INTERNATIONAL JOURNAL FOR NUMERICAL METHODS IN FLUIDS, VOL. 7, 1159–1189 (1987)

SPECTRAL METHODS AND THEIR IMPLEMENTATION TO SOLUTION OF AERODYNAMIC AND FLUID MECHANIC PROBLEMS*

C. L. STREETT

NASA Langley Research Center, Hampton, Virginia 23665, U.S.A.

SUMMARY

Fundamental concepts underlying spectral collocation methods, especially pertaining to their use in the solution of partial differential equations, are outlined. Theoretical accuracy results are reviewed and compared with results from test problems. A number of practical aspects of the construction and use of spectral methods are detailed, along with several solution schemes which have found utility in applications of spectral methods to practical problems. Results from a few of the successful applications of spectral methods to problems of aerodynamic and fluid mechanic interest are then outlined, followed by a discussion of the problem areas in spectral methods and the current research under way to overcome these difficulties.

KEY WORDS Spectral Method Collocation Computational Fluid Dynamics

INTRODUCTION

Spectral methods may be viewed as an extreme development of the class of discretization schemes known by the generic name of the method of weighted residuals (MWR).[1] The key elements of the MWR are the trial functions (also called the expansion or approximating functions) and the test functions (also known as weight functions). The trial functions are used as the basis functions for a truncated series expansion of the solution, which, when substituted into the differential equation, produces the residual. The test functions are used to enforce the minimization of the residual.

The choice of trial function is what distinguishes the spectral methods from the finite-element and finite-difference methods. The trial functions for spectral methods are infinitely differentiable global functions. (Typically they are tensor products of the eigenfunctions of singular Strum–Liouville problems.) In the case of finite-element methods, the domain is divided into small elements and a trial function is specified in each element. The trial functions are thus local in character and well suited for handling complex geometries. The finite-difference trial functions are likewise local.

The choice of test function distinguishes between the Galerkin, collocation and tau approaches. In the Galerkin approach, the test functions are the same as the trial functions; whereas, in the collocation approach, the test functions are translated Dirac delta functions. In other words, the Galerkin approach is equivalent to a least-squares approximation, whereas the collocation approach requires the differential equation to be satisfied exactly at the collocation points. Spectral tau methods are close to Galerkin methods but they differ in the treatment of boundary conditions.

The collocation approach is the simplest of the MWR, and appears to have been first used by Slater[2] in his study of electronic energy bands in metals. A few years later, Barta[3] applied this

---

*Based on an invited lecture.

method to the problem of the torsion of a square prism. Frazer *et al.*[4] developed it as a general method for solving ordinary differential equations. They used a variety of trial functions and an arbitrary distribution of collocation points. The work of Lanczos[5] established for the first time that a proper choice of trial functions and distribution of collocation points is crucial to the accuracy of the solution. Perhaps he should be credited with laying down the foundation of the orthogonal collocation method. This method was revived by Clenshaw,[6] Clenshaw and Norton[7] and Wright.[8] These studies involved application of Chebyshev polynomial expansions to initial-value problems. Villadsen and Stewart[9] developed this method for boundary-value problems.

The earliest investigations of the spectral collocation method to partial differential equations were those of Kreiss and Oliger[10] (who called it the Fourier method) and Orszag[11] (who termed it pseudospectral). This approach is especially attractive because of the ease with which it can be applied to variable coefficient and even non-linear problems. The essential details will be furnished below.

The Galerkin approach is perhaps the most aesthetically pleasing of the MWR since the trial functions and the test functions are the same. Indeed, the first serious application of spectral methods to PDEs—that of Silberman[12] for meteorological modelling—used the Galerkin approach. However, spectral Galerkin methods only became practical for high-resolution calculations of non-linear problems after Orszag[13] and Eliasen *et al.*[14] developed a transform method for evaluating convolution sums arising from quadratic non-linearities. Even in this case, spectral collocation methods retain a factor of two in speed. For more complicated non-linear terms, high-resolution spectral Galerkin methods are still impractical.

The tau approach is the most difficult to rationalize within the context of the MWR. Lanczos[5] developed the spectral tau method as a modification of the Galerkin method for problems with non-periodic boundary conditions. Although it too is difficult to apply to non-linear problems, it has proven quite useful for constant-coefficient problems or subproblems, e.g. for semi-implicit time-stepping algorithms.

In the following discussion of the application of spectral methods to the solution of PDEs, rigorous mathematical detail will be avoided. Rather, emphasis will be given to the practical aspects of their implementation to solution of realistic problems. A few details on the fundamentals of spectral discretization will be given, and their performance on various test problems will be shown. Next, a number of solution schemes which have found utility in applications of spectral methods to practical problems will be discussed. Results from a few of the successful applications of spectral methods to problems of aerodynamic and fluid mechanics interest will then be outlined, followed by a discussion of the problem areas in spectral methods and the current research underway to ovecome these difficulties.

FUNDAMENTALS

In this section, some classical aspects of spectral methods are laid out; References 1 and 5 are used here as a guide.

Consider a function $U(x)$, where $U(x) \in C^m (m > 1)$ and is 2π-periodic. The trigonometric sum $U_n(x)$, which interpolates $U(x)$ on $2n$ evenly spaced points in $[0, 2\pi]$, is given by

$$U_n(x_j) = U(x_j)$$

for

$$x_j = \frac{\pi j}{n}, \quad 0 \leqslant j \leqslant 2n - 1, \tag{1}$$

$$U_n(x) = \frac{a_0}{2} + \sum_{k=1}^{n-1} (a_k \cos kx + b_k \sin kx) + \frac{a_n}{2} \cos nx,$$

where the coefficients are established by

$$\begin{aligned} a_k &= \frac{1}{n}\sum_{j=0}^{2n-1} U(x_j)\cos kx_j, \\ & \qquad\qquad\qquad\qquad 0 \leqslant k \leqslant n, \\ b_k &= \frac{1}{n}\sum_{j=0}^{2n-1} U(x_j)\sin kx_j. \end{aligned} \tag{2}$$

The formulae for the coefficients a_k and b_k are simple, since the functions $\cos kx$ and $\sin kx$ are discretely orthogonal on the $\{x_j\}$:

$$\begin{aligned} &\sum_{j=0}^{2n-1} \cos rx_j \cos sx_j = \begin{cases} 0, & r \neq s, \\ n, & n > r = s > 0, \\ 2n, & r = s = 0, n, \end{cases} \\ &\sum_{j=0}^{2n-1} \cos rx_j \sin sx_j = 0, \quad 0 \leqslant r, s \leqslant n, \\ &\sum_{j=0}^{2n-1} \sin rx_j \sin sx_j = \begin{cases} 0, & r \neq r, \\ n, & n > r = s > 0, \\ 2n, & r = s = 0, n. \end{cases} \end{aligned} \tag{3}$$

It can be shown by a simple integration-by-parts argument[15] that, for $m > 1$, the error in this approximation is

$$|U(x) - U_n(x)| = 0\left(\frac{1}{n^{m-1}}\right). \tag{4}$$

Note that the power of $1/n$ in equation (4), which is generally known in numerical methods as the 'order' of the fit, is not a constant inherent to the method, but is tied directly to the smoothness of the function being fit. This is different from the use of local polynomials to interpolate a function, where the order of the fit is related to the order of the polynomial used. (Finite-difference representations may be thought of as being derived from local polynomial interpolations.) Thus, for the trigonometric interpolation, if $U(x) \in C^\infty$, then $U_n(x)$ approaches $U(x)$ everywhere faster than any (finite) power of $1/n$. This is termed 'exponential' or 'infinite-order' accuracy; the latter term is unfortunate, since many confuse the term 'infinite-order accuracy' with 'infinite accuracy', which certainly is not implied.

Recall that for a function to be periodic and C^∞, the right and left limits of the function and every derivative exist and are equal at every point, including the endpoints of the period, which are essentially the same point. Thus, if a function is C^∞ in $[0, 2\pi]$, but only the functional values and no higher derivatives are equal at the endpoints, then the function considered as being 2π-periodic is only C^0. Convergence of a trigonometric interpolation, then, would be quite slow, e.g. less than first order. Thus, if an exponential-order interpolation of a function with arbitrary end conditions is required, the interpolation must be singular in some sense. For instance, consider a non-periodic function $V(y) \in C^m$ on $[-1, 1]$. Define a new independent variable

$$\phi \equiv \cos^{-1} y \tag{5}$$

and a new function $\tilde{V}(\phi)$ such that

$$\tilde{V}(\phi(y)) = V(y), \quad y \in [-1, 1]. \tag{6}$$

The function $\tilde{V}$ is then defined on $\phi \in [0, \pi]$. Equation (6) may be rewritten as

$$\tilde{V}(\phi) = V(\cos\phi), \quad \phi \in [0, \pi]. \tag{7}$$

Now

$$V'(\phi) = V'(\cos\phi)(-\sin\phi),$$
$$\tilde{V}'''(\phi) = V'''(\cos\phi)(-\sin^3\phi) + 3V''(\cos\phi)(\cos\phi\sin\phi) + V'(\cos\phi)(\sin\phi), \quad \text{etc.} \tag{8}$$

Therefore, if $V \in C^m$, then $\tilde{V}$ is C^m and half-periodic on $[0, \pi]$. A trigonometric interpolation $V_n(y)$ involving only 'cosine' terms is applicable to such a function; i.e.

$$V_n(y) = \frac{a_0}{2} + \sum_{j=1}^{n} a_j \cos(j\phi(y)) \tag{9}$$

$$= \frac{a_0}{2} + \sum_{j=1}^{n} a_j \cos(j \cos^{-1} y).$$

An error bound on $V_n(y)$ similar in form to equation (4) may be found.

The 'singularity' of this interpolation is seen in the distribution of interpolation points y_j corresponding to the evenly spaced ϕ_j:

$$y_j = \cos\frac{\pi j}{n}, \quad 0 \leqslant j \leqslant n. \tag{10}$$

Since $dy/d\phi = 0$ at $\phi = 0, \pi$, the node points y_j pack near the boundaries $y = -1$ and 1; this 'packing' serves to overpower the slow convergence of the interpolation near the boundaries.

The interpolation $V_n(y)$ generated by the change of variables artifice in equation (6) may be interpreted differently. The points y_j from equation (4) are seen to be the zeros of the $(n+1)$th Chebyshev polynomial $T_{n+1}(y)$, where the Chebyshev polynomials are defined by

$$T_n(y) = 2^{-(n-1)} \cos(n \cos^{-1} y). \tag{11}$$

Thus, the trigonometric interpolation from equation (9) is identical with an interpolatory series in Chebyshev polynomials. This shows that a Chebyshev series also exhibits exponential convergence.

The construction of an approximation series may be looked at from another viewpoint which is useful in the utilization of other basis functions, i.e. Legendre polynomials. Represent an n-term approximating series for a function $W(z)$ as

$$W_n(z) = \sum_{j=1}^{n} c_j B_j(z), \tag{12}$$

where the $\{B_j(z)\}$ are the set of basis functions and the c_j the series coefficients representing $W(z)$. The approximating series may be defined by requiring that the L_2-norm of the error, with weight function $w(z)$,

$$\| W(z) - W_n(z) \|_2^2 \equiv \int_a^b w(z)(W(z) - W_n(z))^2 \, dz, \tag{13}$$

be a minimum. For a particular choice of $\{B_j(z)\}$, define the function of $\{\hat{c}_j\}$:

$$\Psi(\hat{c}_1, \hat{c}_2, \ldots, \hat{c}_n) = \int_a^b \left(w(z) W(z) - \sum_{i=1}^{n} \hat{c}_i B_i(z) \right)^2 dz. \tag{14}$$

Minimization of Ψ with respect to $\{\hat{c}_j\}$ then yields the desired coefficients $\{c_j\}$. Now, at the minimum of Ψ, the coefficients $\{\hat{c}_j\}$ must satisfy

$$\left.\frac{\partial \Psi}{\partial \hat{c}_j}\right|_{\hat{c}_j=c_j} = 0, \quad j=1,2,\ldots,n, \tag{15}$$

which yields

$$\sum_{i=1}^{n} c_i \int_a^b w(z)B_i(z)B_j(z)\,\mathrm{d}z = \int_a^b w(z)W(z)B_j(z)\,\mathrm{d}z, \quad j=1,2,\ldots,n. \tag{16}$$

If the basis functions $\{B_i\}$ are orthogonal with respect to this inner product, e.g.

$$\int_a^b w(z)B_i(z)B_j(z)\,\mathrm{d}z = 0, \quad i \neq j, \tag{17}$$

then

$$c_j = \int_a^b w(z)W(z)B_j(z)\,\mathrm{d}z \Bigg/ \int_a^b w(z)B_j^2(z)\,\mathrm{d}z, \quad j=1,2,\ldots,n. \tag{18}$$

For this approximation-series construction technique to be useful in this context, the integrals in equation (18) above must be approximated by finite weighted sums of values of the integrand at specific collocation points. The Gauss–Lobatto form[16] is usually used, as this form includes the endpoints of the interval in the weighted sum. Sobolev space arguments confirm exponential-order convergence for these approximation series, for Chebyshev or Legendre basis functions.[17]

In order to use the above series representations in the context of solution of differential equations, an operator which approximates differentiation must be developed. Represent schematically an interpolation series for a function $W(z)$ as before as

$$W_n(z) = \sum_{j=1}^{n} c_j B_j(z), \tag{19}$$

where the $\{B_j(z)\}$ are the set of basis functions and the c_j the series coefficients representing W. If $W(z)$ is periodic, then sines and cosines are natural basis functions; if general end conditions are to be met, then the Chebyshev polynomials may be used as basis functions. The coefficients c_j are evaluated from convolution sums in either case; these sums involve the values of W at the set of interpolation or 'collocation' points, $\{z_j\}$. Note then that

$$W_n'(z) = \sum_{j=1}^{n} c_j B_j'(z), \tag{20}$$

i.e. that the derivative of the interpolation series involves the derivatives of the basis functions; the coefficients c_j are not affected.

