JOURNAL OF COMPUTATIONAL PHYSICS 58, 96–112 (1985)

Improved Spectral Multigrid Methods for Periodic Elliptic Problems

ACHI BRANDT*

Institute for Computational Studies, Fort Collins, Colorado 80522

SCOTT R. FULTON

Department of Atmospheric Science, Colorado State University, Fort Collins, Colorado 80523

AND

G. D. TAYLOR

Department of Mathematics, Colorado State University, Fort Collins, Colorado 80523

Received September 22, 1983; revised April 10, 1984

The spectral multigrid method for periodic elliptic problems is examined. Several modifications are introduced, including a midpoint pseudospectral discretization which eliminates the need for filtering the highest Fourier mode and new relaxation schemes for isotropic and anisotropic problems. Numerical results are presented demonstrating substantial increases in efficiency and accuracy over previous methods. © 1985 Academic Press, Inc.

1. INTRODUCTION

For elliptic problems with smooth solutions, spectral discretizations can give highly accurate approximations with relatively few degrees of freedom (grid points or basis functions) [1]. However, solving the resulting discrete equations efficiently is complicated. In general, the matrices involved are full, and indirect (iterative) methods must be designed so that fast transforms may be used effectively. Several such methods have been proposed [2, 3].

One of the most promising techniques for solving spectral equations is the multigrid procedure introduced by Zang *et al.* [4, 5]. Like conventional multigrid methods for finite difference discretizations [6, 7], this procedure involves cycling between different levels of discretization (grids) to reduce the error on all scales efficiently. The key elements involved are the relaxation scheme used to smooth the

* On leave from the Department of Applied Mathematics, The Weizmann Institute of Science, Rehovot, Israel, 76100.

0021-9991/85 $3.00

error on each grid, the interpolation processes used for transfers between grids, and the cycling algorithm which decides when the grid transfers take place. These elements have all been discussed in detail in [4, 5] for both periodic and Dirichlet problems.

The purpose of this paper is to introduce for the periodic case some refinements which can substantially improve the overall efficiency of the spectral multigrid method. In Section 2 we show that by evaluating the fluxes at the midpoints between the usual collocation points a "midpoint" pseudospectral discretization is obtained which eliminates the need for filtering the highest modes, resulting in improved accuracy and efficiency. For the isotropic case considered in [4], a weighted residual relaxation scheme with a greatly improved smoothing rate is introduced in Section 3.1. It is shown that by properly scaling the relaxation parameters, the smoothing rates obtained for constant coefficient problems also hold even when the coefficients are not constant. Some comparisons of choices for coarse grid operators and residual transfers are also made. In Section 3.2 we introduce another relaxation scheme, based on defect corrections, which is efficient even for highly anisotropic problems. Our results are summarized in Section 4.

2. Fourier Pseudospectral Discretizations

Spectral methods approximate the solution of a problem by a truncated series expansion in terms of some known basis functions, with the expansion coefficients determined by a projection such as Galerkin, tau, or collocation. Here only the collocation projection will be considered, as it is the most readily applicable to equations with variable coefficients; spectral methods based on collocation are often referred to as pseudospectral methods. For periodic problems, Fourier basis functions are appropriate. In this section we describe and contrast two slightly different Fourier pseudospectral approximations.

For illustration purposes, consider the one-dimensional problem

$$\frac{d}{dx}\left[a(x)\frac{du}{dx}\right]=f(x) \qquad (0\leqslant x\leqslant L) \tag{2.1}$$

with periodic boundary conditions, where a and f are regarded as known. The essence of the Fourier pseudospectral method is to approximate the solution u by a truncated Fourier series U in x which satisfies (2.1) at a finite set of discrete points in $[0, L]$ known as collocation points. The approximate solution is represented by its values at the collocation points and its truncated series expansion is used to compute its derivatives exactly. When a and f are infinitely differentiable and periodic, U converges to u asymptotically faster than any algebraic power of the number of terms in the expansion, a phenomenon known as exponential convergence.

Choosing the collocation points as $x_j = jh$ $(j = 0,..., N-1)$ with $h = L/N$ and setting $U_j = U(x_j)$, the truncated Fourier series for U may be represented by the discrete Fourier transform pair

$$U_j = \sum_{p=-N/2+1}^{N/2} \hat{U}_p e^{2\pi ijp/N}, \tag{2.2a}$$

$$\hat{U}_p = \frac{1}{N} \sum_{j=0}^{N-1} U_j e^{-2\pi ijp/N}. \tag{2.2b}$$

Some subtleties involved in this representation are discussed in the Appendix. The expression for the derivative $U'(x)$ at the collocation points corresponding to (2.2a) is

$$U'_j = \sum_{p=-N/2+1}^{N/2-1} \left(\frac{2\pi}{L}\right) ip\hat{U}_p e^{2\pi ijp/N}. \tag{2.3}$$

The highest mode $\hat{U}_{N/2}$ has been filtered from (2.3); two alternative explanations for this may be found in [4] and the Appendix. Using (2.2) and (2.3), one can construct a pseudospectral approximation to (2.1) as follows. The unknown values U_j are related to the values U'_j of the derivative at the collocation points by (2.3), where the spectral coefficients $\hat{U}_p$ are obtained using (2.2b). The derivative values U'_j are then multiplied pointwise by the values $a_j = a(x_j)$, and the result is differentiated in the same way and set equal to $f(x)$ at the collocation points. This results in a full linear system for determining the values U_j. Since the highest mode is filtered in the process of taking the derivatives, this approach will be referred to as the "filtered discretization"; it is the discretization used in [4].

