JOURNAL OF COMPUTATIONAL PHYSICS 30, 167–180 (1979)

The Accurate Solution of Poisson's Equation by Expansion in Chebyshev Polynomials

DALE B. HAIDVOGEL

Center for Earth and Planetary Physics, Pierce Hall, Harvard University, Cambridge, Massachusetts 02138

AND

THOMAS ZANG

ICASE, NASA Langley Research Center, Hampton, Virginia 23665

Received September 14, 1977; revised January 30, 1978

A Chebyshev expansion technique is applied to Poisson's equation on a square with homogeneous Dirichlet boundary conditions. The spectral equations are solved in two ways—by alternating direction and by matrix diagonalization methods. Solutions are sought to both oscillatory and mildly singular problems. The accuracy and efficiency of the Chebyshev approach compare favorably with those of standard second- and fourth-order finite-difference methods.

1. INTRODUCTION

Many physical problems require a numerical solution of the two-dimensional Poisson's equation. On certain domains, such as rectangles, the so-called fast Poisson solvers that are now widely available (cf. the reviews by Dorr [1] and Swarztrauber [2]) provide a rapid solution of the standard five-point difference approximation to the partial differential equation. However, the resolution of these methods is inherently limited by their algebraic convergence, i.e., provided that the solution of the continuous problem has continuous and bounded fourth partial derivatives (see, for instance, Isaacson and Keller [3, pp. 445–452]), the maximum error of the discrete approximation with $N + 1$ grid points in each direction decays as N^{-2}. Because of this relatively slow rate of convergence, the storage requirements can be quite severe and the computational time rather long for very accurate calculations.

In recent years spectral methods (cf. Gottlieb and Orszag [4]) have proven to be one way to manage the resolution problem, at least for smooth solutions in simple geometries. As discussed by Gottlieb and Orszag, the essence of the spectral approach is the expansion of the solution into a (truncated) series of smooth functions. The convergence of this approximation is governed by the rate of decay of the expansion

0021-9991/79/020167–14$02.00/0

coefficients. The order of magnitude of these coefficients can be estimated by repeated application of integration by parts. If the basis functions are suitably chosen, the boundary terms from each integration by parts will vanish and thus one power of $1/k$ will be added to the estimate of the kth coefficient. These integrations by parts can be repeated so long as the solution can continue to be differentiated. In particular, if the solution is infinitely differentiable, then the expansion coefficients will decrease faster than any finite power of $1/k$. Thus, the error made by retaining only a finite number N of the terms in the series will itself decrease faster than any finite power of $1/N$—the convergence of the spectral approximation will have an exponential rather than an algebraic character. On the other hand, if the solution has only a finite number of derivatives, then the spectral approximation is expected to converge algebraically, that is, like some finite power of $1/N$.

In this paper we describe the application of the particular spectral method based on an expansion in Chebyshev polynomials to Poisson's equation in a square with homogeneous Dirichlet boundary conditions:

$$\begin{aligned} \Delta u(x, y) &= f(x, y), \qquad |x|, |y| < 1, \\ u(x, y) &= 0, \qquad |x| = 1, |y| = 1. \end{aligned} \tag{1}$$

For these basis functions the boundary terms arising from repeated integrations by parts necessarily vanish. In this crucial respect, Chebyshev polynomials are superior to trigonometric functions since the corresponding boundary terms of the latter do not all vanish, not even for homogeneous boundary conditions. Thus an expansion in Chebyshev polynomials can be much more rapidly convergent than an expansoin in, say, sine functions. After presenting a general formulation of the method, we will give several examples, illustrating both the exponential convergence for an infinitely differentiable solution and the algebraic convergence for a solution with some unbounded derivatives. The accuracy and speed of this spectral method will be compared with those of the standard finite-difference approximation.

2. Formulation

In the Chebyshev spectral approach to the two-dimensional Poisson equation, both u and f are approximated by truncated double Chebyshev series,

$$\begin{aligned} u(x, y) &= \sum_{n=0}^{N} \sum_{m=0}^{N} a_{nm} T_n(x)\, T_m(y), \\ f(x, y) &= \sum_{n=0}^{N} \sum_{m=0}^{N} f_{nm} T_n(x)\, T_m(y), \end{aligned} \tag{2}$$

where the polynomials $T_n(x)$ can be expressed as

$$T_n(x) = \cos(n \cos^{-1} x). \tag{3}$$

Following Orszag's [5] use of the recurrence relations for derivatives of Chebyshev polynomials, the partial differential equation (1) can be used to obtain the relationship

$$\frac{1}{c_n}\sum_{\substack{p=n+2\\ p\equiv n(\mathrm{mod}\ 2)}}^{N} p(p^2-n^2)\,a_{pm}+\frac{1}{c_m}\sum_{\substack{q=m+2\\ q\equiv m(\mathrm{mod}\ 2)}}^{N} q(q^2-m^2)\,a_{nq}=f_{nm},$$

