DELFT UNIVERSITY OF TECHNOLOGY

REPORT 03-01

On a class of preconditioners for solving the Helmholtz equation

Y. A. Erlangga, C. Vuik, C. W. Oosterlee

ISSN 1389-6520

Reports of the Department of Applied Mathematical Analysis

Delft 2003

On a class of preconditioners for solving the Helmholtz equation*

Y. A. Erlangga, C. Vuik, C. W. Oosterlee

January 21, 2003

Abstract

In 1983, Bayliss, Goldstein, and Turkel [2] proposed a preconditioner based on the Laplace operator for solving the discrete Helmholtz equation efficiently with CGNR. The preconditioner is especially effective for low wavenumber cases where the linear system is slightly indefinite. Laird [11] proposed a preconditioner where an extra term is added to the Laplace operator. This term is similar to that in the Helmholtz equation but with reversed sign. In this paper, both approaches are further generalized to a new class of preconditioners, the so-called "Shifted Laplace" preconditioners of the form $\Delta\phi - \alpha k^2\phi$ with $\alpha \in \mathbb{C}$. Numerical experiments for various wavenumbers indicate the effectiveness of the preconditioner in terms of numbers of iterations and arithmetic operations. The preconditioner is evaluated in combination with GMRES, Bi-CGSTAB, and CGNR.

Keywords: Helmholtz equation, preconditioners, GMRES, Bi-CGSTAB

1 Introduction

In this paper, the time-harmonic wave equation in 2D heterogeneous media is solved numerically. The underlying equation governs wave propagations and scattering phenomena arising in acoustic problems in many areas, e.g., in aeronautics, marine technology, geophysics, and optical problems. In particular, we look for solutions of the Helmholtz equation discretized by using finite difference discretizations. Since the number of gridpoints per wavelength should be sufficiently fine to result in acceptable solutions, for very high wavenumbers the discrete problem becomes extremely large, prohibiting the use of direct

*This research is financially supported by the *Dutch Ministry of Economic Affairs* under the project BTS01044 "Rigorous modelling of 3D wave propagation in inhomogeneous media for geophysical and optical problems"

methods. Iterative methods are the interesting alternative. However, Krylov subspace methods are not competitive without a good preconditioner. In this paper, we consider a class of preconditioners to improve the convergence of the Krylov subspace methods.

Various authors contributed to the development of powerful preconditioners for Helmholtz problems. The work in [2] and the follow-up investigation in [9] can be considered as the start for the class of preconditioners we are interested in. A generalization has been recently proposed in [11]. In [2, 9, 11], the preconditioners are constructed based on the Laplace operator. In [11], this operator is perturbed by a real-valued linear term. This surprisingly straightforward idea leads to very satisfactorily convergence. Furthermore, the preconditioning matrix allows the use of SSOR, ILU, or multigrid to approximate the inversion within an iteration.

In this paper, we will generalize the approach in [2, 9, 11]. We give theoritical and numerical evidence that introducing a *complex* perturbation to the Laplace operator can result in a better preconditioner than using a real-valued perturbation. We call the resulting class of preconditioners "Shifted Laplace" preconditioners. This class of preconditioners is simple to construct and is easy to extend to inhomogeneous media.

There are various other types of preconditioners for general indefinite linear systems, e.g. [6, 8, 12, 14]. In particular for Helmholtz problems, [8] proposed a class of preconditioners (so called AILU) based on a parabolic factorization of the Helmholtz operator. In [12] another approach is pursued by perturbing the real part of the matrix to make it less indefinite. An interesting alternative is also described in [14], where a preconditioner based on the separation of variables is proposed. This preconditioner effectively accelerates the convergence for high wavenumbers.

This paper is organized as follows. In Section 2 we describe the mathematical model and the discretization used to solve wave propagation problems. Iterative methods used to solve the resulting linear system and the preconditioner will be discussed in Sections 3 and 4 respectively. In Section 5, we present the Shifted Laplace preconditioners and show theoretically the convergence of this type of preconditioner. Numerical results are then presented in Section 6.

2 Mathematical model

We solve wave propagations in a two dimensional medium with inhomogeneous properties in a unit (scaled) domain governed by the Helmholtz equation

$$\Delta\phi + k^2(x,y)\phi = f,\ \Omega = [0,1]^2, \tag{1}$$

where $\Delta \equiv \partial^2/\partial x^2 + \partial^2/\partial y^2$, the Laplace operator, and $k(x,y) \in \mathbb{R}$ is the wavenumber, which may depend on the spatial position in the domain. We consider a so-called "open problem", i.e., outgoing waves penetrate at least at one boundary without (spurious) reflections. To satisfy this condition, a radiation-type condition is imposed. Several formulations have been developed to model the non-reflecting condition at the boundary [1, 3, 4]. In

this paper, the first order Sommerfeld condition is chosen of the form

$$\frac{\partial \phi}{\partial n} - \mathrm{i}k\phi = 0, \text{ on a part of } \Gamma = \partial\Omega \tag{2}$$

with n an outward direction normal to the boundary. Eventhough (2) may not be sufficiently accurate for inclined outgoing waves [4], it is state-of-the-art in industrial codes, easy to implement in our discretization, and requires only a few gridpoints. We anticipate for possible reflections by considering a sufficiently large domain enabling any wave reflections to be immediately damped out and therefore be localized in the neighborhood of the boundaries.

To find numerical solutions of (1), the equation is discretized using the second-order difference scheme, in x-direction:

$$\frac{\partial^2 \phi}{\partial x^2} = \frac{1}{\Delta x^2}\left(\phi_{i-1} - 2\phi_i + \phi_{i+1}\right) + \mathcal{O}(\Delta x^2), \tag{3}$$

and similar in y-direction. The first order derivative in (2) is discretized with the first order forward scheme

$$\frac{\partial \phi}{\partial n} = \frac{1}{\Delta n}\left(\phi_{i+1} - \phi_i\right). \tag{4}$$

Substituting (3) and (4) into (1) and (2), one obtains a linear system

$$Ap = b, \; A \in \mathbb{C}^{N\times N}, \tag{5}$$

where A is a large, sparse symmetric matrix, and N is the number of gridpoints. Matrix A is complex-valued and indefinite for large values of k. Throughout this paper, we say "A is indefinite" if A has eigenvalues with a positive real part and eigenvalues with negative real part [6].

3 Krylov subspace method

For a large, but sparse matrix, Krylov subspace methods are very popular. The methods are developed based on a construction of iterants in the subspace

$$\mathcal{K}^j(A, r_0) = \mathrm{span}\{r_0, Ar_0, A^2 r_0, \cdots, A^{j-1} r_0\}, \tag{6}$$

where $\mathcal{K}^j(A, r_0)$ is the j-th Krylov subspace associated with A and r_0 (see, e.g., [16]).

