JOURNAL OF COMPUTATIONAL PHYSICS 54, 489–507 (1984)

Spectral Multigrid Methods for Elliptic Equations II

THOMAS A. ZANG*

College of William and Mary, Williamsburg, Virginia 23185

YAU SHU WONG† AND M. YOUSUFF HUSSAINI†

Institute for Computer Applications in Science and Engineering, NASA Langley Research Center, Hampton, Virginia 23665

Received May 17, 1983

A detailed description of spectral multigrid methods is provided. This includes the interpolation and coarse-grid operators for both periodic and Dirichlet problems. The spectral methods for periodic problems use Fourier series and those for Dirichlet problems are based upon Chebyshev polynomials. An improved preconditioning for Dirichlet problems is given. Numerical examples and practical advice are included.

I. INTRODUCTION

The motivation for applying a pseudospectral discretization to elliptic problems is to obtain a highly accurate approximation with a small number of collocation points. The major advantage that this sort of discretization often offers over standard finite difference or finite element techniques is greatly reduced storage requirements. At the NASA Ames Symposium on Multigrid Methods, we proposed a spectral multigrid approach to solving the discrete equations which arise from applying pseudospectral approximations to variable-coefficient self-adjoint elliptic equations [1]. The focus of that preliminary report was on problems with periodic boundary conditions. We demonstrated that the number of multigrid iterations necessary to achieve convergence was independent of the size of the problem. The tentative results given in [1] for problems with Dirichlet boundary conditions were not so satisfactory because the required number of multigrid iterations increased with the number of grid points. The purpose of this paper is to fill in some of the details omitted from [1] because of page constraints, to describe an improved version of spectral multigrid for Dirichlet problems, and to offer some practical advice for implementation.

* Supported under NASA Grant NAG1-109.

† Supported under NASA Contracts NAS1-17130 and NAS1-17070 while in residence at ICASE, NASA Langley Research Center, Hampton, Va. 23665.

0021-9991/84 $3.00

II. Spectral Multigrid on a Simple Model Problem

The fundamentals of spectral multigrid (SMG) are perhaps easiest to grasp for the simple model problem

$$-\frac{d^2u}{dx^2}=f \tag{1}$$

on $(0, 2\pi)$ with periodic boundary conditions. We will examine this trivially solvable problem in detail for the benefit of those unfamiliar with either spectral or multigrid methods. The standard collocation points are

$$x_j=\frac{2\pi j}{N}, \qquad j=0, 1,..., N-1. \tag{2}$$

Let $f_j=f(x_j)$ and let u_j be the approximation to $u(x_j)$. The discrete Fourier coefficients of u_j are

$$\hat{u}_p=\frac{1}{N}\sum_{j=0}^{N-1} u_j e^{-2\pi ijp/N}, \qquad p=-\frac{N}{2},-\frac{N}{2}+1,...,\frac{N}{2}-1. \tag{3}$$

The inverse relationship can be written

$$u_j=\sum_{p=-N/2}^{(N/2)-1} \hat{u}_p e^{ipx_j}. \tag{4}$$

Thus, a sensible approximation to the left-hand side of Eq. (1) at the collocation points is

$$\sum_{p=-N/2}^{(N/2)-1} p^2\hat{u}_p e^{ipx_j}. \tag{5}$$

The pseudospectral approximation to Eq. (1) may be represented by

$$LU=F, \tag{6}$$

where

$$U=(u_0, u_1,..., u_{N-1}), \tag{7}$$

$$F=(f_0, f_1,..., f_{N-1}), \tag{8}$$

and

$$L=C^{-1}DC. \tag{9}$$

The matrix C represents the discrete Fourier transform; its elements are

$$C_{pj} = \frac{1}{N} e^{-2\pi ipj/N}. \tag{10}$$

Clearly,

$$(C^{-1})_{jp} = e^{2\pi ijp/N}. \tag{11}$$

The diagonal matrix D represents the second derivative in transform space,

$$D_{pq} = p^2 \delta_{p,q}. \tag{12}$$

In Eqs. (5), (10)–(12) the indices p and q have the range indicated in Eq. (3); refer to Eq. (2) for the range of j.

A Richardson's iterative scheme [2] for solving Eq. (6) is

$$V \leftarrow V + \omega(F - LV), \tag{13}$$

where V is the current approximation to U and ω is a relaxation parameter. The eigenfunctions of L are

$$\xi_j(p) = e^{2\pi ijp/N}, \tag{14}$$

with the corresponding eigenvalues

$$\lambda(p) = p^2. \tag{15}$$

The ranges of j and p are the same as above. The index p has a natural interpretation as the frequency of the eigenfunction.

The error at any stage of the iterative process is $V - U$; it can be resolved into an expansion in the eigenvectors of L. Each iteration reduces the pth error component to $\nu(\lambda_p)$ times its previous value, where

$$\nu(\lambda) = 1 - \omega\lambda. \tag{16}$$

The optimal choice of ω results from minimizing $|\nu(\lambda)|$ for $\lambda \in [\lambda_{\min}, \lambda_{\max}]$, where $\lambda_{\min} = 1$ and $\lambda_{\max} = N^2/4$. (One need not worry about the $p = 0$ eigenfunction since it corresponds to the mean level of the solution, which is at one's disposal for this problem.) The optimal relaxation parameter for this single-grid procedure is

$$\omega_{SG} = \frac{2}{\lambda_{\max} + \lambda_{\min}}. \tag{17}$$

It produces the spectral radius

$$\rho_{SG} = \frac{\lambda_{\max} - \lambda_{\min}}{\lambda_{\max} + \lambda_{\min}}. \tag{18}$$

Unfortunately, $\rho_{SG} \simeq 1 - 8/N^2$, which implies that $O(N^2)$ iterations are required to achieve convergence.

