

A MULTIGRID METHOD ENHANCED BY KRYLOV SUBSPACE ITERATION FOR DISCRETE HELMHOLTZ EQUATIONS*

HOWARD C. ELMAN[†], OLIVER G. ERNST[‡], AND DIANNE P. O'LEARY[§]

Abstract. Standard multigrid algorithms have proven ineffective for the solution of discretizations of Helmholtz equations. In this work we modify the standard algorithm by adding GMRES iterations at coarse levels and as an outer iteration. We demonstrate the algorithm's effectiveness through theoretical analysis of a model problem and experimental results. In particular, we show that the combined use of GMRES as a smoother and outer iteration produces an algorithm whose performance depends relatively mildly on wave number and is robust for normalized wave numbers as large as 200. For fixed wave numbers, it displays grid-independent convergence rates and has costs proportional to the number of unknowns.

Key words. Helmholtz equation, multigrid, Krylov subspace methods

AMS subject classifications. Primary, 65N55, 65F10, 65N22; Secondary, 78A45, 76Q05

PII. S1064827501357190

1. Introduction. Multigrid algorithms are effective for the numerical solution of many partial differential equations, providing a solution in time proportional to the number of unknowns. For some important classes of problems, however, standard multigrid algorithms have not been useful, and in this paper we focus on developing effective multigrid algorithms for one such class, the discrete Helmholtz equation.

Our main interest lies in solving exterior boundary value problems of the form

$$(1.1) \quad -\Delta u - k^2 u = f \quad \text{on } \Omega \subset \mathbb{R}^d,$$

$$(1.2) \quad Bu = g \quad \text{on } \Gamma \subset \partial\Omega,$$

$$(1.3) \quad \frac{\partial u}{\partial n} = Mu \quad \text{on } \Gamma_\infty \subset \partial\Omega,$$

that arise in the modeling of time-harmonic acoustic or plane-polarized electromagnetic scattering by an obstacle. The boundary Γ represents the scattering obstacle, and the boundary operator B can be chosen so that a Dirichlet, Neumann, or Robin boundary condition is imposed. The original unbounded domain is truncated to the finite domain Ω by introducing the artificial boundary Γ_∞ on which the radiation boundary condition (1.3) approximates the outgoing Sommerfeld radiation condition. Depending on what type of radiation condition is chosen, M can be either a (local) differential operator or a global integral operator coupling all points on Γ_∞ (see [16]). The data for the problem are given by the right-hand side f and the boundary data

*Received by the editors June 16, 1999; accepted for publication (in revised form) June 17, 2001; published electronically December 4, 2001.

<http://www.siam.org/journals/sisc/23-4/35719.html>

[†]Department of Computer Science and Institute for Advanced Computer Studies, University of Maryland, College Park, MD 20742 (elman@cs.umd.edu). This author's work was supported in part by the National Science Foundation under grant DMS9423133.

[‡]Department of Mathematics and Computer Science, TU Bergakademie Freiberg, 09596 Freiberg, Germany (ernst@math.tu-freiberg.de). This author's work was supported in part by the National Science Foundation under grant ASC9704683 and by the University of Maryland Institute for Advanced Computer Studies.

[§]Department of Computer Science and Institute for Advanced Computer Studies, University of Maryland, College Park, MD 20742 (oleary@cs.umd.edu). This author's work was supported in part by National Science Foundation grants NSF CCR-95-03126 and NSF CCR-97-32022.

g. In the most common case, $f \equiv 0$ and $-g$ is the boundary data of an incident plane wave. The critical parameter is the wave number k , which is positive in the case of unattenuated wave propagation. Due to the radiation boundary condition, the solution of (1.1)–(1.3) is a complex-valued function $u : \Omega \rightarrow \mathbb{C}$.

Discretization of (1.1)–(1.3) by finite differences or finite elements leads to a linear system of equations

$$(1.4) \quad \mathbf{A}u = \mathbf{f}$$

in which the coefficient matrix \mathbf{A} is complex-symmetric, i.e., not Hermitian. Moreover, for large values of the wave number k , it becomes highly indefinite.

This indefiniteness has prevented multigrid methods from being applied to the discrete Helmholtz equation with the same success as these methods have enjoyed for symmetric positive-definite problems. Some proposed multilevel strategies for the Helmholtz equation impose restrictions on the coarse grid, requiring that these grids be sufficiently fine for the algorithm to be convergent [2, 6, 10, 31, 35]; such restrictions limit the utility of these techniques. It is also possible to precondition the indefinite Helmholtz problem with preconditioners for the leading (second order) term [4, 5, 17, 36], although the effectiveness of this approach is limited for problems with large wave numbers. Our aim in this work is to identify the difficulties arising in a standard multigrid iteration for the Helmholtz equation and to analyze and test techniques designed to address these difficulties. In particular, in section 2, we show there are difficulties with both of the main multigrid components, smoothing and coarse grid correction. Standard smoothers such as Jacobi or Gauß–Seidel relaxation become unstable for indefinite problems since there are always error components—usually the smooth ones—that are amplified by these smoothers. The difficulties with the coarse grid correction are usually attributed to the poor approximation of the Helmholtz operator on very coarse meshes, since such meshes cannot adequately resolve waves with wavelength $\lambda = 2\pi/k$ of which the solution primarily consists. We show, however, that although the coarse grid correction is inaccurate when coarse grid eigenvalues do not agree well with their fine-grid counterparts, coarse meshes can still yield useful information in a multigrid cycle.

Our approach for smoothing is to use standard damped Jacobi relaxation when it works reasonably well (on fine enough grids) and then to replace this with a Krylov subspace iteration when it fails as a smoother. Earlier works by Bank [2] and Brandt and Ta'asan [12] have employed relaxation on the normal equations in this context. Krylov subspace smoothing, principally using the conjugate gradient method, has been considered by a variety of authors [3, 8, 9, 30, 32].

For coarse grid correction, we identify the type and number of eigenvalues that are handled poorly during the correction, and we remedy the difficulty by introducing an acceleration for multigrid; that is, we use multigrid as a preconditioner for an outer Krylov subspace iteration. This approach has been used by many authors; see, e.g., [7, 22, 25, 26, 31, 35]. In some settings, where it is used for positive-definite systems [7, 22], the Krylov subspace method accelerates the multigrid computations, but multigrid alone is also rapidly convergent and improvements in performance are not dramatic. A significant point in the present study is that multigrid does a poor job of eliminating some modes from the error. As a result, it converges slowly and even diverges in some cases, and the outer Krylov subspace iteration is needed for the method to be robust. (This phenomenon is also observed for the convection-diffusion equation in [25, 26].) Any Krylov subspace method is an option for both the smoother and the outer iteration; we use GMRES [28].

A different approach for adapting multigrid to the Helmholtz equation, based on representing oscillatory error components on coarse grids as the product of an oscillatory Fourier mode and a smooth amplitude—or ray—function, has been developed by Brandt and Livshits [11]. The standard V-cycle is augmented by so-called ray cycles, in which the oscillatory error components are eliminated by approximating the associated ray functions in a multigrid fashion. This wave-ray methodology has also been combined by Lee et al. [23] with a first-order system least-squares formulation for the Helmholtz equation. Results in [11, 23] suggest that these techniques may be somewhat more efficient than the methods of this paper, but they are considerably more difficult to implement, and it is unclear whether they can be generalized to handle variable coefficient problems.

An outline of the paper is as follows. In section 2, we perform a model problem analysis, using a one-dimensional problem to identify the difficulties encountered by both smoothers and coarse grid correction, and supplementing these observations with an analysis of how dimensionality of the problem affects the computations. In section 3, we present the refined multigrid algorithms and test their performance on a set of two-dimensional benchmark problems on a square domain. In particular, we demonstrate the effectiveness of an automated stopping criterion for use with GMRES smoothing, and we show that the combined use of GMRES as a smoother and outer iteration produces an algorithm whose performance depends relatively mildly on wave number and is robust for wave numbers as large as 200. In section 4, we show the performance of the multigrid solver on an exterior scattering problem. Finally, in section 5, we draw some conclusions.

2. Model problem analysis. Most of the deficiencies of standard multigrid methods for solving Helmholtz problems can be seen from a one-dimensional model problem. Therefore, we consider the Helmholtz equation on the unit interval $(0, 1)$ with homogeneous Dirichlet boundary conditions

$$(2.1) \quad -u'' - k^2 u = f, \quad u(0) = u(1) = 0.$$

This problem is guaranteed to be nonsingular only if k^2 is not an eigenvalue of the negative Laplacian, and we will assume here that this requirement holds. The problem is indefinite for $k^2 > \pi^2$, which is the smallest eigenvalue of the negative Laplacian.

Finite difference discretization of (2.1) on a uniform grid containing N interior points leads to a linear system of equations (1.4) with the $N \times N$ coefficient matrix $\mathbf{A} = \mathbf{A}^h = (1/h^2) \text{tridiag}(-1, 2, -1) - k^2 \mathbf{I}$, where $h = 1/(N + 1)$ denotes the mesh width and \mathbf{I} denotes the identity matrix. Under the assumptions on k above, it is well known (see [29]) that for sufficiently fine discretizations, the discrete problems are also nonsingular. We also assume that all coarse grid problems are nonsingular.

The eigenvalues of \mathbf{A} are

$$(2.2) \quad \lambda_j = \frac{2(1 - \cos j\pi h)}{h^2} - k^2 = \frac{4}{h^2} \sin^2 \frac{j\pi h}{2} - k^2, \quad j = 1, \dots, N,$$

and the eigenvectors are

$$(2.3) \quad \mathbf{v}_j = \sqrt{2h} [\sin ij\pi h]_{i=1}^N, \quad j = 1, \dots, N.$$

The choice of Dirichlet boundary conditions in (2.1) allows us to perform Fourier analysis using these analytic expressions for eigenvalues and eigenvectors. In experiments described in section 3, we will examine how our observations coincide with

performance on problems with radiation conditions, which are nonsingular for all k [20]. Aspects of the algorithm that depend on the dimensionality of the problem will be considered at the end of this section.

2.1. Smoothing. For the smoothing operator, we consider damped Jacobi relaxation, defined by the stationary iteration

$$\mathbf{u}_{m+1} = \mathbf{u}_m + \omega \mathbf{D}^{-1} \mathbf{r}_m = \mathbf{u}_m + \omega \mathbf{D}^{-1} \mathbf{A} \mathbf{e}_m,$$

where $\mathbf{r}_m = \mathbf{f} - \mathbf{A} \mathbf{u}_m$ and $\mathbf{e}_m = \mathbf{A}^{-1} \mathbf{f} - \mathbf{u}_m$ denote the residual and error vectors at step m , respectively. $\mathbf{D} = (2/h^2 - k^2) \mathbf{I}$ denotes the matrix consisting of the diagonal of \mathbf{A} , and ω is the damping parameter. The associated error propagation matrix is $\mathbf{S}_\omega = \mathbf{I} - \omega \mathbf{D}^{-1} \mathbf{A}$, and the eigenstructure of this matrix governs the behavior of the error $\mathbf{e}_{m+1} = \mathbf{S}_\omega \mathbf{e}_m$. Since \mathbf{D} is a multiple of the identity matrix, \mathbf{S}_ω is a polynomial in \mathbf{A} and hence shares the same system of orthonormal eigenvectors (2.3). The eigenvalues of \mathbf{S}_ω are

$$(2.4) \quad \mu_j = 1 - \omega \left(1 - \frac{\cos j\pi h}{1 - \frac{1}{2}k^2 h^2} \right), \quad j = 1, \dots, N.$$

Thus, the eigenvalue μ_j of \mathbf{S}_ω is the damping factor for the error component corresponding to the eigenvalue λ_j of \mathbf{A} .

