

Turbulence:

V&V and UQ Analysis of a Multi-scale Complex System

Presented by Parviz Moin, Curtis W. Hamman and Gianluca Iaccarino

1 Introduction

Turbulent motions of liquid and gases are ubiquitous and impact almost every aspect of our life, from the formation of hurricanes to the mixing of a cappuccino [1, 2, 3, 4, 5]. Nearly all human endeavors must contend with turbulent transport: turbulence is the rule, not the exception, in fluid dynamics. Energy, transportation, and the environment function on length and time scales where turbulence rapidly develops: even a simple, slow stroll leaves behind a turbulent wake with a wide range of scales (see table 2 for examples). With the advent of faster computers, numerical simulation of turbulent flows is becoming more practical and more common [6, 7]. In this short note, we describe the fundamental physics and numerics needed for accurate turbulent flow simulations. Several illustrative examples of such simulations are presented. In addition, we show how simulations of turbulent flows contribute to breakthroughs in clean energy technologies, reduce dependence on volatile oil reserves and develop carbon-free sources of energy [8]. We also show how increased industrial competitiveness in high-tech sectors such as transportation and aerospace is made possible by high-performance computing, especially where physical experiments are impossible, dangerous or inordinately costly to perform [8].

1.1 Aircraft and jet engines

Aerospace industry was one of the early pioneers to apply computational fluid dynamics (CFD) to the design process. The primary motivation has always been to reduce expensive physical testing (e.g., wind tunnel tests) and the number of prototypes built during the design cycle. In the early seventies, computations (using linearized inviscid approximations of flow equations) of transonic wings convincingly demonstrated the resulting cost savings of CFD as well as the potential limitations imposed by ground-test facilities, such as wind tunnel walls. Higher fidelity approximations were then applied to simulations of complex three-dimensional aircraft beginning with non-linear inviscid approximations and culminating with Reynolds-Averaged Návier-Stokes (RANS) methods in the 1990's. Each successive stage of approximation allowed a simulation to capture a new class of physics. Effective use of high-fidelity methods, however, became practical only when the requisite computer power and algorithms for that stage became available [2, 6].





Figure 1: Effect of high-performance computing on wind tunnel development testing for Boeing aircraft. The number of wing prototypes built and tested each year in wind tunnels during the design phase of several Boeing aircraft are shown (black bar-graph, linear scale). Peak computer speed (in floating point operations per second, flops) of top supercomputers shown (gray circles, log scale). In thirty years since 1980, available computational power has increased more than 8 orders of magnitude at a rate faster than $10\times$ every 5 years [2]. Use of available computing resources and technology has helped reduce the need for costly physical prototyping and wind tunnel testing [9, 7, 6]. Further reductions, however, require more than just powerful computers as indicated by the testing plateau experienced during the B777 and B787 wing designs [9, 10]. The low-fidelity algorithms and physical models used in these designs reached a plateau. Confidence in such simulations could not grow with increased computer power owing to irreducible epistemic modeling errors. To certify future aircraft designs with less reliance on physical testing, high-fidelity methods, such as LES, are needed that reduce modeling errors and uncertainty in proportion to computer speed.


As figure 1 shows, the predictive capability of mainstream CFD in the design cycle seems to have reached a plateau. Owing to the use of CFD, the number of wind tunnel tests in the development of Boeing jet liners reduced dramatically from 77 tests in the 1980s (B757 and B767) to 11 in the 1990s (B777). This trend has not continued with the development of the newest Boeing model (B787) where the required number of tests has remained at 11 (an earlier Boeing projection for this figure was 5) [9, 10]. This is in spite of the fact that the available computer power in the same period has increased by nearly four orders of magnitude. That is, although the design cycle time (a highly coveted factor in aerospace design) has decreased significantly, the computations have not achieved the state of predictability and accuracy required to reduce the required number of wind tunnel tests. A similar plateau is observed in the jet engine industry where the number of expensive high-pressure rig tests required for engine certification has not declined despite the rapid growth

State-of-the-art in 1997



Schematic (top) of a coaxial combustor with air-methane co-flow and an instantaneous snapshot of the mixture fraction (bottom) from the first LES of turbulent combustion in a research combustor performed at Stanford University [11]. The simulations used a flamelet and progress-variable approach (FPVA) with over 2 million grid points computed on the ASCI Red platform at Sandia National Laboratories. The highest-fidelity simulations of reactive flow in 1997 were restricted to simple geometries and gas-phase only physics.

State-of-the-art in 2007



Flow inside a realistic Pratt & Whitney gas turbine engine combustor (1 sector). Normalized temperature contours shown for a mid-plane at cruise conditions (high Reynolds number). Liquid fuel is injected and breaks up into drops that evaporate and react with co-flowing air. Unstructured grids were used for the diffuser surrounding the combustion chamber, the injectors, swirlers, dilution holes, etc., which constitute a geometrically very complex configuration [12]. Combustor simulation integrated into a complete jet engine simulation, see figure 2.

Table 1: Evolution of high-fidelity computational capability.

in computing capability [7]. To certify future aircraft and jet engine designs with less reliance on physical testing, higher fidelity methods than the current industry standard are needed.

The RANS approach remains the industry standard in CFD for aerospace and jet engine applications [9, 7, 13]. Although derived from first principles, the averaging process introduces unclosed terms to model all scales of turbulence but the mean flow. These unclosed terms are modeled phenomenologically, and not based on first principles. As a result, RANS models have not proven to be sufficiently universal and predictive when applied to situations outside the domain of calibration. The RANS approach suffers from fundamental (epistemic) uncertainties in the functional forms of the turbulence closure models, which at least for the current models used in engineering practice, cannot be made more accurate by calibration against experimental or direct numerical simulation (DNS) data. Since RANS is limited by the (irreducible) accuracy of the turbulence closure models, the advances in computing hardware shown in figure 1 have not significantly improved the predictive capability of engineering design computations nor reduced reliance on physical testing. The effectiveness

pressure turbine blades at the combustor outlet are all important unsteady phenomena not simply accounted for by truncated boundary conditions and low-fidelity methods. One of the first high-fidelity simulations of an integrated full jet engine system to achieve this predictive capability was performed in the Center for Turbulence Research at Stanford University in collaboration with Pratt & Whitney as a part of the original ASCI Academic Strategic Alliances Program [12]. Figure 2 shows the approach taken to couple different component simulations with differing levels of fidelity. Overarching, interdisciplinary, high-fidelity simulations such as this are needed in the design and construction of aircraft and jet engines if further reductions in physical testing are to be made.

