RELAXATION SCHEMES FOR CHEBYSHEV SPECTRAL MULTIGRID METHODS*

Yimin Kang and Scott R. Fulton
Department of Mathematics and Computer Science
Clarkson University
Potsdam, NY

SUMMARY

Two relaxation schemes for Chebyshev spectral multigrid methods are presented for elliptic equations with Dirichlet boundary conditions. The first scheme is a pointwise-preconditioned Richardson relaxation scheme and the second is a line relaxation scheme. The line relaxation scheme provides an efficient and relatively simple approach for solving two-dimensional spectral equations. Numerical examples and comparisons with other methods are given.

INTRODUCTION

For limited-area problems with general (non-periodic) boundary conditions, Chebyshev spectral methods give exponential convergence for smooth solutions. However, except in some very simple cases (e.g., one-dimensional constant-coefficient problems), Chebyshev approximations usually lead to full linear systems which cannot be solved efficiently by direct methods, and iterative methods must be used. Unfortunately, designing efficient iterative methods for discrete spectral equations has proven difficult, especially for problems with non-constant coefficients (ref. 1). Perhaps the most promising technique to date for solving spectral discretizations of elliptic problems is the spectral multigrid method (ref. 2, 3). However, the best relaxation schemes known today are complicated to apply. In this paper we introduce two simpler relaxation schemes and investigate their performance.

As prototype problems we consider one- and two-dimensional elliptic equations with Dirichlet boundary conditions on simple geometric domains. In one dimension we consider

$$\begin{aligned} -u''(x) &= f(x), \qquad |x| < 1, \\ u(+1) &= a, \qquad u(-1) = b. \end{aligned} \tag{1}$$

The two-dimensional prototype problem is

$$\begin{aligned} -\triangle u(x,y) &= f(x,y), \qquad |x|, |y| < 1, \\ u(x,y) &= g(x,y), \qquad |x| = 1, |y| = 1. \end{aligned} \tag{2}$$

We discretize these problems by Chebyshev collocation. For example, for the two-dimensional problem (2), the solution $u(x,y)$ is approximated by a set of discrete values $\bar{u}_{j,k}$ on the Chebyshev

*Work supported by the National Science Foundation under Grant No. ATM-9118966.

grid $\{(\bar{x}_j, \bar{y}_k) = (\cos(j\pi/N_x), \cos(k\pi/N_y)) \,|\, 0 \le j \le N_x, 0 \le k \le N_y\}$, with the requirement that problem (2) be satisfied on this grid, i.e.,

$$\begin{aligned} -(\bar{u}^{(xx)}_{j,k} + \bar{u}^{(yy)}_{j,k}) &= f(\bar{x}_j, \bar{y}_k), \qquad 1 < j < N_x, 0 < k < N_y \\ \bar{u}_{j,k} &= g(\bar{x}_j, \bar{y}_k), \qquad j = 0, j = N_x, k = 0, k = N_y \end{aligned} \tag{3}$$

where $\bar{u}^{(xx)}_{j,k}$ and $\bar{u}^{(yy)}_{j,k}$ are values of the second-order derivatives of the Chebyshev approximation $\sum_{m=0}^{N_x} \sum_{n=0}^{N_y} \hat{u}_{mn} T_m(x) T_n(y)$ to $u(x,y)$ on the Chebyshev grid. For simplicity, we will assume here that $N_x = N_y = N$; however, the codes described in this paper do not require this.

The discrete problem (3) can be expressed in form of a linear system

$$A\bar{U} = \bar{F} \tag{4}$$

Unfortunately, the matrix A, formulated by Chebyshev collocation approximations, is full and non-symmetric. For two-dimensional problems, direct methods (like Gaussian elimination) would require $O(N^6)$ operations for factorization and $O(N^4)$ for the subsequent solution, which is far too much work to be practical. Thus, iterative methods must be used.

THE POINTWISE PRECONDITIONED RICHARDSON RELAXATION SCHEME

The most efficient method available today for solving (4) and its generalizations to other elliptic problems is the spectral multigrid method of Zang et al. (ref. 2, 3), which employs finite-difference preconditioned Richardson iteration as the relaxation scheme in a multigrid context. Preconditioned Richardson relaxation for (4) takes the form

$$V \longleftarrow V + \omega H(F - AV), \tag{5}$$

where V is the current approximation to $\bar{U}$, ω is a relaxation parameter, and H is the preconditioner. The criteria for choosing a preconditioner H are:

- H should give fast multigrid convergence,
- H should be easy and cheap to generate or apply.

The finite-difference preconditioning of Zang et al. (ref. 2, 3) gives fast convergence, but applying it requires solving (or nearly solving) a finite-difference discretization on the nonuniform Chebyshev grid. This procedure is complicated and expensive. Are there alternatives which are simpler and still effective? Achi Brandt (personal communication, 1983) has suggested that pointwise preconditioning based on the (variable) Chebyshev mesh spacing might work well. In this section, we investigate the performance of this simple preconditioner when applied to the problem (4).