We now consider the effects of damped Jacobi smoothing on three levels of grids: fine, coarse, and intermediate.

2.1.1. Fine grids. The fine-grid mesh size is determined by accuracy requirements on the discretization, and this allows us to make certain assumptions on the size of h versus k on the fine grid. Recall that the wavelength λ associated with a time-harmonic wave with wave number $k > 0$ is given by $\lambda = 2\pi/k$. The quantity

$$\frac{\lambda}{h} = \frac{2\pi}{kh} = \frac{2\pi}{k}(N+1)$$

is the number of mesh points per wavelength. We will enforce the condition

$$(2.5) \quad \lambda/h \geq 10 \quad \text{or, equivalently,} \quad kh \leq \pi/5,$$

which is a rule of thumb for approximability of second order discretizations commonly used in engineering computations [19]. We also note that, for reasons of stability, a bound on the quantity $h^2 k^3$ is also required [20]; for high wave numbers this bound is more restrictive than the bound on kh .

As a consequence of (2.5), the quantity multiplying the smoothing parameter ω in (2.4) will vary between about $-1/4$ and $9/4$ for $j = 1, \dots, N$, and plain Jacobi smoothing ($\omega = 1$) results in a slight amplification of the most oscillatory modes as well as of the smoothest modes. One can adjust ω so that the most oscillatory mode is damped, and this is the case as long as $\omega < \omega_1 := (4 - 2k^2 h^2)/(4 - k^2 h^2)$. For \mathbf{S}_ω to be an effective smoother, ω is usually chosen to maximize damping for the oscillatory half of the spectrum. This leads to the choice

$$(2.6) \quad \omega_0 = \frac{2 - k^2 h^2}{3 - k^2 h^2},$$

which is equal to the familiar optimal value of $2/3$ for the Laplacian [24, p. 11] when $k = 0$ and equals 0.61 when $\lambda/h = 10$. However, the smoothest mode is amplified for

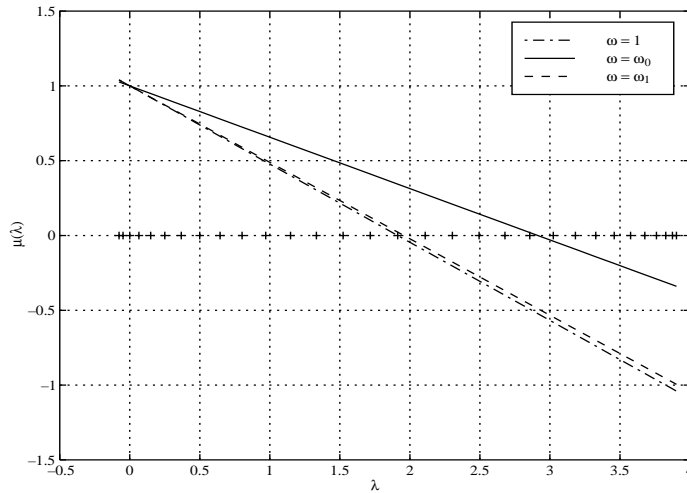


FIG. 2.1. The damping factors for the damped Jacobi relaxation plotted against the eigenvalues of \mathbf{A} (+) for $\omega = 1$, $\omega = \omega_0$, and $\omega = \omega_1$ ($N = 31$, $k = 3\pi$).

any positive choice of ω when the discrete problem is indefinite, and this is the case for the discrete Helmholtz operator (1.4) when $k^2 > \pi^2$. As can be seen from (2.4), more smooth-mode eigenvalues of \mathbf{S}_ω become larger than one in magnitude as h is increased, thus making damped Jacobi—as well as other standard smoothers—increasingly more unstable as the mesh is coarsened.

Figure 2.1 shows the damping factors μ_j for each of the eigenvalues λ_j of \mathbf{A} for wave number $k = 3\pi$ on a grid with $N = 31$. The maximal amplification occurs for the smoothest mode, corresponding to the leftmost eigenvalue of \mathbf{A} . When $\omega = \omega_0$ this amplification factor is approximately equal to

$$(2.7) \quad \rho = \rho(kh) = \frac{1}{1 - \frac{1}{3}k^2h^2}.$$

Figure 2.2 shows how ρ varies with kh . Limiting this largest amplification factor, say, to $\rho \leq 1.1$, would lead to the mesh size restriction $kh \leq 0.52$, somewhat stronger than (2.5). One also observes that, for $kh > \sqrt{6}$, this mode is once again damped.

In summary, the situation on the finest grids is similar to the positive-definite case, except for the small number of amplified smooth modes whose number and amplification factors increase as the mesh is coarsened.

2.1.2. Very coarse grids. As the mesh is coarsened, the eigenvalues of \mathbf{A} that correspond to the larger eigenvalues of the underlying differential operator disappear from the discrete problem, while the small ones—those with smooth eigenfunctions—remain. This means that, for a fixed k large enough for the differential problem to be indefinite, there is a coarsening level below which all eigenvalues are negative. For the model problem (2.1), this occurs for $kh > 2 \cos(\pi h/2)$ for any fixed $k > \pi$. In this (negative-definite) case, the damped Jacobi iteration is convergent for $\omega \in (0, \omega_2)$, with

$$\omega_2 = \frac{2 - k^2h^2}{2 \sin^2(\pi h/2) - \frac{1}{2}k^2h^2},$$

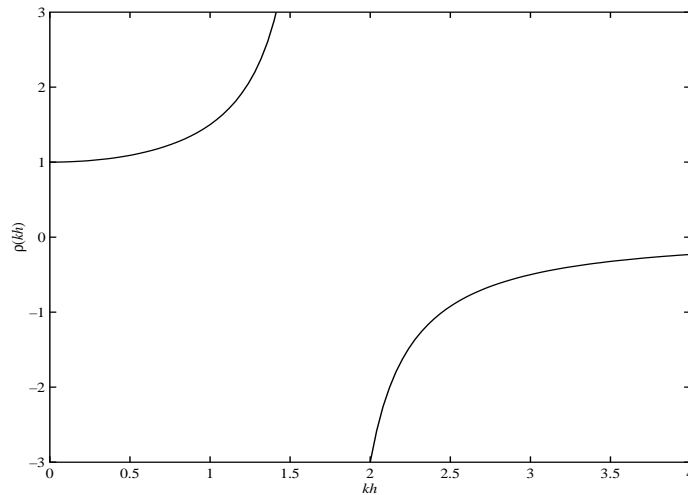


FIG. 2.2. The variation of the damping/amplification factor of the smoothest mode as a function of kh for $\omega = \omega_0$.

and the spectral radius of \mathbf{S}_ω is minimized for $\omega = 1$. This would permit the use of (undamped) Jacobi as a smoother on very coarse grids, but we shall not make use of this.

2.1.3. Intermediate grids. What remains is the difficult case: values of kh for which the problem is not yet negative definite but for which a large number of smooth modes are amplified by damped Jacobi relaxation. Jacobi smoothing and other standard smoothers are therefore no longer suitable, and it becomes necessary to use a different smoothing procedure. In [12] and [18] it was proposed to replace classical smoothers with the Kaczmarz iteration, which is Gauß–Seidel relaxation applied to the symmetric positive-definite system $\mathbf{A}\mathbf{A}^*\mathbf{v} = \mathbf{f}$ for the auxiliary variable \mathbf{v} defined by $\mathbf{A}^*\mathbf{v} = \mathbf{u}$. This method has the advantage of not amplifying any modes, but it suffers from the drawback that the damping of the oscillatory modes is very weak. In the following section we propose using Krylov subspace methods such as GMRES for smoothing. These methods possess the advantage of reducing error components on both sides of the imaginary axis without resorting to the normal equations.

2.2. Coarse grid correction. The rationale behind coarse grid correction is that smooth error components can be well represented on coarser grids, and hence a sufficiently good approximation of the error can be obtained by approximating the fine-grid residual equation using the analogous system on a coarser mesh. This assumes both that the error consists mainly of smooth modes and that the solution of the coarse grid residual equation is close to its counterpart on the fine grid. In this section, we present an analysis of what goes wrong for the Helmholtz problem.

2.2.1. Amplification of certain modes. Assume the number of interior grid points on the fine grid is odd, and consider the next coarser mesh, with $n = (N - 1)/2$ interior points. We identify \mathbb{R}^N and \mathbb{R}^n , respectively, with the spaces of grid functions on these two meshes that vanish at the endpoints, and we indicate the mesh such vectors are associated with using the superscripts h and H . Let $\mathbf{e}^h = \mathbf{u} - \mathbf{u}^h$ denote the fine-grid error, let $\mathbf{r}^h = \mathbf{f} - \mathbf{A}^h\mathbf{u}^h$ denote the residual, and let $H = 2h$ denote the

coarse mesh size. Let the coarse-to-fine transformation be given by the interpolation operator $\mathbf{I}_H^h : \mathbb{R}^n \rightarrow \mathbb{R}^N$,

$$[\mathbf{I}_H^h \mathbf{w}^H]_i := \begin{cases} \mathbf{w}_{i/2}^H, & i \text{ even,} \\ \frac{1}{2}[\mathbf{w}_{(i-1)/2}^H + \mathbf{w}_{(i+1)/2}^H], & i \text{ odd,} \end{cases} \quad i = 1, \dots, N.$$

The following indication of what can go wrong with the (exact) coarse grid correction was given in [12]: consider a fine-grid error $\mathbf{e}^h = \mathbf{v}^h$ consisting of only the smoothest eigenvector \mathbf{v}^h of \mathbf{A}^h with associated eigenvalue λ^h . The fine-grid residual is thus given by $\mathbf{r}^h = \mathbf{A}^h \mathbf{e}^h = \lambda^h \mathbf{v}^h$, and, since we are assuming that \mathbf{v}^h is smooth, its restriction $\hat{\mathbf{r}}^H := \mathbf{I}_h^H \mathbf{r}^h = \lambda^h \mathbf{I}_h^H \mathbf{v}^h$ to the coarse grid will again be close to an eigenvector of the coarse grid operator \mathbf{A}^H but with respect to a slightly different eigenvalue λ^H . The coarse grid version of the correction is

$$\mathbf{e}^H = (\mathbf{A}^H)^{-1} \hat{\mathbf{r}}^H = \lambda^h (\mathbf{A}^H)^{-1} \mathbf{I}_h^H \mathbf{v}^h \approx \frac{\lambda^h}{\lambda^H} \mathbf{I}_h^H \mathbf{v}^h.$$

Hence the error on the fine grid after the correction is

$$(2.8) \quad \mathbf{e}^h - \mathbf{I}_H^h \mathbf{e}^H \approx \mathbf{v}^h - \frac{\lambda^h}{\lambda^H} \mathbf{I}_H^h \mathbf{I}_h^H \mathbf{v}^h = \left(1 - \frac{\lambda^h}{\lambda^H}\right) \mathbf{v}^h,$$

where we have assumed that the smooth mode \mathbf{v}^h is invariant under restriction followed by interpolation. This tells us that, under the assumption that the restrictions of smooth eigenvectors are again eigenvectors of \mathbf{A}^H , the quality of the correction depends on the ratio λ^h/λ^H . If the two are equal, then the correction is perfect, but if the relative error is large, the correction can be arbitrarily bad. This occurs whenever one of λ^h, λ^H is close to the origin and the other is not. Moreover, if λ^h and λ^H have opposite signs, then the correction is in the wrong direction.

