

A Primer on Direct Numerical Simulation of Turbulence – Methods, Procedures and Guidelines

Gary N. Coleman and Richard D. Sandberg

*Aerodynamics & Flight Mechanics Research Group
School of Engineering Sciences, University of Southampton, SO17 1BJ, UK*

Technical Report AFM-09/01b (March 2010)

1. INTRODUCTION

Direct Numerical Simulation (DNS) is the branch of CFD devoted to high-fidelity solution of turbulent flows. DNS differs from conventional CFD in that the turbulence is explicitly resolved, rather than modelled by a Reynolds-averaged Navier-Stokes (RANS) closure. It differs from large-eddy simulation (LES) in that all scales, including the very smallest ones, are captured, removing the need for a subgrid-scale model. DNS can thus be viewed as a numerical experiment producing a series of non-empirical solutions, from first principles, for a virtual turbulent flow (see Figure 1). Its great strength is the ability to provide complete knowledge, unaffected by approximations, at all points within the flow, at all times within the simulation period. DNS is therefore ideal for addressing basic research questions regarding turbulence physics and modelling. This ability, however, comes at a high price, which prevents DNS from being used as a general-purpose design tool.

The defining characteristics of DNS stem from the distinctive characteristics of turbulence. Because turbulence is inherently unsteady and three-dimensional, DNS requires time-dependent calculations within a three-dimensional domain.* These two features are shared with LES (and therefore LES/RANS hybrid strategies such as detached eddy simulation (DES)). The unique feature of DNS is associated with the manner in which turbulence is affected by viscosity. This is responsible for the two chief drawbacks of DNS – its extreme computational cost, and severe limitation on the maximum Reynolds number that can be considered.

*DNS can also be used to study problems involving, for example, laminar-to-turbulent transition or two-dimensional geophysical phenomena. We limit our attention here to cases where the turbulence is fully developed and contains vorticity fluctuations in all three directions.

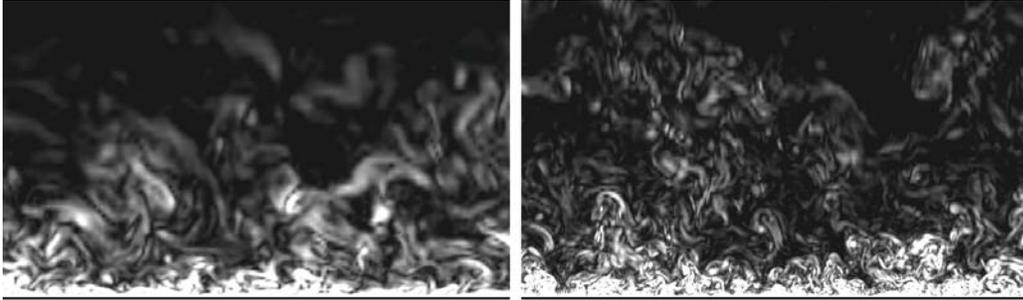


Figure 1. Vorticity magnitude contours from spectral Ekman-layer DNS at two Reynolds numbers Re_E (Spalart *et al.* 2008). Vertical planes shown are normal to the free-stream velocity, and cover the same area measured in local boundary-layer thickness. Reynolds number of right-hand-side flow ($Re_E = 2828$) is twice that of left-hand-side flow ($Re_E = 1414$). Courtesy of Dr. R. Johnstone, University of Southampton.

2. THE REYNOLDS NUMBER CONSTRAINT

Turbulence contains eddies with a wide range of sizes. These eddies interact with each other in a non-linear fashion through their induced velocity fields, changing the orientation and shape of their neighbours. As first described by Richardson (1922) and quantified by Kolmogorov (1941), the net effect of this change-of-shape (i.e. straining) process is to ‘cascade’ kinetic energy from the largest to the smallest scales of the turbulence. As a result, the largest eddies are the most energetic, and their size, shape and speed are set by the details of the flow configuration, and not directly affected by the viscosity of the fluid. The size of the smallest eddies, on the other hand, is determined both by how much energy enters the cascade at the large scales and by the viscosity. The primary role of viscosity is to define the scale at which the energy is dissipated. The Reynolds number of the flow thus determines how small the smallest scales are, relative to the largest eddies.

This behaviour, known as *Reynolds-number similarity*,[†] can be observed in Figure 1, which presents DNS results from the same boundary-layer flow at two Reynolds numbers that differ by a factor of two. This illustrates the challenge faced by DNS, which must use a domain large enough to comfortably include the largest naturally-occurring eddies while using a grid fine enough to fully resolve the dissipation scales. Using today’s most capable computers, this can only be done for Reynolds numbers orders of magnitude smaller than found, for example, in full-scale aeronautical flows.

The manner in which the cost of a DNS scales with Reynolds number can be precisely determined for the case of homogeneous isotropic turbulence, for which the size and the speed of the largest eddies can be characterised by a single length and time scale, ℓ_{LE} and t_{LE} , respectively. Because the largest contribution to the turbulence kinetic energy (per unit

[†]Reynolds-number similarity is the reason film makers can use small-scale-model special effects, since they yield turbulent flows whose largest eddies are roughly correct. That these special effects do not look quite right – at least to the discerning fluid-dynamicist! – compared to the full-scale flow, is a symptom of the truncated range of scales in the scale-model flow, associated with the vastly different Reynolds numbers.

mass) $k_T = \frac{1}{2}q^2$ is made by the largest eddies, their characteristic velocity is proportional to $q = \sqrt{2k_T}$, and the rate (i.e. flux) of energy $\dot{\epsilon}_{LE}$ leaving the largest eddies scales with q^2/t_{LE} . And since, for this flow, $\dot{\epsilon}_{LE}$ is proportional to ϵ , the rate at which energy is currently being dissipated at the smallest scales, it is reasonable to assume that t_{LE} goes like q^2/ϵ , and thus (since $q \sim \ell_{LE}/t_{LE}$) that ℓ_{LE} is proportional to q^3/ϵ . Consequently, the intrinsic Reynolds number of the turbulence $Re_T = q\ell/\nu$ is $q^4/\nu\epsilon$. Following Kolmogorov, assuming the smallest scales of turbulence are universal and isotropic, and thus depend solely on ϵ and the kinematic viscosity ν , leads on dimensional grounds to definition of the turbulence microscale $\eta = (\nu^3/\epsilon)^{1/4}$. The number of grid points N in each direction required for DNS of homogeneous isotropic turbulence will then be of the order $\ell_{LE}/\eta = (q^3/\epsilon)(\epsilon^{1/4}/\nu^{3/4}) = (q^4/\nu\epsilon)^{3/4} = Re_T^{3/4}$. The total number of grid points, $N^3 \sim \mathcal{O}[(\ell_{LE}/\eta)^3]$, will then scale with $Re_T^{9/4}$. Factoring in the change in timestep (required for example to maintain the same CFL number), the total computational effort is $\mathcal{O}[Re_T^3]$. When viewed in terms of operations per grid point, the picture becomes slightly worse: if an algorithm requiring $N \log N$ operations in each spatial direction were used (a reasonable estimate for an efficient spectral method), the total operational count would be $\mathcal{O}[Re_T^3(\frac{3}{4} \log Re_T)^3]$. Therefore, doubling the Reynolds number from a currently attainable value would increase the computational cost (i.e. CPU time, memory) by roughly a factor of 11! Assuming that the current trend of doubling computing power every 18 months continues and numerical algorithms scale perfectly for even larger grid counts, this implies that for this flow the Reynolds number can be doubled only every five to six years. Although the exact Reynolds-number scaling of the cost will vary from flow to flow, and will not always be as stringent as for homogeneous isotropic turbulence,[‡] the general point still holds that we are currently far away, and will remain so for the foreseeable future, from being able to perform DNS for Reynolds numbers typical of full-scale engineering and especially geophysical applications.

3. NUMERICAL CONSIDERATIONS

The obligation of having to resolve all spatial and temporal scales of the turbulence requires that numerical errors be monitored and controlled. As a result, DNS has historically not used commercial CFD packages, but specially-written codes, optimised for the flow-types of interest. We now consider the issues that must be addressed by authors and users of DNS codes. The need for DNS algorithms to be *efficient* – that is, to have a high ratio of accuracy to computational cost – is particularly important.