The One-Dimensional Case

Formulation

As an analogue of the Gauss-Seidel relaxation for a finite-difference method, the pointwise preconditioning for the Chebyshev discretization takes the form

$$v_j \longleftarrow v_j + \omega \frac{h_j^2}{2} r_j, \tag{6}$$

where $h_j = (\bar{x}_{j-1} - \bar{x}_{j+1})/2$ is the effective grid size at the point $\bar{x}_j$, r_j is the the residual $R = F - AV$ at $\bar{x}_j$, and ω is a relaxation parameter to be chosen to accelerate the convergence. Note that (6) is equivalent to choosing the preconditioning matrix H in (5) as a diagonal matrix

$$H = \text{diag}\left(1, \frac{h_1^2}{2}, \ldots, \frac{h_{N-1}^2}{2}, 1\right). \tag{7}$$

Analysis

The evolution of the error $E = V - \bar{U}$ in the Richardson relaxation (5) is described by

$$E \longleftarrow (I - \omega HA)E. \tag{8}$$

Therefore, the convergence factor for (5) on a single grid is

$$\sigma_{SG} = \rho(I - \omega HA),$$

where ρ denotes the spectral radius. Likewise, the multigrid smoothing factor for (5), when used as a smoother in a multigrid method (e.g., ref. 4), is

$$\bar{\mu} = \rho(G(I - \omega HA)), \tag{9}$$

where G represents the perfect coarse-grid correction, i.e., set all low modes of the error to zero.

For the simple preconditioning (7), our numerical computations show that the eigenvalues of the matrix HA are all positive real numbers. The maximum eigenvalue is $\lambda_{\max} \approx 5.0$, the middle is $\lambda_{\text{mid}} \approx 1.5$, and the minimum is $\lambda_{\min} \approx O(N^{-2})$. The formulas of Zang et. al. (ref. 2, 3) then give a good approximation to the optimal ω and $\bar{\mu}$, namely,

$$\omega \approx \frac{2}{\lambda_{\max} + \lambda_{\text{mid}}} \approx 0.325, \qquad \bar{\mu} \approx \frac{\lambda_{\max} - \lambda_{\text{mid}}}{\lambda_{\max} + \lambda_{\text{mid}}} \approx 0.6. \tag{10}$$

Indeed, computing the smoothing factor directly from (9) using $\omega = 0.325$, we find that $\bar{\mu} \leq 0.6$ for all $N \leq 512$.

To take into account the effects of grid transfers (omitted in the smoothing analysis above), we use the following two-grid analysis. The evolution of the error E in one two-grid $V(n_1, n_2)$-cycle (where n_1 and n_2 specify the number of relaxation sweeps before and after the coarse-grid

correction, respectively) is described by the matrix

$$T = (I - \omega H A_f)^{n_2}(I - P A_c^{-1} R A)(I - \omega H A_f)^{n_1}. \tag{11}$$

Here, R represents the fine-to-coarse grid transfer (we use injection), P represents the coarse-to-fine grid transfer (we use Chebyshev interpolation), and A_f and A_c represent the discrete operator matrix in (4) on the fine and coarse grids, respectively. Note that (11) assumes that the coarse-grid problem is solved exactly.

We computed the two-grid convergence factor $\sigma_{TG} = \rho(T)$ for $N \leq 512$ using different values of ω, and the numerical results show that $\omega = 0.325$ again gives the optimal convergence factor (or very close to it). Using that constant value, we find that the smoothing factor per sweep $\mu_s = (\sigma_{TG})^{1/(n_1+n_2)}$ satisfies

$$0.5 \leq \mu_s \leq 0.6$$

for all $N \leq 512$. A similar analysis for the one-dimensional Helmholtz problem

$$\lambda\, u(x) - u''(x) = f(x) \tag{12}$$

shows that with various choices of λ and boundary conditions (Dirichlet, Neumann and mixed), an appropriate pointwise preconditioner also yields the smoothing factor per sweep $\mu_s \leq 0.6$.

We have developed FORTRAN-77 routines to implement the Chebyshev multigrid method using the pointwise preconditioner as described above. The code has been used to solve the problem (12) with various choices of $u(x)$, λ, and boundary conditions. The observed convergence factor per sweep μ_s is smaller than 0.60 for all cases tested, in agreement with the analysis presented above.

The Two-Dimensional Case

Formulation

We note that Gauss-Seidel relaxation for the second-order centered finite difference approximation to (2) can be written as

$$u_{j,k} \longleftarrow u_{j,k} + \frac{h^2}{4} r_{j,k},$$

where $r_{j,k}$ is the finite-difference residual. A natural analogue for the Chebyshev collocation discretization (3) is

$$\bar{v}_{j,k} \longleftarrow \bar{v}_{j,k} + \omega \left(\frac{1}{2/h_j^2 + 2/h_k^2} \right) \bar{r}_{j,k}, \tag{13}$$

where h_j and h_k are the grid sizes at the point $(\bar{x}_j, \bar{y}_k)$, $\bar{r}_{j,k} := \bar{f}_{j,k} - [-(\bar{v}_{j,k}^{(xx)} + \bar{v}_{j,k}^{(yy)})]$ is the residual of Chebyshev discretization, and ω is a relaxation parameter to be chosen to accelerate the convergence. Clearly, the iteration (13) is a special case of the Richardson iteration (5), with

a diagonal preconditioner H with diagonal entries $(H)_{jk,jk} = (2/h_j^2 + 2/h_k^2)^{-1}$. This preconditioner is easy and fast to apply. Does it gives a fast convergence? Unfortunately, the following analysis shows that the answer is no.