We now go beyond existing analysis and examine which eigenvalues are problematic in this sense for finite differences; a similar analysis can also be performed for linear finite elements. Consider the coarse grid eigenfunctions $\mathbf{v}_j^H = [\sin ij\pi H]_{j=1}^n$. To understand the effects of interpolation of these grid functions to the fine grid, we must examine both the first n fine-grid eigenfunctions $\{\mathbf{v}_j^h\}_{j=1}^n$ and their *complementary modes* $\{\mathbf{v}_{N+1-j}^h\}_{j=1}^n$; these are related by $[\mathbf{v}_{N+1-j}^h]_i = (-1)^{i+1} [\mathbf{v}_j^h]_i$. As is easily verified, there holds [13]

$$(2.9) \quad \mathbf{I}_H^h \mathbf{v}_j^H = c_j^2 \mathbf{v}_j^h - s_j^2 \mathbf{v}_{N+1-j}^h, \quad j = 1, \dots, n,$$

with $c_j := \cos j\pi h/2$ and $s_j := \sin j\pi h/2$, $j = 1, \dots, N$.

If full weighting is used for the restriction operator $\mathbf{I}_h^H : \mathbb{R}^N \rightarrow \mathbb{R}^n$, we have componentwise

$$[\mathbf{I}_h^H \mathbf{u}^h]_i := \frac{1}{4} \left([\mathbf{u}^h]_{2i-1} + 2 [\mathbf{u}^h]_{2i} + [\mathbf{u}^h]_{2i+1} \right), \quad i = 1, \dots, n,$$

and the relation $\mathbf{I}_H^h = 2 (\mathbf{I}_h^H)^\top$. The following mapping properties are easily established:

$$(2.10) \quad \mathbf{I}_h^H \mathbf{v}_j^h = \begin{cases} c_j^2 \mathbf{v}_j^H, & j = 1, \dots, n, \\ 0, & j = n + 1, \\ -c_j^2 \mathbf{v}_{N+1-j}^H, & j = n + 2, \dots, N, \end{cases}$$

with c_j and s_j as defined above.

If \mathbf{A}^H denotes the coarse grid discretization matrix, then the corrected iterate $\tilde{\mathbf{u}}^h := \mathbf{u}^h + \mathbf{I}_H^h(\mathbf{A}^H)^{-1}\mathbf{r}^H$ possesses the error propagation operator $\mathbf{C} := \mathbf{I} - \mathbf{I}_H^h(\mathbf{A}^H)^{-1}\mathbf{I}_h^H\mathbf{A}^h$. Denoting the eigenvalues of \mathbf{A}^h and \mathbf{A}^H by $\{\lambda_j^h\}_{j=1}^N$ and $\{\lambda_j^H\}_{j=1}^n$, respectively, we may summarize the action of \mathbf{C} on the eigenvectors using (2.9) and (2.10) as follows.

THEOREM 2.1. *The image of the fine-grid eigenfunctions $\{\mathbf{v}_h^h\}_{j=1}^N$ under the error propagation operator \mathbf{C} of the exact coarse grid correction is given by*

$$(2.11) \quad \mathbf{C}\mathbf{v}_j^h = \begin{cases} \left(1 - c_j^4 \frac{\lambda_j^h}{\lambda_j^H}\right) \mathbf{v}_j^h + s_j^2 c_j^2 \frac{\lambda_j^h}{\lambda_j^H} \mathbf{v}_{N+1-j}^h, & j = 1, \dots, n, \\ \mathbf{v}_{n+1}^h, & j = n + 1, \\ \left(1 - c_j^4 \frac{\lambda_j^h}{\lambda_{N+1-j}^H}\right) \mathbf{v}_j^h + s_j^2 c_j^2 \frac{\lambda_j^h}{\lambda_{N+1-j}^H} \mathbf{v}_{N+1-j}^h, & j = n + 2, \dots, N. \end{cases}$$

As a consequence, the two-dimensional spaces spanned by a smooth mode and its complementary mode are invariant under \mathbf{C} : $\mathbf{C}[\mathbf{v}_j^h, \mathbf{v}_{N+1-j}^h] = [\mathbf{v}_j^h, \mathbf{v}_{N+1-j}^h]\mathbf{C}_j$ with

$$(2.12) \quad \mathbf{C}_j := \begin{bmatrix} 1 - c_j^4 \frac{\lambda_j^h}{\lambda_j^H} & c_j^2 s_j^2 \frac{\lambda_{N+1-j}^h}{\lambda_j^H} \\ s_j^2 c_j^2 \frac{\lambda_j^h}{\lambda_j^H} & 1 - s_j^4 \frac{\lambda_{N+1-j}^h}{\lambda_j^H} \end{bmatrix}, \quad j = 1, \dots, n.$$

The following result shows the dependence of the matrices \mathbf{C}_j on kh .

THEOREM 2.2. *Using the notation defined above, there holds*

$$(2.13) \quad \mathbf{C}_j = \begin{bmatrix} s_j^2 \left(1 - \frac{k^2 c_j^2}{\lambda_j^H}\right) & c_j^2 \left(1 + \frac{k^2 c_j^2}{\lambda_j^H}\right) \\ s_j^2 \left(1 + \frac{k^2 s_j^2}{\lambda_j^H}\right) & c_j^2 \left(1 - \frac{k^2 s_j^2}{\lambda_j^H}\right) \end{bmatrix}, \quad j = 1, \dots, n.$$

Moreover,

$$(2.14) \quad \lim_{kh \rightarrow 0} \mathbf{C}_j = \begin{bmatrix} s_j^2 & c_j^2 \\ s_j^2 & c_j^2 \end{bmatrix}, \quad \lim_{kh \rightarrow \infty} \mathbf{C}_j = \begin{bmatrix} s_j^2(1 + c_j^2) & s_j^2 c_j^2 \\ s_j^2 c_j^2 & c_j^2(1 + s_j^2) \end{bmatrix}, \quad j = 1, \dots, n.$$

Proof. Both (2.13) and (2.14) are simple consequences of (2.12) and the representation (2.2) of the eigenvalues λ_j^h . \square

Application of the error propagation operator to a smooth mode \mathbf{v}_j^h gives

$$\mathbf{C}\mathbf{v}_j^h = \mathbf{C}[\mathbf{v}_j^h, \mathbf{v}_{N+1-j}^h] \begin{pmatrix} 1 \\ 0 \end{pmatrix} = [\mathbf{v}_j^h, \mathbf{v}_{N+1-j}^h] \mathbf{C}_j \begin{pmatrix} 1 \\ 0 \end{pmatrix}.$$

If the entries of the first column of \mathbf{C}_j are small, then this mode is damped by the coarse grid correction. However, if the (1,1)-entry is large, then this mode is amplified; and if the (2,1)-entry is large (somewhat less likely), then the smooth mode is corrupted by its complementary mode. As seen from (2.13), these difficulties occur whenever λ_j^H is small in magnitude. From the limits (2.14), it is evident that no such problems arise in the symmetric positive-definite case (a fact that is well known), but they also do not occur when kh is very large, i.e., when the coarse grid Helmholtz

operator is negative definite. These observations can be extended by returning to (2.8) and using (2.2), wherein it holds that

$$(2.15) \quad \frac{\lambda_j^h}{\lambda_j^H} = \frac{4s_j^2/h^2 - k^2}{4s_j^2c_j^2/h^2 - k^2} = 1 + \frac{s_j^4}{s_j^2c_j^2 - (kh/2)^2}.$$

That is, the coarse grid correction strongly damps smooth error modes for either very small or very large values of kh , but it may fail to do so in the intermediate range where $s_j^2c_j^2 \approx (kh/2)^2$ for some index j associated with a smooth mode.

We also note that in the limit $k = 0$ the eigenvalues of C_j are 0 and 1, so that C_j is a projection, and in this case the projection is orthogonal with respect to the inner product induced by the symmetric and positive-definite operator A^h . The projection property is lost for $k > 0$, since the coarse grid operator as we have defined it fails to satisfy the Galerkin condition $A^H = I_h^H A^h I_h^h$. (The Galerkin condition is, however, satisfied, e.g., for finite element discretizations with interpolation by inclusion.) Moreover, regardless of the type of discretization, the term A^h -orthogonality ceases to make sense once k is sufficiently large that A^h is indefinite.

2.2.2. Number of sign changes. In this section, we determine the number of eigenvalues that undergo a sign change during the coarsening process. In light of the discussion above, this number gives an indication of the number of smooth modes that are not eliminated from the error by coarse grid correction. Since these modes are not properly handled by the multigrid algorithm, another construction is needed to reduce the error associated with them. In the next section, we will introduce an *outer* Krylov subspace iteration designed for this task.

This is the only aspect of the algorithm that significantly depends on the dimensionality of the problem. Thus, here we are considering the Helmholtz equation (1.1) on the d -dimensional unit cube $(0,1)^d$, $d = 1, 2$ or 3 , with homogeneous Dirichlet boundary conditions. We consider standard finite differences (second order three-point, five-point, or seven-point discretizations of the Laplacian in one, two, or three dimensions, respectively), as well as the class of low order finite elements consisting of linear, bilinear, or trilinear elements.

We first state the issue more precisely using finite differences. In d dimensions, the eigenvalues of the discrete operator on a grid with mesh size h and N grid points in each direction are

$$(2.16) \quad \lambda_{\mathcal{I}}^h = \sum_{i=1}^d \frac{4}{h^2} \left(\sin^2 \frac{j_i \pi h}{2} \right) - k^2, \quad \mathcal{I} = \{j_1, \dots, j_d\}, \quad j_i = 1, \dots, N.$$

For any fixed multi-index \mathcal{I} , this eigenvalue is a well-defined function of h that converges to the corresponding eigenvalue of the differential operator as $h \rightarrow 0$. Our concern is the indices for which this function changes sign, for these are the troublesome eigenvalues that are not treated correctly by some coarse grid correction. As the mesh is coarsened, the oscillatory modes ($j_i > N/2$ for some i) are not represented on the next coarser mesh, but the smooth-mode eigenvalues $\{\lambda_{\mathcal{I}}^H\}$ are slightly shifted to the left with respect to their fine-grid counterparts $\{\lambda_{\mathcal{I}}^h\}$, and some of these eigenvalues change sign at some point during the coarsening process.