$$0\leqslant n\leqslant N-2,\qquad 0\leqslant m\leqslant N-2, \tag{4}$$

where $c_0=2$ and $c_n=1$ for $n\geqslant 1$, between the coefficients of the solution and the source. The notation $p\equiv n\ (\mathrm{mod}\ 2)$ means that the sum includes every other term, e.g., $p=n+2$, $n+4$, etc. The boundary conditions and the property $T_n(\pm 1)=(\pm 1)^n$ imply

$$\sum_{\substack{p=0\\ p\equiv 0(\mathrm{mod}\ 2)}}^{N} a_{pm}=\sum_{\substack{p=1\\ p\equiv 1(\mathrm{mod}\ 2)}}^{N} a_{pm}=0,\qquad 0\leqslant m\leqslant N, \tag{5}$$

$$\sum_{\substack{q=0\\ q\equiv 0(\mathrm{mod}\ 2)}}^{N} a_{nq}=\sum_{\substack{q=1\\ q\equiv 1(\mathrm{mod}\ 2)}}^{N} a_{nq}=0,\qquad 0\leqslant n\leqslant N. \tag{6}$$

Equations (5) and (6) give only $4N$ independent equations since the four corners of the square have been counted twice. Hence, when the boundary conditions are combined with Eq. (4), the $(N+1)\times(N+1)$ coefficients a_{nm}, for $0\leqslant n, m\leqslant N$, are completely determined. As explained by Orszag [5], this approach of truncating the exact, infinite Chebyshev expansion for $u(x, y)$ by dropping the equations for the highest modes from Eq. (4) and determining them directly from the boundary conditions amounts to Lanczos' tau method [6].

Equations (4)–(6) can be written more compactly as

$$Aa+aA^T=f, \tag{7}$$

where a and f denote matrices with the entries a_{nm} and f_{nm}, $0\leqslant n, m\leqslant N-2$, and A is a matrix representing the Chebyshev approximation to $\partial^2/\partial x^2$ with homogeneous boundary conditions, i.e., Aa represents the first term in Eq. (4) with the boundary conditions of Eq. (5) used to eliminate $a_{N-1,m}$ and $a_{N,m}$ for $0\leqslant m\leqslant N-2$. The y-derivatives in Eq. (4) appear above in the term aA^T, where A^T denotes the transpose of the matrix A.

The successful implementation of this scheme requires efficient numerical procedures for converting between the physical and spectral representations of u and f and for solving the linear system of equations for a_{nm}.

The Fast Fourier Transform [7] is well suited to the first task since the Chebyshev expansion is, in effect, a cosine sum in terms of the variables $(\cos^{-1}x, \cos^{-1}y)$. We thus focus just upon the points $(x_i, y_j)=(\cos \pi i/N, \cos \pi j/N)$, $0\leqslant i, j\leqslant N$, and view Eq. (2) as a pair of discrete Fourier series. Then, for suitable values of N such

as powers of 2, the Fast Fourier Transform enables us to evaluate the solution $u(x, y)$ rapidly at these discrete points via

$$u\left(\cos\frac{\pi i}{N}, \cos\frac{\pi j}{N}\right) = \sum_{n=0}^{N}\sum_{m=0}^{N} a_{nm}\cos\frac{n\pi i}{N}\cos\frac{m\pi j}{N}, \qquad 0 \leqslant i, j \leqslant N, \tag{8}$$

and to approximate efficiently the Chebyshev coefficients of the source $f(x, y)$ as

$$f_{nm} = \frac{4}{N^2\bar{c}_n\bar{c}_m}\sum_{i=0}^{N}\sum_{j=0}^{N}(\bar{c}_i\bar{c}_j)^{-1} f\left(\cos\frac{\pi i}{N}, \cos\frac{\pi j}{N}\right)\cos\frac{n\pi i}{N}\cos\frac{m\pi j}{N}, \quad 0 \leqslant n, m \leqslant N, \tag{9}$$

where $\bar{c}_0 = \bar{c}_N = 2$ and $\bar{c}_i = 1$, $1 \leqslant i \leqslant N-1$.

3. An ADI Solution Technique

The choice of a procedure for solving the linear system of equations (7) is less clear-cut. We present here several alternatives. In some situations an attractive method is alternating direction implicit (ADI) iteration [8] which can provide an exact, as well as a reasonably fast approximate, solution to Eq. (7). In the ADI scheme an initial guess a^0 is made, and the iteration proceeds according to the prescription

$$(\omega_\nu I + A)\, a^{\nu+\frac{1}{2}} = a^\nu(\omega_\nu I - A^T) + f, \tag{10}$$

$$a^{\nu+1}(\omega_\nu I + A^T) = (\omega_\nu I - A)\, a^{\nu+\frac{1}{2}} + f, \tag{11}$$

where I is the identity matrix and ω_ν is the iteration parameter for the νth stage.

