JoT JOURNAL OF TURBULENCE

http://jot.iop.org/

Computing high-Reynolds-number turbulence: will simulations ever replace experiments?*

Javier Jiménez

School of Aeronautics, Universidad Politécnica, 28040 Madrid, Spain
and
Centre for Turbulence Research, Stanford University, Stanford, CA 94305, USA
E-mail: jimenez@torroja.dmt.upm.es

Received 23 April 2003
Published 12 June 2003

Abstract. The cost and future prospects of direct and large-eddy numerical simulations are discussed in the context of their use in the development of turbulence theory. It is concluded that even today simulations are comparable to experiments in many aspects, including the Reynolds number, and that simulations at asymptotically large Reynolds numbers should become available within the next decade. For isotropic or free-shear turbulent flows, this should be equivalent to an infinitely large Reynolds number, but to achieve the same result in wall-bounded flows requires the development of a theory for the structure of the logarithmic layer. This is the key obstacle for practical large-eddy simulations of wall flows, and it is argued that its development would be better supported by numerical simulations than by experiments.

PACS numbers: 47.11.+j, 47.27.-i, 47.27.Jv, 47.27.Nz

* This article is a modified version of the original article from the Proceedings of the 5th International Symposium on Engineering Turbulence Modelling and Measurements, Mallorca, 16–18 September 2002, which appeared in 'Engineering Turbulence Modelling and Measurements 5', Elsevier Science, ISBN 0-08-044114-9, ed W Rodi and N Fueyo.

 PII: S1468-5248(03)62586-5 1468-5248/03/000001+14$30.00

Contents

JoT 4 (2003) 022

1. Introduction

The title of this paper asks whether numerical simulations will ever replace experiments in the study of turbulent flows, which is of course a question which requires qualification. Few tools in the history of science have ever been fully replaced. There are always problems which can be best treated using older techniques, and there is usually no need to renounce what is known. There have been, however, many cases in which new tools have become cheaper or more effective for a particular task, and the question that we will investigate here is whether, in any foreseeable future, direct or large-eddy simulations (LES) of turbulence will become more convenient than experiments for some particular purposes.

There should be little doubt by now that careful numerical simulations of the Navier–Stokes equations are just a different kind of experiment, and that there is no reason to expect them to be less accurate than laboratory ones [1]. Some of their drawbacks, such as the artificial nature of some boundary conditions, are no worse than the artificial nature of the walls in a wind tunnel and the artificiality can, in both cases, be circumvented by careful design. It is not always clear that what we simulate numerically is exactly the same thing that we would measure in the laboratory, but there is no reason why that should be the case. Simulations and experiments are equally valid idealizations against which to test our theories or designs. In many senses simulations are actually better characterized than experiments, since they do not usually suffer from ambiguities about boundary, entry, or initial conditions.

There is also little doubt that, once a flow has been successfully simulated, it can be observed more thoroughly than in a laboratory experiment. The cost of setting up a simulation is high, although not necessarily higher than that of a large wind tunnel, but the instrumentation problems are simpler, and almost any observation that can be imagined can be made. While the results of laboratory experiments are often constrained by the instrumentation technology, those from numerical simulations are mainly limited by the ability of the researcher to ask the right questions.

It has moreover been argued that the main advantage of simulations over laboratory experiments is the ease with which they can be adapted to perform 'thought' experiments, in which the equations of motion or the boundary conditions are changed to, in effect, 'take the

system apart' [2]. This has always been a useful device in physics, and 'thought' experiments have often been used to constrain physical theories by asking what would happen if a given 'implausible' modification were implemented. Think for example of potential inviscid flows. For such arguments to be useful, however, we must be able to answer the questions that we have posed, and it is in that respect that many of the idealized systems of classical physics differ from complex ones, such as turbulence. While it might be possible with pencil, paper, and a lot of imagination, to decide what happens when two observers try to synchronize their clocks under certain conditions, the answer is often harder when trying to decide the outcome of a particular modification to a turbulent flow. Numerical simulations allow us to obtain answers to our thought experiments.

While those are clear advantages of simulations over laboratory experiments, the argument over their relative merits has often centred on the different question of which systems can be studied with each technique. It is often claimed that experiments can be run at higher Reynolds numbers than simulations. What we would like to discuss in this paper is when, if ever, that limitation is likely to be removed. Since the argument is essentially over cost, that issue will be discussed first, centring initially on isotropic turbulence. This will be complemented by a discussion of the particular challenges of wall-bounded flows and by a description of some recent simulations in that field, as an illustration of the present status of the technology. Some reflections are then offered on the effect of simulations on the future development of turbulent research.

2. The problem of turbulence

No discussion of a methodology can be initiated without defining what problem is being considered, and turbulence, being both a poorly understood physical phenomenon and one with profound practical implications, is a problem that needs to be defined more carefully than many others. The point of view in this paper is that physics drives engineering, although other approaches are possible and in some cases desirable, and that the problem that needs to be solved is the fundamental physical understanding of turbulence.

The book by Townsend [4] is a good reference for the classical theory used in the rest of this paper, and contains a figure which summarizes the essence of the turbulence problem. The version in figure 1 shows the spectral distributions of the turbulent energy and of the dissipation. Each spectrum has its peak at the wavenumbers where the respective property resides. The energy is at large scales, and the dissipation is at small ones. What defines turbulence is that the energy and the gradients are at different scales, and that there is an intermediate gap in which both quantities are small, and through which one is transported into the other.