The following theorem gives a bound, as a function of k , on the maximal number eigenvalue sign changes occurring on all grids.

THEOREM 2.3. *For finite difference discretization of the Helmholtz equation with Dirichlet boundary conditions on the unit cube in d dimensions ($d = 1, 2, 3$), the*

number of eigenvalues that undergo a change in sign during the multigrid coarsening process is bounded above by

$$(2.17) \quad \begin{cases} k \left(\frac{1}{2} - \frac{1}{\pi} \right) \approx 0.18 k, & d = 1, \\ k^2 \left(\frac{1}{8} - \frac{1}{4\pi} \right) \approx 0.045 k^2, & d = 2, \\ k^3 \left(\frac{1}{24\sqrt{3}} - \frac{1}{6\pi^2} \right) \approx 0.0072 k^3, & d = 3. \end{cases}$$

For the finite element discretizations, the number of sign changes is bounded above by

$$(2.18) \quad \begin{cases} k \left(\frac{1}{\pi} - \frac{1}{\sqrt{12}} \right) \approx .030 k, & d = 1, \\ k^2 \left(\frac{1}{4\pi} - \frac{1}{24} \right) \approx .038 k^2, & d = 2, \\ k^3 \left(\frac{1}{6\pi^2} - \frac{1}{216} \right) \approx .012 k^3, & d = 3. \end{cases}$$

Proof. For finite differences, let η_{fine}^- denote the number of negative eigenvalues on some given fine grid, and let η_{lim}^- denote the number of negative eigenvalues of the continuous Helmholtz operator. Because eigenvalues (2.16) with the same index \mathcal{I} shift from right to left with grid coarsening, it follows that

$$(2.19) \quad \eta_{\text{lim}}^- \leq \eta_{\text{fine}}^-;$$

this is an equality for all fine enough grids, as the discrete eigenvalues tend to the continuous ones. To identify η_{lim}^- , consider the continuous eigenvalues

$$\begin{aligned} \lambda_\ell &= \pi^2 \ell^2 - k^2, & \ell \in \mathbb{N} & \quad (d = 1), \\ \lambda_{\ell,m} &= \pi^2 (\ell^2 + m^2) - k^2, & \ell, m \in \mathbb{N} & \quad (d = 2), \\ \lambda_{\ell,m,n} &= \pi^2 (\ell^2 + m^2 + n^2) - k^2, & \ell, m, n \in \mathbb{N} & \quad (d = 3). \end{aligned}$$

It is convenient to view the indices of these eigenvalues as lying in the positive orthant of a d -dimensional coordinate system. The negative eigenvalues are contained in the intersection of this orthant with a d -dimensional sphere of radius k/π centered at the origin. Let \mathcal{N} denote this intersection, and let $\hat{\mathcal{N}}$ denote the d -dimensional cube enclosing \mathcal{N} . The number of indices in $\hat{\mathcal{N}}$ is $\lfloor k/\pi \rfloor^d$, and the number in \mathcal{N} is $\rho \lfloor k/\pi \rfloor^d$, where

$$\rho = \begin{cases} \left(\frac{k}{\pi} \right) / \left(\frac{k}{\pi} \right) = 1, & d = 1, \\ \frac{1}{4} \pi \left(\frac{k}{\pi} \right)^2 / \left(\frac{k}{\pi} \right)^2 = \frac{\pi}{4}, & d = 2, \\ \frac{1}{8} \cdot \frac{4}{3} \pi \left(\frac{k}{\pi} \right)^3 / \left(\frac{k}{\pi} \right)^3 = \frac{\pi}{6}, & d = 3 \end{cases}$$

is the ratio of the volume of \mathcal{N} to that of $\hat{\mathcal{N}}$. It follows that

$$(2.20) \quad \eta_{\text{lim}}^- = \begin{cases} k \cdot \frac{1}{\pi}, & d = 1, \\ k^2 \cdot \frac{1}{4\pi}, & d = 2, \\ k^3 \cdot \frac{1}{6\pi^2}, & d = 3. \end{cases}$$

Now consider the eigenvalues of discrete problems. Again, since sign changes occur from right to left with coarsening, the mesh size that yields the maximum number of negative eigenvalues is the smallest value h for which the discrete operator

is negative semidefinite. With N mesh points in each coordinate direction, this is equivalent to

$$d \sin^2 \frac{N\pi h}{2} = \left(\frac{kh}{2}\right)^2, \quad d = 1, 2, 3.$$

Thus, $h = 2\sqrt{d}/k$, and

$$\eta_{\max}^- = \left(\frac{k}{2\sqrt{d}}\right)^d = \begin{cases} k \cdot \frac{1}{2}, & d = 1, \\ k^2 \cdot \frac{1}{8}, & d = 2, \\ k^3 \cdot \frac{1}{24\sqrt{3}}, & d = 3. \end{cases}$$

Combining (2.19) with the fact that $\eta_{\text{fine}}^- \leq \eta_{\max}^-$, it follows that

$$\eta_{\max}^- - \eta_{\text{fine}}^- \leq \eta_{\max}^- - \eta_{\text{lim}}^-.$$

The latter difference, shown in (2.17), is then a bound on number of sign changes.

For finite elements, we are concerned with the eigenvalues of the coefficient matrix \mathbf{A}^h , but it is also convenient to consider the associated operator \mathcal{A}^h defined on the finite element space V^h . The eigenvalues of \mathcal{A}^h are those of the generalized matrix eigenvalue problem

$$(2.21) \quad \mathbf{A}^h \mathbf{u}^h = \sigma^h \mathbf{M}^h \mathbf{u}^h,$$

where \mathbf{M}^h is the mass matrix. These eigenvalues tend to those of the continuous operator. Moreover, since V^H is a subspace of V^h , the Courant–Fischer min-max theorem implies that eigenvalues σ^h and σ^H with a common index shift to the right with coarsening (or to the left with refinement). In addition, since \mathbf{M}^h is symmetric positive definite, Sylvester’s inertia theorem implies that the number of negative eigenvalues of \mathbf{A}^h is the same as that of (2.21). It follows from these observations that the maximal number of negative eigenvalues of \mathbf{A}^h is bounded above by the fine-grid limit η_{lim}^- .

This is also a bound on the number of sign changes. It can be improved by examining the eigenvalues of \mathbf{A}^h more closely. Using the tensor product form of the operators, we can express these eigenvalues as

$$(2.22) \quad \lambda_{j_1, j_2, j_3}^h = \kappa_{j_1} \mu_{j_2} \mu_{j_3} + \mu_{j_1} \kappa_{j_2} \mu_{j_3} + \mu_{j_1} \mu_{j_2} \kappa_{j_3} - k^2 \mu_{j_1} \mu_{j_2} \mu_{j_3},$$

where $h = 1/(N + 1)$, the indices j_1, j_2, j_3 run from 1 to N , and

$$(2.23) \quad \kappa_j = \frac{1}{h}(2 - 2 \cos j\pi h), \quad \mu_j = \frac{h}{6}(4 + 2 \cos j\pi h).$$

Consider the requirement $\lambda_{j_1, j_2, j_3}^h \leq 0$ for all indices j_1, j_2, j_3 , so that \mathbf{A}^h is negative semidefinite. This is equivalent to

$$\frac{\kappa_{j_1}}{\mu_{j_1}} + \frac{\kappa_{j_2}}{\mu_{j_2}} + \frac{\kappa_{j_3}}{\mu_{j_3}} \leq k^2.$$

Since the expression κ_j/μ_j is monotonically increasing with j , the largest eigenvalue in d dimensions equals zero if

$$k^2 = d \cdot \frac{\kappa_N}{\mu_N} = d \cdot \frac{6}{h^2} \frac{1 - \cos N\pi h}{2 + \cos N\pi h} \approx d \cdot \frac{12}{h^2},$$

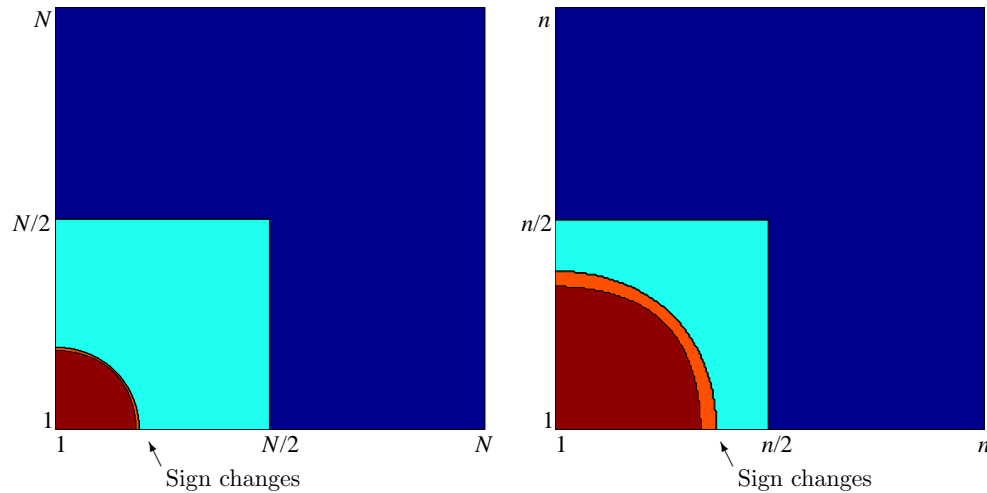


FIG. 2.3. Indices of eigenvalues undergoing a sign change during coarsening of an $N \times N$ finite element grid with $kh = \pi/5$, $kH = 2\pi/5$ (left) and during further coarsening of the next coarser ($n \times n$) grid with $kh = 2\pi/5$, $kH = 4\pi/5$ (right).

i.e., $h = \sqrt{12d}/k$. For this value of h , there are $\eta_{\max}^- = (k/\sqrt{12d})^d$ negative eigenvalues, and, on coarser meshes, the problem remains negative definite. Consequently, none of these η_{\max}^- quantities undergo a sign change, giving the bound $\eta_{\lim}^- - \eta_{\max}^-$ of (2.18). \square

Figure 2.3 gives an idea of how sign changes are distributed for bilinear elements in two dimensions. At levels where the changes take place, the indices of the eigenvalues lie in a curved strip in the two-dimensional plane of indices. Typically, there is one level where the majority of sign changes occur. As k is increased and h decreased correspondingly via (2.5), the shape of these strips remains fixed, but the number of indices contained in them grows like $O(h^{-d}) = O(k^d)$. (Note, however, that (2.5) is not needed for the analysis.) The behavior for finite differences is similar.