The error bound $|W'(z) - W_n'(z)|$ is obvious from equation 4; indeed, if $W(z)$ is analytic, then the series $W_n'(z)$ will exhibit exponential-order convergence to $W'(z)$.

A specific form for the derivatives of the basis functions $B_j'(z)$ must be found. To accomplish this, represent $B_j'(z)$ as a series

$$B_j'(z) = \sum_{k=1}^{n} d_{jk} B_k(z). \tag{21}$$

The coefficients d_{jk} may be considered as elements of an operator D which approximates the derivative of the set of basis functions as a series of the basis functions. Therefore.

$$W_n'(z) = \sum_{j=1}^{n} c_j \sum_{k=1}^{n} d_{jk} B_k(z). \tag{22}$$

This may be rewritten as

$$W'_n(z) = \sum_{l=1}^{n} e_l B_l(z) \tag{23}$$

by manipulating the nested series; the e_l are thus the coefficients of the series in $B_l(z)$ representing $W'(z)$.

The techniques for utilizing these high-order representations of a function and its derivative in numerical solution of a differential equation were briefly described in the Introduction. The collocation method, wherein the approximation polynomial is required to satisfy the differential equation exactly at the collocation points (the quadrature points of the discrete approximation to equation (18), is the most often used for a number of reasons. First, non-linearities are formulated in a straightforward manner, using the function and its derivatives directly in the 'physical' space rather than formulated in difficult convolution sums in the 'coefficient' space. Similarly, boundary conditions may be directly and straightforwardly implemented, as the usual choice of collocation points includes the boundary points at which conditions must be set. Second, a number of efficient solution schemes are suggested by the pointwise physical space representation, as will be discussed in the fourth section. Finally, model problem results indicate an accuracy advantage of collocation over the Galerkin and tau methods. For these reasons, the following discussions will be restricted to the implementation of the collocation technique.

MODEL PROBLEM RESULTS

In this section, a number of model problems are discretized using spectral collocation techniques; comparison with corresponding finite-difference results shows the accuracy of this method.

The first example is provided by the model problem

$$u_t + u_x = 0, \tag{24}$$

with periodic boundary conditions on $[0, 2\pi]$ and the initial condition

$$u(x, 0) = \sin(\pi \cos x).$$

The exact solution is

$$u(x, t) = \sin[\pi \cos(x - t)\pi]. \tag{25}$$

Equation (24) was discretized in the x-direction using Fourier series, which is a natural choice for a periodic problem, and using second- and fourth-order finite differences. The time-advancement scheme was a classical fourth-order Runge–Kutta; the time step was taken so small for these comparisons that time-discretization errors were negligible. Shown in Table I are the maximum

Table I. Maximum error for a 1-D periodic problem

N	Fourier spectral	Second-order finite difference	Fourth-order finite difference
8	1·62 (−1)	1·11 (0)	9·62 (−1)
16	4·97 (−4)	6·13 (−1)	2·36 (−1)
32	1·03 (−11)	1·99 (−1)	2·67 (−2)
64	9·55 (−12)	5·42 (−2)	1·85 (−3)
128		1·37 (−2)	1·18 (−4)

errors at $t = 1$ for the various discretization methods. Because the solution is infinitely smooth, the error estimate of the Fundamentals section indicates that the spectral method should converge faster than any power of $1/N$; indeed, this behaviour is seen. Note that the error for the $N = 64$ spectral result is so small that it is affected by the machine round-off error (single precision on a CDC CYBER 175). In most practical applications, the benefit of the spectral method is not the extraordinary accuracy available for large N, but rather the small size of N necessary for a moderately accurate solution.

To illustrate the use of Chebyshev spectral discretization, boundary and initial conditions were chosen so that

$$u(x,t) = \sin \alpha\pi(x - t) \tag{26}$$

was a solution of equation (24). Time discretization was the same as in the first problem. Shown in Table II are the results for Chebyshev and second-order finite-difference discretizations. Also included is discretization by Fourier series, which is inappropriate for this non-periodic problem. The Chebyshev spectral method shows the expected exponential-order convergence, where the error of the Fourier discretization remains $O(1)$.

The final model problem example is that of Laplace's equation

$$u_{xx} + u_{yy} = 0 \tag{27}$$

on $[-1,1] \times [-1,1]$, with boundary conditions

$$u(x,-1) = u(x,1) = u(-1,y) = 0, \quad u(1,y) = \cos\frac{\pi y}{2}.$$

Shown in Table III are the maximum errors from Chebyshev and second-order finite-difference discretizations of this problem. Again, the exponential-order convergence of the spectral method is apparent.

Table II. Maximum error for a 1-D Dirichlet problem

N	Chebyshev spectral	Fourier spectral	Finite difference
4	1·49 (0)	1·85 (0)	1·64 (0)
8	6·92 (−1)	1·92 (0)	1·73 (0)
16	1·50 (−4)	2·27 (0)	1·23 (0)
32	3·45 (−11)	2·28 (0)	3·34 (−1)
64	9·55 (−11)	2·27 (0)	8·44 (−2)

Table III. Maximum error for Laplace equation problem

N	Chebyshev spectral	Finite difference
7	1·35 (−5)	9·15 (−3)
9	8·84 (−8)	5·54 (−3)
11	4·02 (−10)	3·48 (−3)
13	1·31 (−12)	1·50 (−3)
16	1·31 (−12)	1·10 (−3)
20		6·81 (−4)

SOLUTION SCHEMES FOR SPECTRAL DISCRETIZATIONS

Efficient solution of the equations from spectral discretizations is essential if the high-order accuracy of these methods is to be realized in practice. Direct solution of these equations is rarely feasible, thus iterative techniques are required. As a consequence of the ability of spectral methods to capture accurately a wide bandwidth of information, the eigenvalue spread of a spectral operator is large, typically of the order of the square of a corresponding low-order finite-difference operator. Therefore, an explicit method would converge quite slowly and the convergence rate would deteriorate very rapidly with mesh refinement. Implicit methods are not efficient, since spectral operators are full rather than banded as in the case of finite difference.

An alternative to pure explicit iterative schemes is a preconditioned scheme, where explicit iteration of some form is driven not by the spectral residual, but by the residual obtained after some processing is applied to reduce its eigenvalue spread. As will be shown later, this process can also be thought of as an approximate implicit scheme, which can give some insight into a relevant preconditioning operator.

Preconditioned iteration schemes for spectral collocation discretization of second-order equations are well known and proven. For instance, the time-accurate incompressible Navier–Stokes simulations of certain fluid mechanical phenomena, in which Chebyshev collocation is used in two co-ordinate directions, require the solution of a number of Helmholtz or Poisson equations per time step.[18] The preconditioning operator for these second-order equations is the low-order central finite-difference operator, using the Chebyshev collocation points as its mesh. Orszag[19] originally suggested such preconditioning and provided some analysis for the case of Fourier discretization. Included here are some results which indicate the efficiency of this technique for Chebyshev discretizations.

The steady-state compressible Euler or Navier–Stokes equations at high Reynolds numbers, however, are advection dominated. A spectral solution technique for such equations must therefore deal with operators which are predominantly first order. Application of preconditioning equivalent to that used for second derivatives will not work, however. Elementary analysis of preconditioning first-order Fourier discretization with central finite difference on the collocation mesh indicates that the convergence rate of such a scheme would be unacceptable; the eigenvalue spread of this preconditioned operator is unbounded. A similar situation appears in the case of Chebyshev discretization, as will be shown here.

A preconditioning scheme for first-order Chebyshev collocation operators is proposed herein, in which the central finite-difference mesh is finer than the collocation mesh.[20] Details of the proper techinques for transferring information between the meshes are given, and the scheme is analysed by examination of the eigenvalue spectra of the preconditioned operators.

Consider the following linear equation:

$$L_{\mathrm{sp}}u = f \tag{28}$$

where the operator L_{sp} is derived from spectral collocation of a differential equation. An iterative scheme is to be used to solve this equation. Given a current estimate of the solution u^n at iterate 'n', a simple Richardson iteration scheme for computing a better estimate u^{n+1} is

$$u^{n+1} = u^n - \omega(L_{\mathrm{sp}}u^n - f), \tag{29}$$

where the scalar relaxation factor ω may be chosen either experimentally or via a requirement that some norm of the residual

$$R^n = L_{\mathrm{sp}}u^n - f \tag{30}$$

be minimized at each step. Rewriting the scheme (29) as

$$u^{n+1} - u^n = \Delta u^n = -\omega R^n \tag{31}$$

shows the explicit nature of this iteration. Preconditioning involves choosing an operator M which is more easily invertible than L_{sp}, and for which the scheme

$$\Delta u^n = -\omega M^{-1} R^n \tag{32}$$

converges more rapidly than the scheme (31). The convergence rate of such schemes is quantified in the following way. Expand (32) as

$$u^{n+1} = (I - \omega M^{-1} L_{sp}) u^n - \omega M^{-1} f, \tag{33}$$

where I is the identity operator. Subtracting the discrete solution to equation (28) (the desired solution u) from both sides of (33), and adding to the right-hand side

$$\omega M^{-1} (L_{sp} u - f), \tag{34}$$

which is equal to zero by (28), yields an equation for the discrete error:

$$(u^{n+1} - u) = (I - \omega M^{-1} L_{sp})(u^n - u). \tag{35}$$

For the preconditioner scheme (32) to be convergent, all norms of $u^n - u$ must decrease; thus,

$$\| I = \omega M^{-1} L_{sp} \| < 1 \tag{36}$$

is required. Given a preconditioning operator M, the relaxation factor ω is used to satisfy equation (36).

From equation (36) may be seen what constitutes an effective preconditioner: if the eignevalues of $M^{-1} L_{sp}$ are clustered in a unit circle centred at 1 in the complex plane, then all error components will converge at nearly the same rate, and an optimal ω may be chosen which will yield rapid convergence. A slow scheme is characterized by a wide spread of eigenvalues from such a clustered pattern.

As an aside, a preconditioned iteration scheme may be looked at as an approximate implicit scheme in the following way. Ideally, given an estimate u^n, one wants the residual at the next iterate to be zero; expand

$$R^{n+1} = L_{sp}(u^n + \Delta u^n) - f = 0 \tag{37}$$

or

$$L_{sp} \Delta u = -R^n. \tag{38}$$

However, L_{sp} is difficult to invert; approximate it on the left by a more easily inverted operator $\omega^{-1} M$, giving

$$\Delta u^n = -\omega M^{-1} R^n, \tag{39}$$

which is identical to equation (32). Thus, the better $\omega^{-1} M$ approximates L_{sp}, the faster the scheme will converge.

The inverse of the preconditioning operator may be obtained and used directly, as implied by equation (32), or the preconditioning equation

$$M \Delta u^n = -\omega R^n \tag{40}$$

may be inverted iteratively. This procedure is beneficial when the spectral operator L_{sp} is non-linear, necessitating a different operator M at each iterate, or when M is still expensive to invert. An example of the latter case is when equation (28) results from spectral discretization of a PDE in two or more dimensions. The corresponding finite-difference discretization operator works well as a preconditioner (with the modifications described in the next section for first-order operators), but may still be expensive to invert due to its size. An iterative technique for computing the Δu^n is more efficient than direct inversion for this case, using, for instance, an approximate-factorization scheme (equation (32)) to converge rapidly.