A second, and perhaps more natural, approach avoids the filtering of the highest mode by evaluating the flux $a(x)\,U'(x)$ at the midpoints $x_{j+1/2} = (j+1/2)\,h$ between the collocation points. For this approach, the derivative relation (2.3) is replaced by

$$U'_{j+1/2} = \sum_{p=-N/2+1}^{N/2} \left(\frac{2\pi}{L}\right) ip\hat{U}_p e^{\pi ip/N} e^{2\pi ijp/N}, \tag{2.4}$$

where the phase factor $e^{\pi ip/N}$ comes about in the shift from x_j to $x_{j+1/2}$. In the resulting pseudospectral discretization of (2.1) the flux $a(x)\,U'(x)$ is evaluated at the midpoints $x_{j+1/2}$, and the result is differentiated and transferred back to the collocation points x_j by a formula analogous to (2.4) except with the phase factor $e^{\pi ip/N}$ replaced by $e^{-\pi ip/N}$. With this approach, referred to as the "midpoint discretization," the information in the highest mode is retained.

Consider now the two-dimensional problem

$$\mathscr{L}u = (au_x)_x + (cu_y)_y = f \tag{2.5}$$

on the domain $[0, L_x] \times [0, L_y]$ with periodic boundary conditions. This is the problem to be considered in the remainder of this paper. We introduce the grid

$$\Omega_h = \{(x_j, y_k) = (jh_x, kh_y), 0 \leqslant j < N_x, 0 \leqslant k < N_y\} \tag{2.6}$$

of collocation points with $h = (h_x, h_y) = (L_x/N_x, L_y/N_y)$ and N_x and N_y even. The pseudospectral discretizations described above are easily extended to (2.5); indeed, since the spectral coefficients are used only in computing the derivatives, all operations are essentially one dimensional, with the solution being in spectral space in at most one direction (x or y) at any step. Note that in the filtered discretization, both the highest modes in x and the highest modes in y must be removed. Either of the approximations described above leads to a discrete equation of the form

$$\mathcal{L}^h U^h = F^h, \tag{2.7}$$

where $\mathcal{L}^h$ is a linear operator and U^h and F^h are grid functions consisting of the values of the approximate solution $U(x, y)$ and right-hand side $f(x, y)$ on the grid Ω_h.

To compare the two discretizations consider the isotropic test case introduced in [4], in which the coefficients and analytical solution for (2.5) with $L_x = L_y = 2\pi$ are given by

$$\begin{aligned} a(x, y) &= c(x, y) = 1 + \varepsilon e^{\cos(x+y)}, \\ u(x, y) &= \sin\left(\pi \cos x + \frac{\pi}{4}\right) \sin\left(\pi \cos y + \frac{\pi}{4}\right). \end{aligned} \tag{2.8}$$

Here ε measures the departure of the coefficients a and c from constant, and the right-hand side f is obtained from (2.5). The corresponding truncation error $\tau^h = F^h - \mathcal{L}^h u$, measured in the norm

$$\|\tau^h\| = \left\{ \sum_{j=0}^{N_x - 1} \sum_{k=0}^{N_y - 1} [\tau^h(x_j, y_k)]^2 h_x h_y \right\}^{1/2}, \tag{2.9}$$

is presented in Table I (the numbers in parentheses indicate powers of 10). On the finest grid (32×32) the two discretizations give about the same accuracy when $\varepsilon = 0$; this is reasonable since the solution u is infinitely differentiable and thus the coefficients of the filtered modes (63 of the 1024 total modes, or about 6%) are negligible. However, for nonzero ε, aliasing increases the importance of these modes and the midpoint discretization is significantly more accurate. We note that this increase in accuracy holds only in the region of exponential convergence; in cases where exponential convergence is not obtained it is questionable whether to use spectral methods at all. Further comparison of the two discretizations will be made below in the context of spectral multigrid methods.

TABLE I

Truncation Error for the Isotropic Test Case

Grid size ($N_x \times N_y$)	Discretization	Truncation error $\|\tau^h\|$			
		$\varepsilon = 0.0$	$\varepsilon = 0.1$	$\varepsilon = 0.2$	$\varepsilon = 0.3$
4×4	midpoint	2.5 (+1)	2.9 (+1)	3.3 (+1)	3.7 (+1)
	filtered	0.0	3.9 (0)	7.8 (0)	1.2 (+1)
8×8	midpoint	6.1 (0)	6.9 (0)	7.7 (0)	8.6 (0)
	filtered	5.8 (0)	6.5 (0)	7.8 (0)	9.5 (0)
16×16	midpoint	2.6 (−2)	3.0 (−2)	3.6 (−2)	4.2 (−2)
	filtered	2.6 (−2)	5.9 (−2)	1.1 (−1)	1.6 (−1)
32×32	midpoint	2.2 (−9)	4.9 (−9)	8.8 (−9)	1.3 (−8)
	filtered	2.2 (−9)	6.3 (−8)	1.2 (−7)	1.9 (−7)

