Available online at www.sciencedirect.com
SCIENCE DIRECT®
Journal of Computational Physics 203 (2005) 358–370
JOURNAL OF COMPUTATIONAL PHYSICS
www.elsevier.com/locate/jcp

A complex Jacobi iterative method for the indefinite Helmholtz equation

G. Ronald Hadley *,[1]

Sandia National Laboratories, P.O. Box 5800, Albuquerque, NM 87185-5800, USA

Received 16 July 2004; received in revised form 17 September 2004; accepted 18 September 2004
Available online 5 November 2004

Abstract

An iterative procedure is described for the solution of the indefinite Helmholtz equation that is a two-step generalization of classic Jacobi iteration using complex iteration parameters. The method converges for well-posed problems at a rate dependent only upon the grid size, wavelength and the effective absorption seen by the field. The use of a simple Jacobi preconditioner allows the solution of 3D problems of interest in waveguide optics in reasonable runtimes on a personal computer with memory usage that scales linearly with the number of grid points. Both the iterative method and the preconditioner are fully parallelizable.

MSC: 65F10; 65N12; 65N22; 78-08; 78A45
Keywords: Helmholtz; Iterative methods; Jacobi

1. Introduction

The indefinite Helmholtz equation occurs repeatedly in the description of virtually all problems involving wave motion modelled in the frequency domain; i.e. single frequency wave motion with time-independent amplitudes. Two prominent examples are the propagation of sound waves in water, and the propagation of light through a dielectric medium. In both these cases, problems of current interest commonly require three-dimensional (3D) modelling with complicated material geometries and/or problem

* Tel.: +1 505 844 4015; fax: +1 505 844 8985.
E-mail address: grhadle@sandia.gov.
[1] Sandia is a multiprogram laboratory operated by Sandia Corporation, a Lockheed Martin Company, for the United States Department of Energy's National Nuclear Security Administration under contract DE-AC04-94AL85000.

0021-9991/$ - see front matter
doi:10.1016/j.jcp.2004.09.015

boundaries, in other words, numerical models. But the grid size for such models is restricted to some fraction of the wavelength, and thus 3D models for practical problems usually result in matrix equations too large to be solvable by direct inversion. A similar situation in potential theory was dealt with half a century ago by the introduction of a number of attractive iterative methods for the solution of Poisson's equation including successive over-relaxation (SOR), alternating direction implicit (ADI) [1,2] and (more recently) multigrid algorithms [3], just to name a few. This success was in large part due to the definite nature of the Laplacian operator, i.e. it's eigenvalues are all of one sign. In contrast, it has been known for some time that the indefinite nature of the Helmholtz operator describing wave motion renders most of these classical methods non-convergent. The one exception is multigrid, which although convergent, is limited in usefulness because solutions on grids too coarse to adequately resolve the wavelength become poor approximations, and solutions on grids finer than about one-tenth of a wavelength unnecessary. Thus the dynamic range of available grids (and therefore the feature that gives multigrid methods their advantage) is limited to less than an order-of-magnitude.

More recent modelling activities directed towards addressing this problem have gravitated towards the collection of methods known as Krylov subspace methods [4]. These techniques seek to minimize some function of the residual vector in an optimal manner. Included in this category are well-known methods such as conjugate gradients (CG) and generalized minimum residuals (GMRES), among others. Although effective in some cases, they tend to be complicated and difficult to use. In many cases, the use of preconditioners is a requirement for convergence, so that the search of a suitable preconditioner adds substantially to the burden of the modeler.

This article describes the generalization of one of the classical iterative techniques, namely Jacobi iteration [5], to make it convergent for the indefinite Helmholtz equation. Jacobi's original method has enjoyed little use except as a smoother for multigrid schemes, even for Laplace's equation, due to its unacceptably slow convergence rate and poor scaling with problem size. For the Helmholtz equation it is nonconvergent in its original formulation. When used as a two-step method with complex iteration parameter, however, it will be shown below to converge at a rate dependent only upon the grid size, wavelength and effective absorption coefficient. The method is simple to understand and code, and is immediately suitable for parallel computing. In the sections to follow, we develop the basic theory of convergence for a simple 2D model problem, discuss the use of preconditioners, and apply the method to a 3D example problem describing reflection of a dielectric waveguide mode at an etched air gap.

2. Discussion of the method

We begin by considering the simpler case of the indefinite Helmholtz equation in two dimensions

$$\left[\nabla^2 + k^2\right] H(x,y) = 0, \quad \Re(k^2) > 0, \quad \Im(k^2) > 0, \tag{1}$$

where the restriction on the real part of k^2 is necessary for the description of wave motion; i.e. it is what is meant by the adjective "indefinite". The second restriction on the imaginary part is necessary for the problem to be well-posed in the following sense: waves generated in the problem interior or introduced at a Dirichlet boundary must be absorbed somewhere, either in an absorptive material contained within the problem region or at a transparent boundary. Otherwise, no steady state solution in terms of wave amplitudes is possible. It does not preclude the presence of gain or lossless materials in specific regions, only that loss must predominate in some overall sense. If the loss occurs at a boundary, then the restriction above is understood to be an "effective" material loss. The point is that the restrictions above do not represent any real limitations of the method, but are necessary to investigate convergence rates.