Each of the two steps in the single iteration described by Eqs. (10) and (11) is fairly simple to manage. To illustrate, consider the first equation. Clearly, it can be decomposed into $N+1$ separate systems, one for each value of m. To solve these systems, it is convenient to reintroduce $a_{N-1,m}$ and $a_{N,m}$ by means of Eq. (5) and then to manipulate Eq. (10) into the form (see, for instance, Haidvogel [9, pp. 137–140])

$$\omega_\nu\frac{c_{n-2}}{4n(n-1)}a^{\nu+\frac{1}{2}}_{n-2,m} + a^{\nu+\frac{1}{2}}_{n,m} - \omega_\nu\frac{1}{2(n^2-1)}a^{\nu+\frac{1}{2}}_{n,m} + \omega_\nu\frac{1}{4n(n+1)}a^{\nu+\frac{1}{2}}_{n+2,m} \tag{12}$$

$$= \frac{c_{n-2}}{4n(n-1)}h^\nu_{n-2,m} - \frac{1}{2(n^2-1)}h^\nu_{n,m} + \frac{1}{4n(n+1)}h^\nu_{n+2,m} \quad (2 \leqslant n \leqslant N),$$

where

$$h^\nu = a^\nu(\omega_\nu I - A^T) + f \tag{13}$$

and it is understood that the last term on each side of Eq. (12) is dropped for $n \geqslant N-3$ and the next to last term for $n \geqslant N-1$.

Note that both here and in Eq. (5) the even and odd terms (in n) are uncoupled and thus they may be handled separately. The first step of the iteration, then, produces $2(N-1)$ systems in at most $(N+1)/2$ variables. Each of these systems consists of a

tridiagonal set of equations accompanied by a row of 1's arising from the boundary conditions. These quasitridiagonal systems, as we shall call them, are solved by a Gaussian elimination process tailored to their special structure. In practice, we have found pivoting to be unnecessary for this process. In addition, the LU factorization, independent of m, need be performed only once (for $m = 0$ for instance). Hence, the bulk of the work occurs in the $N + 1$ separate "backsolving" calculations. One complete iteration of the ADI process (comprising one sweep in each direction) requires approximately $2N^2$ divisions, $20N^2$ multiplications, and $18N^2$ additions. The solution of the quasitridiagonal equations themselves accounts for all the divisions, $6N^2$ multiplications, and $8N^2$ additions. The rest of the operations occur in the evaluation of h^ν and in the computation of the right-hand side of Eq. (12). Note that the matrix multiplication $a^\nu A^T$ occurring in Eq. (13) can be accomplished in $O(N^2)$ operations. This matrix expression corresponds to the second term in Eq. (4), from which it is apparent that the requisite sums can be evaluated recursively.

The final practical consideration is the choice of the iteration parameters ω_ν. As is the case for ADI applied to a finite-difference approximation of Poisson's equation (see, for instance, Isaacson and Keller [3, pp. 475–478]), the exact solution to the algebraic equation (7) can be obtained after $N - 1$ stages if ω_ν, $\nu = 0, 1, \ldots, N - 2$ are chosen to be the eigenvalues λ_ν of A. In fact, the exact ADI solution can be obtained in $\frac{1}{2}N$ iterations rather than $N - 1$. This is accomplished by using the even-odd uncoupling of the coefficients a_{nm} to factor Eq. (7) into four separate equations: n and m even; n odd, m even; n even, m odd; and n and m odd. Each of these equations can be solved exactly by using just the even or just the odd eigenvalues, as appropriate, as iteration parameters.

We also observe that an approximate ADI solution can be obtained by some other choice of the iteration parameters. The spectrum of A furnishes some hints for this. The smaller eigenvalues of A are clearly just good approximations to $\lambda_n = -\frac{1}{4}n^2\pi^2$, the eigenvalues of $\partial^2 u/\partial x^2 = \lambda u$ with $u(-1) = u(1) = 0$. At the other extreme, Gershgorin's theorem implies that the very largest eigenvalues increase as N^4. Empirically, we find that λ_{N-2} and λ_{N-3} are within a few percent of each other, with $\lambda_{N-2} \sim 0.303N^4$, while λ_{N-4} is a factor of 20 smaller. This suggests that a good approximate solution to Eq. (7) can be obtained by setting $\omega_0 = \lambda_{N-2}$, $\omega_1 = \lambda_{N-3}$ and then distributing the remaining ω_ν, $\nu = 2, 3, \ldots, I - 1$, between λ_{N-4} and λ_0, where a total number of I iterations of the ADI process are used. For the purpose of picking the $I - 2$ iteration parameters in the range spanned by λ_{N-4} and λ_0 we have used the procedure described by Varga [10], in which

$$\omega_\nu = \lambda_{N-4}(\lambda_0/\lambda_{N-4})^{(2\nu-3)/(2I-4)}, \qquad \nu = 2, 3, \ldots, I - 1. \tag{14}$$

In the special case $I - 2 = 2^k$, the Wachspress [11] parameters are another possible choice.