3. Relaxation Schemes for Spectral Multigrid Methods

Having formulated the discrete equation (2.7) we now consider multigrid methods for its solution. Central to these methods is the relaxation scheme used to smooth the error on each grid. In [4] the Richardson (or Euler) scheme

$$\bar{u}^h = u^h - \omega r^h \tag{3.1}$$

was used. Here u^h and $\bar{u}^h$ are approximations to the true (discrete) solution U^h on the grid Ω_h before and after the relaxation, respectively, $r^h = F^h - \mathcal{L}^h u^h$ is the residual, and ω is a relaxation parameter. If ω is the same for all sweeps on a given grid the scheme is called stationary; if ω is allowed to vary from one sweep to the next the scheme is called nonstationary. Since the major part of the work involved in (3.1) is in evaluating the residual r^h, we consider below two alternatives to (3.1) which utilize the residual more effectively.

3.1. *Isotropic Problems*

The problem (2.5) is called isotropic on the grid Ω_h if

$$\frac{a}{h_x^2} = \frac{c}{h_y^2}, \tag{3.2}$$

which will be assumed throughout this section. For such problems the residual can be used more effectively by updating the value of u^h at a point using not only the value of r^h at that point, as in (3.1), but also values of r^h at neighboring points [6]. This is the idea of the weighted residual relaxation scheme

$$\bar{u}^h = u^h - \omega \begin{pmatrix} \gamma & \beta & \gamma \\ \beta & \alpha & \beta \\ \gamma & \beta & \gamma \end{pmatrix} r^h. \tag{3.3}$$

Here α, β, and γ are weights to be used at surrounding points, with (3.3) interpreted as

$$\begin{aligned} \bar{u}^h_{jk} = u^h_{jk} - \omega[&\alpha r^h_{jk} + \beta(r^h_{j-1,k} + r^h_{j+1,k} + r^h_{j,k-1} + r^h_{j,k+1}) \\ &+ \gamma(r^h_{j-1,k-1} + r^h_{j+1,k+1} + r^h_{j-1,k+1} + r^h_{j+1,k-1})], \end{aligned} \tag{3.4}$$

where j and k specify the collocation points (x_j, y_k) with appropriate periodicity.

3.1.1. *Smoothing Analysis*

To analyze the smoothing properties of (3.3) we use local mode analysis [7]. When the coefficients a and c are constant the discrete Fourier modes

$$E_\theta(x_j, y_k) = \exp\left[i\left(\frac{x_j \theta_x}{h_x} + \frac{y_k \theta_y}{h_y} \right) \right] \tag{3.5}$$

are the eigenfunctions of $\mathscr{L}^h$ with corresponding eigenvalues

$$\lambda_\theta = -\left(\frac{a\theta_x^2}{h_x^2} + \frac{c\theta_y^2}{h_y^2} \right). \tag{3.6}$$

Here $\theta = (\theta_x, \theta_y) = 2\pi(p/N_x, q/N_y)$ is the discrete wavenumber with p and q integers such that $|p| \leqslant N_x/2$ and $|q| \leqslant N_y/2$. Then if the error $u^h - U^h$ before the relaxation (3.3) has a component $A_\theta E_\theta$, the error $\bar{u}^h - U^h$ after relaxation will have the component $\bar{A}_\theta E_\theta$ with

$$\bar{A}_\theta = A_\theta\{1 + \omega\lambda_\theta[\alpha + 2\beta(\cos\theta_x + \cos\theta_y) + 4\gamma \cos\theta_x \cos\theta_y]\}. \tag{3.7}$$

Choosing the relaxation parameter ω as

$$\omega = \frac{h_x^2}{a} = \frac{h_y^2}{c} \tag{3.8}$$

in (3.7) results in the convergence factor

$$\mu(\theta) = \frac{\bar{A}_\theta}{A_\theta} = 1 - (\theta_x^2 + \theta_y^2)[\alpha + 2\beta(\cos\theta_x + \cos\theta_y) + 4\gamma \cos\theta_x \cos\theta_y]. \tag{3.9}$$

Now the modes representable on the grid Ω_h have $|\theta| = \max(|\theta_x|, |\theta_y|) \leqslant \pi$; of these, the modes not representable on the coarser grid $\Omega_{h/\rho}$ $(0 < \rho < 1)$ are the high

wavenumbers $|\theta| > \rho\pi$. Including $|\theta| = \rho\pi$ as high wavenumbers for convenience we define the multigrid smoothing factor as

$$\bar{\mu} = \max_{\rho\pi \leqslant |\theta| \leqslant \pi} |\mu(\theta)|, \tag{3.10}$$

i.e., the factor by which all high wavenumber errors are reduced by one relaxation sweep. We seek to choose the residual weights α, β, and γ so as to minimize $\bar{\mu}$.