Rigorous methods are required for validation and verification of the outputs from these complex codes and numerical experiments. Verification is the process by which we demonstrate that a flow solver correctly solves its governing mathematical equations. A properly verified code is free of programming errors affecting the theoretical order-of-accuracy of the numerical algorithm [15]. Validation that numerical results conform to reality often involves comparison with physical experiment. Figure 2, for example, shows a partial validation of the jet engine simulations with high-pressure rig tests for the temperature profile at the combustor exit. A single sample of simulation data that exactly matches the physical experiment alone is, however, insufficient. We do not have precise knowledge of the inputs (e.g. boundary conditions) and physical models so that the simulation output predictions are also uncertain. An engine designer would like to, for example, estimate the likelihood the temperature at the combustor exit will exceed some permissible threshold, e.g. thermal fatigue. The objective is then to put uncertainty bars on numerical output that reflect the output statistical response to uncertainties in computational inputs such as boundary conditions and turbulence models [16]. An example of UQ applied to a complete aeroacoustic problem with uncertainty bars is presented in §4.6. Validation, verification and uncertainty quantification (UQ) are critical steps to improving confidence of numerical simulations and establishing credible predictions of turbulence.

Computer simulation tools have achieved wide spread use in the design and analysis of engineering devices and for scientific discovery. This has shortened the overall product design cycle and provided a refined understanding of the operating behavior of complex engineering systems. As a consequence, numerical simulations have led to a reduction in physical prototyping, testing and costs [7, 9, 6]. In spite of this considerable success, it remains difficult to provide objective confidence levels in the operability and design of complex engineering systems by numerical simulation alone. The complexity arises from uncertain inputs and turbulence models. As a result, especially in the area of reliability and safety, physical testing remains the dominant mechanism for certification of new devices. To move beyond the current plateau and effectively utilize available high-performance computing resources, widespread use of uncertainty quantification and high-fidelity simulation methods, such as LES, are necessary enabling tools in the design and analysis of complex engineering systems. With increasing computer performance, computer simulations are expected to replace expensive physical testing of design prototypes. Resolving more scales in turbulence is one way to achieve this goal.

2 Scales in turbulence



Configuration	Velocity	Length	Reynolds Number R_ℓ	Length Scales ℓ/η
	1 m/s	0.15 m	1×10^4	10^3
	40 m/s	0.075 m	2×10^5	10^4
	250 m/s	0.3 m	5×10^6	10^5
	1.5 m/s	1000 m	1×10^8	10^6

Table 2: Typical velocity and integral length scales in different turbulent flows: a person walking slowly, a pitcher throwing a fastball, an aircraft (wing) flying at cruise conditions, and cumulus cloud thermal updrafts. At higher Reynolds number, the range of scales becomes wider. Correspondingly, simulations of higher Reynolds number flow must either resolve or model these smaller scale turbulent motions.

Turbulent flows transport momentum, energy and material by simply carrying macroscopic fluid parcels to a new location. Rustling leaves stirred by the wind, a boiling pot of porridge and puffy cumulus clouds make this turbulent transport process readily visible to the naked eye [17, 18, 19]. But, your eyes only capture part of the story: the large-scale swirling, fluctuating eddies of size L and velocity u that transport most of the momentum, heat and moisture in a time $T_t \sim L/u$. These are the large billows in a cumulus cloud that turnover on timescales of about ten minutes (see table 2: $L/U = (1000 \text{ m})/(1.5 \text{ m/s}) \approx 600 \text{ s}$). If you look closer (e.g. out an airplane window as you pass above the clouds), you may see even smaller-scale corrugations (or eddies) embedded in these large scale billows (perhaps,




Figure 3a: Development of turbulence in a zero-pressure-gradient flat-plate boundary layer visualized by isosurface of Q -criterion colored by streamwise velocity [20, 21]. Three spanwise numerical rakes (the red, blue and green cylinders) record the instantaneous streamwise velocity signals shown in the upper-left plot (with matching colors). The streamwise position of the rakes is marked by a red line in the lower-right plot of the time and spanwise-averaged skin-friction profiles along the plate. Reynolds number $R_\theta \approx 600$.

similar to those seen in thermal updrafts shown in figure 6). Even smaller eddies (about a millimeter in dimension) also transport momentum, heat and moisture in such cumulus clouds. At a microscopic level, molecules also transport momentum by colliding with neighboring molecules, which tends to smooth out any macroscopic velocity gradients. Viscosity ν is a measure of how fast this intermolecular diffusion occurs. At very small length scales, viscosity can quickly smooth out velocity fluctuations by dissipating this small-scale energy into heat. But, over long distances L , viscosity acts slowly in a time of order $T_m \sim L^2/\nu$ [3]. The ratio of a molecular time scale to a large-eddy time scale $T_m/T_t \sim uL/\nu \equiv R$ is a Reynolds number.

Turbulent flows always occur at high Reynolds number ($R \gg 1$). Since $T_m/T_t \gg 1$ when $R \gg 1$, turbulence transports and mixes momentum, energy and material far faster than molecular diffusion alone. Conceptually, large eddies break-up into smaller and smaller eddies creating steeper velocity gradients until the action of viscosity damps out these rapid changes in velocity. As a result of this eddy cascade, widely separated fluid parcels are quickly brought close together where they can mix diffusively and react chemically. In the governing Návier-Stokes equations, non-linear inertial terms generate successively smaller-scale fluctu-




Figure 3b: See figure 3a for caption. This figure shows Reynolds number $R_\theta \approx 1200$ flow over a flat plate boundary layer. Note how the boundary layer thickens and length scales become larger [20, 21]. Simulations of H/K-type transition performed on Intrepid at the Argonne Leadership Computing Facility (ALCF) and with dedicated time on the LLNL BGL machine for improved statistical sampling.

ations, but the viscous terms prevent an infinite cascade to small scales by dissipation (long before length scales comparable with the mean free path are reached). Although turbulent motions themselves are not typically much smaller than a millimeter, turbulence acts to make inhomogeneities more vulnerable to the effects of molecular diffusion.

The range of scales or eddy sizes in a turbulent flow increases with Reynolds number. The fluid velocity field $\mathbf{u}(\mathbf{x}, t)$ becomes increasingly irregular in space and time. As figure 3 shows [20, 21], the Reynolds number grows with downstream distance along a flat plate. At the leading edge of the flat plate, a thin laminar (non-turbulent) boundary layer develops by the diffusion of momentum across the layer by transverse convection and viscosity alone. The laminar momentum deficit at the wall diffuses outward. At first, small perturbations are quickly damped by viscosity but, inevitably, the most unstable disturbances amplify, begin to grow and generate smaller scales (e.g. see the Tollmien-Schlichting waves, which appear as early sinusoidal fluctuations in the spanwise-averaged skin-friction coefficient, and the Λ -shaped vortices visible in figure 3a). Eventually a continuum of smaller scales are generated by these instabilities causing the laminar boundary layer to transition, as evidenced by the sudden jump in skin-friction, into a turbulent boundary layer. As the flow becomes turbulent, velocity fluctuations (in the upper left inset) have finer scales and the largest flow




Figure 4a: Instantaneous streamwise velocities at three different heights from the wall at Reynolds number $Re_\theta = 900$ from the boundary-layer simulation of [23]. Variables made dimensional assuming air flowing with free-stream velocity $U_\infty = 15$ m/s along a wall, which corresponds to the velocity and time scales seen in the thin boundary layer on automobiles that is only a few millimeters thick.

structures (or eddies) continue to grow downstream as the boundary layer thickens. These turbulent fluctuations increase the rate at which freestream momentum is transferred toward the wall (e.g. by engulfing freestream fluid at the edge of the boundary layer). The wall is stationary and, as one approaches the wall, the fluid velocity must also be zero (i.e. air does not penetrate nor slip relative to a solid surface). As a result, this enhanced turbulent mixing converts freestream momentum into a tangential shear force at the wall known as the skin-friction drag [22], which is shown in the lower-right inset of figure 3.

