

ABSOLUTE VALUE PRECONDITIONING FOR SYMMETRIC INDEFINITE LINEAR SYSTEMS

EUGENE VECHARYNSKI* AND ANDREW KNYAZEVA†

Abstract. We introduce a novel strategy for constructing symmetric positive definite (SPD) preconditioners for linear systems with symmetric indefinite coefficient matrices. The strategy is motivated by the observation that the preconditioned minimal residual method with the inverse of the absolute value of the coefficient matrix as a preconditioner converges to the exact solution of the system in at most two steps. Neither the exact absolute value of the coefficient matrix, nor its exact inverse are computationally feasible to construct in general. However, as the proof of concept, we provide two practical examples of SPD preconditioners, which are based on the suggested approach, called absolute value preconditioning. The first example is for strictly (block) diagonally dominant coefficient matrices, where we propose using the inverse to the absolute value of the (block) diagonal as the preconditioner. Our second example is less intuitive. We consider a model problem with a shifted discrete negative Laplacian, and suggest a geometric multigrid preconditioner, where the inverse to the absolute value of the coefficient matrix 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 multigrid preconditioner for moderately small shifts.

Key words. Preconditioning, linear system, preconditioned minimal residual method, polar decomposition, matrix absolute value, geometric multigrid

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

1. Introduction. We consider a system of linear equations

$$(1.1) \quad Ax = b, \quad A = A^* \in \mathbb{R}^{n \times n}, \quad b \in \mathbb{R}^n,$$

where the coefficient matrix A is nonsingular and symmetric indefinite, i.e., the spectrum of A contains both positive and negative eigenvalues.

For a relatively small problem size n , the *exact* (up to the effects of round-off errors) solution of (1.1) can be efficiently found using a *direct* method; see, e.g., the survey in [7]. However, in most cases, the computational cost of such a method does not *optimally* scale with respect to $n \rightarrow \infty$, i.e., the amount of involved computations does not grow *proportionally* to the increasing number of matrix elements. Therefore, if the problem size is significantly large, the application of a direct method may become infeasible. Additionally, it is often required to find only an *approximate* solution of (1.1), as opposed to the *exact* solution targeted by a direct solver.

The above arguments motivate an *iterative* technique, which can be expected to provide a (nearly) optimal complexity and, instead of the exact solution of (1.1), delivers a sequence of its approximations. In this manuscript, it is assumed that the coefficient matrix A is extremely large and possibly sparse. We focus only on *iterative* methods for solving linear system (1.1). Let us remark that, in the suggested framework, under “solving a linear system” we understand “*approximately* solving a linear system,” i.e., finding a satisfactory approximation to the exact solution of (1.1).

Linear systems with large, sparse, symmetric, and indefinite coefficient matrices 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

*Department of Computer Science and Engineering, University of Minnesota, 200 Union Street S.E., Minneapolis, MN 55455 (eugenev@cs.umn.edu).

†Department of Mathematical and Statistical Sciences, University of Colorado Denver, P.O. Box 173364, Campus Box 170, Denver, CO 80217-3364 (andrew.knyazev@ucdenver.edu).

equations of fluid and solid mechanics; see, e.g., [3] and references therein. In acoustics, large sparse symmetric indefinite linear systems are obtained after discretizing the Helmholtz equation for certain media types and boundary conditions. Often the need of solving indefinite problem (1.1) comes as an auxiliary task within other computational routines, such as the inner step in the interior point methods in linear and nonlinear optimization, see, e.g., [3, 24], or solution of the correction equation in the Jacobi-Davidson method [29] for a symmetric eigenvalue problem.

There is a number of iterative methods developed *specifically* to solve symmetric indefinite linear systems, ranging from modifications of the Richardson's iteration, e.g., [8, 22, 27], to the optimal Krylov subspace minimal residual method delivered through short-term recurrent schemes, such as, e.g., the MINRES algorithm [25]. It is known, however, that in practical problems the coefficient matrix A in (1.1) is very ill-conditioned which, along with the location of the spectrum to both sides of the origin, makes the straightforward application of the existing techniques inefficient due to extremely slow convergence. In order to improve the convergence, it is possible to introduce a *preconditioner* $T \in \mathbb{R}^{n \times n}$ and formally replace (1.1) with

$$(1.2) \quad TAx = Tb.$$

If T is properly chosen, then an iterative method for (1.2) can exhibit a better convergence behavior compared to the original scheme applied to (1.1). The *preconditioned matrix* TA is never explicitly computed.

If T is *not* symmetric positive definite (SPD), then the matrix TA of preconditioned linear system (1.2), in general, is not symmetric with respect to any inner product; see, e.g. [26, Theorem 15.2.1]. Thus, the introduction of a non-SPD preconditioner leads to the substitution of original *symmetric* problem (1.1) by generally *nonsymmetric* problem (1.2). As a result, specialized methods for *symmetric* linear systems, suitable for (1.1), may no longer be applicable to (1.2), and, hence, should be replaced by iterative schemes for *nonsymmetric* linear systems; e.g., GMRES or GMRES(m) [28], Bi-CGSTAB [35], QMR [13], etc.

Though known to be effective for a number of applications, the approach based on the choice of a non-SPD preconditioner, which leads to solving nonsymmetric problem (1.2), has several disadvantages. First, no short-term recurrent scheme that delivers an *optimal* Krylov subspace method is typically available for a nonsymmetric linear system; see [11]. In practice, this means that implementations of the relevant optimal methods (e.g., GMRES) require an increasing amount of work and storage at every new step, and 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, e.g., the (estimated) spectrum of the preconditioned coefficient matrix; see the corresponding results for GMRES and GMRES(m) in [17, 37]. This makes it difficult to predict the computational cost of the selected approach.

If T is chosen to be SPD, i.e., $T = T^* > 0$, then the matrix TA of preconditioned linear system (1.2) is symmetric (indefinite), however, with respect to the T^{-1} -inner product defined by $(u, v)_{T^{-1}} = (u, T^{-1}v)$ for any $u, v \in \mathbb{R}^n$. Here (\cdot, \cdot) denotes the Euclidean inner product, i.e., $(u, v) = v^*u$, in which the matrices A and T are symmetric. Due to this symmetry preservation, linear system (1.2) can be solved using an *optimal* Krylov subspace method which admits a *short-term recurrent* implementation. Moreover, the convergence of the method can be fully estimated in terms of the spectrum of the preconditioned matrix TA .

In the light of the discussion above, the choice of a properly defined SPD preconditioner for a symmetric indefinite linear system can be regarded as natural and favorable. We advocate the use of SPD preconditioning for problem (1.1) and propose a novel strategy for constructing SPD preconditioners.

Given an SPD preconditioner T , we consider solving symmetric indefinite linear system (1.1) with the *preconditioned minimal residual method*, implemented in the form of the PMINRES algorithm [16, 25]. In the absence of round-off errors, at step i , the method constructs an approximation $x^{(i)}$ to the solution of (1.1) of the form

$$(1.3) \quad x^{(i)} \in x^{(0)} + \mathcal{K}_i(TA, Tr^{(0)}),$$

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

$$(1.4) \quad \|r^{(i)}\|_T = \min_{u \in \mathcal{AK}_i(TA, Tr^{(0)})} \|r^{(0)} - u\|_T.$$

Here,

$$\mathcal{K}_i(TA, Tr^{(0)}) = \text{span} \left\{ Tr^{(0)}, (TA)Tr^{(0)}, \dots, (TA)^{i-1}Tr^{(0)} \right\}$$

is the Krylov subspace generated by the matrix TA and the vector $Tr^{(0)}$, and

$$A\mathcal{K}_i(TA, Tr^{(0)}) = \text{span} \left\{ (AT)r^{(0)}, \dots, (AT)^i r^{(0)} \right\}$$

is the corresponding Krylov residual subspace; $\|v\|_T^2 = (v, v)_T$ for any $v \in \mathbb{R}^n$, and $x^{(0)}$ is the initial guess. The PMINRES implementation of the optimal preconditioned minimal residual method (1.3)–(1.4) for (1.1) is based on a *short-term recurrence*. It can be viewed as the MINRES algorithm applied to (1.2) with occurring Euclidean inner products replaced by T^{-1} -based inner products, written in a way to avoid computations involving T^{-1} . The conventional convergence rate bound for method (1.3)–(1.4) can be found, e.g., in [16], and relies solely on the distribution of eigenvalues of the preconditioned matrix TA .

