NASA Contractor Report 172202

ICASE

NASA-CR-172202
19830026391

Preconditioners for the Spectral Multigrid Method

Timothy N. Phillips

Thomas A. Zang

M. Yousuff Hussaini

Contract No. NAS1-17070 and NAS1-17130
September 1983

INSTITUTE FOR COMPUTER APPLICATIONS IN SCIENCE AND ENGINEERING
NASA Langley Research Center, Hampton, Virginia 23665

Operated by the Universities Space Research Association

National Aeronautics and
Space Administration

Langley Research Center
Hampton, Virginia 23665

Preconditioners for the Spectral Multigrid Method

Timothy N. Phillips
Institute for Computer Applications in Science and Engineering

Thomas A. Zang
NASA Langley Research Center

M. Yousuff Hussaini
Institute for Computer Applications in Science and Engineering

Abstract

The systems of algebraic equations which arise from spectral discretizations of elliptic equations are full and direct solutions of them are rarely feasible. Iterative methods are an attractive alternative because Fourier transform techniques enable the discrete matrix-vector products to be computed with nearly the same efficiency as is possible for corresponding but sparse finite difference discretizations. For realistic Dirichlet problems preconditioning is essential for acceptable convergence rates. A brief description of Chebyshev spectral approximations and spectral multigrid methods for elliptic problems is given. A survey of preconditioners for Dirichlet problems based on second-order finite difference methods is made. New preconditioning techniques based on higher order finite differences and on the spectral matrix itself are presented. The preconditioners are analyzed in terms of their spectra and numerical examples are presented.

Research was supported by the National Aeronautics and Space Administration under NASA Contracts No. NAS1-17070 and No. NAS1-17130 while the first and third author were in residence at ICASE, NASA Langley Research Center, Hampton, VA 23665.

1. Introduction

Spectral methods involve representing the solution to a problem in terms of a truncated series of smooth global functions which are known as trial functions. Specifically, these functions are eigenfunctions of a singular Sturm-Liouville problem (Gottlieb and Orszag (1977)). This global character distinguishes spectral methods from finite difference and finite element methods. It is also responsible for their superior approximation properties, yielding accurate solutions with substantially fewer grid points than are required by finite difference methods.

Test functions are used to minimize the residual that results from the substitution of the series expansion of the solution into the differential equation. The choice of test functions distinguishes between the spectral Galerkin and spectral collocation methods. In the Galerkin approach the test functions are the same as the trial functions whereas in the spectral collocation method the test functions are shifted Dirac delta functions. For many problems, especially nonlinear ones, the spectral collocation method is the easiest to implement, and the most efficient as well (Hussaini, Kopriva, Salas and Zang (1983)). The present discussion is confined to the spectral collocation method, and all future references to the spectral method will mean the spectral collocation method.

The matrices representing the discrete spectral collocation operator corresponding to elliptic problems (Zang, Wong, and Hussaini (1982)) are usually full. Even in the constant coefficient linear case, direct inversion of these matrices is usually expensive. Iterative schemes are a practical necessity. The condition number of the spectral matrices is large, and therefore effective preconditioning is necessary.

In recent years research in the field of elliptic equations has focused on multigrid methods. Basically, the multigrid method is a numerical strategy to solve partial differential equations by switching between finer and coarser levels of discretization. The characteristic feature of the method is the combination of a smoothing step and a coarse grid correction (Brandt (1977), Hackbusch (1980)). During the smoothing step the residuals are not necessarily decreased but smoothed. In the following correction step the discrete solution is improved by means of an auxiliary equation on a coarser grid. This results in an iterative method that is usually very fast and efficient.

2. Formulation of the Problem

We consider the self-adjoint elliptic equation

$$\frac{\partial}{\partial x}\left(a(x,y,u)\,\frac{\partial u}{\partial x}\right)+\frac{\partial}{\partial y}\left(b(x,y,u)\,\frac{\partial u}{\partial y}\right)=f(x,y) \tag{2.1}$$

with Dirichlet boundary conditions in the domain $[-1,1]\times[-1,1]$. A proper representation of the solution to this Dirichlet problem employs Chebyshev polynomials. The details of implementing Chebyshev collocation methods for this type of problem have been given by Zang, Wong, and Hussaini (1982) and by Hussaini, Salas, and Zang (1983). This discretization of equation (2.1) may be written as

$$L_{sp}(V) = F, \tag{2.2}$$

where V is the vector of unknowns at the collocation points, F is the vector of the values of the right-hand side at the collocation points and

L_{sp} is the vector-valued (generally nonlinear) operator representing a spectral discretization of the left-hand side.

For nonlinear problems iterative schemes for (2.2) are a necessity. Iterative methods appear to be preferable to direct methods even for constant coefficient linear problems since the matrices representing L_{sp} are full. No fast direct methods are known for these systems. Iterative methods are appealing because the standard implementation of spectral discretizations employs Fast Fourier Transforms which reduces the cost of evaluating the left-hand side of (2.2) to $O(N \log N)$ operations, even for nonlinear problems, where N is the total number of unknowns. Iterative methods also have a clear advantage over direct methods in terms of storage.

Many iterative schemes can be described within the framework of defect correction. Let H be some approximation to the Jacobian J_L of $L_{sp}(V)$ and let V^n be the latest approximation to V. The simplest iterative scheme is the preconditioned Richardson's method

$$V^{n+1} = V^n - \omega_n H^{-1}[L_{sp}(V^n) - F], \tag{2.3}$$

where ω_n is a relaxation parameter.

In practice the preconditioning matrix H is constructed to be an approximation to J_L having the following properties:

(i) H has a sparse matrix structure;

(ii) H is easily invertible.

One choice for H considered here is an approximate LU decomposition of a finite difference approximation to (2.1), i.e.,

$$H = LU$$

where L and U are lower and upper triangular matrices respectively.

Although in this section we have considered the general nonlinear equation (2.1) we shall restrict ourselves to the linear case in the remainder of this paper. Treatment of the nonlinear case will be the subject of future work.