Figure 4 shows a time history $u(t)$ of the streamwise velocity from a turbulent boundary layer as in figure 3. The local velocity fluctuates about its mean (time-averaged) velocity by 30% or more of $\bar{u}(x, y)$. The fluctuating velocity, i.e. the difference $\mathbf{u}' = \mathbf{u} - \bar{\mathbf{u}}$, varies irregularly whereas the mean $\bar{\mathbf{u}}(x, y)$ is stationary. Disordered fluctuations conceptually correspond to a superposition of turbulent eddies of different sizes ℓ each having a characteristic velocity $u(\ell)$ and time scale $T_\ell = \ell/u(\ell)$. Both the velocity $u(\ell)$ and timescale $\tau(\ell)$ tend to decrease as the eddy size (or scale of motion) ℓ decreases. Small eddies then appear as




Figure 4b: See previous for caption. Fluctuating streamwise velocities shown.

low-amplitude, rapid wiggles in turbulent velocity profiles (see figures 3-4) whereas large eddies correspond to the large-amplitude, low-frequency undulations. Large eddies break-up into smaller ones thereby transferring energy to smaller scales. This cascade continues until the smallest eddies with sufficiently small Reynolds number $R(\ell) = u(\ell)\ell/\nu$ can breakup no further as viscous stresses rapidly dissipate their kinetic energy in the form of heat.

Turbulent motions consume energy. This energy may come from mean flow velocity gradients, pressure gradients, instability, or other production mechanisms. For a flat plate boundary layer (see figures 3-4), the velocity difference between the freestream and no-slip wall supplies energy to the large eddies, which sustains smaller scale eddies. Mean flow energy is converted to large-scale eddies that are themselves unstable (i.e. their eddy Reynolds number $R_\ell \gg 1$) and they breakdown into smaller scales that then break-up into smaller eddies until the smallest scale is reached. If this energy source were removed (e.g. by translating the wall at the freestream speed with a moving belt), then the turbulence would dissipate and decay. The rate of dissipation ε is therefore determined by the energy of the largest eddies. These eddies have energy of order u_0^2 and timescale $T_0 = \ell_0/u_0$. The rate of energy supply to small-scale eddies is then of order $u_0^2/T_0 = u_0^3/\ell$. Since the supply rate should equal the dissipation rate, we have $\varepsilon \sim u_0^3/\ell_0$. Since small-scale eddies fluctuate rapidly in comparison to (relatively slow) large eddies, one may assume that small-scale




Figure 5: Cross-sectional slices of the temperature field between a lower heated plate and cold upper wall (known as Rayleigh-Bénard turbulence). The right figure shows flow at a higher Reynolds number than the left figure; both figures show the same cross-sectional area. At higher Reynolds number (given the same fluid and distance between walls), much smaller scales of motion are generated that more rapidly mix the flow compared to more viscous flow. Lower Reynolds number flows consequently have relatively “coarse” small-scale structure.

eddies are (statistically) independent of the large-scale flow, as long as this scale separation is maintained. The size of the smallest eddies should then depend only on the viscosity ν and the rate at which they are supplied energy, ε , from the large-scale motion. From these two parameters, we can form the following length, velocity and time scales for the small (Kolmogorov-scale) eddies: $\eta \equiv (\nu^3/\varepsilon)^{1/4}$, $u_\eta \equiv (\varepsilon\nu)^{1/4}$, $\tau_\eta \equiv (\nu/\varepsilon)^{1/2}$.

The range of scales in turbulence, as given by the size ratio between the energy-containing and small-scale turbulent eddies, is then

$$\frac{\ell_0}{\eta} = \left(\frac{\varepsilon\ell_0^4}{\nu^3} \right)^{1/4} = \left(\frac{u_0^3\ell_0^3}{\nu^3} \right)^{1/4} = \text{R}_\ell^{3/4}, \quad (1)$$

where $\text{R}_\ell = u_0\ell_0/\nu$ is the Reynolds number based on the turbulent fluctuating velocity and length scale of the energy-containing turbulent eddies (for reference, $\tau_0/\tau_\eta = \text{R}_\ell^{3/4}$ and $u_0/u_\eta = \text{R}_\ell^{1/4}$). From this relationship, we can estimate the largest and smallest length scales for different flow configurations as shown in table 2. A slow saunter generates eddies as large as one’s girth and smaller than a millimeter while the higher Reynolds number flow past aircraft wings generate eddies as small as a micron. When $\text{R}_\ell \approx 1$, the energy-containing and small-scales are about the same size. Viscosity is then able to act directly on the energy-containing eddies causing the turbulence to dissipate and decay (e.g. when airplane engines are turned off).




Figure 6: Isovolume rendering of thermal updrafts in air. The Reynolds number corresponds to the same flow conditions given in the text except the radiator is now a constant temperature floor heater. When the thermal plumes rise and come close to the cold wall, they spread out laterally creating flattened but corrugated plume tops. The corrugations arise because very fine-scale (cold) thermals are trying to break through the hot plume.

As the Reynolds number increases, the separation in scales widens (see equation 1). Given two turbulent flows in the same device (i.e. with the same large-eddy size ℓ), the higher Reynolds number flow has a much finer scale structure than the other. Figure 5 shows two such flows at different Reynolds numbers (with the same ν and ℓ). The small-scale turbulence is visualized by temperature fluctuations (white is hot, black is cold and gray is in between). The abundance of small-scale detail and enhanced mixing shows that the flow on the right is at a much higher Reynolds number (in fact, about 1000 times higher based on the typical speed for a rising hot thermal plume). While this example is from a simulation, you may have seen such small scale structure in a roaring campfire (where unburned carbon particles give off a yellow-orange glow) or along the trembling, jittery horizon on a sweltering summer day (made visible by index-of-refraction gradients from local temperature variations).

While the large-eddies are largely the same size in the two flows in figure 5, the velocity is much higher and the time scales much faster in the flow on the right. The higher Reynolds number flow on the right of figure 5 therefore transfers heat much more rapidly between the hot lower wall and cold upper wall by turbulent mixing but also has a higher dissipation rate. For example, a radiator standing 10 cm high that is 10 °K hotter than the ambient air in a room 5 m on a side would generate turbulent eddies as large as the room and as small as 1 cm. A simulation of such a room at these conditions would generate thermal updrafts

like those shown in figure 6. These turbulent eddies, patches of zig-zagging and swirling fluid moving randomly about, are able to exchange energy and heat the room in just a few minutes. On the other hand, if the air were still and not turbulent, more than 100 hours would pass before the room would warm up [3]. Turbulent flows, however, are always dissipative and, therefore, require a continuous supply of energy, e.g. in the form of a mean shear or buoyancy, to make up for viscous losses. The turbulence energy absorbed due to mean shear or buoyancy cascades farther downscale to eddy sizes able to dissipate this energy. At high Reynolds numbers, almost all of the turbulent energy and shear is accounted for in eddies much larger than these small dissipative scales. Thus, the momentum and energy in the large-eddies can be balanced accurately at high Reynolds number with only an approximate modeling of the small dissipating eddies. An important counter-example is the very small-scale motions within an extremely thin layer adjacent to a wall, which is an important part of the energy balance and momentum transport. Simulating all scales of turbulence can be extremely demanding on computer memory and speed. But given sufficient computer power, numerical simulation of turbulent flows from essentially first principles are possible by directly resolving the large-eddies and modeling the more universal small-scale eddies.