3.1. Spatial Discretisation

Central to the success of any DNS is the ability to faithfully reproduce all spatial variations of the dependent variables. There are a number of strategies that DNS codes have employed to do this, including finite-volume, finite-element, discrete-vortex and B-spline methods. However, for reasons that shall soon be clear, DNS has to this point been dominated by

[‡]For example, for a boundary-layer simulation only a fraction of the flow will be turbulent and thus adhere to the Reynolds-number scaling. An increase in Reynolds number would not affect the resolution in the freestream region.

spectral and finite-difference schemes. We shall briefly consider each in turn.

3.1.1. Spectral Methods The Reynolds-number constraint and the need to minimise numerical errors using the available high-performance computing (HPC) resources naturally led the first practitioners of DNS (e.g. Orszag & Patterson 1972; Kim *et al.* 1987) to chose spectral methods to account for spatial variations (see Gottlieb & Orszag 1977, Hussaini & Zang 1987, or Canuto *et al.* 1988 for an overview). Results from a recent spectral DNS are shown in Figure 1. Spectral methods approximate the flow variables as linear combinations ('expansions') of global basis functions that involve complex-exponential or orthogonal-polynomial eigensolutions $\phi_n(x)$ ($n = 0, 1, 2, \dots$) of an appropriate Sturm-Liouville problem over the interval $x_1 \leq x \leq x_2$. They thus satisfy the orthogonality relationship

$$\int_{x_1}^{x_2} w(x)\phi_n(x)\phi_m(x)dx = A_n\delta_{nm}, \quad (1)$$

where A_n is a positive n -dependent coefficient and δ_{nm} is the Kronecker delta (i.e. unity when $n = m$ and zero otherwise). The weight function $w(x) \geq 0$ defines the resolution characteristics of the basis functions, since in regions where $w(x)$ is largest, $\phi_n(x)$ must have more zero crossings in order for the weighted product of $\phi_n(x)$ and $\phi_m(x)$ to integrate to zero. The spatial-discretisation procedure will be demonstrated here as if the Navier-Stokes equations $\mathcal{L}(\mathbf{u}) = \mathbf{0}$ involve only time t and one spatial direction x , which might be a cartesian or spherical/polar coordinate. * Each dependent variable $u(x)$ is replaced, over the domain $x_1 \leq x \leq x_2$, by

$$u(x) \leftarrow u_M(x, t) = \sum_{n=0}^{M-1} \alpha_n(t) \psi_n(x), \quad (2)$$

where α_n are the M expansion coefficients, and, for example, $\psi_n(x) = g(x)\phi_n(x)$ or $h(x)d\phi_n/dx$, where $g(x)$ and $h(x)$ might be low-order polynomials that ensure each $\psi_n(x)$ satisfies the appropriate boundary conditions at $x = x_1$ and x_2 . When $u(x)$ can be assumed to vary periodically over the distance Λ , it is natural to use a Fourier-series representation, with $\psi_n(x) = \phi_n(x) = \exp(ik_nx)$, such that

$$u_M(x, t) = \sum_{n=-M/2}^{+M/2} \alpha_n(t) e^{ik_nx}, \quad (3)$$

where $i = \sqrt{-1}$, $k_n = 2\pi n/\Lambda$ and the weight function $w(x) = 1$ (indicating uniformly spaced resolution over $\Lambda = x_2 - x_1$). † For other (wall-bounded) flows, $\phi_n(x)$ might be Chebyshev or Jacobi polynomials, which posses $w(x)$ weightings that concentrate resolution capability towards the $x = x_1$ and/or $x = x_2$ boundaries, with $\psi_n(x)$ defined (via $g(x)$ and $h(x)$) to e.g. automatically satisfy the boundary conditions and (for incompressible flows) the divergence-free condition (cf. Spalart *et al.* 1991).

*The generalisation to three dimensions is straightforward, involving expansions of expansion coefficients.

† Since $\alpha_{M/2} = \alpha_{-M/2}$ (Press *et al.* 1986), equation (3) is shown with $M + 1$, rather than M , coefficients, in order to make explicit for later use that the magnitude of the maximum wavenumber $k_{\max} = |k_{M/2}| = \pi/(\Lambda/M)$.

Replacing $u(x)$ with $u_M(x)$ in $\mathcal{L}(u) = 0$, introduces the nonzero residual $\mathcal{R}_M = \mathcal{L}(u_M)$. To minimise the error, the product of the residual and each of the basis functions is set to zero, such that

$$\int_{x_1}^{x_2} w(x) \psi_m^* \mathcal{R}_M dx = 0 \quad \text{for } m = 0, 1, 2, \dots, M-1, \quad (4)$$

where ψ_m^* is the complex-conjugate of $\psi_m(x)$. The result is the semi-discrete $M \times M$ system of coupled first-order ordinary differential equations in the dependent variable $\vec{\alpha}(t) = [\alpha_0(t), \alpha_1(t), \dots, \alpha_{M-1}(t)]^T$, with

$$\mathcal{A} \frac{d\vec{\alpha}}{dt} = \mathcal{B} \vec{\alpha} + \vec{f}, \quad (5)$$

and where \mathcal{A} and \mathcal{B} are narrow-banded $M \times M$ matrices, [‡] whose elements (which are independent of time) are defined by (1) and $w(x)$. The quantity \vec{f} in (5) accounts for the nonlinear (e.g. convection) terms in \mathcal{L} . To improve efficiency, these are evaluated using a collocation/pseudo-spectral approximation (Gottlieb & Orszag 1977). The process of evaluating the nonlinear terms in this manner opens the calculation to errors introduced by the numerical quadrature. Since they consist of products of expansions in ϕ_n up to order M , the nonlinear terms will effectively involve an expansion up to order $2M$. Consequently, if the number of collocation points N within $x_1 \leq x \leq x_2$ is not larger than the number of expansion coefficients M , the quadrature will not be exact, leading to *aliasing* errors.[§] The algorithm is completed by applying a time-marching scheme to (5), to yield $\vec{\alpha}(t)$ and therefore $u_M(x, t)$ at discrete times (see section 3.2).

Spectral methods have a number of significant advantages. Chief among these is the fact that when the basis-function variations correctly ‘fit’ the dependent-variable variations (in terms of smoothness, boundary conditions, and regions of most rapid spatial change) the magnitude of the expansion coefficients α_n goes to zero as $M \rightarrow \infty$, and do so such that the error decreases faster than any power of $1/M$ (Gottlieb & Orszag 1977). Therefore, in contrast to a p^{th} -order finite-difference approximation, whose error is proportional to $(\Delta x)^p \propto (1/M)^p$, a spectral method converges with exponential or ‘infinite-order’ accuracy. The spatial accuracy of spectral DNS can thus be easily determined by confirming that the magnitude of the expansion coefficients $|\alpha_n| \rightarrow 0$ as $n \rightarrow M$, which provides a quality-control check that is only available to non-spectral methods via a grid-refinement study. The quality of the domain size can also be deduced from the low- n variation of α_n . Another attractive feature of Fourier- and Chebyshev-based methods is the ability to employ fast transforms when computing u_M from α_n , and when using the collocation/pseudo-spectral procedure to calculate \vec{f} in (5), both of which must be done at each time step.

[‡]The banded nature of \mathcal{A} and \mathcal{B} allows the time advancement to be effected quite efficiently. When the basis functions are built from orthogonal polynomials (e.g. Chebyshev or Jacobi), the matrix inversion involved in solving (5) is often able to use the Thomas algorithm for a tri- or penta-diagonal system. When the basis functions are complex exponentials, \mathcal{A} and \mathcal{B} are diagonal, the Fourier coefficients are completely decoupled, and the time-advance procedure is trivial.