3. An Heuristic Discussion of Spectral Multigrid

Of course, far better schemes are available than the simple Richardson's method. Multigrid methods have demonstrated their ability to accelerate many types of relaxation schemes (see, e.g., the conference proceeedings edited by Hackbusch and Trottenberg (1982)). Multigrid methods have recently been developed for spectral discretizations of the linear version of (2.1) by Zang, Wong, and Hussaini (1982, 1983). Streett, Zang, and Hussaini (1983) have shown that these techniques are extremely effective for the nonlinear potential flow problem of transonic aerodynamics. Preconditioning techniques play a crucial role in these spectral multigrid schemes. Our purpose here is to investigate additional choices for the preconditioning matrix.

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

$$-\frac{d^2u}{dx^2} = f \tag{3.1}$$

on $[0,2\pi]$ with periodic boundary conditions. The Fourier approximation to the left-hand side of (3.1) at the collocation points $x_j = 2\pi j/N$ is

$$\sum_{p=-N/2+1}^{N/2-1} p^2 \hat{u}_p e^{ipx_j}, \tag{3.2}$$

where $\hat{u}_p$ are the Fourier coefficients of u. The eigenfunctions of this approximation are

$$\xi_j(p) = e^{2\pi ijp/N}, \tag{3.3}$$

with the corresponding eigenvalues

$$\lambda(p) = p^2, \tag{3.4}$$

where $j = 0,1,\cdots,N-1$ and $p = -N/2+1,\cdots,N/2-1$. The index p has a natural interpretation as the frequency of the eigenfunction.

Consider now the iterative process described by (2.3) with H taken to be the identity matrix, i.e., without any preconditioning. The iteration matrix, C, of this scheme is given by

$$C = I - \omega L_{sp}.$$

The iterative scheme is convergent if the eigenvalues, λ, of L_{sp} satisfy

$$|1 - \omega\lambda| < 1.$$

The best choice of ω is that for which

$$(1 - \omega \lambda_{max}) = - (1 - \omega \lambda_{min}),$$

where $\lambda_{max}(= N^2/4)$ and $\lambda_{min}(= 1)$ are the largest and smallest eigenvalues of L_{sp} respectively, for then the largest values of $\mu = 1 - \omega\lambda$ are equal in magnitude and have opposite sign (see Fox (1964)). 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

$$\omega_{SG} = \frac{2}{\lambda_{max} + \lambda_{min}} . \tag{3.5}$$

It produces the spectral radius

$$\mu_{SG} = \frac{\lambda_{max} - \lambda_{min}}{\lambda_{max} + \lambda_{min}} . \tag{3.6}$$

Unfortunately, $\mu_{SG} \simeq 1 - 8/N^2$, 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 above. 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 settles for balancing the middle-frequency one ($|p| = N/4$) with the highest frequency one ($|p| = N/2$), and hence damps effectively only those modes which cannot be resolved on coarser grids. In (3.5) and (3.6), λ_{min} is replaced by $\lambda_{mid} = \lambda(N/4)$. The optimal relaxation parameter in this context is

$$\omega_{MG} = \frac{2}{\lambda_{max} + \lambda_{mid}} . \tag{3.7}$$

The multigrid smoothing factor

$$\mu_{MG} = \frac{\lambda_{max} - \lambda_{mid}}{\lambda_{max} + \lambda_{mid}} \tag{3.8}$$

measures the damping rate of the high-frequency modes. Alternatively, we may write

$$\mu_{MG} = \frac{\kappa_{MG} - 1}{\kappa_{MG} + 1}$$

where $\kappa_{MG} = \lambda_{max}/\lambda_{mid}$ is known as the multigrid condition number. 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. Table I compares the single-grid and multigrid damping factors for N = 64. 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. For the case illustrated in Table I the high-frequency error reduction in the multigrid context is roughly 250 times as fast as the single-grid reduction for N = 64.

Table I. Dampling Factors for N = 64

1	.9980	.9984
2	.9922	.9938
4	.9688	.9750
8	.8751	.9000
12	.7190	.7750
16	.5005	.6000
20	.2195	.3750
24	.1239	.1000
28	.5298	.2250
32	.9980	.6000

Morchoisne (1979) and Orszag (1980) have proposed a preconditioning for spectral methods which amounts to using a low-order finite difference approximation for H. Let $H^{(2)}$, $H^{(4)}$ and L_{sp} denote second-order, fourth-order and spectral discretizations of the operator $-d^2/dx^2$. The eigenvalues of these discretizations are given below:

$$\lambda_k^{(2)} = \frac{2[1 - \cos(k\Delta x)]}{(\Delta x)^2} ,$$

$$\lambda_k^{(4)} = \frac{\cos(2k\Delta x) - 16\cos(k\Delta x) + 15}{6(\Delta x)^2} ,$$

$$\lambda_k^{(\infty)} = k^2 .$$

The effective eigenvalues of the preconditioned iterations based on $\left(H^{(2)}\right)^{-1} L_{sp}$ and $\left(H^{(4)}\right)^{-1} L_{sp}$ are then given by

$$\Lambda_k^{(2)} = (k^2)\left(\lambda_k^{(2)}\right)^{-1}$$

$$= \frac{(k\Delta x)^2}{2[1 - \cos(k\Delta x)]} ,$$

$$\Lambda_k^{(4)} = (k^2)\left(\lambda_k^{(4)}\right)^{-1}$$

$$= \frac{6(k\Delta x)^2}{\cos(2k\Delta x) - 16\cos(k\Delta x) + 15} .$$

Table II. Properties of Finite Difference Preconditioning for the Model Problem

Finite Difference Order	Λ_{min}	Λ_{mid}	Λ_{max}	μ_{SG}	μ_{MG}
2	1.00	1.23	2.47	0.424	0.336
4	1.00	1.06	1.85	0.298	0.273
6	1.00	1.02	1.63	0.240	0.231

Similar results for even higher-order finite difference preconditionings are straightforward but tedious. The key properties of this class of preconditioning are given in Table II. It will be seen in subsequent sections that these results on the spread of the eigenvalues of the preconditioned systems agree well with those obtained computationally in two dimensions for Dirichlet problems.