Classical theory tells us that, to a good approximation, the two peaks are independent of each other, and that they only interact through the intermediate inertial cascade, which is universal. If we could compute a flow in which those two peaks were really separated, we would have computed asymptotic turbulence and, in a sense, there would be nothing else to do. Once in possession of the results of such a simulation, we would be able to test our theories in detail against them.

The same separation holds for the Reynolds stress, whose cospectrum is included in figure 1. It is also concentrated at large scales, and decays faster towards the smaller scales than the energy spectrum. In most cases, if we correctly compute the spectral energy peak, we have also correctly computed the distribution of the shear stresses.

Because of our emphasis on physical insight, our only concern in this paper will be with computations in which at least some of the turbulent scales are resolved. Direct numerical simulations (DNS) make no assumptions about the physics, and resolve all the scales of the flow, including the energy and dissipation spectral peaks. LES rely on the classical result

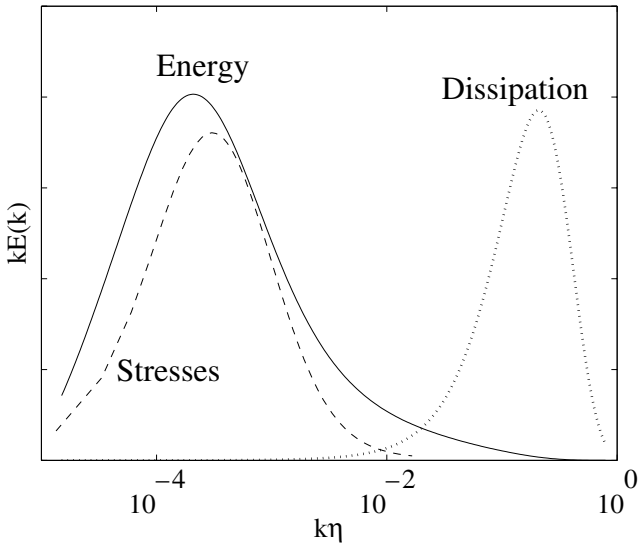

Figure 1. Premultiplied spectra for the kinetic energy, Reynolds stresses, and dissipation in a turbulent flow. The data are for the logarithmic region in a boundary layer from [3].

that the inertial energy cascade is independent of the dissipation mechanism, and replace the dissipation peak by a suitable model, in the hope that the dynamics of the large scales would not be disturbed. We will also speak briefly about the dual small-scale simulations (small-eddy simulation (SES)), in which the aim is to understand the dissipation mechanics independently of the details of the energy injection method.

The last two simulation techniques rely on theoretical assumptions which are unproven. The independence of the dissipation from the details of the energy injection, and vice versa, are both highly plausible models, but need theoretical or experimental confirmation. This assumed independence, and the related nature of the large-Reynolds-number limit, are the central problems of turbulence theory and, because of their connection with the basic assumptions of LES, they are also crucially important for industrial simulation techniques. In the absence of a proper theory the only hope of settling them is either by laboratory or by numerical experimentation. The subject of this paper is whether the latter can ever be expected to substitute for the former in this task.

Reynolds-averaged modelling will not be considered here, although it is obviously of great practical importance. Modelling, at least in the form in which it is used in applications, is the final result of turbulence theory, but it is not an experimental problem. It does not seek to support our understanding of how turbulence works, but rather to use the results after a proper theory is available.

3. The computational cost of the Reynolds number

3.1. Direct numerical simulations

The resolution requirements of direct simulations have often been discussed. A recent summary can be found in [5]. The number of required grid points can be estimated from the ratio L/η between the integral scale L, where the energy resides, and the Kolmogorov length η, where the dissipation is maximum. We know from Kolmogorov's theory that this ratio is proportional to $Re_\lambda^{3/2}$, where Re_λ is the microscale Reynolds number. In figure 1, where $Re_\lambda \approx 1500$, the ratio

is $L/\eta \approx 5000$, and the two spectral peaks are separated by a factor of about 1000. If we want to do things correctly, it is not enough to compute from one peak to the other. There is energy to the left of the energy peak, typically at scales which are a factor of 10 larger, and dissipation on scales of the order of 1.5η, which is three times smaller than the peak of the dissipation spectrum. Taking into account that at least two grid points are needed for each of the smallest relevant scales, the necessary grid diameter is of the order of

$$N = 1.7\, Re_\lambda^{3/2}. \tag{1}$$

It is clear from figure 1 that the Reynolds number in that flow is approximately the minimum for which the two spectral peaks are well separated. Applying to it equation (1) implies a grid of the order of $100\,000^3$ points for an asymptotically turbulent simulation.

We may save something by reducing the resolution slightly at both ends of the spectrum, or by using a slightly smaller Reynolds number. We could then perhaps reduce the diameter of our grid to 10^4, instead of 10^5, but the order of magnitude of the number of points that we need to compute 'true' turbulence is 10^{12}–10^{15}. Those are very large numbers, and to decide whether we will be able to compute a problem of that size in the foreseeable future we have to consider the performance and evolution of computers.