Again for the sake of simplicity and concreteness we shall examine the finite difference form of Eq. (1) using a uniform Cartesian grid:

$$\left[\delta_x^2 + \delta_y^2 + k^2\right] H = 0, \tag{2}$$

where the finite difference operators are the common centered forms

$$\delta_x^2 H \equiv \frac{H_{i+1,j} + H_{i-1,j} - 2H_{i,j}}{\Delta x^2}, \tag{3}$$

$$\delta_y^2 H \equiv \frac{H_{i,j+1} + H_{i,j-1} - 2H_{i,j}}{\Delta y^2}, \tag{4}$$

and the subscripts refer to the x,y grid positions as usual. The standard Jacobi iterative method is defined by

$$2\alpha\left(\frac{1}{\Delta x^2} + \frac{1}{\Delta y^2}\right)\left(H_{i,j}^{n+1} - H_{i,j}^{n}\right) = \left(\delta_x^2 + \delta_y^2 + k^2\right) H^n, \tag{5}$$

where the superscripts refer to iteration level and α is the iteration parameter. (In Eq. (5), a more standard form of the so-called point Jacobi method would include the term k^2 in the factor on the left-hand-side. This term has been omitted in order to keep the latter from becoming zero for coarse grids. However, including it would not change any of the basic conclusions below.) A common way of investigating the convergence properties of such a scheme is to look at the ratio of the amplitudes of a single Fourier mode

$$A\, \mathrm{e}^{\mathrm{i}\left[(i-1)k_x\Delta x + (j-1)k_y\Delta y\right]}, \tag{6}$$

between successive iterations:

$$\frac{A^{n+1}}{A^n} = 1 + \frac{\xi - 2\sin^2\delta}{\alpha}, \tag{7}$$

where we have defined the spatial frequency parameter

$$\sin^2\delta \equiv \frac{\sin^2\frac{k_x\Delta x}{2}}{1 + \frac{\Delta x^2}{\Delta y^2}} + \frac{\sin^2\frac{k_y\Delta y}{2}}{1 + \frac{\Delta y^2}{\Delta x^2}} \tag{8}$$

and

$$\xi \equiv \frac{k^2}{2\left[\frac{1}{\Delta x^2} + \frac{1}{\Delta y^2}\right]}. \tag{9}$$

It is apparent that there is no value of α that will keep the ratio in (7) within the unit circle for all spatial frequencies unless $\xi = 0$ (Laplace's equation). In that case

$$\frac{A^{n+1}}{A^n} = 1 - \frac{2\sin^2\delta}{\alpha}. \tag{10}$$

The convergence rate (10) is unacceptably slow for the lowest spatial frequencies, and consequently Jacobi iteration has never been seriously considered as an iterative method for Laplace's or Poisson's equation. However, setting $\alpha = 2$ leads to rapid damping of the highest spatial frequencies, and so Jacobi iteration has sometimes been employed as a smoother in multigrid schemes where the low spatial frequencies are converged on a coarser grid.

Returning to the Helmholtz equation, we see that classical Jacobi iteration is not convergent. This can be remedied, however, by considering pairs of Jacobi steps with different complex α's. Then

$$\frac{A^{n+2}}{A^n} = \left(1 + \frac{\xi - 2\sin^2\delta}{\alpha_1}\right)\left(1 + \frac{\xi - 2\sin^2\delta}{\alpha_2}\right). \tag{11}$$

To confront the problem of the indefiniteness of the Helmholtz eigenvalue $\xi - 2\sin^2\delta$, we take $\alpha_2 = -\alpha_1^*$. Then (dropping the subscripts on the α's) we have

$$\frac{A^{n+2}}{A^n} = 1 - \frac{2\mathrm{i}\alpha_\mathrm{i}\left(\xi - 2\sin^2\delta\right) + \left(\xi - 2\sin^2\delta\right)^2}{|\alpha|^2}, \tag{12}$$

where from now on we will use the notation $\alpha = \alpha_\mathrm{r} + \mathrm{i}\alpha_\mathrm{i}$. Eq. (12) shows clearly that the indefinite eigenvalue now either appears squared or multiplying a purely imaginary number. To examine the stability of the scheme further, we write $\xi = \xi_\mathrm{r} + \mathrm{i}\xi_\mathrm{i}$ and define $\psi \equiv \xi_\mathrm{r} - 2\sin^2\delta$, resulting in the expression

$$\frac{A^{n+2}}{A^n} = 1 - a - b\mathrm{i}, \tag{13}$$

where

$$a \equiv \frac{\psi^2 - \xi_\mathrm{i}^2 - 2\alpha_\mathrm{i}\xi_\mathrm{i}}{|\alpha|^2}, \tag{14}$$

$$b \equiv \frac{2\psi(\alpha_\mathrm{i} + \xi_\mathrm{i})}{|\alpha|^2}. \tag{15}$$