When only the single parameter α is considered (i.e., $\beta = \gamma = 0$) the weighted residual scheme (3.3) reduces to the Richardson scheme (3.1), except for the definition of ω. In this case the optimal α and corresponding μ can be obtained analytically as

$$\alpha = \frac{2}{(2+\rho^2)\pi^2}, \qquad \bar{\mu} = \frac{2-\rho^2}{2+\rho^2}. \tag{3.11}$$

For the standard mesh ratio $\rho = \frac{1}{2}$, (3.11) gives $\bar{\mu} = \frac{7}{9}$ as obtained for small h in [4], where the maximum in (3.10) did not include the filtered modes $|\theta| = \pi$. When more than one parameter is considered the minimization problem for $\bar{\mu}$ must be solved numerically. Table II shows the optimal choices for two parameters (α, β) and three parameters (α, β, γ), along with the corresponding smoothing factors, for the mesh ratio $\rho = \frac{1}{2}$. These results were obtained using the differential-correction algorithm [8], and show that the weighted residual scheme improves the smoothing factor dramatically. Further testing shows that three parameters are probably best, as including more parameters (i.e., utilizing the residual at even more surrounding points) leads to only minimal reductions in the smoothing factor.

3.1.2. *Implementation*

The weighted residual relaxation scheme was tested in a simple multigrid procedure based on the one described in [4, 5]. Since the details of the method were given in those papers, only a brief description will be included here.

TABLE II

Optimal Parameters and Smoothing Factors for the Weighted Residual Relaxation Scheme

Number of parameters	Parameters considered	Optimal values	Corresponding $\bar{\mu}$
1	α	0.9006 (−1)	0.7778
2	α	0.1491	0.4718
	β	0.3024 (−1)	
3	α	0.2240	0.1058
	β	0.7000 (−1)	
	γ	0.2800 (−1)	

Four grids were used, with 4, 8, 16, and 32 points in each direction. For simplicity a simple V-cycle control structure was used, with fixed numbers N_d and N_u of relaxations on each grid in the downward (fine to coarse) and upward (coarse to fine) branches, respectively. The results reported here are for $N_d = 2$ and $N_u = 1$, which were optimal for the weighted residual scheme; $N_d = 3$ and $N_u = 0$ were used with nonstationary Richardson relaxation (3 parameters). Initial approximations on all grids were zero.

The grid transfers employed were based on Fourier interpolation; they consist of transforming the grid function, setting the higher Fourier coefficients to zero or appending additional zero coefficients as appropriate, and transforming back to the new grid. These transfers are described in detail in [4, 5]. Two points deserve mention here. First, the grid transfers (like the residual calculation) can be carried out one direction (x or y) at a time, so that only one-dimensional transforms are needed. Second, when using the midpoint discretization, care must be taken with the last Fourier coefficient (see the Appendix). As implied by (A5), the last coefficient on a coarse grid must be divided by 2 when transferring to a finer grid; conversely, when transferring from a fine grid to a coarse grid the last coefficient for the coarse grid must be multiplied by 2. Note that with the filtered discretization this consideration does not arise, but the highest Fourier mode in each direction then must be filtered from the residual each time it is computed.

The analysis in the previous section is rigorous when the coefficients a and c are constant. Numerical experiments show that when a and c are not constant (but still equal) and the relaxation scheme (3.1) or (3.3) is not changed, the convergence rate degrades significantly; in fact, the method often diverges. A partial fix suggested in [5] is underrelaxation, i.e., decreasing the value of ω by 25–50%. In this manner convergence may be obtained, but the rate is still significantly degraded. A more effective choice is to evaluate ω *pointwise* by (3.8) as suggested by the local mode analysis. This pointwise "scaling" by the coefficient a (or c) is used in all results reported here; with it there is practically no degradation of convergence rates for nonconstant a and c. For the weighted residual method one can choose to scale either the relaxation parameter ω or the residual r^h; we refer to the resulting schemes as the scaled weighted residual and weighted scaled residual relaxation schemes, respectively. Note that for the midpoint discretization the coefficient a is stored at the midpoints in x; simple linear interpolation to get corresponding values at the collocation points for scaling works well.

3.1.3. *Numerical Results*

The spectral multigrid method was tested by starting with an initial approximation of zero and making V-cycles repeatedly until the norm of the residual after a cycle was less than the norm of the truncation error. The convergence rate can be measured by the convergence factor per V-cycle

$$\mu_V = \left(\frac{\|u^h - U^h\|_n}{\|u^h - U^h\|_0} \right)^{1/n} \tag{3.12}$$

where the subscripts "0" and "n" denote values before the first and after the last V-cycle, respectively, and n is the total number of cycles executed. This convergence factor can be related to the work done in several ways.

The standard multigrid work unit is the amount of work involved in one relaxation sweep on the finest grid. Counting work in the usual way gives $W = (1 + \frac{1}{4} + \frac{1}{16} + \frac{1}{64}) \times (N_d + N_u) \doteq 3.984$ work units per V-cycle. Admittedly this measure does not take into account the work of residual transfers and interpolation, nor the fact that for the spectral method the work increases as $N_x N_y \log(N_x N_y)$. However, the convergence factor per work unit $\mu_W = (\mu_V)^{1/W}$ should be close to the smoothing factor $\bar{\mu}$ (helpful in checking the code) and provides an estimate of the efficiency which is independent of both the computer and the programmer. Table III shows the convergence factor μ_W obtained for the isotropic test case (2.8) using the weighted residual and Richardson relaxation schemes for both the midpoint and filtered discretizations. The Richardson schemes were implemented as described in [4]. For the filtered discretization one could in principle reduce the smoothing factor μ slightly by excluding the filtered modes $|\theta| = \pi$ from the maximum in (3.10); in practice this is rather awkward, as the parameters α, β, and γ then depend on h, and was done only for the stationary Richardson scheme. Recalling that ε measures the departure of the coefficients a and c from constant we see that the scaling introduced above works well, with the weighted scaled residual scheme producing the best results.