Denote by N the number of points along the edge of our three-dimensional grid. We need something like 10 variables per grid point, and at least $\log_2(N)$ operations per variable in each time step. Note that this last estimate is true as much for spectral codes as for good finite-difference schemes using multi-grid algorithms to solve the Poisson matrices. The memory needed is then

$$M \approx 40N^3 \text{ (byte)}, \tag{2}$$

and the number of operations per step is

$$Q \approx 10N^3 \log_2(N) \text{ (flops s)}, \tag{3}$$

where we use flops (floating-point operations per second) as a unit of computer speed. Both requirements constrain our computer; the first one in terms of central storage, and the second one in terms of speed.

Consider now the cost of running the largest problem that would fit into a computer with a central storage of size M. We can estimate the computer time needed for each time step as

$$Q/S = 4 \log_2(N)[M \text{ (byte)}/S \text{ (flops)}], \tag{4}$$

where S is the computer speed. For $N \approx 10^3$–10^4, this becomes

$$Q/S \approx [M/S] \text{ (min)}. \tag{5}$$

The parameter M/S is a figure of merit of the computer, and has historically been approximately equal to one; large computers tend to have 1 Gbyte of central memory for each Gflops of computational speed. That means that the computer time needed to run a single step of the largest possible turbulence DNS does not depend too much on the problem size, and is always of the order of 1 min. This agrees reasonably well with the personal experience of the author during the last 20–30 years. Recent experience with modern parallel machines and with grids of the order of 10^9 points, using sustained processor speeds taken from [6], gives an estimate of 2.5 flops h for each grid point and per time step, which is equivalent to adding a factor of about three in front of equation (5).

We need to be careful with this estimate, because the parameter M/S is decided by computer manufacturers, and it could change in the future. Recent discussions of future computer architectures routinely use values of M/S of the order of 10^{-2} [7]. Those new architectures are driven by the requirements of high-speed graphics, rather than by simulations, and the

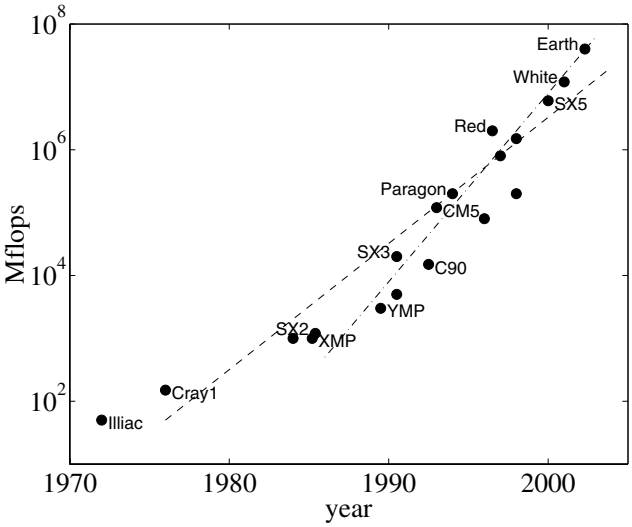

Figure 2. Historical development of computer speed in recent decades. The two straight lines correspond to the popular version of Moore's law that single-processor speed doubles every 18 months, which results in a factor of 100 per decade, and to the faster recent historical rate for parallel machines of 10^3 per decade.

Table 1. Speed and memory requirements for different isotropic turbulence simulations, and the expected date when they would be possible.

N	Re_λ	Memory	Speed	CPU time	Year
10^3	70	50 Gbyte	50 Gflops	1 month	1993
10^4	300	50 Tbyte	50 Tflops	1 year	2003
10^5	1500	50 Pbyte	50 Pflops	10 years	2015?
10^5			500 Pflops	1 year	2018?

speed–memory combinations needed for the two applications are different. The hierarchical memory structure imposed by the growing disparity between the speed of processors and that of memory chips is also making large memories harder to implement, and some of the largest computers which have recently come online have relatively low memory ratios. While ASCI White and ASCI Q, which were the two most powerful machines in the United States in 2002, still have $M/S \approx 1$ (at $S \approx 8$ Tflops), the Earth Simulator in Japan delivers sustained speeds of 35 Tflops, but only has 10 Tbyte of central memory.

This trend is dangerous for computational fluid mechanics, because memory is as much a limiting resource as speed for the simulations mentioned above, and it is important that the simulation community transmits this message to manufacturers. The rest of this section assumes the historical value of $M/S \approx 1$, but it should be kept in mind that this assumption rests on an active effort by the fluid mechanics community to steer computer centres towards our requirements.

To estimate how long it takes to compute a given Reynolds number, and the likely date at which it could become possible, we need to decide how many time steps have to be run. This is fixed by how long it takes to compile good statistics. A good rule of thumb, because of the numerical CFL condition, is that we need $2N$ time steps for a fluid particle to cross the computational box, and at least $20N$ steps for the statistics to converge. Using those estimates and equation (5) we get table 1.

We had 50 Gflops supercomputers around 1993, and a 40 Tflops machine in 2002. Computer speed has increased historically by a factor of 100–1000 every 10 years (figure 2), and computer scientists seem confident that the same trend will hold for some time. Nobody knows whether that will be true in 40 years, but for the moment the formula works and the proper technologies keep appearing at the right pace.

We can anticipate a 50 Pflops computer by the year 2010–2020, and we will then be able to start computing our first incontrovertible example of turbulence, with a real separation between the energy and the dissipation scales. After that we should be able to test for the first time whether our theories agree in detail with experiments or not. In practice we may have to wait a little longer, because the total times given in table 1 are for dedicated use of the full machine, but it is clear that we should be well on our way to computing asymptotic turbulence by 2015.