As stated earlier, preconditioning of second-order equations is relatively easy. Use of second-order central finite-difference operators, defined directly on the collocation points using unequal mesh formulae, gives excellent performance for iterative solution. The eigenvalues of the preconditioned operator remain real and are positive and bounded with the number of modes. To illustrate this, the maximum eigenvalues of the Chebyshev second-order operator with Dirichlet boundary conditions imposed, with and without this finite-difference preconditioning, are shown in Table IV. Note that as this norm of L_{sp} alone goes like N^4, it remains bounded for the preconditioned operator.

This fortuitous circumstance does not carry over for first-order operators, however. A simple example will show why preconditioning using central finite differences on the collocation mesh is inadequate. For the model scalar problem $U_x = f$ with periodic boundary conditions on $[0, 2\pi]$, the eigenvalues of the Fourier collocation operator $\partial/\partial x$ are $i\alpha\Delta x$, where α is the wavenumber and Δx is the constant collocation mesh spacing. The product $\alpha\Delta x$ falls in the range $0 < |\alpha\Delta x| < \pi$. The corresponding eigenvalues for the central difference operator are $i\cdot\sin(\alpha\Delta x)$. Note that as $\alpha\Delta x \rightarrow \pi$, the ratio of these eigenvalues is unbounded. This ratio corresponds to the eigenvalues of the preconditioned operator, denoted $M^{-1}L_{sp}$ in the previous section. Such a preconditioned scheme is thus unconditionally unstable, with unbounded growth of the highest-wavenumber error components. The use of a finite preconditioning grid averts this unbounded component by introducing some natural dissipation in the preconditioner at the highest wavenumber of the spectral operator.

Because of the difference between the spectral and preconditioner meshes, a scheme is needed for transferring information between them. In multigrid terminology, a prolongation operator is needed to transfer the spectral residual, which acts as the source term of the preconditioning equation (40), and a restriction operator is needed to compute the solution updates on the spectral mesh. Naturally, it is desirable to transfer as much information as possible from the spectral residual to the preconditioning equation; spectral interpolation is therefore used for the prolongation operator. In the restriction operation, however, aliasing of correction components, with wavenumber higher than that of the spectral mesh, must be avoided. Therefore, spectral restriction cannot be used; low-order Lagrange interpolation is used here.

The preconditioning scheme proposed here proceeds as follows for the spectral discretization of a simple model problem $U_x = f$ At each iterate, compute the residual (equation (30)) defined on

Table IV. Maximum eigenvalues of second-order Chebyshev–Dirichlet operator, without and with preconditioning

N	$\rho(L_{sp})$	$\rho(M^{-1}L_{sp})$
9	$0{\cdot}2144 \times 10^3$	2·1314
17	$0{\cdot}3175 \times 10^4$	2·3058
33	$0{\cdot}4994 \times 10^5$	2·3884
45	$0{\cdot}1780 \times 10^6$	2·4103

N_{sp} points in the domain. This information is transferred to the (finer) preconditioning mesh via the spectral interpolation operator $I_{\mathrm{sp}}^{\mathrm{FD}}$. Denoting

$$\tilde{R} = I_{\mathrm{sp}}^{\mathrm{FD}} R, \tag{41}$$

the preconditioning equation becomes

$$\tilde{M}\,\Delta u = \tilde{R}, \tag{42}$$

where Δu is the update on the preconditioning mesh. For the model problem $U_x = f$ in $x \in [-1, 1]$ considered here, the preconditioner $\tilde{M}$ must approximate the first derivative operator via central finite differences. One possibility is

$$\tilde{M} = \delta_x^0 - \varepsilon \delta_{xx}^0 \tag{43}$$

using standard divided-difference notation. The second-difference term is required to avert the odd–even uncoupling of the pure central first-difference operator and to aid in its inversion. Some type of artificial viscosity is essentially always required when inverting this type of operator. Note that this artificial viscosity does not affect the spectral solution, being confined to the preconditioner.

The final step of this preconditioning scheme is to carry the update information to the spectral mesh. The iterate thus concludes with

$$\Delta u = \tilde{I}_{\mathrm{FD}}^{\mathrm{sp}}\,\Delta u$$

and

$$u^{n+1} = u^n + \omega\,\Delta u, \tag{44}$$

where the operator $\tilde{I}_{\mathrm{FD}}^{\mathrm{sp}}$ uses low-order Lagrange interpolation. For analysis purposes, the above sequence may be collected into a single operator

$$M^{-1} = \tilde{I}_{\mathrm{FD}}^{\mathrm{sp}} \tilde{M}^{-1} I_{\mathrm{sp}}^{\mathrm{FD}}. \tag{45}$$

This is, in effect, the preconditioning operator applied in the above scheme.

In practice, on a real problem of interest, the finite-difference operator M is also too large to invert directly at reasonable cost. Preconditioning operators are therefore inverted iteratively in all but the simplest one-dimensional test problems. Experience with both first- and second-order operators indicates that full machine-zero convergence of this iterative inversion is not necessary; depending on the nature of the iterative scheme, only two-orders-of-magnitude reduction in the residual of the preconditioner is necessary to obtain convergence of the overall scheme.

To investigate the effectiveness of the above preconditioning scheme, the eigenvalues of the operator $M^{-1} L_{\mathrm{sp}}$ were computed for a number of combinations of N_{sp} and N_{FD} for the model problem

$$U_x = \sin \pi x, \quad x \in [-1, 1], \tag{46}$$

$$U(-1) = 0.$$

Some basic characteristics of the eigenvalue spectra for this problem are as follows. The eigenvalues of spectral operator L_{sp} alone are dominantly imaginary, with the magnitude of the largest growing as N_{sp}^2. When the spectral operator is preconditioned with the central finite-difference operator on the same mesh, the eigenvalues of the overall operator become dominantly real, which is a desirable feature for use in an iterative scheme. The real parts of these eigenvalues are all positive, the smallest being near 1 and the largest of the order of 50 for the grids investigated. The eigenvector associated with this largest eigenvalue is highly oscillatory, as expected from the discussion in the previous section. As N_{FD} is increased relative to N_{sp}, this eigenvalue pattern generally collapses onto the point (1, 0) in the complex plane, with both real and imaginary parts decreasing. This collapse is at first rapid as N_{FD} is increased from N_{sp}, then slows as the eigenvalue with the largest real part moves to the interior of the unit circle centred at 1. There is a single exception to this clustered pattern, that of a small ($O(0{\cdot}1)$) eigenvalue which remains essentially fixed for all $N_{\mathrm{FD}} < N_{\mathrm{sp}}$. Its eigenvector indicates that this eigenvalue is associated with the different conditions used at $x = 1$ by M and L_{sp}. In general, though, the eigenvalues of the preconditioned operator are strongly clustered; thus one would expect rapid convergence of the iterative scheme.

In Table V are shown the maximum and minimum eigenvalues for the preconditioned operator from the test problem described above, for various combinations of N_{sp} and N_{FD}. The range of N_{sp} considered covers what we expect to be required in spectral discretizations of practical aerodynamic problems. As can be seen, the maximum eigenvalue is large when $N_{\mathrm{FD}} = N_{\mathrm{sp}}$ and drops rapidly as N_{FD} is increased. The operator appears well conditioned; that is, the eigenvalues are tightly grouped inside the unit circle centred at 1, when $N_{\mathrm{FD}} \geqslant 1{\cdot}5N_{\mathrm{sp}}$ for all of the grids considered. The spectra for $N_{\mathrm{FD}} = N_{\mathrm{sp}}$, $1{\cdot}2N_{\mathrm{sp}}$ and $1{\cdot}5N_{\mathrm{sp}}$ are also shown in Figures 1 and 2 for $N_{\mathrm{sp}} = 45$ and 60. These results were produced with first-order Lagrange interpolation in the restriction operator and with the artificial viscosity coefficient used in M fixed at 10^{-3}. Virtually identical results were obtained for $\varepsilon = 10^{-4}$ and for second-order restriction.

Convergence of preconditioned iteration schemes may be accelerated by a number of means. One powerful technique is multigrid. This technique has been extensively developed for finite-difference and finite-element discretizations[21] and has recently been applied to spectral discretizations.[22-24] Briefly put, multigrid methods take advantage of a property shared by a wide variety of relaxation schemes—potential efficient reduction of the high-frequency error components, but unavoidable slow reduction of the low-frequency components.

Let us consider just the interplay between two grids. A general, non-linear fine-grid problem can be written

Table V. Maximum and minimum eigenvalues for various N_{sp}, N_{FD} preconditioning equation (46) (real, imaginary parts)

N_{sp}	N_{FD}	$\lambda_{\max}$		$\lambda_{\min}$	
9	9	21·48,	0·0	1,	0·0
	14	1.208,	±0·046	0·339,	±0·310
	18	1·054,	±0·018	0·173,	0·0
17	17	47·46,	0·0	1	0·0
	21	2·233,	0·0	0·276,	0·0
	26	1·364,	0·0	0·218,	0·0
	34	1·136,	0·0	0·082,	0·0
33	33	56·28,	0·0	1,	0·0
	40	2·178,	0·0	0·097,	0·0
	55	1·365,	±0·013	0·056,	0·0
	66	1·348,	0·0	0·049,	0·0
45	45	37·42,	0·0	1,	±0·002
	54	2·456,	0·0	0·073,	0·0
	68	1·53,	0·0	0·049,	0·0
	75	1·390,	±0·062	0·045,	0·0

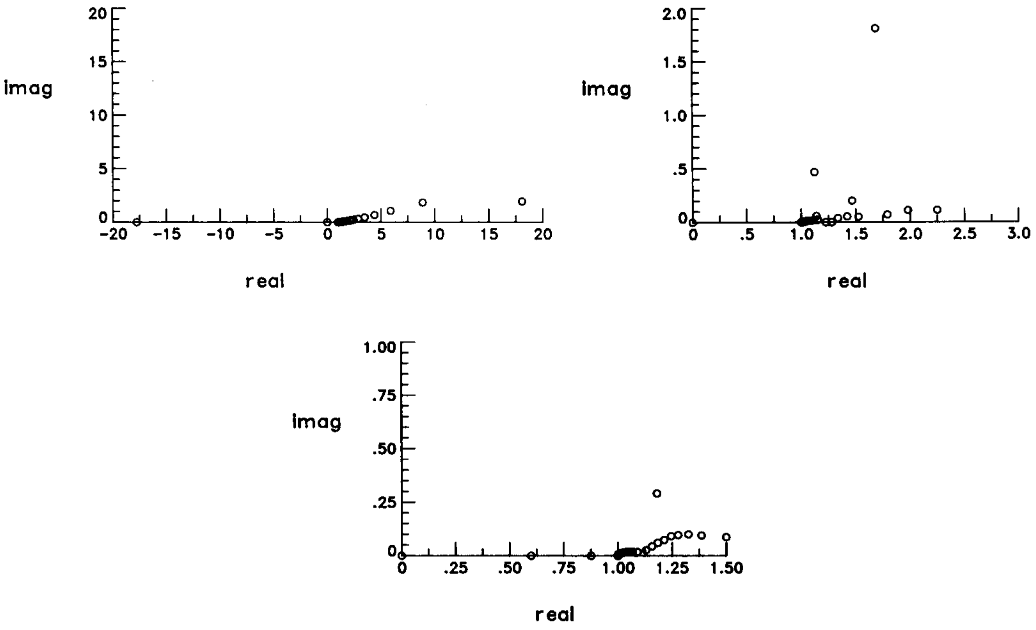

Figure 1. Eigenvalue spectrum of Chebyshev first-order operator, preconditioned with central finite differences, $N_{\text{sp}} = 45$; (a) $N_{\text{FD}} = N_{\text{sp}}$, (b) $N_{\text{FD}} = 1{\cdot}2\, N_{\text{sp}}$, (c) $N_{\text{FD}} = 1{\cdot}5\, N_{\text{sp}}$