To obtain a more direct measure of efficiency it is convenient to define a time unit as the actual execution time required to evaluate the midpoint residual on the finest grid. The convergence factor per time unit $\mu_T = (\mu_V)^{1/T}$, where T is the number of time units per V-cycle, then includes the work of grid transfers and other overhead (excluding setup work) but depends on both the computer and the skill of

TABLE III

Smoothing Factor $\bar{\mu}$ and Convergence Factor per Work Unit μ_W for the Isotropic Test Case

Relaxation scheme	Discretization	$\bar{\mu}$	Convergence per work unit μ_W			
			$\varepsilon = 0.0$	$\varepsilon = 0.1$	$\varepsilon = 0.2$	$\varepsilon = 0.3$
Weighted scaled residual	midpoint	0.1058	0.1040	0.1184	0.1196	0.1116
	filtered		0.1135	0.1486	0.1794	0.2020
Scaled weighted residual	midpoint	0.1058	0.1040	0.1277	0.1462	0.1583
	filtered		0.1135	0.1641	0.2071	0.2453
Nonstationary Richardson	midpoint	0.5994	0.5972	0.6143	0.6272	0.6372
	filtered		0.6118	0.6209	0.6330	0.6430
Stationary Richardson	midpoint	0.7778	0.7654	0.7752	0.7820	0.7874
	filtered	0.7510	0.7384	0.7294	0.7264	0.7268

the programmer. The results presented here were obtained on a CRAY-1 computer using assembly language FFTs and vectorized code where possible; making the grid functions explicitly periodic by including two extra grid points in each direction made the code considerably simpler and easier to vectorize. The time unit obtained was 3.149 msec, with (midpoint) Richardson and weighted residual relaxations taking about 1.15 and 1.21 time units, respectively, per sweep on the finest grid. Table IV shows the convergence factor μ_T, along with the number n of V-cycles required to reach truncation error, for the same cases as considered in Table III. With the filtered discretization the residual must be filtered each time it is computed, increasing by about 17% the time required per relaxation sweep on the finest grid and leading to larger values of μ_T than with the midpoint discretization. Since the execution time is proportional to $\log(\mu_T)$ these results show that the weighted residual schemes are about 4 and 8 times faster than the nonstationary and stationary Richardson schemes, respectively, in the context of this spectral multigrid method.

Two modifications of the spectral multigrid method described above were investigated, both of which are related to the idea of compatible coarsening [7]. The first concerns the construction of the coarse grid operators. For the results presented above the coefficients a and c were evaluated on the finest grid and transferred to the coarser grids by Fourier interpolation (taking into account the location of the midpoints), resulting in what Zang *et al.* [5] refer to as the "filtered" coarse grid operators. Alternatively, one can simply evaluate a and c on the coarser grids directly, thereby obtaining the "unfiltered" coarse grid operators. The second modification involves replacing the Fourier transfer of residuals by injection, i.e., evaluating the residual on the fine grid and transferring it to the coarse grid by

TABLE IV

Convergence Factor per Time Unit μ_T and Number of V-Cycles n Required to Reach Truncation Error for the Isotropic Test Case

Relaxation scheme	Discretization	Convergence per time unit μ_T/number of V-cycles n			
		$\varepsilon = 0.0$	$\varepsilon = 0.1$	$\varepsilon = 0.2$	$\varepsilon = 0.3$
Weighted scaled residual	midpoint	0.43/3	0.46/3	0.46/3	0.46/3
	filtered	0.50/3	0.54/3	0.58/3	0.60/3
Scaled weighted residual	midpoint	0.43/3	0.46/3	0.49/3	0.50/3
	filtered	0.49/3	0.56/3	0.60/3	0.63/4
Nonstationary Richardson	midpoint	0.82/11	0.83/11	0.84/11	0.84/11
	filtered	0.85/12	0.85/11	0.86/11	0.87/11
Stationary Richardson	midpoint	0.90/20	0.91/21	0.91/21	0.91/21
	filtered	0.90/21	0.90/18	0.90/17	0.90/17

TABLE V

Comparison of Residual Transfers and Coarse Grid Operators

Residual transfers	Coarse grid operators	Discretization	Convergence per work unit μ_W			
			$\varepsilon = 0.0$	$\varepsilon = 0.1$	$\varepsilon = 0.2$	$\varepsilon = 0.3$
Fourier	filtered	midpoint	0.1040	0.1184	0.1196	0.1116
		filtered	0.1135	0.1486	0.1794	0.2020
	unfiltered	midpoint	0.1040	0.1132	0.1193	0.1138
		filtered	0.1135	0.1535	0.1856	0.2085
Injection	filtered	midpoint	0.1184	0.1525	0.1649	0.1774
		filtered	0.1180	0.1640	0.1884	0.2020
	unfiltered	midpoint	0.1184	0.1520	0.1597	0.1762
		filtered	0.1180	0.1638	0.1883	0.2028

simply using the values at points common to both grids. Table V shows the convergence factors μ_W obtained with these two modifications using the weighted scaled residual relaxation scheme for the isotropic test case (2.8). The changes in convergence factors are fairly small, as are the changes in overall execution time (not shown). Fourier residual transfers give slightly better results than injection, with little difference between the filtered and unfiltered coarse grid operators.