In this work, we describe a novel approach for constructing SPD preconditioners for symmetric indefinite linear systems. First, in Section 2, we present an *ideal* SPD preconditioner for problem (1.1), which we define as the inverse of the absolute value of the coefficient matrix. Next, motivated by this ideal preconditioner, we introduce a notion of an *absolute value preconditioner*, and refer to the new preconditioning strategy as the *absolute value preconditioning*. The rest of the paper deals with the question if absolute value preconditioners can be efficiently constructed in practice.

In Section 3, we provide simple examples of such preconditioners for linear systems with strictly (block) diagonally dominant coefficient matrices. In Section 4, we suggest a different example, a geometric multigrid (MG) absolute value preconditioner for a model problem resulting from a discretization of a shifted negative Laplace operator on a unit square. The efficiency of this preconditioner for moderately small shifts is demonstrated in our numerical tests, also in Section 4. In particular, we compare the proposed absolute value preconditioner to several other SPD preconditioners, which result from the known ideas of using the inverted Laplacian and the so-called Bunch-Parlett factorization for preconditioning symmetric indefinite linear systems.

Throughout this paper, we assume the exact arithmetic. The choice of the real vector spaces has been made to simplify the presentation. The generalization of the results to the complex case is straightforward. We note that the results presented in this work are partially based on the PhD thesis of the first co-author [36], defended at the University of Colorado Denver, under the supervision of the second co-author.

2. Absolute value preconditioning for symmetric indefinite linear systems. Let $A \in \mathbb{R}^{n \times n}$ be a symmetric matrix with an eigendecomposition $A = V\Lambda V^*$, where V is an orthogonal matrix of eigenvectors, and $\Lambda = \text{diag}\{\lambda_j\}$, $j = 1, \dots, n$, is a diagonal matrix of eigenvalues of A . We consider the factorization of the form

$$(2.1) \quad A = |A| \text{sign}(A) = \text{sign}(A) |A|,$$

where $|A| = V|\Lambda|V^*$ is an absolute value of A (matrix absolute value), and $\text{sign}(A) = V\text{sign}(\Lambda)V^*$ is a sign of A (matrix sign); $|\Lambda| = \text{diag}\{|\lambda_j|\}$, $\text{sign}(\Lambda) = \text{diag}\{\text{sign}(\lambda_j)\}$. Factorization (2.1) is, in fact, a *polar decomposition*, see, e.g., [19], of the symmetric matrix A , with the symmetric positive (semi) definite factor $|A|$ and the orthogonal factor $\text{sign}(A)$. If, additionally, A is nonsingular, then $|A|$ is SPD.

The following theorem regards the inverted absolute value of the coefficient matrix, i.e., $|A|^{-1}$, as an SPD preconditioner for symmetric indefinite linear system (1.1).

THEOREM 2.1. *Preconditioned minimal residual method (1.3)–(1.4) with the preconditioner $T = |A|^{-1}$ converges to the exact solution of (1.1) in at most two steps.*

Proof. Minimization property (1.4) at a step i of a preconditioned minimal residual method can be equivalently written as

$$(2.2) \quad \|r^{(i)}\|_T = \min_{p \in \mathcal{P}_i, p(0)=1} \|p(AT)r^{(0)}\|_T,$$

where \mathcal{P}_i is a set of all polynomials of degree at most i . Then, according to decomposition (2.1), the choice $T = |A|^{-1}$ results in the matrix $AT = \text{sign}(A)$ with only two distinct eigenvalues: -1 and 1 . Hence, the minimal polynomial of AT is of the second degree. Thus, by (2.2), $\|r^{(i)}\|_T = 0$ for at most $i = 2$. \square

Theorem 2.1 implies that $T = |A|^{-1}$ is an *ideal SPD preconditioner* for symmetric indefinite linear system (1.1). Indeed, a preconditioner T can be considered ideal if it delivers the preconditioned matrix TA with the corresponding minimal polynomial of the least possible degree; e.g., see related discussions in [14, 23]. The degree of the minimal polynomial gives the number of iterations typically required by a Krylov subspace method to guarantee the convergence to the exact solution of (1.1). Since the symmetric matrix A has both positive and negative eigenvalues, then so does TA , provided that T is SPD. This implies that the degree of the minimal polynomial of TA is at least 2. Therefore, an ideal SPD preconditioner for symmetric indefinite linear system (1.1), in general, ensures the convergence to the exact solution in two steps. The one step convergence can occur for special choices of the initial guess. We also note that Theorem 2.1 holds not only for preconditioned minimal residual method (1.3)–(1.4), but for all other methods, where convergence to the exact solution is determined by the degree of the minimal polynomial of the preconditioned matrix.

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 reasonably inexpensive SPD preconditioners, which resemble, but are not equal to, $|A|^{-1}$, and are able to significantly accelerate the convergence of an underlying iterative method. We refer to such a preconditioning strategy as the *absolute value preconditioning* and define *absolute value preconditioners* as following.

DEFINITION 2.2. *We call an SPD preconditioner T for a symmetric indefinite linear system (1.1) an **absolute value preconditioner**, if it satisfies the relation*

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

with constants $\delta_1 \geq \delta_0 > 0$, such that the ratio $\delta_1/\delta_0 \geq 1$ is reasonably small. If linear system (1.1) represents a hierarchy of mesh problems, then the ratio must be independent of the problem size n , i.e., matrices $|A|$ and T^{-1} are spectrally equivalent [9].

Let us remark that Definition 2.2 of the absolute value preconditioner is informal, since no precise assumption is made of 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 resembles an ideal SPD preconditioner $|A|^{-1}$ up to a positive scaling. If A is SPD, then $|A| = A$, and, for mesh problems, Definition 2.2 is consistent with the well known concept of spectrally equivalent preconditioning for linear systems with SPD coefficient matrices; see [9].

The following theorem provides bounds for eigenvalues of the preconditioned matrix TA in terms of the spectrum of $T|A|$.

THEOREM 2.3. *Let us be given a nonsingular symmetric indefinite $A \in \mathbb{R}^{n \times n}$ and an SPD $T \in \mathbb{R}^{n \times n}$, and 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 the proof 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

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

Further, 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 can write

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

Moreover, since taking maxima in vectors $v \in S$, and then minima in subspaces $S \in S^j = \{S \subseteq \mathbb{R}^n : \dim(S) = j\}$, of all sides of (2.7) preserves the inequality, we have

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

By the Courant-Fischer theorem (see, e.g., [19, 26]) 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 the matrix TA :

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

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 $\phi(v)$, defined in (2.12), allows us to conclude from (2.8) that

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

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

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

Combining (2.9) and (2.13), we obtain (2.4). This completes the proof. \square

Theorem 2.3 suggests two useful implications, which are given by the corresponding theorems below. In particular, the following result describes the location of spectrum of the preconditioned matrix TA in terms of δ_0 and δ_1 in (2.3).

THEOREM 2.4. *Let us be given a nonsingular symmetric indefinite $A \in \mathbb{R}^{n \times n}$, an SPD $T \in \mathbb{R}^{n \times n}$, and constants $\delta_1 \geq \delta_0 > 0$, such that (2.3) holds. Then*

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

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

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

The next theorem 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 .