3 Numerical simulation of turbulent flows

Experiments, theory and numerical simulations of turbulence play important roles in engineering design and scientific discovery. They provide complementary information, and analysis of complex turbulent flows requires one to use all tools at his or her disposal. The evolution of simulations as a productive tool for turbulence research and development has, however, had a significant impact upon how other tools are used. The primary scientific tools for turbulence simulation are direct numerical simulation (DNS) and large-eddy simulation (LES). DNS provides data that can be obtained in no other way enabling scientific insight into turbulence physics, modeling and control, but DNS as a general engineering tool for high Reynolds number flows is rather limited. Large-eddy simulation (LES), on the other hand, is intended to be useful for the study of turbulence physics at high Reynolds numbers and for predicting flows of technical interest. This is achieved by directly computing the details of the largest scales of motion (i.e. those responsible for the primary transport) and using a simple model for the smaller scales. As table 2 shows, turbulence involves a wide range of scales; the task of the simulator is to resolve and/or model these scales while accounting for the continuum scale physics on a discrete numerical grid with finite resolution.

3.1 Direct simulation

A direct numerical simulation (DNS) of turbulence dispenses with turbulence models and solves the three-dimensional, time-dependent Navier-Stokes equations for specified initial and boundary conditions [24]. Direct simulations of turbulent flows must resolve both the smallest and largest dynamically significant turbulent eddies and encompass the entire flow domain. This is a very challenging requirement for high Reynolds number flows as table 2

and figure 5 illustrate. Since all turbulent length and time scales are resolved – from the largest down to the smallest scales, which decrease rapidly with increasing Reynolds number – DNS is limited to relatively low Reynolds number flows by the available computational resolution we can afford with today’s most capable supercomputers [2].

We can estimate the cost for DNS in the simplest possible case in which the turbulence is statistically homogeneous and isotropic, depending on neither position nor orientation. The simulation must capture the viscous dissipation, which peaks around ten times the Kolmogorov scale η falling off rapidly for smaller scales, and comfortably resolve the largest eddies of size ℓ . This requires that we select a suitable mesh size $\Delta x \sim \eta$ and box size $\mathcal{L} \gtrsim \ell$. The box size \mathcal{L} must be large enough to represent the energy-containing motions and the grid spacing Δx must be small enough to adequately resolve the dissipative scales with the chosen numerical discretization (see §3.3 for further discussion). We may need even more grid points as the domain size may extend much further than the largest-eddy size. If we use periodic boundary conditions, our turbulence is fully correlated on opposite sides of the domain; therefore, the computational domain size \mathcal{L} must be sufficiently larger than ℓ to assure any error due to the finite period is small (by making sure two-point correlations decay to negligible magnitude within the period) and to contain enough large eddies to sufficiently sample all possible fluid motions for the purposes of computing turbulence statistics.

To the correct order-of-magnitude, the range of spatial scales we must resolve in each direction is then proportional to $\ell/\eta \sim \text{R}_\ell^{3/4}$ from equation 1 so that, in three dimensions, we expect the total number of grid points $N^3 \sim (\text{R}_\ell^{3/4})^3 = \text{R}_\ell^{9/4}$. For time-accurate simulations, we must also resolve the low-frequency, undulating large-eddies and the high-frequency, vorticity-bearing small-scale eddies whose time scale separation grows as $\tau_\ell/\tau_\eta \sim \text{R}_\ell^{1/2}$. For time-accurate *and stable* simulations, however, most commonly employed numerical algorithms require a slightly more restrictive timestep (or CFL) condition such that the total number of timesteps $M \sim \text{R}_\ell^{3/4}$. To a first approximation, the total computational work (e.g. floating-point and memory access operations) required is then proportional to $N^3 M \sim \text{R}_\ell^3$, i.e. the computational cost increases as the cube of the Reynolds number. As a result, there is always a limit to the maximum attainable Reynolds number.

Since the computational work scales as the cube of Reynolds number, the next $1000\times$ jump in computational concurrency (which the petascale to exascale roadmap promises) will only provide at most a factor of $10\times$ gain in Reynolds number by direct simulation. Since $\ell/\eta \sim \text{R}_\ell^{3/4} \sim (N^3 M)^{1/4}$, this means that $10^4\times$ the computational concurrency is needed for a $10\times$ gain in the resolvable range of scales ℓ/η in the same wall-clock time, i.e. each extra decade of resolution requires $10^4\times$ the concurrency. By direct simulation where we resolve all such scales, this means that each successive example in table 2 requires $10^4\times$ more capable resources. With today’s most powerful computers, we can directly resolve $\ell/\eta \sim \mathcal{O}(10^3)$ range of scales. This has enabled the direct simulation of several low-Reynolds number engineering flows with simplified configurations such as low-pressure turbine passages in jet engines (see figure 7) [25, 12, 7, 6]. If $10^4\times$ as much computational power is available, one may resolve additional scales (e.g. $10 \times \ell/\eta$) or add more physics (e.g. transport $\mathcal{O}(100)$ reacting chemical species with $2 \times \ell/\eta$).




Figure 7: Instantaneous contours of velocity magnitude from a three-dimensional direct numerical simulation of unsteady turbine cascade flow with upstream passing wakes [25]. Upstream, the flow velocity $U_{\text{ref}} = 1.0$ makes an angle $\alpha = 37.7^\circ$ with the x -axis and the Reynolds number $R = U_{\text{ref}}L/\nu = 1.48 \times 10^5$. The low-pressure turbine stage installed in jet engines occur upstream of the propelling nozzle, supplying power to the fan and first compressor stages. The Reynolds numbers are in a range where DNS is practical (but, out of reach for routine design calculations [7]). The flow is turned by more than 100° and accelerated by a factor of ~ 2 producing significant streamwise straining. The flow features revealed by 3D, unsteady numerical simulations are not available experimentally and may lead to new insights with important consequences for turbomachinery flows [25].

DNS has advanced in both directions towards multi-scale, multi-physics problems. DNS is used to study turbulence physics, guide the development of turbulence models, assess turbulence closure theories, investigate ways to control turbulence, and predict flows of technical interest at low Reynolds number (and Prandtl/Schmidt numbers not far from unity if heat/scalar transport are of interest) [24]. But, as new computational platforms become available, somewhat higher Reynolds numbers (with more complex physics and improved statistical sampling) are accessible to DNS as table 3 shows. We now examine how this trend toward increasingly multi-scale, multi-physics problems persists in the historical record of landmark direct numerical simulations (during the transition from 1 megaflop \rightarrow 10 gigaflop \rightarrow 100 teraflop \rightarrow 1 exaflop).