This slow convergence is the outcome of balancing the damping of the lowest frequency eigenfunction with that of the highest frequency one in the minimax problem described after Eq. (16). The multigrid approach takes advantage of the fact that the low-frequency modes ($|p| < N/4$) can be represented just as well on coarser grids. It balances the middle-frequency eigenfunction ($|p| = N/4$) with the highest frequency one ($|p| = N/2$), and, hence, damps effectively only those modes which cannot be resolved on coarser grids. In Eqs. (17) and (18), $\lambda_{\min}$ is replaced with $\lambda_{\text{mid}} = \lambda(N/4)$. The optimal relaxation parameter in this context is

$$\omega_{MG} = \frac{2}{\lambda_{\max} + \lambda_{\text{mid}}}. \tag{19}$$

The multigrid smoothing factor

$$\mu_{MG} = \frac{\lambda_{\max} - \lambda_{\text{mid}}}{\lambda_{\max} + \lambda_{\text{mid}}} \tag{20}$$

measures the damping rate of the high-frequency modes. In this example $\mu_{MG} = 0.60$, independent of N. The price of this effective damping of the high-frequency errors is that the low-frequency errors are hardly damped at all. However, on a grid with $N/2$ collocation points, the modes for $|p| \in [N/8, N/4]$ are now the high-frequency ones. They get damped on this grid. Still coarser grids can be used until relaxations are so cheap that one can afford to damp all the remaining modes, or even to solve the discrete equations exactly.

The spectral multigrid approach requires the use of a sequence of grids (or levels). Denote these levels by the index k, where $k = 2, 3, \ldots, K$. Level k consists of N_k collocation points, where $N_k = 2^k$. Equations (2)–(12) apply on each level with N_k replacing N. Throughout this paper the symbol k will be used solely to denote a multigrid level.

On each level we must define a discrete problem, a relaxation scheme, and interpolation operators. The discrete problems will be denoted by

$$L^k V^k = F^k. \tag{21}$$

On the finest level, $L^k = L$, $F^k = F$, and $V^k = U$, the solution to Eq. (6). The relaxation scheme will be confined here to be Richardson iteration

$$v^k \leftarrow v^k + \omega^k (F^k - L^k v^k), \tag{22}$$

where v^k is the approximation to V^k and ω^k is the relaxation parameter. The interpolation operator R^k represents the fine-to-coarse restriction of residuals from level k to level $k-1$,

$$F^{k-1} = R^k (F^k - L^k v^k). \tag{23}$$

The interpolation operator P^k represents the prolongation of corrections from level $k-1$ to level k,

$$v^k \leftarrow v^k + P^k v^{k-1}. \tag{24}$$

Appropriate choices for the coarse-grid operators L^k and the interpolation operators are discussed in the next section. Many choices are possible for the scheduling algorithm which controls the transfer between grids. We will return to this issue in the section on numerical examples.

III. Interpolation and Coarse Grid Operators

We will focus on the one-dimensional problem

$$\frac{d}{dx}\left\{a(x)\frac{du}{dx}\right\} = f(x), \tag{25}$$

on either $(0, 2\pi)$ as in the periodic case or on $(-1, 1)$ as in the Dirichlet case. We will occasionally refer to equations from [1], denoting them by the prefix I, for example, Eq. (I.5).

Fourier Series

The natural interpolation operators represent trigonometric interpolation. They were defined in [1] by Eqs. (I.31) and (I.32). Useful explicit representations of the restriction and prolongation operators (with the superscript k suppressed) are

$$R_{jl} = \frac{1}{N} \sum_{q=-(N/4)+1}^{(N/4)-1} e^{2\pi i q(2j-l)/N} \tag{26}$$

and

$$P_{jl} = \frac{2}{N} \sum_{q=-(N/4)+1}^{(N/4)-1} e^{2\pi i q(j-2l)/N}. \tag{27}$$

These summations may be performed in closed form to yield

$$R_{jl} = \frac{1}{N} S_{2j-l} \tag{28}$$

and

$$P_{jl} = \frac{2}{N} S_{l-2j}, \tag{29}$$

where

$$\begin{aligned} S_r &= \frac{N}{2} - 1, && r \equiv 0 \pmod{N}, \\ &= \sin\left(\frac{\pi r}{2}\right) \cot\left(\frac{\pi r}{N}\right) - \cos\left(\frac{\pi r}{2}\right), && \text{otherwise.} \end{aligned} \tag{30}$$

In analyzing the coarse-grid operator, Eqs. (26) and (27) are more useful than Eqs. (28)–(30). Moreover, as noted in [1], the interpolation can be implemented efficiently by fast Fourier transforms rather than by using Eqs. (28)–(30) in matrix-vector multiplications. By the way, the definition of C given here in Eq. (10) differs slightly (by a factor $\sqrt{N}$) from the definition used in [1]. Note that except for a factor of 2, P and R are adjoint.