The remedy suggested in [12] for these difficulties consists of maintaining an approximation of the eigenspace V^H of the troublesome eigenvalues. A projection scheme is then used to orthogonalize the coarse grid correction against V^H , and the coefficients of the solution for this problematic space are obtained separately. Since it involves an explicit separate treatment of the problematic modes, this approach is restricted to cases where there are only very a small number of these.

Finally, we note that although the closed forms for the eigenvalues studied here are restricted to rectangular domains, we expect the trends displayed to be general. For example, the direction of shifts of eigenvalues, derived from the form of the matrices for finite differences and from the Courant–Fischer theorem for bilinear elements, will be the same for general domains. This assertion is also largely borne out by experiments on a nonsquare domain shown in section 4.

3. Incorporation of Krylov subspace methods. In view of the observations about smoothing in section 2.1 and coarse grid correction in section 2.2, we modify the standard multigrid method in the following way to treat Helmholtz problems:

- To obtain smoothers that are stable and still provide a strong reduction of oscillatory components, we use Krylov subspace iteration such as GMRES as smoothers on intermediate grids.

- To handle modes with eigenvalues that are either close to the origin on all grids—and hence belong to modes not sufficiently damped on any grid—or that cross the imaginary axis and are thus treated incorrectly by some coarse grid corrections, we add an outer iteration; that is, we use multigrid as a preconditioner for a GMRES iteration for (1.4).

We will demonstrate the effectiveness of this approach with a series of numerical experiments. In all tests the outer iteration is run until the stopping criterion

$$\|\mathbf{r}_m\|/\|\mathbf{r}_0\| < 10^{-6}$$

is satisfied, where $\mathbf{r}_m = \mathbf{f} - \mathbf{A}\mathbf{u}_m$ is the residual of the m th GMRES iterate and the norm is the vector Euclidean norm. The multigrid algorithm is a V-cycle in all cases; the smoothing schedules are specified below.

3.1. GMRES accelerated multigrid. We begin with an experiment for the one-dimensional Helmholtz equation on the unit interval with forcing term $f = 0$ and inhomogeneous Dirichlet boundary condition $u(0) = 1$ on the left and Sommerfeld condition on the right. We discretize using linear finite elements on a uniform grid, where the discrete right-hand side \mathbf{f} is determined by the boundary conditions. We apply both a V-cycle multigrid algorithm and a GMRES iteration preconditioned by the same V-cycle multigrid method. The smoother in these tests is one step of damped Jacobi iteration for both presmoothing and postsmoothing, using $\omega = (12 - 4k^2h^2)/(18 - 3k^2h^2)$, the analogue of (2.6) for finite element discretization that provides maximal damping of oscillatory modes. The initial guess was a vector with normally distributed entries of mean zero and variance one, generated by the Matlab function `randn`.

Table 3.1 shows the iteration counts for increasing numbers of levels beginning with fine grids containing $N = 256, 512,$ and 1024 elements, for wave numbers $k = 4\pi$ and $k = 8\pi$, which correspond to two and four wavelengths in the unit interval, respectively.

TABLE 3.1

Iteration counts for multigrid V-cycle as an iteration and as a preconditioner for GMRES applied to the one-dimensional model Helmholtz problem with damped Jacobi smoothing. A dash denotes divergence of the iteration.

# levels	256 elements		512 elements		1024 elements	
	$k = 4\pi$	$k = 8\pi$	$k = 4\pi$	$k = 8\pi$	$k = 4\pi$	$k = 8\pi$
2	7	3	7	3	7	3
3	7	5	7	5	7	5
4	7	6	7	6	7	5
5	7	6	7	6	7	6
6	16	8	7	6	7	7
7	–	9	13	8	7	10
8	–	11	–	9	9	12
9	–	11	–	11	–	16
10	–	–	–	11	–	16

We observe first that both methods display typical h -independent multigrid behavior until the mesh size on the coarsest grid reaches $kh \approx \pi/2$. (With 256 elements, this occurs for $k = 4\pi$ with 6 levels, coarsest mesh $h = 1/8$, and for $k = 8\pi$ with 5 levels, coarsest $h = 1/16$.) At this point both methods require noticeably more iterations, the increase being much more pronounced in the stand-alone multigrid case.

TABLE 3.2

Iteration counts for the two-dimensional problem for fine grids with $k = 4\pi$ and $k = 8\pi$ on 128×128 and 256×256 meshes. A dash denotes divergence of the iteration.

# levels	128×128 elements		256×256 elements		512×512 elements	
	$k = 4\pi$	$k = 8\pi$	$k = 4\pi$	$k = 8\pi$	$k = 4\pi$	$k = 8\pi$
2	12	7	12	7	12	6
3	12	7	12	7	12	6
4	12	7	22	11	12	7
5	13	8	—	33	12	7
6	—	15	—	64	12	7
7	—	19	—	64	12	8
8	—	—	—	19	—	15

When yet coarser levels are added, multigrid diverges, whereas the multigrid preconditioned GMRES method again settles down to an h -independent iteration count, which does, however, increase with k .

Table 3.2 shows the same iteration counts for the two-dimensional Helmholtz problem on the unit square with a second order absorbing boundary condition (see [1, 14]) imposed on all four sides and discretized using bilinear quadrilateral finite elements on a uniform mesh. Since the problem cannot be forced with a radiation condition on the entire boundary, in this and the remaining examples of section 3, an inhomogeneity was imposed by choosing a discrete right-hand side consisting of a random vector with mean zero and variance one, generated by `randn`. The initial guess was identically zero. (Trends for problems with smooth right-hand sides were the same.) In addition, for all two-dimensional problems, we use two Jacobi pre- and postsmoothing steps whenever Jacobi smoothing is used. The damping parameter ω is chosen to maximize damping of the oscillatory modes. For the grids on which we use damped Jacobi smoothing this optimum value was determined to be $\omega = 8/9$. The results show the same qualitative behavior as for the one-dimensional problem, in that multigrid begins to diverge as coarse levels are added while the GMRES-accelerated iteration converges in an h -independent number of iterations growing with k , although with a larger number of iterations than in the one-dimensional case.

A natural question is whether corrections computed on the very coarse grids, in particular those with meshes containing fewer than ten points per wavelength ($2\pi/k < 10$), make any contribution at all towards reducing the error. We investigate this in a sequence of experiments with GMRES accelerated multigrid with $k = 8\pi$, where we omit all calculations—be they smoothing or direct solves—on an increasing number of coarse levels.

The results, for a one-dimensional example with 512 elements and a two-dimensional example with 256^2 elements are in Table 3.3. The leftmost entries of the table show the iteration counts when no coarse grid information is used, i.e., for GMRES with preconditioning by two steps of Jacobi iteration. Reading from left to right, subsequent entries show the counts when smoothings on a succession of coarser grids are included, but no computations are done at grid levels below that of the coarsest grid. For the rightmost entry, a direct solve was done on the coarsest mesh; this is a full V-cycle computation. The results in two dimensions indicate improved performance for coarse grids with mesh width less than $1/16$, which corresponds to four points per wave; performance degrades, but not dramatically so, for coarser meshes. For the one-dimensional test, the coarse grid at level two, which has only two points per wavelength, still accelerates the outer iteration. These results also show that multigrid

TABLE 3.3

Iteration counts when varying amount of course grid information is used. Table entries further to the right result from using more multigrid levels than those to the left. Test problems are $k = 8\pi$ with 512 elements in the one-dimensional example and 256^2 elements in the two-dimensional example.

		No V-Cycle					Full V-cycle			
		↓					↓			
1D	h^{-1} for coarsest grid	512	256	128	64	32	16	8	4	
	GMRES iterations	152	78	42	25	18	17	16	16	
2D	h^{-1} for coarsest grid	256	128	64	32	16	8	4	2	
	GMRES iterations	387	173	89	57	52	64	63	63	

is dramatically superior to simple Jacobi-preconditioned GMRES.

These results show that, although multigrid by itself may diverge, it is nevertheless a powerful enough preconditioner for GMRES to converge in an h -independent number of steps. Two additional questions are whether replacing the unstable Jacobi smoother with a Krylov subspace iteration leads to a convergent multigrid method and how sensitive convergence behavior is as a function of the wave number k . We address the former in the following section.

3.2. GMRES as a smoother. In this section we replace the unstable Jacobi smoother with GMRES smoothing. We use GMRES on all levels j where $kh_j \geq 1/2$ and continue using damped Jacobi relaxation when $kh_j < 1/2$. This choice is motivated by the discussion at the end of section 2.1.1, and it ensures that the largest amplification factor for the Jacobi smoother does not become too large. The results of section 2.1.2 show that we could switch back to Jacobi smoothing for very coarse meshes, but we have not explored this option.

3.2.1. Nonconstant preconditioners. This introduces a slight complication with regard to the *outer* GMRES iteration when multigrid is used as a preconditioner. The inner GMRES smoothing steps are not linear iterations, and therefore a different preconditioner is being applied at every step of the outer iteration. A variant of GMRES able to accommodate a changing preconditioner (known as flexible GMRES (FGMRES)) is due to Saad [27]. It requires the following minor modification of the standard (right preconditioned) GMRES algorithm: if the orthonormal basis of the $(m+1)$ st Krylov space $\mathcal{K}_{m+1}(\mathbf{A}\mathbf{M}^{-1}, \mathbf{r}_0)$ in the case of a constant preconditioner \mathbf{M} is denoted by $\mathbf{V}_{m+1} = [\mathbf{v}_1, \dots, \mathbf{v}_{m+1}]$, then the Arnoldi relation $\mathbf{A}\mathbf{M}^{-1}\mathbf{V}_m = \mathbf{V}_{m+1}\tilde{\mathbf{H}}_m$ holds with an $(m+1) \times m$ upper Hessenberg matrix $\tilde{\mathbf{H}}_m$. If the preconditioning and matrix multiplication step

$$\mathbf{z}_m := \mathbf{M}^{-1}\mathbf{v}_m, \quad \mathbf{w} := \mathbf{A}\mathbf{z}_m,$$

is performed with a changing preconditioner $\mathbf{M} = \mathbf{M}_m$, this results in the modified Arnoldi relation

$$\mathbf{A}\mathbf{Z}_m = \mathbf{V}_{m+1}\tilde{\mathbf{H}}_m,$$

where $\mathbf{Z}_m = [\mathbf{z}_1, \dots, \mathbf{z}_m]$. The residual vector is now minimized over the space $\text{span}\{\mathbf{z}_1, \dots, \mathbf{z}_m\}$, which need no longer be a Krylov space. This requires storing the vectors $\{\mathbf{z}_j\}$ in addition to the orthonormal vectors $\{\mathbf{v}_j\}$, which form a basis of $\text{span}\{\mathbf{A}\mathbf{z}_j : j = 1, \dots, m\}$.

TABLE 3.4

Manually optimized GMRES smoothing schedule for the two-dimensional model Helmholtz problem: “J” denotes Jacobi smoothing and “D” denotes a direct solve. The FGMRES algorithm uses the multigrid V-cycle as a preconditioner.