[§]For incompressible flows, aliasing errors will be avoided if $N \geq 3M/2$ (Spalart *et al.* 1991). For compressible-flow DNS, for which the Navier-Stokes equations involve triple products of dependent variables, a straightforward remedy is not available, and the only recourse is to make N/M as large as possible. Some incompressible-flow algorithms (e.g. Kim *et al.* 1987) are also not able to remove aliasing errors in all directions, because e.g. they use linear combinations of the M expansion coefficients to enforce the no-slip boundary conditions.

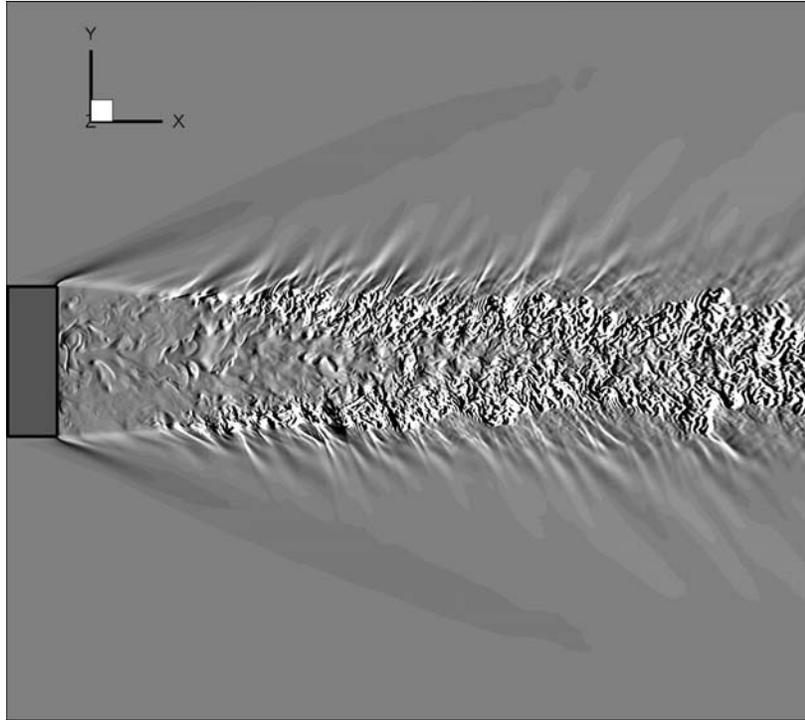


Figure 2. Instantaneous contours of streamwise density gradients from high-order finite-difference DNS of a supersonic axisymmetric wake at Mach 2.46 and $Re = 100,000$ (based on freestream velocity and diameter of body) (Sandberg 2008).

The disadvantages of spectral methods include the inability to consider non-simple flow geometries and the special treatments required to enforce inflow/outflow boundary conditions (see section 3.3). Due to their use of global basis functions, and the need to access the entire domain in each direction, spectral methods tend not to perform well on large distributed-memory parallel systems. They also cannot accurately represent flow discontinuities, and therefore are not well suited for DNS of high-speed compressible flows containing shock waves. These shortcomings led to the development of the high-order finite-difference schemes discussed in the next section.

3.1.2. Finite-Difference Methods Because of their ease of implementation, suitability to parallelization, and possible high-order accuracy, finite-difference (FD) schemes for DNS have become increasingly popular, especially for the emerging area of computational aeroacoustics (CAA) (see Figure 2). A wide range of options is available. Low-order FD methods allow complex geometries and irregular grids, and in this sense are similar to

finite-volume methods.[¶] The computational efficiency of low-order (especially upwind) FD approximations (FDA) is often unacceptable, requiring many more grid points than a spectral method to achieve the same accuracy. This can be demonstrated by applying a low-order FDA on a uniform grid with constant spacing $\Delta x = \Lambda/N$ for $j = 0, 1, \dots, N$ to a periodic function

$$u(x) = \sum_{k=-k_{\max}}^{+k_{\max}} \hat{u}(k) e^{ikx_j}, \quad (6)$$

where $k_{\max} = \pi/\Delta x$ (see (3)) is taken to be large enough to capture all spatial variations in $u(x)$. At $x_j = j\Delta x$, the function is given by $u_j = u(x_j) = \sum_k \hat{u}(k) e^{ikx_j}$ (the implicit $\pm k_{\max}$ limits on the k summation will hereinafter be suppressed). We consider the accuracy of the one-sided backward-difference approximation at x_j ,

$$\left. \frac{\Delta u}{\Delta x} \right|_{x_j} = \frac{u_j - u_{j-1}}{\Delta x}, \quad (7)$$

by applying it to the periodic function at x_j , resulting in

$$\left. \frac{\Delta u}{\Delta x} \right|_{x_j} = \sum_k \frac{1}{\Delta x} [1 - e^{-ik\Delta x}] \hat{u}(k) e^{ikx_j}. \quad (8)$$

By analogy to the exact result, $du/dx = \sum_k ik\hat{u}(k) e^{ikx}$, equation (8) can be written

$$\left. \frac{\Delta u}{\Delta x} \right|_{x_j} = \sum_k ik' \hat{u}(k) e^{ikx_j}, \quad (9)$$

where k' is the *modified wavenumber*, which for this one-sided FDA is

$$k'(k) = \sin(k\Delta x)/\Delta x + i[\cos(k\Delta x) - 1]/\Delta x. \quad (10)$$

The real part k'_r of k' is plotted as the solid line in Figure 3. The deviation from the exact (spectral) result $k'_r = k$ becomes larger as k approaches $k_{\max} = \pi/\Delta x$, which implies that derivatives of all but the very largest wavelengths ($k < 0.2\pi/\Delta x$) will be spuriously represented by this low-order FDA.

This becomes even clearer when (7) is used to discretise the linear convection/wave equation,

$$\frac{\partial u}{\partial t} + c \frac{\partial u}{\partial x} = 0, \quad (11)$$

where c is the constant wave speed. This equation is a model for the material-derivative terms in the Navier-Stokes equations, and is especially relevant to CAA, for which sound-wave propagation must be accurately represented. Applying the backward-difference approximation to the $c \partial u / \partial x$ term amounts to using a first-order upwind-difference scheme. The exact (in time and space) solution is (in terms of the Fourier modes)

$$\hat{u}(k, t) = \hat{u}(k, 0) e^{-ickt}, \quad (12)$$

[¶]Finite-volume methods have the advantage of in-built conservation properties, and the ability to produce solutions on irregular/unstructured grids without using a coordinate transformation.

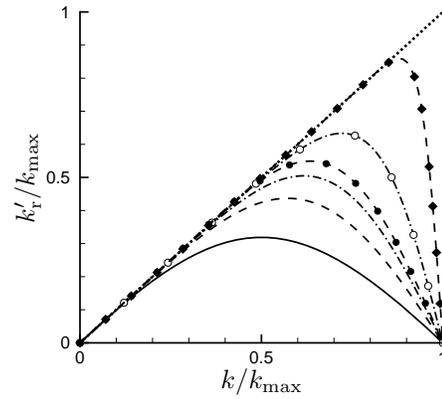


Figure 3. Modified wave numbers for first-derivative finite-difference approximations: —, explicit 1st-order one-sided/backward-difference & 2nd-order CDA $(P, Q) = (1, 0)$; ----, 4th-order explicit CDA $(2, 0)$; - · - · -, 6th-order explicit CDA $(3, 0)$; ○ · · · · ○, 6th-order compact CDA $(2, 1)$; ● · · · · ●, 4th-order explicit DRP CDA $(3, 0)$ (Tam & Webb 1993); ◆ · · · · ◆, 4th-order optimised compact CDA $(3, 2)$ (Kim 2007).