THEOREM 2.5. *Let us be given a nonsingular symmetric indefinite $A \in \mathbb{R}^{n \times n}$ and an SPD $T \in \mathbb{R}^{n \times n}$, and 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. Follows directly from bounds (2.4). \square

Theorem 2.4 suggests that the ratio $\delta_1/\delta_0 \geq 1$ of the constants from (2.3) measures the quality of the absolute value preconditioner T in terms of the convergence speed of the preconditioned minimal residual method, which is determined by the spectrum of TA . Additionally, Theorem 2.5 prompts that a “good” absolute value 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 the preconditioned minimal residual method, for linear system (1.1).

At the same time, we recall that the costs of the construction and application of T should preferably be similar to the costs of the matrix-vector multiplication by the coefficient matrix A . Below we discuss several possible approaches to construct an absolute value preconditioner T for linear system (1.1) with a general symmetric indefinite coefficient matrix. Without loss of generality, we assume that T can be accessed only through a matrix-vector multiplication. In other words, given a vector $r \in \mathbb{R}^n$, we define a preconditioner as a mapping $r \mapsto w = Tr$, where the matrix T may not be *explicitly* constructed and stored. Iterative schemes are normally designed to handle such *matrix-free* preconditioning.

One of the approaches to construct an absolute value preconditioner is to approximately solve for z the equation $|A|z = r$. However, the coefficient matrix $|A|$ is generally not available. Thus, the problem of approximately solving the linear system $|A|z = r$ can be replaced by the problem of finding a vector w which approximates the action of the *matrix function* $f(A) = |A|^{-1}$ on the vector r , i.e., $w \approx f(A)r = |A|^{-1}r$, moreover the construction of w requires only knowledge of A , not $|A|$.

The latter constitutes a well established task in matrix function computations, which is standardly fulfilled by a Krylov subspace method, e.g., [15, 18]. Our numerical experience shows that the convergence, with respect to the (outer) iteration count, of a linear solver can be significantly improved with this approach, but the computational costs of approximating $f(A)r = |A|^{-1}r$, e.g., by the Lanczos method in [4], are too high for their application in the context of absolute value preconditioners.

We also note that a Krylov subspace method for constructing $w = Tr \approx |A|^{-1}r$ may result in the preconditioner T , which is symmetric positive *semidefinite* and *variable*, whereas the assumption for T to be *SPD* and *non-variable* is typically crucial for the conventional convergence analyses of iterative linear solvers. In particular, the use of variable preconditioners leads to the loss of the global optimality of a Krylov subspace method for (1.1); see, e.g., the corresponding convergence bounds for the conjugate gradients method in [20].

Another approach to construct an absolute value preconditioner is to use a preconditioning technique, which aims to deliver T satisfying (2.3) with the ratio δ_1/δ_0 reasonably small (and independent of n for mesh problems), however, which possibly does not correspond to a solver for equation $|A|z = r$, i.e., the norm $\|T - |A|^{-1}\|$ may not necessarily be small. In Section 4, we demonstrate, for a model problem, that such a construction of an efficient absolute value preconditioner is indeed possible, e.g., if based on MG techniques.

The discussion above concerns the construction of absolute value preconditioners for a general symmetric indefinite linear system (1.1). In the next section, we consider

a special case, where the coefficient matrix A has a (block) diagonal approximation C , so that $|C|^{-1}$ is available at a small computational cost. A reasonable absolute value preconditioner can then be expected to be given by $T = |C|^{-1}$.

3. Absolute value preconditioning for strictly (block) diagonally dominant matrices. For certain types of symmetric indefinite matrices, absolute value preconditioners can be easily constructed. For example, if $A = \{a_{ij}\}$ is *strictly diagonally dominant*, i.e.,

$$(3.1) \quad |a_{ii}| > \sum_{\substack{j=1 \\ j \neq i}}^n |a_{ij}|, \quad i = 1, \dots, n,$$

then the preconditioner T may be chosen as a diagonal matrix with elements $|a_{ii}|^{-1}$. Assumptions similar to (3.1) appear, e.g., in electronic structure calculations using plain-wave discretization; see [31].

In the following theorems we do not establish (2.3), but rather directly prove a statement similar to (2.14).

THEOREM 3.1. *Let $A = \{a_{ij}\} \in \mathbb{R}^{n \times n}$ be a strictly diagonally dominant symmetric indefinite matrix, such that*

$$(3.2) \quad \delta |a_{ii}| \geq \sum_{\substack{j=1 \\ j \neq i}}^n |a_{ij}|, \quad i = 1, \dots, n,$$

for a fixed $\delta \in [0, 1)$. Let $T = \text{diag}\{|a_{11}|^{-1}, |a_{22}|^{-1}, \dots, |a_{nn}|^{-1}\}$. Then

$$(3.3) \quad \Lambda(TA) \subset \{y \in \mathbb{R} : |y + 1| \leq \delta\} \cup \{y \in \mathbb{R} : |y - 1| \leq \delta\}.$$

Proof. We observe that the matrix $TA = \{a_{ij}/|a_{ii}|\}; i, j = 1, \dots, n$. Since A is symmetric indefinite and T is SPD, the spectrum of TA resides to both sides of the origin on the real axis. In particular, the Gershgorin circle theorem, see, e.g., [19, Theorem 6.1.1], suggests that all eigenvalues of TA are located in the union of n intervals $\bigcup_{i=1}^n G_i$, where

$$(3.4) \quad G_i \equiv \left\{ y \in \mathbb{R} : |y - \text{sign}(a_{ii})| \leq \frac{1}{|a_{ii}|} \sum_{\substack{j=1 \\ j \neq i}}^n |a_{ij}| \right\}, \quad i = 1, \dots, n.$$

From (3.2), we notice that

$$\frac{1}{|a_{ii}|} \sum_{\substack{j=1 \\ j \neq i}}^n |a_{ij}| \leq \delta, \quad i = 1, \dots, n.$$

Thus,

$$\begin{aligned} G_i &\subseteq \{y \in \mathbb{R} : |y - \text{sign}(a_{ii})| \leq \delta\} \\ &= \{y \in \mathbb{R} : |y + 1| \leq \delta\} \cup \{y \in \mathbb{R} : |y - 1| \leq \delta\}, \quad i = 1, \dots, n, \end{aligned}$$

which completes the proof. \square

Theorem 3.1 shows that for a strictly diagonally dominant symmetric indefinite coefficient matrix A , with δ sufficiently close to zero in (3.3), a reasonable absolute value preconditioner is given by $T = \text{diag}\{|a_{11}|^{-1}, |a_{22}|^{-1}, \dots, |a_{nn}|^{-1}\}$. The spectrum of TA , in this case, is located around -1 and 1 . In fact, a similar result holds for strictly *block* diagonally dominant matrices as we now show.

A block partitioned matrix $A = \{A_{ij}\} \in \mathbb{R}^{n \times n}$, where $A_{ij} \in \mathbb{R}^{n_i \times n_j}$, $i, j = 1, \dots, s$, is strictly *block* diagonally dominant (relative to the partitioning $\{A_{ij}\}$) if the diagonal blocks A_{ii} are nonsingular, and if

$$(3.5) \quad (\|A_{ii}^{-1}\|)^{-1} > \sum_{\substack{j=1 \\ j \neq i}}^s \|A_{ij}\|, \quad i = 1, \dots, s.$$

The matrix norms in (3.5) are defined as spectral norms, i.e., $\|A_{ij}\| = \sigma_{\max}(A_{ij})$ and $\|A_{ii}^{-1}\| = \sigma_{\min}^{-1}(A_{ii})$, where $\sigma_{\max}(A_{ij})$ and $\sigma_{\min}(A_{ij})$ are the largest and the smallest singular values of A_{ij} , respectively.

Inequality (3.5) can be viewed as a direct generalization of definition (3.1) of the strict diagonal dominance to the case of block partitioned matrices. Several other generalizations can be found, e.g., in [12, 32]. In particular, the following extension of the Gershgorin circle theorem holds.