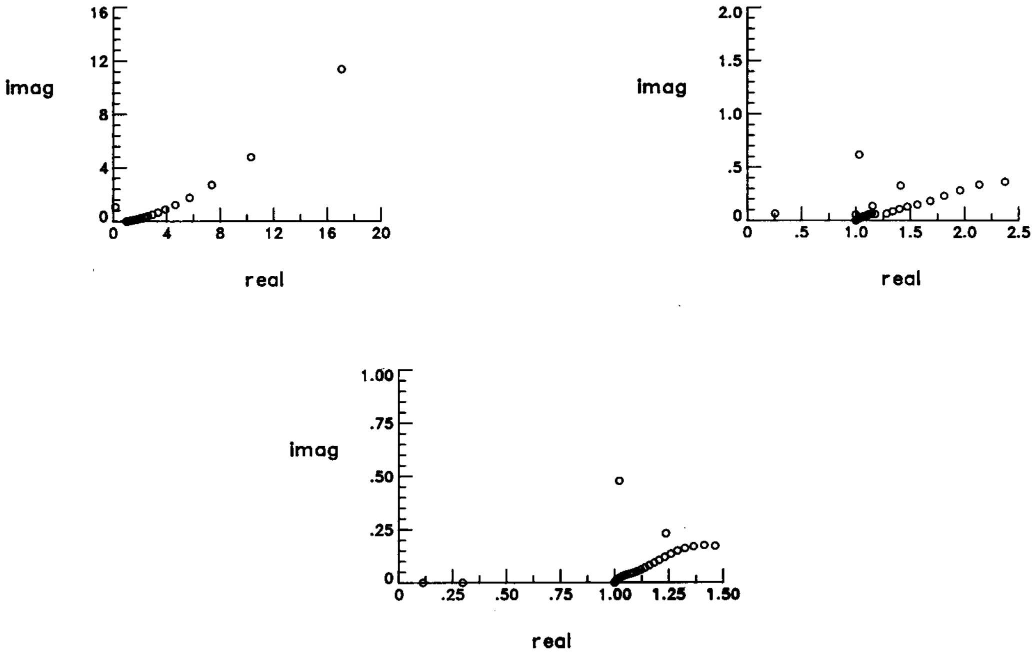

Figure 2. Eigenvalue spectrum of Chebyshev first-order operator, preconditioned with central finite differences, $N_{\text{sp}} = 60$; (a) $N_{\text{FD}} = N_{\text{sp}}$, (b) $N_{\text{FD}} = 1{\cdot}2\, N_{\text{sp}}$, (c) $N_{\text{FD}} = 1{\cdot}5\, N_{\text{sp}}$

$$L^{\mathrm{f}}(U^{\mathrm{f}}) = F^{\mathrm{f}}. \tag{47}$$

The shift to the coarse grid occurs 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 $V^{\mathrm{f}} - U^{\mathrm{f}}$ has been sufficiently reduced. The related coarse-grid problem is

$$L^{\mathrm{c}}(U^{\mathrm{c}}) = F^{\mathrm{c}}, \tag{48}$$

where

$$F^{\mathrm{c}} = R(F^{\mathrm{f}} - L^{\mathrm{f}}(V^{\mathrm{f}})) + L^{\mathrm{c}}(RV^{\mathrm{f}}). \tag{49}$$

The restriction operator R interpolates a function from the fine grid to the coarse grid. The coarse-grid operator and solution 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 corrected via

$$V^{\mathrm{f}} \leftarrow V^{\mathrm{f}} + P(V^{\mathrm{c}} - RV^{\mathrm{f}}). \tag{50}$$

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

A complete multigrid algorithm requires specific choices of the interpolation operators, the coarse-grid operators and the relaxation schemes. These issues are discussed at length in References 22–24 for both Fourier and Chebyshev multigrid methods. Numerous linear, variable-coefficient examples are also provided there. What basically has been concluded about spectral multigrid techniques is that extraordinary acceleration of convergence may be achieved, even for difficult problems,[24] but at the expense of fairly high overhead in the algorithm. Also, spectral multigrid tends to be more sensitive to the proper transfer of information between grids than for a corresponding finite-difference multigrid algorithm, and non-Dirichlet boundary conditions have a serious impact on convergence.

A somewhat more robust acceleration technique is based on dynamically choosing an 'optimum' relaxation parameter, e.g. ω in equation (29). Two techniques have been used with considerable success. The first is based on requiring that the residual at the $(n+1)$th time level be minimum with respect to an L_2 inner product. This leads to a preconditioned scheme of the form

$$\Delta u^{n+1} = -M^{-1}(L_{\mathrm{sp}}u^n - f) = M^{-1}R^n, \tag{51}$$

$$u^{n+1} = n^n + \omega^{n+1}\,\Delta u^{n+1}, \tag{52}$$

where

$$\omega^{n+1} = \omega^{n+1}_{\mathrm{MR}} = (R^n, L_{\mathrm{sp}}\,\Delta u^{n+1})(L_{\mathrm{sp}}\,\Delta u^{n+1}, L_{\mathrm{sp}}\,\Delta u^{n+1}). \tag{53}$$

A second choice for ω^{n+1} comes from requiring that the residual at the new time level $(n+1)$ be orthogonal to R^n under this inner product, i.e.

$$\omega^{n+1}_{\mathrm{OR}} = (R^n, R^n)\;(R^n, L_{\mathrm{sp}}\,\Delta n^{n+1}). \tag{54}$$

The subscripts 'MR' and 'OR' in equations (53) and (54) stand for 'minimum residual' and 'orthogonal residual' respectively. For linear problems, these algorithms can be formulated so as to require no more evaluations of $L_{\mathrm{sp}}u$ than a standard preconditioned iteration scheme. The minimum-residual scheme has been found to be slightly faster than the orthogonal-residual scheme for problems with smooth solutions. It has the tendency to 'fall asleep,' however, and not converge for more oscillatory problems; the orthogonal-residual scheme is more robust for such problems. In practice, both methods can achieve order-of-magnitude reductions in residual per iteration for Chebyshev-discretized Poission equations.

APPLICATIONS

In this section, a few representative examples of the successful applications of spectral methods to aerodynamic and fluid mechanic problems will be given. The first example involves solution of the full-potential equations for subsonic or transonic flow about an arbitrary lifting airfoil. The aspects of geometric complexity, requirement of a fast solution scheme, and the appearance of discontinuous weak solutions had to be dealt with in this work. Solution of the boundary-layer equations will be outlined next. Accurate and fast solution schemes were devised for versions of the boundary-layer equations from the Blasius form to the two-dimensional inverse form for computing weakly separated boundary layers. The final application to be described involves solution of the incompressible Navier–Stokes equations. The particular algorithm to be outlined here will be shown in the context of predicting the Taylor–Couette vortex flow in a finite-length geometry.

In the discretization of the potential equation, the first step is to define a reduced potential by substracting freestream and boundary surface terms from the total potential function to allow the use of homogeneous boundary conditions. After a conformal mapping that transfers the airfoil shape to a circle, the potential equation becomes

$$\frac{\partial}{\partial R}\left(R\rho\frac{\partial G}{\partial R}\right)+\frac{\partial}{\partial \Theta}\left(\frac{\rho}{R}\frac{\partial G}{\partial \Theta}\right)=0, \tag{55}$$

where G is the reduced potential, R and Θ are the computational polar co-ordinates and ρ is the fluid density. The reduced potential is periodic in Θ and satisfies

$$\partial G/\partial R=0 \quad \text{at} \quad R=1, \tag{56}$$

$$G\rightarrow 0 \quad \text{as} \quad R\rightarrow\infty, \tag{57}$$

and the Kutta condition. The density is given by the isentropic relation

$$\rho=\left(1-\frac{\gamma-1}{2}M_{\infty}^{2}(q_r^2+q_\theta^2-1)\right)^{1(\gamma-1)}, \tag{58}$$

where the ratio of specific heats is denoted by γ and M_{∞} is the Mach number at infinity. The velocity components in the physical (r,θ) plane are

$$q_r=\frac{1}{h}\frac{\partial\Phi}{\partial R}, \quad q_\theta=\frac{1}{Rh}\frac{\partial\Phi}{\partial\Theta}, \tag{59}$$

and the Jacobian between the complex physical plane ($z=r\mathrm{e}^{\mathrm{i}\theta}$) and the complex computational plane ($\sigma=R\mathrm{e}^{\mathrm{i}\Theta}$) is

$$h=\left|\frac{\mathrm{d}z}{\mathrm{d}\sigma}\right|. \tag{60}$$

The mapping is numerically generated via a Fourier series. Further details are provided in Reference 25.

The spectral method employs a Fourier series representation in Θ. Constant grid spacing in Θ corresponds to a convenient dense spacing in the physical plane at the leading and trailing edges. The domain in R (with a large, but finite outer cut-off) is mapped onto the standard Chebyshev domain $[-1,1]$ by an analytical stretching transformation that clusters the collocation points

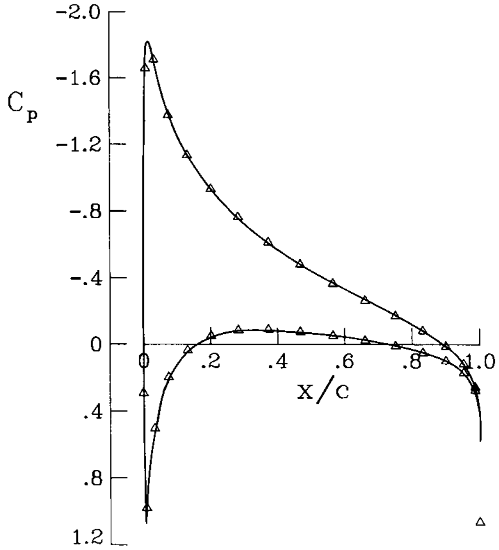

Figure 3. Surface pressure coefficient distribution about NACA0012 airfoil; subcritical test case, spectral solution (symbols) versus FLO36 (full curve)

near the airfoil surface. The stretching is so severe that the ratio of the largest-to-smallest radial intervals is typically greater than 1000.

The flow past an NACA0012 airfoil at 4° angle of attack and a freestream Mach number of 0·5 is a challenging subsonic and thus elliptic case. Nevertheless, the spectral solution on a relatively coarse grid captures all the essential details of the flow. The surface pressure coefficient from the spectral code AFSP[24], using 16 points in the radial (R) direction and 32 points in the azimuthal (Θ) direction, is displayed in Figure 3. The symbols denote the solution at the collocation points. For comparison, the result from the finite-difference, multigrid, approximate-factorization code FLO36[26] is shown as a full curve. The grid used in the benchmark finite-difference calculation is so fine (64×384 points) that the truncation error is well below plotting accuracy. The FLO36 and AFSP results are identical to plotting accuracy. The spectral computation on this mesh yields a lift coefficient with truncation error less than 10^{-4}. Spectral solutions on a 16×32 grid are thus of more than adequate resolution and accuracy for subsonic flows.

The truncation error of the spectral solution decays rapidly with mesh refinement. In Figure 4 is shown the error in the lift coefficient predicted by the spectral method (AFSP) against the average mesh spacing in the computational plane, as compared to results from FLO36. The full curves indicate the relative error of a second-order method, with the same accuracy as each method on its finest mesh. The superior accuracy of the spectral solution is not surprising, as the solution is smooth.

In Figure 5 are shown convergence histories from FLO36, the multigrid spectral code MGAFSP and the finite-difference, approximate-factorization, single-grid code TAIR[27]. Meshes which yield approximately equivalent accuracy were chosen. The surface pressure results are the same to plotting accuracy, the lift coefficient is converged in the third decimal place and the predicted drag coefficient is less than 0·001. (Actually, the spectral result is an order of magnitude more accurate than these limits, but the TAIR result barely meets them.) Figure 6 demonstrates the improvement produced by the spectral multigrid scheme over the spectral single-grid method (AFSP). There is well over an order-of-magnitude gain in efficiency.