Unlike the original system, which has a condition number scaling as N^2, the preconditioned system has a condition number which is independent of N. The fourth-order finite difference operator offers around a 20% improvement in convergence rate over the second-order operator. This is partially offset by the additional cost of inverting the finite difference operator. The higher-order preconditionings are of doubtful utility.

The preconditionings are clearly most effective for the longer wavelength eigenfunctions, as reflected by how close Λ_{mid} is to 1. In fact, for the fourth-order version, Λ_{mid} is already so close to 1 that the multigrid convergence rate is only slightly faster than the single-grid rate. The advantage of multigrid for these preconditioned systems only shows up in two dimensions. Unlike the one-dimensional case, the inversion of the two dimensional operator is nontrivial. It seems advisable in these situations to

do only an approximate inversion of the finite difference operator, for example, by using an incomplete LU decomposition of H as the actual preconditioner. The outcome of this choice is that while Λ_{max} can be kept well under control, exhibiting a very slow growth with N, Λ_{min} plunges precipitously to zero. Fortunately Λ_{mid} remains virtually unchanged. Thus multigrid is attractive for these two-dimensional problems.

We describe the multigrid process by considering the interplay between two grids. The fine grid problem can be written in the form

$$L^f U^f = F^f.$$

The decision to switch to the coarse grid is made after the fine grid approximation V^f has been sufficiently smoothed by the relaxation process, i.e., after the high-frequency content of the error $V^f - U^f$ has been sufficiently reduced. The auxiliary equation on the coarse grid is

$$L^c U^c = F^c,$$

where

$$F^c = R[F^f - L^f V^f].$$

The restriction operator R interpolates a function from the fine grid to the coarse grid. The coarse grid and correction are denoted by L^c and U^c, respectively. After an adequate approximation V^c to the coarse grid problem has been obtained, the fine grid approximation is updated using

$$V^f \leftarrow V^f + PV^c.$$

The prolongation operator P interpolates a function from the coarse grid to the fine grid.

The natural prolongation operator in this application represents trigonometric interpolation. We describe this process below.

On the coarse grid the discrete Fourier coefficients of the corrections u_j at the collocation points x_j are computed using

$$\hat{u}_p = \frac{1}{N_c} \sum_{j=0}^{N_c-1} u_j e^{-ipx_j}, \qquad p = -N_c/2, \cdots, N_c/2 - 1.$$

The fine grid approximation is then updated using

$$u_j \leftarrow u_j + \sum_{p=-N_c/2}^{N_c/2-1} \hat{u}_p e^{ip\tilde{x}_j},$$

where $\tilde{x}_j$, $j=0,1,\cdots,N_f - 1$, are the fine grid collocation points.

The restriction operator is constructed in a similar fashion. It turns out that except for a factor of 2, P and R are adjoint. In the Chebyshev case we force R to be the adjoint of P.

4. A Survey of Second-Order Finite Difference Preconditionings

In this section we compare several types of preconditioning based on incomplete LU decomposition (see Meijerink and van der Vorst (1981)) of the matrix which represents the standard five-point second-order finite difference approximation to the differential equation (2.1). Two such preconditionings were discussed in Zang, Wong, and Hussaini (1983). The first, denoted by H_{LU}, has L identical to the lower triangular portion of H_{FD}, the finite difference matrix, and U chosen so that the two super diagonals of LU

agree with those of H_{FD}. The second preconditioning H_{RS} has the diagonal elements of L altered from those of H_{FD} so as to ensure that the row sums of H_{RS} and H_{FD} are identical. We introduce a third type of preconditioning based on the strongly implicit method of Stone (1968).

The elements of L and U in all these instances can be easily computed by simple recursive formulae. The construction of the factors L and U are described in detail for Stone's method.

A five-point approximation to (2.1) at the pth mesh point can be written as

$$B_p\, u_{p-N} + D_p\, u_{p-1} + E_p\, u_p + F_p\, u_{p+1} + H_p\, u_{p+N} = q_p, \qquad (4.1)$$

or, in matrix form, as

$$H_{FD}\, \underset{\sim}{u} = \underset{\sim}{q}\ .$$

Stone's idea was to modify the matrix H_{FD} by a 'small' matrix M so that:

(i) the factorization of $\tilde{H} = H_{FD} + M$ into the product LU involves much less work than the standard LU decomposition of H_{FD}:

(ii) the elements of L and U are easily calculated;

(iii) $\|M\| << \|H_{FD}\|$.

As a step to satisfying these criteria the factors L and U were chosen to have three non-zero diagonals, as shown in Fig. 1, corresponding to the diagonals B, D, and E, and E, F, and H respectively of H_{FD}.

The product LU has seven non-zero diagonals, the additional two being immediately interior to the B and H diagonals. The matrix M is taken to comprise these extra two diagonals. The elements of L and U can be computed from their relationships with the elements of H_{FD}.