Analysis

Computational results indicate that the eigenvalues of the matrix HA are all positive real numbers. Again, good approximations to the optimal ω and $\bar{\mu}$ can be obtained by

$$\omega \approx \frac{2}{\lambda_{\max} + \lambda_{\text{qua}}}, \qquad \bar{\mu} \approx \frac{\lambda_{\max} - \lambda_{\text{qua}}}{\lambda_{\max} + \lambda_{\text{qua}}}, \tag{14}$$

where $\lambda_{\max}$ is the maximum eigenvalue and λ_{qua} is the quarter eigenvalue (ref. 1). More precise values of the optimal ω and $\bar{\mu}$, can be obtained by actually computing the spectral radius $\rho(G(I - \omega HA))$ for different choices of ω and comparing the results. For $N \leq 32$, the eigenvalues $\lambda_{\max}$ and λ_{qua}, ω and $\bar{\mu}$ computed by (14) and the optimal ω and $\bar{\mu}$ are listed in Table I. Since $\bar{\mu}$ is large and increases with N, these results suggest that the pointwise preconditioner (13) will not be a good multigrid smoother.

Table I also lists the two-grid smoothing factors per sweep $\mu_s = (\rho(T))^{1/(n_1+n_2)}$ computed from the matrices in (11) for $N \leq 32$ using $\omega = 0.36$. These results again show that the pointwise preconditioning (13) does not give fast convergence.

We have implemented the pointwise preconditioning (13) in a multigrid solver written in Fortran 77. Computational results from a number of test cases confirm the above analysis: we conclude that the pointwise preconditioning does not give fast convergence.

Table I. Multigrid Analysis of Two-Dimensional Pointwise Preconditioning

	Eigenvalues of HA		By (14)		By computation		
N	$\lambda_{\max}$	λ_{qua}	ω	$\bar{\mu}$	ω_{opt}	$\bar{\mu}$	μ_s
4	3.00	1.83	0.41	0.24	0.35	0.28	0.51
8	4.10	1.26	0.37	0.53	0.35	0.52	0.68
16	4.57	0.95	0.36	0.66	0.36	0.75	0.80
32	4.76	0.78	0.36	0.72	0.36	0.82	0.88

THE LINE RELAXATION METHOD

The poor performance of pointwise preconditioning in two dimensions can be understood in terms of the anisotropy introduced by the nonuniform Chebyshev collocation grid. Since the mesh spacing varies with x and y, at any given point (x, y) the coupling in the discrete operator in (3) may be stronger in x or in y. In finite-difference multigrid methods, point relaxation performs poorly in such anisotropic cases, and the cure is to use alternating direction line relaxation. Thus, it is reasonable to try an analogous approach for the Chebyshev discretization.

Formulation

To formulate the line relaxation method, we express the discrete problem (3) in the matrix form

$$(\mathcal{H} + \mathcal{V})\bar{U} = \bar{F}, \tag{15}$$

where $\mathcal{H}$ and $\mathcal{V}$ correspond to the horizontal part $(-\partial^2/\partial x^2)$ and vertical part $(-\partial^2/\partial y^2)$ of the Laplacian operator, respectively. Starting from an approximation V^{old} to the solution $\bar{U}$, one sweep of (alternating direction) line relaxation based on (15) consists of the following two parts:

1. Sweep along the x-direction. On each grid line parallel to x-axis, use the values of V^{old} except those on the current line, and solve for values on the current line by solving (15). This can be expressed in the matrix form as

$$(\mathcal{H} + \mathcal{V}_d)V^{\text{mid}} = \bar{F} - \mathcal{V}_o V^{\text{old}}, \tag{16}$$

   where $\mathcal{V}_d$ and $\mathcal{V}_o$ denote the diagonal and off-diagonal parts of the matrix $\mathcal{V}$, respectively. Note that the entries of $\mathcal{V}_d$ are known (ref. 1) and $\mathcal{V}_d$ is a constant on each grid line parallel to the x-axis. Thus, the system (16) can be decoupled into $(N-1)$ one-dimensional discrete problems, each of which is a Chebyshev collocation approximation to a Helmholtz equation on an interior grid line parallel to x-axis; the x-directional sweep consists of solving these equations.

2. Sweep along the y-direction. The y-direction sweep is basically the same as the x-direction sweep except that we now work on grid lines that are parallel to y-axis and use values of V^{mid} instead of V^{old}. The equation we need to solve is

$$(\mathcal{H}_d + \mathcal{V})V^{\text{new}} = \bar{F} - \mathcal{H}_o V^{\text{mid}}, \tag{17}$$

   where $\mathcal{H}_d$ and $\mathcal{H}_o$ are the diagonal and off-diagonal parts of $\mathcal{H}$. As in the x-direction sweep, the two-dimensional problem (17) is solved by solving $(N-1)$ one-dimensional Helmholtz equations.

It turns out that as it stands, the line relaxation (16)–(17) is not a good multigrid smoother; however, this can be fixed as follows. Let $C^{\text{mid}} = V^{\text{mid}} - V^{\text{old}}$ and $C^{\text{new}} = V^{\text{new}} - V^{\text{mid}}$ denote the corrections for V^{old} and V^{mid}, and $R^{\text{old}} = \bar{F} - AV^{\text{old}}$ and $R^{\text{mid}} = \bar{F} - AV^{\text{mid}}$ denote the residuals of V^{old} and V^{mid}, respectively. Rewriting equations (16) and (17) as correction equations and

introducing a relaxation parameter ω (to be determined by analysis to accelerate the convergence), we obtain

$$(\mathcal{H} + \mathcal{V}_d)C^{\text{mid}} = \omega R^{\text{old}}, \qquad (\mathcal{H}_d + \mathcal{V})C^{\text{new}} = \omega R^{\text{mid}} \tag{18}$$

We refer to (18) as the collocation version of the line relaxation method.