The basic algorithm within this class is the Conjugate Gradient method (CG) which has the nice properties that it uses only three vectors in memory and minimizes the error in the A-norm. However, the algorithm mainly performs well if the matrix A is symmetric, and positive definite. In cases where one of these two properties is violated, CG may break down. For indefinite linear systems, CG can be applied to the normal equations since the resulting linear system becomes (positive) definite. Upon application of CG to the normal equations, CGNR [16] results. Using CGNR, the iterations are guaranteed to converge.

The drawback is that the condition number of the normal equations equals the square of the condition number of A, slowing down the convergence drastically.

Some algorithms with short recurrences but without the minimizing property are constructed based on the bi-Lanczos algorithm [16]. Within this class, BiCG [5] exists and its modifications: CGS [19] and Bi-CGSTAB [20]. In BiCG, the Krylov subspace is constructed from the orthogonalization of two residual vectors based on actual matrix A and its transpose A^T. Accordingly, one extra matrix/vector multiplication and one transpose operation are needed. In CGS, the extra transpose operation can be avoided. One can accelerate the convergence by squaring the polynomial. Whenever the convergence is smooth, CGS converges twice as fast as BiCG. However, if the BiCG iteration diverges, CGS also diverges twice as fast as BiCG. To stabilize CGS, rather than taking the square of the polynomial, another polynomial can be chosen and multiplied with the polynomial of A. This results in Bi-CGSTAB. In many cases, Bi-CGSTAB exhibits a smooth convergence behavior and often converges faster than CGS. Also within this class are QMR [7] and COCG [21].

MINRES [13] can also be used to solve indefinite symmetric linear systems, as well as its generalization to the nonsymmetric case, GMRES [17, 16]. Both algorithms have the minimization property but GMRES uses long reccurences. GMRES has the advantage that theoretically the algorithm does not break down unless convergence has been reached. The main problem in GMRES is that the amount of storage increases as the iteration number increases. Therefore, the application of GMRES may be limited by the computer storage. To remedy this problem, a restarted version, GMRES(m), can be utilized [17]. Since restarting removes the previous convergence history, GMRES(m) is not guaranteed to converge. There is no specific rule to determine the restart parameter m. In cases characterized by superlinear convergence, m should often be chosen very large which makes restarting much less attractive. Another way to remedy the storage problem in GMRES is by including a so-called "inner iteration" as in GMRESR [22] and FGMRES [15].

Since the convergence theory of GMRES is well established, in our numerical experiments mainly full GMRES is used. Of course, experiments then become very restrictive (in this paper, upto $k = 30$) and for large problems, restarts become necessary. We also compute solutions using Bi-CGSTAB and compare the convergence results with GMRES. For completeness, since the underlying theory on the preconditioners is developed based on the normal equations [11], we also include the convergence results using CGNR in the last experiment.

4 Preconditioner

To improve the convergence of iterative methods, a preconditioner should be incorporated.

By left preconditioning, one solves a linear system premultiplied by a preconditioning matrix M^{-1},

$$M^{-1}Ap = M^{-1}b. \tag{7}$$

Often, right preconditioning is used, i.e.,

$$AM^{-1}\tilde{p} = b, \tag{8}$$

where $\tilde{p} = Mp$. Both preconditionings show typically a very similar convergence behavior. However, for left preconditioning GMRES computes the residuals based on the preconditioned system. In contrast, for right preconditioning GMRES computes the actual residuals. This difference may affect the stopping criterion to be used (see discussions in [16]).

The best choice for M^{-1} is the inverse of A, which is impractical. If A is SPD, one can approximate A^{-1} by one iteration of SSOR or multigrid. However, most practical wave problems result in an indefinite linear system, for which SSOR or multigrid are not guaranteed to converge (and do not converge).

In general, one can distinguish two approaches for constructing preconditioners: matrix-based and operator-based. Within the first class lie, e.g., incomplete LU (ILU) factorizations. Several ILU techniques have been developed with different choices of the tolerated fill-in in the sparsity pattern of A, e.g., zero fill-in ILU, or ILU with drop tolerance. Another different approach but falling into this category is the approximate inverse (see, e.g., [16]). An example of an operator-based preconditioner is analytic ILU (AILU) [8], which is based on the continuous Helmholtz operator.

In the next sections, we will briefly discuss some preconditioners for Helmholtz problems.

4.1 ILU preconditioner

An ILU preconditioner can be constructed by performing Gauss elimination and dropping some elements based on certain criteria. One can, e.g., drop all elements except for those in the same diagonals as the original matrix. This leads to ILU(0). ILU(p) allows fill-in in p additional diagonals. One can also drop elements which are smaller than a specified value, giving ILU(tol). In applications involving M-matrices, this class of preconditioners is sufficiently effective. However, preconditioners from this class are not effective for general indefinite problems. Reference [8] shows some results in which ILU-type preconditioners are used to solve the Helmholtz equation using QMR. For high wavenumbers k, ILU(0) converges slowly, while ILU(tol) encounters storage problems (and also slow convergence). For sufficiently high wavenumbers k, the cost to construct the ILU(tol) factors may become very high.

Instead of constructing the ILU factors from A, the Helmholtz operator $\mathcal{L}_h = \Delta + k^2$ can be used to set up *ILU-like* factors in so-called analytic ILU (AILU) [8]. Starting with the Fourier transform of the analytic operator in one direction, one constructs parabolic factors of the Helmholtz operator consisting of a first order derivative in one direction and a non-local operator. To remove the non-local operator, a localized approximation is proposed, involving optimization parameters. Finding a good approximation for inhomogeneous problems is the major difficulty in this type of preconditioner. This is because the

method is sensitive with respect to small changes in these parameters. The optimization parameters depend on $k(x, y)$.

4.2 Shifted Laplace preconditioner

Another approach is found in *not* looking for an approximate inverse of the discrete indefinite operator A, but merely looking for a form of M, for which $M^{-1}A$ has satisfactory properties for Krylov subspace acceleration. A first effort to construct a preconditioner in such a way is in [2]. An easy-to-construct $M = \Delta$ preconditioner is incorporated for CGNR. One SSOR iteration is used whenever operations involving M^{-1} are required. The subsequent work on this preconditioner with multigrid was done in [9].

Instead of the Laplace operator as the preconditioner, [11] investigates possible improvements if an extra term $-k^2$ is added to the Laplace operator. So, the Helmholtz equation with reversed sign is proposed as the preconditioner M. This preconditioner is then used in CGNR. One multigrid iteration is employed whenever M^{-1} must be computed. Instead of the normal equations, our findings suggest that GMRES can solve the preconditioned linear system efficiently in less arithmetic operations, despite of a storage problem for high k. However, the latter problem can be overcome, e.g., by applying GMRES(m) or GMRESR. Bi-CGSTAB, which is considered a good alternative except for more matrix/vector multiplications, does not perform satisfactorily [11]. (See also results in Section 6.)

In the next section, we concentrate on this type of preconditioners and present a generalization.