The pseudospectral evaluation of the left-hand side of Eq. (25) can be expressed as

$$MAM\,U, \tag{31}$$

where

$$M = C^{-1}DC, \tag{32}$$

$$A_{jl} = a(x_j)\,\delta_{j,l}, \tag{33}$$

and in a slight change of notation the diagonal matrix D which represents the first derivative in wavenumber space is given by

$$\begin{aligned} D_{pp} &= ip, && |p| < \frac{N}{2}, \\ &= 0, && p = -\frac{N}{2}. \end{aligned} \tag{34}$$

The reason for setting $D_{pp} = 0$ for $p = -N/2$ is given in [1]. Equation (31) costs only $O(N \ln N)$ operations to evaluate when the fast Fourier transform is used. A simple, efficient, and effective choice for the coarse-grid operator L^{k-1} is

$$L^{k-1} = M^{k-1}A^{k-1}M^{k-1}, \tag{35}$$

where A^k is the diagonal matrix given by

$$A^k_{jl} = \tilde{a}^k(x_j)\,\delta_{j,l}, \tag{36}$$

with $\tilde{a}^K(x_j) = a(x_j)$, and for $k = 3, 4, \ldots, K$,

$$\tilde{a}^{k-1} = R^k\tilde{a}^k. \tag{37}$$

In other words, the variable coefficient to be used on the coarser grid $k-1$ is a

filtered version of the coefficient on level k. Otherwise, the coarse-grid operators are the natural pseudospectral approximations on those levels.

As has been stressed especially by Nicolaides [3], Hackbusch [4], and Wesseling [5], it seems desirable to use

$$L^{k-1} = R^k L^k P^k, \tag{38}$$

with R^k the adjoint of P^k. The choice made above in Eqs. (35)–(37) does not satisfy Eq. (38), except for special $a(x)$ such as $a(x) \equiv 1$. Indeed, one can show that the coarse-grid operator so produced is equal to the right-hand side of Eq. (38) plus some additional terms which are due to aliasing effects. This is a simple, but lengthy calculation. Here one should be sure to use Eqs. (26) and (27) as well as orthogonality relations such as

$$\begin{aligned} \sum_{j=0}^{N-1} e^{2\pi i j p/N} &= N, \qquad p \equiv 0 \pmod N, \\ &= 0, \qquad \text{otherwise.} \end{aligned} \tag{39}$$

One can achieve a better approximation to the property of Eq. (38), that is, the aliasing terms are far fewer, by a rather simple modification of the pseudospectral method. This technique is known as the two-thirds rule. It consists of discarding the upper third of the frequency spectrum. On a grid with N collocation points Eq. (34) is replaced with

$$\begin{aligned} D_{pp} &= ip, \qquad |p| < N/3, \\ &= 0, \qquad N/3 \leqslant |p| \leqslant N/2, \end{aligned} \tag{40}$$

and the interpolation operators become

$$R_{jl} = \frac{1}{N} \sum_{q=-(N/6)+1}^{(N/6)-1} e^{2\pi i q(2j-l)/N} \tag{41}$$

and

$$P_{jl} = \frac{2}{N} \sum_{q=-(N/6)+1}^{(N/6)-1} e^{2\pi i q(j-2l)/N}. \tag{42}$$

The price of this modification is that one-third of the collocation points are wasted. Thus the two-thirds rule version of SMG would have to produce a substantial improvement in the convergence rate in order to compensate for its reduced accuracy. Although no such examples have yet emerged from our numerical experiments, the two-thirds rule option may eventually prove to be of some use.

Chebyshev Series

The cosine transform matrix C and the Chebyshev differentiation matrix D as given in Eqs. (I.45) and (I.47) will be left unchanged from [1]. The analogs to Eqs. (26)–(30) are

$$R_{jl} = \frac{2}{N\bar{c}_l} \sum_{q=0}^{N/2}{}'' \bar{c}_q^{-1} \cos\left(\frac{2\pi jq}{N}\right) \cos\left(\frac{\pi lq}{N}\right) \tag{43}$$

and

$$P_{jl} = \frac{4}{N\tilde{c}_l} \sum_{q=0}^{N/2}{}'' \tilde{c}_q^{-1} \cos\left(\frac{2\pi lq}{N}\right) \cos\left(\frac{\pi jl}{N}\right), \tag{44}$$

where

$$\begin{aligned} \bar{c}_q &= 2, \quad && q = 0 \text{ or } N, \\ &= 1, && 1 \leqslant q \leqslant N-1, \end{aligned} \tag{45}$$

$$\begin{aligned} \tilde{c}_q &= 2, \quad && q = 0 \text{ or } N/2, \\ &= 1, && 1 \leqslant q \leqslant N/2 - 1 \end{aligned} \tag{46}$$

and

$$R_{jl} = \frac{2}{N\bar{c}_j} (\bar{Q}_{2j-l} + \bar{Q}_{2j+l}) \tag{47}$$

$$P_{jl} = \frac{4}{N\tilde{c}_l} (Q_{j-2l} + Q_{j+2l}), \tag{48}$$

where

$$\begin{aligned} \bar{Q}_r &= (N/4) + \tfrac{1}{4}, && r \equiv 0 \pmod N, \\ &= \frac{1}{4} + \frac{1}{2} \cos\left(\left(\frac{1}{2} + \frac{1}{N}\right)\frac{\pi r}{2}\right) \sin\left(\frac{\pi r}{4}\right) \csc\left(\frac{\pi r}{N}\right), && \text{otherwise}, \end{aligned} \tag{49}$$

$$\begin{aligned} Q_r &= N/4, && r \equiv 0 \pmod N, \\ &= \frac{1}{4} - \frac{1}{4} \cos\left(\frac{\pi r}{2}\right) + \frac{1}{2} \cos\left(\left(\frac{1}{2} + \frac{1}{N}\right)\frac{\pi r}{2}\right) \sin\left(\frac{\pi r}{4}\right) \csc\left(\frac{\pi r}{N}\right), \\ & && \text{otherwise}. \end{aligned} \tag{50}$$