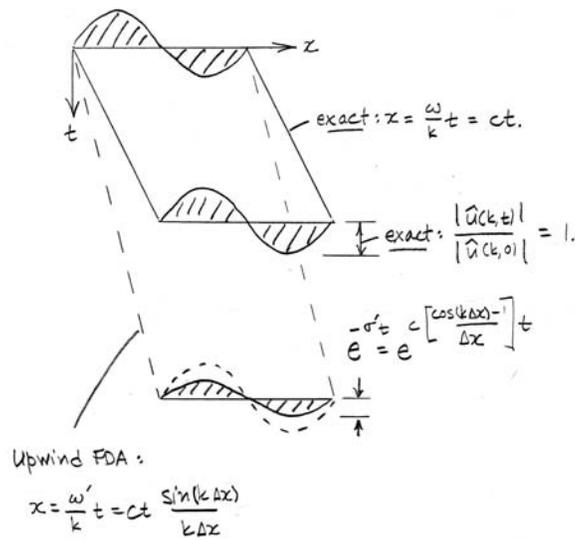


Figure 4. Amplitude and dispersion error for one-sided upwind FDA.

while the semi-discrete (in space) approximation is

$$\hat{u}(k, t) = \hat{u}(k, 0) e^{-ick't} = \hat{u}(k, 0) e^{-\sigma' t} e^{-i\omega' t}, \quad (13)$$

where $\sigma' = -ck'_i$ and $\omega' = ck'_r$ are the modified amplification and radial frequency, respectively, and $k' = k'_r + ik'_i$. (The temporal variation is kept exact here, in order to isolate the effect of the spatial differencing.) We can now quantify the amplitude/dissipation and phase/dispersion errors introduced by the upwind FDA, by comparing to the exact results, $\sigma = 0$ and $\omega = ck$. The modified decay rate is $\sigma' = c[(1 - \cos(k\Delta x))/\Delta x]$ and the modified phase speed is $c' = c[\sin(k\Delta x)/k\Delta x]$, which again indicates that only the $k \rightarrow 0$ modes, for which $\sigma' \rightarrow 0$ and $c' \rightarrow c$, are treated appropriately. The closer k becomes to $k_{\max} = \pi/\Delta x$, the faster the amplitude decays (as $\sigma' \rightarrow 2c/\Delta x$) and the slower the wave form convects (as $c' \rightarrow 0$); see Figure 4. This example illustrates the low computational efficiency, compared to spectral schemes, of low-order FD methods, and how they require many more grid points to achieve the same level of accuracy. Although currently available HPC resources allow meaningful DNS to be done (and in some cases, involving very complex geometries, can *only* be done) using low-order FDA, the loss in resolution per grid point represents a significant reduction in parameter space for quantities like Reynolds number and domain size. As a result, high-order finite-difference methods for DNS have become increasingly popular, especially for compressible turbulence and CAA.

Although their coding is more involved, especially for distributed-memory parallel systems, high-order FDA schemes can provide an excellent compromise between accuracy and flexibility for flows involving realistic geometries, such as aerofoils. The most general form of a central-finite-difference approximation (CDA) of a first derivative du/dx on an evenly spaced grid can be written

$$u'_j + \sum_{\ell=0}^Q b_\ell (u'_{j+\ell} + u'_{j-\ell}) = \frac{1}{\Delta x} \left(\sum_{\ell=0}^P a_\ell (u_{j+\ell} - u_{j-\ell}) \right) + \mathcal{O}(\Delta x^p) \quad \text{for } j = 0, 1, \dots, N, \quad (14)$$

where $u'_j = (\Delta u/\Delta x)_{x_j}$ is the FDA at $x_j = j\Delta x$, and $\mathcal{O}(\Delta x^p)$ is the truncation error of the scheme. The coefficients a_ℓ and b_ℓ are typically derived using Taylor-series expansion and are chosen to give the largest possible exponent $p = 2(P + Q)$, in order to minimize the truncation error and thus to maximise the formal order of accuracy. When $Q = 0$, the derivative is discretised with a standard, or explicit, finite-difference stencil, which only includes function values (not derivatives) at neighbouring nodes. In general, schemes with a small P require fewer operations per grid point than wider stencils, but at the cost of lower accuracy, with $p = 2P$. For simulations of turbulence and noise generation/propagation, higher-order methods are more computationally efficient. For these cases, finite differences with large stencils are preferred.

CDA with $Q \neq 0$ are termed *compact* or Padé schemes (Lele 1992). In contrast to explicit schemes, choosing e.g. $Q = 1$ with a five-point stencil ($P = 2$) increases the accuracy from 4th to 6th order. Because the derivatives at all grid points x_j are required simultaneously, a banded matrix system must be inverted, which increases the computational cost. As seen from our previous example, the formal order-of-accuracy is not the only criterion by which the performance of a finite-difference stencil should be judged. The dissipation and dispersion errors can be much more significant. In traditional CFD, dissipative errors often help maintain stability and accelerate convergence to a steady state. However, for DNS (in particular for

compressible turbulence and especially CAA applications), dissipation must be minimised to avoid, for example, attenuation of acoustic waves that propagate over long distances.^{||} Dispersion errors are also of crucial importance, because of their potential for introducing unphysical behaviour in the acoustic field (for CAA) or the evolution of the vortical structure of the turbulence.

We have already seen, for the first-order upwind scheme discussed earlier, how dissipation and dispersion errors are determined by the modified wavenumber profile $k'(k)$ of the FDA in question, and the transport equation that the FDA is used to discretise. The modified wavenumber for an arbitrary first-derivative CDA scheme can be determined by substituting $u'_j = \sum_k i k' \hat{u}(k) e^{i k x_j}$ into (14), and using the trigonometric relations $e^{i\theta} + e^{-i\theta} = 2 \cos \theta$ and $e^{i\theta} - e^{-i\theta} = 2i \sin \theta$, to find

$$k' = \frac{\left(\frac{1}{\Delta x}\right) \sum_{\ell=0}^P 2a_\ell \sin(\ell k \Delta x)}{1 + \sum_{\ell=0}^Q 2b_\ell \cos(\ell k \Delta x)}. \quad (15)$$

Note that since it has no imaginary component, ^{**} k' will introduce no dissipation/amplitude error when applied to the convective term of the linear wave equation (11); significant dispersion errors, however, are possible, since the modified frequency $\omega' = c k'(k)$ associated with (15) approaches zero as $k \Delta x \rightarrow \pi$.

The modified wavenumbers for several CDA schemes, both explicit and compact, are included in Figure 3. Although the 2nd-order CDA does not spuriously introduce an imaginary wavenumber, its real component $k'_r(k)$ is (and therefore its wave-equation dispersion characteristics are) identical to that of the first-order FDA considered above (solid line in Figure 3). Increasing the width of the stencil has a noticeable positive effect on the resolution of the explicit schemes, as evidenced by the upward shift of the maximum k' as the order increases, from 2nd to 4th to 6th. Note that a compact scheme with the same formal order of accuracy performs significantly better in terms of wavenumber resolution characteristics, albeit at a higher computational cost (compare the dash-dot curves with and without open symbols).

Given that the formal order of accuracy is not necessarily the most important feature of a FD scheme, it is logical to use the coefficients a_ℓ and b_ℓ in (14) to optimise the wavenumber resolution characteristics, rather than minimise the truncation error. Noteworthy examples of this exercise are shown in Figure 3, which involve sacrificing formal order of accuracy in order to push the maximum k' towards the exact $k_{\max} = \pi/\Delta x$ limit. This can be done for both explicit ($Q = 0$) and compact ($Q \geq 1$) schemes. An example of both is shown in Figure 3 (solid symbols). The optimised $k'(k)$ profile for the 4th-order compact scheme is particularly impressive. Because an increase in maximum k' allows a grid to be coarsened by the same factor to resolve the same flow features, these modified schemes represent a significant increase in computational efficiency, compared to a standard/explicit CDA using the same resolution, which can outweigh the additional cost of the more complex algorithm.

The discussion so far has been restricted to central interior schemes on uniform grids. FDA are also needed for the boundaries. Typically, boundary stencils are one-sided, so

^{||} Minimising dissipation is crucial for CAA, since the energy of the acoustic field is typically orders of magnitude smaller, and occurs at much different wavelengths, than that of the hydrodynamic field.