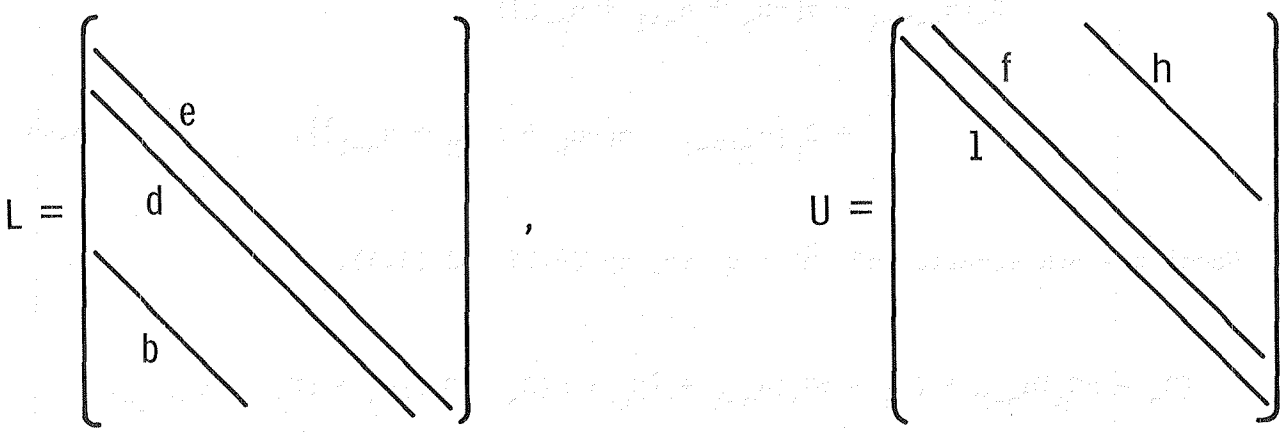

Figure 1

Analogous with (4.1) the pth component of $M\underset{\sim}{u}$ can be written in the form

$$C_p\, u_{p-N+1} + G_p\, u_{p+N-1} \,. \tag{4.2}$$

Stone decided to diminish the magnitude of this term by subtracting from it a closely equivalent expression obtained by Taylor series expansions. Using these expansions it is easily shown that

$$u_{p-N+1} \underset{\sim}{-} - u_p + u_{p+1} + u_{p-N}$$

and (4.3)

$$u_{p+N-1} \underset{\sim}{-} - u_p + u_{p+N} + u_{p-1} \,.$$

Stone then introduced a parameter α, $0 < \alpha < 1$, and defined the pth component of $M\underset{\sim}{u}$ to be

$$C_p\{u_{p-N+1} - \alpha(-u_p + u_{p+1} + u_{p-N})\}$$

$$+ G_p\{u_{p+N-1} - \alpha(-u_p + u_{p+N} + u_{p-1})\}. \qquad (4.4)$$

Hence the pth equation of $\tilde{H}\underset{\sim}{u} = \underset{\sim}{q}$ is, by (4.1) and (4.4),

$$(B_p - \alpha C_p)u_{p-N} + (D_p - \alpha G_p)u_{p-1} + \{E_p + \alpha(C_p + G_p)\}u_p + (F_p - \alpha C_p)u_{p+1}$$

$$+ (H_p - \alpha G_p)u_{p+N} + C_p\, u_{p-N+1} + G_p\, u_{p+N-1} = q_p\,. \qquad (4.5)$$

The relationships between the elements of $\tilde{H} = H_{FD} + M$ and those of L and U are then given by

$$b_p = B_p/(1 + \alpha f_{p-N})$$

$$d_p = D_p/(1 + \alpha h_{p-1})$$

$$e_p = E_p + \alpha(b_p\, f_{p-N} + d_p\, h_{p-1}) - (b_p\, h_{p-N} + d_p\, f_{p-1}) \qquad (4.6)$$

$$f_p = (F_p - \alpha b_p\, f_{p-N})/e_p$$

$$h_p = (H_p - \alpha\, d_p\, h_{p-1})/e_p\,,$$

for $p = 1,\cdots,(N-1)^2$. Any terms with non-positive subscripts that occur in (4.6) are replaced by zero. For any fixed value of α, $0 < \alpha < 1$, (4.6) defines an incomplete LU-decomposition of H_{FD}. We denote this factorization by $H_{ST}(\alpha)$.

The eigenvalues of the iteration matrices $\tilde{H}^{-1} L_{sp}$ corresponding to these types of preconditioning have been computed numerically by the QR algorithm (see Wilkinson (1965)). The extreme eigenvalues for H_{FD}, H_{LU}, and H_{RS} are given in Zang, Wong, and Hussaini (1983), and those for H_{FD}, $H_{ST}(1.0)$ and $H_{ST}(0.9)$ in Table III. A few of the eigenvalues at the lower end of the spectrum have small imaginary parts while the rest are completely real.

Table III. Extreme Eigenvalues for Preconditioned Chebyshev Operator

N	$H_{FD}^{-1} L$		$H_{ST}^{-1}(1.0)L$		$H_{ST}^{-1}(0.9)L$	
	λ_{min}	λ_{max}	λ_{min}	λ_{max}	λ_{min}	λ_{max}
4	1.00	1.76	1.01	1.64	0.99	1.65
8	1.00	2.13	0.78	2.04	0.80	2.07
16	1.00	2.31	0.62	2.28	0.55	2.33
24	1.00	2.36	0.58	2.95	0.36	2.41

Table IV. Single-grid Condition Number

N	$H_{LU}^{-1} L$	$H_{RS}^{-1} L$	$H_{ST}^{-1}(1.0)L$	$H_{ST}^{-1}(0.9)L$
4	1.85	1.72	1.63	1.67
8	3.91	2.71	2.61	2.59
16	11.62	4.07	3.67	4.24
24	24.66	5.22	5.14	6.69

Table V. Multigrid Condition Number

N	$H_{LU}^{-1} L$	$H_{RS}^{-1} L$	$H_{ST}^{-1}(1.0)L$	$H_{ST}^{-1}(0.9)L$
8	1.79	2.07	1.70	1.71
16	2.12	2.92	2.09	2.08
24	2.26	3.79	2.81	2.15

In order to examine the effectiveness of these preconditionings from the multigrid point of view, we need to know the smallest high-frequency eigenvalue. The numerical results indicate that this is 1.22 for H_{LU} and 1.45 for H_{RS}, independent of N. For H_{ST} this value lies between 1.10 and 1.21. Tables IV and V contain the single-grid and multigrid condition numbers respectively for H_{LU}, H_{RS}, $H_{ST}(1.0)$ and $H_{ST}(0.9)$. Here we see that $H_{ST}(0.9)$ is more effective as a preconditioner for multigrid iterations. The multigrid condition number for $H_{ST}(1.0)$ lies between those for H_{LU} and H_{RS}.