3.2. *Anisotropic Problems*

When (3.2) does not hold, the problem (2.5) is said to be anisotropic on the grid Ω_h. The weighted residual scheme can be extended to this case, but the optimal weights α, β, and γ then depend on a and c. These weights could be determined for any given problem, but the smoothing factors obtained can be quite large. In general, this approach is tiresome at best.

A more robust relaxation scheme based on defect corrections can be described as follows. Each relaxation sweep consists of computing the residual $r^h = F^h - \mathscr{L}^h u^h$, solving

$$\mathscr{L}_D^h V^h = r^h \tag{3.13}$$

for the approximate correction V^h, and updating the approximation u^h via

$$\bar{u}^h = u^h + \omega V^h. \tag{3.14}$$

Here ω is a relaxation parameter and $\mathscr{L}_D^h$ is an approximation to $\mathscr{L}^h$, which we will take as the five-point second-order finite-difference operator

$$(\mathscr{L}_D^h V^h)_{jk} = \frac{a_{j+1/2,k}(V^h_{j+1,k} - V^h_{j,k}) - a_{j-1/2,k}(V^h_{j,k} - V^h_{j-1,k})}{h_x^2} + \frac{c_{j,k+1/2}(V^h_{j,k+1} - V^h_{j,k}) - c_{j,k-1/2}(V^h_{j,k} - V^h_{j,k-1})}{h_y^2}. \tag{3.15}$$

3.2.1. *Smoothing Analysis*

For the defect correction method described above, local mode analysis gives the convergence factor

$$\mu(\theta) = 1 - \omega\left[\frac{ah_y^2\theta_x^2 + ch_x^2\theta_y^2}{2ah_y^2(1-\cos\theta_x) + 2ch_x^2(1-\cos\theta_y)}\right], \tag{3.16}$$

assuming a and c are constant. The inequalities $2(1-\cos\phi) \leqslant \phi^2 \leqslant (\pi^2/4) \times 2(1-\cos\phi)$ for $|\phi| \leqslant \pi$ lead to the lower and upper bounds 1 and $\pi^2/4$ for the quantity in brackets in (3.16) on the region $0 \leqslant |\theta| \leqslant \pi$. Thus the maximum of $|\mu(\theta)|$ on this region is minimized by choosing

$$\omega = \frac{8}{\pi^2+4}, \tag{3.17}$$

which yields the corresponding smoothing factor

$$\bar{\mu} = \frac{\pi^2-4}{\pi^2+4} \doteq 0.4232. \tag{3.18}$$

This smoothing factor could be reduced somewhat by choosing ω to minimize $|\mu(\theta)|$ over $\rho\pi \leqslant |\theta| \leqslant \pi$ as in (3.10), or by weighting the residuals at surrounding points as before. However, in either case the optimal parameters then depend on a and c and the resulting improvement in $\bar{\mu}$ is small (especially for a/c much different than one). Therefore the choice (3.17) seems the most useful in practice.

3.2.2. *Implementation*

When using (3.13)–(3.14) as a relaxation scheme within a spectral multigrid method, (3.13) must only be solved for the high wavenumbers, and only approximately at that. Thus one can "solve" (3.13) by simply making a small number of sweeps of an appropriate relaxation method, obtaining an approximation $v^h \approx V^h$ from an initial approximation of zero; this amounts to a preconditioning very similar to that used for Chebyshev spectral methods by Zang *et al.* [4, 5]. For the results presented here we chose alternating direction line relaxation to accommodate general a and c [6]. This consists of first relaxing along lines of constant y by solving

$$\frac{a_{j-1/2,k}}{h_x^2}\,\bar{v}^h_{j-1,k} - \left(\frac{a_{j-1/2,k}+a_{j+1/2,k}}{h_x^2} + \frac{c_{j,k-1/2}+c_{j,k+1/2}}{h_y^2}\right)\bar{v}^h_{j,k} + \frac{a_{j+1/2,k}}{h_x^2}\,\bar{v}^h_{j+1,k}$$

$$= r^h_{jk} - \frac{c_{j,k-1/2}}{h_y^2}\,v^h_{j,k-1} - \frac{c_{j,k+1/2}}{h_y^2}\,v^h_{j,k+1} \tag{3.19}$$

for $\bar{v}^h$, and then doing an analogous sweep along lines of constant x. Vectorization was achieved by solving first for the odd ("white") lines and then for the even ("black") lines, resulting in alternating direction zebra (ADZ) relaxation. The linear systems involved are symmetric and tridiagonal (with periodicity), and were solved in parallel using a modified Cholesky decomposition of the form $A = LDL^T$, where D is diagonal and L is lower triangular with ones along the main diagonal. By computing the factorizations once and storing them the execution time per ADZ sweep (about 0.95 time unit on the finest grid) could be cut in half; for simplicity this was not done here. Since the midpoint discretization is quite natural for this scheme, the filtered discretization was not tested.