THEOREM 3.2 ([12, 32]). *Let $A = \{A_{ij}\} \in \mathbb{R}^{n \times n}$ be a block partitioned matrix with $A_{ij} \in \mathbb{R}^{n_i \times n_j}$; $i, j = 1, \dots, s$. Let us define*

$$(3.6) \quad G_i \equiv \Lambda(A_{ii}) \cup \left\{ y \in \mathbb{C} \setminus \Lambda(A_{ii}) : (\|yI_i - A_{ii}\|)^{-1} \leq \sum_{\substack{j=1 \\ j \neq i}}^s \|A_{ij}\| \right\}$$

for $i = 1, \dots, s$, where $\Lambda(A_{ii}) \subset \mathbb{C}$ is the spectrum of the submatrix A_{ii} , I_i is the n_i -by- n_i identity matrix, and $\|\cdot\|$ denotes the spectral norm¹. Then the spectrum of A is enclosed in the union of sets (3.6), i.e., $\Lambda(A) \subset \bigcup_{i=1}^s G_i$.

We further use Theorem 3.2 to show that for a strictly *block* diagonally dominant symmetric indefinite $A = \{A_{ij}\}$ in (1.1), the preconditioner T may be chosen as a block diagonal matrix with blocks $|A_{ii}|^{-1}$ on the diagonal.

THEOREM 3.3. *Let $A = \{A_{ij}\} \in \mathbb{R}^{n \times n}$, $i, j = 1, \dots, s$, be a strictly *block* diagonally dominant symmetric indefinite matrix, such that*

$$(3.7) \quad \delta (\|A_{ii}^{-1}\|)^{-1} \geq \sum_{\substack{j=1 \\ j \neq i}}^s \|A_{ij}\|, \quad i = 1, \dots, s,$$

for a fixed $\delta \in [0, 1)$. Let $T = \text{diag}\{|A_{11}|^{-1}, |A_{22}|^{-1}, \dots, |A_{ss}|^{-1}\}$. Then

$$(3.8) \quad \Lambda(TA) \subset \{y \in \mathbb{R} : |y + 1| \leq \delta\} \cup \{y \in \mathbb{R} : |y - 1| \leq \delta\}.$$

¹ Matrix norms in (3.5) and (3.6) can be defined more generally as conventional operator norms, induced by the vector norms on the corresponding n_i - (and n_j -) dimensional spaces; see [12, 32]. The choice of the spectral norm has been made only for the purposes of the current paper.

Proof. The proof is similar to that of Theorem 3.1. Since A is symmetric indefinite and T is SPD, the spectrum of TA is real. Given the block diagonal structure of T , the preconditioned matrix TA has the block form $TA = \{|A_{ii}|^{-1} A_{ij}\}; i, j = 1, \dots, s$. Since each submatrix A_{ii} of A is symmetric, the i -th diagonal block ($|A_{ii}|^{-1} A_{ii}$) of TA is equal to $\text{sign}(A_{ii}) \in \mathbb{R}^{n_i \times n_i}$, and $\Lambda(\text{sign}(A_{ii})) = \{-1, 1\}$. By Theorem 3.2, $\Lambda(TA) \subset \bigcup_{i=1}^s G_i$, where

$$G_i \equiv \{-1, 1\} \cup \left\{ y \in \mathbb{R} \setminus \{-1, 1\} : (\|yI_i - \text{sign}(A_{ii})\|^{-1})^{-1} \leq \sum_{\substack{j=1 \\ j \neq i}}^s \| |A_{ii}|^{-1} A_{ij} \| \right\}.$$

For any $y \in \mathbb{R} \setminus \{-1, 1\}$, the matrix $yI_i - \text{sign}(A_{ii})$ is symmetric and nonsingular, with eigenvalues $y \pm 1$. Hence, by definition of the spectral norm, we have

$$(3.9) \quad \| (yI_i - \text{sign}(A_{ii}))^{-1} \|^{-1} = \min \{|y + 1|, |y - 1|\}, \quad i = 1, \dots, s.$$

Next, the submultiplicativity of a matrix norm gives

$$(3.10) \quad \sum_{\substack{j=1 \\ j \neq i}}^s \| |A_{ii}|^{-1} A_{ij} \| \leq \| |A_{ii}|^{-1} \| \left(\sum_{\substack{j=1 \\ j \neq i}}^s \| A_{ij} \| \right), \quad i = 1, \dots, s.$$

The definition of the spectral norm implies that $\| |A_{ii}|^{-1} \| = \| A_{ii}^{-1} \|$. Thus, from (3.7) and (3.10), we obtain

$$(3.11) \quad \sum_{\substack{j=1 \\ j \neq i}}^s \| |A_{ii}|^{-1} A_{ij} \| \leq \delta, \quad i = 1, \dots, s.$$

Combining (3.9) and (3.11), we get

$$\begin{aligned} G_i &\subseteq \{-1, 1\} \cup \{y \in \mathbb{R} \setminus \{-1, 1\} : \min \{|y + 1|, |y - 1|\} \leq \delta\} \\ &= \{y \in \mathbb{R} : |y + 1| \leq \delta\} \cup \{y \in \mathbb{R} : |y - 1| \leq \delta\}, \quad i = 1, \dots, s. \end{aligned}$$

This completes the proof. \square

Theorem 3.3 shows that for symmetric indefinite coefficient matrices, which are well approximated by their block diagonal submatrices with relatively small block sizes, absolute value preconditioners can be efficiently constructed.

Finally, let us again remark that the (block) diagonal preconditioners, which result from Theorem 3.1 and Theorem 3.3, have been introduced without explicit validation of Definition 2.2 of an absolute value preconditioner, i.e., without a formal calculation of constants δ_0 and δ_1 in (2.3). If δ in (3.2)–(3.3) or (3.7)–(3.8) is sufficiently small, we still refer to these preconditioners as absolute value preconditioners.

In the next (and last) section, we present a different example of the absolute value preconditioning.

4. MG absolute value preconditioning for a model problem. Let us consider the following real boundary value problem,

$$(4.1) \quad \begin{aligned} -\Delta u(x, y) - c^2 u(x, y) &= f(x, y), \quad (x, y) \in \Omega = (0, 1) \times (0, 1), \\ u|_{\Gamma} &= 0, \end{aligned}$$

where $\Delta = \frac{\partial^2}{\partial x^2} + \frac{\partial^2}{\partial y^2}$ is the Laplace operator (Laplacian), and Γ denotes the boundary of Ω . Problem (4.1) is a particular instance of the Helmholtz equation with Dirichlet boundary conditions, where $c \in \mathbb{R}$ is a wavenumber.

After introducing a uniform grid of the size h in both directions and using the standard 5-point finite-difference stencil to discretize continuous problem (4.1), one obtains the corresponding discrete problem (1.1) of the form

$$(4.2) \quad (L - c^2 I)x = b,$$

where the coefficient matrix $A = L - c^2 I$ represents a discrete negative Laplace operator L , satisfying the Dirichlet boundary condition at the grid points on the boundary, shifted by a scalar c^2 times the identity matrix I . The right-hand side b in (4.2) is the vector of function values of $f(x, y)$ calculated at the grid points (numbered in the lexicographical order). The solution of system (4.2) provides an approximation to the solution of the boundary value problem (4.1), evaluated at the grid points. Below, we call (4.2) the *model problem*.

Assuming that c^2 is different from any eigenvalue of the SPD negative Laplacian L and is greater than its smallest, however less than its largest, eigenvalue, i.e., $\lambda_{\min}(L) < c^2 < \lambda_{\max}(L)$, where $\lambda_{\min}(L) = 2\pi^2 + \mathcal{O}(h^2)$ and $\lambda_{\max}(L) = 8h^{-2} + \mathcal{O}(1)$, we conclude that $A = L - c^2 I$ is nonsingular symmetric indefinite. Thus, as an SPD preconditioner T for linear system (4.2) one can choose an absolute value preconditioner; see Definition 2.2 with $A = L - c^2 I$.