^{**} A property unique to central-difference schemes on uniform grids.

maintaining good wavenumber resolution characteristics is difficult (see Kim 2007 for a discussion of available options). Furthermore, finite-difference methods require general curvilinear coordinates when they are applied to complex geometries. This introduces metric terms to map from (a uniformly spaced grid in) computational space to physical space. Grid stretching is also routinely used, in particular for wall-bounded flows and/or for aeroacoustic problems where the acoustic wavelength is much larger than the hydrodynamic scales in the turbulent source region, which permits considerable coarsening of the grid towards the far field. The above Fourier-based modified wavenumber analysis cannot be directly applied to these cases, since the grid metrics introduce non-constant coefficients in the governing equations. For more details and analysis of the wavenumber resolution characteristics on non-uniform and curvilinear meshes, the reader is referred to Colonius & Lele (2004).

3.2. Temporal Discretisation

Two general methods of time discretisation can be distinguished: (i) explicit methods, where all spatial derivatives needed to advance the solution in time are evaluated at earlier time steps; (ii) implicit schemes, where the spatial derivatives are approximated using information at the new time step. In addition, a combination of the two approaches have been used for reasons that will become apparent soon, with, for example, an explicit method applied to the convective terms and an implicit method to the viscous terms. Whichever time-marching scheme is chosen, it needs to be both accurate (i.e. possess good frequency resolution characteristics, similar to the wavenumber equivalent for spatial schemes) and stable.

Accuracy of the temporal discretisation is paramount because the wide range of length scales found in turbulence is associated with a similarly wide range of timescales. In analogy to the length scales from section 2, the time scales of the largest and smallest eddies t_{LE} and t_η are respectively proportional to q^2/ϵ and $(\nu/\epsilon)^{1/2}$ (Kolmogorov 1941). The turbulence Reynolds number $Re_T = q^4/\nu\epsilon$ is therefore proportional to $(t_{LE}/t_\eta)^2$ which demonstrates that the Reynolds-number constraint previously discussed also applies to the time-integration of the Navier–Stokes equations. A system of differential equations yielding solutions with highly disparate time scales is termed *numerically stiff*.* When dealing with stiff systems, both accuracy and stability considerations require that the product of the timestep and the largest eigenvalue of the semi-discrete system produced by the spatial discretisation must be kept below a certain threshold. This threshold depends on the details of both the spatial and time-marching scheme used.

Accuracy and stability are tightly connected, making it impossible to treat either requirement individually. Generally, explicit time-integration schemes can be devised with p -th order of accuracy, and similar to optimization of spatial FD schemes for improved wavenumber accuracy, time-advancement schemes can also be optimized for frequency resolution rather than the highest possible order of accuracy. However, as will be shown with an example below, the allowable timesteps for explicit methods must be below a certain value to ensure numerical stability. It is often the case that the maximum permitted timestep is well below the accuracy threshold required to resolve all physical turbulence scales. For these flows,

*Formally, numerically stiff systems are those whose largest and smallest eigenvalues of the discretised system have very different magnitudes.

the DNS code can be run with a variable timestep set by the stability (e.g. constant-CFL) condition, in order to avoid wasting HPC resources using an unnecessarily small (constant) value.

Runge-Kutta methods (of various order) are usually the explicit scheme of choice for DNS, since they offer a good compromise between accuracy and stability. Implicit schemes on the other hand are usually unconditionally stable, but at the expense of having lower-order accuracy. In addition, implicit schemes generally require some type of iterative or approximate factorization methods and the higher computational cost due to the complexity of the algorithms typically outweigh any advantages gained due to larger time steps, in particular if high frequencies need to be resolved.

An example of the interplay between the spatial and temporal discretisations can be observed by again considering spatially periodic solutions to the wave equation, in this case a generalised version that includes diffusion-induced damping:

$$\frac{\partial u}{\partial t} + c \frac{\partial u}{\partial x} = \nu \frac{\partial^2 u}{\partial x^2}, \quad (16)$$

where c and ν are the constant convection speed and diffusion coefficient, respectively. Besides serving as a one-dimensional analogue of the Navier–Stokes equations, this linear model again allows its solutions to be written in terms of a single Fourier component at arbitrary wavenumber, $\hat{u}(k)$, which satisfies the Fourier-transformed equivalent of (16),

$$\frac{d\hat{u}}{dt} = \lambda \hat{u}, \quad (17)$$

where λ is defined by the spatial discretisation. For the exact (i.e. spectral) solution, $\lambda = -ick - \nu k^2$; for a FDA, $\lambda = -ick' - \nu(k'')^2$, and k' and k'' are respectively the modified wavenumbers defined by whichever finite-difference schemes are used for the first and second derivatives.

Let us examine the implications of using an explicit Euler time-marching method to integrate (17), for either the exact or semi-discrete/FDA case.[†] The solution at time $t_n = n\Delta t$, for $n = 0, 1, 2, \dots$, is thus given by

$$\hat{u}|_n = (1 + \Delta t \lambda) \hat{u}|_{n-1} = (1 + \Delta t \lambda)(1 + \Delta t \lambda) \hat{u}|_{n-2} = \dots = [1 + \Delta t \lambda]^n \hat{u}_0, \quad (18)$$

where Δt is the timestep increment and $\hat{u}|_n = \hat{u}(k, t_n)$. Note that the solution at t_n is given by the product of the initial value \hat{u}_0 and a factor raised to the power n . Therefore, for this scheme to remain stable, the magnitude of the term in square brackets must not exceed unity, such that

$$(1 + \lambda_r \Delta t)^2 + (\lambda_i \Delta t)^2 \leq 1, \quad (19)$$

where $\lambda = \lambda_r + i\lambda_i$, and the real and imaginary components are (in general) respectively $\lambda_r = ck'_i - \nu(k''_i)^2$ and $\lambda_i = -ck'_i - \nu(k''_i)^2$. The stability requirement can thus be written

$$[1 + (ck'_i - \nu(k''_i)^2) \Delta t]^2 + [(ck'_i + \nu(k''_i)^2) \Delta t]^2 \leq 1. \quad (20)$$

[†]This choice is motivated solely by its instructional value. Because of its low-order accuracy and stability limitations, the explicit-Euler scheme would not be appropriate for production DNS.

We consider two cases. When the exact/Fourier spectral discretisation is used, the modified wavenumbers are replaced by the actual wavenumber k , with $k'_r = k$, $(k''_r)^2 = k^2$ and $k'_i = (k''_i)^2 = 0$. The critical condition is defined by the maximum resolvable wavenumber k_{\max} , which for this spectral method is $\pi/\Delta x$. Therefore,

$$\left[1 - \pi^2 \frac{\nu \Delta t}{\Delta x^2}\right]^2 + \left[\pi \frac{c \Delta t}{\Delta x}\right]^2 \leq 1. \quad (21)$$

It is instructive to compare this result to a finite-difference counterpart. If, for example, central differences were used to approximate both the first and second derivatives in (16), we know from (15), and via the same procedure used earlier to analyse the one-sided backward difference scheme, that the modified wavenumbers are

$$k'_r = \sin(k\Delta x)/\Delta x, \quad (k''_r)^2 = 2(1 - \cos(k\Delta x))/\Delta x^2, \quad (22)$$

with $k'_i = (k''_i)^2 = 0$.[‡] Because the maximum wave numbers for the CDA scheme are $k'_{r,\max} = 1/\Delta x$ and $(k''_i)_{\max}^2 = 4/\Delta x^2$ (see (15) and Figure 3), the stability condition is now

$$\left[1 - 4 \frac{\nu \Delta t}{\Delta x^2}\right]^2 + \left[\frac{c \Delta t}{\Delta x}\right]^2 \leq 1. \quad (23)$$

Equations (21) and (23) both underline the role of viscosity in aiding the stability of a CFD algorithm. (Note that if $\nu = 0$, both the spectral and CDA implementations of the explicit-Euler schemes would be unconditionally unstable.) They also contain a reminder that the diffusion stability condition, $\Delta t \propto \Delta x^2/\nu$ can be more important than the CFL condition, $\Delta t \propto \Delta x/c$, for viscous flows in regions where the grid is very fine (as it must be, for example, to resolve the small-scale turbulence adjacent to no-slip walls; cf. Figure 1). Note that if the timestep is constrained by the diffusion stability condition, it scales with the square of the grid spacing (versus with Δx for the CFL limit) potentially making it prohibitively small. This has important implications for the time-marching treatments used in DNS of wall-bounded flows, including boundary layers and channels. Consequently, many incompressible DNS of wall-bounded flows have used implicit treatments for the linear viscous terms (typically Crank-Nicholson). For compressible flows, however, the non-constant coefficients due to variations in density and temperature make an implicit treatment of the viscous terms much less straightforward, and the advantage of combining explicit and implicit time-integration schemes is not as obvious.