In Table III we see that the maximum eigenvalue of $H_{ST}^{-1}(0.9)L_{sp}$ grows more slowly than that of $H_{ST}^{-1}(1.0)L_{sp}$ with increasing N. For values of N greater than 24 the eigenvalue calculation using the QR algorithm becomes too expensive. Using the power method the maximum eigenvalue of $H_{ST}^{-1}(0.9)L_{sp}$ for $N = 32$ was calculated and found to be 2.45. Thus the maximum eigenvalue of $H_{ST}^{-1}(0.9)L_{sp}$ exhibits slow growth with N. This means that the multigrid condition number does not increase drastically with N. However, it is doubtful whether the choice of α in Stone's algorithm is problem independent and this may detract from the robustness of this preconditioning.

To complete this section we look at one other type of incomplete LU-decomposition which is due to Wesseling (1982a). In this decomposition the sparsity of the factors L and U differ from those in Fig. 1 by the addition of extra non-zero diagonals c and g which are located immediately interior to the b and h diagonals respectively. The main diagonal of U is specified to be unity. The elements of L and U are computed from those of H_{FD} (see (4.1)) recursively as follows:

$$
\begin{aligned}
b_p &= B_p \\
c_p &= -b_p\, f_{p-N} \\
d_p &= D_p - b_p\, g_{p-N} \\
e_p &= E_p - b_p\, h_{p-N} - c_p\, g_{p-N+1} - d_p\, f_{p-1} \\
f_p &= \left(F_p - h_{p-N+1}\, c_p\right)/e_p \\
g_p &= -\, d_p\, h_{p-1}/e_p \\
h_p &= H_p/e_p \,,
\end{aligned}
\tag{4.7}
$$

for $p = 1,\cdots,(N-1)^2$. Quantities that are not defined are replaced by zero. This preconditioning will be known as the seven-diagonal preconditioning and we denote it by H_{SD}. The error matrix M contains two non-zero diagonals. It can be shown that the norm of the error matrix of the seven-diagonal preconditioning is smaller than those corresponding to H_{LU} and H_{RS}.

In Table VI we give details of the seven-diagonal preconditioning for different values of N. Included in this table are the extreme eigenvalues of $H_{SD}^{-1} L_{sp}$ and the single-grid and multigrid condition numbers. The smallest high frequency eigenvalue was found to be 1.23 independent of N. The results indicate that H_{SD} is the most effective preconditioner considered to date.

Table VI. Details of Seven-diagonal Preconditioning

N	λ_{min}	λ_{max}	κ_{SG}	κ_{MG}
4	1.00	1.76	1.77	-
8	0.85	2.16	2.54	1.76
16	0.46	2.38	5.22	1.93
24	0.25	2.47	9.81	2.01

5. Higher-order Finite Difference Preconditioning

Here we consider the possibility of choosing the preconditioning matrix $\tilde{H}$ to be a fourth-order finite difference representation of (2.1). For a general non-uniform grid a compact nine-point finite difference approximation cannot be constructed since the set of equations for the coefficients of the scheme is inconsistent. Instead an approximation was constructed based on fourth-order finite difference formulae for each of the second derivatives separately. The coefficients of the function values at the grid points are functions of the mesh lengths which define the non-uniform mesh and are given in the Appendix. At internal points the approximation to the second derivative reduces to the following in the case when the mesh is uniform:

$$\frac{d^2u}{dx^2} = \frac{1}{12h^2}\{-u_{i-2} + 16u_{i-1} - 30u_i + 16u_{i+1} - u_{i+2}\} + O(h^4). \qquad (5.1)$$

The approximation at points one mesh length from the boundary is constructed using the same number of points as interior equations thus maintaining the same order of accuracy. However, in doing this the symmetry of the sparsity pattern of the coefficient matrix is destroyed. As before, it is an incomplete LU decomposition of this finite difference matrix that is used to precondition (2.3). An algorithm due to Wesseling (1982b) was used to perform this decomposition. The factors L and U have the same sparsity pattern as the lower and upper triangular portions of H_{FD} respectively and the corresponding diagonal elements of L and U are equal. We let H_W denote this preconditioning.

Table VII. Extreme Eigenvalues for Preconditioned Chebyshev Operator

N	$H_{FD}^{-1} L_{sp}$		$H_W^{-1} L_{sp}$	
	λ_{min}	λ_{max}	λ_{min}	λ_{max}
8	1.00	1.59	0.51	1.75
16	1.00	1.79	0.19	2.03
24	1.00	1.83	0.09	2.13

Table VIII. Condition Numbers

N	Single-grid	Multigrid
8	3.43	1.73
16	10.68	2.01
24	23.67	2.11

Table VII contains the largest and smallest eigenvalues of $H_{FD}^{-1} L_{sp}$ and $H_W^{-1} L_{sp}$. Table VIII presents the single-grid and multigrid condition numbers for the matrix $H_W^{-1} L_{sp}$. We note that the multigrid condition numbers presented here compare favourably with those obtained using second-order finite difference preconditioning. The smallest high frequency eigenvalue was found to be 1.01, independent of the value of N. A disadvantage of this preconditioning is that more work is needed to perform the decomposition.

6. Spectral Preconditioning

The final type of preconditioning we investigate is that based on the spectral matrix, L_{sp}, itself. As stated earlier the spectral matrix is full and hence costly to invert. Let H_5 be the matrix containing five diagonals of L_{sp}, the positions of which correspond to the non-zero diagonals of the finite difference matrix based on the five point formula. Let the matrix H_9, illustrated in Fig. 2, be the corresponding matrix with nine non-zero diagonals.

Experimentally, we found that several eigenvalues of $H_5^{-1} L_{sp}$ were negative, thus showing that this matrix is not positive definite. The eigenvalues of $H_9^{-1} L_{sp}$ were also computed and found to be positive. Therefore, H_5 was discarded as a potential preconditioner and the usefulness of H_9 examined further.