It is not practical to implement the collocation version because there are no fast solvers available for the collocation approximations, even for one-dimensional problems. However, in the multigrid context, a relaxation scheme functions as a smoother rather than a solver: instead of solving each problem exactly, we only need to smooth out the error, i.e., reduce high modes in the error. Therefore, it is reasonable to replace the one-dimensional problems in (18) by approximate versions which can be solved efficiently. We consider two alternatives as follows.

In the first, we replace the collocation discretizations of the one-dimensional Helmholtz equations in (18) by tau discretizations. Tau approximations have the same exponential convergence as collocation method, but can be solved directly in $O(N \log N)$ operations. This leads to the tau version of the line relaxation method, and the total work of one x or y-direction sweep is $O(N^2 \log N)$. As we will see below, this tau version turns out to be an efficient multigrid smoother.

In the second, we replace the collocation discretizations of the one-dimensional Helmholtz equations in (18) by finite-difference discretizations. This leads to the finite-difference version of the line relaxation method, which has two obvious advantages over the tau version. First, it is faster because it eliminates the transforms required in tau version, thus reducing the operation count for solving each one-dimensional problem from $O(N \log N)$ to $O(N)$. Second, it can be extended to solve more generalized problems, e.g., problems with variable coefficients. As we will see below, this finite-difference version also turns out to be an efficient multigrid smoother, even in the case of variable coefficients.

Analysis

As in the case of the pointwise preconditioned Richardson relaxation, we can analyze the performance of the line relaxation methods described above by computing the eigenvalues of the corresponding interation matrices. Because the tau version cannot be expressed in matrix form like (18), we will only do the analysis for the collocation and finite-difference versions. Note that the tau and collocation versions are nearly the same, so the analysis for collocation version should give a good prediction for the performance of the tau version. In this section, we will give details of the analysis for finite-difference version and only list results for collocation version.

Smoothing Analysis

For the finite-difference version of the line relaxation iteration, the error evolution is described by

$$E^{\text{mid}} \longleftarrow [I - \omega(\mathcal{H}^{fd} + \mathcal{V}_d)^{-1}(\mathcal{H} + \mathcal{V})]E^{\text{old}}, \tag{19}$$

$$E^{\text{new}} \longleftarrow [I - \omega(\mathcal{H}_d + \mathcal{V}^{fd})^{-1}(\mathcal{H} + \mathcal{V})]E^{\text{mid}}. \tag{20}$$

where $\mathcal{H}^{fd}$ and $\mathcal{V}^{fd}$ are the finite-difference analogues of the collocation discretization matrices $\mathcal{H}$ and $\mathcal{V}$, respectively. Therefore, the error evolution matrix for one relaxation is

$$S = [I - \omega(\mathcal{H}_d + \mathcal{V}^{fd})^{-1}(\mathcal{H} + \mathcal{V})][I - \omega(\mathcal{H}^{fd} + \mathcal{V}_d)^{-1}(\mathcal{H} + \mathcal{V})]. \tag{21}$$

The matrices $S_{\mathcal{H}} = (\mathcal{H}^{fd} + \mathcal{V}_d)^{-1}(\mathcal{H} + \mathcal{V})$ and $S_{\mathcal{V}} = (\mathcal{H}_d + \mathcal{V}^{fd})^{-1}(\mathcal{H} + \mathcal{V})$ have the same eigenvalues (since x and y can be interchanged in the Laplacian operator), so we can focus on just the x-direction sweep (19). The eigenvalues of $S_{\mathcal{H}}$ are all positive real numbers, so we can use formulas (14) to obtain approximate values of ω and $\bar{\mu}$ (squaring $\bar{\mu}$ to represent the effect of both the x and y sweeps). These values are listed in Table II for $N \leq 32$, along with the optimal relaxation parameter ω and corresponding multigrid smoothing factor $\bar{\mu} = \rho(GS)$ computed directly. These results suggest that for large values of truncation number N, $\omega_{opt} \approx 0.6$ and $\bar{\mu} \leq 0.5$, independent of the grid size. Corresponding results for the collocation version are listed in Table III.

Multigrid Analysis

For a multigrid $V(n_1, n_2)$-cycle, if we use zeros as initial guesses on all coarse grids (which is a natural choice because the coarse-grid solution is a correction to the solution on the next finer grid), then we can write out the error evolution matrix explicitly as

$$M = S^{n_2}\,[I - P\,G\,R\,(\mathcal{H} + \mathcal{V})]\,S^{n_1}. \tag{22}$$

This represents a procedure of n_1 pre-relaxations (S^{n_1}) followed by a coarse-grid-correction ($I - P\,G\,R\,(\mathcal{H} + \mathcal{V})$) and then n_2 post-relaxations (S^{n_2}). The matrix S is the error evolution matrix of one relaxation on the finest grid defined in (21). The central part $I - PGR(\mathcal{H} + \mathcal{V})$ represents the coarse-grid-correction, where R represents the fine-to-coarse grid transfer (we use injection) and P represents the coarse-to-fine grid transfer (we use Chebyshev interpolation). The matrix G is defined on the next coarser grid as follows: on the coarsest grid, $G = (\mathcal{H} + \mathcal{V})^{-1}$ (which means the coarsest grid problem is solved exactly); otherwise,

$$G = [I - M] * (\mathcal{H} + \mathcal{V})^{-1}, \tag{23}$$

which represents a multigrid solution procedure on that grid. Note that (23) is actually a recursive definition, since the matrix M in (23) includes another matrix G on the next coarser grid.