When the ADI scheme is used to obtain the exact solution of Eq. (7), some preprocessing is required in order to find the eigenvalues λ_n. However, the cost of this overhead is minor—less than 10% of the remaining cost of solving Eq. (7). If the

eigenvalues are stored and used for repeated solutions of Poisson's equation, this overhead is entirely trivial. It is even less for the approximate ADI method, for only the largest several eigenvalues are needed, and even they can be approximated well by a power series in $1/N$ (starting with N^4).

4. A Matrix Diagonalization Solution Technique

In this section we describe another method of solving the matrix equation (7). This alternative scheme is an order of magnitude more efficient than ADI, but it requires a large amount of preprocessing. As demonstrated by Murdock [12] and Haidvogel, Schulman, and Robinson [13], however, it is nicely suited to problems in which repeated solutions of Poisson's equation are required, as for instance, in time-dependent problems for which the forcing function f changes at each time step.

As has been demonstrated above, matrix equations of the form (10) or (11) can be rewritten as quasitridiagonal systems and inverted quite easily. As it stands, this is not true of Eq. (7) in which the x- and y- derivative terms are treated simultaneously rather than individually as in ADI. However, Eq. (7) can be cast in quasitridiagonal form by the following diagonalization procedure. Consider B, the $(N+1) \times (N+1)$ matrix of coefficients representing the Chebyshev approximation to $\partial^2/\partial x^2$ [Eq. (4), first term] with the homogeneous boundary conditions (5). Corresponding to B are $N-1$ nonzero eigenvalues λ_ν, $\nu = 0, 1, \ldots, N-2$, and the matrix e of eigenvector coefficients $e_{n\nu}$. Together, these satisfy the equation

$$Be = e\Lambda,$$

where Λ is the $(N-1) \times (N-1)$ matrix whose diagonal elements are the eigenvalues λ_ν. Both e and Λ can be calculated via standard eigensystem routines; we have used EISPACK [14] for this purpose. In terms of the $T_n(x)$, therefore,

$$V_\nu(x) = \sum_{n=0}^{N} e_{n\nu} T_n(x)$$

are the eigenfunctions of the discrete $\partial^2/\partial x^2$ operator.

The partial diagonalization of Eq. (7) is achieved by expanding $u(x, y)$ and $f(x, y)$ in the eigenfunctions $V_\nu(x)$:

$$u(x, y) = \sum_{n=0}^{N-2} \sum_{m=0}^{N} b_{nm} V_n(x)\, T_m(y)$$

and

$$f(x, y) = \sum_{n=0}^{N-2} \sum_{m=0}^{N} g_{nm} V_n(x)\, T_m(y).$$

The coefficients g_{nm} and f_{nm} $(0 \leqslant n \leqslant N-2,\ 0 \leqslant m \leqslant N)$ are related by

$$\tilde{g} = \hat{e}^{-1} \tilde{f}, \tag{16}$$

where $\tilde{g}$ and $\tilde{f}$ are the matrices of coefficients g_{nm} and f_{nm} and $\hat{e}^{-1}$ is the $(N-1) \times (N-1)$ inverse of the matrix $\hat{e}$ (e minus its last two rows). The loss of information associated with truncating the matrix of eigenvector coefficients e_{nm} at $n = N-2$ is a further result of the tau approximation.

In the new representation, the equation analogous to (7) is

$$\Lambda b + bB^T = \tilde{g}. \tag{17}$$

It is identical in form to either of the ADI sweeps—Eq. (10) or (11)—except that the matrix $\omega_\nu I$ is replaced by Λ. Thus, it too can be decomposed into $2(N-1)$ quasi-tridiagonal systems in at most $\frac{1}{2}(N+1)$ unknowns. The Chebyshev coefficients can be recovered from the solution of Eq. (7) by first performing the matrix multiplication

$$\tilde{a} = \hat{e}b, \tag{18}$$

where $\tilde{a}$ is the matrix a_{nm} ($0 \leqslant n \leqslant N-2$, $0 \leqslant m \leqslant N$), and then applying Eq. (6) to fix $a_{N,m}$ and $a_{N-1,m}$ for $0 \leqslant m \leqslant N$.

Apart from the preprocessing necessary to determine the λ_ν and $e_{n\nu}$ for a given N (and a given set of boundary conditions), this methodology provides quite an efficient way to determine the exact solution to problem (1). In comparison to ADI, for which $4(N-1)$ quasitridiagonal systems must be solved for each of the $\frac{1}{2}N$ iterations required for the exact (discrete) solution, the eigenvector technique requires only $2(N-1)$ solutions of quasitridiagonal systems. Thus, the matrix b can be found from Eq. (17) in $O(N^2)$ operations, a gain in efficiency of a factor N. This is partially offset by the overhead involved in computing the coefficient matrices $\tilde{g}$ and $\tilde{a}$ from Eqs. (16) and (18)—that is, in transforming spectral fields into the eigenvector representation and back again. These transformations require $O(N^3)$ operations, and hence account for most of the work in this method. For Poisson's equation with Dirichlet boundary conditions, the eigenfunctions $V_\nu(x)$, though functions of $T_n(x)$, involve only the odd or only the even polynomials. As a result, the components of $\hat{e}$ and $\hat{e}^{-1}$ are alternatingly zero and nonzero. This property can be used to speed up the matrix multiplications by a factor of 2. Even including these linear transformations, the matrix diagonalization technique appears to offer roughly an order of magnitude increase in efficiency over the ADI method (see Tables V and VI in the next section).