5 Spectral properties of Shifted Laplace preconditioners

In this section we provide some analysis to understand the performance of the Shifted Laplace preconditioners. The analysis is based on eigenvalue properties of the preconditioned system. It is often that the eigenvalue distribution can help in understanding the behavior of CG-like iterations. Since the spectra of $M^{-1}A$ and AM^{-1} are identical, we concentrate on left preconditioning.

5.1 Real Shifted-Laplace preconditioner

The preconditioners in [2, 11] can be motivated as follows. Consider the continuous 1D Helmholtz equation, subject to discretization. For simplicity, suppose that both boundary conditions are either Dirichlet or Neumann conditions.

We first consider the eigenvalues for the 1D Helmholtz operator without any preconditioning. Eigenvalues of this standard problem, denoted by λ^s, are found to be

$$\lambda_n^s = k_n^2 - k^2, \quad k_n = n\pi, \; n \in \mathbb{N}\backslash\{0\}. \tag{9}$$

In (9), k_n is the natural frequency of the system. (We use n to indicate the eigenmodes). If one considers the modulus of the eigenvalues (which in this case is simply their absolute value), it is easily seen that $|\lambda|$ becomes unbounded if either n or k are large. If the l_2-condition number $\kappa = |\lambda_{\max}/\lambda_{\min}|$ is used to evaluate the quality of eigenvalue clustering, one concludes that for any sufficiently small $\lambda_{\min}$ the condition number is extremely large.

Now, suppose an operator of the form

$$\frac{d^2}{dx^2} - \alpha k^2,\ \alpha \geq 0 \in \mathbb{R} \tag{10}$$

is used as a preconditioner, constructed with the same discretization stencil and boundary conditions. The following generalized eigenvalue problem is obtained, i.e.,

$$\left(\frac{d^2}{dx^2} + k^2\right)\phi_v = \lambda\left(\frac{d^2}{dx^2} - \alpha k^2\right)\phi_v,\ x \in [0,1] \subseteq \mathbb{R}. \tag{11}$$

For (11), we find the eigenvalues to be

$$\lambda_n = \frac{k_n^2 - k^2}{k_n^2 + \alpha k^2} = \frac{1 - (k/k_n)^2}{1 + \alpha(k/k_n)^2},\ \ n \in \mathbb{N}\backslash\{0\}. \tag{12}$$

For $n \to \infty$, $\lambda_n \to 1$, i.e. the eigenvalues are bounded above by one. Examining the low eigenmodes, for $k_n \to 0$, we obtain $\lambda \to -1/\alpha$. This eigenvalue remains below one unless $\alpha \leq 1$. The maximum eigenvalue can thus be written as

$$|\lambda_{\max}| = \max\left(|\frac{1}{\alpha}|, 1\right),\ \ \alpha \geq 0 \in \mathbb{R}. \tag{13}$$

To estimate the minimum eigenvalue, one can use a simple but rough analysis as follows. It is assumed that the minimum eigenvalue is very close (but not equal) to zero. This assumption indicates a condition $k_j \approx k$ as obtained from (12). To be more precise, let $k_j = k + \epsilon$, where ϵ is any small number. If this relation is substituted into (12), and if higher order terms are neglected, and $\epsilon k \ll k^2$ is assumed, then we find

$$\lambda_{\min} = \frac{2}{1+\alpha}\left(\frac{\epsilon}{k}\right). \tag{14}$$

From (14), the minimum eigenvalue can be very close to zero as α goes to infinity. The condition number of the preconditioned Helmholtz operator now reads

$$\kappa = \begin{cases} \frac{1}{2}(1+\alpha)k/\epsilon & \text{if } \alpha \geq 1, \\ \frac{1}{2\alpha}(1+\alpha)k/\epsilon & \text{if } 0 \leq \alpha \leq 1. \end{cases} \tag{15}$$

If $\alpha \geq 1$, κ is a monotonically increasing function with respect to α. The best choice is $\alpha = 1$, which gives minimal κ. If $0 \leq \alpha \leq 1$, κ is a monotonically decreasing function with respect to α. κ is minimal in this range if $\alpha = 1$. In the limit sense we find that

$$\lim_{\alpha \downarrow 1} \kappa = \lim_{\alpha \uparrow 1} \kappa = k/\epsilon, \tag{16}$$

which is the minimum value of κ for $\alpha \geq 0 \in \mathbb{R}$.

5.2 Generalization to complex α

The analysis on 1D Shifted Laplace preconditioners for $\alpha \in \mathbb{R}$ gives $\alpha = 1$ as the optimum case. The nice property of the real Shifted Laplace operator, at least in 1D, is that the eigenvalues have an upper bound. However, this property does not guarantee that the eigenvalues are favourably distributed. There is still the possibility that one or some eigenvalues (which are extremely small) can be very close to zero. We can improve the preconditioner by preserving the upper boundedness and at the same time shifting the minimum eigenvalue as far as possible from zero. In this section, we generalize α to be complex-valued.

Consider the minimum eigenvalue $\lambda_{\min}$ obtained from the 1D problem (14). We have shifted this eigenvalue away from zero by adding some real values to λ. In general, this addition will shift all eigenvalues, which is undesirable. An alternative is multiplying the eigenvalues by a factor. From (12) the relation between eigenvalues for $\alpha = 0$ and $\alpha = 1$ reads

$$(\lambda_{\alpha=1})_n = \frac{1}{1 + (k/k_n)^2} (\lambda_{\alpha=0})_n \,. \tag{17}$$

Equation (17) indicates that $\lambda_{\alpha=0}$ is scaled by a factor $1 + (k/k_n)^2$. Now we shift the minimum eigenvalue as far as possible away from zero. Using (14), we obtain the following relation:

$$(\lambda_{\alpha=1})_{\min} = \frac{1}{2} (\lambda_{\alpha=0})_{\min} \,. \tag{18}$$

We generalize the same process to the complex plane: shifting the eigenvalues along the complex axis away from zero. We introduce a complex coefficient of the form $\alpha + \mathrm{i}\beta$, and consider a more general complex-valued Shifted Laplace operator

$$\frac{d^2}{dx^2} - (\alpha + \mathrm{i}\beta)k^2, \ \alpha \geq 0 \in \mathbb{R}, \ \beta \in \mathbb{R}. \tag{19}$$

Eigenvalues of the premultiplied equation, denoted by λ^c, are

$$\lambda^c = \frac{k_n^2 - k^2}{k_n^2 + (\alpha + \mathrm{i}\beta)k^2} \Rightarrow |\lambda^c|^2 = \frac{(k_n^2 - k^2)^2}{(k_n^2 + \alpha k^2)^2 + \beta^2 k^2}. \tag{20}$$

Evaluating $\lambda_{\max}$ and $\lambda_{\min}$ as in (13) and (14) one finds

$$|\lambda^c_{\max}|^2 = \max\left(\frac{1}{\alpha^2 + \beta^2}, 1\right), \ \ |\lambda^c_{\min}|^2 = \frac{4}{(1+\alpha)^2 + \beta^2} \left(\frac{\epsilon}{k}\right)^2. \tag{21}$$