Tables II and III also list computed values of smoothing factor per sweep $\mu_s = (\rho(M))^{1/(n_1+n_2)}$ for the case $\omega = 0.6$, $n_1 = 2$, and $n_2 = 1$. These results suggest that the smoothing factor of the line relaxation method is less than 0.5, independent of the grid size. Note that while we could also use Chebyshev restriction instead of injection for the fine-to-coarse grid transfer R, our numerical experience shows very little difference between these two choices.

Table II. Analysis of the Finite-Difference Version

	Eigenvalues of $S_{\mathcal{H}}$		By (14)		By computation		
N	$\lambda_{\max}$	λ_{qua}	ω	$\bar{\mu}$	ω_{opt}	$\bar{\mu}$	μ_s
4	1.995	1.000	0.669	0.110	0.58	0.110	0.181
8	2.513	1.000	0.569	0.186	0.60	0.168	0.293
16	2.780	0.995	0.530	0.224	0.60	0.271	0.364
32	2.898	0.815	0.539	0.315	0.60	0.366	0.421

Table III. Analysis of the Collocation Version

	Eigenvalues of $S_{\mathcal{H}}$		By (14)		By computation		
N	$\lambda_{\max}$	λ_{qua}	ω	$\bar{\mu}$	ω_{opt}	$\bar{\mu}$	μ_s
4	1.651	1.000	0.754	0.060	0.68	0.120	0.302
8	2.322	0.922	0.616	0.186	0.60	0.216	0.328
16	2.701	0.810	0.570	0.290	0.58	0.326	0.380
32	2.869	0.700	0.560	0.370	0.60	0.410	0.428

Computational Results

We have implemented the tau and finite-difference versions of the line relaxation scheme described above in a Chebyshev collocation multigrid solver for the two-dimensional Helmholtz problem

$$\begin{aligned} \lambda u(x,y) - \triangle u(x,y) &= f(x,y), \qquad |x|,|y| < 1, \\ u(x,y) &= g(x,y), \qquad |x| = 1, |y| = 1, \end{aligned}$$

with various choices of f, g, and λ. For both versions, the observed convergence factor per sweep is less than 0.5 for all cases tested, in agreement with the analysis above. The finite-difference version turns out to have slightly better convergence factors than the tau version, but the difference is minor.

Comparisons with Other Methods

In this section we compare the line relaxation spectral multigrid method developed above to two other methods for solving the two-dimensional prototype problem (2). The first is a conventional finite-difference multigrid method; the second is a matrix diagonalization technique. We do not compare with the method of Zang et. al. (ref. 3) since the details presented in that paper were not enough to allow programing the method. All computations are done on a SUN SPARCstation2 using double precision; the machine round-off error is about 2.22×10^{-16}.

Conventional Finite-Difference Multigrid Method

The finite-difference discretization is the usual second-order five-point scheme on a uniform grid. The finite-difference multigrid method uses Gauss-Seidel (Red-Black) iteration as a relaxation scheme, the fine-to-coarse grid transfer is half-injection, the coarse-to-fine grid transfer is bilinear interpolation, and the multigrid V-cycle algorithm is used.

According to computations, the average execution time of one $V(2,1)$-cycle of the finite-difference multigrid method is approximately $(0.56 \times 10^{-4})\, N^2$ seconds, and $(0.21 \times 10^{-3})\, N^2 \log_2 N$ seconds for line relaxation spectral multigrid method. Therefore, for the same grid sizes, one $V(2,1)$-cycle of the finite-difference multigrid method is approximately $3.75 \log_2 N$ times faster than the line relaxation spectral multigrid method.

However, because spectral methods have exponential convergence and finite-difference methods only have polynomial convergence, when high accuracy is required, finite-difference multigrid methods must use much bigger grid sizes than spectral methods. The result is that the line relaxation spectral multigrid method is faster than finite-difference when high accuracy is required. As a specific example, consider the prototype problem (2) with true solution $u(x,y) = e^{2x+y}\cos(\pi(x+4y+0.25))$. The relation between accuracy and execution time required to achieve that accuracy is plotted in Figure 1 for both methods. We can see that when low accuracy is required, the finite-difference multigrid method is much faster than the line relaxation spectral multigrid method, but the situation is reversed when high accuracy is required. The crossover point for this problem is at an accuracy of about one percent error. The same conclusion would hold for finite-difference methods of higher (fixed) orders, although the crossover point would shift. Variable-order finite-difference methods could be expected to perform more like the spectral method, at a cost of considerable complexity.

Matrix Diagonalization Technique

The matrix diagonalization technique is introduced in (ref. 5) as a direct solver for the Chebyshev spectral approximation to the Poisson equation with Dirichlet boundary conditions. This technique requires a preprocessing step, which involves computing the eigenvalues and eigenvectors of a one-dimensional operator matrix ($O(N^3)$ operations), and a solution step, which involves one-dimensional matrix multiplications ($O(N^3)$ operations).