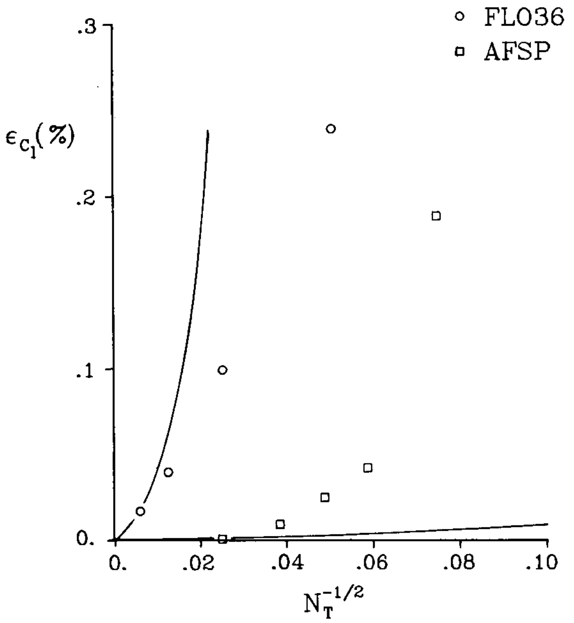

Figure 4. Relative error in lift versus average mesh spacing in computational plane; subcritical test case, spectral and finite-difference methods: NACA0012, $M_\infty = 0{\cdot}5$, $\alpha = 4°$

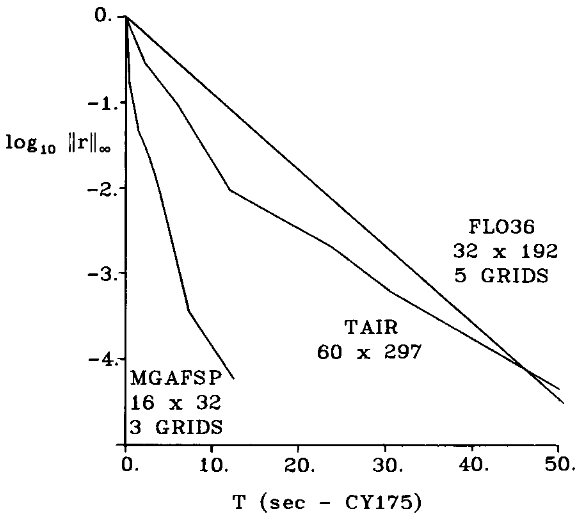

Figure 5. History of maximum residual versus machine time; subcritical test case, spectral multigrid and finite-difference methods

The potential flow problem is much more difficult whenever the flow field contains both supersonic (hyperbolic) and subsonic (elliptic) regions. Nevertheless, the spectral multigrid algorithm that succeeded for the subsonic flow case requires only a minor modification in order to succeed for the transonic (mixed) problem as well.

The most expedient technique for dealing with the mixed elliptic–hyperbolic nature of the

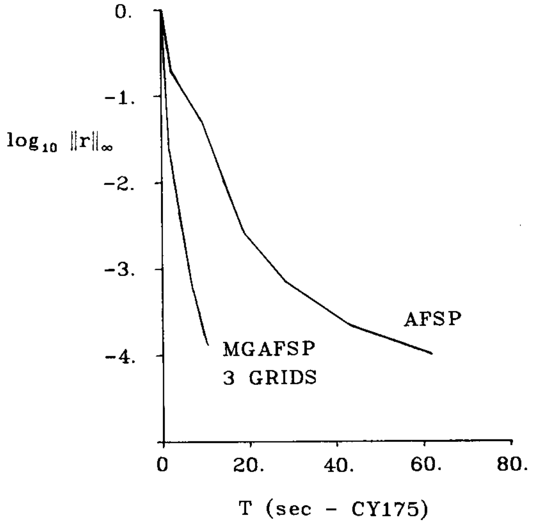

Figure 6. History of maximum residual versus machine time; subcritical test case, spectral single-grid and multigrid methods

transonic problem is to use the artificial density approach of Hafez *et al.*[28] The original artificial density is

$$\tilde{\rho} = \rho - \mu \overleftarrow{\delta} \rho \tag{61}$$

with

$$\mu = \max\{0, 1 - 1/M^2\}, \tag{62}$$

where M is the local Mach number and $\delta\rho$ is an upwind first-order undivided) difference. The spectral calculations employed a higher-order artificial density formula. The spectral method also required a weak filtering technique to deal with some high-frequency oscillations generated by the shock. Details are available in Reference 25.

A lifting transonic case is provided by the NACA0012 airfoil at $M_\infty = 0{\cdot}75$ and 2° angle of attack. A shock appears only on the upper surface for these conditions and is rather strong for a potential calculation; the normal Mach number ahead of the shock is about 1·36. Lifting transonic cases are especially difficult for spectral methods, since the solution will always have significant content in the entire frequency spectrum: the shock populates the highest frequencies of the grid and the lift is predominantly on the scale of the entire domain. An iterative scheme therefore must be able to damp error components across the spectrum.

The surface pressure coefficient distribution predicted by the spectral method for this supercritical case is shown in Figure 7 for a coarse-grid (16 × 48) solution and a fine-grid (30 × 80) solution. Note that the shock location and pre-shock pressure level are computed accurately on the coarse grid. The lift from the coarse-grid solution is also accurate to about 2 per cent, an error level for which the finite-difference code FLO36 requires a 32 × 192 grid.

The application of spectral methods to the solution of the boundary-layer equations involved the steady, two-dimensional, incompressible form of the equations, written in Görtler variables.

Momentum

$$\frac{\partial^2 f}{\partial \eta^2} - v\frac{\partial f}{\partial \eta} - \beta(f^2 - 1) - 2\xi f\frac{\partial f}{\partial \xi} = 0 \quad \text{(momentum)}, \tag{63}$$

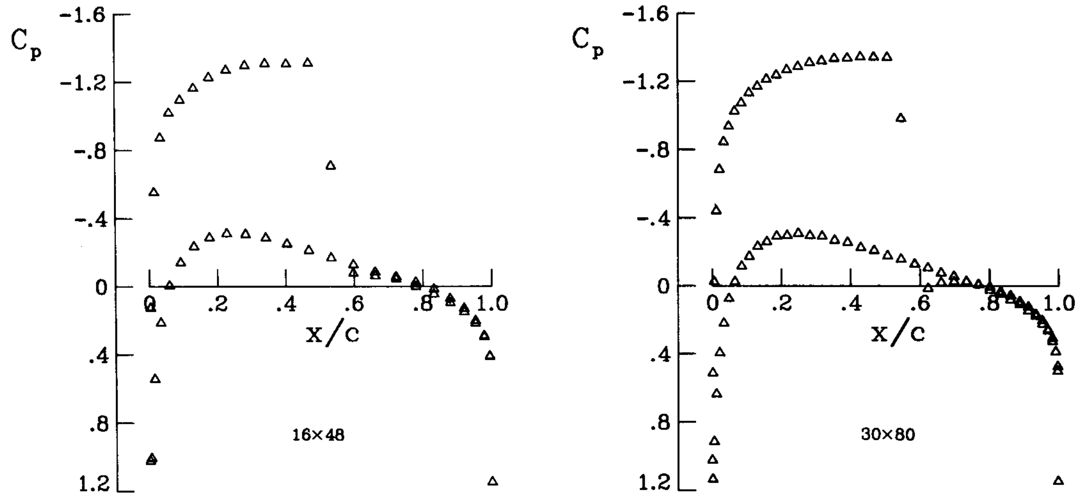

Figure 7. Surface pressure coefficient distribution about NACA0012 airfoil; supercritical test case: $M_\infty = 0{\cdot}75$, $\alpha = 2°$

Continuity

$$\frac{\partial v}{\partial \eta} + f + 2\xi \frac{\partial f}{\partial \xi} = 0 \quad \text{(continuity)}. \tag{64}$$

Equations (63) and (64) define a parabolic system of equations. The boundary conditions in η are

$$\begin{aligned} f &= 0 \quad \text{(no slip)} \\ v &= 0 \quad \text{(no transpiration)} \quad \text{at } \eta = 0, \\ f &= 1 \quad \text{as} \quad \eta \to \infty. \end{aligned} \tag{65}$$

The independent variables ξ and η are given in terms of the physical variables x and y by

$$\xi = \int_0^x U_e(\phi)\,\mathrm{d}\phi, \tag{66}$$

$$\eta = \frac{U_e(x)}{2\xi} y, \tag{67}$$

where $U_e(x)$ is the imposed inviscid edge velocity. The dependent variables f and v are the normalized streamwise velocity and stream function respectively:

$$f = \frac{\bar{u}}{U_e}, \tag{68}$$

$$v = \frac{2\xi}{U_e}\bar{v} + \frac{2\xi f(\partial\eta/\partial x)}{U_e}, \tag{69}$$

where $\bar{u}$ and $\bar{v}$ are the physical velocities in the x- and y-directions respectively. The external pressure gradient parameter β is given by

$$\beta = \frac{2\xi}{U_e}\frac{\mathrm{d}U_e}{\mathrm{d}\xi}. \tag{70}$$

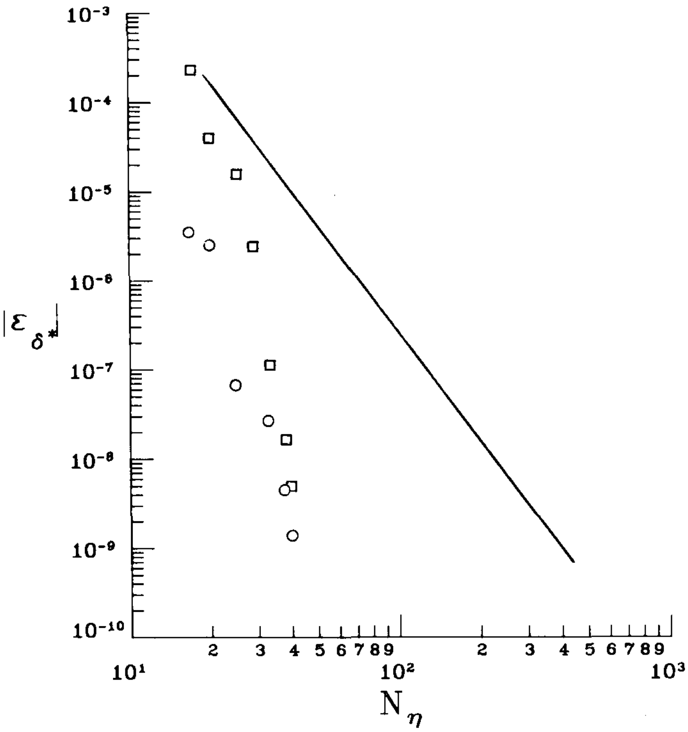

Figure 8. Relative error in skin friction versus number of grid points; Blasius flow, $\beta = 0$: spectral Chebyshev (○) and Legendre (□) methods and fourth-order box scheme (———)

Both Chebyshev- and Legendre-based collocation discretization were studied,[29] to assess their relative accuracy. Solution of the resulting equations is by preconditioned Richardson iteration. The coupling between the momentum and continuity equations is maintained in the update scheme; this was found to be critical for rapid convergence.

The simplest boundary-layer flow, Blasius flow, is obtained by setting $(\partial/\partial\xi)$ and $\beta = 0$ in equations (63) and (64). Figure 8 shows that the error in the displacement thickness predicted by the spectral method for the Blasius equations decays far more rapidly with increasing grid points than for a fourth-order box scheme. Note also the concave-down shape of the spectral error decay, characteristic of exponential-order accuracy.

The so-called inverse form of the boundary-layer equations[30] is useful for computing flows with moderate separated regions. In this form, the displacement thickness or skin friction is prescribed, rather than the external pressure gradient; the pressure gradient is computed in the solution of the equations. This procedure eliminates the Goldstein singularity in skin friction at the separation point. For the case shown in Figure 9, the prescribed displacement thickness yields a fairly large separated region as indicated by the closed streamlines. For this problem, the accuracy of the spectral method was checked by comparing the predicted skin friction at the point $\xi = 1{\cdot}61$ (indicated by the arrow in Figure 9) with the result from a second-order finite-difference code. For four-decimal-place accuracy, the spectral method required 40 points in the normal direction and 26 points in the streamwise direction; the finite-difference method needed 240 and 200 points, respectively, to resolve this sensitive flow to four-decimal-place accuracy. Moreover, the spectral solution requires only 10 per cent of the CPU time taken by the finite-difference method.