The first decade of (marginally) resolved scales from isotropic turbulence were computed on a CDC 6600 (~ 1 megaflop) at the National Center for Atmospheric Research c. 1970

Year	Authors	Simulation Type	# of Grid Points	Data Size
1972	Orszag & Patterson [26]	Isotropic Turbulence	32768	1 MB*
1981	Rogallo [27]	Homogeneous Turbulence	2×10^6	40 MB*
1987	Rogers & Moin [28]	Homogeneous Turbulence	2×10^6	80 MB
1987	Kim, Moin & Moser [29]	Plane Channel Flow	4×10^6	150 MB
1988	Spalart [30]	Turbulent Boundary Layer	11×10^6	422 MB
1991	Jiménez, Wray, Saffman & Rogallo[31]	Homogeneous Turbulence	134×10^6	1.3 GB*
1992	Lee, Lele & Moin [32]	Isotropic Shock-Turbulence	0.8×10^6	32 MB
1994	Le & Moin [33]	Backward-Facing Step	10×10^6	360 MB
1997	Freund, Lele & Moin [34]	Supersonic Mixing Layer	13×10^6	480 MB
2000	Freund, Lele & Moin [34]	Supersonic Jet and Noise	22×10^6	844 MB
2003	Earth Simulator [35]	Isotropic Turbulence	69×10^9	1310 GB*
2006	Hoyas & Jiménez [36]	Plane Channel Flow	18×10^9	683 GB
2008	Wu & Moin [37]	Turbulent Pipe Flow	630×10^6	24 GB
2009	Larsson & Lele [38, 39]	Isotropic Shock-Turbulence	153×10^6	5.7 GB
2009	Wu & Moin [23, 40]	Turbulent Boundary Layer	210×10^6	7.8 GB

Table 3: Selected landmark direct numerical simulations over several decades. Estimated memory requirements (assuming five storage variables in double-precision at each grid point) for each simulation are given (some authors were more constrained by memory and instead used specific memory-aware algorithms, e.g. Rogallo [27] and Kim, Moin & Moser [29]; in these cases, the true value may be overestimated. † These computations used zonal grids; all others used rectangular, Cartesian grids (except those by Freund, Lele & Moin [34], which used a cylindrical coordinate system). * These computations are known to have used single precision arithmetic. For example, Yokokawa, Itakura, Uno, Ishihara and Kaneda only saved fields in single-precision to reduce I/O time and save space on the Earth Simulator (many Terabytes of data were collected). Data size number shown reflects this.

[26, 41]. Within the next ten years c. 1976–1981, Rogallo [27] made substantial algorithmic improvements to conserve every last word of memory on the 64-processor ILLIAC-IV machine at NASA-Ames and extend the simulations of homogeneous isotropic turbulence to just over one decade of well-resolved scales. Ten years later c. 1991, Jiménez, Wray, Saffman and Rogallo [31] computed more than two well-resolved decades of turbulence scales using the Intel Touchstone Delta prototype (a ~ 10 gigaflop, 512 processor architectural successor to the 64-processor parallel ILLIAC-IV platform) installed at Caltech. On the Earth Simulator (~ 40 teraflops) c. 2002, three decades of resolved scales in homogeneous isotropic turbulence were computed. Pushing DNS to higher Reynolds number is limited by the available computational resources. At very high Reynolds numbers, all dynamically active turbulence scales cannot be simulated at the same time: some must be discarded or modeled as in large-eddy simulation, see §3.2.

Such limitations are not unique to the direct numerical simulation of turbulence. Wind tunnel operating costs also grow as the cube of the Reynolds number (for the same reason a ship traveling doubly fast consumes eight times the power). Such energy considerations




Figure 8: DNS of water flowing over a backward-facing step with varying air injection rates [42]. The effect of air injection rate on stability of the air layer was investigated. Air-layer drag reduction was confirmed by direct simulation. At high air injection flow rates ($q_{\text{air}} = 0.00357 \text{ m}^2/\text{s}$ per unit span), a stable air layer forms that significantly reduces the drag; compared to the case without air-injection, the total drag is reduced by more than 99%. At lower injection rates, the air-layer begins to break-up showing intermittent spots of water reaching the wall that increase the drag. The air-water interface (zero level set function) is shown. Flow is spanwise periodic. (a) Bottom view of air-water interface for $q_{\text{air}} = 0.00357 \text{ m}^2/\text{s}$. Detailed view of air-water interface for $q_{\text{air}} = 0.00357 \text{ m}^2/\text{s}$ at two different downstream positions: (b) $x/X_r = 0.5$; (c) $x/X_r = 0.75$. (d) Time-averaged skin-friction coefficient for three different air injection rates. Reynolds number based on step height is $Re_h = 22,860$ (air/water: $\rho_g/\rho_l \approx 832$, $We_h = 586$); experimental data from Jovic & Driver (1995) for single-phase backward-facing step flow at different Reynolds numbers shown. Skin-friction coefficient measurements from two-phase flow experiments were not available for validation, but the phase interface structure and transition from unstable to stable air-layers are comparable. Simulations used 271 million grid points; unstable cases run on Intrepid at the ALCF; stable cases run on the LLNL BGL machine [42].

impose significant restrictions on testing time and Reynolds number in large wind tunnels [2]. This large, fixed operational cost dedicated to leadership-class, national user facility wind tunnels ultimately led to widespread decommissioning as priorities changed, budgets decreased and funding for electricity bills diminished. Power consumption and facility costs of today's leadership-class computing facilities are comparable and expected to grow as more powerful computers become available. Supercomputers are, however, general purpose scientific instruments able to serve a broad and diverse base of users and applications. The fact that such facilities are more than just numerical wind tunnels suggests that a more substantial power ceiling and investment may be sustainable.

Many of the canonical turbulent flows first studied by laboratory experiment are now reliably simulated by computer (turbulent channels, turbulent boundary layers, mixing lay-

ers, subsonic and supersonic jets, backward-facing steps, etc.) [43, 44]. Figure 8 shows data from direct numerical simulations of two-phase flow past a backward-facing step. This configuration is designed to investigate the effects of air-layer drag reduction for marine vehicles (skin-friction drag accounts for over 60% of the total drag of a typical cargo ship). Water flows on a flat plate over a backward-facing step where air is injected through a slot at the step. Experiments have demonstrated drag reduction by means of injection of air between the water and solid wall. In the simulations, the air-layer with high injection rates were stabilized across the entire domain whereas lower air injection rates led to instability of the air-layer. As the air-layer breaks up, water is able to penetrate toward the wall. These water spots on the solid wall led to much larger skin-friction drag. These simulations provided important data in this transitional regime concerning air-layer stability and air-layer drag reduction [42].