These results give the following condition numbers

$$\kappa^2 = \begin{cases} \frac{1}{4}\left(1 + \frac{1+2\alpha}{\alpha^2+\beta^2}\right)(k/\epsilon)^2, & \alpha^2 + \beta^2 \leq 1, \\ \frac{1}{4}\left((1+\alpha)^2 + \beta^2\right)(k/\epsilon)^2, & \alpha^2 + \beta^2 \geq 1. \end{cases} \tag{22}$$

Since $\alpha^2 + \beta^2$ is non-negative, for any given α taking the circle $\alpha^2 + \beta^2 = 1$ in the first expression in (22) provides the smallest κ^2. Likewise, for any given α, κ^2 is minimal for

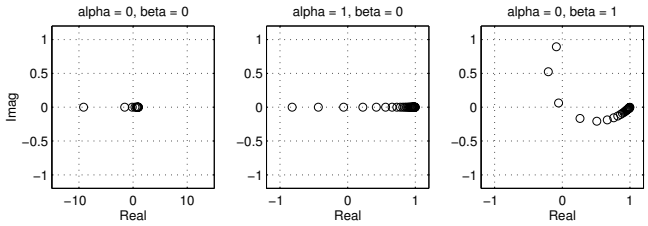

Figure 1: *Generalized eigenvalues of the continuous 1D Helmholtz equation,* $k = 10$

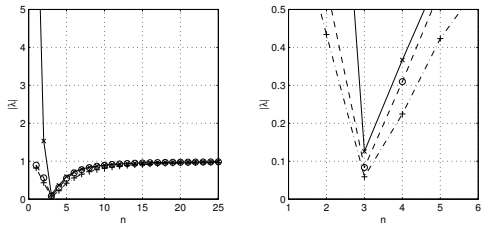

Figure 2: *The modulus of eigenvalues of the continuous 1D Helmholtz equation.* $k = 10$ *and* $h^{-1} = 100$ *for various preconditioners:* M_0 $(\times)$, M_1 $(+)$, M_i $(\circ)$

the second expression in (22) whenever $\alpha^2 + \beta^2 = 1$. (One can verify that there is no other circle giving κ^2 lower than that on the circle with radius one. This can be seen, e.g., by introducing condition $\alpha^2 + \beta^2 = 1 + \epsilon_1,\ \epsilon_1 \geq 0$). With condition $\alpha^2 + \beta^2 = 1$, κ is minimal if one takes $\alpha = 0$, implying $\beta = 1$. This combination gives the lowest condition number possible for the Shifted-Laplace preconditioner for the 1D model problem.

Figure 1 shows spectra of the preconditioned systems of the 1D Helmholtz problem using $M_{\alpha=0,\beta=0}$, $M_{\alpha=1,\beta=0}$, and $M_{\alpha=0,\beta=1}$ for our 1D problem. For simplicity, we denote these preconditioners as M_0, M_1, and M_i, respectively. Figure 1 shows that the preconditioner M_i clusters the eigenvalues stronger than M_1 and pushes the eigenvalues in the negative real plane towards the imaginary axis. This clustering may improve the performance of the preconditioned iterative methods. However, with this preconditioner there is still a possibility that some eigenvalues lie very close to zero causing unsatisfactory numerical performance. To estimate the position of these minimum eigenvalues, we consider the real part of (20). Similar as in (14), one finds that

$$\mathrm{Re}(\lambda^c_{\min}) = \epsilon/k. \tag{23}$$

This estimate is the same as the estimate for M_1 and smaller than that for M_0. However, the modulus $|\lambda^c_{\min}| = \sqrt{2}(\epsilon/k) > |\lambda^{\alpha=1}_{\min}| = \epsilon/k$ because of the imaginary shift (see Figure 2). Because of the same upper bound as M_1, M_i may perform better than M_0 and M_1.

In Figure 2, a comparison of the modulus of eigenvalues for $k = 10$ is shown, indicating boundedness of eigenvalues of M_1 and M_0 near $|\lambda| = 0$. The right-hand figure zooms in to show the minimum $|\lambda|$. Evidently, M_i has small eigenvalues with the modulus slightly

larger than M_1, but smaller than M_0.

5.3 Spectrum of the discrete Helmholtz equation

We extend the analysis to the discrete formulation of (1). Suppose that the Helmholtz equation is discretized, we arrive at the linear system $Ap = b$.

Matrix A can be splitted into two parts: the Laplace component B and the additional diagonal term $k^2 I$ so that $A = B + k^2 I$ and therefore

$$\left(B + k^2 I\right) p = b. \tag{24}$$

In this analysis, we use only Dirichlet or Neumann conditions at the boundaries in order to keep the matrix A real-valued. We precondition (24) using $M = B - (\alpha + \mathrm{i}\beta)k^2 I$, constructed with the same boundary conditions as for A. This gives

$$\left(B - (\alpha + \mathrm{i}\beta)k^2 I\right)^{-1} \left(B + k^2 I\right) p = \left(B - (\alpha + \mathrm{i}\beta)k^2 I\right)^{-1} b. \tag{25}$$

The generalized eigenvalue problem of (25) is accordingly

$$\left(B + k^2 I\right) p_v = \lambda_v \left(B - (\alpha + \mathrm{i}\beta)k^2 I\right) p_v. \tag{26}$$

Both systems (25) and (26) are indefinite if k^2 is larger than the smallest eigenvalue of B. In such a case, the convergence is difficult to estimate. Therefore, the subsequent analysis will be based on the normal equations formulation of the preconditioned matrix system (as in [11]).

Denote the eigenvalues of B as $0 < \mu_1 \leq \mu_2 \leq \cdots \leq \mu_n$. We find for the eigenvalues of the four following cases:

$$\lambda\left(A^* A\right) = \left(\mu_i - k^2\right)^2, \tag{27}$$

$$\lambda\left(\left(M_0^{-1}A\right)^* \left(M_0^{-1}A\right)\right) = \left(\frac{\mu_i - k^2}{\mu_i}\right)^2 = \left(1 - \frac{k^2}{\mu_i}\right)^2, \tag{28}$$

$$\lambda\left(\left(M_1^{-1}A\right)^* \left(M_1^{-1}A\right)\right) = \left(\frac{\mu_i - k^2}{\mu_i + k^2}\right)^2 = \left(1 - \frac{2k^2}{\mu_i + k^2}\right)^2, \tag{29}$$

$$\lambda\left(\left(M_{\mathrm{i}}^{-1}A\right)^* \left(M_{\mathrm{i}}^{-1}A\right)\right) = \left(\frac{\mu_i - k^2}{\mu_i + \mathrm{i}k^2}\right) \overline{\left(\frac{\mu_i - k^2}{\mu_i + \mathrm{i}k^2}\right)} = 1 - \frac{2\mu_i k^2}{\mu_i^2 + k^4}. \tag{30}$$