3.2. Asymptotic simulations and LES

Beyond that we may keep relying indefinitely on computer power to increase our Reynolds number, or we may decide that we do not need to do anything else. In fact, such a 10^{15}-points simulation could perhaps be considered as a checking exercise for a cheaper asymptotic LES, which could be accessible today. The argument is that, if we believe in classical theory, once we have computed the whole spectral energy peak, and the whole Reynolds stress peak, classical theory implies that we should not need to do much more. In a sense, we would have done the ultimate LES. All that it would require would be some dissipation mechanism, not necessarily a physically plausible one. The energy-containing scales would be asymptotically independent of the Reynolds number. It follows from figure 1 that the width of the energy peak is only about 10^4, and we should be able to save a factor of 10 in grid diameter with respect to the full DNS, because we would not need to resolve the dissipation peak very well. So, with a grid which is only $10\,000^3$ it should be possible to do asymptotic LES, and it should be essentially the same as using an infinite number of points. Looking at table 1, such a simulation could be started today, and it could easily be finished within the next 10 years, even without the use of whole dedicated machines.

Something else that needs to be done is an asymptotic SES, which could be defined as a simulation which resolves the whole dissipation peak, even if the energy injection at the large scales is done artificially. This is something that is not very practical from the engineering point of view, since engineers are usually interested in the energy or in the stresses but, from the physical point of view we need to answer questions such as whether the small scales are independent of the character of the forcing, whether turbulent flow near a wall is independent from the outside flow, or whether molecular mixing is independent of the larger scales. Questions about the origin of intermittency, which would not be addressed by the LES mentioned above, are probably more easily settled in this way than by full direct simulations at ever increasing Reynolds numbers. The same is true for applied problems in which intermittency is important, such as the generation of contaminants in turbulent combustion.

Asymptotic SES is cheaper than asymptotic LES because, looking at the spectra, the dissipation peak is narrower than the energy one. A grid of 1000^3 should be enough. This is actually what most simulations of isotropic turbulence have been doing for some time, since they tend to be well resolved at the small-scale end, and worry less about the forcing. The latest simulations in this series are the 4096^3 isotropic boxes in [8], at a nominal $Re_\lambda \approx 1200$, and they can easily be considered asymptotic. There is some extra physics that needs to be done in this area, and we have to learn how to do optimal forcing, and how important it is, but the grids needed are already accessible, and the first of the several 'definitive' turbulence computations that need to be done has already been accomplished.

4. The problem with walls

The conclusions of the previous section do not apply in the neighbourhood of walls, especially regarding the existence of an asymptotic LES. The argument was that the only modes that need to be computed in LES are those which cannot be described by an isotropic Kolmogorov cascade, and that there is only a fixed finite number of those modes. The rest of the problem can be accommodated by some universal model. In wall-bounded turbulence the first assumption remains true, but the second one fails. The anisotropic modes are roughly those for which the Reynolds stress cospectrum does not vanish, and we saw in figure 1 that they are confined to scales larger than a given fraction of the integral scale. This is still true as we approach a wall across the turbulent logarithmic layer [9], but the difference in that case is that the integral scale decreases continuously as we approach the wall.

Consider the number of non-universal modes which need to be computed in the LES of a cubic fluid volume whose side δ is determined by the geometry of the flow, such as the boundary layer thickness. In the neighbourhood of the wall the integral length decreases linearly as $L \approx \kappa y$, where κ is the Kármán constant and y is the distance to the wall. The eddies remain anisotropic as long as their size is larger than some fraction, $\Delta x \sim y$, of the integral scale, and this trend continues until y is of the order of the viscous wall unit, which is the wall equivalent of the Kolmogorov viscous scale. The number of anisotropic modes in a slab of thickness $\mathrm{d}y$ is then

$$\mathrm{d}N \sim \delta^2 \,\mathrm{d}y/\Delta x^3, \tag{6}$$

and their total number is given by the integral

$$N_T \sim \int_{y_0}^{\infty} \delta^2 \,\mathrm{d}y/y^3 \sim \delta^2/y_0^2, \tag{7}$$

where y_0 is some inner wall distance that determines the number of modes. In the absence of a good model for anisotropic turbulence, we can only choose this limit as a fixed number y_0^+ of viscous wall units, so $y_0 = \nu y_0^+/u_\tau$, where ν and u_τ are the kinematic viscosity and the friction velocity. The number of anisotropic modes then becomes

$$N_T \sim (u_\tau \delta/\nu)^2 = Re_\tau^2. \tag{8}$$

The quantity Re_τ is proportional to Re_λ^2, and the estimate (8) is only slightly lower than the estimate for the full direct numerical simulation of the whole flow, $N_T \sim Re_\tau^{9/4}$. Both estimates increase without limit with the Reynolds number, and there is not in this case a fixed number of modes beyond which we can consider a LES to be asymptotic. Our only hope is for a DNS which is large enough for the energy and the dissipation peaks to be 'asymptotically' far apart. Note that the estimate in (8) is not linked to any particular LES model, being just a count of the number of 'non-Kolmogorov' modes per unit volume of wall turbulence. These modes depend on more parameters than just the rate of energy dissipation, and they are unlikely to be modelled correctly by any simple approximations to the subgrid cascade.