There is a variety of existing preconditioning techniques for linear systems, which are based on the knowledge of the coefficient matrix A ; see, e.g., the survey in [2]. Many popular approaches rely on the sparsity of A , and aim to construct sparse preconditioners, which (partially) preserve the structure of the coefficient matrix. Since we propose to use the absolute value of A to derive a preconditioner, it is interesting to check if $|A|$ can be as sparse as A , at least after dropping small elements.

Figure 4.1 demonstrates the “spy” plots of the absolute value of the shifted Laplacian from our model problem (4.2) for different shift values, after dropping the entries of magnitude less than 1% of the largest element of $|L - c^2 I|$. We observe that although the fill increases for a larger shift value, the matrix $|L - c^2 I|$, up to the dropping, remains sparse and structured; moreover, the symmetry and positive definiteness can be preserved for sufficiently small drop tolerances. This brings a hope that methods for constructing sparse preconditioners, e.g., based on incomplete factorizations or sparse inverse approximations of A , can be extended to constructing absolute value preconditioners, even though $|A|$ is not explicitly known. Below, however, we do not follow this option, but use the MG approach instead.

4.1. The MG absolute value preconditioner. In this section, we combine principles underlying (geometric) MG methods, e.g., [6, 33], with the idea of the absolute value preconditioning to construct a preconditioner for symmetric indefinite system (4.2) with a slightly or moderately indefinite operator $L - c^2 I$, i.e., for the model problem with a relatively small shift value c^2 .




FIG. 4.1. Sparsity patterns of $|L - c^2 I|$ after dropping the entries of magnitude less than 1% of the largest element; $n = 225$.

Along with the (fine) grid of the mesh size h , underlying discrete problem (4.2), let us consider a (coarse) grid of a mesh size $H > h$. We denote the discretization of the negative Laplacian on this grid by L_H , I_H represents the identity operator of the corresponding dimension. Further, we assume that the fine-level absolute value $|L - c^2 I|$ and its inverse are not computable, while the inverse of the coarse-level operator $|L_H - c^2 I_H|$ can be efficiently constructed, e.g., by the full eigendecomposition. 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.

We suggest the following scheme as an example of the two-grid absolute value preconditioner for model problem (4.2).

ALGORITHM 4.1 (The two-grid absolute value preconditioner).

Input r , output w .

1. Presmoothing. Apply ν smoothing steps:

$$(4.3) \quad w^{(i+1)} = w^{(i)} + M^{-1}(r - Lw^{(i)}), \quad i = 0, \dots, \nu - 1, \quad w^{(0)} = 0,$$

where the (nonsingular) matrix M defines the choice of a smoother, $\nu \geq 1$. Set $w^{pre} = w^{(\nu)}$.

2. Coarse grid correction. Restrict the vector $r - Lw^{pre}$ to the coarse grid, multiply it by the inverted coarse-level absolute value $|L_H - c^2 I_H|$, and then prolongate the result back to the fine grid. This delivers the coarse grid correction, which is added to w^{pre} to obtain the corrected vector w^{cgc} :

$$(4.4) \quad w_H = |L_H - c^2 I_H|^{-1} R(r - Lw^{pre}),$$

$$(4.5) \quad w^{cgc} = w^{pre} + Pw_H,$$

where P and R are prolongation and restriction operators, respectively.

3. Postsmoothing. Apply ν smoothing steps:

$$(4.6) \quad w^{(i+1)} = w^{(i)} + M^{-*}(r - Lw^{(i)}), \quad i = 0, \dots, \nu - 1, \quad w^{(0)} = w^{cgc},$$

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

In (4.4) we assume that the coarse-grid operator $|L_H - c^2 I_H|$ is invertible, i.e., c^2 is different from any eigenvalue of L_H . The number of smoothing steps in (4.3)

and (4.6) is the same; the presmoothing is defined by the nonsingular matrix M , while the postsmoothing is delivered by M^* .

Algorithm 4.1 can be viewed as a modified two-grid cycle for linear system $|L - c^2I|z = r$. Here, the computationally expensive or unavailable absolute value of the fine-grid operator, $|L - c^2I|$, is replaced by the easily accessible negative Laplacian $L = |L|$ at smoothing steps (4.3) and (4.6), as well as in the restricted residual in (4.4). The operator $L - c^2I$ appears only at the coarse grid (4.4).

If used repeatedly, Algorithm 4.1 does not solve any linear system. However, as seen later, its use as a preconditioner (along with the respective MG extension described below) significantly accelerates the convergence of an iterative scheme applied to model problem (4.2) with a relatively small shift value; moreover, the convergence is independent of the mesh size.

Two-grid Algorithm 4.1 implicitly constructs a mapping $r \mapsto w = T_{tg}r$, where the operator $T = T_{tg}$ has the following structure:

$$(4.7) \quad T_{tg} = (I - M^{-*}L)^\nu P |L_H - c^2I_H|^{-1} R (I - LM^{-1})^\nu + F,$$

with $F = L^{-1} - (I - M^{-*}L)^\nu L^{-1} (I - LM^{-1})^\nu$. The fact that the constructed preconditioner $T = T_{tg}$ is SPD follows directly from the observation that the first term in (4.7) is symmetric positive semidefinite, provided that $P = \alpha R^*$ for some nonzero scalar α , while the second term F is symmetric and positive definite, if the spectral radii of $I - M^{-1}L$ and $I - M^{-*}L$ are less than 1. The latter condition, in fact, requires the pre- and postsmoothing iterations (4.3) and (4.6) to represent convergent methods for system (4.2) with $c = 0$ and $b = r$ (i.e., for the discrete Poisson's equation) on their own. We note that the above argument for the operator $T = T_{tg}$ to be SPD essentially repeats the corresponding pattern to justify symmetry and positive definiteness of a two-grid preconditioner applied within an iterative scheme, e.g., the preconditioned conjugate gradient method (PCG), to solve a system of linear equations with an SPD coefficient matrix; see, e.g., [5, 30].

Now let us consider a hierarchy of $m + 1$ grids numbered by $l = m, m - 1, \dots, 0$ with the corresponding mesh sizes $\{h_l\}$ in the decreasing order ($h_m = h$ corresponds to the finest, and h_0 to the coarsest, grid). For each level l we define the discretization $L_l - c^2I_l$ of the differential operator in (4.1), where L_l is the discrete negative Laplacian on grid l , and I_l is the identity of the same size.

In order to extend the two-grid absolute value preconditioner given by Algorithm 4.1 to the *multigrid*, instead of inverting the absolute value $|L_H - c^2I_H|$ in step 2 (formula (4.4)), we recursively apply the algorithm to the restricted vector $R(r - Lw^{pre})$. This pattern is then followed, in the V-cycle "fashion", on all levels, with the exact inversion of the absolute value of the shifted discrete negative Laplace operator on the coarsest grid.

If started from the finest grid $l = m$, the following scheme gives the multilevel extension of the two-grid absolute value preconditioner defined by Algorithm 4.1. We note that the subscript l is introduced to match the occurring quantities to the corresponding grid.

ALGORITHM 4.2 (AVP-MG(r_l): the MG absolute value preconditioner).

Input r_l , *output* w_l .

1. Presmoothing. Apply ν smoothing steps:

$$(4.8) \quad w_l^{(i+1)} = w_l^{(i)} + M_l^{-1}(r_l - L_l w_l^{(i)}), \quad i = 0, \dots, \nu - 1, \quad w_l^{(0)} = 0,$$

where the (nonsingular) matrix M_l defines the choice of a smoother on level l , $\nu \geq 1$. Set $w_l^{pre} = w_l^{(\nu)}$.