For all cases, we find for the minimal and maximal eigenvalue, if $k \in \mathbb{R}$, $0 < k^2 < \mu_1$,

$$\begin{aligned}
\lambda\left((A^*A)\right)_{\min} &= \left(\mu_1 - k^2\right)^2, \\
\lambda\left((A^*A)\right)_{\max} &= \left(\mu_n - k^2\right)^2, \qquad (31)\\
\lambda\left(\left(M_0^{-1}A\right)^*\left(M_0^{-1}A\right)\right)_{\min} &= \left(1 - \frac{k^2}{\mu_1}\right)^2, \\
\lambda\left(\left(M_0^{-1}A\right)^*\left(M_0^{-1}A\right)\right)_{\max} &= \left(1 - \frac{k^2}{\mu_n}\right)^2, \qquad (32)\\
\lambda\left(\left(M_1^{-1}A\right)^*\left(M_1^{-1}A\right)\right)_{\min} &= \left(1 - \frac{2k^2}{\mu_1 + k^2}\right)^2, \\
\lambda\left(\left(M_1^{-1}A\right)^*\left(M_1^{-1}A\right)\right)_{\min} &= \left(1 - \frac{2k^2}{\mu_n + k^2}\right)^2, \qquad (33)\\
\lambda\left(\left(M_i^{-1}A\right)^*\left(M_i^{-1}A\right)\right)_{\min} &= 1 - \frac{2\mu_1 k^2}{\mu_1^2 + k^4}, \\
\lambda\left(\left(M_i^{-1}A\right)^*\left(M_i^{-1}A\right)\right)_{\max} &= 1 - \frac{2\mu_n k^2}{\mu_n^2 + k^4}. \qquad (34)
\end{aligned}$$

Since $k^2/\mu_1 < 1$, one easily sees that

$$\lambda\left((M_0^{-1}A)^*(M_0^{-1}A)\right)_{\min} > \lambda\left((M_1^{-1}A)^*(M_1^{-1}A)\right)_{\min}.$$

As $n \to \infty$, one finds also that

$$\lim_{\mu_n \to \infty} \lambda\left((M_0^{-1}A)^*(M_0^{-1}A)\right)_{\max} = \lim_{\mu_n \to \infty} \lambda\left((M_1^{-1}A)^*(M_1^{-1}A)\right)_{\max} = 1.$$

With respect to the l_2-condition number, it becomes evident that for frequencies lower than $\sqrt{\mu_1}$, M_0 may be better than M_1. For M_{i}, one can compute that

$$\begin{aligned}
\lambda\left(\left(M_i^{-1}A\right)^*\left(M_i^{-1}A\right)\right)_{\min} / \lambda\left(\left(M_0^{-1}A\right)^*\left(M_0^{-1}A\right)\right)_{\min} &= \frac{(\mu_1 + k^2)^2}{\mu_1^2 + k^4} > 1, \\
\lim_{\mu_n \to \infty} \lambda\left(\left(M_i^{-1}A\right)^*\left(M_i^{-1}A\right)\right)_{\max} &= 1.
\end{aligned}$$

Furthermore, $\lambda\left((M_iA)^*(M_iA)\right)_{\min} \leq \lambda\left((M_0A)^*(M_0A)\right)_{\min}$. So, for frequencies lower than $\sqrt{\mu_1}$, M_0 may perform better than any other choice. However, the main focus is on high wavenumbers which are of practical interests.

If now $\mu_1 < k^2 < \mu_n$, we encounter an indefinite problem. However, for the standard A^*A, one finds

$$\begin{aligned}
&\lambda\left(A^*A\right)_{\min} = \left(\mu_{m_1} - k^2\right)^2, \text{ where } |\mu_{m_1} - k^2| \leq |\mu_i - k^2|, \ \forall i, \\
&\lambda\left(A^*A\right)_{\max} = \left(\mu_n - k^2\right)^2,
\end{aligned} \qquad (35)$$

which is always positive definite. In this case, the eigenvalues are unbounded either for large μ_n or large k. For the preconditioned system $(M_0^{-1}A)^*(M_0^{-1}A)$ one finds

$$\begin{aligned}
\lambda\left((M_0^{-1}A)^*(M_0^{-1}A)\right)_{\min} &= \left(\frac{\mu_{m_2}-k^2}{\mu_{m_2}}\right)^2, \\
&\quad \text{where } |\frac{\mu_{m_2}-k^2}{\mu_{m_2}}| \le |\frac{\mu_i-k^2}{\mu_i}|,\ \forall i, \\
\lambda\left((M_0^{-1}A)^*(M_0^{-1}A)\right)_{\max} &= \max\left(\left(\frac{\mu_n-k^2}{\mu_n}\right)^2, \left(\frac{\mu_1-k^2}{\mu_1}\right)^2\right). \qquad (36)
\end{aligned}$$

In this case, there will be a possible boundedness for large μ_n, i.e., for $\mu_n \to \infty$, $\lambda_n = 1$ as long as k is finite (because $\lim_{k\to\infty}((\mu_i - k^2)/(\mu_i))^2 = \infty$). Furthermore, $\lim_{\mu_1\to 0}((\mu_1 - k^2)/(\mu_1))^2 = \infty$. Therefore, $\lambda_{\max}$ can become extremely large, which makes M_0 less favorable for preconditioning.

For the preconditioned system $(M_1^{-1}A)^*(M_1^{-1}A)$, one finds that

$$\begin{aligned}
\lambda\left((M_1^{-1}A)^*(M_1^{-1}A)\right)_{\min} &= \left(\frac{\mu_{m_3}-k^2}{\mu_{m_3}+k^2}\right)^2, \\
&\quad \text{where } |\frac{\mu_{m_3}-k^2}{\mu_{m_3}+k^2}| \le |\frac{\mu_i-k^2}{\mu_i+\mu_{m_3}}|,\ \forall i, \\
\lambda\left((M_1^{-1}A)^*(M_1^{-1}A)\right)_{\max} &= \max\left(\left(\frac{\mu_n-k^2}{\mu_n+k^2}\right)^2, \left(\frac{\mu_1-k^2}{\mu_1+k^2}\right)^2\right). \qquad (37)
\end{aligned}$$

From (37), it is found that

$$\lim_{\mu_n\to\infty}\left(\frac{\mu_n-k^2}{\mu_n+k^2}\right)^2 = \lim_{\mu_1\to 0}\left(\frac{\mu_1-k^2}{\mu_1+k^2}\right)^2 = \lim_{k\to\infty}\left(\frac{\mu_i-k^2}{\mu_i+k^2}\right)^2 = 1. \qquad (38)$$

For all possible extreme cases, the preconditioned system $M_1^{-1}A$ is always bounded above by one, i.e. the eigenvalues are always clustered. We can conclude that in the indefinite case M_1 may be better than M_0.