3.2.3. *Numerical Results*

The defect correction relaxation scheme described above was tested in the spectral multigrid method for the anisotropic test case

$$a(x, y) = 1 + \eta \cos x, \qquad c(x, y) = 1 + \eta \cos y,$$
$$u(x, y) = \sin\left(\pi \cos x + \frac{\pi}{4}\right) \sin\left(\pi \cos y + \frac{\pi}{4}\right). \tag{3.20}$$

Here η measures the amount of anisotropicity, with $(a/c)_{\max} = (c/a)_{\max} = (1+\eta)/(1-\eta)$. The convergence factor per work unit μ_W obtained is shown in Table VI and compared with that obtained using the other relaxation schemes. For these results

TABLE VI

Smoothing Factor $\bar{\mu}$ and Convergence Factor per Work Unit μ_W for the Anisotropic Test Case

Relaxation scheme	$\bar{\mu}$	Convergence per work unit μ_W				
		$\eta = 0.0$	$\eta = 0.2$	$\eta = 0.4$	$\eta = 0.6$	$\eta = 0.8$
Multigrid						
Defect correction	0.4232	0.4553	0.4570	0.4609	0.4683	0.4843
Scaled weighted residual	0.1058	0.1040	0.3943	1.110	2.872	11.67
Nonstationary Richardson	0.5994	0.5972	0.5944	1.138	1.924	3.802
Stationary Richardson	0.7778	0.7654	0.7643	1.030	1.680	3.252
Single grid (32×32)						
Defect correction	0.4232	0.9611	0.9584	0.9550	0.9512	0.9480

only one ADZ sweep was used in solving (3.13), resulting in 2.15 time units per defect correction sweep on the finest grid. By using more ADZ sweeps the values of μ_W can be reduced approximately 5%, but this improvement is not worth the computational cost. The defect correction method achieved a convergence factor per time unit μ_T of about 0.83, with truncation error reached in 9 V-cycles. Further testing shows that the convergence rates reported here may deteriorate somewhat for very large anisotropicity $[(a/c)_{\max} \gtrsim O(10^2)]$.

The analysis in Section 3.2.1 indicates that when the coefficients a and c are constant and (3.13) is solved to sufficient accuracy, (3.18) is in fact the maximum convergence factor $|\mu(\theta)|$ for all modes θ. This suggests that the defect correction scheme may be useful as a solution method on a single fine grid. Indeed, if (3.13) is solved exactly, then this is the "spectral iteration" method of Orszag [2] and McCrory and Orszag [9]. However, the approximate solution of (3.13) by one ADZ sweep as above is not adequate on a single fine grid; the resulting convergence factors shown in the last line of Table VI are fairly large, and only slight improvement is obtained by using more ADZ sweeps. The problem here is that ADZ sweeps give an accurate solution only for the high wavenumbers. This suggests that an efficient way to solve (3.13) approximately would be to make a single cycle of a standard (finite-difference) multigrid method. In this way multigrid techniques could be used "interior" to the defect correction method, rather than "exterior" to it as described above. The relative effectiveness of these two approaches has not been examined.

4. Conclusions

We have examined the spectral multigrid method of [4, 5] and studied several modifications which significantly improve its accuracy and efficiency for periodic elliptic problems. We conclude that

(1) the midpoint discretization described in Section 2 is more accurate than the standard filtered discretization, and is more efficient since it eliminates the need to filter the residual each time it is computed;

(2) for isotropic problems the weighted residual relaxation scheme introduced in Section 3.1 yields a smoothing factor considerably better than that of the Richardson schemes with very little increase in computational work;

(3) scaling the residual pointwise by the coefficients preserves the constant coefficient convergence factor even when the coefficients are not constant;

(4) anisotropic problems can be solved efficiently by the spectral multigrid method using the defect correction relaxation scheme introduced in Section 3.2.

The results presented here further substantiate the usefulness of spectral multigrid methods.

APPENDIX

Truncated Fourier series are usually represented on a computer by discrete Fourier transforms (computed via the FFT algorithm). The distinction between the two is often inconsequential and hence overlooked, but in cases where it does matter it can lead to some confusion. In this Appendix we consider this problem in detail.

Let $f(x)$ be a function of period L expressible as the truncated Fourier series

$$f(x) = \sum_{p=-N/2}^{N/2} \tilde{f}_p e^{2\pi i p x/L}, \tag{A1}$$

where the coefficients are given by

$$\tilde{f}_p = \frac{1}{L}\int_0^L f(x)\, e^{-2\pi i p x/L}\, dx \qquad \left(|p| \leqslant \frac{N}{2}\right) \tag{A2}$$

and N is even. To obtain suitable discrete forms of (A1) and (A2) one introduces the points $x_j = jh$ ($j = 0,..., N-1$) with $h = L/N$. Clearly the N values $f(x_j)$ ($j = 0,..., N-1$) do not suffice to completely represent $f(x)$, as (A1) involves the $N+1$ degrees of freedom $\tilde{f}_{-N/2}, \tilde{f}_{-N/2+1},..., \tilde{f}_{N/2}$. Indeed, evaluating (A1) at x_j gives

$$f_j = f(x_j) = \sum_{p=-N/2+1}^{N/2-1} \tilde{f}_p e^{2\pi i j p/N} + (-1)^j(\tilde{f}_{-N/2} + \tilde{f}_{N/2}), \tag{A3}$$

showing that the lost information is precisely the difference $\tilde{f}_{N/2} - \tilde{f}_{-N/2}$. Since this difference does not contribute to the values f_j it is convenient to take it to be zero, i.e.,

$$\tilde{f}_{-N/2} = \tilde{f}_{N/2}. \tag{A4}$$

Then defining new spectral coefficients $\hat{f}_p$ by

$$\hat{f}_p = \begin{cases} \tilde{f}_p & |p| < \dfrac{N}{2} \\ 2\tilde{f}_p & |p| = \dfrac{N}{2} \end{cases}, \tag{A5}$$