Equations (43) and (44) represent the "obvious" restriction and interpolation operators. Both may be implemented efficiently by fast cosine transforms. Unlike the Fourier series case, however, R and P are not adjoint (even aside from a constant multiple) unless the boundary conditions happen to be homogeneous Dirichlet. A common choice in finite difference multigrid algorithms and a natural one in finite

element cases is to force R to be adjoint to P. Our own computational experience with Chebyshev SMG leads us to endorse this strategy. For residual transfers, then, we recommend

$$R_{jl} = \frac{2}{N\tilde{c}_j} \sum_{q=0}^{N/2} \tilde{c}_q^{-1} \cos\left(\frac{2\pi jq}{N}\right) \cos\left(\frac{\pi lq}{N}\right), \tag{51}$$

which reduces to

$$R_{jl} = \frac{2}{N\tilde{c}_j} (Q_{2j-l} + Q_{2j+l}). \tag{52}$$

However, the construction of the filtered-coefficient version of the coarse-grid operator via Eqs. (35)–(37) should still be based on the restriction formula in Eq. (43).

IV. An Improved Preconditioning for Dirichlet Problems

Consider the self-adjoint elliptic equation

$$\frac{\partial}{\partial x}\left[a(x, y)\frac{\partial u}{\partial x}\right] + \frac{\partial}{\partial y}\left[b(x, y)\frac{\partial u}{\partial y}\right] = f, \tag{53}$$

on $(-1, 1) \times (-1, 1)$ with Dirichlet boundary conditions. The appropriate pseudospectral approximation employs Chebyshev polynomials. The collocation points (x_j, y_l) satisfy

$$x_j = \cos(\pi j/N), \qquad y_l = \cos(\pi l/N), \quad j, l = 1, 2, \ldots, N-1. \tag{54}$$

Let $\mathcal{N} = (N-1)^2$ denote the total number of degrees of freedom. The pseudospectral approximation leads to a discrete set of equations like Eq. (6). A detailed description of the matrix L representing the Chebyshev discretization of Eq. (53) is given in [1].

It is apparent from Eq. (18) that the convergence rate of Richardson's iteration on a single grid is governed by the ratio of the largest-to-smallest eigenvalues of L. This ratio will be referred to as the single-grid condition number. The multigrid condition number, on the other hand, is the ratio of the largest eigenvalue to the smallest high-frequency eigenvalue. It controls the smoothing rate (see Eq. (20)). The estimates given in [1] for these eigenvalues are $\lambda_{\max} = O(N^4)$, $\lambda_{\text{mid}} = O(N^2)$, and $\lambda_{\min} \simeq \pi^2/4$. The implication is that effective preconditioning is essential for multigrid as well as for single-grid iterative schemes.

Preconditioned Richardson iteration can be expressed as

$$v \leftarrow v + \omega H^{-1}(F - Lv), \tag{55}$$

TABLE I

Extreme Eigenvalues for Preconditioned Chebyshev Operator in Two Dimensions

	$H_{FD}^{-1}L$		$H_{LU}^{-1}L$		$H_{RS}^{-1}L$	
N	λ_{min}	λ_{max}	λ_{min}	λ_{max}	λ_{min}	λ_{max}
4	1.000	1.757	0.929	1.717	1.037	1.781
8	1.000	2.131	0.582	2.273	1.061	2.877
16	1.000	2.305	0.224	2.603	1.043	4.241
24	1.000	2.361	0.111	2.737	1.031	5.379

where H is a preconditioning matrix. An obvious choice for H is a finite difference approximation H_{FD} to the differential operator in Eq. (52). In more than one dimension, these finite difference approximations are themselves costly to invert. An attractive alternative is to use instead an approximate LU-decomposition of H_{FD}, that is, H is taken as the product of a lower triangular matrix $\mathscr{L}$ and an upper triangular matrix $\mathscr{U}$. In one such type of preconditioning (originally proposed by Buleev [6] and Oliphant [7] for finite difference discretizations of Eq. (53)), denoted by H_{LU}, $\mathscr{L}$ is identical to the lower triangular portion of H_{FD} and $\mathscr{U}$ is chosen so that the two super diagonals of LU agree with those of H_{FD}. In [1], a similar decomposition, denoted by H_{RS}, was proposed in which the diagonal elements of $\mathscr{L}$ were altered from those of H_{FD} to ensure that the row sums of H_{RS} and H_{FD} are identical. Incomplete LU-decompositions have been used by Wesseling and Sonneveld [8] for multigrid solutions of finite difference discretizations.