256 × 256, $k = 4\pi$, $o_{\max} = 6$			
# levels	Smoothing schedule	MG	FGMRES
6	J J J J 13 D	8	6
7	J J J J 13 16 D	8	6
8	J J J J 13 16 1 D	8	6
9	J J J J 13 16 1 0 D	8	6
128 × 128, $k = 8\pi$, $o_{\max} = 7$			
# levels	Smoothing schedule	MG	FGMRES
4	J J 25 D	9	7
5	J J 25 39 D	9	7
6	J J 25 39 24 D	9	7
7	J J 25 39 16 2 D	9	7
8	J J 25 39 16 2 0 D	9	7
256 × 256, $k = 8\pi$, $o_{\max} = 7$			
# levels	Smoothing schedule	MG	FGMRES
5	J J J 23 D	9	7
6	J J J 23 42 D	9	7
7	J J J 23 39 0 D	9	7
8	J J J 23 38 0 0 D	9	7
9	J J J 23 38 0 0 0 D	9	7
256 × 256, $k = 16\pi$, $o_{\max} = 10$			
# levels	Smoothing schedule	MG	FGMRES
4	J J 34 D	12	10
5	J J 38 29 D	12	10
6	J J 35 20 6 D	12	10
7	J J 35 20 6 3 D	12	10
8	J J 35 20 6 3 0 D	12	10
9	J J 35 20 6 3 0 0 D	12	10
256 × 256, $k = 32\pi$, $o_{\max} = 18$			
# levels	Smoothing schedule	MG	FGMRES
3	J 34 D	21	18
4	J 39 39 D	22	18
5	J 35 33 5 D	23	18
6	J 35 33 5 3 D	23	18
7	J 35 33 5 3 2 D	23	18
8	J 35 33 5 3 2 0 D	23	18
9	J 35 32 6 5 3 0 0 D	23	18

3.2.2. Hand-tuned smoothing schedules. Numerical experiments with a fixed number of GMRES smoothing steps at every level did not result in good performance. To get an idea of an appropriate smoothing schedule, we proceed as follows. For given k , we calculate the number o_{\max} of FGMRES iterations needed with j -level multigrid preconditioning, where we use Jacobi smoothing on all grids for which $kh_i < 1/2$ and do a direct solve at the next coarser grid, making j grids in all. We then replace the direct solve on the coarsest grid of the j -level scheme with GMRES smoothing on this grid, coupled with a direct solve on the next coarser grid, and determine the smallest number m_j of GMRES smoothing steps required for the outer iteration to converge in o_{\max} steps. For example, for the first line of Table 3.4, six outer FGMRES steps were needed for a 5-level scheme, and then $m_5 = 13$ GMRES smoothing steps were needed for the outer iteration of the new 6-level preconditioner to converge in six steps. When the number m_j has been determined, we could fix the number of

GMRES smoothing steps to m_j on this grid, add one coarser level, determine the optimal number of GMRES smoothing steps on the coarser grid, and continue in this fashion until the maximal number of levels is reached. This approach is modified slightly by, whenever possible, trying to reduce the number of smoothings on finer levels once coarser levels have been added. This is often possible, since replacing the exact solve on the coarsest grid with several GMRES smoothing steps often has a regularizing effect, avoiding some damage possibly done by an exact coarse grid correction in modes whose eigenvalues are not well represented on the coarse grid. This hand-tuning procedure gives insight into the best possible behavior of this algorithm.

In contrast to classical linear smoothers, whose damping properties for different modes is fixed, the damping properties of GMRES depend on the initial residual. In particular, since GMRES is constructed to minimize the residual, it will most damp those modes that lead to the largest residual norm reduction. For this reason, we will favor postsmoothing over presmoothing to prevent the unnecessary damping of smoother modes that should be handled by the coarse grid correction. We do include two GMRES presmoothing steps to avoid overly large oscillatory components in the residual prior to restricting it to the next lower level, which could otherwise lead to spurious oscillatory error components being introduced by the coarse grid correction.

The results are shown in Table 3.4. The entry “D” denotes a direct solve on the corresponding level, and “J” indicates that damped Jacobi smoothing was used on this level. Looking at the smoothing schedules, we observe a “hump” in the number of GMRES smoothing steps on the first two levels on which GMRES smoothing is used. Below this, the number decreases and is often zero for the coarsest levels. However, GMRES smoothing still helps on levels which are extremely coarse with regard to resolution of the waves: in the case $k = 32\pi$, for instance, performing three GMRES smoothing steps on level 4 (which corresponds to $1/2$ point per wavelength) still improves convergence.

We remark that the number of outer iterations in all these tests, for both preconditioned FGMRES and stand-alone MG, is the same as for the corresponding two-grid versions of these methods, so we cannot expect faster convergence with respect to the wave number k . We also note that the number of iterations for multigrid is very close to that for FGMRES with multigrid preconditioning. We believe this is because the relatively large number of GMRES smoothing steps on intermediate levels eliminates lower frequency errors, and this mitigates the effects of axis crossings. We will return to this point in section 3.4.

3.3. A stopping criterion based on L^2 -sections. Hand tuning as in the previous section is clearly not suitable for a practical algorithm. In this section, we develop a heuristic for finite element discretizations that automatically determines a stopping criterion for the GMRES smoother. This technique is based on an idea introduced in [32].

We briefly introduce some standard terminology for multilevel methods applied to second order elliptic boundary value problems on a bounded domain $\Omega \subset \mathbb{R}^2$ (see [34]). We assume a nested hierarchy of finite element spaces

$$V_1 \subset V_2 \subset \cdots \subset V_J \subset V = H^1(\Omega)$$

in which the largest space V_J corresponds to the grid on which the solution is sought. We require the L^2 -orthogonal projections $Q_\ell : V \rightarrow V_\ell$, defined by

$$(Q_\ell u, v) = (u, v) \quad \forall v \in V_\ell, \quad \ell = 1, \dots, J,$$

where (\cdot, \cdot) denotes the L^2 -inner product on Ω . Let $\Phi_\ell = \{\phi_1^{(\ell)}, \dots, \phi_{n_\ell}^{(\ell)}\}$ denote the basis of the finite element space V_ℓ of dimension n_ℓ used in defining the stiffness and mass matrices. By the nestedness property $V_\ell \subset V_{\ell+1}$, there exists an $n_{\ell+1} \times n_\ell$ matrix $\mathcal{I}_\ell^{\ell+1}$ whose columns contain the coefficients of the basis Φ_ℓ in terms of the basis $\Phi_{\ell+1}$, so that, writing the bases as row vectors,

$$[\phi_1^{(\ell)}, \dots, \phi_{n_\ell}^{(\ell)}] = [\phi_1^{(\ell+1)}, \dots, \phi_{n_{\ell+1}}^{(\ell+1)}] \mathcal{I}_\ell^{\ell+1}.$$

The stopping criterion we shall use for the GMRES smoothing iterations is based on the representation of the residual r_ℓ of an approximate solution \tilde{u}_ℓ of the level- ℓ equation as the sum of differences of L^2 -projections,

$$r_\ell = (I - Q_{\ell-1})r_\ell + (Q_{\ell-1} - Q_{\ell-2})r_\ell + \dots + (Q_2 - Q_1)r_\ell + Q_1r_\ell,$$

which we refer to as *residual sections*. The following result for coercive problems, which was proven in [32], shows that the error $u_\ell - \tilde{u}_\ell$ is small if each appropriately weighted residual section is small.

THEOREM 3.1. *Assume the underlying elliptic boundary value problem is H^1 -elliptic and $H^{1+\alpha}$ -regular with $\alpha > 0$. Then there exists a constant c independent of the level ℓ such that the $H^1(\Omega)$ -norm of the error on level ℓ is bounded by*

$$(3.1) \quad \|u_\ell - \tilde{u}_\ell\|_1 \leq c \left(\|Q_1r_\ell\| + \sum_{j=2}^{\ell} h_j^\alpha \|(Q_j - Q_{j-1})r_\ell\| \right).$$

The boundary value problem (1.1)–(1.3) under consideration is not H^1 -elliptic and therefore does not satisfy the assumptions of this theorem. We have found, however, that the bound (3.1) suggests a useful stopping criterion: terminate the GMRES smoothing iteration on level ℓ as soon as the residual section $(Q_\ell - Q_{\ell-1})r_\ell$ has become sufficiently small. To obtain a formula for the computation of these sections, assume the residual r_ℓ is represented by the coefficient vector \mathbf{r}_ℓ in terms of the dual basis of Φ_ℓ . The representation of $Q_{\ell-1}r_\ell$ with respect to the dual basis of $\Phi_{\ell-1}$ is then given by the coefficient vector $\mathcal{I}_\ell^{\ell-1}\mathbf{r}_\ell \in \mathbb{C}^{n_{\ell-1}}$, where $\mathcal{I}_\ell^{\ell-1} := (\mathcal{I}_{\ell-1}^\ell)^\top$. Returning to the representation with respect to the basis $\Phi_{\ell-1}$ requires multiplication with $\mathbf{M}_{\ell-1}^{-1}$, so that we obtain

$$\|Q_{\ell-1}r_\ell\|^2 = (Q_{\ell-1}r_\ell, Q_{\ell-1}r_\ell) = (\mathcal{I}_\ell^{\ell-1}\mathbf{r}_\ell)^\top \mathbf{M}_{\ell-1}^{-1} \mathcal{I}_\ell^{\ell-1}\mathbf{r}_\ell.$$

If the sequence of triangulations underlying the finite element spaces V_ℓ is quasi-uniform, then the mass matrix of level ℓ is uniformly equivalent to the identity scaled by h^d , where d denotes the dimension of the domain. For the case $d = 2$ under consideration, this means that the Euclidean inner product on the coordinate space \mathbb{C}^{n_ℓ} , denoted by $(\cdot, \cdot)_E$, when scaled by h_ℓ^2 , is uniformly equivalent (with respect to the mesh size) to the L^2 -inner product on V_ℓ . Therefore, the associated norms satisfy

$$ch_\ell^2\|\mathbf{v}_\ell\|_E^2 \leq \|v_\ell\|^2 = \mathbf{v}_\ell^\top \mathbf{M}_\ell \mathbf{v}_\ell \leq Ch_\ell^2\|\mathbf{v}_\ell\|_E^2 \quad \forall v_\ell \in V_\ell,$$

where \mathbf{v}_ℓ is the coordinate vector of v_ℓ with respect to Φ_ℓ . Using this norm equivalence it is easily shown that

$$\begin{aligned} c\|(\mathcal{I} - h_\ell^2/h_{\ell-1}^2 \mathcal{I}_{\ell-1}^\ell \mathcal{I}_\ell^{\ell-1})\mathbf{r}_\ell\|_E &\leq h_\ell\|(I - Q_{\ell-1})r_\ell\| \\ &\leq C\|(\mathcal{I} - h_\ell^2/h_{\ell-1}^2 \mathcal{I}_{\ell-1}^\ell \mathcal{I}_\ell^{\ell-1})\mathbf{r}_\ell\|_E \end{aligned}$$

for some constants c and C uniformly for all levels ℓ . As a result, the residual sections may be computed sufficiently accurately without the need for inverting mass matrices.