(A3) reduces to

$$f_j = \sum_{p=-N/2+1}^{N/2} \hat{f}_p e^{2\pi i j p/N}, \tag{A6}$$

with the inverse given by

$$\hat{f}_p = \frac{1}{N}\sum_{j=0}^{N-1} f_j e^{-2\pi i j p/N}. \tag{A7}$$

Equations (A6) and (A7) constitute the discrete Fourier transform pair.

The continuous form of the derivative is

$$f'(x) = \sum_{p=-N/2}^{N/2} \left(\frac{2\pi}{L}\right) ip\tilde{f}_p e^{2\pi ipx/L}. \tag{A8}$$

To get an analogous discrete form one can evaluate (A8) at the points x_j, obtaining

$$f'_j = f'(x_j) = \sum_{p=-N/2}^{N/2} \left(\frac{2\pi}{L}\right) ip\tilde{f}_p e^{2\pi ijp/N}. \tag{A9}$$

However, in view of (A4) the contributions for $p = -N/2$ and $p = +N/2$ cancel each other, so using (A5), (A9) yields the discrete form

$$f'_j = \sum_{p=-N/2+1}^{N/2-1} \left(\frac{2\pi}{L}\right) ip\hat{f}_p e^{2\pi ijp/N} \tag{A10}$$

and the information in the highest modes ($|p| = N/2$) is simply lost. Note that if the discrete Fourier coefficients $\hat{f}'_p$ of the derivative are obtained from the discrete coefficients $\hat{f}_p$ by the "natural" formula

$$\hat{f}'_p = \left(\frac{2\pi}{L}\right) ip\hat{f}_p \qquad \left(p = -\frac{N}{2}+1,\ldots,\frac{N}{2}\right) \tag{A11}$$

then the highest mode must be filtered (i.e., one must set $\hat{f}'_{N/2} = 0$) to avoid introducing a spurious contribution into (A10).

An alternative approach which retains the information in the highest modes is to evaluate the derivative at the midpoints $x_{j+1/2} = (j+1/2)h$ ($j = 0,\ldots, N-1$) between the points x_j. From (A8) one obtains

$$f'_{j+1/2} = f'(x_{j+1/2}) = \sum_{p=-N/2}^{N/2} \left(\frac{2\pi}{L}\right) ip\tilde{f}_p e^{\pi ip/N} e^{2\pi ijp/N}, \tag{A12}$$

which can be reduced using (A5) to the discrete form

$$f'_{j+1/2} = \sum_{p=-N/2+1}^{N/2} \left(\frac{2\pi}{L}\right) ip\hat{f}_p e^{\pi ip/N} e^{2\pi ijp/N}. \tag{A13}$$

With this approach the "natural" spectral differentiation formula (A11) is valid; the need for filtering has been eliminated at the cost of a slight increase in operation count [(A13) involves a multiplication by a complex number, where (A10) involves a multiplication by a pure imaginary number].

ACKNOWLEDGMENTS

We wish to acknowledge helpful discussions with Professor Thomas Zang and our colleagues Wayne Schubert, Paul Ciesielski, and Mark DeMaria. We are also indebted to Odilia Panella for her help in preparing the manuscript. This work was supported by the Office of Naval Research under Contracts N00014-83-K-0068 and N00014-C-79-0793 and the National Science Foundation under Grant ATM-82-07563. Acknowledgment is also made of the National Center for Atmospheric Research, which is sponsored by the National Science Foundation, for computer time used in this research.

REFERENCES

1. D. GOTTLIEB AND S. A. ORSZAG, "Numerical Analysis of Spectral Methods," NSF-CBMS Monograph No. 26, Soc. Ind. and Appl. Math., Philadelphia, 1977.
2. S. A. ORSZAG, *J. Comput. Phys.* **37** (1980), 70.
3. Y. S. WONG, T. A. ZANG, AND M. Y. HUSSAINI, "Efficient Iterative Techniques for the Solution of Spectral Equations," Symposium on Spectral Methods in Partial Differential Equations, ICASE, NASA Langley Research Center, Hampton, Va., August 16–18, 1982.
4. T. A. ZANG, Y. S. WONG, AND M. Y. HUSSAINI, *J. Comput. Phys.* **48** (1982), 485.
5. T. A. ZANG, Y. S. WONG, AND M. Y. HUSSAINI, "Spectral Multigrid Methods for Elliptic Equations II," *J. Comput. Phys.*, in press.
6. A. BRANDT, *Math. Comput.* **31** (1977), 333.
7. A. BRANDT, *in* "Multigrid Methods" (W. Hackbusch and U. Trottenberg, Eds.), Springer-Verlag, New York, 1982.
8. E. H. KAUFMAN, JR. AND G. D. TAYLOR, *Internat. J. Numer. Methods Engrg.* **9** (1975), 297.
9. R. L. MCCRORY AND S. A. ORSZAG, *J. Comput. Phys.* **37** (1980), 93.