An approximate LU-decomposition of H_9 is used to precondition (2.2), i.e., H is taken to be the product of a lower triangular matrix L and an upper triangular matrix U. We perform this decomposition according to the criterion described in Section 5 for constructing H_{LU}, i.e., the factors L and U are chosen so that L is identical to the lower triangular portion

of H_9 and U is such that the four super diagonals of LU agree with those of H_9. The sparsity pattern of the factors L and U is shown in Fig. 3. We denote this preconditioning by H_{sp}.

The elements of the matrices L and U can be computed recursively as follows:

$$
\begin{aligned}
a_p &= A_p \\
b_p &= B_p - a_p\, k_{p-2N} \\
c_p &= C_p \\
d_p &= D_p - c_p\, f_{p-1} \\
e_p &= E_p - d_p\, f_{p-1} - c_p\, g_{p-2} - b_p\, k_{p-N} - a_p\, \ell_{p-2N} \\
f_p &= \left(F_p - d_p\, g_{p-1}\right)/e_p \\
g_p &= G_p/e_p \\
k_p &= \left(K_p - b_p\, \ell_{p-N}\right)/e_p \\
\ell_p &= L_p/e_p,
\end{aligned}
\tag{6.1}
$$

for $p = 1,\cdots,(N-1)^2$. Any terms with non-positive subscripts occurring in (6.1) are replaced by zero.

Table IX contains the largest and smallest eigenvalues of $H_9^{-1} L_{sp}$ and $H_{sp}^{-1} L_{sp}$. Table X presents the single-grid and multigrid condition numbers for the matrix $H_{sp}^{-1} L_{sp}$. The smallest high frequency eigenvalue was found to be 0.90, independent of the value of N. For N = 32 we were able to calculate that the maximum eigenvalue of $H_{sp}^{-1} L_{sp}$ was 1.52. We conclude from these results that H_{sp} is likely to be an effective preconditioner in

multigrid applications since the observed behaviour of the condition number compares favourably with that of the other preconditioners considered in this paper.