4.1. Direct numerical simulations of wall flows

What prevents us from designing an asymptotic simulation of a wall-bounded turbulent flow, in the same spirit as those suggested above for isotropic ones, is the lack of a theory for the structure of the logarithmic layer. This intermediate layer, connecting the viscous structures near the wall with the larger ones of the core region, plays the same role in wall-bounded flows as the inertial range in homogeneous turbulence. It is generally believed that its dynamics are universal in the same sense as that of the inertial range, but it differs from the latter in that it transports energy as much in the physical space as across different scales. The details of this

Table 2. Characteristics of some representative channel-flow simulations.

Re_τ	L_x/δ	L_z/δ	Points	Year
180	12	6	5 M	1987 [11]
590	6	3	40 M	1997 [12]
550	25	12	600 M	2001 [13]
950	25	9	4 G	2003 [14]
1 900	3	1.5	450 M	2003 [14]
10 000	12	6	900 T	2015?

transport are not well understood, and there is even some controversy on whether the mean velocity profile of this layer is really logarithmic [10].

An equally important region is the near-wall buffer layer, where the effect of viscosity begins to be noticeable, and which is the wall-bounded equivalent of the near-dissipative range of isotropic turbulence. The difference in this case is that the dissipative structures in the lower buffer region are also responsible for a large fraction of the turbulent energy production, and determine, among other things, the overall friction coefficient.

As with most other aspects of turbulence we have no quantitative theory for either the buffer or the logarithmic layers. The difference between wall-bounded flows and isotropic or free-shear turbulence is that, while the inertial and dissipative scales of the latter are essentially passive recipients of the energy generated at the larger scales, the buffer and logarithmic layers of wall-bounded flows are crucial participants in the turbulence dynamics. While we can ignore the effect of the smaller scales in isotropic turbulence, except for their role as a dissipative energy dump, we cannot even compute the mean velocity profile of wall-bounded flows without understanding how energy is generated and transmitted by the buffer and logarithmic layers.

It is for this reason that wall-bounded flows were among the first to be simulated directly, and why the effort to extend those simulations to higher Reynolds numbers and to different aspects of the flow continues to this day. Some representative simulations are collected in table 2, which also tracks the evolution of direct numerical simulations in general. All the simulations in table 2, except perhaps the projected last one, use essentially the same pseudospectral code at comparable resolutions [11], and they are therefore broadly comparable.

There is a large spread of Reynolds numbers and box sizes in table 2, reflecting both the computational power available at the different dates, and the different aspects of the flow that were addressed by each simulation. The Reynolds number of the earliest simulation was too low to allow much beyond the study of the near-wall layer, but [11] was the seminal contribution to the field, and had to balance scale separation with a computational domain large enough for discarding possible effects of the box size on the large scales.

The simulations in [12] and the smallest box in [14] were designed to push the Reynolds number at the expense of box size, and can probably be best classified as SESs of the inner logarithmic layer. The large-box simulations in [13, 14] continue the line initiated in [11] of balanced resolution of the small and large scales.

For the same scale separation, wall-bounded simulations are more expensive than isotropic ones, both because the inhomogeneity of the flow results in slightly larger grid requirements, and because the presence of a mean flow velocity requires shorter time steps. The $Re_\tau = 950$ channel in [14] is the largest wall-flow simulation attempted at the time of writing this paper, but it only uses 6% as many grids points as the isotropic simulation in [8].

The flow in the logarithmic layer is very anisotropic and contains very long structures [15]–[19], with appreciable energy at lengths of 20–30 times the boundary layer thickness. In principle

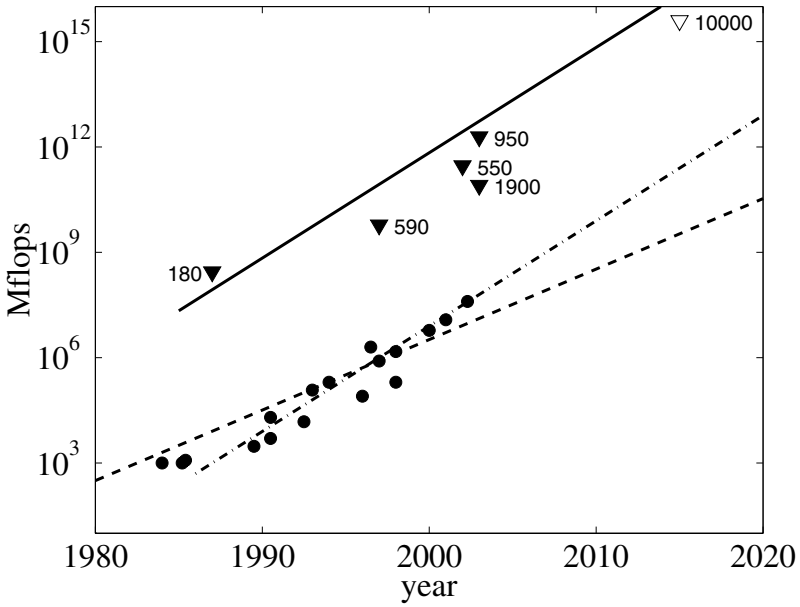

Figure 3. Computational costs of selected simulations of channel flows. The circles are computers in figure 2, with the trend lines for their computational speed in flops. The triangles are simulations with their total cost in flops s. The solid line is one day of dedicated time in a machine following the historical trend for parallel computers.

those structures have to be accommodated by the numerical grid, which should also be fine enough to resolve the dissipative structures near the wall. The required resolution is of the order of $\Delta x^+ = 8$ and $\Delta z^+ = 4$ [11]. The Chebyshev polynomials used in the wall-normal direction of most pseudospectral simulations overresolve the wall region when a reasonable resolution $\Delta y^+ \approx 8$ is used near the channel centre.