A quick examination of (13) shows the stability condition to be

$$b^2 < a(2 - a). \tag{16}$$

Inserting the definitions (14) and (15) into (16) and defining the normalized quantities

$$\alpha_\mathrm{i}' \equiv \frac{\alpha_\mathrm{i}}{|\alpha|}, \quad \psi' \equiv \frac{\psi}{|\alpha|}, \quad \xi_\mathrm{i}' \equiv \frac{\xi_\mathrm{i}}{|\alpha|}, \tag{17}$$

results in the inequality

$$4\psi'^2\left(\alpha_\mathrm{i}' + \xi_\mathrm{i}'\right)^2 < \left(\psi'^2 - \xi'^2 - 2\alpha_\mathrm{i}'\xi_\mathrm{i}'\right)\left(2 - \psi'^2 + \xi_\mathrm{i}'^2 + 2\alpha_\mathrm{i}'\xi_\mathrm{i}'\right). \tag{18}$$

In typical problems ξ_i' is usually small, so we proceed by neglecting the square of ξ_i' compared with other quantities. After some rearrangement of terms, this results in

$$\alpha_\mathrm{i}'^2 < \frac{1}{2} - \frac{\psi'^2}{4} - \alpha_\mathrm{i}'\xi_\mathrm{i}'\left(1 + \frac{1}{\psi'^2}\right). \tag{19}$$

Now we shall soon see that α_i' is always negative, and since ξ_i' has been restricted to be positive, it is sufficient to demand that

$$\alpha_\mathrm{i}'^2 < \frac{1}{2} - \frac{\psi'^2}{4}. \tag{20}$$

The most restrictive condition occurs for the maximum value of ψ'^2, which is approximately $4/|\alpha|^2$, corresponding to the highest spatial frequencies allowed by the grid. Thus the stability condition simplifies to

$$\alpha_\mathrm{i}^2 < \alpha_\mathrm{r}^2 - 2, \tag{21}$$

which immediately requires $\alpha_\mathrm{r}^2 > 2$. So, as is common with iterative methods, the stability condition results from the prevention of unstable growth of the highest spatial frequencies. Given this restriction, what

convergence rate can be expected for real problems? Dominant spatial frequencies in a converged problem are those in the vicinity of $\psi = 0$ since these are the frequencies that correspond to satisfaction of the original Helmholtz equation. For this value of ψ, the convergence rate from (12) is

$$\frac{A^{n+2}}{A^n} = 1 + \frac{\xi_i^2 + 2\alpha_i\xi_i}{|\alpha|^2} \tag{22}$$

$$\approx 1 + \frac{2\alpha_i\xi_i}{|\alpha|^2} \tag{23}$$

Thus we see that for convergence to occur, α_i must be negative. Next we must determine the complex number α that, subject to the stability constraint (21), provides the highest convergence rate. It is seen from (23) that $|\alpha_i|$ should be as large as possible, so we look along the stability boundary. It then follows from simple differentiation that the optimum parameter is $\alpha = (\sqrt{3}, -1)$ which corresponds to $\alpha_i/|\alpha|^2 = -1/4$ and a convergence rate of

$$\frac{A^{n+2}}{A^n} = 1 - \frac{\xi_i}{2} = 1 - \frac{\Im(k^2)\Delta x^2}{8} \tag{24}$$

for $\Delta x = \Delta y$.

Thus, we see from (24) that convergence is assured provided that the "effective" imaginary part of k^2 is positive; i.e. that there is net loss. We emphasize that this does not require that any particular material in the problem be lossy, and is a restriction essentially on the problem eigenvalues, which are influenced also by the boundary conditions. In addition, we see from (24) that there is a penalty for fine zoning, since the convergence rate decreases as the square of the grid size, in addition to the increased computational time per step due to the added grid points. In practice, however, the necessary grid resolution is approximately fixed by the wavelength, and further grid refinements are of limited benefit. In other words, $k\Delta x$ is roughly constant from problem to problem, with the total run time then scaling more-or-less linearly with problem size.

The specific form of Eq. (5) also points to two major advantages of this method: parallelization and low memory requirement. Since the updated field at a certain grid point depends only upon nearby fields at the previous iteration step (the method is a point relaxation scheme), it is manifestly parallelizable. This fact promises considerable time savings for large problems when compared with more complex schemes that do not parallelize. Also, memory requirements are low even for large problems, because only enough memory is needed to store field components and coefficients. This results in a memory requirement proportional to the number of grid points.