From the point of view of designing and performing a DNS, perhaps the most significant message of (21) and (23) is regarding the relationship between the stability-based time-step limitation and the accuracy of the spatial discretisation. The maximum Δt given by both the CFL and viscous/diffusion criteria is significantly larger for the CDA than spectral method (by factors of π and $\pi^2/4$, respectively). However, it is important to realise that this extra stability is a result of diminished spatial resolution, related to the maximum modified wavenumbers of the CDA schemes – i.e. $1/\Delta x$ and $4/\Delta x^2$ versus $\pi/\Delta x$ and $\pi^2/\Delta x^2$, for the Fourier method;

[‡]We have already commented in section 3.1.2 upon the equivalence of the real parts of the modified wavenumber profiles $k'(k)$ for the one-sided and central-difference approximations, but the lack of spurious k'_i provided by the CDA, and the implications for amplitude/dissipation errors associated with the convection term.

see (22). In other words, the FD algorithm is able to take larger time steps because it is not faithfully representing the $k \rightarrow k_{\max}$ wavenumbers. One should be aware of the potential for a spatial scheme to unphysically enhance a DNS code's stability characteristics.

The stability criteria obtained from one-dimensional linear model problems usually provide a reliable estimate of what is required in practice for nonlinear and multi-dimensional simulations. In many cases however, in particular when applying optimized methods and high-order-accurate schemes, spurious oscillations in the high wavenumber range can occur. In addition to insufficient grid resolution, several factors can cause this undesired behaviour. First, a mismatch between interior and boundary FD schemes can make the overall algorithm unstable. This is in particular true for compact FD schemes combined with non-compact boundary schemes. To ensure stability of the entire scheme, a stability analysis needs to be performed of the overall method and not just the scheme used for the interior derivatives. Second, nonlinearities can cause a method that was stable for a linear equation to be unstable. Several approaches have been proposed to treat the nonlinear terms to ensure long-time stability, such as entropy splitting (Sandham *et al.* 2002), higher-entropy conservation (Honein & Moin 2004), skew-symmetric splitting (Kennedy & Gruber 2008), or explicit filtering (Visbal & Gaitonde 2001).

3.3. Boundary & Initial Conditions

The unsteady three-dimensional nature of turbulence leads to some special problems when it comes to defining and enforcing appropriate boundary conditions for DNS. Given that the most appropriate boundary conditions are Navier-Stokes solutions, the problem is especially challenging for spatially developing flows, such as boundary layers, that are fully turbulent at the inflow boundary of the domain, since these require complete, physically realistic, turbulent histories for each of the dependent variables to be prescribed at each point on the inflow plane. (Flows for which the turbulence does not convect across the domain boundaries can have other boundary-condition complications; see below.) One of the most attractive features of Fourier spectral methods is their ability to sidestep this problem: applying them in directions in which the turbulence is statistically homogeneous, for which it is natural to assume spatial periodicity (provided the domain is large enough), automatically produces conditions whose history and spatial structure fully satisfy the governing equations. This is true both when the homogeneous direction does and does not coincide with the direction of the net mean flow (for example, in both the longitudinal and lateral directions of plane-channel/Poiseuille or Couette flow). As a result, the early (Fourier/spectral) DNS work focused on parallel flows, the first of which were statistically stationary (e.g. fully developed plane-channel flow) and realisable in the laboratory. Later, the parallel-flow geometry was used for DNS of time-developing analogues of spatially developing cases (boundary layers, plane mixing layers, wakes, etc). For these time-developing flows, the task of defining a sequence of physical turbulent inflow conditions is replaced by specifying a single turbulent-field initial condition (revealing another advantage of using the time-developing parallel-flow analogue).

There are a number of strategies currently available for generating the required turbulent inflow conditions for non-parallel flows. Fourier schemes can be altered to accommodate spatial flows by adding non-physical terms to the governing equations that are active only in *fringe* or *sponge* regions at the downstream end of the periodic domain. These terms artificially 'process' the turbulence so that it 'forgets' the outflow state, and re-enters the

upstream domain suited to the local flow conditions (see Spalart & Watmuff 1993 for details).[§] However, for the many flows for which Fourier methods are not viable (due for example to complex geometries or rapid streamwise variations), other discretisation approaches, such as finite-difference or finite-volume methods, must be used, and the inflow boundary conditions must be explicitly specified. One option is to let the flow transition to turbulence within the flow domain, by specifying a relevant base flow at the inlet (e.g. a laminar profile) superposed with suitable perturbations (either low-amplitude disturbances that feed linear instabilities, or larger-amplitude fluctuations tuned to efficiently trigger nonlinear bypass transition.[¶]) All of these approaches involve sacrificing a fraction of the domain to allow the turbulence to reach the desired fully physical state. In contrast, one can use a companion/precursor DNS to obtain the turbulent inflow data, or employ a self-contained ‘recycling’ technique, whereby turbulent results from within the domain are extracted, rescaled and specified at the inflow plane at each time step (Lund *et al.* 1998). Each of these approaches represents different compromises between minimising, on the one hand, the non-physical region of the domain used to achieve a fully developed state, and, on the other, the computational cost.

Outflow and far-field boundary conditions can also require some care, especially for external flows. (There is usually no ambiguity about the boundary conditions for internal, wall-bounded/duct flows, which typically involve a combination of no-slip or transpiration conditions at walls, and zero-derivative or periodic conditions at outflows.^{||}) For most external flows, the actual physical domain will be much larger than the largest affordable computational domain. The challenge is to devise boundary conditions that allow for flow disturbances (vorticity, acoustic or entropy waves) to cross the finite/truncated-domain boundaries without generating excessive spurious waves that are reflected back to the region of interest. The quality of DNS for CAA in particular is tied to an ability to meet this requirement. Various so-called *non-reflecting boundary conditions* have been derived using characteristics-based boundary conditions (Thompson 1987). In applications where these boundary conditions are unable to sufficiently attenuate spurious reflections, additional fringe/sponge layers, involving filtering, grid-stretching, additional convection terms, or combinations thereof, can be employed.

The quality of the *initial conditions* needed by a DNS varies widely. For stationary flows, the only benefit of specifying a fully physical turbulent initial condition is to minimise the time it takes to overcome an initial transient; because the long-time flow statistics do not depend on the initial condition, totally unphysical conditions (e.g. random velocity fluctuations) can be specified. **

[§]In addition to representing spatially developing flows, Fourier schemes can also be adapted to exactly represent a compact region of turbulence in an unbounded/infinite domain (see Corral & Jiménez 1995).

[¶]These larger-amplitude disturbances could be, for example, based on a synthetic/analytic time-dependent model of near-wall turbulence structures (Sandham *et al.* 2003), or the result of applying *digital filters* to unsteady random perturbations in order to impose the desired length- and time-scales at the inflow (Xie & Castro 2008).

^{||}Very complex surface geometries, such as roughness elements, can be approximated using *immersed boundary methods* (Goldstein *et al.* 1993).

**A key requirement is to determine when the flow reaches the beginning of its statistically steady state, after which time-averaging of data should begin. This can be achieved by monitoring higher-order statistics and histories of global quantities, such as plane-averaged shear stress, bulk mass flux, or integrated turbulence kinetic energy. To avoid unnecessary computational cost, it is usually wise to accelerate the arrival of the stationary state by beginning the simulation on a spatial grid that is significantly coarser than that required to resolve the flow in question. This could be done in a number of steps, increasing the resolution after the under-resolved flow has equilibrated, as

For flows where it is important to impose physical fully turbulent initial conditions (such as the parallel-flow time-developing analogues of spatial flows mentioned earlier; cf. Coleman *et al.* 2003), they can be obtained using counterparts of the turbulence inflow generation strategies (except the recycling techniques) mentioned above.