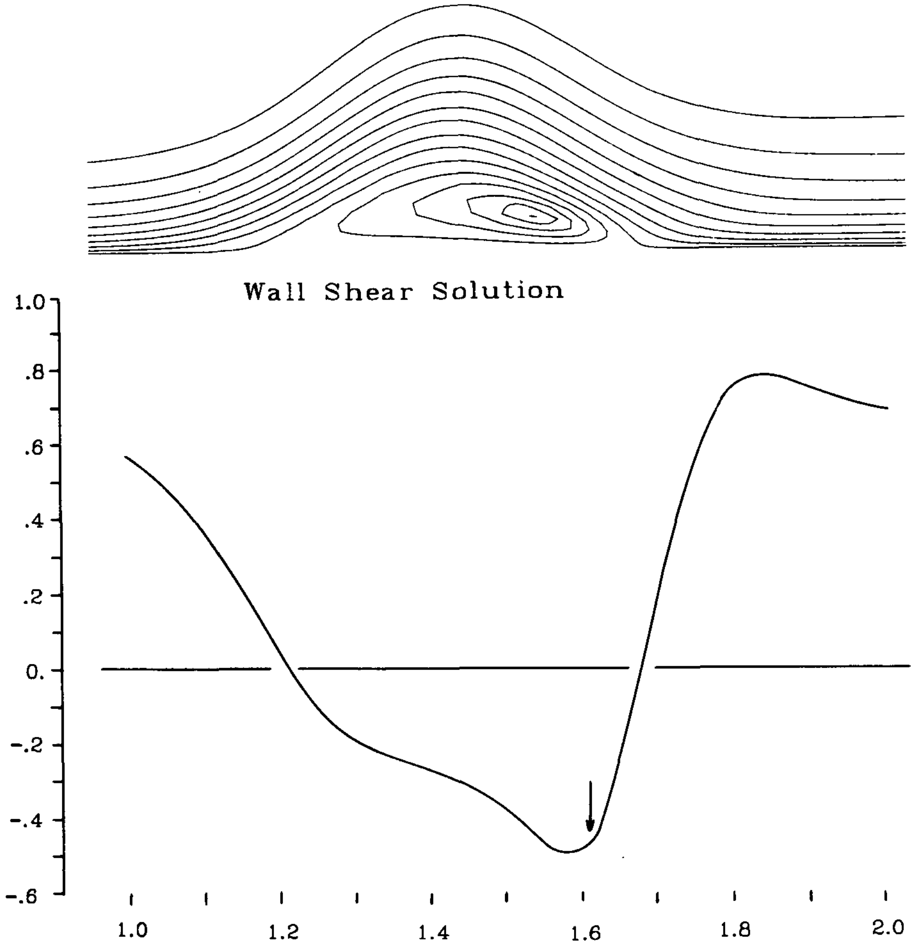

Figure 9. Spectral solution of inverse boundary-layer equations for separated flow; (a) streamlines, (b) skin friction

Most of the current interest in spectral methods stems not from their successful, but limited, applications to aerodynamic problems, but from the strides in fluid mechanical simulations which they have made possible. The simple geometries for which these simulations have generally been carried out and the computational effort required for high-resolution solutions of the time-dependent Navier–Stokes equations yields an area tailor-made for the application of spectral methods. An extensive review of some of these applications is contained in Reference 31. The purpose here is not to provide a comprehensive review of past work, but merely to provide illustrative examples of the applications of spectral methods. The simulation of bifurcation of states in Taylor–Couette flow[18] will serve as such an example, for solution of the incompressible Navier–Stokes equations by spectral methods.

The Taylor experiment on Couette flow between coaxial circular cylinders has been the subject of numerous theoretical and experimental studies.[32] This flow is rich in complex phenomena; so rich, in fact, that they are still being discovered,[33] and our understanding of them is far from complete. In a typical Taylor experiment, the inner cylinder rotates with a constant angular velocity while the outer cylinder, along with the top and bottom walls, are kept at rest. The relevant geometric parameters are the radius ratio, which is the ratio of the radii of the inner and outer cylinders, and the aspect ratio, which is the ratio of the length of the annulus to its width. The dynamic parameter is the Reynolds number based on the angular velocity of the inner cylinder and the annulus width. The Taylor–Couette flow is strongly dependent on all of these parameters. The purpose of our continuing research effort is to solve the unsteady Navier–Stokes equations by a highly accurate spectral collocation method with a view to elucidate the underlying processes leading to laminar–turbulent transition in the Taylor–Couette flow. The main result of these studies is a second-order transition from a two-cell flow, symmetric under reflection about the midplane, to an asymmetric single-cell flow that ensues with increasing Reynolds number beyond a certain critical value.

For these simulations, the incompressible time-dependent Navier–Stokes equations are written in conservative form for a cylindrical geometry. No-slip boundary conditions are enforced for all radial and axial walls, and periodicity is imposed in the azimuthal direction. The appropriate discretizations are, therefore, Chebyshev collocation in the radial and axial (r, z) directions and Fourier collocation in the azimuthal (θ) direction. A splitting scheme is used for the time-accurate solution of these discrete equations.[18] The splitting used is a two-step process in which the momentum equations without pressure terms are advanced to an intermediate time, then incompressibility is enforced in an inviscid 'pressure correction' step. Boundary conditions must be carefully handled in both steps, to ensure stability and accuracy; see Reference 18 for details. The momentum equation is advanced in time using a mixed implicit/explicit scheme, with low-storage, third-order Runge–Kutta time stepping for the non-linear advection terms and Crank–Nicolson on the viscous terms to avoid the severe time-step limitation of an explicit method. Preconditioned MR iteration is used to solve the Helmholtz equations which result from this time discretization. Additional efficiency is gained by solving these equations in the wavenumber space of the Fourier discretization, which effectively decouples the three-dimensional equations into a set of planar equations.

The results presented here pertain to the axisymmetric two-cell/one-cell bifurcation, which occurs when the Taylor apparatus has an aspect ratio up to about 1·5. The form of the bifurcation depends sensitively on this parameter; experiments[34] show that this transition can change from supercritical to subcritical with variations in the aspect ratio of as little as 8 per cent.

Most of the results are obtained either by making quasi-static changes in the Reynolds number and allowing the stable, dominant mode to settle, or by slowly sweeping through a Reynolds number range, monitoring the change in a particular mode. Of course, using a time-accurate code to simulate the bifurcations of steady-state solutions is quite inefficient, owing to the extremely small growth rates near the bifurcation points. However, the eventual aim of this work is to simulate the turbulence and broadband structure exhibited by Taylor–Couette flow at moderate Reynolds numbers; and the code was developed with these time-dependent flows in mind. The ability to simulate accurately the sensitive, steady-state bifurcations at lower Reynolds numbers is an excellent test for the numerical method.

Moreover, by using a time-dependent computation to investigate the steady-state bifurcations, we can obtain information on the path which the system follows as states exchange stability. This is illustrated by the following result. A 17×17 grid is used in all these simulations with a few calculations on a 25×25 mesh for accuracy check. The time step in these simulations corresponds to a maximum Courant number of about 0·2; the time step is limited by accuracy and not by stability. For the geometry of Benjamin and Mullin[35] with a radius ratio 0·615 and aspect ratio 1·05, the symmetric two-cell mode is allowed to stabilize at a relatively low Reynolds number $(\mathrm{Re} = 62)$. The Reynolds number is then raised impulsively to 165, above the experimental bifurcation boundary of about 150, and the growth of the one-cell asymmetric mode is observed. Random machine round-off error of the order of 10^{-14} provides the initial energy for the mode.

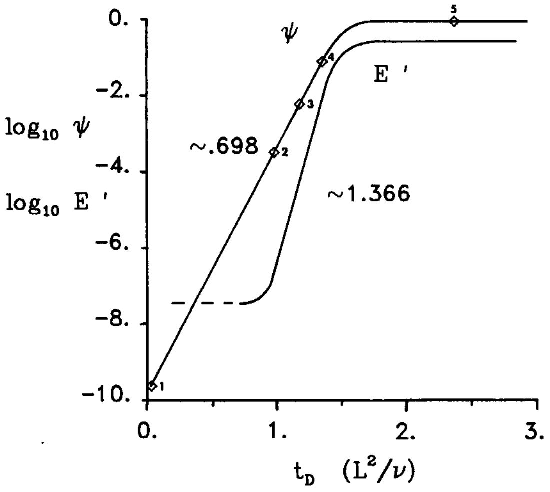

Figure 10. Growth of order parameter and perturbation energy in finite-length Taylor–Couette flow: $\Gamma = 1{\cdot}05$, $\eta = 0{\cdot}615$, $Re = 150$

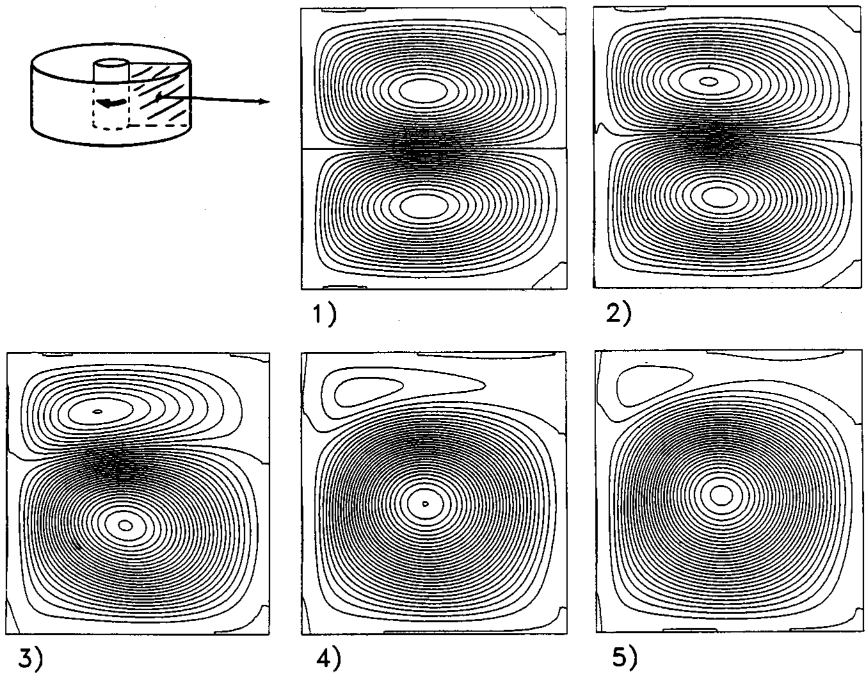

Figure 11. Progression of streamlines in crossflow plane, finite-length Taylor–Couette flow

The order parameter used here to quantify the asymmetry of the mode is due to Lucke *et al.*[36]:

$$\psi = \frac{\displaystyle\int \mathrm{d}r\,\mathrm{d}z(u(r,z) - u(r,z))}{\displaystyle\int \mathrm{d}r\,\mathrm{d}z(|u(r,z)| + |u(r,-z)|)}. \tag{71}$$

The integrals were performed by spectral collocation. The logarithm of this parameter is shown in Figure 10 plotted against time in units of the diffusion time scale L^2/ν. As can be seen, the initial instability leading to the one-cell mode appears to be linear; that is, exponential growth with time is observed, with only the later stages being modulated by non-linear effects. Also shown in Figure 10 is a plot of the logarithm of the disturbance energy against time. After an initial period, the disturbance energy grows at a rate which is within 2 per cent of twice the growth rate of ψ, as expected.

Streamlines in a cross-sectional plane at various stages in the two-cell/one-cell exchange are shown in Figure 11. The locations of these intermediate states on the ψ against time curve are indicated in Figure 10. Note that the progression between states is smooth, without abrupt collapse or alteration of the flow field structure.

PROBLEM AREAS AND CURRENT RESEARCH

The previous section has briefly demonstrated the breadth of applicability of spectral methods. There appear to be, however, two major problem areas which somewhat limit the utility of spectral methods in aerodynamic problems. The first lies in the requirement that the physical domain, however complicated, must be mapped onto a simple computational domain via a global, infinite-order smooth mapping. The necessity of high-order smoothness and a global functional description of the mapping precludes the use of the many powerful and flexible algebraic grid generators which are used in finite-difference and finite-element calculations. One potential solution to this problem is to use a so-called 'elliptic' or 'hyperbolic' grid generator,[37,38] wherein the mapping equations are discretized by an appropriate spectral method. Such a technique could be used to map a smooth physical domain onto a computational domain for spectral solution.

However, if there exist geometric discontinuities, or if there are severe resolution requirements in an isolated area of the domain, global mapping techniques are inadequate. These difficulties may be alleviated by the use of a multi-domain technique, in which the physical domain is split into a small number of subdomains which are independently discretized. The individual discretizations must be interfaced in such a way as to preserve overall spectral accuracy.

A number of spectral domain decomposition techniques have appeared in the literature. One popular technique is the spectral element method.[39,40] In this technique, discretization in individual elements is carried out via a variational statement; for advection–diffusion equations, a mixed collocation–Galerkin discretization is resorted to. Since the flexible and rapid solution schemes for pure collocation discretizations are inapplicable to such a statement, it appears that the number of nodes used per domain is limited in practice in the spectral element method.

Other spectral multi-domain techniques based on collocation involve an explicit statement of continuity of a derivative of the solution across the domain interface. Although adequate in theory, these methods in practice can produce a local disturbance near the interface, especially when resolution requirements are severe, and when the discretizations of the adjoining domains are radically different. Additionally, these techniques require the spectral discretization to include boundary points lying on the interface, precluding the use of a staggered mesh.

The multi-domain technique of Macaraeg and Streett[41] does not suffer from these difficulties. In this method, a global flux balance statement, generated by an appropriate integration of the governing equation, provides the necessary interface relation; see Reference 41 for details.

