

Fourier/Chebyshev Methods for the Incompressible Navier–Stokes Equations in Infinite Domains

ROQUE CORRAL

School of Aeronautics, U. Politécnica, Pl. Cardenal Cisneros 3, 28040 Madrid, Spain

AND

JAVIER JIMÉNEZ

School of Aeronautics, U. Politécnica, Pl. Cardenal Cisneros 3, 28040 Madrid, Spain, Center for Turbulence Research, Stanford University, Stanford, California and NASA Ames Research Center, Ames, Iowa

Received June 14, 1993; revised January 5, 1995

A fully spectral numerical scheme is presented for the unsteady, high Reynolds number, incompressible Navier–Stokes equations, in domains which are infinite or semi-infinite in one dimension. The domain is not mapped, and standard Fourier or Chebyshev expansions can be used. The handling of the infinite domain does not introduce any significant overhead. The scheme assumes that the vorticity in the flow is essentially concentrated in a finite region, which is represented numerically by standard spectral collocation methods. To accommodate the slow exponential decay of the velocities at infinity, extra expansion functions are introduced, which are handled analytically. A detailed error analysis is presented, and two applications to direct numerical simulation of turbulent flows are discussed in relation with the numerical performance of the scheme. © 1995 Academic Press, Inc.

1. INTRODUCTION

Spectral collocation methods have proved in many occasions their advantages for the direct numerical simulation of incompressible turbulent flows. Their high numerical accuracy and the lack of appreciable numerical viscosity induced by the spatial discretization are invaluable at the high Reynolds numbers which are typical of turbulence or transition. Unfortunately, their use has been mostly limited to spatially periodic situations [1] or to domains which are bounded in their non-periodic directions, like channels [1–2]. The former are well adapted to Fourier representations, while the latter can be treated efficiently by mixed Fourier/Chebyshev schemes. In both cases, the transforms are done using an FFT algorithm. They typically consume 70–80% of the computational time of the problem, while most of the rest is spent in the solution of an elliptic sub-step, typically a Poisson equation, which is introduced by the incompressible limit. In both Fourier and Chebyshev schemes, this second part is relatively fast. The matrix to be inverted is diagonal in the first case and tridiagonal in the second.

There are many flows of scientific or technical interest which are neither periodic nor bounded in one or more of their dimensions. Examples are wall boundary layers, free shear layers, and jets. In all of them there is at least one transverse dimension in which the flow is unbounded and has to be matched to an irrotational free stream. It is still possible to use spectral methods in those cases by means of distorted grids and Chebyshev expansions [3], but the tridiagonal structure of the Poisson matrices is destroyed, and their solution is complicated. Malik *et al.* [4] address this difficulty using preconditioned iterative techniques for solving the implicit part of their algorithm, but they cannot avoid an unnecessary concentration of points at infinity. Alternatively, expansions in mapped Jacobi polynomials can be used [5–6]. They are well suited to infinite domains, but they have no fast transform equivalent to the FFT. In any case, the introduction of the boundary condition at infinity results in a substantial increase in the amount of work involved.

Spalart *et al.* [5–6] noted that, if the vorticity formulation is used for the equations of motion, the irrotational stream behaves linearly and can be treated in a much simpler way than the nonlinear rotational part. Their expansions include extra exponential terms, whose coefficients are marched independently and whose mission is to represent these irrotational fluctuations. The rotational part of the flow is usually localized in a fairly compact part of the field (e.g., near the wall) and decays very fast at large distances.

In this paper we extend their idea by noting that, if the spectral representation is only needed for compact quantities which decay very fast at infinity, as is the case of vorticity, it is possible to consider the problem in a finite domain and to extend those variables periodically across the numerical boundaries with negligible discontinuity, in such a way that they can be represented as Fourier series and treated by standard spectral methods. The quantities that contain an irrotational component, such as the velocities, can then be expanded in Fourier series

plus exponential terms. We will show below that the coefficients of those extra terms need not be marched and that they can be computed directly from the Fourier part of the expansion. The result is a compact, fast code, that respects the efficiency of the FFT and that retains the simple structure of the Poisson matrices. The overhead introduced by the infinite domain is negligible.

The computational problem and the implementation of the boundary conditions are discussed in the next three sections. An analysis of the approximation errors is presented next, followed by some numerical experiments to test the robustness of the scheme to off-design situations, such as those in which the vorticity does not remain confined away from the numerical boundaries. The scheme is shown to be robust, in the sense that it handles smoothly, although incorrectly, conditions whose vorticity has not decayed completely at the edge of the computational domain.

We have implemented the scheme in two practical cases: a two-dimensional free mixing layer, periodic in one dimension and doubly infinite in the other [7], and a three-dimensional boundary layer, periodic in two dimensions and semi-infinite in the third [8]. The former is implemented as a double Fourier expansion, while the latter uses two Fourier directions and a Chebyshev expansion to account for the boundary condition at the wall. Both codes have been tested extensively, and the computational experience derived from their use is discussed briefly in the last section. The boundary layer code has been run both for turbulent simulations that can be compared to experimental results and for transition calculations, checked against known linear instability eigenvalues. The two-dimensional mixing layer has been checked against previous calculations. We have run a parallel implementation of the shear layer code up to Peclet numbers of 3×10^5 . The simplicity of the new scheme was important for the successful parallelization of the code.