Both types of preconditioning can be computed by a simple recursion. Let (x_j, y_l) be an interior point of the grid. Suppose that the finite difference approximation to Eq. (53) at this point is given by

$$\bar{b}_{j,l}u_{j,l-1} + \bar{d}_{j,l}u_{j-1,l} + \bar{e}_{j,l}u_{j,l} + \bar{f}_{j,l}u_{j+1,l} + \bar{h}_{j,l}u_{j,l+1} = f_{jl}. \tag{56}$$

The lower triangular matrix $\mathscr{L}$ has the nonzero elements $\tilde{b}_{j,l}$, $\tilde{d}_{j,l}$, and $\tilde{e}_{j,l}$ and the

TABLE II

Condition Number for Preconditioned Chebyshev Operator in Two Dimensions

	Single-grid		Multigrid	
N	$H_{LU}^{-1}L$	$H_{RS}^{-1}L$	$H_{LU}^{-1}L$	$H_{RS}^{-1}L$
4	1.85	1.72		
8	3.91	2.71	1.79	2.07
16	11.62	4.07	2.12	2.92
24	24.66	5.22	2.26	3.79

upper triangular matrix $\mathscr{U}$ has unit diagonal plus the nonzero elements $\tilde{f}_{j,l}$ and $\tilde{h}_{j,l}$, where

$$\begin{aligned} \tilde{b}_{j,l} &= \bar{b}_{j,l}, \\ \tilde{e}_{j,l} &= \bar{e}_{j,l} - \tilde{b}_{j,l}\tilde{h}_{j,l-1} - \bar{d}_{j,l}\tilde{f}_{j-1,l} - \alpha(\tilde{b}_{j,l}\tilde{f}_{j,l-1} + \bar{d}_{j,l}\tilde{h}_{j-1,l}), \\ \tilde{f}_{j,l} &= \bar{f}_{j,l}/\tilde{e}_{j,l}, \qquad \tilde{h}_{j,l} = \bar{h}_{j,l}/\tilde{e}_{j,l}. \end{aligned} \tag{57}$$

The H_{LU} result uses $\alpha = 0$ and H_{RS} uses $\alpha = 1$. Straightforward modifications are made near the boundaries.

The eigenvalues of the iteration matrices $H^{-1}L$ corresponding to these three types of preconditioning have been computed numerically by the QR algorithm [9] for the constant coefficient, Poisson's equation. The extreme ones are given in Table I. In all cases, the region between 1 and 2.4 is fairly uniformly populated with eigenvalues. In the H_{LU} version there are a few (roughly 15%) eigenvalues between λ_{min} and 1; likewise, about 20% of the eigenvalues of the H_{RS} preconditioning fall between 2.4 and λ_{max}. A few of the smaller eigenvalues have small imaginary parts. The remaining eigenvalues are real. In order to assess the effectiveness of these preconditionings in multigrid calculations, one also needs to know the smallest high-frequency eigenvalue. The numerical results indicate that this is 1.22 for H_{FD} and H_{LU} and 1.45 for H_{RS}, essentially independent of N. The relevant condition numbers are given in Table II. Both H_{LU} and H_{RS} require only $O(\mathscr{N})$ operations to invert. Thus, we reach the striking conclusion that although H_{RS} is more effective for single-grid iterations, H_{LU} is noticeably superior in multigrid applications.

Beyond $N = 24$ computations of the complete eigenvalue spectra are impractical since the full two-dimensional matrix then takes over a million words of storage. The multigrid condition numbers and smoothing rates given in Table III are based on iterative calculations of the extreme eigenvalues of $H_{LU}^{-1}L$ for $N = 32$ and $N = 64$ and the empirical formulae (based on least squares fits to the data for $N \leqslant 64$),

$$\lambda_{max} \simeq 1.381N^{1/8}, \qquad \lambda_{min} \simeq 28.37N^{-7/4}, \tag{58}$$

for $N > 64$.

The more important of these is the former and it is accurate to better than 1% for

TABLE III

N	Multigrid Condition No.	1-Parameter Smoothing Rate	3-Parameters Smoothing Rate
8	1.863	0.301	0.194
16	2.134	0.362	0.236
32	2.324	0.398	0.262
64	2.525	0.433	0.287
128	2.763	0.469	0.313

$N \geqslant 16$. These results suggest that $O(\mathcal{N}^{17/16} \ln \mathcal{N})$ operations are required for convergence of the SMG method based on the H_{LU} preconditioning. This is only slightly worse than the best possible result of $O(\mathcal{N} \ln \mathcal{N})$. The 1-parameter smoothing rates are based on a stationary Richardson iteration, whereas the 3-parameter smoothing rates are based on nonstationary Richardson iteration employing 3 distinct parameters. Comparing the present smoothing rates with those given in [1] for the H_{RS} preconditioning it is apparent that the use of the H_{LU} preconditioning reduces the number of SMG iterations for $N \geqslant 32$ by at least a factor of 2.

V. Computational Experience

The global character of pseudospectral approximations sharply distinguishes them from local approximations such as finite difference ones. Admittedly, this character makes spectral methods more complicated to implement, but it also is responsible for their superior approximation properties. One should expect that somewhat different considerations are important in SMG codes than in finite difference ones. Here we report on the performance we have obtained with several variants of SMG on two-dimensional problems and offer some general advice on their use.