Some mention should be made of the necessary preprocessing step in which the eigenvalues and eigenvectors of the discrete approximation to $\partial^2/\partial x^2$ (subject to the usual boundary conditions) are determined. First, these operations are quite costly, requiring in excess of two orders of magnitude more time than the solution of Poisson's equation itself (Table VI). This technique, therefore, is efficient only for problems in which the number of required Poisson solutions is much larger than the number of preprocessing steps. Second, the accuracy of the resulting solution to the algebraic equations is limited by the accuracy of the preprocessing calculations. In particular, determining $\hat{e}^{-1}$ could conceivably be difficult for a poorly conditioned

matrix $\hat{e}$. For the problem at hand, we have in fact calculated the condition number K, given by

$$K^2 = \max_{i,j} \frac{\lambda_i(\hat{e}\,\hat{e}^T)}{\lambda_j(\hat{e}\,\hat{e}^T)}$$

for the matrix $\hat{e}$ [15]. For $N = 32$, $K \simeq 18$, indicating that the accuracy of our calculations should be limited only by computational round-off error. Indeed, the exact (discrete) solutions given by ADI and the matrix diagonalization method differ only in their least significant digits.

Finally, we note that a third way of solving Eq. (7) is by the method based on Schur decomposition devised by Bartels and Stewart [16]. Although the computational time of this algorithm turned out to be somewhat shorter than that of ADI, we encountered the loss of from three to five digits in numerous examples. These digits can be recovered through the use of iterative refinement, but the computational time then exceeds that for ADI (which does not appear to suffer from round-off error), at least for $N \leqslant 64$.

5. Results

We present here a comparison for three model problems between the Chebyshev expansion technique and both a second and a fourth-order finite-difference scheme. The second-order method is the usual five-point approximation to Poisson's equation, solved by means of the cyclic reduction routines of the NCAR package [17]. Fourth-order results were obtained from these by the method of deferred corrections [18] in which the result, say u_1, of the five-point scheme is improved by subtracting off the lowest-order error remaining in the solution. Thus, the corrected $u(x, y)$ is the solution (on the same grid) to the equation $\Delta u = f + (1/12)\, h^2(\Delta f - 2\partial^4 u_1/\partial x^2\, \partial y^2)$, where h is the mesh size and centered differences are used to approximate the derivative terms on the right-hand side.

For the finite-difference methods the maximum pointwise error on the grid between the exact and calculated solutions is our measure of error. For comparison, the Chebyshev coefficients a_{nm} obtained from the solution of Eq. (7) (by either the alternating direction or matrix diagonalization methods) were Fourier transformed according to Eq. (8), yielding the spectral solution at the points $(x_i, y_j) = (\cos \pi i/N, \cos \pi j/N)$. The maximum error at these discrete locations was then computed.

The first example—$f = -32\pi^2 \sin 4\pi x \sin 4\pi y$—tests the Chebyshev approximation to a moderately oscillatory but otherwise well-behaved (infinitely differentiable) solution—$u = \sin 4\pi x \sin 4\pi y$. This example will also serve to illustrate the convergence of the approximate iterative method for solving Eq. (7). The ADI process was begun with an initial guess a^0 obtained by Fourier transforming the known solution u by Eq. (9) and increasing those Chebyshev coefficients by 20%. Both complete ($I = \frac{1}{2}N$) and partial ($I = 6$, 10, 14, and 18) iterations were examined.

Consider first the complete iteration case, $I = \frac{1}{2}N$, for which the Chebyshev equation (7) is solved exactly. The entries in the last column of Table I for $N = 16$, 24, and 32 indicate the exponential convergence rate of the Chebyshev expansion. (The less impressive improvement in the last two rows presumably reflects the 14-digit single-precision accuracy of the machine on which these calculations were performed—a CDC Cyber 175.) This rate of convergence has its source in the coefficients of the exact solution: $a_{nm} = 4(c_n c_m)^{-1} J_n(4\pi) J_m(4\pi) \sin \frac{1}{2}n\pi \sin \frac{1}{2}m\pi$, where $J_n(z)$ denotes the Bessel function of the first kind. The properties of Bessel functions as the order n (or m) is increased imply that a_{nm} decrease exponentially fast.