The simulations in [13, 14] were designed to test the scaling of the logarithmic-layer structures. What they show is that there is an inner spectral peak at $\lambda_x^+ \times \lambda_z^+ \approx 800 \times 100$, which is strongest near the wall and scales in wall units, and an outer one which peaks at the upper edge of the logarithmic layer at $\lambda_x \times \lambda_z \approx 4\delta \times \delta$, scaling in outer units. This separation of viscous dissipation structures near the wall and energy-carrying ones in the channel core is the equivalent in wall-bounded flows of the separation of the energy and dissipation peaks in figure 1.

Using the grid resolution mentioned above, the number of collocation points for a spectral box of size $L_x\delta \times \delta \times L_z\delta$ is

$$N_T = 0.012 L_x L_z \, Re_\tau^3. \tag{9}$$

Experience dictates that reasonable statistics need simulations that run for at least ten wash-out times, where a wash-out is defined as the time taken by a fluid particle at the centreline to cross the computational box. The mean velocity of those particles also limits the numerical time step. If the CFL of the simulation is set at about 0.5, the number of time steps per wash-out is

$$N_t = 0.4 L_x \, Re_\tau. \tag{10}$$

Using the estimates in section 3 for the computational cost per grid point, the total cost of a wall-bounded simulation is approximately

$$Q_T = 40 L_x^2 L_z N_w \, Re_\tau^4 \text{ flops s}, \tag{11}$$

where N_w is the number of desired wash-outs. This estimate is plotted in figure 3, together with the evolution of computer power. It is interesting that most of the simulations in table 2 required a total time expenditure equivalent to somewhat less than one dedicated CPU-day of the largest machine available at the time, although all of them took considerably longer in the machines in which they were run.

We can now address the question of the date at which an asymptotic DNS of a wall-bounded flow could be run. While the cost of wall-bounded simulations is larger than for isotropic ones, the fact that the energy and the dissipation peaks are located at different wall distances suggests that we may reach some asymptotic behaviour with lower scale separations than in the isotropic case. The simulations in [14], for which the separation is about 5–10, show that the coupling between the outer and the wall spectral peaks is weak, and that even the relatively small Reynolds numbers of those simulations allow tentative conclusions to be drawn about their asymptotic interaction. Assume that we require a scale separation factor of 50 between the large- and small-scale spectral peaks, whose ratio is $4\delta/800^+ = 5 \times 10^{-3}\, Re_\tau$. This would require $Re_\tau \approx 10^4$. Such a simulation has been included in table 2 and in figure 3. If we assume the same level of effort as has been used for historical channels, it could be comfortably run by 2015, and would provide a data set against which to test theories for the logarithmic layer and for the interaction between the wall- and outer-layer structures. If a group were ready to invest one CPU-month instead of one CPU-day, the same simulation could be initiated by 2007.

A truly asymptotic simulation, with a scale separation of 500 between peaks, should be possible in one CPU-month by 2020. The Reynolds number of such a simulation, $Re_\tau = 10^5$, would be in the high range of industrial flows, and beyond the range of most laboratory experiments.

4.2. Thought experiments

Wall-bounded flows, and specially the near-wall sublayer and buffer regions, have been the driving problem behind many of the 'thought' numerical experiments mentioned in the introduction. The fact that excess turbulent energy is generated very near the wall suggests that the dynamics of that region can be studied in isolation from the rest of the flow, with the interactions with the outer layer appearing as lower-order corrections.

One of the first thought experiments in the field was the 'minimal' simulation introduced in [20], in which the spatially periodic computational box was made just long and wide enough to contain a single near-wall structure. The result was the replacement of the chaotic superposition of near-wall structures in real turbulent flows with an ordered periodic array of identical units which could then be studied individually. It was shown that the near-wall flow was not greatly disturbed by this substitution. Since the lateral dimensions of the minimal box were not large enough to accommodate the outer-flow structures, these experiments also gave initial support to the idea of an autonomous wall region.

Similar techniques, together with the use of very low Reynolds numbers and different computational techniques, led to the isolation of permanent and weakly periodic individual structures in Poiseuille [21]–[23] and Couette flows [24]–[26], with dimensions and characteristics very similar to those observed in near-wall turbulence. Not surprisingly the basic structure turns out to be a pair of staggered quasi-streamwise vortices flanking a low-velocity streak, very similar to those statistically identified from full simulations of turbulent boundary layers [27].

All these numerical experiments can be considered extreme cases of the SES discussed above for isotropic turbulence, since they study the small scales in the hope that their interaction with the larger ones is only secondary. An experiment specifically designed to test that hypothesis in wall-bounded flows was the 'autonomous' flow introduced in [2], in which the wall was made to evolve in the absence of any outer turbulence. It was found to be enough to keep an undisturbed

region below $y^+ = 50$ to retain turbulent characteristics which were virtually indistinguishable from those of full flows, and the small differences observed gave information on the interactions between the inner and outer regions of real wall turbulence [28].