The operation count for a single relaxation sweep is fundamentally different for SMG—$O(\mathcal{N} \ln \mathcal{N})$ rather than $O(\mathcal{N})$. Thus, the standard multigrid accounting device of assessing the cost of a coarse-grid relaxation as one-fourth of the work on a grid with half the mesh size is inappropriate. We prefer to make our comparisons in terms of actual machine time rather than Brandt's work units [10]. This choice has the virtue of including all auxiliary effort such as the various interpolations, but it also has the disadvantage of depending on the quality of the programmer and the computer.

The specific measure to be used is the equivalent smoothing rate, denoted by μ_e and defined as follows. In some preliminary calculations, the average time τ_0 required for a single fine-grid relaxation is determined. For an actual multigrid calculation let r_1 and r_2 be the residuals after the first and last fine-grid relaxations, respectively, and let τ be the total CPU time. Then

$$\mu_e = \left[\frac{r_2}{r_1}\right]^{1/((\tau/\tau_0)-1)}. \tag{59}$$

In all the runs reported here, the finest level $K = 5$. The four types of schedules that were examined are described in Table IV. In schedules A and B, the problem was first solved on level 2; then that solution was interpolated to level 3 as the initial guess for a multigrid iteration involving levels 2 and 3; then the converged level 3 solution was interpolated to level 4 as its initial guess, etc. This strategy is commonly referred to as the full multigrid approach. The other two schedules simply began on level 5. Most schedules were run in the so-called accommodative mode, that is, the anticipated smoothing rates (e.g., Table III) were used in a dynamic determination of when to

TABLE IV

Description of Multigrid Schedules

Schedule	First Level	Control Mode	Lower Level Problem
A	2	Accommodative	Unfiltered
B	2	Accommodative	Filtered
C	5	Accommodative	N/A
D	5	Fixed	N/A

shift between levels. Schedule D used the simple fixed schedule of performing just one sweep through the parameter sequence on a given level before interpolating to another one. This schedule is known as a V-cycle. All runs employed the correction scheme [10] and used random numbers for the initial guess. The difference between schedules A and B lies in the right-hand side used to define the lower level problems. In the unfiltered version the pointwise values of $f(x)$ were used, but in schedule B, the lower level right-hand sides were obtained by applying the appropriate restriction operator to the finest level right-hand side. The lower level problem distinction is not applicable to schedules C and D.

Periodic Problems

The test problems have the form of Eq. (53) with the coefficients

$$a(x, y) = b(x, y) = 1 + \varepsilon e^{\cos(\beta(x+y))}, \tag{60}$$

and the exact solution

$$u(x, y) = \sin(\alpha\pi \cos x + \pi/4) \sin(\alpha\pi \cos y + \pi/4) - u_{00}, \tag{61}$$

where

$$u_{00} = \tfrac{1}{2} J_0^2(\alpha\pi) \tag{62}$$

guarantees that the mean value of $u(x, y)$ vanishes. The source term $f(x, y)$ is

TABLE V

Parameters of the Periodic Test Problems

Problem No.	ε	α	β	ω/ω_*
1	0.00	1	1	1.00
2	0.10	1	1	0.75
3	0.20	2	2	0.50

TABLE VI

Influence of Coarse-Grid Operator on Fourier SMG

Coarse-Grid Operator	Equivalent Smoothing Rate		
	Problem 1	Problem 2	Problem 3
Unfiltered	0.76	0.78	0.81
Filtered	0.76	0.78	0.81

adjusted accordingly. The parameters of several test problems are listed in Table V. The last column relates the relaxation parameter actually employed to the optimal parameter for the $\varepsilon = 0$ problem as given in Eq. (I.37).

The influence of the coarse-grid operator is indicated in Table VI. The filtered coarse-grid operator is defined by Eqs. (35)–(37). The unfiltered one replaces Eq. (37) with the pointwise values of $a(x_j)$. Schedule C was used for all runs. The filtered operator presents essentially no improvement. Although we find this result puzzling in light of the corresponding results for Dirichlet problem, we did not pursue it further because there are few applications for purely periodic boundary conditions.

The dependence upon scheduling is given in Table VII. These runs used nonstationary Richardson iteration with three distinct relaxation parameters as described in [1]. The filtered coarse-grid operators were also employed. The most striking result is the distinct superiority of schedule B on problem 1. The explanation for this behaviour lies in the very special relationship that exists for the constant-coefficient problem between the interpolation operators R^k and P^k and the operators L^k: The eigenfunctions of L^{k-1} are a subset (in fact precisely the low-frequency subset) of the eigenfunctions of L^k. The prolongation operator P^k leaves the eigenfunctions of L^{k-1} unchanged. Thus, this interpolation introduces no spurious high-frequency components. A similar relationship holds for the restriction operator.