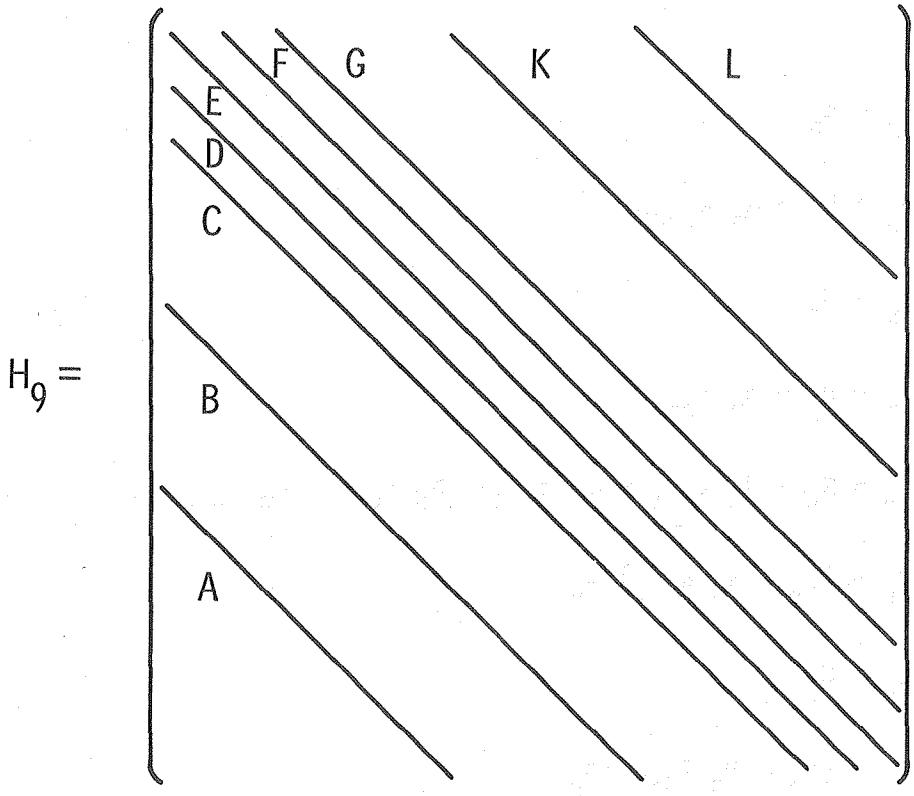

Figure 2

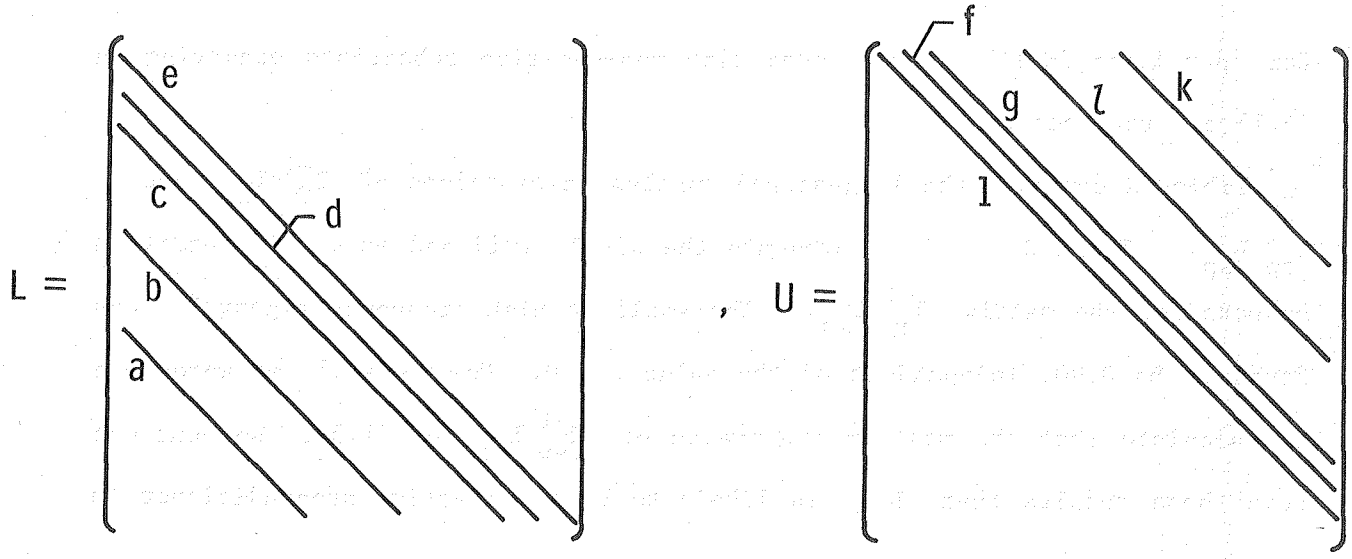

Figure 3

Table IX. Extreme Eigenvalues for the Preconditioned Chebyshev Operator

N	$H_9^{-1} L_{sp}$		$H_{sp}^{-1} L_{sp}$	
	λ_{min}	λ_{max}	λ_{min}	λ_{max}
8	0.63	1.19	0.39	1.31
16	0.23	1.33	0.12	1.47
24	0.11	1.37	0.05	1.51

Table X. Condition Numbers

N	Single-grid	Multigrid
8	3.40	1.44
16	12.53	1.63
24	27.87	1.69

7. Numerical Results

Here we investigate the performance of various preconditioning ideas developed in this paper within the SMG method. The operation count for a relaxation sweep of SMG is $O(N \log N)$ where $N = (N-1)^2$ compared with $O(N)$ for finite difference methods. Thus we make our comparisons in terms of machine time instead of the work units of Brandt (1977).

The measure used is the equivalent smoothing rate, defined by μ_e. This was introduced in earlier work on SMG by Zang, Wong, and Hussaini (1983) and is 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 we define μ_e by

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

There are many variants of the multigrid method. In the one used here we first solve the problem on the coarsest grid, that solution is then interpolated to the next finer level to serve as the initial approximation for a multigrid iteration involving these two levels, etc. The sizes of the grids on the coarsest and finest levels are 4×4 and 32×32 respectively. Internal checks based on the anticipated smoothing rates were used in governing decisions to switch levels. Non-stationary Richardson iteration employing three distinct parameters was used for relaxation. We used the correction scheme of Brandt (1977) with random numbers for the initial guess. On lower levels the right-hand sides were obtained by applying the appropriate restriction operator to the finest level right-hand side.

The test problems are specified by

$$a(x,y) = b(x,y) = 1 + \varepsilon \exp\left(\cos(\beta\pi(x+y))\right),$$

$$u(x,y) = \sin\left(\sigma\pi x + \pi/4\right) \sin\left(\sigma\pi y + \pi/4\right).$$

The parameters of the test problem are given in Table XI. Problem 1 has constant coefficients and is also well-resolved by the Chebyshev collocation method. On the other hand, for Problem 3 the coefficients of the equation oscillate so rapidly that the finest grid cannot resolve them. The converged solution of the finest collocation equations has an error of order 1. This is

a test of whether the Chebyshev SMG method is robust enough to converge on such a problem. Table XII contains the values of the equivalent smoothing rates for various preconditioners.

Table XI. Parameters of the Test Problems

Problem	ε	σ	β
1	0.00	1	1
2	0.20	2	2
3	1.00	5	10

Table XII. Equivalent Smoothing Rates

Problem	H_{LU}	$H_{ST}(0.9)$	H_{SD}	H_{sp}
1	0.26	0.16	0.12	0.24
2	0.60	0.65	0.47	0.54
3	0.76	0.84	0.76	0.82

The results show the seven-diagonal finite difference preconditioning and the spectral preconditioning are comparable in terms of efficiency and give an improvement over the other preconditioners except on Problem 3 where the performance of all the preconditioners is about the same. However, since in our computations to date the spectral matrix is never stored, the elements of H_{sp} are more expensive to compute than H_{SD}. Notice the superb performance of Stone's preconditioning with $\alpha = 0.9$ on the constant coefficient problem. This performance is not maintained on the more difficult problems which demonstrates that the choice of the parameter α is problem dependent. The seven-diagonal preconditioning also performs extremely well on

the constant coefficient problem. We did not experiment with the fourth-order finite difference preconditioning since to compute the approximation in the non-constant coefficient case was thought to be too laborious.

8. Conclusions

In this paper we have developed efficient iterative techniques for solving the algebraic equations which arise from the discretization of a self-adjoint elliptic equation using the spectral collocation method. An advantage of using spectral methods is that they possess superior approximation properties compared with finite difference and finite element methods. In practice, this means one can obtain the same accuracy with fewer mesh points than are needed for finite difference or finite element methods. Acceleration of the basic iterative scheme has been enhanced by employing the multigrid method. The need for preconditioning has been demonstrated and efficient preconditioners for multigrid iterations presented.

Appendix

Higher-order Finite Differences on Non-uniform Grids

Suppose that we require an approximation to the second derivative of u at point C in Fig. 4. We assume that the approximation has the form

$$\frac{d^2u}{dx^2} \simeq au_A + bu_B + cu_C + du_D + eu_E, \qquad (A.1)$$

where the coefficients a, b, c, d and e are to be determined.

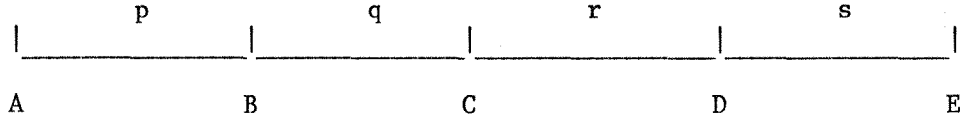

Figure 4

The right-hand side of (A.1) is expanded in a Taylor's series about the point C. The coefficient of the second derivative is set to unity while those of u and the first, third, and fourth derivatives are set to zero. This results in a system of five linear equations for the coefficients appearing in (A.1). The accuracy of the approximation is $O(h^3)$ where h is a typical mesh length.