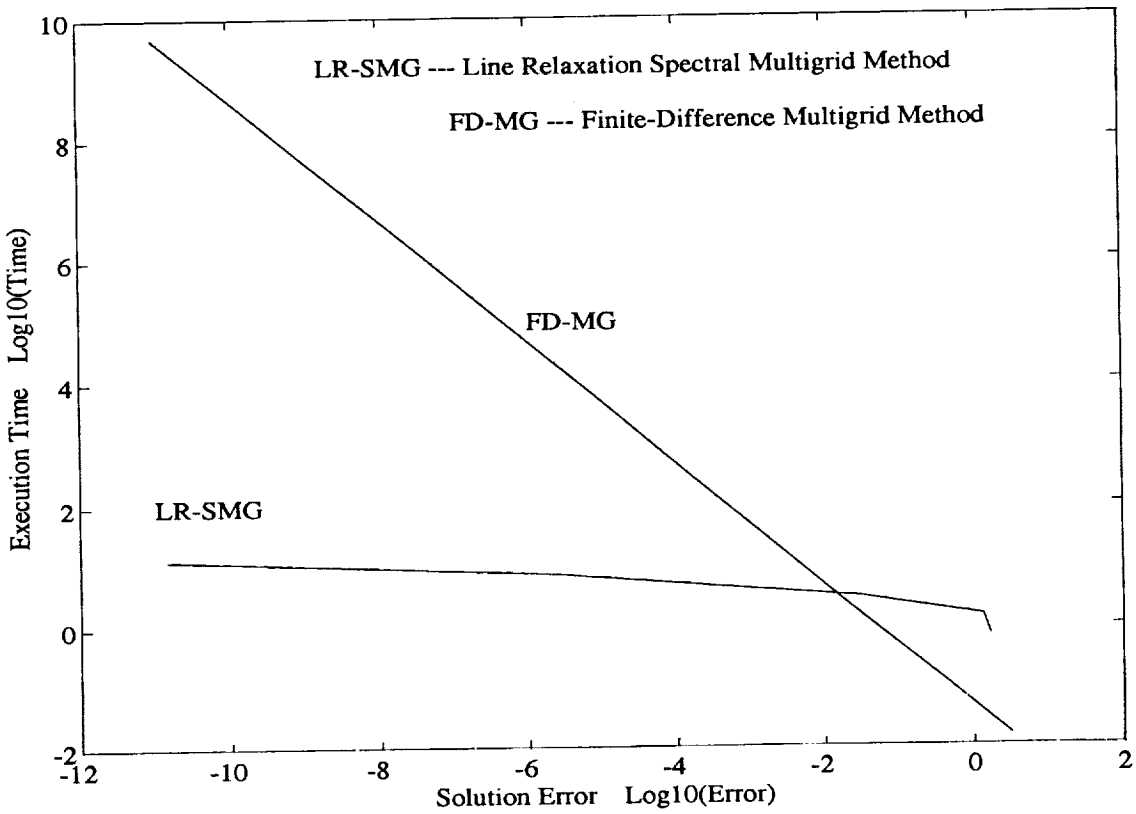

Figure 1. Execution time: LR-SMG vs FD-MG

To compare execution times, we note that the line relaxation spectral multigrid method usually takes approximately 10 V-cycles to solve to the level of machine precision. Thus, Figure 2 compares the execution time of 10 V-cycles of the line relaxation spectral multigrid method with the execution time of solving the same problem directly by the matrix diagonalization method (including the preprocessing step). These results show that the matrix diagonalization method is quite fast for small grid sizes, but as the grid size grows, it becomes slower than the line relaxation spectral multigrid method. This is because the line relaxation spectral multigrid method is an $O(N^2 \log N)$ method, while the matrix diagonalization method requires $O(N^3)$ operations (even without the preprocessing step).

The matrix diagonalization technique is very efficient for problems with constant coefficients, especially when repeated solutions are required. However, this technique can only handle problems with constant coefficients. As shown below, the line relaxation spectral multigrid method is able to solve problems with non-constant coefficients.