2. THE COMPUTATIONAL SCHEME

The governing equations for an incompressible flow can be written as

$$\frac{\partial \mathbf{v}}{\partial t} = \mathbf{H} + \frac{1}{\text{Re}} \nabla^2 \mathbf{v}, \quad (2.1)$$

$$\text{div } \mathbf{v} = 0, \quad (2.2)$$

where \mathbf{v} is the velocity vector and \mathbf{H} includes the nonlinear convective terms and the pressure gradient. Re is a Reynolds number, which we assume to be large. We are interested in the case in which the flow is periodic in all directions except one, along which the computational domain is at least semi-infinite. We will call this last direction y and use a standard Fourier representation for the flow in (x, z) ,

$$\mathbf{v}(x, y, z) = \sum \sum \hat{\mathbf{v}}_{jk}(y) e^{i(j\alpha x + k\beta z)}, \text{ etc.} \quad (2.3)$$

The result is an evolution equation for each individual Fourier component.

Assume now that most of the fluid is initially irrotational, that the vorticity is concentrated within a finite range of y , and that the entrainment at infinity vanishes or is incoming. Because vorticity is essentially convected by the fluid, its support remains approximately bounded for all times, with only viscous diffusion to spread it into the irrotational region. For high Reynolds numbers we can assume that the flow remains irrotational for large values of y . The velocity in an irrotational incompressible flow satisfies a Laplace equation, and each of its (j, k) Fourier components decays exponentially for large y as

$$\hat{v}_{jk}(y) \sim e^{-\gamma|y|}, \quad \gamma^2 = (j\alpha)^2 + (k\beta)^2. \quad (2.4)$$

Especially for those components with the lowest wave numbers, this decay is slow and the numerical domain has to be extended to large distances before the velocity can be considered to have decayed to its value at infinity. For example, if we consider a computational box $(0, L_x) \times (-L, L) \times \dots$, the lowest wavenumber in x is $\alpha = 2\pi/L_x$, and the magnitude of the velocities at $y = \pm L$ is of the order of $\exp(-2\pi L/L_x)$, which can only be made small at the expense of using very wide boxes such that $L \gg L_x$. Because of the compactness of the vorticity, much of the computational box is then irrotational.

On the other hand, if the equations are written in terms of the vorticity, the decay of the important variables is much faster, and the domain can be kept smaller, provided that some procedure is implemented to handle the slow velocity decay in the irrotational region. In two dimensions, this can be done by using a vorticity-stream function (ω, ψ) formulation,

$$\frac{\partial \omega}{\partial t} = h_\omega + \frac{1}{\text{Re}} \nabla^2 \omega, \quad (2.5)$$

where h_ω includes the convective term, $\mathbf{v} \cdot \nabla \omega$, and all the body forces, but is independent of the pressure gradient. In addition, each integration step involves an elliptic equation

$$\nabla^2 \psi = -\omega, \quad u = \partial \psi / \partial y, \quad v = -\partial \psi / \partial x, \quad (2.6)$$

which enforces continuity. Since the boundary conditions are normally expressed in terms of velocity, they are applied on this second equation and can usually be reduced to $\psi \rightarrow 0$ as $y \rightarrow \infty$. An equivalent formulation for three-dimensional flow is given in [2] and also involves an evolution step for a vorticity-like vector, plus a Poisson equation that relates it to the velocities.

Assume that the flow is doubly infinite in y and periodic in all other directions. We can Fourier transform the flow variables in all the periodic dimensions, and we would like to do the same in y , but we are prevented from doing so by the infinite domain. Assume, on the other hand, that we are only interested




FIG. 1. Behavior at large distances of e^{-y^2} (dashed line) and of the solution of $d^2 f / dy^2 - f = e^{-y^2}$ with $f(\pm\infty) \rightarrow 0$ (solid).

in representing the vorticity variables, which are essentially zero beyond the finite range $y \in (-L, L)$. From the point of view of these variables, we can restrict our numerical domain to $(-L, L)$, assume a periodic extension in y and use a Fourier transform valid within the domain. The errors associated to this Fourier expansion are caused by the lack of continuity of the variables, or of their derivatives, at the two boundaries but, since the vorticities decay very fast, both the variables and their derivatives can be assumed to vanish at the boundary, and those errors are negligible.

The same is not true of the stream function, or of the velocities. Consider the Fourier transform of the three-dimensional equivalent of the Poisson equation (2.6),

$$\frac{\partial^2 \hat{\psi}_{\mathbf{k}}}{\partial y^2} - |\mathbf{k}|^2 \hat{\psi}_{\mathbf{k}} = -\hat{\omega}_{\mathbf{k}}(y), \quad (2.7)$$

where \mathbf{k} is the wave vector in the (x, z) plane. As $y \rightarrow \pm\infty$, $\hat{\omega}_{\mathbf{k}} \rightarrow 0$, and the solutions of (2.7) approach those of the homogeneous equation, $\exp(\pm|\mathbf{k}|y)$, which decay much slower than the right-hand side (Fig. 1). Since the exact behavior of the solution at large distances is known, it can be included explicitly in the numerical scheme, so that the remaining part of the solution decays fast enough for the use of standard expansions. To do this we note that we can construct, within $y \in (-L, L)$, solutions to Eq. (2.7) which satisfy any consistent set of boundary conditions and, in particular, solutions which are periodic in y with period $2L$. Then, since any solution of (2.7) can be written as the sum of any particular solution and of some linear combination of the homogeneous exponentials, we should be

able to find an expansion in terms of a Fourier part and of the two known exponentials.