These partial-flow experiments are typically much less costly than full simulations. Even the very large autonomous boxes used in [28] to test the origin of the large-scale organization of the near-wall region contain only a few million grid points. The same is true of the small-box full-depth simulations cited in table 2, which can be considered as minimal flow units for the logarithmic layer. Because their computational boxes are four times shorter and narrower than even full boxes of moderate dimensions, it follows from equation (11) that these simulations are 64 times cheaper than full ones, which translates into about 6 years of computer development. A simulation in a box of size $3\delta \times 2\delta \times 1.5\delta$, at the 'asymptotic' $Re_\tau = 10^4$, would only cost 15 CPU-days in present machines.

Because these simulations refer to 'physically wrong' systems, in which some aspect of the real flow has been purposely removed, these thought experiments have to be used with care, and only when the nature of the simplification is well understood and tested against a true flow. The core region of all the minimal channels, whether those designed to study the near-wall region at low Reynolds numbers, or the higher-Reynolds-number ones intended for the logarithmic layer, are typically wrong, and it would be incorrect, for example, to use those systems to characterize global flow quantities such as the change of the friction coefficient due to a particular control algorithm. With proper precautions, however, conceptual simulations are probably the method of choice to answer specific questions in turbulence research.

5. Simulations versus experiments: conclusions

It is interesting at this point to get back to the initial question of whether simulations can be considered as substitutes for experiments. We have seen several cases in which direct or LES will reach asymptotic Reynolds numbers in one or two decades, and other cases in which conceptual experiments have already given answers to questions which would be difficult to test in the laboratory. All these examples refer to canonical flows in simple geometries. Adding general geometry to a simulation code, or stiff physics such as mixing at high Schmidt numbers, can add orders of magnitude to its computational cost. There is therefore a large class of problems for which experiments will remain cheaper than simulations for the foreseeable future.

There is however a certain sense in which simulations are already comparable with experiments.

Consider isotropic turbulence. There are now well-characterized isotropic turbulence experiments at $Re_\lambda \approx 1000$, but when the original simulation in [29] was published in 1991, at $Re_\lambda \approx 140$, its Reynolds number was higher than that of the isotropic experiment commonly used as reference at the time [30]. One of the currently open questions in isotropic turbulence is the behaviour of intermittency at large Reynolds numbers. A transition in the growth of the flatness of the velocity derivatives was reported at $Re_\tau \approx 700$ in [31], but it is not found in the compilation of data in [32]. The discrepancy remains unresolved to this day, in spite of attempts by various groups, due to a difficulty which should be familiar to any simulator: the instrumental resolution required to span the ratio between the largest and the smallest flow scales at those the Reynolds numbers. Although the required post-processing has still not been performed, it is interesting to note that the sequence of simulations in [8, 33], using essentially similar codes, forcing schemes, and numerical resolutions, already spans the Reynolds number range in question, and could throw light on the controversy in a way that experiments are having difficulties doing.

It is also worth remembering that most of the laboratory experiments cited in section 4.1 to support the existence of large anisotropic scales were done in pipes at $Re_\tau \approx 2 \times 10^3$, which is

comparable to the Reynolds numbers of some of the simulations in table 2. Even today it is hard to find reliable laboratory experiments regarding the structural properties of turbulent channels beyond $Re_\tau \approx 1000$. The experiments on very-near-wall turbulence, where most measurement schemes have instrumental problems, have been calibrated for years against simulations, not the other way around.

There are therefore already instances in which brute-force DNS matches or surpasses experiments, even from the point of view of the Reynolds number. Before closing this paper it might be useful to reflect a moment on what the far future might bring, even if only because of our responsibility as educators. The far future is the $100\,000^3$ isotropic computation, or the $Re_\tau = 10^4$ channel, that we labelled above as asymptotic. Both could be run by 2015–2020. Before that we will have run the ultimate LES, at $10\,000^3$, which could be started today, and the ultimate SES which has probably already been done. These will all be long, one-of-a-kind research simulations. But to extrapolate further in the future, as computers keep getting faster and we decide that we do not need to compute anything bigger. Even if we take the conservative estimate of an increase in computer speed by a factor of 100 every decade, we will be able to run our ultimate LES in 2 days by 2020, in 30 min by 2030, and in 20 s by 2040. After that, engineering turbulence theory, at least for free-shear flows in simple configurations, becomes free. We will be able to answer all theoretical questions about turbulent energy and Reynolds stresses by pushing a button and waiting for 1 min. It will still take some computer time to do the research DNS, although not much, but we need to consider whether we would really want to do them. What happens after we are able to answer all the practical questions without theoretical insight? Perhaps more to the point, would it make sense to keep funding turbulence theory at that stage?

We have no space here to address all the implications of the last two questions, but note that, more than to the relation of simulations with experiments, they point to real issues about what the role of turbulence theory will be in 40 years, and about whether it would make sense from an engineering point of view to try to develop good turbulence models beyond those needed for rough design estimates.

Acknowledgments

Partial funding for this work was provided by the Spanish CICYT contract BFM2000-1468 and by the US ONR grant N0014-00-1-01416. It is a pleasure to acknowledge many useful discussions with J C del Álamo and R D Moser.

