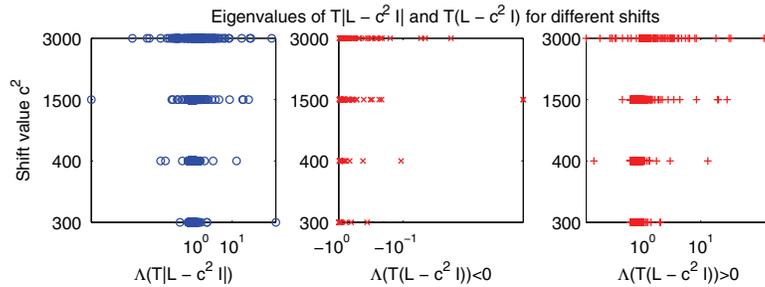


FIG. 4.2. Spectrum of $T|L - c^2 I|$ (left), negative eigenvalues of $T(L - c^2 I)$ (center), and positive eigenvalues of $T(L - c^2 I)$ (right); $n = 16129$.

TABLE 4.5

Number of eigenvalues λ_j that equal the upper/lower bound in (2.4) up to machine precision.

	$c^2 = 300$	$c^2 = 400$	$c^2 = 1500$	$c^2 = 3000$
Upper	0	15	1	115
Lower	0	0	10	0

correspond to the shift values c^2 , while the columns match the mesh size h . The cell in the intersection contains the numbers of steps performed to achieve the decrease by the factor 10^{-8} in the error 2-norm with the choices of the “switching” parameter $\delta = 1/3$ and $\delta = 3/4$.

As previously, the size of the coarsest grid has been set according to Table 4.1 with $\delta = 1$. We conclude that the convergence does not slow with the decrease of h ; thus, PMINRES preconditioned by Algorithm 3.2 is optimal. Note that for larger shifts, $c^2 = 1500$ and $c^2 = 3000$, mesh-independent convergence occurs for h sufficiently small when the “switching” pattern is stabilized, i.e., $B_l = L_l$ on a few finer grids and $B_l = p_{m_l}(L_l - c^2 I_l)$ on the coarser grids that follow.

Table 4.4 shows that as c^2 grows, the increase in the iteration count is mild and essentially linear. As expected, the smaller value of δ , which leads to the construction of the polynomial approximations earlier on finer levels, results in a higher accuracy of the AV preconditioner.

Finally, in Figure 4.2 we plot the eigenvalues of $T|L - c^2 I|$ and $T(L - c^2 I)$ for different shift values using the logarithmic scale; $n = 16129$. As suggested by Corollary 2.5, clusters of eigenvalues of $T|L - c^2 I|$ are preserved in the spectrum of the preconditioned matrix $T(L - c^2 I)$. Almost all eigenvalues of $T(L - c^2 I)$ are clustered around -1 and 1 , with only a few falling outside the clusters. We note that the clustering and the condition number $\kappa(T|A|)$ deteriorate as c^2 increases from 300 to 3000, which is compatible with the results in Table 4.4.

The spectra computed in Figure 4.2 allow validating numerically the tightness of bounds (2.4) in Theorem 2.3 for the MG AV preconditioner. In Table 4.5, we report the number of eigenvalues λ_j of $T(L - c^2 I)$ that satisfy either the upper or the lower bound up to machine precision. The table shows that the bound is numerically sharp.

5. Conclusions. We propose a new approach for SPD preconditioning for symmetric indefinite systems based on the idea of implicitly constructing approximations to the inverse of the system matrix absolute value. An MG example of such a preconditioner is presented for a real-valued Helmholtz problem. Our experiments demonstrate that PMINRES with the new MG AV preconditioner leads to an efficient

iterative scheme, which has modest memory requirements and outperforms traditional GMRES based methods if available memory is tight.

Acknowledgments. The authors thank Michele Benzi, Yvan Notay, and Joe Pasciak for their comments on the draft of the manuscript.