2. Coarse grid correction. Restrict the vector $r_l - L_l w_l^{pre}$ to the grid $l - 1$. If $l = 1$, then multiply the restricted vector by the inverted coarse-level absolute value $|L_0 - c^2 I_0|$,

$$(4.9) \quad w_0 = |L_0 - c^2 I_0|^{-1} R_0 (r_1 - L_1 w_1^{pre}), \text{ if } l = 1.$$

Otherwise, recursively apply AVP-MG to the restricted vector,

$$(4.10) \quad w_{l-1} = \text{AVP-MG}(R_{l-1} (r_l - L_l w_l^{pre})), \text{ if } l > 1.$$

Prolongate the result back to the fine grid. This delivers the coarse grid correction, which is added to w_l^{pre} to obtain the corrected vector w_l^{cgc} :

$$(4.11) \quad w_l^{cgc} = w_l^{pre} + P_l w_{l-1},$$

where w_{l-1} is given by (4.9)–(4.10). The operators R_{l-1} and P_l define the restriction from the level l to $l - 1$ and the prolongation from the level $l - 1$ to l , respectively.

3. Postsmoothing. Apply ν smoothing steps:

$$(4.12) \quad w_l^{(i+1)} = w_l^{(i)} + M_l^{-*} (r_l - L_l w_l^{(i)}), \quad i = 0, \dots, \nu - 1, \quad w_l^{(0)} = w_l^{cgc},$$

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

The described MG absolute value preconditioner implicitly constructs a mapping $r \mapsto w = T_{mg} r$, where the operator $T = T_{mg}$ has the following structure:

$$(4.13) \quad T_{mg} = (I - M^{-*} L)^\nu P T_{mg}^{(m-1)} R (I - L M^{-1})^\nu + F,$$

with F as in (4.7) and $T_{mg}^{(m-1)}$ defined according to the recursion below,

$$(4.14) \quad \begin{aligned} T_{mg}^{(l)} &= (I_l - M_l^{-*} L_l)^\nu P_l T_{mg}^{(l-1)} R_{l-1} (I_l - L_l M_l^{-1})^\nu + F_l, \\ T_{mg}^{(0)} &= |L_0 - c^2 I_0|^{-1}, \quad l = 1, \dots, m - 1, \end{aligned}$$

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

In (4.13), we skip the subscript in the notation for the quantities associated with the finest level $l = m$. The structure of the multilevel preconditioner $T = T_{mg}$ in (4.13) is similar to that of the two-grid preconditioner $T = T_{tg}$ in (4.7), with $|L_H - c^2 I_H|^{-1}$ replaced by the recursively defined operator $T_{mg}^{(m-1)}$ in (4.14). If the assumptions on the fine-grid operators M , M^* , R and P , sufficient to ensure that the two-grid preconditioner in (4.7) is SPD, remain valid throughout the coarser levels, i.e., $P_l = \alpha R_{l-1}^*$, and the spectral radii of $I_l - M_l^{-1} L_l$ and $I_l - M_l^{-*} L_l$ are less than 1, $l = 1, \dots, m - 1$, then the symmetry and positive definiteness of the MG preconditioner $T = T_{mg}$ in (4.13) can be obtained from the same property of the two-grid operator through relations (4.14). We remark that preconditioner (4.13)–(4.14) is non-variable, i.e., it can preserve the global optimality of an underlying solver.

We do not theoretically prove that (2.3), with $A = L - c^2 I$, holds independently of h , thus formally establishing that Algorithm 4.2 represents an absolute value preconditioner. Instead, we verify (2.3) numerically in Section 4.3.

4.2. Alternative SPD preconditioners. We consider two *known* strategies for constructing SPD preconditioners for model problem (4.2) here. The first one is based on the inverted Laplacian, while the second one relies on the Bunch-Parlett factorization of the coefficient matrix. In Section 4.3, we compare preconditioners, resulting from these approaches, to the absolute value preconditioner of Algorithm 4.2.

4.2.1. The inverted Laplacian preconditioner. The absolute value preconditioner given by Algorithm 4.2 can be viewed as a modified MG V-cycle for equation $|L - c^2 I|z = r$, where the absolute value of the shifted discrete negative Laplace operator on finer levels is replaced by the computationally inexpensive discrete negative Laplacian. Alternatively, Algorithm 4.2 can be interpreted as a modified MG V-cycle for the discrete Poisson's equation. Specifically, if coarse-level step (4.9) in Algorithm 4.2 is replaced by

$$(4.15) \quad w_0 = (L_0)^{-1} R_0 (r_1 - L_1 w_1^{pre}),$$

then one gets the standard V-cycle for the discrete Poisson's equation, which delivers an approximation of L^{-1} .

The idea of using an (approximate) inverse of the discrete negative Laplacian L as a preconditioner for the corresponding Helmholtz problem is well known. Introduced in Turkel et al. [1], it remains an object of active research; see, e.g., [10, 21, 34]. Since linear system (4.2) formally corresponds to the discrete Helmholtz equation, the strategy based on (approximately) inverting L represents one of the few established approaches to construct SPD preconditioners for (4.2). We use this approach in the numerical experiments of Section 4.3.

4.2.2. Preconditioning based on the Bunch-Parlett factorization. This preconditioning is suggested in [14], primarily, in the context of linear systems arising in optimization. It is based on decomposing the matrix A in (1.1) into the product

$$(4.16) \quad A = Q^* U^* D U Q,$$

where Q is a permutation, U is upper triangular, and D is (symmetric) block diagonal with block sizes 1 or 2. Factorization (4.16) is often referred to as the *Bunch-Parlett factorization*, and is known to exist for any symmetric A ; see [7].

Given decomposition (4.16) of the coefficient matrix and the spectral factorization $D = W \Theta W^*$ of the block diagonal term, a “*perfect*” SPD preconditioner for problem (1.1) is defined in [14] as

$$(4.17) \quad T = (Q^* U^* \bar{D} U Q)^{-1}, \quad \bar{D} = W |\Theta| W^*,$$

where the orthogonal W and the diagonal $\Theta = \text{diag}\{\theta_j\}$ are the matrices of eigenvectors and eigenvalues of D , respectively; $|\Theta| = \text{diag}\{|\theta_j|\}$. One can check, using (4.16), that the preconditioned matrix TA , with a nonsingular symmetric indefinite A and SPD T as in (4.17), has exactly two distinct eigenvalues: -1 and 1 .

Following the argument in the proof of Theorem 2.1, this observation allows us to conclude that choice (4.17) of the preconditioner T enforces minimal residual method (1.3)–(1.4) to converge to the exact solution of (1.1) in at most two steps. Thus, along with the preconditioner $T = |A|^{-1}$ introduced in Section 2, expression (4.17) represents an alternative instance of an *ideal* SPD preconditioner for a symmetric indefinite linear system. Both of these approaches guarantee the location of all eigenvalues of the preconditioned matrix at -1 and 1 , but the corresponding

eigenspaces are not the same, i.e., the resulting matrices TA are different for each of the two options.

If the coefficient matrix A in (1.1) is very large, then the computational cost of Bunch-Parlett factorization (4.16) for constructing exact ideal preconditioner (4.17) is prohibitive. The strategy chosen in [14] relies on the structure of the underlying problem, associating A with its approximation, such that the latter admits an efficient Bunch-Parlett factorization. The obtained factors are then used to construct a preconditioner, which can be viewed as an approximation of (4.17).

We propose fitting the construction of *preconditioners based on Bunch-Parlett factorization*, i.e., those resembling (4.17), into the same MG framework as in Section 4.1. In particular, for model problem (4.2), this can be done by following the steps of Algorithm 4.2, with the only difference that, instead of inverting the matrix absolute value on the coarsest grid, one computes the Bunch-Parlett factorization of the coarse-level operator, i.e., $L_0 - c^2 I_0 = Q_0^* U_0^* D_0 U_0 Q_0$, and replaces (4.9) by

$$(4.18) \quad w_0 = (Q_0^* U_0^* \bar{D}_0 U_0 Q_0)^{-1} R_0 (r_1 - L_1 w_1^{pre}), \quad \bar{D}_0 = W_0 |\Theta_0| W_0^*,$$

where $D_0 = W_0 \Theta_0 W_0^*$ is the eigendecomposition of the block diagonal D_0 . Similar to (4.13)–(4.14), it can be verified that the resulting preconditioner is SPD under mild assumptions on the restriction, prolongation, and smoothers. In Section 4.3, we compare this preconditioner to the MG absolute value and the inverted Laplacian preconditioners described above.