Many physical problems, such as turbulent combustion, jet atomization, particle/polymer-laden turbulent flows, require one to address the coupling of turbulent flow with the molecular structure of matter. In such flows, the dynamically active scales may extend beyond that of single-phase, non-reacting turbulence, e.g. one may need a mesh fine enough to resolve the thin inner structure of flame fronts. Direct resolution of even smaller scales may be intractable. Instead, detailed numerical simulations are carried out where the Navier-Stokes equations are solved with additional constitutive or mesoscopic closure models to account for detailed chemical kinetics and transport as in turbulent combustion [45, 46]; gas-liquid interface dynamics or the break-up and atomization of polydisperse droplets (with point-particle, ensemble-averaged or mesoscopic kinetic models to account for unresolved particle wakes, boundary layers or drag/heat transfer) [18, 43, 47]; and the configurational response of entangled polymers in turbulent shear flow [48, 49]. Detailed numerical simulations coupled to advanced chemical mechanisms provide a framework to study many important turbulence-chemistry interactions in canonical and laboratory-scale turbulent flames.

Supersonic flow in scramjets, for example, requires that air and fuel mix on a molecular level and react chemically before leaving the engine. Due to the very short residence times in such high speed flows, combustion ignition, instability, extinction and other non-equilibrium phenomena all hinge upon accurate modeling of turbulence-chemistry interactions amongst many species (i.e. reactants and products) [45]. Detailed numerical simulation with advanced finite-rate chemistry models are therefore needed to resolve several outstanding scientific and modeling questions for such highly non-equilibrium flows [11, 50]. Fundamental studies of compressible reacting flows have relied upon simplified/reduced chemical mechanisms and limited closures for the chemical source term in the species transport equations. The computational cost and stiffness of the governing equations were too great for more realistic chemical mechanisms, and only in recent years has reliable DNS with finite-rate chemistry for reacting flows become possible [45]. A recent example of supersonic combustion in a temporal mixing layer with a detailed, finite-rate chemical mechanism for hydrogen/oxygen combustion (9 species and 29 reactions) is shown in figures 9 and 10 [50]. This simulation database was later used for validation of a reduced-order chemical mechanism, known as the flamelet and progress-variable approach (FPVA), in supersonic reacting flows. Figure 11




Figure 9: Instantaneous OH mass fraction (top) and vorticity magnitude (bottom) in a plane normal to span.




Figure 10: Supersonic temporal mixing layer Q-criterion isosurfaces. Density contours in background; centerline strain-rate projected below.

shows results from this study, which led to the development of improved FPVA models that account for the low-dimensional manifolds discovered in the supersonic reacting mixing layer by detailed numerical simulation [50]. We anticipate future computing platforms will enable further detailed assessments of turbulence-chemistry interaction models for more complex fuels in high Reynolds number reacting flows.

The general trend for simulations (both DNS and LES) is that with greater availability of computing resources more complex physics problems are solved that accumulate much more data. The number of grid points (a measure of the memory used and the time spent analyzing data) has grown, on average, by about three orders of magnitude in the previous three decades. This does not account for the integration time (which is a measure of statistical convergence and is a measure of how much full field data is saved). The total amount of field data saved has grown in proportion. A billion grid point 1024^3 DNS will need to save at least three variables per grid point (i.e. the three velocity components) and perhaps two scalars; at eight bytes each (i.e. double precision) this totals to $1024^3 \times 5 \approx 40$ Gbytes *per field*. Each field is only a single instantaneous snapshot of the flow. Since each simulation involves many fields, the volume of data is enormous.

Given this increase in data to analyze, the time taken for data analysis and resources used for data storage continues to increase. Often a single desktop workstation does not have enough memory to process the entire flow field data; personal mini-supercomputers with large memory nodes are particularly helpful for post-processing turbulence data. A typical DNS database may contain three velocity components and perhaps pressure. To compute derived quantities, such as vorticity, from this data one must perform numerical operations on the database in a manner consistent with the original numerical method that generated the data. Non-linear products may require dealiasing and statistics involving derivatives may not be




Figure 11: Temperature contours shown from a detailed numerical simulation (left) and simulation (right) using a reduced chemical mechanism (a flamelet and progress-variable approach or FPVA) are shown. The DNS provided valuable validation data for the FPVA simulations. The DNS showed that, in reacting mixing layers, there exist intrinsic low-dimensional manifolds in the supersonic regime for several chemical species (e.g. OH shown, middle). FPVA assumes the existence of such manifolds. This DNS validation database helped extend the FPVA model to the supersonic regime.

accurately computed at all from a given DNS database. Many simulators are reluctant to release their DNS data to others for fear of improper post-processing, but the potential for extracting additional information from old simulations is often so great that much of the scientist-programmer’s time may be sapped doing service calculations for other scientists. As a result, most DNS database archives are simply a few text files of averaged statistics or sparsely documented binary data dumps. Support for the database archival process and making data and computational resources for post-processing more widely available is just as important as the simulation itself.

Simulations of this type may take one group of people one year to compute (the recent turbulent boundary layer simulations [36, 23, 20, 21] required about one year to reach statistically converged solutions on readily available computing systems). Post-processing of the data, however, may take twice as long as the simulations themselves (as there are many avenues of scientific inquiry to explore). Extra simulations and statistics then follow. As a result, the “lifetime” of simulation where there are many groups actively exploring the data may last as long as five to ten years. Archiving and making the data accessible to the wider community of scientific research is a high priority during this time. This is achieved by making the data easily accessible (e.g. with parallel/portable storage systems and dedicated large-memory computing hardware), properly documenting the data/file format and any limitations, and ideally providing software tools, languages or interfaces to derive quantities of interest from the database. One unique approach is to make the data accessible via web-enabled database queries (the JHU turbulence database being one example: <http://turbulence.pha.jhu.edu/>). The Summer Program at the Center for Turbulence Research also provides an example of how to accelerate scientific data discovery and communicate such results to society by bringing the scientist, simulator and numerical database together in one place with ample computing resources for post-processing and collaboration (see <http://www.stanford.edu/group/ctr/SummerProgram/> for more details). Figure 12 shows



Figure 12: Front covers from the biannual Center for Turbulence Research Summer Program Proceedings highlight the historical progress of turbulence simulation (computational capability increases from tens of megaflops \rightarrow tens of gigaflops \rightarrow terascale era \rightarrow petascale era). From left to right: 1988, 2D/3D simulations of spatial/temporal mixing layers investigated by several authors; 1992, “The structure of intense vorticity in homogeneous isotropic turbulence” by J. Jiménez, A.A. Wray, P.G. Saffman and R.S. Rogallo [31] directly resolved two full decades of turbulence using 512 cores on the Intel Touchstone Delta machine; 1998, several participants used laptop personal computers to carry out some of the work on several different DOE/NASA/DOD/home parallel computing resources (showing an increased availability of computing power near the turn of the century) and turbulent combustion [11] was the largest group in the program; 2004, solar interiors, radiative transfer from large pool fires and other complex multiphysics problems were simulated by a diverse group of researchers from solar physics, magneto-hydrodynamics, combustion, acoustics, multiphase flow, LES and RANS communities demonstrating the ubiquity of turbulence in cross-disciplinary research; 2006, fundamental modeling issues for multiphase flow [42] and numerical methods for predictive science were explored including the results from a six-day dedicated simulation of rotating, sheared turbulence on 65,536 cores of the LLNL BG/L (the fastest supercomputer at the time), unfortunately, due to a sign mistake in an input parameter, the results did not have the expected impact; 2010, several high-fidelity DNS/LES numerical simulations and uncertainty quantification of multi-physics turbulent flows in complex domains were addressed in the proceedings such as the front cover image from “DNS analysis of a $Re=40,000$ swirl burner” by V. Moureau, P. Domingo, L. Vervisch and D. Veynante [46].

selected front covers of the CTR Summer Program Proceedings to emphasize the multi-scale, multi-physics and complex geometry evolution of turbulence simulation in the past quarter century. Simulation databases are a powerful tool in turbulence research; databases have proven to be effective catalysts for interaction among researchers.