3. Model problem convergence

We purport to test the convergence rate predictions derived in the previous section by numerical solution of a simple model problem. The problem parameters are detailed in Fig. 1. The interior is a single lossy material, and all boundary conditions are Dirichlet in nature, with a nonzero Gaussian field at the bottom boundary to act as a source. The loss range has been set low enough to allow significant boundary reflections, but high enough to enable convergence in a reasonable number of iterations. The computed field amplitude profiles are shown in Fig. 2. Using the optimum value of α as computed above, the model problem was solved by complex Jacobi iteration for two grid sizes ($\Delta x = \Delta y$) as a function of material loss. In each case the iterations were terminated when a field-weighted residual of 10^{-6} was satisfied at each grid point. The resulting comparison between the predictions of Eq. (24) and the observed iteration count is shown in Fig. 3 (a pair of Jacobi steps is counted as a single iteration). As can be seen, the agreement

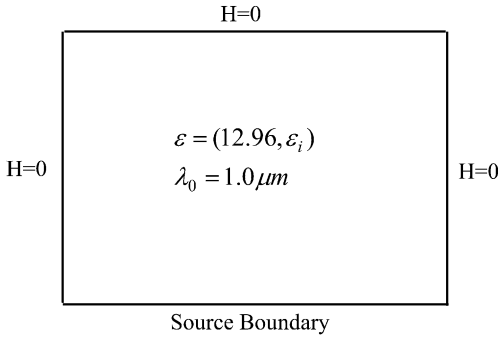

Fig. 1. Schematic of model problem.

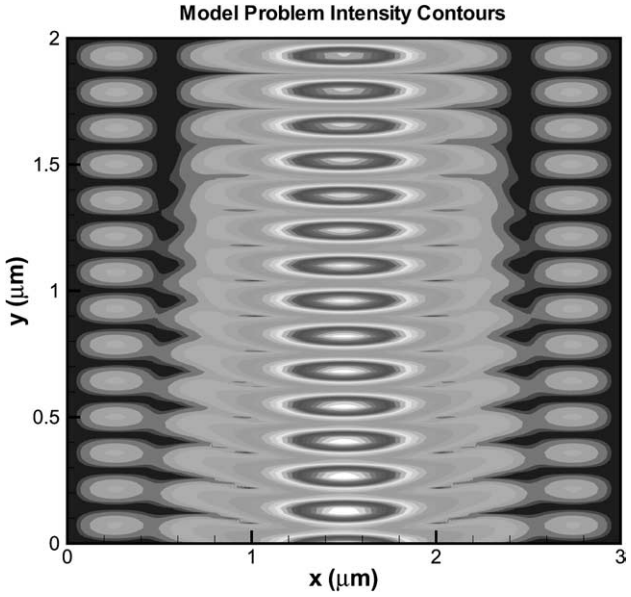

Fig. 2. Amplitude profiles for converged model problem. Input at the lower boundary is a Gaussian with field radius 0.5 μm.

between the stability analysis and the code performance is quite good, with the analysis accurately predicting the influence of both grid size and loss. Keep in mind that although the number of iterations is large, individual steps are simple and execute quickly. Overall run times for a realistic 3D problem are discussed in a later section.

4. Preconditioning

Like most finite difference methods, the range of possible eigenvalues of the differenced Helmholtz operator is much larger than the actual range encountered in a converged problem, because of the existence of non-physical high spatial frequencies that are resolved by the grid. As was seen in Section 2, the convergence parameter is restricted by requiring the stability of these frequencies even though they will not appear

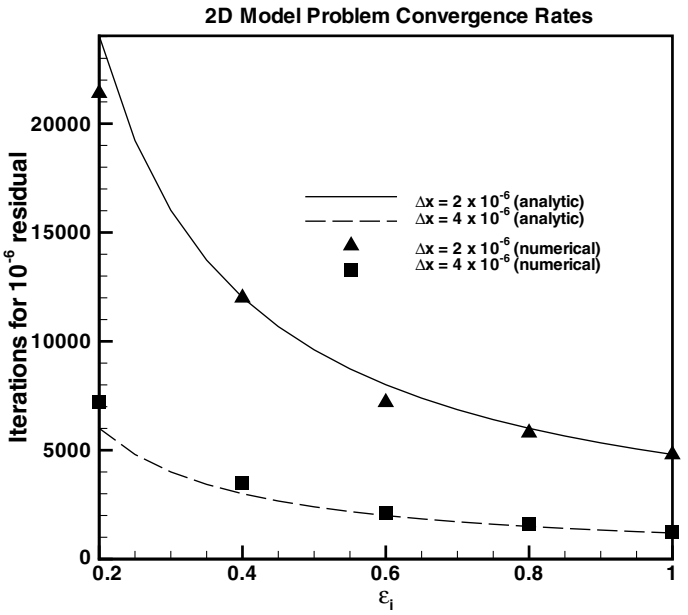

Fig. 3. Iterations required for convergence of the model problem vs. the imaginary part of the dielectric constant (no preconditioning). The curves reflect the analysis described in Section 2.