Note that the mode $|\mathbf{k}| = 0$ is special and has to be treated differently. The homogeneous solutions for it are $\hat{\psi}_0 = 1$ and $\hat{\psi}_0 = y$ and have to be matched to boundary conditions that are related to the velocities in the free streams.

A complete code to compute the evolution of Eqs. (2.5)–(2.6) would contain a time marching step for Eq. (2.5), followed by a Poisson solver for Eq. (2.6). The latter can be formulated as a set of individual Helmholtz equations (2.7), which are solved by the mixed scheme outlined in the previous paragraph and which will be explained in detail in Section 4.

There is another Poisson solver associated to each time marching step. It is usually convenient to treat the viscous term implicitly, and each step has the structure

$$\left(1 - \frac{\Delta t}{\text{Re}} \nabla^2\right) \omega^{(n+1)} = \omega^{(n)} + \Delta t h_{\omega}^{(n)}, \quad (2.8)$$

with boundary conditions on $\omega(y = \pm\infty)$. However, the exponential factors associated to the Helmholtz operator in the left-hand side are $\pm(\text{Re}/\Delta t)^{1/2}$ and are usually so large ($\sim 10^3$) that they can be assumed to decay completely at $\pm L$. As a consequence, the periodic solution to the Helmholtz equation can be taken to be the full solution, and no exponential correction is needed. Note that this assumption is needed if the compactness of the vorticity is going to be maintained in time. In practice, even for the solutions of Eq. (2.7), the exponents are large enough that the tails are negligible for all but the few lowest x and z harmonics.

3. SEMI-INFINITE DOMAINS

The same scheme can be used in semi-infinite domains. In this case, the vorticities are expanded in terms of Chebyshev polynomials to accommodate the effect of the boundary at $y = 0$. Note that there is no special discontinuity of the variables at $y = L$, which represents infinity and that there is no need for a singular grid deformation at that point. As a consequence, only half of the Chebyshev collocation points are needed, located at

$$y_j/L = 1 - \cos(j\pi/2N), \quad j = 0 \dots N. \quad (3.1)$$

This saves approximately one half of the storage space and of the transformation cost and avoids an unnecessary accumulation of grid points near the boundary representing infinity.

This expansion is equivalent to using a full Chebyshev expansion of the variables in a larger domain $y \in (0, 2L)$, keeping only half of the spectral components. There are many possible choices of which components to select, but the two obvious ones are to use only even or odd polynomials. The implication is that the function being expanded is either even or odd about

TABLE I

Type and Spatial Representation along Each Axis for Different Variables in a Three-Dimensional Boundary Layer Code

Variable	Type	Representation
u_1, u_3	Even	F-F-E
u_2	Odd	F-F-E
ω_1, ω_3	Odd	F-F-C
ω_2	Even	F-F-C

Note. F \equiv Fourier, E \equiv exponential, C \equiv Chebyshev.

the ‘‘mid-point’’ at $y = L$ and that either its derivative or its value should vanish at that point. While the choice is immaterial for variables which have actually decayed to zero at $y = L$, it has some effect in practical cases. Thus, if a variable ω is being expanded in even Chebyshev polynomials, the expansion assumes that $\partial\omega/\partial y(L) = 0$, and any non-zero value of the actual derivative at that boundary appears as a discontinuity which induces a proportional error. Conversely, the errors in odd expansions are proportional to the value of the variable itself at the boundary. The basic assumption of our numerical scheme is that both errors are negligible.

In problems involving several variables, the choice of even or odd expansions is not completely free. Odd operators, such as $\partial/\partial y$, change the parity of a function so that, if an even expansion is used for ω , the one for $\partial\omega/\partial y$ should be odd and vice versa. As an example, Table I contains the parity of the expansions used for the different velocity and vorticity components in our three-dimensional turbulent boundary layer code. Note that the Helmholtz equation (2.7) involves only variables with uniform parity.

The last column of Table I contains the ‘‘most spectral’’ representation that is possible for each variable. Thus, while the vorticities can be expressed in fully spectral mode, the velocities can only be expanded in the two homogeneous directions and need extra exponential terms in the y -direction.

The exponential terms are always associated with the solution of the Poisson step and, as is shown in the next section, can be computed algebraically when the boundary conditions are imposed on the velocities. This step can be considered as part of the inverse (Fourier to physical) y -transform for the velocities and can be achieved in a number of operations that is only proportional to the number of points. The performance of the FFT is not degraded.