In order to get the full benefit of this property, however, the lower level problems (used for obtaining initial guesses on the higher level problems) must have alias-free

TABLE VII

Influence of Scheduling on Fourier SMG

Schedule	Equivalent Smoothing Rate		
	Problem 1	Problem 2	Problem 3
A	0.56	0.70	0.82
B	0.50	0.70	0.82
C	0.61	0.70	0.80
D	0.70	0.76	0.83

TABLE VIII

Parameters of the Dirichlet Test Problems

Problem No.	ε	α	β	ω/ω_*
1	0.00	1	1	1.00
2	0.20	2	2	1.00
3	1.00	5	10	0.90

right-hand sides. Consider how the simple model problem described by Eq. (1) behaves in transform space for an elementary multigrid scheme which uses only levels 2 and 3. The (transformed) level 3 equations are

$$p^2\hat{u}_p = \hat{f}_p, \qquad p = -4, -3, \ldots, 3. \tag{63}$$

Suppose that a fully converged level 2 solution is used for the initial guess on level 3. If the level 2 problem is defined by the restricted values of $f(x_j)$ (as in schedule B), then its equations are

$$p^2\hat{u}_p = \hat{f}_p, \qquad p = -2, -1, 0, 1. \tag{64}$$

The resulting interpolated initial guess for the level 3 problem will have $\hat{u}_p$ for $p = -2, -1, 0, 1$, precisely correct and $\hat{u}_p$ for $p = -4, -3, 2, 3$, will be 0. Thus, the only errors in the level 3 solution will be in the high-frequency modes and there will be no need to make any coarse-grid corrections, i.e., no true multigridding will occur. On the other hand, suppose that the level 2 problem is defined by the pointwise values of $f(x_j)$ (as in schedule A). Then the level 2 equations will be

$$p^2\hat{u}_p = \tilde{f}_p, \qquad p = -2, -1, 0, 1, \tag{65}$$

where

$$\begin{aligned} \tilde{f}_{-2} &= \hat{f}_{-2} + \hat{f}_2, \\ \tilde{f}_{-1} &= \hat{f}_{-1} + \hat{f}_3, \\ \tilde{f}_0 &= \hat{f}_0 + \hat{f}_{-4}, \\ \tilde{f}_1 &= \hat{f}_1 + \hat{f}_{-3}. \end{aligned} \tag{66}$$

The last term in each of these equations is, of course, the alias of its preceding term. When this converged level 2 solution is interpolated to level 3 for use as the initial guess on the fine level, there will be errors in the low-frequency modes. Hence, coarse-grid corrections will have to be made. Consequently, schedule A consumes more computer time than schedule B.

The superiority of schedule B does not extend to nontrivial cases, as represented here by problems 2 and 3. Since the eigenfunctions of the discrete operators are no

longer simple trigonometric functions, they are not preserved by the interpolation operators.

Schedule D is clearly marked as inferior to schedule C. Although schedule C operates in the accommodative mode, in nearly every case there are 6 relaxations on a level before restriction to a coarser level occurs. Schedule D uses only half as many relaxations before restriction occurs. (Recall that a nonstationary Richardson iteration with 3 parameters is employed, so that the number of relaxations is necessarily a multiple of 3.) Part of the increased efficiency of schedule C arises from the less frequent use of interpolations.

A Fourier SMG program has some subtleties that deserve mention. They are connected with the zero eigenvalues of the discrete operator that arise from the $p = 0$ and $p = -N/2$ diagonal entries of the Fourier differentiation operator (see Eq. (34)). The associated mean-value and highest frequency eigenfunctions are undamped by the iterative scheme. For the constant-coefficient case, a sufficient precaution is to filter these components out of the right-hand side and the initial guess. In variable-coefficient problems one must ensure that none of the highest frequency component enters the solution during a relaxation.

Dirichlet Problems

The test problems are specified by

$$a(x, y) = b(x, y) = 1 + \varepsilon e^{\cos(\beta\pi(x+y))}, \tag{67}$$

$$u(x, y) = \sin(\alpha\pi x + \pi/4)\sin(\alpha\pi y + \pi/4). \tag{68}$$

The parameters of the test problems are given in Table VIII. Problem 1 not only has constant coefficients but it is also so well resolved by the Chebyshev pseudospectral method that its discretization error on level 5 is well below the round-off error of the CDC Cyber 175 (14 digits). Problem 3 is at the other extreme. The coefficients of the equation oscillate so rapidly that the level 5 grid cannot resolve them. Instead the converged solution of the level 5 collocation equations has an error of order 1. This case is included as a test of whether the Chebyshev SMG method is robust enough to converge on such a problem.

The difference between the two choices of the coarse-grid operators is shown in Table IX. Both versions are identical on problem 1. The unfiltered coarse-grid

TABLE IX

Influence of Coarse-Grid Operator on Chebyshev SMG

Coarse-Grid Operator	Equivalent Smoothing Rate		
	Problem 1	Problem 2	Problem 3
Unfiltered	0.62	0.66	Divergent
Filtered	0.62	0.62	0.76

TABLE X

Influence of Scheduling on Chebyshev SMG

Schedule	Equivalent Smoothing Rate		
	Problem 1	Problem 2	Problem 3
A	0.26	0.58	0.77
B	0.25	0.57	0.75
C	0.51	0.59	0.70
D	0.60	0.67	0.74

operators produce a slow method on problem 2. The extra work occurs on the coarser levels where the smoothing is less effective. On problem 3 the unfiltered coarse-grid operators lead to a divergent method. We find it curious that the filtered coarse-grid operators failed to produce a simular improvement in Fourier SMG. Perhaps the difference lies in the use of preconditioning for Chebyshev SMG. (The filtered coefficients are, of course, used in the finite difference preconditioning as well.) Note that the variable coefficients of problem 3 are extremely oscillatory. Our unpreconditioned Fourier SMG method cannot handle anything remotely as difficult.