in the final solution. One way to view preconditioning, then, is that it preferentially damps the high spatial frequencies so as to allow a more aggressive convergence parameter, which in turn will lead to a faster problem convergence. Consistent with this description, we will utilize a classical Jacobi iteration step for the Laplace operator as our preconditioner, with a convergence parameter chosen to optimally damp the high spatial frequencies. From Eq. (5), with the convergence parameter denoted by β, if the single Jacobi step for the Helmholtz equation is written in the form

$$H^{n+1} = \left(1 + \frac{A}{4\alpha}\right) H^n, \tag{25}$$

where A is the normalized Helmholtz operator

$$A \equiv \frac{L + k^2}{\frac{1}{2}\left(\frac{1}{\Delta x^2} + \frac{1}{\Delta y^2}\right)}, \tag{26}$$

$$L \equiv \delta_x^2 + \delta_y^2, \tag{27}$$

then we can define a preconditioned complex Jacobi two-step algorithm by

$$H^{n+2} = \left[1 - \frac{1}{4\alpha^*}\left(1 + \frac{L}{4\beta}\right)A\right]\left[1 + \frac{1}{4\alpha}\left(1 + \frac{L}{4\beta}\right)A\right] H^n. \tag{28}$$

Performing the same stability analysis on (28) that was previously done for the non-preconditioned algorithm results in

$$\frac{A^{n+2}}{A^n} = 1 - \frac{\widetilde{\psi}^2 - \widetilde{\xi}_i^2 - 2\alpha_i\widetilde{\xi}_i + 2i\widetilde{\psi}\left(\alpha_i + \widetilde{\xi}_i\right)}{|\alpha|^2}, \tag{29}$$

where

$$\widetilde{\psi} \equiv \left(1 - \frac{2\sin^2\delta}{\beta}\right)\psi, \tag{30}$$

$$\widetilde{\xi}_{\mathrm{i}} \equiv \left(1 - \frac{2\sin^2\delta}{\beta}\right)\xi_{\mathrm{i}}. \tag{31}$$

Since the form of (29) is identical to the non-preconditioned result (13), we expect the solution for an optimum value of α to be similar. However, it is not identical because the maximum value of $\widetilde{\psi}^2$ encountered in the stability analysis now depends on β. Thus we must determine

$$f(\beta) \equiv \max_{\delta\in[0,\pi/2]} \left\{ \frac{1}{2}\left(1 - \frac{2\sin^2\delta}{\beta}\right)^2 \left(\xi_{\mathrm{r}} - 2\sin^2\delta\right)^2 \right\}. \tag{32}$$

Since ξ_{r}^2 is expected to be small, the expression in (32) reduces to two possible candidates: (1) a local maximum around $\sin^2\delta = \beta/4$ or (2) the right endpoint $\sin^2\delta = 1$. Thus,

$$f(\beta) = \frac{1}{2}\max\left\{ \frac{1}{16}\left(1 - \frac{\xi_{\mathrm{r}}}{\beta}\right)^2 (\xi_{\mathrm{r}} - \beta)^2, \left(1 - \frac{2}{\beta}\right)^2 (\xi_{\mathrm{r}} - 2)^2 \right\}. \tag{33}$$

Since the first function is monotone increasing over the range of interesting values of β and the second decreasing, this becomes (again neglecting ξ_{r})

$$f(\beta) = \begin{cases} 2\left(1 - \frac{2}{\beta}\right)^2, & \beta < 4\left(\sqrt{2} - 1\right), \\ \beta^2/32, & \beta > 4\left(\sqrt{2} - 1\right), \end{cases} \tag{34}$$

and the analogue of Eq. (21) for the stability condition is

$$\alpha_{\mathrm{i}}^2 < \alpha_{\mathrm{r}}^2 - f(\beta). \tag{35}$$

Now the convergence rate for spatial frequencies of interest is

$$\frac{A^{n+2}}{A^n} = 1 + \frac{2\alpha_{\mathrm{i}}\xi_{\mathrm{i}}(1 - 4\xi_{\mathrm{r}}\beta)}{|\alpha|^2} \tag{36}$$

$$\approx 1 + \frac{2\alpha_{\mathrm{i}}\xi_{\mathrm{i}}}{|\alpha|^2} \tag{37}$$

So, as before, the optimum corresponds to a maximization of $\alpha_{\mathrm{i}}/|\alpha|^2$ subject to the stability constraint (35). This leads to the optimum

$$\alpha_{\mathrm{opt}} = \left(\sqrt{\frac{3f}{2}}, -\sqrt{\frac{f}{2}}\right), \tag{38}$$

corresponding to

$$\frac{\alpha_{\mathrm{i}}}{|\alpha|^2} = \frac{1}{\sqrt{8f(\beta)}}. \tag{39}$$

Thus for the fastest convergence, we should choose β to minimize f. This occurs at the intersection of the two functions in (33) at the value $\beta = 4(\sqrt{2} - 1) \approx 1.657$. For this value, $f = 0.086$, $\alpha_{\mathrm{opt}} = (0.36, -0.21)$, and $\alpha_{\mathrm{i}}/|\alpha|^2 = 1.207$.