The pseudospectral step for the computation of the nonlinear terms proceeds as usual. The convective terms have the general form $\mathbf{v} \cdot \nabla \omega$. Both factors are transformed to physical representation, including the exponential terms for \mathbf{v} , and multiplied. Since $\nabla \omega$ is compact at large y , so is the product, and the

full nonlinear term can be transformed back to Fourier or to Chebyshev by extending it periodically across the edge of the computational domain. There is no need for extra functions in that expansion. In contrast, there is no obvious way of transforming the velocity back into spectral representation, but none is needed.

There exists an extra difficulty when computing the Navier–Stokes equations in wall-bounded flows using a vorticity formulation. The physical boundary conditions cannot be imposed directly on the Helmholtz equations, because the former are expressed in terms of velocities, while the latter involve variables that contain derivatives of at least one degree higher than those in the boundary conditions. In the 2D formulation, Eq. (2.5) requires a boundary condition for ω at the wall while the physical boundary condition is that $u = v = 0$. In the 3D case the same problem exists, because one of the Helmholtz equations is written for $\nabla^2 v$, where v is the velocity component normal to the wall, while the boundary conditions at the wall involve v and its first derivative.

The approach that we have followed takes advantage of the linear character of the Helmholtz equation to split the problem in two Dirichlet sub-problems. The final solution is built as a linear combination of the solutions to these problems, in such a way that the desired boundary conditions are satisfied. A similar procedure for a 3D channel can be found in [2]. During this process the exponential terms have to be taken into account when determining the constants for the final linear combination.

For the sake of clarity, we develop in some detail the problem of a two-dimensional time-developing boundary layer on a flat plate. The plate is located at $y = 0$, and Eqs. (2.5)–(2.6) have to be solved with the boundary conditions

$$\begin{aligned} \psi = 0, \quad \psi_y = 0 \quad \text{at } y = 0, \\ \psi_y \rightarrow 1, \quad \omega \rightarrow 0 \quad \text{as } y \rightarrow \infty. \end{aligned} \quad (3.2)$$

Once the variables are expanded in Fourier series in x , each component satisfies two coupled equations, one of which is equivalent to the implicit step (2.8), while the other is the familiar Poisson equation (2.7), which can be written as

$$\begin{aligned} \hat{\omega}'' - k_\tau^2 \hat{\omega} &= -\hat{g}, \\ \hat{\psi}'' - k^2 \hat{\psi} &= -\hat{\omega}, \end{aligned} \quad (3.3)$$

where $\hat{\omega}$ and $\hat{\psi}$ are the Fourier components for wavenumber k , the modified wavenumber k_τ is defined by $k_\tau^2 = k^2 + \text{Re}/\Delta t$, and \hat{g} is related to the right-hand side of Eq. (2.8). These equations satisfy boundary conditions similar to (3.2) in which the condition at infinity is substituted by

$$\hat{\psi}'(\infty) = 0 \quad (3.4)$$

for all wavenumbers except $k = 0$. The trouble is that no

boundary condition is available for the vorticity at $y = 0$, while there is one condition too many for the stream function. To circumvent this difficulty two problems are solved for each wave number. In the first one, Eqs. (3.3) are solved with homogeneous boundary conditions

$$\hat{\psi}_1(0) = \hat{\psi}'_1(\infty) = \hat{\omega}_1(\infty) = 0, \quad \hat{\omega}_1(0) = 0. \quad (3.5)$$

In the second problem, the first equation in (3.3) is substituted by its homogeneous equivalent, while the last boundary condition in (3.5) is made inhomogeneous,

$$\hat{\omega}_2'' - k^2 \hat{\omega}_2 = 0, \quad \omega_2(0) = 1. \quad (3.6)$$

It is clear that any linear combination of the form

$$\hat{\psi} = \hat{\psi}_1 + c \hat{\psi}_2 \quad (3.7)$$

satisfies Eqs. (3.3) and that c can be determined so that $\hat{\psi}'(0) = 0$.

The three-dimensional problem is treated in the same way. In this case the variables of integration are the vorticity normal to the wall, ω_y , and the Laplacian of the component of the velocity, $\nabla^2 v$. The equation for the normal vorticity does not present any special problem, and $\omega_y(0) = 0$ is imposed at the wall, while the equation for $\nabla^2 v$ needs that one condition, $v_y(0) = 0$, be enforced in two steps.

This difficulty is common to all vorticity formulations of wall bounded flows and has been addressed in terms very similar to those described here in [10, 11] for two-dimensional flows and in [2] for a three-dimensional one. Of interest in the present context is that the introduction of exponential terms does not interfere with those well-known procedures.