In [32], it was suggested that the GMRES smoothing iteration for a full multigrid cycle be terminated as soon as the residual section on the given level is on the order of the discretization error on that level. For the problem under consideration here, we shall use the relative reduction of L^2 -sections as a stopping criterion, so that roughly an equal error reduction for all modes is achieved in one V-cycle. On the first level on which GMRES smoothing is used, we have the additional difficulty that many eigenvalues may be badly approximated on the next-coarser level. For this reason, it is better to also smooth the oscillatory modes belonging to the next lower level and base the stopping criterion on the residual section $(I - Q_{\ell-2})r_\ell$ instead; we will use this “safer” choice on all levels. Numerical experiments with optimal smoothing schedules have shown the relative reduction of this residual section to scale like kh_ℓ , so that we arrive at the stopping criterion

$$(3.2) \quad \left\| \mathbf{r} - \frac{h_\ell^2}{h_{\ell-2}^2} \mathcal{I}_{\ell-1}^\ell \mathcal{I}_{\ell-2}^{\ell-1} (\mathcal{I}_{\ell-1}^\ell \mathcal{I}_{\ell-2}^{\ell-1})^\top \mathbf{r} \right\|_E \leq \gamma kh_\ell.$$

A complete description of the multigrid V-cycle algorithm starting on the finest level ℓ is as follows.

ALGORITHM 3.1. $\tilde{\mathbf{u}}_\ell = MG(\mathbf{u}_\ell^{(0)}, \mathbf{f}_\ell)$. *Multigrid V-cycle with GMRES smoothing on coarse levels*

```

if  $\ell = 1$ 
     $\tilde{\mathbf{u}}_\ell := \mathbf{A}_\ell^{-1} \mathbf{f}_\ell$ 
else
    if  $kh_\ell < 1/2$ 
        perform  $m_1$  steps of damped Jacobi smoothing to obtain  $\mathbf{u}_\ell^{(1)}$ 
    else
        perform 2 steps of GMRES smoothing to obtain  $\mathbf{u}_\ell^{(1)}$ 
    endif
     $\mathbf{u}_\ell^{(2)} := \mathbf{u}_\ell^{(1)} + \mathcal{I}_{\ell-1}^\ell MG(0, \mathcal{I}_{\ell-1}^{\ell-1} (\mathbf{f}_\ell - \mathbf{A}_\ell \mathbf{u}_\ell^{(1)}))$ 
    if  $kh_\ell < 1/2$ 
        perform  $m_2$  steps of damped Jacobi smoothing to obtain  $\tilde{\mathbf{u}}_\ell$ 
    else
        perform GMRES smoothing until stopping criterion (3.2) is satisfied
        or  $m = m_{\max}$  to obtain  $\tilde{\mathbf{u}}_\ell$ 
    endif
endif

```

In the multigrid V-cycle, Algorithm 3.1 is used recursively beginning with the finest level and iterated until the desired reduction of the relative residual is achieved on the finest level. In the FGMRES variant, Algorithm 3.1 represents the action of the inverse of a preconditioning operator being applied to the vector \mathbf{f}_ℓ .

3.4. Experiments with automated stopping criterion. We now show how the multigrid solver and preconditioner perform with the automated stopping criterion for GMRES smoothing. Each method is applied to the two-dimensional Helmholtz problem on the unit square with a second order absorbing boundary condition and random right-hand side data. In these tests, we used $\gamma = 0.1$ in (3.2), and we also imposed an upper bound m_{\max} on the number of GMRES smoothing steps, terminating the smoothing if the stopping criterion is not satisfied after m_{\max} steps; we tested two

TABLE 3.5

Iteration counts for multigrid and multigrid-preconditioned FGMRES for various fine-grid sizes and wave numbers. In all cases, GMRES smoothing is performed on levels for which $kh > 1/2$ and the smoothing is terminated by the L^2 -section stopping criterion or when m_{\max} smoothing steps are reached. A dash denotes divergence.

$N \setminus k$	Multigrid						MG-preconditioned FGMRES					
	2π	4π	8π	16π	32π	64π	2π	4π	8π	16π	32π	64π
	$m_{\max} = 40$											
64	12	12	13				7	8	9			
128	12	12	13	16			7	8	9	13		
256	12	12	13	17	27		7	8	9	13	20	
512	12	12	13	16	27	78	7	8	9	13	21	36
	$m_{\max} = 20$											
64	12	12	13				7	8	9			
128	12	12	13	22			7	8	9	16		
256	12	12	13	21	201		7	8	10	16	37	
512	12	12	13	21	331	—	7	8	10	16	36	80

values, $m_{\max} = 40$ and $m_{\max} = 20$. At fine-grid levels, where damped Jacobi smoothing is used, the number of presmoothings and postsmoothings was $m_1 = m_2 = 2$.

We present three sets of results. Table 3.5 shows iteration counts for a variety of wave numbers and mesh sizes. Table 3.6 examines performance in more detail by showing the automatically generated smoothing schedules for two wave numbers, $k = 8\pi$ and $k = 32\pi$. Finally, to give an idea of efficiency, Table 3.7 shows an estimate for the operation counts (multiplications) required for the problems treated in Table 3.6.

TABLE 3.6

Smoothing schedules with automated stopping criterion for selected parameters.

$k = 8\pi, m_{\max} = 40$											
	Grid	# levels	Smoothing schedule								Iterations
FGMRES	64×64	6	J	20	17	11	2	D		9	
	128×128	7	J	J	19	16	11	2	D	9	
MG	64×64	6	J	16	17	11	2	D		13	
	128×128	7	J	J	18	16	11	2	D	13	
$k = 8\pi, m_{\max} = 20$											
	Grid	# levels	Smoothing schedule								Iterations
FGMRES	64×64	6	J	19	17	12	2	D		9	
	128×128	7	J	J	18	16	11	2	D	9	
MG	64×64	6	J	16	17	11	2	D		13	
	128×128	7	J	J	18	16	11	2	D	13	
$k = 32\pi, m_{\max} = 40$											
	Grid	# levels	Smoothing schedule								Iterations
FGMRES	256×256	8	J	34	38	22	1	0	0	D	20
	512×512	9	J	J	33	38	22	1	0	0	D
MG	256×256	8	J	32	36	18	1	0	0	D	27
	512×512	9	J	J	32	37	19	1	0	0	D
$k = 32\pi, m_{\max} = 20$											
	Grid	# levels	Smoothing schedule								Iterations
FGMRES	256×256	8	J	20	20	18	1	0	0	D	37
	512×512	9	J	J	20	20	18	1	0	0	D
MG	256×256	8	J	20	20	16	1	0	0	D	201
	512×512	9	J	J	20	20	16	1	0	0	D

TABLE 3.7

Operation counts (in millions) for selected parameters with $m_{\max} = 40$.

Grid	$k = 8\pi$		$k = 32\pi$	
	MG	FGMRES	MG	FGMRES
64×64	13.2	13.3		
128×128	24.0	22.1		
256×256	61.2	43.2	1091.2	971.1
512×512	196.6	148.1	1418.1	1377.8

We make the following observations on these results:

- For low wave numbers, the number of iterations of stand-alone multigrid is close to that for FGMRES. The difference increases as the wave number increases, especially for the case $m_{\max} = 20$. For large enough k , multigrid fails to converge, whereas MG-preconditioned FGMRES is robust. This behavior is explained by the results of section 2.2.2. For large wave numbers, the increased number of amplified modes eventually causes multigrid to fail; a larger number of smoothing steps mitigates this difficulty, presumably by eliminating some smooth errors. The (outer) FGMRES iteration handles this situation in a robust manner.
- The automated stopping criterion leads to smoothing schedules close to those obtained by hand tuning (see Table 3.4), and correspondingly similar outer iteration counts.
- The operation counts shown in Table 3.7 suggest that MG-preconditioned FGMRES is more efficient than stand-alone multigrid even when the latter method is effective. Operation counts for MG-preconditioned FGMRES with $m_{\max} = 20$ (not shown) indicate that the costs for this strategy, which uses more outer iterations but fewer smoothing steps than when $m_{\max} = 40$, are essentially the same, although the larger number of outer FGMRES steps requires more memory to store more fine-grid vectors.
- For fixed wave number, outer iteration counts are mesh independent, so that standard “multigrid-like” behavior is observed. The costs per unknown (for fixed k and smoothing schedule) also display textbook multigrid efficiency, i.e., they are constant. However, because Jacobi smoothing is less expensive than GMRES smoothing, during the initial stages of mesh refinement the operation counts grow at less than a linear rate.
- The growth in outer iteration counts with increasing wave number is slower than linear in k . The operation counts increase more rapidly, however, because of the increased number of smoothing steps required for larger wave numbers.

We also note that GMRES has nontrivial storage requirements; we expect other Krylov subspace methods with lower storage requirements (e.g., QMR [15] or Bi-CGSTAB [33]) to perform in a similar manner.

4. Application to an exterior problem. As a final example we apply the algorithm to an exterior scattering problem for the Helmholtz equation as given in (1.1)–(1.3). The domain Ω consists of the exterior of an ellipse bounded externally by a circular artificial boundary Γ_∞ on which we impose the exact nonlocal Dirichlet-to-Neumann (DtN) boundary condition (see [21]). The source function is $f = 0$; forcing is due to the boundary condition on the boundary Γ of the scatterer, given by

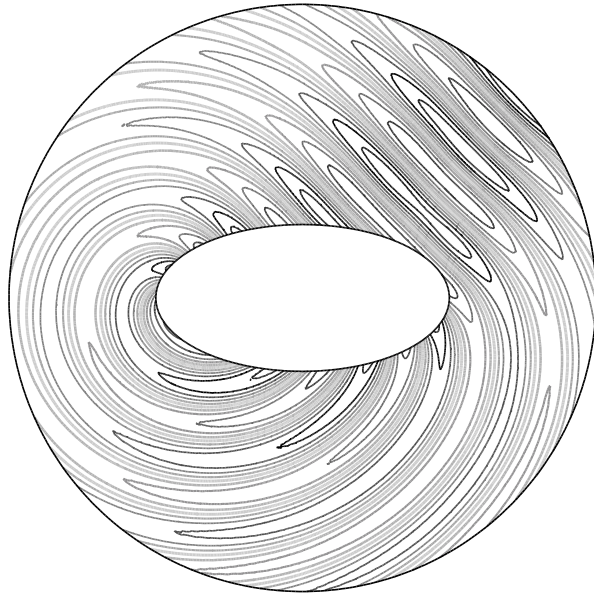


FIG. 4.1. Contour plot of the solution of the Dirichlet problem with wave number $k = 8\pi$.

$$u(x, y) = g(x, y) \quad \text{or} \quad \frac{\partial u(x, y)}{\partial n} = \frac{\partial g(x, y)}{\partial n}, \quad (x, y) \in \Gamma,$$

with data $g(x, y) = -e^{ik(x \cos \alpha + y \sin \alpha)}$ representing a plane wave incident at angle α to the positive x -axis. The solution u represents the scattered field associated with the obstacle and incident field g ; the resulting total field $u+g$ then satisfies a homogeneous Dirichlet or Neumann boundary condition on Γ , respectively. An angle of incidence $\alpha = \pi/4$ was chosen to avoid a symmetric solution. The problems were discretized using linear finite elements beginning with a very coarse mesh which is successively refined uniformly to obtain a hierarchy of nested finite element spaces. The finest mesh, obtained after five refinement steps, contains 32768 degrees of freedom. Several combinations of k and h were tested, where in each case $kh < 0.5$ on the finest mesh. Figure 4.1 shows a contour plot of the solution u of the Dirichlet problem for $k = 8\pi$. The computations make use of the PDE Toolbox of the Matlab 5.3 computing environment.