We see that the effect of the preconditioner has been to relax the high-spatial-frequency stability constraint and allow much more aggressive values of $\alpha_i/|\alpha|^2$ (by more than a factor of 5), resulting in much faster convergence. We note that the derivations in this section have been somewhat approximate in order to keep the algebraic expressions simple and illustrate trends. Exact analyses are of questionable value anyway, since convergence parameters will ultimately be chosen based on experience. The purpose of these formulas is to act as a guide to convergence behavior and parameter optimization.

The preconditioner just described has been tested on the model problem discussed in the previous section as a check on the convergence rates predicted using the simple analysis. The results, using the optimized values derived above, are plotted in Fig. 4. We note first that the iteration count is down by more than a factor of 4 from the non-preconditioned tests. Since somewhat less than twice the numerical effort has been expended, this represents a real computational savings. Also, we see that, as before, there is good agreement between the predicted and observed convergence rates, providing confidence in our analysis and understanding of the method.

At this point, little effort has been expended investigating the effects of other preconditioners, of which there are many. So it is entirely possible that a different preconditioner than the one described here might give considerably improved performance. It must be kept in mind, however, that in order to match up well with the complex Jacobi iterative method, any preconditioner should be fast and simple to code, and should be suitable for parallel computation, as is the Jacobi preconditioner. Otherwise, much of the advantage of this method would be lost.

It should also be mentioned that, although all the analyses reported here employ a regular cartesian grid, the method has been tested successfully with other less regular grids. In fact, the model problem described here has been solved on an irregular triangular grid [6]. Using the same parameter values as reported above, the observed convergence rates followed the same trends as shown in Fig. 3, but were somewhat higher.

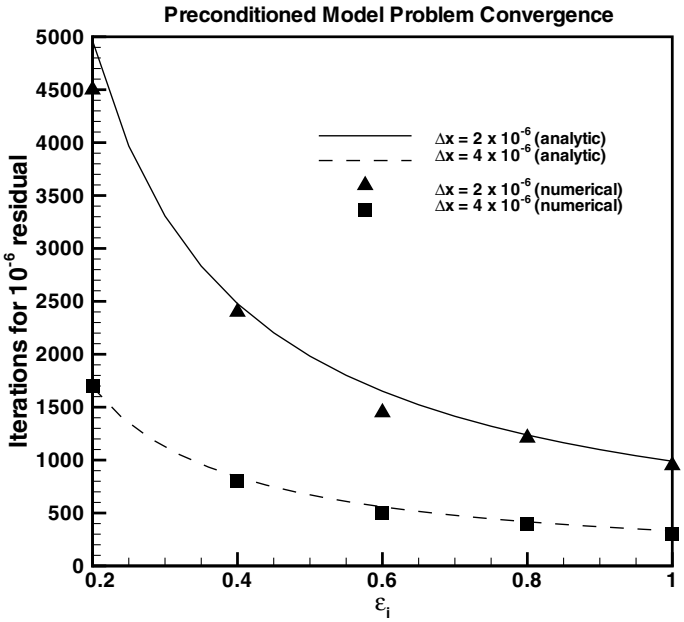

Fig. 4. Iterations required for convergence of the model problem with preconditioning vs. the imaginary part of the dielectric constant. The curves reflect the analysis described in Section 4.

5. Extension to three dimensions

Although the derivations and examples have so far been restricted to two dimensions for reasons of clarity and simplicity, it is clear that the real gains of this approach lie in its ability to provide solutions in three dimensions where the memory requirements of direct matrix inversion are often prohibitive. As might be expected, the general behavior of the method in 3D is essentially the same as 2D with some slight modifications. Classical Jacobi iteration for the Helmholtz equation in 3D may be written

$$2\alpha\left(\frac{1}{\Delta x^2}+\frac{1}{\Delta y^2}+\frac{1}{\Delta z^2}\right)\left(H_{i,j,k}^{n+1}-H_{i,j,k}^{n}\right)=\left(L+k^2\right)H^n, \tag{40}$$

where L is now the three-dimensional Laplace operator $\delta_x^2+\delta_y^2+\delta_z^2$. The two-step complex Jacobi iterative method then becomes

$$H^{n+2}=\left[1-\frac{A}{4\alpha^*}\right]\left[1+\frac{A}{4\alpha}\right]H^n, \tag{41}$$

where now

$$A\equiv\frac{L+k^2}{\frac{1}{2}\left(\frac{1}{\Delta x^2}+\frac{1}{\Delta y^2}+\frac{1}{\Delta z^2}\right)}. \tag{42}$$

Performing the same stability analysis as before on (41) yields a similar result:

$$\frac{A^{n+2}}{A^n}\approx 1+\frac{2\alpha_i\xi_i}{|\alpha|^2}, \tag{43}$$

except that now

$$\xi\equiv\frac{k^2}{2\left(\frac{1}{\Delta x^2}+\frac{1}{\Delta y^2}+\frac{1}{\Delta z^2}\right)}. \tag{44}$$

So we see that a 3D problem with equal grid spacings in all directions will converge at a rate that is 33% slower than an identical 2D problem due to the difference in the denominators between (44) and (9). Otherwise, all the formulas developed to describe convergence behavior for the 2D problem apply immediately to the 3D case, including the effects of preconditioning.