TABLE I

Maximum absolute error of the Chebyshev approximation when $\Delta u = f = -32\pi^2 \sin 4\pi x \sin 4\pi y$, as a function of the number of modes N in each direction and the number of ADI iterations I.

N \ I	6	10	14	18	$\frac{1}{2}N$
16	3.26×10^{-2}				3.33×10^{-2}
24	3.92×10^{-3}	1.55×10^{-5}			6.89×10^{-6}
32	8.18×10^{-3}	6.06×10^{-5}	2.45×10^{-7}		4.77×10^{-11}
48	1.43×10^{-3}	3.19×10^{-4}	5.81×10^{-6}	4.74×10^{-8}	1.90×10^{-12}
64	1.22×10^{-3}	7.65×10^{-4}	1.13×10^{-5}	4.65×10^{-7}	8.67×10^{-13}

[a] The case $I = \frac{1}{2}N$ gives the exact solution to the spectral equations.

For the partial iterations, the parameters ω_ν were distributed according to Eq. (14). The results in Table I indicate that with an initial error of 20%, little if any savings can be made by using an approximate rather than an exact ADI iteration. (We note here that for an initial error of 100%, i.e., with $a^0 = 0$, the errors listed in Table I for the partial iterations would be no more than a factor of 5 larger). The approximate ADI method may be useful, however, when a good initial guess is available, as sometimes occurs when Poisson's equation is one part of a time-dependent calculation.

The comparison with the finite difference methods is reported in Table II. For the well-behaved solution of this first example, the Chebyshev approximation is markedly superior. This is illustrated most dramatically by the improvement between $N = 16$ and $N = 32$: the error of the second-order scheme is reduced by a factor of 4, that of the fourth-order scheme by a factor of 16, and that of the spectral method by a factor of 10^9.

As one would expect from the theory of spectral methods, this extremely rapid convergence is representative of the performance of the Chebyshev expansion for solutions which are infinitely differentiable. We have obtained equally impressive results for other analytic solutions such as six of the test problems considered by Houstis and Papatheodorou [19]. In all but one of those cases 16 Chebyshev polynomials in each direction already produced an approximate solution which is more

TABLE II

Maximum absolute error of the second-order finite difference (FD2), deferred correction (FD4), and spectral (alternating direction or matrix diagonalization: ADI/MD) approximations to $\Delta u = -32\pi^2 \sin(4\pi x) \sin(4\pi y)$ as a function of the number of degrees of freedom N in each direction.

N	FD2	FD4	ADI/MD
16	2.34×10^{-1}	2.81×10^{-2}	3.33×10^{-2}
32	5.30×10^{-2}	1.63×10^{-3}	4.77×10^{-11}
64	1.30×10^{-2}	9.97×10^{-5}	8.67×10^{-13}
128	3.22×10^{-3}	6.21×10^{-6}	2.00×10^{-12}

accurate than could be obtained from an $N = 2048$ second-order finite difference scheme.

We next tested the Chebyshev approximation to the problem $\Delta u = 1$, a case for which the solution u does not oscillate (in fact, it is monotonic, assuming its minimum value of $u \simeq -0.294$ at the origin and increasing toward zero at the boundaries) but does have some singular behavior. The irregularity occurs near each of the four corners, where the partial differential equation insists that $\Delta u = 1$ but the boundary conditions require that Δu vanish. This means that $u = O(r^2 \ln r)$ as $r \to 0$, where r is the distance from a corner [20]. The exact solution can be expressed as the doubly infinite cosine series

$$u(x, y) = \frac{64}{\pi^4} \sum_{\substack{n=1 \\ n \equiv 1 (\mathrm{mod}\ 2)}}^{\infty} \sum_{\substack{m=1 \\ m \equiv 1 (\mathrm{mod}\ 2)}}^{\infty} (-1)^{\frac{1}{2}(n+m)} \frac{\cos \frac{1}{2} n\pi x \cos \frac{1}{2} m\pi y}{nm(n^2 + m^2)}.$$

(The task of computing this sum numerically can be eased greatly by performing one of the sums by means of residue calculus.) The results for this example are presented in Table III. Note that both difference schemes exhibit second-order convergence, in contrast to the fourth-order behavior of the spectral approximation. The deterioration of the convergence rates of the last two methods is clearly a manifestation of the corner singularities. They invalidate the higher-order terms of the asymptotic expan-

TABLE III

Same as Table II except for the problem $\Delta u = 1$.