3.1.1 Governing equations and discrete conservation principles

For incompressible flow with constant viscosity, the integral statements for mass, momentum energy and entropy (i.e. the Navier-Stokes equations) reduce to the following set of four

equations for the three velocity components u_i ($i = 1, 2, 3$) and pressure p ,

$$\frac{\partial u_i}{\partial x_i} = 0 , \quad (2a)$$

$$\frac{\partial u_i}{\partial t} + \frac{\partial u_i u_j}{\partial x_j} = -\frac{1}{\rho} \frac{\partial p}{\partial x_i} + \nu \frac{\partial^2 u_i}{\partial x_j \partial x_j} . \quad (2b)$$

Direct numerical simulations seek to compute the evolution of all significant scales of motion from these differential equations without any turbulence models. In their differential form, these equations conserve mass, momentum and kinetic energy (in the inviscid limit). For example, if we multiply both sides of equation 2b by u_i , we have

$$\frac{dK}{dt} + \frac{\partial T_i}{\partial x_i} = -2\nu S_{ij} S_{ij} , \quad (3)$$

where $K \equiv \frac{1}{2} u_i u_i$ is the kinetic energy, $T_i \equiv u_i p / \rho - 2\nu u_j S_{ij}$ is the flux of energy, and $S_{ij} = \frac{1}{2} (\partial_j u_i + \partial_i u_j)$ is the rate-of-strain tensor. The second term is responsible for the transport of kinetic energy from one region to another. In the inviscid limit ($\nu = 0$), the right-hand side is identically zero and integration over a fixed control volume gives $\frac{d}{dt} \int_V K dV + \int_S T_j n_j dS = 0$, which is a statement of local conservation, i.e. the rate of change of K inside the control volume is balanced by only the net flux through the boundary and includes no sources or sinks. Such terms are in conservative form, i.e. they do not produce or consume kinetic energy but just move it around. The (non-conservative and non-positive) term on the right of equation 3 represents viscous dissipation and is responsible for the conversion of mechanical energy into heat (note that $\varepsilon \equiv 2\nu \bar{s}_{ij} \bar{s}_{ij}$ where $s_{ij} = S_{ij} - \bar{S}_{ij}$).

Credible numerical simulations must be careful to not introduce any additional non-physical, numerical dissipation, which can have a detrimental effect on the resolution of turbulence structures. Numerical algorithms must enforce the kinetic energy conservation principle by ensuring no spurious contributions to the volume integral of kinetic energy are made at the discrete level. Development of such non-dissipative numerical methods is paramount and recent advances are discussed in §3.3; use of such numerical methods for DNS also applies to large-eddy simulation, which is discussed in the following section.

3.2 Large-eddy simulation

If one wants to simulate complex turbulent flows at high Reynolds numbers, there is at present only one way: large-eddy simulation (LES). The basic idea is that one computes the large-scale turbulence directly and models the small-scales. The motivation for this is that large eddies extract energy from the mean flow, are highly anisotropic, vary from flow to flow and transport the bulk of the turbulent momentum and energy whereas the small eddies dissipate this energy (carrying a trivial fraction of the total turbulent energy) and are relatively isotropic or “universal” (meaning easily modelled) in most circumstances (a major exception being the flow very near walls). We like to separate the formulation of LES from the numerical method used for its solution [17, 14, 51]; therefore, the large-eddy field




Figure 13: A sinusoidal function with superimposed random noise (20% amplitude) on an unstructured grid with local grid refinement: unfiltered (left) and filtered (right) [51].

and small-scale turbulence model are formally defined without reference to the solution grid. The large-eddy field is called the *resolved* field and the small-eddy field the *residual* field (which is often called the *sub-grid* field if one is not so careful as to make this distinction). To formally define the resolved field, we require a precise definition of a (low-pass) filtering operation to separate the residual turbulence field from the instantaneous velocity field and compute the associated non-linear convective term. For this purpose, we define the filtered field \bar{f} by means of an integral normalized filter

$$\bar{f}(\mathbf{x}, t) = \int_{\Omega} G(\mathbf{x}, \mathbf{x}'; \Delta_f) f(\mathbf{x}') \, d\mathbf{x}' , \quad (4)$$

where the G is the filter kernel and Δ_f is a characteristic filter width. The instantaneous velocity is then decomposed into filtered and residual fields by $u_i = \bar{u}_i + u'_i$; for example, figure 13 compares the instantaneous and resolved field obtained from filtering on a complex, unstructured grid. We note that one can also filter in time as well as space with a finite residual time scale analogous to the finite filter width (where, in appropriate limits, one can recover both the DNS and RANS equations). Filtering in LES is different than conventional averaging used in RANS; in general, $\overline{f'} \neq 0$ and $\overline{\bar{f}} \neq \bar{f}$ since a second smoothing removes additional structure from the resolved field.

To derive tractable equations for numerical solution, it is important that the filter commute with differentiation, i.e. $\overline{\partial_x f} = \partial_x \bar{f}$. Furthermore, if the filter commutes with differentiation (which requires some additional work on unstructured grids [51]), we may then formally decouple the filter and grid scales to achieve grid-independent LES (with refinement) [51, 52]. If we apply a filtering operation that commutes with differentiation as defined by equation 4 to the unfiltered Návier-Stokes and continuity equations, the filtered equations

for incompressible flow with constant viscosity assume the following form

$$\frac{\partial \bar{u}_i}{\partial x_i} = 0, \quad (5a)$$

$$\frac{\partial \bar{u}_i}{\partial t} + \frac{\partial \overline{\bar{u}_i \bar{u}_j}}{\partial x_j} = -\frac{1}{\rho} \frac{\partial \bar{p}}{\partial x_i} - \frac{\partial \tau_{ij}}{\partial x_j} + \nu \frac{\partial^2 \bar{u}_i}{\partial x_j \partial x_j}. \quad (5b)$$

These equations hold only for filters that commute with differentiation. The effect of the residual field appears as

$$\tau_{ij} = \overline{\bar{u}_i u'_j} + \overline{\bar{u}_j u'_i} + \overline{u'_i u'_j}, \quad (6)$$

which must be modeled. Given a model for τ_{ij} the system of equations is closed and can be solved numerically with appropriate methods (see §3.3). Selecting an appropriate, physics-based residual turbulence model is of course a key research question for LES [14]. We wish to stress that the trivial closure $\tau_{ij} = 0$ where numerical dissipation (e.g. upwind differencing) is used to implicitly set the smallest scale, i.e., (implicit) LES with no explicit residual turbulence models, has been demonstrated to be dependent on numerical parameters, and should be avoided. In a proper explicitly-filtered LES, the numerical truncation error should not exceed the contribution of the residual turbulence model (see discussion in §3.3).