4. SPECTRAL SOLUTION OF A HELMHOLTZ EQUATION FOR COMPACT FORCING

We will discuss in detail the case of a double-infinite domain. The semi-infinite Chebyshev case is entirely similar. The problem is to solve the one-dimensional Helmholtz equation

$$\frac{\partial^2 f}{\partial^2 y} - \gamma^2 f = \omega(y), \quad y \in (-\infty, \infty), \quad (4.1)$$

with the boundary conditions $f \rightarrow 0$ as $y \rightarrow \pm\infty$, where the right-hand side, $\omega(y)$, is assumed to be essentially compact,

$$\omega(y) \text{ negligible for } |y| \geq L_1 < L. \quad (4.2)$$

Because of compactness, $\omega(y)$ can be expanded in Fourier series in the interval $y \in (-L, L)$,

$$\omega(y) = \sum \hat{\omega}_k e^{ik\pi y/L}. \quad (4.3)$$

Since the expansion assumes a periodic extension of ω , there is in practice an effective discontinuity at the boundary, $\Delta\omega = \omega(L) - \omega(-L)$, which induces an algebraic tail for the Fourier coefficients, $\Delta\hat{\omega}_k \sim \Delta\omega/k$. The discontinuity is assumed to be small and, in the case of the vorticities, typically decays as $\exp(-L^2)$, or as an exponential with a large exponential factor. On the other hand, the same expansion cannot be used for the solution, f , since it contains slowly decaying terms of the form $\exp(\pm\gamma y)$, which cannot be neglected unless $\gamma L \gg 1$.

Consider the solution to Eq. (4.1) which satisfies periodicity conditions in the interval $(-L, L)$. It can be written as a Fourier series with coefficients $\hat{f}_{L,k} = -\hat{\omega}_k/(\gamma^2 + k^2)$ and, since it can only differ from the full solution by a solution of the homogeneous equation, we can write

$$f(y) = f_L(y) + a_+ e^{\gamma(y-L)} + a_- e^{-\gamma(y+L)}, \quad (4.5)$$

valid only inside $y \in (-L, L)$.

Finally, because of the compactness of the right-hand side, the solution to Eq. (4.1) has to be of the form

$$f(y) \sim e^{-\gamma y}, \quad \text{for } y \geq L_1, \quad (4.6)$$

so that $f'(L) \equiv df/dy(L) = -\gamma f(L)$. This, together with a similar relation at $y = -L$, is enough to determine the two unknown coefficients in (4.5),

$$a_{\pm} = -(\gamma f_L(L) \pm f'_L(L))/2\gamma. \quad (4.7)$$

There is an equivalent expression for the semi-infinite case, in which the functions are expanded in even or odd Chebyshev polynomials.

The computational cost of the exponential correction in (4.5)–(4.7) is just the calculation of $f_L(L)$ and $f'_L(L)$, and the addition of the exponentials themselves at the collocation points. All of these are $O(N)$ operations that do not add an appreciable overhead to the cost of the inverse FFT. The exponential functions, which are the only expensive part of the operation, are pre-computed and stored, as is done with the sine tables for the FFT.

5. THE APPROXIMATION ERROR

Consider the Fourier case. The procedure outlined above provides an exact solution of Eq. (4.1) with a right-hand side that is equal to ω in $(-L, L)$ and which vanishes identically outside. The errors in the computation arise from this neglect of the right-hand side outside the interval and from the approximation of ω and f_L by a truncated spectral series. Most of those errors are common to other spectral methods, and they will not be discussed here (see, e.g., [12]). The only ones which are




FIG. 2. Point-wise approximation error in the solution of $d^2f/dy^2 - 9f = e^{-y^2}$ with $f(\pm\infty) \rightarrow 0$: (a) As a function of resolution. Different curves correspond to different positions of the outer boundary. (b) As a function of the width of the domain for fixed resolution $N/L \approx 10$. The point for $L = 8$ has been excluded because of round-off errors. The dashed line is the predicted behavior $\exp(-L^2)$.

new to our scheme are those due to the finite interval and the ones associated to the computation of the coefficients in front of the exponentials in (4.5).

If the right-hand side decays fast as $y \rightarrow \pm L$, the error due to the finite interval should also be small. Assuming the most severe case of a smooth ω with a characteristic length scale of L , it follows from a simple estimation, using the Green's functions of (4.1), that the magnitude of this error in f is at most of order $\omega(L)$, which is negligible under the assumption that $\omega(L)$ is small.

With respect to the second source of error, note that, since the exponentials themselves are scaled so as to be at most $O(1)$, the point-wise error induced on $f(y)$ is at most of the same order as the error in a_{\pm} which, because of (4.7), is itself of the same order as the error in f_L . Also, since the Fourier expansion of f_L is obtained from a smoothing operation on the expansion of ω , the error in the former is at most proportional to that in the expansion of the latter. We have to consider its variation both with the number of harmonics N , and with the size of the numerical domain L . As we change N , and as long as there is no discontinuity in ω or in its derivatives, the error decays faster than any power of the resolution N/L [12, Section 2]. In practice, this usually means an exponential decrease of the error with increasing N/L , which is maintained until it is dominated by some other source. In our case, another important contribution to the error is the discontinuity discussed in the previous section, due to the lack of periodicity of ω at $y = \pm L$. This contributes a slow algebraic tail to the spectrum, whose magnitude is $O(\omega(L))$, and which eventually becomes dominant at high resolutions.