6. Sample 3D calculation

In the field of diffractive optics, realistic 3D problems involving single-frequency light propagating through regions of space containing piecewise constant media usually require the solution of the Helmholtz equation for all three components of a vector field, either the magnetic field **H** or the electric field **E**. If these components are cartesian, the component equations uncouple except at boundaries between dielectrics where the requirements of the continuity of tangential **E**, normal $\mathbf{D}=\epsilon\mathbf{E}$, and all components of **H** lead to coupling terms. Thus for an **H**-field formulation we require in each region of uniform dielectric constant ϵ that **H** satisfy the Helmholtz equation [7]

$$\left(\frac{1}{\epsilon}\nabla^2+k_0^2\right)\mathbf{H}=\mathbf{0}, \tag{45}$$

where $k_0 = 2\pi/\lambda_0$ is the vacuum wavevector. The derivation of equivalent finite difference equations proceeds by integrating Eq. (45) over each of the eight rectangular volumes in each quadrant enclosed by planes parallel to the cartesian axes and extending from the origin at point (i,j,k) to half a grid spacing in each direction. For generality, the material in each volume is assumed to be described by a unique dielectric constant. The use of Green's theorem then results in a number of area integrals of normal derivatives of various components of **H**. When all the integrals over the eight volumes are added, the normal derivatives over areas distant from the origin are approximated in the usual way in terms of fields at the surrounding grid points, while those over common planes passing through the origin either cancel (in the case of the normal field component) or are manipulated using the continuity of tangential **E** into tangential derivatives that are easily approximated. This procedure thus leads to the following set of coupled Helmholtz equations:

$$\begin{aligned} &A_x H_x + B_x H_y + C_x H_z = 0, \\ &A_y H_y + B_y H_z + C_y H_x = 0, \\ &A_z H_z + B_z H_x + C_z H_y = 0. \end{aligned} \tag{46}$$

In Eqs. (46) the A's B's and C's are finite difference operators that depend on k_0, the grid spacings and dielectric constants. In a region where all eight dielectric constants are equal, the B's and C's are zero, and the A's reduce to a Helmholtz operator proportional to that defined by Eq. (42). If Eq. (46) is rewritten in matrix form as

$$\mathscr{A}\mathbf{H} = \mathbf{0}, \tag{47}$$

with $\mathscr{A}$ scaled by the grid spacings in a manner similar to that shown in (42), then the complex Jacobi iteration procedure generalizes to

$$\mathbf{H}^{n+2} = \left[\mathbf{1} - \frac{\mathscr{A}}{4\alpha^*}\right]\left[\mathbf{1} + \frac{\mathscr{A}}{4\alpha}\right]\mathbf{H}^n. \tag{48}$$

For concreteness, we have chosen to model the reflection of the fundamental TE guided mode of a buried-oxide waveguide at an etched air gap as a function of the gap width. Calculations such as this are important in the field of integrated optics since they provide an estimate of the facet reflectivity for etched-facet semiconductor lasers. The structure is shown in Fig. 5. The mode of interest is computed using a separate eigenmode solver and used as a Dirichlet boundary condition to serve as a source. Absorbing materials are used adjacent to the input boundary and at other lateral boundaries to absorb power that is radiated laterally or reflected back towards the source boundary. At the exit boundary a simple non-adaptive transparent boundary condition is employed. The reflection coefficient is computed from various overlap integrals as reported previously [8]. Transmission is computed by comparing the integrated power flux through the exit boundary with that obtained with zero gap. The resulting reflection coefficient and normalized transmitted power are plotted in Fig. 6 versus gap width.

The number of iterations required for satisfactory convergence of this problem using a Jacobi preconditioner was found to vary with gap width, with narrow gap problems requiring 2500 iterations, and wider gap problems about 4000. However, runtimes in both cases were acceptable, and an average data point was computed in about 20 min on a high-end workstation (1.45 GHz IBM P650 with 8 processors) and an hour on a laptop pc. Thus the iterative method described here is capable of solving non-ideal problems of interest involving coupled equations over a large problem region (280,000 grid points) with modest memory usage (175 MB) and runtimes. The workstation runs employed the aid of some parallel processing, but little effort was expended to optimize the parallelization, as is evident from the modest 3 to 1 speedup over the laptop times, which employed no parallelization at all.