The first example of the application of this technique will illustrate the capability of the method for resolving very high gradients in a solution, while imposing an interface condition which preserves spectral accuracy. Consider the viscous Burgers' equation

$$U_t + \tfrac{1}{2}(U^2)_x = \nu U_{xx}, \qquad x \in [-1, 1],$$
$$U(-1, t) = U(1, t) = 0, \qquad U(x, 0) = -\sin(\pi x). \tag{72}$$

This problem has been studied extensively by a number of authors, using techniques ranging from standard finite difference to single-domain spectral collocation and spectral element.[42] The solution to this problem develops a very steep gradient region in the centre of the domain; the slope at $x = 0$ reaches a maximum, then decreases as the initial energy is dissipated away. For the parameters studied in Reference 42 ($\nu = 0{\cdot}01/\pi$), this maximum is reached at $t = 0{\cdot}5$; a very accurate analytical solution gives a value of 152·00516 for this maximum slope. The evolution of this solution calculated from the present method is shown in Figure 12 at time increments of 0·1.

In the present study of this problem, three domains were used, the middle domain spanning a very small region ($\pm 0{\cdot}05$) around the 'shock'. Additionally, a mapping was applied in the middle domain, to improve resolution, of the form

$$x^t = \sinh(\beta x^c)/\sinh(\beta), \tag{73}$$

where both the computational coordinate x^c and the transformed coordinate x^t are $\varepsilon[-1, 1]$, and β is an $O(1)$ constant chosen to control the packing at $x^t = 0$. The maximum stretching allowable in this mapping is subject to the same restriction as stretchings in single-domain discretizations; e.g. maximum metric ratios of the order of 10^3.

A second-order backward-Euler time-stepping technique was used; a time-step refinement study was performed to extrapolate out time-stepping errors. For this small one-dimensional problem, the algebraic system resulting from the spectral discretization of the equation plus interface conditions was Newton-linearized and solved directly using Gaussian elimination at each time step.

From the comparison study contained in Reference 42, the two methods giving the best accuracy for a given number of grid points were single-domain spectral collocation and spectral element. The collocation scheme used a mesh stretching with a maximum-to-minimum metric ratio of about 100. Beyond this stretching, a degradation in accuracy was found to occur. The spectral element discretization utilized four elements with 16 nodes in each. The behaviour of the error in maximum slope from these methods and the present scheme is shown in Table VI. As can be seen, the present method with just 35 total points (12 points in the outer domain, 13 points in the middle domain, 12 points in the left outer domain; hereafter denoted 12/13/12) yields results of equivalent accuracy to the spectral element and single-domain spectral collocation methods of Reference 42, which both use 64 total points. Further mesh refinements using the present method show

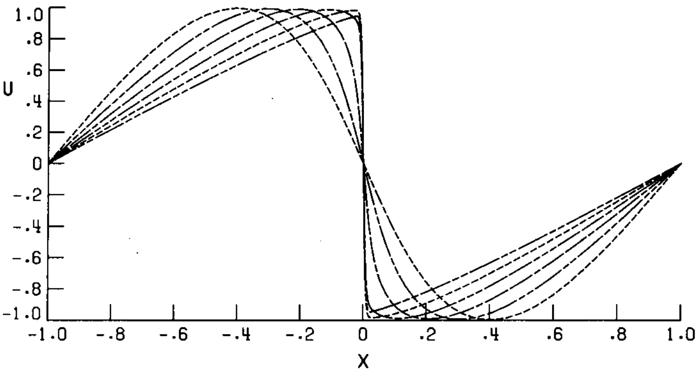

Figure 12. Solution to viscous Burgers' equation with time; $\nu = 0{\cdot}01/\pi$

exponential-order error convergence, as seen in Table VI by the order-of-magnitude decrease in relative error as the mesh is refined to 20/21/20 and again with mesh refinement of 32/33/32. For the same total number of points, the present method is an order of magnitude more accurate than the single-domain collocation or spectral element solutions of Reference 42.

In order to demonstrate the capability of the present method to handle radically different mappings between adjacent domains, a solution to the above viscous Burgers' equation for $\nu = 10^{-4}$ is shown in Figure 13. The maximum slope for this solution is greater than 5000. The discretization used was 12/31/12; the stretching in the middle domain was so severe that the ratio of largest mesh spacing in the outer domains to the smallest in the middle domain is greater than 10^5. A factor-of-5000 magnification of the high-gradient region of this solution is shown in Figure 14. The emphasis in this plot is the oscillation-free resolution of this region. (Linear interpolation between points is used for plotting, making the plot appear somewhat jagged.)

To demonstrate the generality of this multi-domain technique, a solution to the Poisson equation

$$u_{xx} + u_{yy} - \cos\frac{\pi y}{2}\cos\frac{\pi x}{4}, \quad x\in[-2,2], \quad y\in[-1,1], \tag{74}$$

$$u(x,-1) = u(x,1) = u(-2,y) = u(2,y) = 0$$

Table VI. Maximum slope and per cent relative error in maximum slope for viscous Burger's equation (72); comparison of present method with results from Reference 42

Method	Discretization	Maximum slope	Per cent relative error
	3 domains:		
Present	12/13/12	152·03544	1·99 (−2)
Present	20/21/20	152·00011	3·23 (−3)
Present	32/33/32	152·00513	2·14 (−4)
Spectral	4 elements:		
element[42]	16/16/16/16	152·04	2·29 (−2)
Spectral	1 domain:		
collocation[42]	64	152·025	1·31 (−2)
Exact[42]		152·00516	

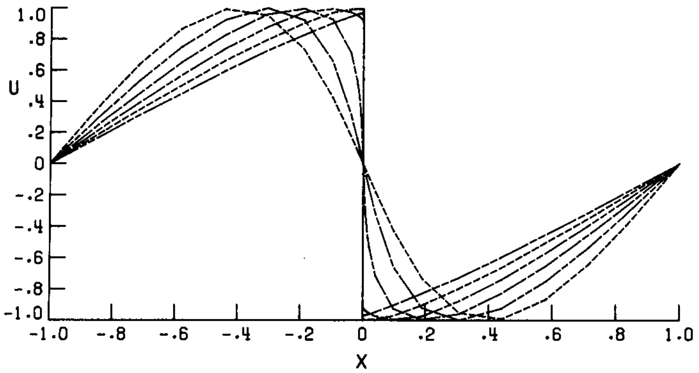

Figure 13. Solution to viscous Burgers' equation with time; $\nu = 10^{-4}$

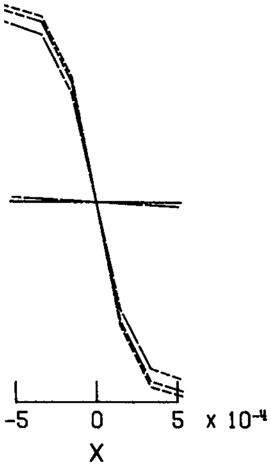

Figure 14. Magnification of central region of Figure 13

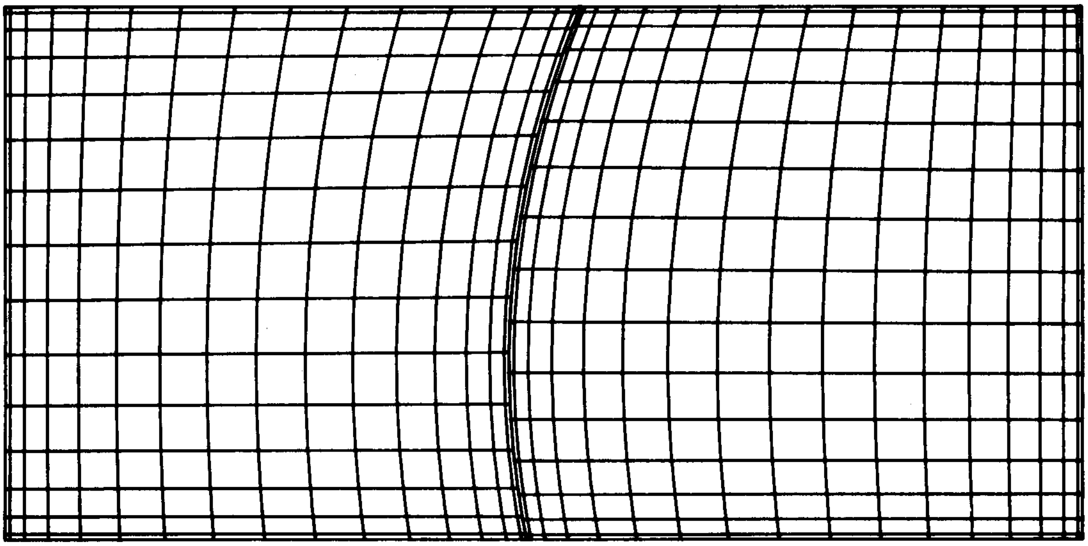

Figure 15. Skewed curvilinear multidomain mesh

was obtained on the skewed two-domain mesh shown in Figure 15. This mesh, containing 17 × 16 and 18 × 17 points in the left and right domains respectively, was generated by first choosing the interface line, in this case a cubic polynomial. Chebyshev distributions with respect to arc length were used to establish the mesh points on the interface, as well as along the domain boundaries at $x = \pm 2$. One curvilinear co-ordinate family was generated by connecting these corresponding points with straight lines. Mesh points along these co-ordinate lines were then established with Chebyshev distributions with respect to arc length, resulting in a sheared non-orthogonal mesh. Equation (74) was written in generalized contravariant flux form; the metrics were evaluated by spectral differentiation of the co-ordinate distributions. The flux component normal to the interface was taken to be continuous in the interface condition. As can be seen in the isolines of the solution shown in Figure 16, the solution is everywhere smooth and regular.

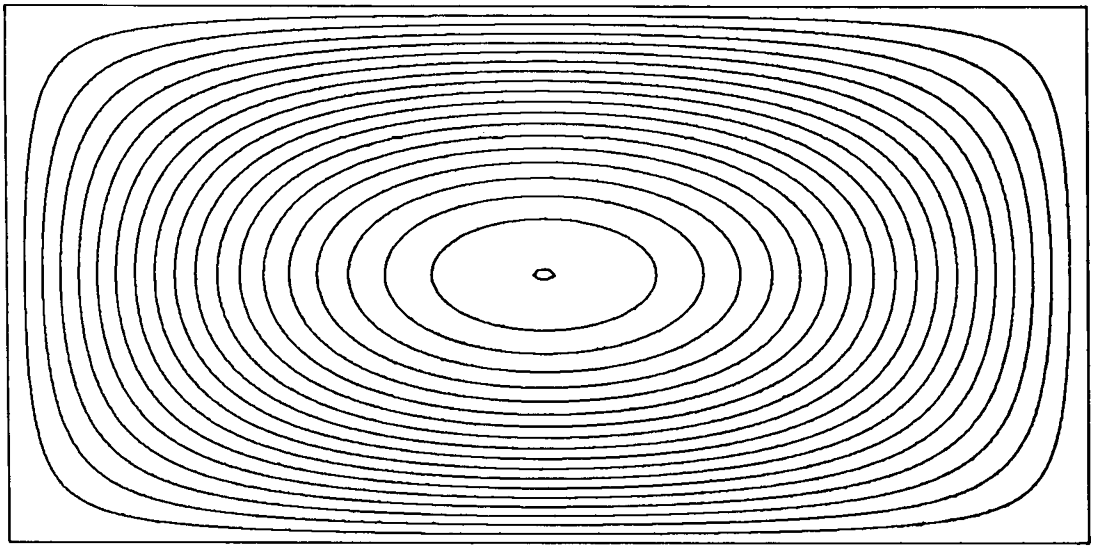

Figure 16. Isolines of solution to Poisson equation on mesh of Figure 15

The second major difficulty in the application of spectral methods to aerodynamic problems occurs when the solution being sought is discontinuous. This situation appears, for instance, in the solution of the compressible Euler equations for supercritical flow with shocks. The difficulty is that the spectral global basis function 'fit' of the discontinuous solution is subject to Gibbs' phenomenon and global low-order error convergence, since the accuracy of a spectral discretization is tied to the smoothness of the function being discretized. A solution may be obtained through the use of an artificial viscosity, as used in finite-difference or finite-element computations; or, equivalently, through the use of a filter. These techniques attenuate the high-frequency oscillations, rendering the solution essentially smooth on the scale of the mesh spacing. This technique was used in the application of spectral methods to supercritical potential flow discussed previously. Although solutions may be obtained by this procedure, spectral high-order accuracy is lost, and thus any advantage of the use of a spectral method is difficult to justify.

One possible solution to this problem is the coupling of a multi-domain technique with a shock-fitting algorithm. In such a method, the position of the interface and the jump conditions across it are part of the solution. Another technique under study is to obtain the oscillatory solution which is characteristic of a spectral discretization of a discontinuous function and post-process it in such a way as to extract a spectrally accurate discontinuous solution. Such methods rely heavily on information content theories of Lax,[43] which state essentially that the oscillations contain information about the tail of the truncated series representing the discontinuous function, and thus contain high-order information about the discontinuity.