3.4. Code Validation & Resolution Guidelines

The fundamental requirement of any DNS is to produce results that can be trusted. Consequently, especially in light of the computational costs involved, it is essential that every DNS code must be properly validated before it enters ‘production’ and then used appropriately thereafter. In this section (adapted from the recommendations given by Sandham (2005)) we offer suggestions about how this might be done.

- A newly developed or modified code must be able to reproduce analytic solutions and asymptotic limits. For example, laminar channel, pipe or boundary-layer flows can be computed and compared to the exact results. A more challenging (and therefore more valuable) test is prediction of growth rates and phase speeds of small disturbances in an appropriate background flow. For fully turbulent cases, rapid-distortion theory (for which only the linear terms in the governing equations are active) often provides a useful benchmark, in terms of closed-form predictions of histories of turbulence statistics. Another powerful method, whose utility does not seem to have been fully appreciated, is to prescribe analytic unsteady, three-dimensional test functions* for all variables, and to modify the governing equations so that they are exactly satisfied by these test functions. This involves adding forcing terms† that cause the left- and right-hand sides to formally agree. The difference between the code results and the test functions, at each point in the domain at each time, can then be monitored, and quantified in terms of the sizes of the grid spacing and time step (S. Chernyshenko, personal communication).
- When assessing the sufficiency of the spatial and temporal resolution for DNS, one cannot follow the standard CFD practice of demonstrating grid and time-step independence by converging towards the same point-wise solution history using different grids and different time-step sizes. This is because of the well-known characteristic of turbulence (often referred to as extreme sensitivity to initial conditions) that small differences between two realisations of turbulence quickly lead to divergence of the individual histories (measured say by velocity fluctuations at the same point) of the two flows. As a result, we must examine the effect of the numerical parameters on time-, plane-, volume- or ensemble-averaged statistics.
- A valuable *a posteriori* check of spatial accuracy of a production DNS is provided by each dependent variable’s (space/time/etc-averaged) wavenumber spectra. These should exhibit decay over several orders of magnitude, with negligible energy in the smallest-scale modes. Spatial spectra are readily available in all directions for fully spectral codes (via the variation of the expansion coefficients versus mode/expansion-

defined by the above criteria.

*The test functions must satisfy the appropriate boundary conditions, and, for incompressible-flow codes, should be divergence free.

†Symbolic equation software is recommended for building the forcing terms from the test functions.

index number, mentioned in section 3.1.1). They can also be obtained in homogeneous directions for finite-difference/finite-volume discretisations that use a uniform grid,[‡] by Fourier transforming results from an instantaneous realisation, and averaging them appropriately. (Recall however that there may be a significant difference between the maximum wavenumber associated with the grid ($\pi/\Delta x$) and the maximum wavenumber captured by the finite-difference approximation; see section 3.1.2.)

- Frequency spectra can also be used to assess the validity of the choice of timestep (by examining the relative magnitude of the energy captured by the time-marching scheme's highest resolvable frequency[§]). Generally, when explicit time-marching schemes are used (as with most compressible-flow codes), the timesteps defined by the stability criteria (section 3.2), are sufficiently small to resolve the smallest time scales of the turbulence. In the case of implicit time-marching schemes, for which numerical stability can be maintained with much larger timesteps, the validity of the time-step must be established on other grounds (see below).
- For both spectral and finite-difference schemes, it is illuminating to compare both the local grid spacing Δx_i with the local Kolmogorov length scale $\eta = (\nu^3/\epsilon)^{1/4}$, and the time step Δt with the Kolmogorov time scale $t_\eta = (\nu/\epsilon)^{1/2}$, where ϵ is computed from the DNS data (i.e. from the products of the spatial derivatives of the velocity fluctuations found in the instantaneous DNS fields). The ratios $\Delta x_i/\eta$ and $\Delta t/t_\eta$ should not be larger than order one.
- For wall-bounded flows, near-wall similarity can be invoked, allowing the resolution to be based on wall units Δx_i^+ , which are obtained by normalizing the grid spacing Δx_i with the kinematic viscosity ν and the friction velocity $u_\tau = (\tau_w/\rho)^{1/2}$, where τ_w is the local wall shear stress, with $\Delta x_i^+ = \Delta x_i u_\tau/\nu$. It has been found from spectral DNS of fully developed wall-bounded turbulence (Kim *et al.* 1987; Spalart 1988) that the full-resolution threshold[¶] requires that the streamwise and spanwise grids (i.e. distance between the evenly spaced Fourier collocation/quadrature points) respectively satisfy the conditions $\Delta x^+ < 15$ and $\Delta z^+ < 8$. The wall-normal grid spacing should expand with distance from the wall, with the requirement that the first point be at $y^+ < 1$ and the first ten points within $y^+ < 10$. Note however, that these estimates were obtained from fully spectral methods and thus will in general not be sufficient for finite-difference schemes, with their inferior wavenumber-resolution characteristics (see section 3.1.2).
- Systematic studies of the size of the computational domain should be performed to verify that all relevant flow features are captured. Two-point correlations tending to zero within the domain give an indication that the domain is large enough. (The low-wavenumber behaviour of the energy spectra multiplied by the wavenumber – specifically, whether or not this quantity approaches zero as the wavenumber does – is another, sometimes more revealing, way to determine the relationship between the largest turbulence scales and the domain size.) If the minimum two-point correlations (at

[‡]Or by interpolating non-uniform grid results onto a uniform grid.

[§]If the constant-CFL condition is used to set the variable timestep, this would require the constant- Δt data to be interpolated from the variable- Δt results.

[¶]Based on reasonable convergence of low-order statistics; see Figure 2 of Spalart (1998) and compare Spalart *et al.* 2008/9.

maximum separation) are not zero within the domain, all scales larger than the domain size are effectively treated as if they were infinitely long.

- Budgets of statistical quantities can be computed and should balance. This is an especially important check, because it reflects the overall quality of the results, as affected by the spatial and temporal resolution, and also the state of convergence of the turbulence statistics (i.e. whether or not enough data has entered the averaging sample). Regarding the resolution implications, a demonstration that the temporal and/or convective changes of, say, turbulence kinetic energy are in good agreement with the sum of the production, dissipation and transport terms that have been computed from the DNS fields, represents fairly compelling evidence that the temporal and especially spatial resolution are sufficient. (Given the uncertainty that can accompany low-order finite-difference/finite-volume codes, this type of diagnostic can go a long way to instil confidence in the results they produce.) The budgets are also a good way to quantify time-discretisation errors, which can appear, for example, in the turbulence kinetic energy budget as either extra dissipation or production, depending upon the scheme's stability characteristics: methods such as the Runge-Kutta algorithm, which linear theory predicts will damp the smallest scales, can introduce numerical dissipation whose magnitude is proportional to the size of the time step (Coleman *et al.* 1992); methods such as the Adams-Bashforth scheme, which tend to energise the smallest scales, can be expected to contribute a non-physical/numerical production to the energy balance. The quality of the statistical sample is also demonstrated by the budgets. The rate at which various statistics converge varies, with higher-order quantities (such as third-order correlations) and those dominated by very-large-scale structures (such as two-point correlations at large separation or spectra at low wavenumber – which involve fewer 'eddy samples' within a given finite domain) tending to converge most slowly. From a practical point of view, an inadequate statistical sample is just as serious a problem as inadequate spatial or temporal resolution, and can seriously limit the utility of the DNS results. The only remedy for poorly converged statistics is to expend the computational resources needed to gather more samples, by either increasing the averaging period for a time average, or the number of experiments entering the ensemble average.
- If unexpected results are encountered, they should be reproduced with a different numerical scheme and/or code.