The problems were solved using both the stand-alone and FGMRES-accelerated versions of multigrid, with GMRES smoothing using the residual section stopping criterion with $\gamma = 0.1$, outer stopping criterion requiring residual reduction by a factor of 10^{-6} as in section 3, and zero initial guess. We used the maximal number of levels in all examples with the exception of the Dirichlet problem for $k = 8\pi$, where we also varied the number of levels from six down to two. The results are shown in Table 4.1. The table gives the wave number k and the length of the ellipse E in wavelengths $\lambda = 2\pi/k$. The third column gives the maximum value of kh on the finest mesh and the fourth column indicates the number of levels used in each computation. The last two columns list the iteration counts.

We observe that the preconditioned iteration performs well in all cases with a growth in number of iterations slower than linear in k . The stand-alone multigrid variant performs less well in comparison, requiring more than 100 steps to converge

TABLE 4.1

Iteration counts for the exterior scattering problem with Dirichlet or Neumann plane wave data on the boundary of an ellipse for various wave numbers, grid sizes, and numbers of levels. A dash denotes divergence.

k	Size $E[\lambda]$	$(kh_{\max})_{\text{fine}}$	# Levels	MG	FGMRES
Dirichlet problem					
2π	1	.10	6	36	13
		.21	5	27	12
		.42	4	26	12
4π	2	.21	6	38	16
		.42	5	27	14
8π	4	.42	6	41	20
			5	—	28
			4	100	26
			3	41	16
			2	41	13
Neumann problem					
2π	1	.10	6	—	21
		.21	5	37	15
		.42	4	21	12
4π	2	.21	6	—	28
		.42	5	53	21
8π	4	.42	6	—	32

in several cases and even diverging in one case. This is particularly the case for the Neumann problem, where the superiority of the preconditioned variant is even more pronounced. For the Neumann problems we also notice a slight growth in iteration counts for fixed k and decreasing h .

5. Conclusions. The results of this paper show that the addition of Krylov subspace iteration to multigrid, both as a smoother and as an outer accelerating procedure, enables the construction of a robust multigrid algorithm for the Helmholtz equation. GMRES is an effective smoother for grids of intermediate coarseness, in that it appears not to amplify any error modes and in addition tends to have a regularizing effect on the contribution to the coarse grid correction coming from smoothing on a given level. The combination of our multigrid algorithm as a preconditioner with FGMRES is effective in handling the deficiencies of standard multigrid methods for the Helmholtz equation, and the outer FGMRES acceleration is necessary, particularly for high wave numbers. In addition, results in the paper indicate that grids too coarse to result in a meaningful discretization of the Helmholtz equation may still provide some useful information for coarse grid corrections. Using an automated stopping criterion based on L^2 -sections of the residual leads to smoothing cycles that are close to hand-tuned optimal smoothing schedules.

An important aspect of our algorithm is that it consists of familiar building blocks and is thus easily implemented. For very large wave numbers for which the discretization must not only keep kh but also k^3h^2 small, the grid hierarchy will contain more grids fine enough to use Jacobi smoothing, thus making the algorithm more efficient. The result is a multigrid method that appears to converge with a rate independent of the mesh size h and with only moderate dependence on the wave number k . Finally, the numerical results show that we are able to effectively solve Helmholtz problems with wave numbers of practical relevance.

REFERENCES

- [1] A. BAMBERGER, P. JOLY, AND J. E. ROBERTS, *Second-order absorbing boundary conditions for the wave equation: A solution for the corner problem*, SIAM J. Numer. Anal., 27 (1990), pp. 323–352.
- [2] R. E. BANK, *A comparison of two multilevel iterative methods for nonsymmetric and indefinite elliptic finite element equations*, SIAM J. Numer. Anal., 18 (1981), pp. 724–743.
- [3] R. E. BANK AND C. C. DOUGLAS, *Sharp estimates for multigrid rates of convergence with general smoothing and acceleration*, SIAM J. Numer. Anal., 22 (1985), pp. 617–633.
- [4] A. BAYLISS, C. I. GOLDSTEIN, AND E. TURKEL, *An iterative method for the Helmholtz equation*, J. Comput. Phys., 49 (1983), pp. 443–457.
- [5] A. BAYLISS, C. I. GOLDSTEIN, AND E. TURKEL, *Preconditioned conjugate gradient methods for the Helmholtz equation*, in Elliptic Problem Solvers II, G. Birkhoff and A. Schoenstadt, eds., Academic Press, Orlando, FL, 1984, pp. 233–243.
- [6] A. BAYLISS, C. I. GOLDSTEIN, AND E. TURKEL, *The numerical solution of the Helmholtz equation for wave propagation problems in underwater acoustics*, Comput. Math. Appl., 11 (1985), pp. 655–665.
- [7] A. BEHIE AND P. FORSYTH, *Comparison of fast iterative methods for symmetric systems*, IMA J. Numer. Anal., 3 (1983), pp. 41–63.
- [8] F. A. BORNEMANN AND P. DEUFLHARD, *The cascadic multigrid method for elliptic problems*, Numer. Math., 75 (1996), pp. 135–152.
- [9] D. BRAESS, *On the combination of the multigrid method and conjugate gradients*, in Multigrid Methods II, W. Hackbusch and U. Trottenberg, eds., Lecture Notes in Math. 1228, Springer-Verlag, Berlin, 1986, pp. 52–64.
- [10] J. H. BRAMBLE, J. E. PASCIAK, AND J. XU, *The analysis of multigrid algorithms for nonsymmetric and indefinite elliptic problems*, Math. Comp., 51 (1988), pp. 389–414.
- [11] A. BRANDT AND I. LIVSHITS, *Wave-ray multigrid method for standing wave equations*, Electron. Trans. Numer. Anal., 6 (1997), pp. 162–181.
- [12] A. BRANDT AND S. TA'ASAN, *Multigrid methods for nearly singular and slightly indefinite problems*, in Multigrid Methods II, W. Hackbusch and U. Trottenberg, eds., Lecture Notes in Math. 1228, Springer-Verlag, Berlin, 1986, pp. 99–121.
- [13] W. L. BRIGGS, *A Multigrid Tutorial*, SIAM, Philadelphia, 1987.
- [14] B. ENQUIST AND A. MAJDA, *Absorbing boundary conditions for the numerical simulation of waves*, Math. Comp., 31 (1977), pp. 629–651.
- [15] 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.
- [16] D. GIVOLI, *Non-reflecting boundary conditions*, J. Comput. Phys., 94 (1991), pp. 1–29.
- [17] C. I. GOLDSTEIN, *Multigrid preconditioners applied to the iterative solution of singularly perturbed elliptic boundary value problems and scattering problems*, in Innovative Numerical Methods in Engineering, Proceedings of the 4th International Symposium, Georgia Institute of Technology, Atlanta, 1986, pp. 97–102.
- [18] W. HACKBUSCH, *Multi-Grid Methods and Applications*, Springer-Verlag, Berlin, 1985.
- [19] I. HARARI AND T. J. R. HUGHES, *Finite element method for the Helmholtz equation in an exterior domain*, Comput. Methods Appl. Mech. Engrg., 87 (1991), pp. 59–96.
- [20] F. IHLENBURG AND I. BABUŠKA, *Finite element solution to the Helmholtz equation with high wave numbers*, Comput. Math. Appl., 30 (1995), pp. 9–37.
- [21] J. B. KELLER AND D. GIVOLI, *Exact non-reflecting boundary conditions*, J. Comput. Phys., 82 (1989), pp. 172–192.
- [22] R. KETTLER, *Analysis and comparison of relaxation schemes in robust multigrid and conjugate gradient methods*, in Multigrid Methods, W. Hackbusch and U. Trottenberg, eds., Lecture Notes in Math. 960, Springer-Verlag, Berlin, 1982, pp. 502–534.
- [23] B. LEE, T. A. MANTEUFFEL, S. F. MCCORMICK, AND J. RUGE, *First-order system least-squares (FOSLS) for the Helmholtz equation*, SIAM J. Sci. Comput., 21 (2000), pp. 1927–1949.
- [24] S. F. MCCORMICK, ED., *Multigrid Methods*, Frontiers Appl. Math. 3, SIAM, Philadelphia, 1987.
- [25] C. W. OOSTERLEE AND T. WASHIO, *An evaluation of parallel multigrid as a solver and a preconditioner for singularly perturbed problems*, SIAM J. Sci. Comput., 19 (1998), pp. 87–110.
- [26] A. RAMAGE, *A multigrid preconditioner for stabilised discretisations of advection-diffusion problems*, J. Comput. Appl. Math., 110 (1999), pp. 187–203.
- [27] Y. SAAD, *Iterative Methods for Sparse Linear Systems*, PWS, Boston, 1995.
- [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] A. SCHATZ, *An observation concerning Ritz-Galerkin methods with indefinite bilinear forms*, Math. Comp., 28 (1974), pp. 959–962.
- [30] V. V. SHAIUROV, *Some estimates of the rate of convergence for the cascadic conjugate-gradient method*, J. Comput. Math. Appl., 31 (1996), pp. 161–171.
- [31] Y. SHAPIRA, *Multigrid techniques for highly indefinite equations*, in Proceedings of the 1995 Copper Mountain Conference on Multigrid Methods, Copper Mountain, CO, 1995; also available online from <http://www.c3.lanl.gov/~yairs/>.
- [32] G. STARKE, *On the Performance of Krylov Subspace Iterations as Smoothers in Multigrid Methods*, manuscript, 1995.
- [33] 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.
- [34] J. XU, *A new class of iterative methods for nonselfadjoint or indefinite problems*, SIAM J. Numer. Anal., 29 (1992), pp. 303–319.
- [35] H. YSERENTANT, *On the multi-level splitting of finite element spaces for indefinite elliptic boundary value problems*, SIAM J. Numer. Anal., 23 (1986), pp. 581–595.
- [36] H. YSERENTANT, *Preconditioning indefinite discretization matrices*, Numer. Math., 54 (1988), pp. 719–734.