A number of methods to extract this information are being explored.[44,45] The first involves a concept of weak convergence. Say a solution is sought to

$$u_t = Lu, \tag{75}$$

where L is a linear differential operator and the initial condition $u(x, 0)$ is discontinuous. The spectral solution to (75), u_N, may not approximate well, pointwise, the exact solution u, but it may be shown[44] that

$$(u_N, \phi) = (u, \phi) + \varepsilon, \tag{76}$$

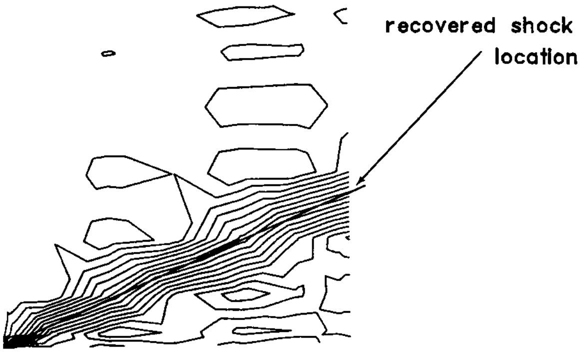

Figure 17. Isolines of viscosity-free solution to Euler equations with shock

where ϕ is a smooth test function and the error ε is

$$\varepsilon \sim \|\phi\|_s/N^s, \quad s, \tag{77}$$

and thus is spectrally small. This is the concept of weak convergence of the spectral solution to the exact solution. A particular construction of ϕ may be used to extract spectrally accurate data from u_N away from the discontinuity:

$$\phi = \phi_1 \phi_2, \tag{78}$$

where ϕ_1 is a spectral approximation to the Dirac delta function and ϕ_2 is a smooth cut-off function, which is zero outside the smooth region around the point under consideration. The inner product of this test function with the spectral solution may be shown[44] to yield a spectrally accurate pointwise approximation to the exact solution. In practice, the accuracy of this technique depends on the distance from the discontinuity and works well only for large of N. The theory for this technique is available for a Fourier–Galerkin discretization of a linear differential operator; extension to the Fourier collocation case requires a modification of the initial data.

Another technique for dealing with a discontinuous solution is to fit a discontinuous function to the oscillatory spectral solution and extract the smooth part of the solution, which is, presumably, high-order accurate. An accurate estimate of the location and strength of the discontinuity may be obtained by this method.[45] Figure 17 shows such an application. A Chebyshev collocation solution to the 2-D Euler equations was obtained on a 9 × 9 mesh with boundary conditions which produce a shock running at an angle through the domain. No artificial viscosity or damping was used in the computation due to the coarseness of the mesh; large oscillations appeared in the solution due to the shock. The oscillatory solution was post-processed by fitting with a step function; the shock location in the domain was recovered to 0·1 per cent of the exact solution, and the constant pre- and post-shock states were computed to 1 per cent. Note that the recovered shock location is two orders of magnitude more accurate than stating that the shock occurred at a particular mesh point, since the average mesh spacing in the field was 1/8. The discontinuity subtraction technique thus yields high-order accurate information on the location and strength of discontinuity in a spectral solution; however, discontinuous higher derivatives still remain. Application of a 'weak convergence filter' could provide a final, spectrally accurate solution.

CONCLUSIONS

Spectral methods have shown their utility on a wide range of aerodynamic and fluid mechanic problems. The primary advantage of their use over traditional finite-difference and finite-element techniques lies in the high-order accuracy shown by these methods on properly formulated problems. This property can be used either to efficiently produce engineering-accuracy solutions or to generate solutions with accuracy unattainable by other methods. The latter situation is typically encountered in simulations of sensitive fluid mechanic phenomena. The further property that spectral methods strongly mimic the continuous solution to the analytic differential equation being discretized has been found to be crucial for bifurcation simulations.

On the negative side, spectral methods are considerably more difficult to apply successfully than finite-difference or finite-element methods. Spectral methods are pathologically sensitive to formulation inconsistencies, due to their strong conformity with the behaviour of analytic differential equations. The use of proper boundary conditions is particularly crucial. Also, since the algebraic systems obtained from spectral discretization tend to be full and ill-conditioned, efficient solution schemes, are essential for the practical application of spectral methods. In addition, the requirement for global smooth mappings complicates their application in complex physical domains. Finally, discontinuous solutions necessitate special formulation and processing of the spectral method if the high-order property is to be maintained. Research in these areas is continuing, with considerable progress being made.

REFERENCES

1. B. A. Finlayson and L. E. Scriven, 'The method of weighted residuals—a review', *Appl. Mech. Rev.*, **19**, 735–748 (1966).
2. J. C. Slater, 'Electronic energy bands in metal', *Phys. Rev.*, **45**, 794–801 (1934).
3. J. Barta, 'Über die Naherungsweise Lösung einiger Zweidimensionaler Elastizitätsaufgaben', *Z. Angew. Math. Mech.*, **17**, 184–185 (1937).
4. R. A. Frazer, W. P. Jones and S. W. Skan, 'Approximation to functions and to the solutions of differential equations', *Report and Memo No. 1799*, Great Britain Aero. Res. Council, London, 1937.
5. C. L. Lanczos, 'Trigonometric interpolation of empirical and analytic functions', *J. Math. Phys.*, **17**, 123–199 (1938).
6. C. W. Clenshaw, 'The numerical solution of linear differential equations in Chebyshev series', *Proc. Cambridge Phil. Soc.*, **53**, 134–149 (1957).
7. C. W. Clenshaw and H. J. Norton, 'The solution of nonlinear ordinary differential equations in Chebyshev series', *Comput. J.*, **6**, 88–92 (1963).
8. K. Wright, 'Chebyshev collocation methods for ordinary differential equations', *Comput. J.*, **6**, 358–365 (1964).
9. J. V. Villadsen and W. E. Stewart, 'Solution of boundary value problems by orthogonal collocation', *Chem. Eng. Sci.*, **22**, 1483–1501 (1967).
10. H.-O. Kreiss and J. Oliger, 'Comparison of accurate methods for the integration of hyperbolic equations', *Report No. 36*, Department of Computer Science, Uppsala University, Sweden, 1971.
11. S. A. Orszag, 'Comparison of pseudospectral and spectral approximations', *Stud. Appl. Math.*, **51**, 253—259 (1972).
12. I. Silberman, 'Planetary waves in the atmosphere', *J. Meteorol.*, **11**, 27–34 (1954).
13. S. A. Orszag, 'Numerical methods for the simulation of turbulence', *Phys. Fluids*, **12**, Supplement II, 250–257 (1969).
14. E. Eliasen, B. Machenauer and E. Rasmussen, 'On a numerical method for integration of the hydrodynamical equations with a spectral representation of the horizontal fields', *Report No. 2*, Department of Meteorology, Copenhagen University, Denmark, 1970.
15. D. Gottlieb and S. A. Orszag, 'Numerical analysis of spectral methods: theory and applications', *CBMS-NSF Regional Conference Series in Applied Mathematics*, SIAM, 1977.
16. P. J. Davis and P. Rabinowitz, *Numerical Integration*, Blaisdell, MA, 1967.
17. C. Canuto and A. Quarteroni, 'Approximation results for orthogonal polynomials in Sobolev spaces', *Math. Comput.*, **38**, No. 157 (1982).
18. C. L. Streett and M. Y. Hussaini, 'Finite length Taylor–Couette flow', *Proc. ICASE/NASA Workshop on Stability of Time Dependent and Spacially Varying Flows*, Hampton, VA, August 1985.
19. S. A. Orszag, 'Spectral methods for problems in complex geometries', *J. Comput. Phys.*, **37**, 70–92 (1980).
20. C. L. Streett and M. G. Macaraeg, 'Preconditioning for first-order spectral discretizations', *NASA TM-87614*, August 1986.
21. W. Hackbusch and U. Trottenberg (eds), *Multigrid Methods, Lecture Notes in Mathematics*, **960**, Springer-Verlag, New York, 1982.
22. T. A. Zang, Y. S. Wong and M. Y. Hussaini, 'Spectral multigrid methods for elliptic equations', *J. Comput. Phys.*, **48**, 485–501 (1982).
23. T. A. Zang, Y. S. Wong and M. Y. Hussaini, 'Spectral multigrid methods for elliptic equations II', *J. Comput. Phys.*, **54**, 489–507 (1984).

24. C. L. Streett, T. A. Zang and M. Y. Hussaini, 'Spectral multigrid methods with applications to transonic potential flow', *ICASE Report No. 83-11*, 1983.
25. C. L. Streett, 'A spectral method for the solution of transonic potential flow about an arbitrary airfoil', *Proc. Sixth AIAA Computational Fluid Dynamics Conference*, Danvers, MA, July 1983.
26. A. Jameson, 'Acceleration of transonic potential flow calculations on arbitrary meshes by the multiple grid method', *AIAA Paper 79-1458*, 1979.
27. T. L. Holst, 'A fast, conservative algorithm for solving the transonic full-potential equation', *AIAA Paper 79-1456*, 1979.
28. M. M. Hafez, J. C. South and E. M. Murman, 'Artificial compressibility methods for numerical solution of transonic full potential equation', *AIAA J.*, **17**, 838–844 (1979).
29. C. L. Streett, T. A. Zang and M. Y. Hussaini, 'Spectral methods for solution of the boundary-layer equation', *AIAA Paper 84-0170*, January 1984.
30. J. E. Carter, 'Inverse boundary-layer theory and comparison with experiment', *NASA TP-1208*, July 1978.
31. M. Y. Hussaini and T. A. Zang, 'Spectral methods in fluid dynamics', *ICASE Report No. 86-25*, 1986; see also *NASA CR-178103*, May 1986.
32. R. C. DiPrima and H. L. Swinney, 'Instabilities and transitions in flow between concentric rotating cylinders', in H. L. Swinney and J. P. Gollub (eds), *Hydrodynamic Instabilities and the Transition to Turbulence*, Springer-Verlag, New York, 1981.
33. G. Pfister, 'Deterministic chaos in rotational Taylor–Couette flow', *Lecture Notes in Physics*, **235**, 199–210 (1985).
34. A. Aitta, G. Ahlers and D. S. Cannel, 'Tricritical phenomena in rotating Couette –Taylor flow', *Phys. Rev. Lett.*, **54**, 673–676 (1985).
35. T. B. Benjamin and T. Mullin, 'Anomalous modes in the Taylor experiment', *Proc. Soc. London A*, **377**, 221–249 (1981).
36. M. Lucke, M. Mihelcic, K. Wingerath and G. Pfister, 'Flow in a small annulus between concentric cylinders', *J. Fluid Mech.*, **140**, 343–353 (1984).
37. J. F. Thompson, F. C. Thames and C. W. Mastin, 'Automatic numerical generation of body-fitted curvilinear coordinate system for field containing any number of arbitrary two-dimensional bodies', *J. Comput. Phys.*, **15**, 299–319 (1974).
38. R. L. Sorenson and J. L. Steger, 'Use of hyperbolic partial differential equations to generate body-fitted coordinates', in *Numerical Grid Generation Techniques, NASA CP-2166*, 1982.
39. A. T. Patera, 'A spectral element method for fluid dynamics: laminar flow in a channel expansion', *J. Comput. Phys.*, **54**, 468–488 (1984).
40. A. T. Patera, 'Fast direct Poisson solvers for high-order finite element discretizations in rectangularly decomposable domains', *J. Comput. Phys.*, **65**, (1986).
41. M. G. Macaraeg and C. L. Streett, 'Improvements in spectral collocation through a multiple domain technique', *Appl. Numer. Methods*, **2**, No. 2, 95–108 (1986).
42. C. Basdevant, M. Deville, P. Haldenway, J. Lacroix, D. Orlandi, J. Quazzani, A. Patera and R. Petret, 'Spectral and finite difference solutions of the Burgers' equation', *Comput. and Fields*, **14**, 23–41 (1986).
43. P. D. Lax, 'Accuracy and resolution in the computations of solutions of linear and nonlinear equations', in *Recent Advances in Dimensional Analysis*, MRC University at Wisconsin, Academic Press, 1978, pp. 107–117.
44. D. Gottlieb and E. Tadmor, 'Recovering pointwise values of discontinuous data within spectral accuracy', *ICASE Report No. 85-3*, January 1985.
45. D. Gottlieb, L. Lustman and C. L. Streett, 'Spectral methods for two-dimensional shocks', *ICASE Report No. 82-38*, November 1982.