Finally, we are looking at the complex shifted preconditioned system with M_i. One finds that

$$\begin{aligned}
\lambda\left((M_i^{-1}A)^*(M_i^{-1}A)\right)_{\min} &= \frac{(\mu_{m_4}-k^2)^2}{\mu_{m_4}^2+k^4}, \\
&\quad \text{where } |\frac{(\mu_{m_4}-k^2)^2}{\mu_{m_4}^2+k^4}| \le |\frac{(\mu_i-k^2)^2}{\mu_i^2+k^4}|,\ \forall i, \\
\lambda\left((M_i^{-1}A)^*(M_i^{-1}A)\right)_{\max} &= \max\left(1-\frac{2\mu_1k^2}{\mu_1^2+k^4}, 1-\frac{2\mu_nk^2}{\mu_n^2+k^4}\right). \qquad (39)
\end{aligned}$$

The following results follow:

$$\lim_{\mu_n \to \infty} \lambda \left((M_i^{-1}A)^*(M_i^{-1}A) \right)_{\max} = \lim_{\mu_1 \to 0} \lambda \left((M_i^{-1}A)^*(M_i^{-1}A) \right)_{\max} = \\ = \lim_{k \to \infty} \lambda \left((M_i^{-1}A)^*(M_i^{-1}A) \right)_{\max} = 1. \tag{40}$$

Hence, the eigenvalues of $(M_i^{-1}A)^*(M_i^{-1}A)$ are always bounded above by one.

To determine the lower bound, we assume that $\lambda_{\min} \approx 0$ implying $\mu_m = k^2 + \epsilon,\ \epsilon > 0$. After substituting this relation to (39), one finds that

$$\lambda \left((M_i^{-1}A)^*(M_i^{-1}A) \right)_{\min} = \frac{1}{2}\frac{\epsilon^2}{k^4}. \tag{41}$$

Comparing to M_1 preconditioning (where $\lambda \left((M_1^{-1}A)^*(M_1^{-1}A) \right)_{\min} = \frac{1}{4}\epsilon^2/k^4$), it follows that $\lambda \left((M_i^{-1}A)^*(M_i^{-1}A) \right)_{\min} = 2\lambda \left((M_1^{-1}A)^*(M_1^{-1}A) \right)_{\min}$. With respect to the l_2-condition number, one finds that

$$\kappa \left((M_i^{-1}A)^*(M_i^{-1}A) \right) = 2 \left(\frac{k^4}{\epsilon^2} \right) \ < \ \kappa \left((M_1^{-1}A)^*(M_1^{-1}A) \right) = 4 \left(\frac{k^4}{\epsilon^2} \right).$$

We conclude that M_i may be a better preconditioner than M_1. M_i also may be better than M_0, which has possibly unbounded eigenvalues (and therefore a very large condition number) if μ is very small and $\mu \ll k$.

6 Numerical results

We provide some numerical results for solving equation (1), and present three cases as the model problems: (i) a 2-D closed-off problem with Dirichlet conditions at all boundaries, (ii) a 2-D open problem in a homogeneous medium with Sommerfeld conditions on a part of the boundary, and (iii) a 2-D open problem in an inhomogeneous medium.

For all cases, we solve the resulting linear system with full GMRES and compare three preconditioners M_0, M_1, and M_i. We set the maximum number of GMRES iterations to 150. Beyond this value, we restart GMRES. As k increases considerably, storing 150 vectors becomes too expensive, requiring a smaller restart parameter. This is the main drawback of using GMRES. Therefore, for the third problem the GMRES convergence is compared to that of CGNR and Bi-CGSTAB. The iteration is terminated at the k-th step if $\|\mathbf{r}_k\|_2/\|\mathbf{b}\|_2 < 10^{-6}$. The step involving M^{-1} is accomplished by using Gauss elimination. In practice this process is very costly. Of course, since M is symmetric and both the real and imaginary parts are positive definite, the LDL^T factorization can always be done (without requiring pivoting) and is unique [10]. We can also approximate M^{-1} using SSOR or multigrid. We do not implement these cheaper processes. The direct computation of M^{-1} can be used as a reference for approximations of M^{-1}.

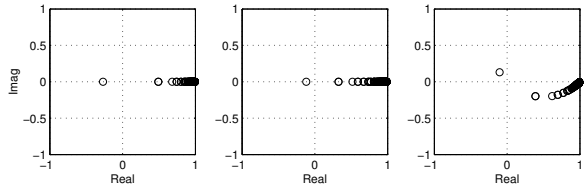

Figure 3: *Some extreme eigenvalues of the preconditioned systems of Problem 1 with* $k = 5$ *and gridsize* $h^{-1} = 20$

6.1 Closed-off problem

We consider a problem in a rectangular homogeneous medium governed by

$$\begin{aligned} \left(\Delta + k^2\right)\phi &= (k^2 - 5\pi^2)\sin(\pi x)\sin(2\pi y), \quad x = [0,1], y = [0,1], \\ \phi &= 0, \qquad\qquad\qquad\qquad\qquad\qquad \text{at the boundaries.} \end{aligned} \tag{42}$$

The exact solution of (42) is $\phi = \sin(\pi x)\sin(2\pi y)$. Different grid resolutions are used to solve the problem with various wavenumbers $k = 2, 5, 10, 15, 20$. $k = 2$ resembles the definite problem. In Figure 3, spectra of the preconditioned system for $k = 5$, a "slightly" indefinite problem, are shown. All spectra are bounded above by one.

Table 1 shows the computational performance in terms of number of iterations and number of arithmetic operations to reach the specified convergence. For low frequencies, all preconditioners show a very satisfactorily comparable performance. It appears that M_0 becomes less effective for increasing values of k, where the number of iterations increases somewhat faster than for M_1 or M_i. This behavior agrees with the theory.

In Table 2, the numerical performance is shown for the preconditioners for different grid resolutions. The preconditioners are sensitive with respect to the grid sizes. For all cases, M_i outperforms the other preconditioners.

6.2 2-D open homogeneous problem

The second problem represents an open problem allowing waves to penetrate the boundaries. We first look at a homogeneous medium in which waves created at the upper surface propagate. We consider

$$\begin{aligned} \Delta\phi + k^2\phi &= f, & \Omega &= [0,1]^2, \\ f &= \delta(x - 1/2)\delta(y), & x &= [0,1],\ y = 0, \\ \phi &= 0, & y &= 0, \\ \frac{\partial\phi}{\partial n} - \mathrm{i}k\phi &= 0, & x &= 0, 1,\ y = 1, \end{aligned} \tag{43}$$

with k constant in Ω. The performance of GMRES with M_0, M_1, and M_i as the preconditioners is compared.