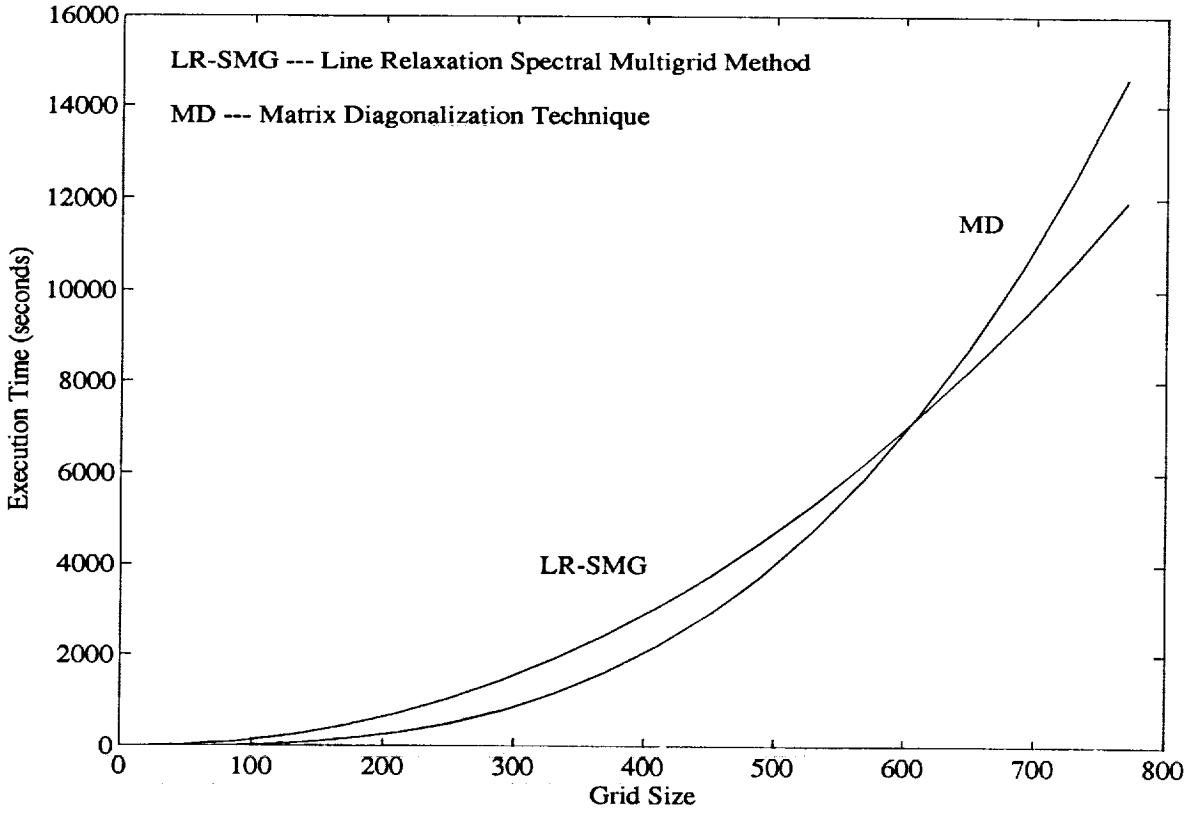

Figure 2. Execution time: LR-SMG vs MD

Extension to Problems With Variable Coefficients

As a test problem with variable coefficients we consider

$$\begin{aligned} -\frac{\partial}{\partial x}\left(a(x,y)\frac{\partial}{\partial x}u(x,y)\right) - \frac{\partial}{\partial y}\left(b(x,y)\frac{\partial}{\partial y}u(x,y)\right) &= f(x,y), \qquad |x|,|y| < 1, \\ u(x,y) &= g(x,y), \qquad |x| = 1, |y| = 1, \end{aligned} \tag{24}$$

where the coefficient functions and the true solution are

$$a(x,y) = b(x,y) = 1 + \varepsilon e^{\cos(\beta\pi(x+y))}, \tag{25}$$

$$u(x,y) = \sin(\alpha\pi x + \frac{\pi}{4})\sin(\alpha\pi y + \frac{\pi}{4}). \tag{26}$$

The parameter ε measures how far the coefficients are away from the constant 1, β measures the oscillation of the coefficients, and α measures the oscillation of the solution.

Implementation of the Line Relaxation Spectral Multigrid Method

The implementation of the finite-difference version of the line relaxation method is basically the same as for the constant coefficient case except for the following:

1. On each grid line, the one-dimensional problem is not a Helmholtz equation anymore. For example, on a gird line $y = \bar{y}_k$ which is parallel to x-axis, we now solve a problem like

$$-\frac{\partial}{\partial x}\left(a(x,\bar{y}_k)\frac{\partial}{\partial x}v(x,\bar{y}_k)\right) - \mathcal{V}_d(x,\bar{y}_k)v(x,\bar{y}_k) = h(x,\bar{y}_k) \tag{27}$$

   by using a second-order finite-difference approximation on the Chebyshev grid.

2. To compute values of $\mathcal{V}_d(\bar{x}_j,\bar{y}_k)$, note that the interior equation in (24) can be rewritten as

$$-a\frac{\partial^2 u}{\partial x^2} - \frac{\partial a}{\partial x}\frac{\partial u}{\partial x} - b\frac{\partial^2 u}{\partial y^2} - \frac{\partial b}{\partial y}\frac{\partial u}{\partial y} = f, \quad |x|,|y| < 1 \tag{28}$$

   and the Chebyshev collocation approximation to (28) can be written as

$$\{-\mathcal{A}\mathcal{D}_{xx} - \mathcal{A}_x\mathcal{D}_x\}\,\bar{U} - \{-\mathcal{B}\mathcal{D}_{yy} - \mathcal{B}_y\mathcal{D}_y\}\,\bar{U} = \bar{F}, \tag{29}$$

   where $\mathcal{A}$ and $\mathcal{B}$ are diagonal matrices containing the values of the coefficients $a(\bar{x}_j,\bar{y}_k)$ and $b(\bar{x}_j,\bar{y}_k)$, $\mathcal{A}_x$ and $\mathcal{B}_x$ are diagonal matrices containing the values of the derivatives $\frac{\partial}{\partial x}a(\bar{x}_j,\bar{y}_k)$ and $\frac{\partial}{\partial y}b(\bar{x}_j,\bar{y}_k)$ (which can be computed from values $a(\bar{x}_j,\bar{y}_k)$ and $b(\bar{x}_j,\bar{y}_k)$), and $\mathcal{D}_x$, $\mathcal{D}_{xx}$, $\mathcal{D}_y$, and $\mathcal{D}_{yy}$ are Chebyshev differentiation matrices. Therefore, $\mathcal{H} = -\mathcal{A}\mathcal{D}_{xx} - \mathcal{A}_x\mathcal{D}_x$ and $\mathcal{V} = -\mathcal{B}\mathcal{D}_{yy} - \mathcal{B}_y\mathcal{D}_y$; generating the diagonal entries of $\mathcal{H}$ and $\mathcal{V}$ is straightforward.