4.3. Numerical experiments. In this section, we numerically examine the MG absolute value preconditioner given by Algorithm 4.2. As an underlying iterative solver we use the MATLAB implementation of the PMINRES algorithm.

Let us highlight that the experiments below aim mainly at testing the preconditioning approach presented in this paper. In our numerical examples, we consider linear system (4.2) with relatively small shifts c^2 , i.e., the coefficient matrix $L - c^2 I$ is only slightly or moderately indefinite. As a proof of concept, we provide an actual example of an absolute value preconditioner for such a problem, however, at the current stage of the research, we do not attempt to deal with linear systems, which are *highly* indefinite, i.e., where c^2 is reasonably large. Indeed, preconditioner in Algorithm 4.2 is designed to handle only moderately small shifts, since the shift appears only at the coarse grid correction and does not affect the choice of the coarse grid. We leave the question of constructing absolute value preconditioners for highly indefinite symmetric linear systems, which result from known physical applications, for further investigation beyond the scope of this paper.

Figure 4.2 illustrates the PMINRES performance applied to model problem (4.2) with different SPD preconditioners. In particular, we compare the MG absolute value preconditioner in Algorithm 4.2 (“AVP-MG”) to the MG inverted Laplacian preconditioner discussed in Section 4.2.1 (“Laplace-MG”) and the MG preconditioner based on the Bunch-Parlett factorization suggested in Section 4.2.2 (“BP-MG”). Runs corresponding to the unpreconditioned iterative scheme (MINRES) and to PMINRES with the *exactly* inverted (using MATLAB “backslash” operator) negative Laplacian as a preconditioner (“Laplace”) are also presented in the figure. In these tests, linear system (4.2) corresponds to problem (4.1) discretized on the grid of size $h = 2^{-7}$ (the fine problem size $n = (2^7 - 1)^2 \approx 1.6 \times 10^4$); $c^2 = 100, 200, 300$ and 400 . The right-hand side vectors b , as well as initial guesses x_0 , are randomly chosen (same for each shift value).




FIG. 4.2. Comparison of several SPD preconditioners for PMINRES applied to the model problem of the size $n = (2^7 - 1)^2 \approx 1.6 \times 10^4$.

In the experiments in Figure 4.2, as well as in the tests below, we define the components for all three of our MG algorithms, i.e., for the absolute value preconditioner (Algorithm 4.2), the inverted Laplacian preconditioner (Algorithm 4.2 with (4.9) replaced by (4.15)), and the preconditioner based on the Bunch-Parlett factorization (Algorithm 4.2 with (4.9) replaced by (4.18)), as following: ω -damped Jacobi iteration as a (pre- and post-) smoother with the damping parameter $\omega = 4/5$, standard coarsening scheme (i.e., $h_{l-1} = 2h_l$), full weighting for the restriction, and piecewise multilinear interpolation for the prolongation; see, e.g., [6, 33]. Unless otherwise is explicitly stated, the coarsest grid is of the size 2^{-4} (the coarse problem size $n_0 = 225$). The number of (pre- and post-) smoothing steps is set to one, i.e., $\nu = 1$.

We note that for sufficiently “rough” coarse grids, the cost of the coarse grid computations for all of the considered MG algorithms, i.e., computations in (4.9), (4.15), and (4.18), is negligible relative to the cost of smoothing and other operations at the fine grid. This means that the costs of the corresponding V-cycles are similar for all our MG preconditioners tested. Thus, our numerical comparison, based on the number of iterations, is representative.

As observed from Figure 4.2, PMINRES with MG absolute value preconditioner noticeably outperforms the method with other SPD preconditioners, which are based on the inverted Laplacian and the Bunch-Parlett factorization. We see that, for the model problem with smaller shifts, the preconditioning approach based on the Bunch-Parlett factorization results in a reasonable SPD preconditioner. However, in

most cases, it is inferior to all other preconditioners considered in Figure 4.2. This behavior may be related to the fact that the inverted negative Laplacian $T = L^{-1}$ and the absolute value $T = |L - c^2 I|^{-1}$ preconditioners share the same eigenvectors with the coefficient matrix $A = L - c^2 I$, while the “perfect” preconditioner (4.17), which is based on the Bunch-Parlett factorization, does not.

Figure 4.2 reveals an interesting phenomenon—the convergence history of the PMINRES method preconditioned by Algorithm 4.2 can be clearly separated into the initial relatively slow convergence phase and the second much faster convergence phase. Moreover, the duration of the first, slower, phase approximately corresponds to the number of negative eigenvalues of $A = L - c^2 I$. PMINRES with MG absolute value preconditioner, given by Algorithm 4.2 with (4.9), significantly accelerates in the second phase, while MINRES with the traditional MG preconditioner for the Laplacian, given by Algorithm 4.2 with (4.15), does not.

Although we do not report the corresponding tests here, it is also observed that the increase in the number of smoothing steps ($\nu = 2$) in Algorithm 4.2 slightly improves the quality of the resulting absolute value preconditioner, however, correspondingly affects its computational cost; see [36]. The same observation applies to the MG preconditioner based on the Bunch-Parlett factorization.




FIG. 4.3. Performance of the MG absolute value preconditioners for the model problem with different shift values. The problem size $n = (2^7 - 1)^2 \approx 1.6 \times 10^4$. The number of negative eigenvalues varies from 0 to 75.

Figure 4.2 demonstrates that the quality of the MG absolute value preconditioner deteriorates with the increase of the shift value. Figure 4.3, which plots the number of PMINRES iterations performed to decrease the 2-norm of the initial error by 10^{-8} for a given value of c^2 , resembles the speed of this deterioration. We see that, for larger shift values, it may be desirable to have a grid of a higher resolution on the coarsest level in Algorithm 4.2.

Let us remind the reader that Algorithm 4.2 with formula (4.9) replaced by (4.15) is a standard geometric MG for the Laplacian. Such a solver is known to have optimal costs, i.e., linearly proportional to n . We have already discussed that Algorithm 4.2 with (4.9) has essentially the same cost as Algorithm 4.2 with (4.15). Therefore, if, in addition, the number of iterations in the iterative solver preconditioned with Algorithm 4.2 does not depend on the problem size, the overall scheme is optimal.

	$h = 2^{-5}$	$h = 2^{-6}$	$h = 2^{-7}$	$h = 2^{-8}$	$h = 2^{-9}$	$h = 2^{-10}$
$c^2 = 100$	14	14	15	14	14	14
$c^2 = 200$	21	21	21	21	21	21
$c^2 = 300$	31	32	31	32	32	30
$c^2 = 400$	40	40	40	39	40	40

TABLE 4.1

Mesh-independent convergence of PMINRES with the MG absolute value preconditioner.

We verify this in Table 4.1, which shows the mesh-independence of the convergence of PMINRES with the MG absolute value preconditioner given by Algorithm 4.2. The rows of the table correspond to the shift values c^2 , while the columns match the mesh size h underlying discrete problem (4.2). The cell in the intersection contains the number of steps performed to achieve the decrease by the factor 10^{-8} in the error 2-norm. The size of the coarsest grid has been kept the same throughout all runs, i.e., $h_0 = 2^{-4}$ ($n_0 = 225$). We conclude that the convergence does not slowdown with the decrease of h , thus, PMINRES preconditioned by Algorithm 4.2 is optimal.

Let us recall that the mesh-independent convergence of PMINRES, as that captured by Table 4.1, is theoretically governed by the ratio δ_1/δ_0 of the constants in (2.3), which is an upper bound for the spectral condition number $\kappa(T|A|) = \mu_n/\mu_1$ of $T|A|$. The definition of an absolute value preconditioner requires that this ratio must be independent of the mesh size, or, equivalently, of the problem size. Theorem 2.4 then implies that δ_1/δ_0 is also a mesh-independent upper bound for the spectral condition number of the preconditioned matrix TA . We do not theoretically bound $\kappa(T|A|)$, but rather evaluate it numerically.