This is confirmed by Fig. 2a, which shows the maximum

pointwise approximation error in the solution of a Helmholtz equation whose right-hand side is a gaussian, solved numerically using the present scheme. At low resolutions the error is dominated by the truncation of the Fourier series representing the right-hand side and decreases exponentially with the square of the resolution $(N/L)^2$. This is the sloping line to the left of the figure. As the resolution increases and this source of error becomes smaller, it is eventually swamped by the errors due to the neglect of the tails of the forcing function outside the interval and to the discontinuity in the derivative of the gaussian as it is extended periodically beyond $y = L$. These errors depend almost exclusively on L and decrease only algebraically with the resolution. These are the approximately horizontal segments leaving the main line towards the right.

The fact that the resolution error decreases faster than exponentially is a peculiarity of the gaussian forcing term, whose Fourier coefficients decay as $\exp(-k^2)$. In practice, the forcing function is not usually as smooth and the error decays slower. The gaussian model is convenient numerically because the existence of an analytical solution to the Helmholtz equation simplifies the computation of the errors.

Figure 2b displays the approximation error as a function of the size of the domain, for constant resolution $N/L = 10$. In this range, the error is dominated by the neglect of the tails of the forcing function and, except for algebraic factors, should be proportional to $\exp(-L^2)$. This behavior is confirmed by the data in the figure. Note that the extremely small errors displayed here are only possible through the use of 128-bit floating point arithmetic.

A consequence of these estimates is that, to keep errors small, L has to be chosen only large enough for the right-hand side




FIG. 3. Decay of the point-wise approximation error with the width of the computational box, for the same equation as in Fig. 2 and for different numbers of harmonics: (a) Exponentials corrections are used in the approximation. (b) Exponentials are not used. Dashed line is the theoretical error due to the neglect of the exponential tail.

of the Helmholtz equations to be adequately represented, in the sense that the discontinuity error at the boundary is less than the spectral truncation error.

This is shown in Fig. 3a which displays a set of practical design curves for the optimization of the numerical algorithm. The problem is how to choose the optimum size of the computational domain for a given fixed number of harmonics. For very short boxes the numerical resolution N/L is high and the error is dominated by the edge effects, which decrease as $\exp(-L^2)$. This is the parabola on the left-hand side of the figure, which is common to all the curves. As the domain is made longer, the mesh becomes coarser and the normal truncation error becomes dominant. In between there is an optimum length for which the error is minimum. For the present example this occurs when $\exp[-(N/L)^2] \approx \exp(-L^2)$, or $L \approx N^{1/2}$. The minimum error is then proportional to $\exp(-N)$.

Figure 3b shows the same curves for a scheme in which the exponential corrections are not applied to the solution. For large boxes, in which the error is controlled by the coarse meshing, both sets are identical. At high resolutions, however, the error is dominated by the neglect of the exponential tails and is always much larger than before. In fact, the truncation error is now proportional to $\exp(-\gamma L)$, and the optimum error occurs where $\gamma L \approx (N/L)^2$ and, for $\gamma = 3$, is proportional to $\exp(-2.1N^{2/3})$.

The benefits of present method are exhibited even for small boxes and few harmonics and not merely as an asymptotic trend. In the figure, for $N = 16$, the use of exponential correction decreases the optimum error by three orders of magnitude, compared to the conventional method. The advantage of the method depends on the magnitude of γL in the Helmholtz

equation, but in the context of the use of spectral methods in fluid mechanics, this is typically of order unity.

The rule of thumb is that the value of the vorticity at the outer boundary should be kept smaller than the amplitude of the Gibbs' ripples induced by the limited resolution. This criterion is easily met in turbulence simulations, where resolution tends to be marginal. On the other hand, the value of the velocity at the outer boundary is not relevant to the approximation, since it is handled exactly by the exponentials. The consequence is that it is possible to choose the boundary of the numerical domain very close to the edge of vorticity distribution, without regard to the irrotational perturbations. In fact, we routinely ran simulations in which vorticity blobs passed very near the outer boundary. This increased the danger that occasional vorticity excursions might impinge on the edge of the domain, violating the basic assumptions of the scheme. The consequences of these events are explored in the next section.

6. ROBUSTNESS

An important characteristic of any code is its ability to respond "gracefully" to off-design situations. In the present case this means cases in which, for some reason, a vortex touches the outer boundary where the vorticity is assumed to have decayed to zero. Although there is no way, within the assumptions of the method, for the code to respond correctly to this situation, it is important to make sure that it recovers from the error with minimum permanent effects.

Such extreme excursions of the vorticity distributions are well known in turbulent flows. In two-dimensional situations,