N	FD2	FD4	ADI/MD
16	9.02×10^{-4}	7.17×10^{-6}	3.52×10^{-5}
32	2.26×10^{-4}	1.79×10^{-6}	2.23×10^{-6}
64	5.67×10^{-5}	4.48×10^{-7}	1.40×10^{-7}
128	1.42×10^{-5}	1.12×10^{-7}	8.72×10^{-9}

sion in h^2 upon which the method of deferred corrections depends and they cause the Chebyshev coefficients to decay algebraically—roughly as N^{-6}—rather than exponentially. As one might expect, the maximum error of these two schemes occurs near the boundary. In fact, regardless of N, the maximum error of the deferred corrections' scheme occurs at $(x, y) = (1 - h, 1 - h)$ and that of the Chebyshev scheme at $(x, y) = (0, \cos \pi/N)$. The deferred corrections method is surprisingly accurate in this example. Apparently, this occurs because the solution has only a single, very gentle extremum. For a problem with both corner singularities and appreciable oscillations, however, the Chebyshev solution should be superior, not only to the second-order method but to the fourth-order one as well. Consider, for instance, the combination of the present right-hand side with that of the preceding example, i.e., $f = 1 + \sin 4\pi x \sin 4\pi y$. In the combined problem, an $N \geqslant 32$ Chebyshev expansion would be sufficient to insure a maximum pointwise error of less than 10^{-6}, whereas in excess of 128 grid points would be required for comparable accuracy of the deferred corrections solution.

As a final example we examine a problem for which the singularity occurs in the interior of the square rather than on the boundary. Consider the function $u(x, y) = v(x)\, v(y)$, where

$$v(x) = \begin{cases} -\frac{1}{4}(x + 1), & x < 0, \\ \frac{1}{2}x^2 - \frac{1}{4}x - \frac{1}{4}, & x \geqslant 0. \end{cases}$$

This solves $\Delta u = f = H(x)\, v(y) + v(x)\, H(y)$, where $H(x)$ is the unit step function centered at the origin. In this case, the second derivatives of the solution are discontinuous along both coordinate axes. The problem we actually address is one which mimics better the uncertainty that discrete methods have about the location of discontinuities in the source f. The step functions are centered at $x = y = \frac{1}{2}h$ [and $v(x)$ is adjusted accordingly] rather than at the origin; here, $h = 2/N$ for the finite difference methods and $h = \sin \pi/N$ for the Chebyshev approximation. The The results are given in Table IV. All three methods display second-order convergence, with the spectral method appearing to converge at a slightly faster rate. For this somewhat extreme example the spectral method offers no advantage over even the usual five-point difference approximation.

TABLE IV

Same as Table II except for the problem $\Delta u = H(x)v(y) + v(x)H(y)$.

N	FD2	FD4	ADI/MD
16	4.98×10^{-4}	1.67×10^{-4}	5.16×10^{-4}
32	1.23×10^{-4}	4.14×10^{-5}	1.26×10^{-4}
64	3.08×10^{-5}	1.03×10^{-5}	3.03×10^{-5}
128	7.66×10^{-6}	2.56×10^{-6}	7.39×10^{-6}

The computational times on a CDC Cyber 175 of the second-order finite-difference method—using the NCAR routine PWSCRT—and the Chebyshev series approximation—using the ADI method—are listed in Table V. Both codes were in FORTRAN. The corresponding execution times for the method of deferred corrections are roughly 220% of those for the second-order scheme. These times pertain just to the cost of solving the appropriate algebraic system of equations. The additional cost of performing two FFT's for the Chebyshev method—using an assembly language version of the one-dimensional FFT—ranges from 0.014 sec for $N = 16$ to 0.864 sec for $N = 128$.

Similar timings on a CDC 7600 for the NCAR code and the eigenvector method of solving Eq. (7) are given in Table VI. For this method the cost of the FFT's is comparable to that for solving the algebraic equations. We note that the preprocessing cost of the eigenvector method is quite substantial, e.g., greater than 8 sec for $N = 32$. However, these preprocessing times are likely higher than necessary since we took extra care to compute the eigenvectors accurately. If the Poisson's equation itself is not being solved to full machine accuracy, then cruder, and therefore less costly, eigenvectors will suffice.

TABLE V

Execution time in seconds for the solution of Poisson's equation on the Cyber 175 by second-order finite-difference (FD2) and the alternating direction spectral (ADI) methods as a function of the number of degrees of freedom N in each direction.

N	FD2	ADI
16	0.009	0.059
32	0.046	0.466
64	0.211	3.689
128	0.968	29.130

TABLE VI

Execution time in seconds for the solution of Poisson's equation on the CDC 7600 by second-order finite-difference (FD2) and matrix diagonalization spectral (MD) methods as a function of the number of degrees of freedom N in each direction.

N	FD2	MD
16	0.010	0.007 (2.7)
32	0.046	0.014 (8.1)
64	0.215	0.254 (77.5)

[a] Approximate preprocessing times associated with the matrix diagonalization technique are given in parentheses.

As both the general theory of the convergence of spectral methods suggests and the first example illustrates, whenever the solution u is infinitely differentiable, the Chebyshev expansion can achieve highly accurate solutions to Poisson's equation far more efficiently than the standard finite-difference methods. Note also that the Chebyshev expansion technique makes less demands upon computer storage since fewer degrees of freedom are needed. Moreover, for solutions whose only singularities are those induced by the corners of the domain the Chebyshev approximation apparently converges faster than finite-difference methods. The advantage of this spectral method is less clear cut, however, when solutions of low accuracy can be tolerated.