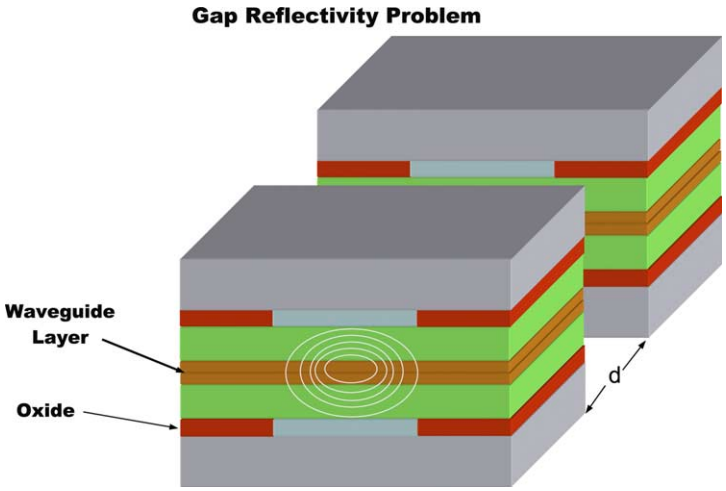

Fig. 5. Geometry for the Sample 3D problem. Problem input is handled as a Dirichlet boundary condition whose profile is the fundamental mode of the input waveguide.

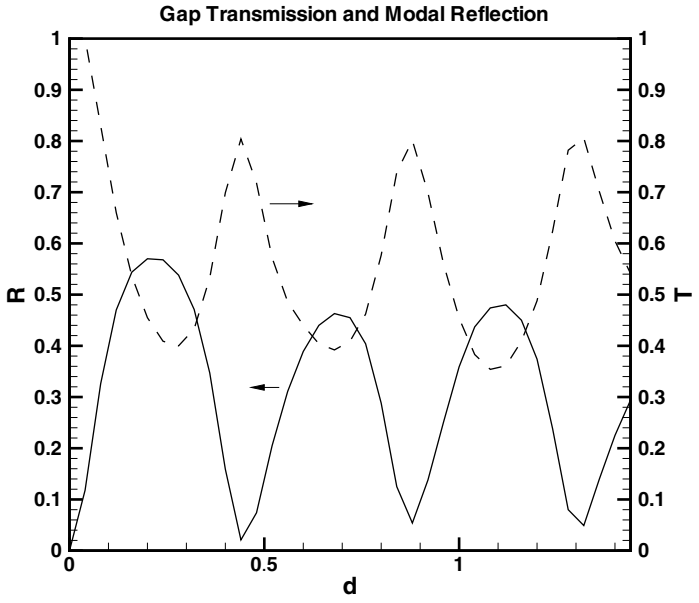

Fig. 6. Computed reflection and transmission coefficients vs air gap width for the structure shown in Fig. 5.

7. Conclusion

We have presented a two-step iterative method for the indefinite Helmholtz equation that converges predictably provided the problem is well-posed; i.e. that adequate absorption is present to provide a sink for whatever sources are introduced so that steady-state solutions exist. A simple convergence analysis has been shown to accurately predict convergence rates for a model problem made up of a uniform lossy dielectric,

as well as the convergence rate enhancement resulting from the application of a Jacobi preconditioner. These formulas not only serve as a guide for selecting optimum convergence and preconditioning parameters, but also provide insight and confidence regarding the overall behavior of the method. Finally, the model has been applied to the solution of a nontrivial problem of importance in diffractive optics involving the resonant reflection of a waveguide mode from an etched air gap. This problem was solved in three dimensions including all three components of the magnetic field vector in about an hour per gap width value on a laptop personal computer. This accomplishment aptly demonstrates the modest runtimes and memory requirements necessary to affect the solution of a nontrivial problem of significant value. Even more efficient execution is expected from parallel processors, since both the iterative method and the Jacobi preconditioner described above are ideally suited for parallel processing.

References

[1] D.M. Young, Iterative Solution of Large Linear Systems, Academic Press, New York, 1971 (Chapter 3).
[2] R.S. Varga, Matrix Iterative Analysis, Prentice-Hall Inc, Englewood Cliffs, NJ, 1962 (Chapter 3).
[3] A. Brandt, Multi-level adaptive solutions to boundary-value problems, Math. Computat. 31 (1971) 333.
[4] C.T. Kelley, Iterative Methods for Linear and Nonlinear Equations, Society for Industrial and Applied Mathematics, Philadelphia, 1995, p. 11.
[5] Ref. [1], p. 71.
[6] G.R. Hadley, Numerical simulation of waveguides of arbitrary cross-section, Int. J. Electron. Commun. 58 (2004) 86.
[7] J.D. Jackson, Classical Electrodynamics, John Wiley and Sons, New York, 1962, p. 241.
[8] G.R. Hadley, Low-truncation-error finite difference representations of the 2-D Helmholtz equation, Int. J. Electron. Commun. 52 (1998) 310.