Table 1: Computational performance of GMRES for 2-D closed-off problem. The preconditioner is the Shifted Laplace operator. 30 gridpoints per wavelength are used. The flops are measured in millions

	M_0		M_1		M_i	
k	Iter	flops	Iter	flops	Iter	flops
2	4	0.008	4	0.008	4	0.010
5	6	0.093	7	0.106	6	0.012
10	10	0.628	11	0.684	10	0.752
15	16	2.223	18	2.481	16	2.623
20	30	7.256	25	6.091	23	5.627
30	38	20.753	34	18.633	31	17.046
40	57	55.133	46	44.677	37	36.157

Table 2: Number of GMRES iterations to solve Problem 1 with various grid resolutions. The preconditioner is the Shifted Laplace operator

	M_0			M_1			M_i		
	h^{-1}			h^{-1}			h^{-1}		
k	50	100	150	50	100	150	50	100	150
5	5	5	5	6	6	6	6	6	5
10	10	10	10	11	11	11	10	10	10
15	17	14	14	19	15	15	17	14	14
20	33	30	21	37	25	21	31	23	19
30	79	55	38	86	57	34	72	52	31

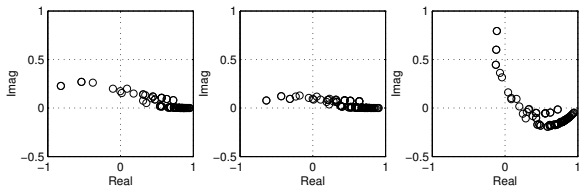

Figure 4: *Spectra of the linear systems from Problem 2 preconditioned with M_0 (left), M_1 (middle), and M_i (right). The radiation condition is replaced by a Dirichlet condition for $k = 5$ with $h^{-1} = 10$ (~ 10 gridpoints per wavelength)*

The Sommerfeld condition is set in the problem to avoid non-physical reflections. However, this may not be important for the preconditioning operator. We prefer to replace Sommerfeld's condition in the preconditioner by simpler a boundary condition. For this purpose, we use the Neumann or Dirichlet conditions in M. To see the effect of imposing different types of boundary conditions on the convergence performance, we analyze the spectra of the preconditioned linear systems.

Figure 4–6 show spectra of the preconditioned linear systems of Problem 2. The Jacobi-Davidson algorithm [18] is used to compute some extreme eigenvalues. Imposing either Dirichlet or Neumann conditions results in a good clustering of eigenvalues. For the example at hand, it shows that imposing Neumann conditions in the preconditioning operator gives better clustering and effectively pushes the negative real eigenvalues towards the imaginary axis (see the effect in Figure 4 and 5 for $k = 5$) compared to the Dirichlet conditions.

Figure 6 gives the spectra of the preconditioned linear system for $k = 5$, where the preconditioning matrix M is constructed using the same physical boundary conditions. With M_i, the eigenvalues are also pushed towards the positive real plane. But, some eigenvalues still remain in the negative real plane. Compared to Figure 3, imposing the same boundary conditions as the physics may lead to a better convergence rate than if the Dirichlet condition is used to replace the Sommerfeld condition. However, this may not be the case if the Sommerfeld condition is replaced by the Neumann condition. From the spectra, it shows that the Neumann condition may be the best option to replace the Sommerfeld condition for constructing M. Our numerical results (which are not shown in this paper) also confirm this conslusion. Therefore, for constructing the preconditioner we choose the Neumann condition to replace the Sommerfeld condition at the corresponding boundary.

Table 3 shows the number of GMRES iterations to solve Problem 2. For all frequencies, M_i outperforms M_0 and M_1. M_0 still performs reasonably well compared to M_i. This is not explained by the theory and may be due to the influence of different boundary conditions imposed in constructing the preconditioning matrix, which is not taken into account in our analysis.

Figure 7 shows the updated residual computed at each iteration as well as the error

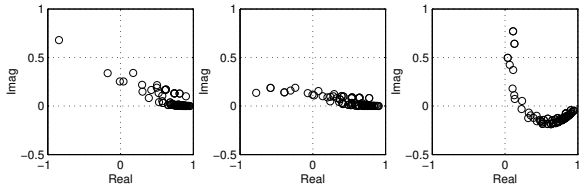

Figure 5: *Spectra of the linear systems from Problem 2 preconditioned with M_0 (left), M_1 (middle), and M_i (right). The radiation condition is replaced by a Neumann condition for $k = 5$ with $h^{-1} = 10$ (~ 10 gridpoints per wavelength)*

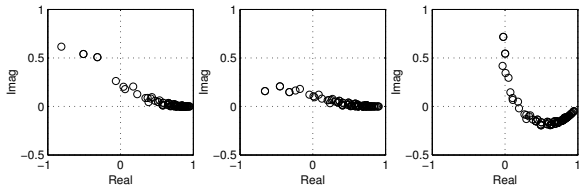

Figure 6: *Spectra of the linear systems from Problem 2 preconditioned with M_0 (left), M_1 (middle), and M_i (right). The Sommerfeld conditions with the physical problems are used to construct the preconditioner. $k = 5$, $h^{-1} = 10$ (~ 10 gridpoints per wavelength)*

for $k = 20$. The residual curve (in the left) indicates slow convergence for the first few iterations and a convergence improvement later on, indicating a superlinear convergence. (For the restarted version, choosing a low restart parameter may cause unacceptably slow convergence or even stagnation.) The error curve in the right figure indicates that a sufficiently small error (based on the l_2-norm) is reached, and that the error convergence closely follows the residual convergence.

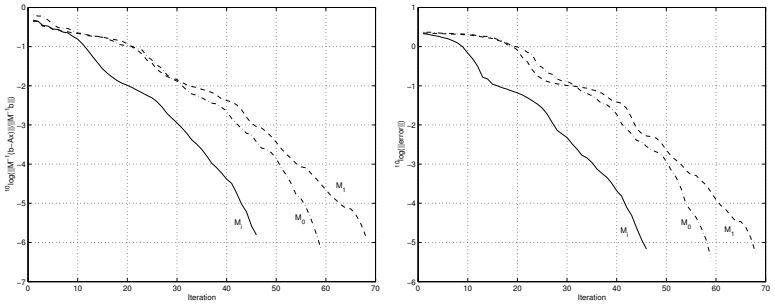

Figure 7: *Convergence history of preconditioned GMRES iterations, $k = 20$*

Table 3: Computational performance of GMRES to solve Problem 2. The preconditioner is the Shifted Laplace preconditioners. 30 gridpoints per wavelength are used

	M_0		M_1		M_i	
k	Iter	flops	Iter	flops	Iter	flops
2	8	0.019	8	0.019	6	0.014
5	12	0.171	14	0.197	11	0.158
10	24	1.476	26	1.596	19	1.179
15	38	5.098	43	5.770	30	4.040
20	59	14.454	68	16.715	46	11.255
30	115	63.984	131	73.479	80	43.997

6.3 2-D open inhomogeneous problem

In this example we repeat the computation of Problem 2 but now in an inhomogeneous medium. The wavenumber varies inside the domain according to

$$k = \begin{cases} k_{\text{ref}} & 0 \leq y \leq 1/3, \\ 1.5k_{\text{ref}} & 1/3 \leq y \leq 2/3, \\ 2.0k_{\text{ref}} & 2/3 \leq y \leq 1.0. \end{cases} \tag{44}$$

The number of gridpoints used is $5 \times k_{\text{ref}}$ (i.e., approximately 30 gridpoints per reference wavelength) in the x and y directions. As the preconditioners are sensitive with respect to the gridsize (refer to Table 3), less gridpoints in layers with $k > k_{\text{ref}}$ may affect the computational performance negatively. Numerical results are presented in Table 4. Here, we compute the solutions using full GMRES, and compare the computational performances with CGNR and Bi-CGSTAB. For the latter iterative methods, we limit the number of iterations to 1000.