REFERENCES

- [1] T. AIRAKSINEN, E. HEIKKOLA, AND A. PENNANEN, *An algebraic multigrid based shifted-Laplacian preconditioner for the Helmholtz equation*, J. Comput. Phys., 226 (2007), pp. 1196–1210.
- [2] A. BAYLISS, C. I. GOLDSTEIN, AND E. TURKEL, *An iterative method for the Helmholtz equation*, J. Comput. Phys., 49 (1983), pp. 443–457.
- [3] M. BENZI, G. H. GOLUB, AND J. LIESEN, *Numerical solution of saddle point problems*, Acta Numer., 14 (2005), pp. 1–137.
- [4] J. H. BRAMBLE, Z. LEYK, AND J. E. PASCIAK, *Iterative schemes for nonsymmetric and indefinite elliptic boundary value problems*, Math. Comp., 60 (1993), pp. 1–22.
- [5] J. H. BRAMBLE AND X. ZHANG, *The analysis of multigrid methods*, in Solution of Equation in \mathbb{R}^n (Part 3), Techniques of Scientific Computing (Part 3), Handb. Numer. Anal. 7, Elsevier, Amsterdam, 2000, pp. 173–415.
- [6] M. BOLLHÖFER, M. J. GROTE, AND O. SCHENK, *Algebraic multilevel preconditioner for the Helmholtz equation in heterogeneous media*, SIAM J. Sci. Comput., 31 (2009), pp. 3781–3805.
- [7] A. BRANDT AND S. TA’ASAN, *Multigrid method for nearly singular and slightly indefinite problems*, in Multigrid Methods II, Lecture Notes in Math. 1228, Springer, Berlin, 1986, pp. 99–121.
- [8] W. L. BRIGGS, V. E. HENSON, AND S. F. MCCORMICK, *A Multigrid Tutorial*, SIAM, Philadelphia, 2000.
- [9] J. R. BUNCH AND B. N. PARLETT, *Direct methods for solving symmetric indefinite systems of linear equations*, SIAM J. Numer. Anal., 8 (1971), pp. 639–655.
- [10] E. G. D’YAKONOV, *Optimization in Solving Elliptic Problems*, CRC Press, Boca Raton, FL, 1996.
- [11] J. ERHEL, F. GUYOMARC, AND Y. SAAD, *Least-Squares Polynomial Filters for Ill-Conditioned Linear Systems*, Report UMSI-2001-32, Minnesota Supercomputer Institute, University of Minnesota, Minneapolis, 2001.
- [12] Y. A. ERLANGGA, C. VUIK, AND C. W. OOSTERLEE, *On a class of preconditioners for solving the Helmholtz equation*, Appl. Numer. Math., 50 (2004), pp. 409–425.
- [13] H. C. ELMAN, O. G. ERNST, AND D. O’LEARY, *A multigrid method enhanced by Krylov subspace iteration for discrete Helmholtz equations*, SIAM J. Sci. Comput., 23 (2001), pp. 1291–1315.
- [14] H. C. ELMAN, D. J. SILVESTER, AND A. J. WATHEN, *Finite Elements and Fast Iterative Solvers with Applications in Incompressible Fluid Dynamics*, Oxford University Press, Oxford, UK, 2005.
- [15] V. FABER AND T. MANTEUFFEL, *Necessary and sufficient conditions for the existence of a conjugate gradient method*, SIAM J. Numer. Anal., 21 (1984), pp. 352–362.
- [16] B. FISCHER, A. RAMAGE, D. J. SILVESTER, AND A. J. WATHEN, *Minimum residual methods for augmented systems*, BIT, 38 (1998), pp. 527–543.
- [17] R. W. FREUND AND N. M. NACHTIGAL, *QMR: A quasi-minimal residual method for non-Hermitian linear systems*, Numer. Math., 60 (1991), pp. 315–339.
- [18] P. E. GILL, W. MURRAY, D. B. PONCELEÓN, AND M. A. SAUNDERS, *Preconditioners for indefinite systems arising in optimization*, SIAM J. Matrix Anal. Appl., 13 (1992), pp. 292–311.
- [19] G. H. GOLUB AND C. F. VAN LOAN, *Matrix Computations*, Johns Hopkins University Press, Baltimore, MD, 1996.
- [20] A. GREENBAUM, V. PTÁK, AND Z. STRAKOŠ, *Any nonincreasing convergence curve is possible for GMRES*, SIAM J. Matrix Anal. Appl., 17 (1996), pp. 465–469.
- [21] E. HABER AND S. MACLACHLAN, *A fast method for the solution of the Helmholtz equation*, J. Comp. Phys., 230 (2011), pp. 4403–4418.
- [22] I. HARARI AND T. J. R. HUGHES, *Finite element method for the Helmholtz equation in an exterior domain*, Comput. Methods Appl. Mech. Engrg., 87 (1991), pp. 59–96.
- [23] N. J. HIGHAM, *Functions of Matrices: Theory and Computation*, SIAM, Philadelphia, 2008.
- [24] R. A. HORN AND C. R. JOHNSON, *Topics in Matrix Analysis*, Cambridge University Press, New York, 1994.