4. OUTLOOK: APPLICATIONS & CHALLENGES

We expect that as HPC capability continues to increase so will the reach of DNS, both in terms of parameter space (especially Reynolds number) and flow complexity. However, if current trends continue, researchers must continue for the foreseeable future to appeal to Reynolds-number similarity when considering the relevance of DNS results to full-scale applications. On the other hand, the ability to perform fundamental studies of 'clean' flows unaffected by numerical, modelling and measurement errors, will continue to make it an attractive complement to other research strategies. The complete control of the initial and boundary conditions, and each term in the governing equations, also leads to profound advantages over laboratory and field studies. While new single, stand-alone, DNS of canonical building-block flows are apt to occur, it is likely that most future work will involve case studies of a

series of physical and non-physical simulations. The latter – which might include artificial initial and/or boundary conditions, unrealistic parameter combinations (e.g. zero gravity, extreme stratification), or extra or missing terms from the governing equations – should be especially powerful, since they can be used to answer basic questions of physics and modelling in a straightforward manner.

Recent advances in HPC hardware have led to additional challenges for DNS algorithms, such as the need to parallelize numerical schemes efficiently to fully exploit systems with a large number of processors. Parallelization is particularly difficult for spectral- and compact-finite-difference methods, for which matrix systems need to be inverted. As a result, despite their inferior resolution characteristics, explicit finite-difference schemes are often employed, since they can be applied much more efficiently using very large numbers ($> 1,000$) of processors. However, because attainable clock-speeds of CPUs have most likely peaked in the last few years, due to thermal and power consumption constraints, further increases in computing power will only be achieved by further massive increases in the number of computing cores, leading to HPC systems with $> 100,000$ cores. Because of the very high clock-speeds of current CPUs, and the size of a typical large-scale DNS, memory access now often defines the total performance, rather than the total number of operations to be performed. This issue will presumably become even more critical in the future. Furthermore, entirely new computing architectures (graphics processing units, cell processors) are being considered for scientific HPC. These developments will undoubtedly pose new challenges for the maintenance of current codes and the development of new efficient numerical methods for DNS.

Another challenge associated with improving HPC performance involves the sheer amount of data that can now be produced.^{||} The issues raised by the need to transfer and store this much data are not trivial. We expect that dedicated post-processing and flow-visualization systems will be required to take full advantage of the results of future DNS studies.

REFERENCES

- Canuto C, Hussaini MY, Quarteroni A & Zang TA. *Spectral Methods in Fluid Dynamics*. Springer-Verlag: 1988.
- Coleman GN, Ferziger JH & Spalart PR. Direct simulation of the stably stratified turbulent Ekman layer. *J Fluid Mech.* 1992; **244**:677–712 (Corrigendum: *J. Fluid Mech.* **252**:721.)
- Coleman GN, Kim J & Spalart PR. Direct numerical simulation of a decelerated wall-bounded turbulent shear flow. *J Fluid Mech.* 2003; **495**:1–18.
- Colonus T & Lele SK. Computational aeroacoustics: progress on nonlinear problems of sound generation. *Prog Aerospace Sci.* 2004; **40**:345–416.
- Corral R & Jiménez J. Fourier/Chebyshev methods for the incompressible Navier-Stokes equations in infinite domains. *J Comp. Phys.* 1995; **121**:261–270.

^{||}The turbulent channel DNS of Hoyas & Jiménez (2006), which required $\mathcal{O}(10^{10})$ grid points, created 25 Tb of raw data.

- Goldstein D, Handler R & Sirovich L. Modeling a no-slip flow boundary with an external force field. *J Comp. Phys.* 1993; **105**:354–366.
- Gottlieb D & Orszag SA. *Numerical Analysis of Spectral Methods: Theory and Applications*. Society for Industrial and Applied Mathematics, 1977.
- Honein AE & Moin P. Higher entropy conservation and numerical stability of compressible turbulence simulations. *J Comp. Phys.* 2004; **201**:531–545.
- Hussaini MY & Zang TA. Spectral methods in fluid dynamics. *Annu. Rev. Fluid Mech.* 1987; **19**:339–367.
- Kennedy CA & Gruber A. Reduced aliasing formulations of the convective terms within the Navier-Stokes equations for a compressible fluid. *J Comp. Phys.* 2008; **227**:1676–1700.
- Kim J, Moin P & Moser RD. Turbulence statistics in fully developed channel flow at low Reynolds number. *J Fluid Mech.* 1987; **177**:133–166.
- Kim JW. Optimised boundary compact finite difference schemes for computational aeroacoustics. *J Comp. Phys.* 2007; **225**:995–1019.
- Kolmogorov AN. The local structure of turbulence in incompressible viscous fluid for very large Reynolds numbers. *Dolk. Akad. Nauk SSSR* 1941; **30**:299–303 (reprinted in *Proc. R. Soc. Lond. A*; 1991; **434**:9–13).
- Lele SK. Compact finite difference schemes with spectral-like resolution. *J Comp. Phys.* 1992; **103**:16–42.
- Lund TS, Wu X & Squires KD. Generation of turbulent inflow data for spatially-developing boundary layer simulations. *J Comp. Phys.* 1998; **140**:233–258.
- Orszag SA & Patterson GS. Numerical simulation of three-dimensional homogeneous isotropic turbulence. *Phys. Rev. Lett.* 1972; **28**:76–79.
- Press WH, Flannery BP, Teukolsky SA & Vetterling WT. *Numerical Recipes: The Art of Scientific Computing*. Cambridge: 1986.
- Richardson LF. *Weather Prediction by Numerical Process*. Cambridge: 1922.
- Sandberg RD. Development of a new compressible Navier-Stokes solver for numerical simulations of flows in turbomachinery. Progress report for *HPC Europa++* Transnational Access Project 1264; 2008.
- Sandham ND. Turbulence simulation. In *Prediction of Turbulent Flows*, Hewitt GF, Vassilicos JC (eds). Cambridge: 2005; 207–235.
- Sandham ND, Li Q & Yee HC. Entropy splitting for high-order numerical simulation of compressible turbulence. *J Comp. Phys.* 2002; **178**:307–322.
- Sandham PR, Yao YF & Lawal AA. Large-eddy simulation of transonic flow over a bump. *Int. J Heat Fluid Flow* 2003; **24**:584–595.
- Spalart PR. Direct simulation of a turbulent boundary layer up to $R_\theta = 1410$. *J Fluid Mech.* 1988; **187**:61–98.
- Spalart PR, Coleman GN & Johnstone R. Direct numerical simulation of the Ekman layer: a step in Reynolds number, and cautious support for a log law with a shifted origin. *Phys. Fluids* 2008; **20**:101507. (See also Retraction in *Phys. Fluids* 2009; **21**:109901.)

- Spalart PR, Moser RD & Rogers MM. Spectral methods for the Navier-Stokes equations with one infinite and two periodic directions. *J Comp. Phys.* 1991; **96**:297–324.
- Spalart PR & Watmuff JH. Experimental and numerical study of a turbulent boundary layer with pressure gradients. *J Fluid Mech.* 1993; **249**:337–371.
- Tam CKW & Webb JC. Dispersion-relation-preserving finite difference schemes for computational acoustics. *J Comp. Phys.* 1993; **107**:262–281.
- Thompson KW. Time dependent boundary conditions for hyperbolic systems. *J Comp. Phys.* 1987; **68**:1–24.
- Visbal MR & Gaitonde DV. Very high-order spatially implicit schemes for computational acoustics on curvilinear meshes. *J Comp. Acoustics* 2001; **9**:1259–1286.
- Xie ZT & Castro IP. Efficient generation of inflow conditions for large-eddy simulation of street-scale flows. *Flow Turbulence Combust.* 2008; **81**:449–470.

FURTHER READING

- Fletcher CAJ. *Computational Galerkin Methods*. Springer-Verlag: 1984.
- Gatski TB, Hussaini MY & Lumley JL (eds). *Simulation and Modeling of Turbulent Flows*. Oxford: 1996.
- Lomax H, Pulliam TH & Zingg DW. *Fundamentals of Computational Fluid Dynamics*. Springer: 2001.
- Moin P. *Fundamentals of Engineering Numerical Analysis*. Cambridge: 2001.
- Moin P & Kim J. Tackling turbulence with supercomputers. *Sci. American* 1997; **276**(1):62–68.
- Moin P & Mahesh K. Direct numerical simulation: a tool in turbulence research. *Annu. Rev. Fluid Mech.* 1998; **30**:539–578.