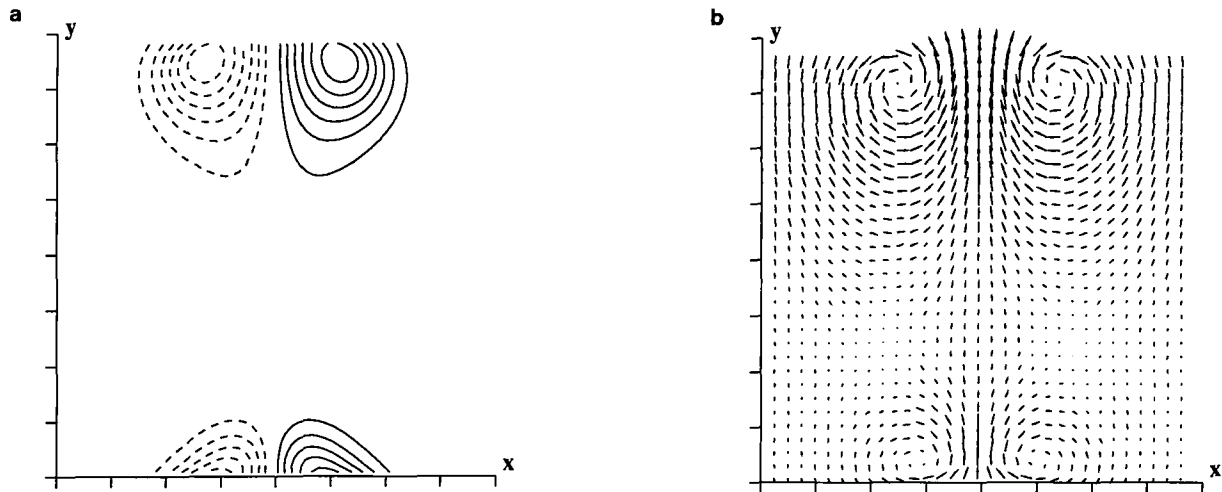


FIG. 4. Vorticity (a) and velocity (b) distributions induced by a vortex pair as it hits the upper boundary of the computational domain. See text for discussion.

especially wakes, they take the form of bound pairs of counterrotating vortices which fly away from the turbulent region as almost independent entities [13]. In three-dimensional situations they are not as well documented in the literature, but we have observed them as hairpin vortices which occasionally extend far above the average location of the edge of the boundary layer. In both cases, the phenomenon appears seldom and seems to be of little dynamical significance. The stray vorticity is quickly dissipated by viscosity and has little effect on the body of the flow. It may therefore be acceptable to treat these occurrences numerically in a way that is not exactly right, as long as it does not interfere strongly with the accuracy of the rest of the computation. In any case, the problem is shared by any method that does not compute the flow to very long distances with uniform resolution.

To test the response of our method to such situations, we integrated the two-dimensional equations (2.5)–(2.6) in a box, periodic in x and extending to $y \in (-\infty, \infty)$. The velocities were assumed to decay at both infinities. The initial condition was chosen as a counterrotating vortex pair (or rather as an x -periodic array of such pairs) arranged so as to move in the y direction. The purpose was to observe the behavior of the code as the pair hit the upper boundary. One frame of the vorticity and velocity fields during the interaction is shown in Fig. 4. The pair crosses cleanly the boundary and reappears on the lower end of the box. The velocity field, however, is incorrect. Each of the two parts of the vortex pair, the one near the upper boundary and the one near the lower one, generates its own velocity field, with streamlines that close on themselves. The vorticity remains bounded and localized.

All this is as it should. The vorticity is periodic, because it is expressed as a Fourier series and remains so during the interaction with the boundary. The vortices near the lower boundary are the periodic image of the vortex pair near the upper

one. The velocity is not periodic, because of the exponential correction. In fact, the stream function generated by Eq. (4.5) is constructed so as to be the one induced by the vorticity inside the computational box, with strictly irrotational flow outside. Therefore, when the vorticity hits the boundary, it behaves as defined by its spectral expansion, with no special continuity errors, but any vorticity that exits the box ceases to have an effect and disappears from the calculation in the vicinity of that boundary. The flow field remains smooth during the interaction, although the velocity is clearly not the one corresponding to the actual vorticity distribution.

The behavior of a Chebyshev expansion is similar. As a vortex approaches the upper boundary, it meets another one from the other half of the Chebyshev domain, whose sign is either equal or opposite, depending on the parity of the expansion. In either case, the induced velocities are only those of the vortices remaining inside the computational box, and they are generally of the correct sign to keep the vortex going out of the domain. In fact, as noted above, several collisions of this kind were observed during the simulation of our three-dimensional boundary layer, and some of them were studied in detail. In all cases the vortices left the domain smoothly.

7. APPLICATION EXPERIENCE

We have implemented this scheme in two codes for the direct numerical simulation of turbulence. The first one is an extension of an older two-dimensional mixing layer, including the passive advection and mixing of a scalar [14]. The purpose of the new set of simulations was to extend the study of the geometry of the mixing interface to a wider range of Peclet numbers than was possible with the older finite differences code in serial machines. The need to recast the code in a parallel computer (a 128-node Intel hypercube at NASA Ames) led us to look

for a simpler implementation. The present scheme proved ideal, since efficient FFTs were available for the Intel machine and the new treatment of the boundary conditions at infinity was straightforward enough as not to interfere with the intrinsic complication of parallel programming. The parallelization strategy and the results are discussed elsewhere [7]. We will only be interested here in the numerical performance.

The code was run on square boxes with uniform collocation grids for a double Fourier transform. The largest grid implemented was 2048×2048 , before dealiasing using the $2/3$ rule. The initial conditions gave rise to four primary vortices that underwent two pairings. During the second pairing, the resulting vorticity structure grew until it filled about 60% of the width of the computational domain. At that moment some ripples appeared near the outer boundaries, but they never exceeded 1–2%. They disappeared shortly after the structure collapsed again. To check whether those ripples were due to limited resolution effects, or to the exponential velocity corrections, the same initial conditions were run without the exponential terms, effectively substituting the problem by a periodic array of parallel shear layers. Although the details of the flow were modified, the magnitude of the ripples stayed essentially unchanged. From the comparison of the two codes it was confirmed that the computational overhead of the exponential corrections was negligible. Finally, the evolution of the integral quantities of the layer were compared with those obtained by the code in [14], with entirely satisfactory agreement.