- [25] A. L. LAIRD AND M. B. GILES, *Preconditioned Iterative Solution of the 2D Helmholtz Equation*, Oxford University Computing Laboratory, Oxford, UK, 2002.
- [26] J. NOCEDAL AND S. WRIGHT, *Numerical Optimization*, Springer, Berlin, 1999.
- [27] D. OSEI-KUFFUOR AND Y. SAAD, *Preconditioning Helmholtz linear systems*, *Appl. Numer. Math.*, 60 (2010), pp. 420–431.
- [28] C. C. PAIGE AND M. A. SAUNDERS, *Solution of sparse indefinite systems of linear equations*, *SIAM J. Numer. Anal.*, 12 (1975), pp. 617–629.
- [29] B. N. PARLETT, *The Symmetric Eigenvalue Problem*, SIAM, Philadelphia, 1998.
- [30] M. J. D. POWELL, *Approximation Theory and Methods*, Cambridge University Press, New York, 1981.
- [31] T. J. RIVLIN, *An Introduction to the Approximation of Functions*, Dover, New York, 1981.
- [32] Y. SAAD, *Filtered conjugate residual-type algorithms with applications*, *SIAM J. Matrix Anal. Appl.*, 28 (2006), pp. 845–870.
- [33] Y. SAAD, *A flexible inner-outer preconditioned GMRES algorithm*, *SIAM J. Sci. Comput.*, 14 (1993), pp. 461–469.
- [34] Y. SAAD, *Iterative Methods for Sparse Linear Systems*, SIAM, Philadelphia, 2003.
- [35] Y. SAAD AND M. H. SCHULTZ, *GMRES: A generalized minimal residual algorithm for solving nonsymmetric linear systems*, *SIAM J. Sci. Statist. Comput.*, 7 (1986), pp. 856–869.
- [36] G. L. G. SLEIJPEN AND H. A. VAN DER VORST, *A Jacobi–Davidson iteration method for linear eigenvalue problems*, *SIAM J. Matrix Anal. Appl.*, 17 (1996), pp. 401–425.
- [37] D. J. SILVESTER AND A. J. WATHEN, *Fast iterative solution of stabilised Stokes systems II: Using general block preconditioners*, *SIAM J. Numer. Anal.*, 31 (1994), pp. 1352–1367.
- [38] O. TATEBE, *The multigrid preconditioned conjugate gradient method*, in *Proceedings of the Sixth Copper Mountain Conference on Multigrid Methods*, 1993, pp. 621–634.
- [39] U. TROTTEBERG, C. W. OOSTERLEE, AND A. SCHÜLLER, *Multigrid*, Academic Press, New York, 2001.
- [40] M. B. VAN GIJZEN, Y. A. ERLANGGA, AND C. VUIK, *Spectral analysis of the discrete Helmholtz operator preconditioned with a shifted Laplacian*, *SIAM J. Sci. Comput.*, 29 (2007), pp. 1942–1958.
- [41] H. A. VAN DER VORST, *Bi-CGSTAB: A fast and smoothly converging variant of Bi-CG for the solution of nonsymmetric linear systems*, *SIAM J. Sci. Statist. Comput.*, 13 (1992), pp. 631–644.
- [42] E. VECHARYNSKI, *Preconditioned Iterative Methods for Linear Systems, Eigenvalue and Singular Value Problems*, Ph.D. thesis, University of Colorado, Denver, 2010.
- [43] E. VECHARYNSKI AND J. LANGOU, *Any admissible cycle-convergence behavior is possible for restarted GMRES at its initial cycles*, *Numer. Linear Algebra Appl.*, 18 (2011), pp. 499–511.
- [44] A. J. WATHEN AND D. J. SILVESTER, *Fast iterative solution of stabilized Stokes systems I: Using simple diagonal preconditioners*, *SIAM J. Numer. Anal.*, 30 (1993), pp. 630–649.
- [45] Y. ZHOU, Y. SAAD, M. L. TIAGO, AND J. R. CHELIKOWSKY, *Self-consistent-field calculation using Chebyshev-filtered subspace iteration*, *J. Comput. Phys.*, 219 (2006), pp. 172–184.