The Chebyshev expansion can be easily extended to handle both inhomogeneous boundary conditions and the Helmholtz equation $\Delta u + \Omega u = f$ for constant Ω. Moreover, it may also be applied to nonseparable problems of the type $\nabla \cdot [a(x, y) \nabla u] = f$ by using the iterative procedure devised by Concus and Golub [21] for finite-difference methods. The spectral method, however, has the added complication of requiring a convolution sum of the Chebyshev coefficients at each iteration. It is advisable to use the transform methods described by Orszag [22] to evaluate this term.

Spectral methods for Poisson's equation, of course, need not be based on Chebyshev polynomials. In some applications an expansion in Legendre polynomials may be more appropriate. This spectral method would also exhibit exponential convergence, but it suffers from the lack of a fast transform such as the FFT. As noted in the introduction trigonometric functions are not as appropriate since the boundary terms arising from the integration by parts estimate of the expansion coefficients do not necessarily vanish. As shown by Skollermo [23], sine series expansions generally exhibit second-order convergence although their convergence rate may be improved by using special solutions. These methods take about as long as the finite-difference methods since the FFT can be employed. Nevertheless, the Fourier series solution will, in general, have a finite order of convergence, even for infinitely differentiable solutions.

ACKNOWLEDGMENTS

The authors are grateful to Professor Steven Orszag for suggesting this approach to Poisson's equation and for his advice throughout this work. Some of the computer time was provided by the National Center for Atmospheric Research, which is supported by the National Science Foundation. The work performed by the first author was supported by the National Science Foundation under Contract Number ID076-00869. The second author was supported by NASA Contract Number NASA1-14101 while in residence at ICASE.

REFERENCES

1. F. W. DORR, *SIAM Rev.* **12** (1970), 248–263.
2. P. N. SWARZTRAUBER, *SIAM Rev.* **19** (1977), 490–501.

3. E. ISAACSON AND H. B. KELLER, "Analysis of Numerical Methods," Wiley, New York, 1966.
4. D. GOTTLIEB AND S. A. ORSZAG, "Numerical Analysis of Spectral Methods: Theory and Application," Society for Industrial and Applied Mathematics, Philadelphia, 1977.
5. S. A. ORSZAG, *J. Fluid Mech.* **50** (1971), 689–703.
6. C. LANCZOS, "Applied Analysis," Prentice–Hall, Englewood Cliffs, N.J., 1956.
7. J. W. COOLEY AND J. W. TUKEY, *Math. Comp.* **19** (1965), 297–301.
8. D. W. PEACEMAN AND H. H. RACHFORD, JR., *SIAM J.* **3** (1955), 28–41.
9. D. B. HAIDVOGEL, Quasigeostrophic regional and general circulation modelling: an efficient pseudospectral approximation technique, *in* "Computing Methods in Geophysical Mechanics" (R. P. Shaw, Ed.), Vol. 25, American Society of Mechanical Engineers, 1977.
10. R. S. VARGA, "Matrix Iterative Analysis," Prentice–Hall, Englewood Cliffs, N.J., 1962.
11. E. L. WACHSPRESS, *J. Soc. Indust. Appl. Math.* **10** (1962), 339–350.
12. J. W. MURDOCK, *AIAA J.* **15** (1977), 1167–1173.
13. D. B. HAIDVOGEL, E. E. SCHULMAN, AND A. R. ROBINSON, to appear in *Rept. Meteorol. Oceanogr.*
14. V. K. KLEMA, B. S. GARBOW, AND C. B. MOLER, "EISPACK, User's Information," Argonne National Laboratory, Applied Mathematics Division, 1973.
15. J. H. WILKINSON, "The Algebraic Eigenvalue Problem," Clarendon, Oxford, 1965.
16. R. H. BARTELS AND G. W. STEWART, *Comm. ACM* **15** (1972), 820–826.
17. P. N. SWARZTRAUBER AND R. A. SWEET, Efficient FORTRAN subprograms for the solution of elliptic partial differential equations, NCAR Technical Note, NCAR-TN/IA-109, 1975.
18. V. PEREYRA, *SIAM J. Numer. Anal.* **4** (1967), 508–533.
19. E. N. HOUSTIS AND T. S. PAPATHEODOROU, Comparison of fast direct methods for elliptic problems, *in* "Advances in Computer Methods for Partial Differential Equations—II" (R. Vichnevetsky, Ed.) IMACS (AICA), Brussels, 1977.
20. R. S. LEHMAN, *J. Math. Mech.* **8** (1959), 727–760.
21. P. CONCUS AND G. H. GOLUB, *SIAM J. Numer. Anal.* **10** (1973), 1103–1120.
22. S. A. ORSZAG, *Stud. Appl. Math.* **50** (1971), 293–327.
23. G. SKÖLLERMO, *Math. Comp.* **29** (1975), 697–711.