In this harder problem, M_i again outperforms M_0 and M_1 indicated by the smaller number of iterations required to reach convergence. Compared to M_0, M_1 shows a less satisfactorily performance, and based on our computational restrictions restart is needed. For GMRES, restarting is needed for $k > 20$.

From Table 4, we also see that the preconditioned Bi-CGSTAB does not perform well for M_0 and M_1, as already indicated in [11]. However, the convergence with M_i as the preconditioner is still satisfactory. Compared to GMRES, Bi-CGSTAB preconditioned by M_i shows better convergence performance (despite of requiring two preconditioning steps within one iteration). If M_i is used as the preconditioner, Bi-CGSTAB can be the alternative to replace full GMRES.

Table 4: Computational performance of GMRES, CGNR, and Bi-CGSTAB to solve Problem 3. The preconditioner is the Shifted Laplace operator. 30 gridpoints per k_{ref} are used

	GMRES			CGNR			Bi-CGSTAB		
k_{ref}	M_0	M_1	M_i	M_0	M_1	M_i	M_0	M_1	M_i
2	6	10	8	12	12	10	6	7	5
5	17	20	14	39	31	23	17	15	10
10	39	47	31	189	88	66	150	56	22
15	73	85	54	647	175	126	685	113	40
20	120	>150	82	>1000	268	194	>1000	177	60
30	>150	>150	141	>1000	502	361	>1000	344	105

From Table 4, one also concludes that CGNR may not be a good iterative method to solve the Helmholtz problem with the Shifted Laplace preconditioners. Since our analysis is made for the normal equations, despite of worse performance compared to BiCGSTAB and GMRES, the results of CGNR confirm our analysis for the preconditioners M_0, M_1, and M_i.

7 Conclusion

In this paper, a class of preconditioners based on the Shifted Laplace operator for the Helmholtz equation has been presented and analyzed. We find that the complex Shifted-Laplace operator leads to the most effective preconditioning matrix within this class of preconditioners. Numerical experiments have been presented to show the effectiveness of the preconditioner. This preconditioner is easy to construct and to extend to inhomogeneous medium cases. Our numerical experiments show that for the latter, this preconditioner performs effectively. With respect to storage and CPU time requirements, we advocate the complex shifted preconditioner in combination with Bi-CGSTAB.

References

[1] A. Bamberger, P. Joly, J.E. Roberts, Second-order absorbing boundary conditions for the wave equation: a solution for corner problem, SIAM J. Numer. Anal. 27(2) (1990) 323–352.

[2] A. Bayliss, C.I. Goldstein, E. Turkel, An iterative method for Helmholtz equation, J. Comput. Phys. 49 (1983) 443–457.

[3] R.W. Clayton, B. Engquist, Absorbing boundary conditions for wave-equation migration, Geophysics 45(5) (1980) 895–904.

[4] B. Engquist, A. Majda, Absorbing boundary conditions for the numerical simulation of waves, Math. Comp. 31 (1977) 629–651.

[5] R. Fletcher, Conjugate gradient methods for indefinite systems, in: G.A. Watson, (Eds.), Proc. of the Dundee Biennal Conference on Numerical Analysis 1974, Springer Verlag, New York, 1975, pp. 73–89.

[6] R.W. Freund, Preconditioning of symmetric, but highly indefinite linear systems, Num. Anal. Manus. 97-3-03, Bell Labs., Murray Hill, NJ, 1997.

[7] R.W. Freund, N.M. Nachtigal, QMR: a quasi minimum residual method for non-Hermitian Linear Systems, Numer. Math. 60 (1991) 315–339.

[8] M.J. Gander, F. Nataf, AILU for Helmholtz problems: a new preconditioner based on the analytic parabolic factorization, J. Comp. Acoustics 9(4) (2001) 1499–1509.

[9] J. Gozani, A. Nachshon, E. Turkel, Conjugate gradient coupled with multigrid for an indefinite problem, in: Advances in Computer Methods for Partial Differential Equations V (1984) 425–427.

[10] N.J. Higham, Factorizing complex symmetric matrices with positive definite real and imaginary parts, Math. Comp. 67(3) (1998) 1591–1599.

[11] A.L. Laird, Preconditioned iterative solution of the 2D Helmholtz equation, First Year's Report St. Hugh's College (2001).

[12] M.M.M. Made, Incomplete factorization-based preconditionings for solving the Helmholtz equation, Int. J. Numer. Meth. Engng. 50 (2001) 1077–1101.

[13] C.C. Paige, M.A. Saunders, Solution of sparse indefinite systems of linear equations, SIAM J. Numer. Anal. 12 (4) (1975) 617–629.

[14] R.-E. Plessix, W.A. Mulder, Separation of variables as a preconditioner for an iterative Helmholtz solver, To appear in Appl. Num. Math. (2003).

[15] Y. Saad, A flexible inner-outer preconditioned GMRES algorithm, SIAM J. Sci. Stat. Comput. 14 (1993) 461–469.

[16] Y. Saad, Iterative methods for sparse linear systems, PWS Publishing Company Boston, 1996.

[17] Y. Saad, M.H. Schultz, GMRES: A generalized minimal residual algorithm for solving nonsymmetric linear systems, SIAM J. Sci. Stat. Comput. 7(12) (1986) 856–869.

[18] G.L.G. Sleijpen, J. Booten, D.R. Fokkema, H.A. van der Vorst, Jacobi-Davidson type methods for generalized eigenproblems and polynomial eigenproblems, BIT 36 (1996) 593–633.

[19] P. Sonneveld, CGS, a fast Lanczos-type solver for nonsymmetric linear systems, SIAM J. Sci. Stat. Comput. 10 (1) (1986) 36–52.

[20] 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. Stat. Comput. 13(2) (1992) 631–644.

[21] H.A. van der Vorst, J.B.M. Melissen, A Petrov-Galerkin type method for solving $Ax = b$, where A is symmetric complex, IEEE Trans. on Magnetics 26(2)(1990) 706–708.

[22] H.A. van der Vorst, C. Vuik, GMRESR: a family of nested GMRES methods, Num. Lin. Algebra Appl. 1(4) (1994) 369–386.