3. On coarse grids, we need to use so-called "filtered" coefficients $a(x,y)$ and $b(x,y)$ to formulate the coarse grid problems; i.e., the coefficients $a(x,y)$ and $b(x,y)$ are evaluated on the finest grid and then transferred to the coarser grids by Chebyshev restriction (ref. 3).

Computational Results

We have run the line relaxation spectral multigrid method for different values of parameters ε, α and β. For $\alpha = 1.0$ and $N_x = N_y = 32$, the smoothing factor is graphed in Figure 3 as a function of ε and α. Here we have chosen to measure the smoothing by the "smoothing factor per work unit" defined by $\mu_w = (r_2/r_1)^{\tau_0/\tau}$, where r_1 and r_2 are residual norms before and after one multigrid V-cycle, τ is the execution time of one cycle and τ_0 is the execution time of one relaxation. These results show that for a wide range of ε and β, the method converges relatively quickly.

In (ref. 3) the same test problem (24) was solved using the Richardson relaxation (5) using two-dimensional finite-difference preconditioning; incomplete LU decomposition was used to approximately solve the finite difference approximation on the Chebyshev grid. With only limited details of the formulation and results of this method, it is difficult to make a complete comparison to the line relaxation method considered here. However, it appears that the line relaxation method gives convergence factors at least as small as those in (ref. 3); moreover, it is simpler.

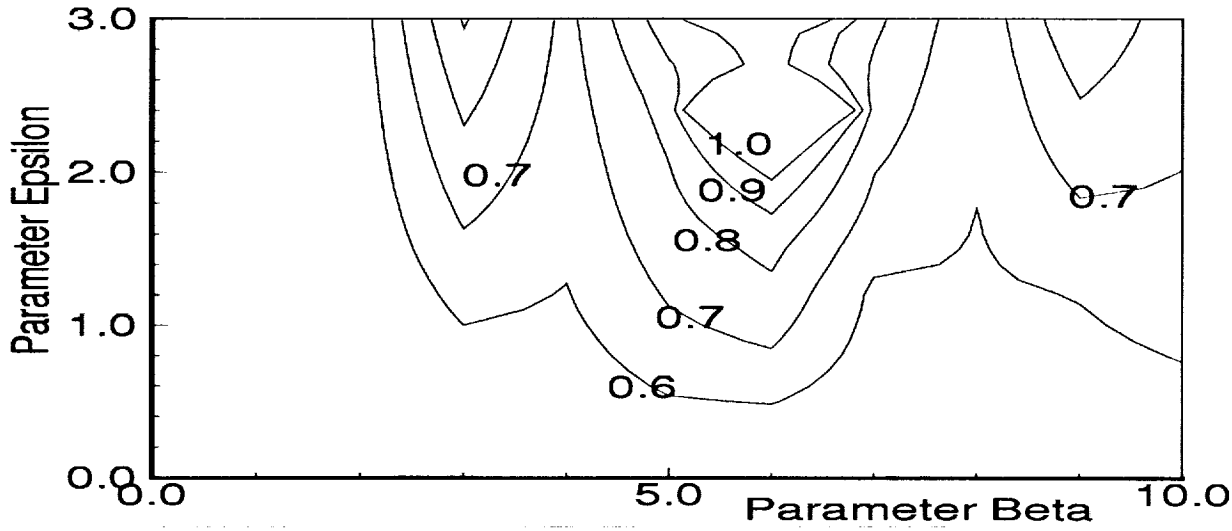

Figure 3. Smoothing factors for problems with variable coefficients.

CONCLUSIONS

The pointwise preconditioning is simple and fast to apply. It is very efficient for one-dimensional problems. Unfortunately, it does not give fast multigrid convergence for two-dimensional problems.

The line relaxation method provides a new approach to accelerate the multigrid Chebyshev spectral method for solving two-dimensional elliptic problems. It is efficient (yielding multigrid smoothing factors no larger than 0.5 per sweep) and inexpensive (requiring $O(N^2 \log N)$ operations per sweep).

When high accuracy is required, the spectral multigrid method using line relaxation is orders of magnitude faster than a conventional finite-difference multigrid method, due primarily to the exponential convergence of the spectral discretization. Compared to other methods for solving the discrete spectral equations, the line relaxation method also has advantages: it is comparable in efficiency to matrix diagonalization and finite-difference preconditioned Richardson relaxation, but can solve problems with variable coefficients which the former cannot, and is simpler than the latter.

REFERENCES

1. Canuto; Hussaini; Quarteroni; and Zang: *Spectral Methods in Fluid Dynamics.* Springer-Verlag, 1988.

2. Zang, T. A.; Wong, Y. S.; and Hussaini, M. Y.: Spectral multigrid methods for elliptic equations. *J. Comp. Phys.*, vol. 48, 1982, pp. 485–501.

3. Zang, T. A.; Wong, Y. S.; and Hussaini, M. Y.: Spectral multigrid methods for elliptic equations II. *J. Comp. Phys.*, vol. 54, 1984, pp. 489–507.

4. Brandt, A.: *Multigrid Techniques: 1984 Guide with Applications to Fluid Dynamics*, 1984. [Available from GMD, Postfach 1240, D-5202 St. Augustin, 1, F.R.G.]

5. Haidvogel, Dale B.; Zang, Thomas: The Accurate Solution of Poisson's Equation by Expansion in Chebyshev Polynomials. *J. Comp. Phys.*, vol. 30, 1997, pp. 167-180.