References

[1] Moin P and Mahesh K 1998 Direct numerical simulation: a tool in turbulence research *Annu. Rev. Fluid Mech.* **30** 539–78
[2] Jiménez J and Pinelli A 1999 The autonomous cycle of near-wall turbulence *J. Fluid Mech.* **389** 335–9
[3] Saddoughi S G and Veeravali S V 1994 Local isotropy in turbulent boundary layers at high Reynolds numbers *J. Fluid Mech.* **268** 333–72
[4] Townsend A A 1976 *The Structure of Turbulent Shear Flow* 2nd edn (Cambridge: Cambridge University Press) p 55
[5] Pope S B 2000 *Turbulent Flows* (Cambridge: Cambridge University Press) pp 344–57
[6] Top 500 supercomputer sites 2002 http://www.top500.org
[7] Messina P, Culler D, Pfeiffer W, Martin W, Oden J T and Smith G 1998 Architecture *Commun. ACM* **41** 36–44
[8] Kaneda Y, Ishihara T, Yokokawa M, Itakura K and Uno A 2003 Energy dissipation rate and energy spectrum in high resolution direct numerical simulations of turbulence in a periodic box *Phys. Fluids* **15** L21–4
[9] Baggett J S, Jiménez J and Kravchenko A G 1997 Resolution requirements in large-eddy simulations of shear flows *CTR Annu. Res. Briefs* 51–66

[10] Barenblatt G I, Chorin A J and Prostokishin V M 2000 Self-similar intermediate structures in turbulent boundary layers at large Reynolds numbers *J. Fluid Mech.* **410** 263–83
[11] Kim J, Moin P and Moser R 1987 Turbulence statistics in fully developed channel flow at low Reynolds number *J. Fluid Mech.* **177** 133–66
[12] Moser R D, Kim J and Mansour N N 1999 Direct numerical simulation of turbulent channel flow up to $Re_\tau = 590$ *Phys. Fluids* **11** 943–5
[13] del Álamo J C and Jiménez J 2003 Spectra of the very large anisotropic scales in turbulent channels *Phys. Fluids.* **15** L41–4
and see also
del Álamo J C and Jiménez J 2001 *CTR Annu. Res. Briefs* 329–42
[14] del Álamo J C, Jiménez J, Zandonade P and Moser R D 2003 The spectrum of turbulent channels up to $Re_\tau = 1900$ *J. Fluid Mech.* submitted
[15] Perry A E and Abell C J 1977 Asymptotic similarity of turbulence structures in smooth- and rough-walled pipes *J. Fluid Mech.* **79** 785–99
[16] Perry A E, Henbest S and Chong M S 1986 A theoretical and experimental study of wall turbulence *J. Fluid Mech.* **165** 163–99
[17] Hites M H 1997 Scaling of high-Reynolds number turbulent boundary layers in the National Diagnostic Facility *PhD Thesis* Illinois Institute of Technology
[18] Jiménez J 1998 The largest structures in turbulent wall flows *CTR Annu. Res. Briefs* 137–54
[19] Kim K C and Adrian R J 1999 Very large-scale motion in the outer layer *Phys. Fluids* **11** 417–22
[20] Jiménez J and Moin P 1991 The minimal flow unit in near-wall turbulence *J. Fluid Mech.* **225** 221–40
[21] Toh S and Itano T 2001 On the regeneration mechanism of turbulence in the channel flow—role of the traveling-wave solution *Proc. IUTAM Symp. on Geometry and Statistics of Turbulence* ed T Kambe, T Nakano and T Miyauchi (Dordrecht: Kluwer) pp 305–10
[22] Jiménez J and Simens M P 2001 Low-dimensional dynamics in a turbulent wall *J. Fluid Mech.* **435** 81–91
[23] Waleffe F 2001 Exact coherent structures in channel flow *J. Fluid Mech.* **435** 93–102
[24] Nagata M 1990 Three-dimensional finite-amplitude solutions in plane Couette flow: bifurcation from infinity *J. Fluid Mech.* **217** 519–27
[25] Waleffe F 1998 Three-dimensional coherent states in plane shear flows *Phys. Rev. Lett.* **81** 4140–3
[26] Kawahara G and Kida S 2001 Periodic motion embedded in plane Couette turbulence: regeneration cycle and burst *J. Fluid Mech.* **449** 291–300
[27] Robinson S K 1991 Coherent motions in the turbulent boundary layer *Annu. Rev. Fluid Mech.* **23** 601–39
[28] Jiménez J, del Álamo J C and Flores O 2003 The large-scale dynamics of near-wall turbulence *J. Fluid Mech.* submitted
[29] Vincent A and Meneguzzi M 1991 The spatial structure and statistical properties of homogeneous turbulence *J. Fluid Mech.* **225** 1–20
[30] Comte-Bellot G and Corrsin S 1971 Simple Eulerian time correlations of full and narrow-band velocity signals in grid-generated isotropic turbulence *J. Fluid Mech.* **48** 273–337
[31] Tabeling P, Zocchi G, Belin F, Maurer J and Willaime H 1996 Probability density functions, skewness, and flatness in large Reynolds number turbulence *Phys. Rev.* E **53** 1613–21
[32] Van Atta C W and Antonia R A 1980 Reynolds number dependence of skewness and flatness factors of turbulent velocity derivatives *Phys. Fluids* **32** 252–7
[33] Jiménez J, Wray A A, Saffman P G and Rogallo R S 1993 The structure of intense vorticity in isotropic turbulence *J. Fluid Mech.* **255** 65–90