If we define $G = p + q$ and $H = r + s$, then the coefficients are given by

$$a = \frac{-2q(H + r) + 2Hr}{[Gp(G + H)(r + G)]},$$

$$b = \frac{2G(H + r) - 2Hr}{[pq(q + r)(H + q)]},$$

$$d = \frac{2H(G + q) - 2Gq}{[rs(q + r)(G + r)]},$$

$$e = \frac{-2r(G + q) + 2Gq}{[Hs(H + q)(H + G)]},$$

$$c = -a - b - d - e.$$

References

BRANDT, A. (1977). Multi-level adaptive solutions to boundary-value problems, Math. Comp. 31, 330-390.

FOX, L. (1964). Numerical Linear Algebra. Oxford, Clarendon Press.

GOTTLIEB, D. and ORSZAG, S. A. (1977). Numerical Analysis of Spectral Methods: Theory and Applications. CBMS-NSF Regional Conference Series in Applied Mathematics No. 26, Philadelphia: SIAM.

HACKBUSCH, W. (1980). Convergence of multigrid iterations applied to difference equations, Math. Comp., 34, 425-440.

HACKBUSCH, W. and TROTTENBERG, U. (eds.) (1982). Multigrid Methods, Proceedings of the Conference held at Köln-Porz, West Germany, Lecture Notes in Mathematics No. 960, Berlin,Springer-Verlag.

HUSSAINI, M. Y., KOPRIVA, D. A., SALAS, M. D., and ZANG, T. A. (1983). Spectral Methods for the Euler Equations, AIAA 6th Computational Fluid Dynamics Conference, Danvers, Massachusetts, USA, Paper No. 83-1942-CP.

HUSSAINI, M. Y., SALAS, M. D. and ZANG, T. A. (1983). Spectral methods for inviscid, compressible flows. In Advances in Computational Transonics (W. G. Habashi, ed), Swansea, Pineridge Press.

MEIJERINK, J. A. and VAN DER VORST, H. A. (1981). Guidelines for the usage of incomplete decompositions in solving sets of linear equations as they occur in practical problems, J. Comput. Phys., 44, 134-155.

MORCHOISNE, Y. (1979) Resolution of Navier-Stokes equations by a space-time pseudospectral method. La Recherche Aerospatial 5, 293-309.

ORSZAG, S. A. (1980) Spectral methods for problems in complex geometries, J. Comput. Phys., 37, 70-92.

STONE, H. L. (1968) Iterative solution of implicit approximations of multi-dimensional partial differential equations, SIAM J. Numer. Anal., 5, 530-558.

STREETT, C. L., ZANG, T. A. and HUSSAINI, M. Y. (1983). Spectral multigrid methods with applications to transonic potential flow. ICASE Report 83-11, NASA Langley Research Center, Virginia, USA.

WESSELING, P. (1982a) A robust and efficient multigrid method. In Multigrid Methods (W. Hackbusch and U. Trottenberg, eds.), Lecture Notes in Mathematics No. 960, Berlin, Springer-Verlag.

WESSELING, P. (1982b) Theoretical and practical aspects of a multigrid method, SIAM J. Sci. Statis. Comput., 3, 387-407.

WILKINSON, J. H. (1965). The Algebraic Eigenvalue Problem. Oxford, Clarendon Press.

ZANG, T. A., WONG, Y. S. and HUSSAINI, M. Y. (1982). Spectral multigrid methods for elliptic equations, J. Comput. Phys., 48, 485-501.

ZANG, T. A., WONG, Y. S. and HUSSAINI, M. Y. (1983). Spectral multigrid methods for elliptic equations II, NASA CR-172131.

1. Report No.	2. Government Accession No.	3. Recipient's Catalog No.
NASA CR-172202		

4. Title and Subtitle	5. Report Date
Preconditioners for the Spectral Multigrid Method	September 1983
	6. Performing Organization Code

7. Author(s)	8. Performing Organization Report No.
T. N. Phillips, T. A. Zang, and M. Y. Hussaini	83-48
	10. Work Unit No.

9. Performing Organization Name and Address	11. Contract or Grant No.
Institute for Computer Applications in Science and Engineering Mail Stop 132C, NASA Langley Research Center Hampton, VA 23665	NAS1-17070 & NAS1-17130
	13. Type of Report and Period Covered
12. Sponsoring Agency Name and Address	Contractor Report
National Aeronautics and Space Administration Washington, DC 20546	14. Sponsoring Agency Code

15. Supplementary Notes

Langley Technical Monitor: Robert H. Tolson
Final Report

16. Abstract

The systems of algebraic equations which arise from spectral discretizations of elliptic equations are full and direct solutions of them are rarely feasible. Iterative methods are an attractive alternative because Fourier transform techniques enable the discrete matrix-vector products to be computed with nearly the same efficiency as is possible for corresponding but sparse finite difference discretizations. For realistic Dirichlet problems preconditioning is essential for acceptable convergence rates. A brief description of Chebyshev spectral approximations and spectral multigrid methods for elliptic problems is given. A survey of preconditioners for Dirichlet problems based on second-order finite difference methods is made. New preconditioning techniques based on higher order finite differences and on the spectral matrix itself are presented. The preconditioners are analyzed in terms of their spectral and numerical examples are presented.

17. Key Words (Suggested by Author(s))	18. Distribution Statement
spectral, multigrid, elliptic equations, preconditioning	64 Numerical Analysis Unclassified-Unlimited

19. Security Classif. (of this report)	20. Security Classif. (of this page)	21. No. of Pages	22. Price
Unclassified	Unclassified	33	A03

For sale by the National Technical Information Service, Springfield, Virginia 22161

NASA-Langley, 1983

End of Document