The second experiment was the implementation of a three-dimensional boundary layer code, periodic in the two directions parallel to the wall. The results are described in [8]. Three-dimensional codes are less forgiving than two-dimensional ones, since they have intrinsic vorticity amplification mechanisms, and numerical errors can lead to blow up by mimicking the generation of real vortices. Part of our interest was to test whether any such numerical instability mechanism might be present. None was found.

The structure of this code is reflected in the discussion of Chebyshev schemes in the previous sections. Typical grids were $36 \times 97 \times 36$ before dealiasing in x and z . No dealiasing was applied to the Chebyshev transform. The code was used to study the low Reynolds number regime near the relaminarization limit. The outer boundary was adjusted roughly at 10 displacement thickness, so that the classical boundary layer thickness filled most of the computational domain. The behavior of the r.m.s. vorticity and velocity fluctuations is shown in Fig. 5, which displays the whole computational domain. Vorticity fluctuations fill more than two thirds of the domain, but decay very fast above that limit. Velocity, on the other hand, decays much slower, and its exponential rate matches well the one predicted from the Helmholtz equation corresponding to the lowest Fourier wavenumber in the computational box. Note that, even at the edge of the domain, the exponential decay is smooth.




FIG. 5. Behavior at large distances from the wall of the r.m.s. velocity and vorticity fluctuations in a three-dimensional boundary layer. The straight line in the velocity plot is the predicted exponential rate $u' \sim \exp(-2\pi y/L_x)$. Note the fast decay of the vorticity and the exponential behavior of the velocities up to the edge of the computational domain. $Re_\nu = 302$: dashes, v'_i ; solids, ω'_i ; grids, $24 \times 97 \times 23$ after dealiasing.

8. CONCLUSIONS

We have presented a modification of the standard Fourier or Chebyshev spectral collocation method to include the handling of boundary conditions extending to infinity along one direction. The scheme is based on the assumption that the flow becomes irrotational at large distances and that the computational domain contains all the vorticity. The irrotational velocity perturbations are treated analytically. We have shown that the resulting overhead, compared to the fully periodic or bounded problems, is negligible, and that the approximation errors are comparable in both cases. We have also proved that the computational boxes can be chosen quite close to the edge of the rotational region without ill effect and that the occasional failures in the basic hypothesis, when a vortex touches the numerical boundary, are handled smoothly by the scheme. Finally we have discussed briefly two applications to flows of practical interest, in which no unexpected numerical problems were identified.

ACKNOWLEDGMENTS

This work was supported in part by the Hermes program of the ESA, under Contract AMD-RDANE 3/88. Some of it, including the parallel implementation of the two dimensional code, was carried out during an extended stay of J.J. in the Center for Turbulent Research at NASA Ames. We are grateful to the NAS computational group for the free use of their parallel machine. The adaptation to a boundary layer was done within the ELFIN program of Brite

Euram. We are indebted to R. Moser for many useful discussions and to C. Martel for coding an early version of the mixing layer code.

REFERENCES

1. R. S. Rogallo and P. Moin, *Annu. Rev. Fluid Mech.* **16**, 99 (1984).
2. J. Kim, P. Moin, and R. Moser, *J. Fluid. Mech.* **177**, 133 (1987).
3. E. Laurien and L. Kleiser, *J. Fluid. Mech.* **199**, 403 (1989).
4. M. R. Malik, T. A. Zang, and M. Y. Hussaini, *J. Comput. Phys.* **61**, 64 (1985).
5. P. R. Spalart, NASA Technical Memorandum 88222, 1986, (unpublished).
6. P. R. Spalart, R. D. Moser, and M. M. Rogers, *J. Comput. Phys.* **96**, 297 (1991).
7. J. C. Agüí and J. Jiménez, Technical Note ETSIA/ME-918, School of Aeronautics, Madrid; *Int. Symp. Parallel CFD, Stuttgart, June 10–12, 1991*.
8. R. Corral and J. Jiménez, in *AGARD Symp. Transition and Turbulence, Chania, Crete, April 18–21, 1994*, p. 21.1.
9. M. Deville, L. Kleiser, and F. Montigny-Rannou, *Int. J. Numer. Methods Fluids* **4**, 1149 (1984).
10. J. Jiménez, *J. Fluid Mech.* **218**, 265 (1990).
11. D. Gottlieb and S. A. Orszag, *Numerical Analysis of Spectral Methods: Theory and Applications*, Regional Conference Series in Applied Mathematics, Vol. 26 (SIAM, Philadelphia, 1977).
12. Y. Couder and C. Basdevant, *J. Fluid Mech.* **173**, 255 (1986).
13. J. Jiménez and C. Martel, *Phys. Fluids A* **3**, 1261 (1991).