	$h = 2^{-5}$	$h = 2^{-6}$	$h = 2^{-7}$
$c^2 = 100$	2.31	3.06	3.33
$c^2 = 200$	3.46	4.52	6.03
$c^2 = 300$	284.53	422.03	465.7
$c^2 = 400$	148.78	136.07	135.27

TABLE 4.2

Spectral condition numbers of $T|L - c^2I|$ for different mesh sizes and shifts; $h_0 = 2^{-4}$.

In Table 4.2, we report explicitly computed (using the MATLAB EIG function) condition numbers of $T|A|$, where $A = L - c^2I$, for different mesh sizes and shifts. Due to computer memory limitations, the smallest grid size we handle is $h = 2^{-7}$. This is apparently not small enough for $\kappa(T|A|)$ to stabilize as a function of $h \rightarrow 0$. Although the existence of a mesh-independent bound on the condition number, as required by Definition 2.2, seems plausible for a fixed c^2 . The case $c^2 = 300$ in Table 4.2 is special since the coarse grid matrix $|L_0 - c^2I_0|$ with $h_0 = 2^{-4}$ has an eigenvalue 299.9, too close to $c^2 = 300$. This results in a poor behavior of the coarse grid solve with the coefficient matrix $|L_0 - c^2I_0|$ in step (4.9). The table shows a moderate growth of $\kappa(T|L - c^2I|)$ with respect to h for $c^2 = 100$ and 200. At the same time, in Table 4.1 we observe no slowdown for these cases. This effect may be attributed to clustering of the spectra of $T|A|$ and TA . To check it, we compute the

spectra.




FIG. 4.4. 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); where $h = 2^{-7}$ and $h_0 = 2^{-4}$.

In Figure 4.4, we plot the eigenvalues of $T|L - c^2 I|$ and $T(L - c^2 I)$ for different shift values using the logarithmic scale. As suggested by Theorem 2.5, clusters of eigenvalues of $T|L - c^2 I|$ are preserved in the spectrum of the preconditioned matrix $T(L - c^2 I)$. First, we observe that almost all eigenvalues of $T(L - c^2 I)$ are clustered around -1 and 1 . The few eigenvalues that fall outside of the clusters may suddenly increase the condition number $\kappa(T|A|)$ with varying h , as seen in Table 4.2. However, they do not noticeably affect PMINRES convergence, as Table 4.1 and Figure 4.2 show. Second, both the clustering and the condition number $\kappa(T|A|)$ worsen as c^2 increases from 100 to 400, as should be expected. This explains the slowdown of PMINRES in Figure 4.2 and Table 4.1 with the increase of c^2 .

Acknowledgments. The authors thank Michele Benzi for interesting and useful comments; in particular, for pointing to the reference [14] and suggesting to consider the extension of the preconditioning based on the Bunch-Parlett factorization to the MG framework. The authors also thank Yvan Notay, who proposed checking the sparsity of $|A|$ for our model problem, which resulted in Figure 4.1.

REFERENCES

- [1] A. BAYLISS, C. I. GOLDSTEIN, AND E. TURKEL, *An iterative method for the Helmholtz equation*, J. Comput. Phys., 49 (1983) pp. 443–457.
- [2] M. BENZI, *Preconditioning Techniques for Large Linear Systems: A Survey*, J. Comput. Phys., 182 (2002), pp. 418–477.
- [3] M. BENZI, G. H. GOLUB, AND J. LIESEN, *Numerical solution of saddle point problems*, Acta Numerica, 14 (2005), pp. 1–137.
- [4] A. BORIÇI, *A Lanczos Approach to the Inverse Square Root of a Large and Sparse Matrix*, J. Comput. Phys., 162 (2000), pp. 123–131.
- [5] J. H. BRAMBLE AND X. ZHANG, *The analysis of multigrid methods*, in Handb. Numer. Anal., Solution of Equation in \mathbb{R}^n (Part 3), Techniques of Scientific Computing (Part 3), Elsevier, pp. 173–415, 2000.
- [6] W. L. BRIGGS, V. E. HENSON, AND S. F. MCCORMICK, *A Multigrid Tutorial*, SIAM, 2000.

- [7] 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.
- [8] D. CALVETTI AND L. REICHEL, *An adaptive Richardson iteration method for indefinite linear systems*, Numer. Algorithms, 12 (1996), pp. 125–149.
- [9] E. G. D'YAKONOV, *Optimization in solving elliptic problems*, CRC-Press, 1996.
- [10] 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.
- [11] 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.
- [12] D. G. FEINGOLD AND R. S. VARGA, *Block diagonally dominant matrices and generalizations of the Gerschgorin circle theorem*, Pac. J. Math., 12 (1962), pp. 1241–1250.
- [13] 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.
- [14] 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.
- [15] G. H. GOLUB AND C. F. VAN LOAN, *Matrix Computations*, The Johns Hopkins University Press, 1996.
- [16] A. GREENBAUM, *Iterative Methods for Solving Linear Systems*, SIAM, 1997.
- [17] 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.
- [18] N. J. HIGHAM, *Functions of Matrices: Theory and Computation*, SIAM, Philadelphia, PA, USA, 2008.
- [19] R. A. HORN AND C. R. JOHNSON, *Topics in Matrix Analysis*, Cambridge University Press, New York, 1994.
- [20] A. V. KNYAZEV, I. LASHUK, *Steepest descent and conjugate gradient methods with variable preconditioning*, SIAM J. Matrix Anal. Appl., 29 (4), pp. 1267–1280.
- [21] A. L. LAIRD AND M. B. GILES, *Preconditioned iterative solution of the 2D Helmholtz equation*, Oxford University Computing Laboratory, Oxford, UK, 02/12, 2002.
- [22] V. I. LEBEDEV, *Iterative methods for solving operator equations with spectrum contained in several intervals*, Zhurnal Vychislitel'noi Matematiki i Matematicheskoi Fiziki, 9 (1969), pp. 17–24, English transl. in USSR Comput. Math. Math. Phys. 9 (1972).
- [23] M. F. MURPHY, G. H. GOLUB, AND A. J. WATHEN, *A Note on Preconditioning for Indefinite Linear Systems*, SIAM J. Sci. Comput., 21 (1999), pp. 1969–1972.
- [24] J. NOCEDAL AND S. WRIGHT, *Numerical Optimization*, Springer, 1999.
- [25] C. C. PAIGE AND M. A. SAUNDERS, *Solution of Sparse Indefinite Systems of Linear Equations*, SIAM J. Numer. Anal., 12 (1975), pp. 617–629.
- [26] B. N. PARLETT, *The symmetric eigenvalue problem*, SIAM, 1998.
- [27] Y. SAAD, *Iterative Solution of Indefinite Symmetric Linear Systems by Methods Using Orthogonal Polynomials over Two Disjoint Intervals*, SIAM J. Numer. Anal., 20 (1983), pp. 784–811.
- [28] 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.
- [29] 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.
- [30] O. TATEBE, *The multigrid preconditioned conjugate gradient method*, in Sixth Copper Mountain Conference on Multigrid Methods, pp. 621–634, 1993.
- [31] M. P. TETER AND D. C. ALLAN, *Solution of Schrödinger's equation for large systems*, Phys. Rev. B, 40 (1989), pp. 12255–12263.
- [32] C. TRETTER, *Spectral theory of block operator matrices and applications*, Imperial College Press, 2008.
- [33] U. TROTTEBERG, C. W. OOSTERLEE, AND A. SCHÜLLER, *Multigrid*, Academic Press, 2001.
- [34] 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.
- [35] 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.
- [36] E. VECHARYNSKI, *Preconditioned Iterative Methods for Linear Systems, Eigenvalue and Singular Value Problems*, PhD thesis, University of Colorado Denver, 2010.
- [37] 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.