The scheduling dependence is given in Table X. The filtered coarse-grid operator was used along with nonstationary Richardson iteration. The same trends are apparent here as for the periodic case, except that schedule B is now only slightly better than schedule A on problem 1. In the Chebyshev method the eigenfunctions of the discrete constant-coefficient operator are not preserved by the interpolation procedures. As before schedule D nearly always produces 6 relaxations prior to restriction. We again see that it is not a good strategy to relax the minimum number of times before restricting.

At the end of the third section, we recommended the restriction operator defined by Eq. (51) for the residual transfer. Two alternatives have been tested. The "obvious" restriction operator of Eq. (43) fails miserably. It does not even work for the constant-coefficient problem. In the accommodative mode the algorithm rapidly settles into a "limit cycle" involving levels 2 and 3: it alternates between these two levels, always arriving on either level with the solution in the same state as at the start of its last visit. The second alternative is to "homogenize" the restrictions and prolongations by forcing the boundary values of the corrections to be zero both before and after the interpolation. Although this bizarre choice was made by accident, it actually works. However, since it has uniformly been slightly less effective than the adjoint choice, there is no good reason to resort to it.

TABLE XI

Finite Difference vs Spectral MG for $\alpha = 1$

N	FDMG Error	SMG Error	FDMG CPU Time	SMG CPU Time
8	2.37 (−2)	4.42 (−5)	0.08	0.35
16	5.73 (−3)	8.53 (−13)	0.24	2.32
32	1.42 (−3)	1.25 (−14)	0.76	5.62
64	3.55 (−4)		2.59	
128	8.88 (−5)		9.71	

VI. Comparison with Finite Difference Multigrid

The convergence properties of spectral methods are dramatically different from those of finite difference schemes. Once a spectral method has enough points to resolve all the features of the solution, that is, with errors of a few percent, then a doubling of the number of grid points in each coordinate direction will produce an answer which is many orders of magnitude better. This makes a comparison of spectral with finite difference methods rather subjective for it can be very sensitive to the desired accuracy.

Nevertheless, the novelty of the SMG approach warrants some comparison with finite difference multigrid (FDMG). Brandt [10] has provided a FDMG program. We used this to solve Eq. (53) on $(-1, 1) \times (-1, 1)$ with the coefficients and solution given by Eqs. (67) and (68). An easy problem is given by $\varepsilon = 0$ and $\alpha = 1$ and a harder problem uses $\varepsilon = 0$ and $\alpha = 5$. This FDMG program uses Gauss–Seidel relaxation and has been tailored to Poisson's equation. These same two problems were also solved with a version of our SMG program which was also tailored to Poisson's equation.

The RMS errors and the CPU times in seconds (on a CDC Cyber 175) for these two methods are shown in Tables XI and XII. The finite difference algorithm used

TABLE XII

Finite Difference vs Spectral MG for $\alpha = 5$

N	FDMG Error	SMG Error	FDMG CPU Time	SMG CPU Time
8	1.68 (0)	1.73 (1)	0.07	0.33
16	1.92 (−1)	2.56 (−2)	0.19	1.19
32	4.11 (−2)	5.25 (−10)	0.53	7.84
64	9.87 (−3)		2.26	
128	2.44 (−3)		9.52	

the correction scheme in the full multigrid accommodative mode. The spectral algorithm used the same strategy, that is, schedule B with a 3-parameter Richardson's iteration. The relaxations were continued until the RMS error was within a few percent of its final value. Some experimentation was conducted to determine this optimal stopping point. The results are self-explanatory.

VII. Perspective

Together with [1] this paper represents a comprehensive description of the fundamentals of the spectral multigrid method for solving pseudospectral discretizations of self-adjoint elliptic equations. The key elements of the coarse-grid operators and the interpolation operators have been described in detail and their efficacy demonstrated in the numerical examples. No doubt the spectral multigrid method will prove as amenable to improvements in the relaxation scheme and preconditioning aspects as have finite difference and finite element multigrid methods. The next big step is the demonstration of the power of this method on a difficult nonlinear problem of engineering interest. This has already been achieved and is reported in [11].

References

1. T. A. Zang, Y. S. Wong, and M. Y. Hussaini, *J. Comput. Phys.* **48** (1982), 485.
2. L. F. Richardson, *Philos. Trans. Roy. Soc. London Ser. A* **210** (1910), 307.
3. R. A. Nicolaides, *J. Comput. Phys.* **19** (1975), 418.
4. W. Hackbusch, *Computing* **20** (1978), 291.
5. P. Wesseling, *SIAM J. Sci. Statist. Comput.* **3** (1982), 387.
6. N. I. Buleev, *Math. USSR-Sb.* **51** (2) (1960), 227.
7. T. A. Oliphant, *Quart. Appl. Math.* **19** (1961), 221.
8. P. Wesseling and P. Sonneveld, Numerical experiments with a multiple grid and a preconditioned Lanczos type method, Lecture Notes in Mathematics, No. 771, p. 543, Springer-Verlag, Berlin/New York, 1980.
9. J. H. Wilkinson, "The Algebraic Eigenvalue Problem," Oxford Univ. Press, London/New York, 1965.
10. A. Brandt, *Math. Comput.* **31** (1977), 333.
11. C. L. Streett, T. A. Zang, and M. Y. Hussaini, "Spectral Multigrid Methods with Applications to Transonic Potential Flow," ICASE Report 83-11, April 1983.