The most commonly used residual scale model for LES is an eddy viscosity model, e.g. where $\tau_{ij}^S = q_R^2 \delta_{ij} / 3 - 2\nu_R S_{ij}$ where q_R^2 is the energy of the resolved turbulence, $\bar{S}_{ij} = \frac{1}{2}(\partial_j \bar{u}_i + \partial_i \bar{u}_j)$ is the strain rate of the resolved field, and ν_R is the effective viscosity of the residual field. Increased transport and dissipation comparable to that Historically [53, 54], most practitioners have used the constant-coefficient Smagorinsky model, which assumes that ν_R is proportional to the product of a turbulent timescale formed from the resolved deformation rate (e.g. $\tau_S = \sqrt{2\bar{S}_{ij}\bar{S}_{ji}}$) and the square of an appropriate length scale (e.g. $\ell_S = C_S \Delta_f$) made proportional to the filter width Δ_f by an ad-hoc proportionality constant (e.g. $C_S \approx 0.23$) that would vary from flow to flow. This may be appropriate when the residual turbulence is always and everywhere in equilibrium with the resolved field, but this is certainly not the case in flows undergoing transition to turbulence, relaminarization, or other rapidly adjusting flows.

For most turbulent flows, a dynamic residual turbulence model is the preferred choice. The key difference of the dynamic procedure applied to the Smagorinsky model [17, 55] is that the eddy viscosity coefficient, C_S , is allowed to vary in both space and time and is dynamically computed using information contained in the resolved turbulence scales using a mathematical identity. This effectively eliminates uncertainties associated with tunable model parameters such as C_S . The dynamic procedure is particularly successful in addressing the limitations of the Smagorinsky model in flows undergoing transition to turbulence, relaminarization and in the viscous near-wall region (e.g. the dynamical model produces the correct limiting behavior of vanishing eddy viscosity near walls while a constant-coefficient Smagorinsky model artificially generates non-zero Reynolds stress in the viscous near-wall region). In the application of LES to chemical reacting flows such as turbulent combustion, a model that allows residual turbulence models to dynamically adjust to chemical reactions occurring on residual time scales faster than the time scale of the resolved eddies is important [11].




Figure 14: Instantaneous spanwise vorticity contours of flow over a NACA 0015 airfoil with synthetic jet actuation (left) without control (middle) and with control (right) [56].




Figure 15: Validation of simulation (lines) against experimental data (symbols) without control (blue) and with control (red) with airfoil surface pressure coefficient (left) and velocity wake at $x/C = 1.2$ (right) [56].

LES provides a tractable method for the simulation of turbulent flows at high Reynolds numbers in complex geometries. For turbulent free-shear flows such as supersonic turbulent jets, the computational cost of LES scales linearly with Reynolds number. With an appropriate near wall-model [57, 12], wall-bounded turbulent shear flows, such as the attached flow on the suction side of an airplane wing, also scale linearly and are not limited to moderate Reynolds number flows as for direct numerical simulation (compare with §3.1 and equation 1) [58, 2]. As a predictive engineering science tool, the cost effectiveness of LES for high-Reynolds number flow on massively parallel computers allows for more complex physics, control and optimization strategies.

For example, You & Moin performed large-eddy simulation of turbulent flow separation over an airfoil to evaluate the effectiveness of synthetic jets as a separation control technique at a chord Reynolds number of 896,000 [56]. A small slot across the entire span connected to a cavity inside the airfoil is used to generate oscillatory synthetic jets. Figure 14 shows spanwise vorticity contours for flow over an NACA 0015 airfoil without and with control produced by the zero-net-mass-flux synthetic jet actuator shown. Detailed flow structures inside the synthetic jet actuator and the synthetic jet cross-flow interaction were simulated




Figure 16: Left: contours of instantaneous streamwise velocity for original airfoil design shape and the best feasible shape after three iterations, which gives 89% noise reduction. Maximum velocity $1.29U_\infty$, minimum $-0.32U_\infty$, 15 contour levels. On the top left: the original shape (gray line) and best feasible shape (black line) for turbulent-flow optimization. Right: the normalized cost function vs. the number of LES evaluations for four iterations until convergence. Each LES computation used about 7 million mesh points.

using an unstructured-grid finite-volume large eddy simulation solver. As in the companion experiment, the large eddy simulations validation results shown in figure 15 confirm that synthetic-jet actuation effectively delays the onset of flow separation and generates a significant increase in the lift coefficient. In automotive applications, similar control strategies are used to reduce drag by modifying the large, separated wake structures behind vehicles [59, 60]. Active control and optimization are both important in the design of complex engineering systems [61, 9].

When designing an airfoil or aircraft wing as in figure 1, engineers often build several design prototypes to search for an optimal design with respect to some objective function or metric, e.g. to minimize drag and maximize lift, subject to certain constraints, e.g. airfoil thickness [9]. This iterative design process can be formalized, automated and carried out with parallel computers [62]. Figure 16 demonstrates a shape optimization algorithm applied to a time-dependent, three-dimensional turbulent flow over an airfoil with a rounded trailing-edge. The optimization procedure deformed both the upper and lower surfaces of the trailing-edge in order to minimize the cost function proportional to the total radiated noise over all frequencies computed by LES. Lift and drag constraints were applied by using a lower-fidelity RANS model to filter out undesirable solutions before computing the cost function with LES. The computational cost of the optimization was significantly reduced by separating the constraint and cost-function evaluations in this way. In the end, an optimal airfoil shape was designed that reduced noise levels by +85% with no decrease in lift or increase in drag after only 8 LES evaluations (and up to +89% after 18 LES runs). We observe that the large scale vortex shedding and wake thickness were also reduced significantly compared to the original leading to significant reductions in tonal and broadband noise. In §4.6, uncertainty




Figure 17: Actual chevron nozzle geometry (upper left inset) and computed far-field sound spectra statistics (right). Grid-convergence and validation with physical experiment shown [63, 64, 65]. Visualization (left) of sound propagated from the FWH surface (cyan) and chevron jet turbulence shown. Redscale temperature contours are displayed inside the FWH surface, grayscale contours show the perturbation pressure in the exterior. Simulations performed on Intrepid at ALCF. All 40 racks (163,480 cores) used for the 528M run.

quantification is applied to a similar problem of trailing-edge noise where LES is used to predict the unsteady turbulent fluctuations responsible for airfoil and fan blade noise.

Recently, LES was used to simulate the supersonic jet flow issuing from the complex geometry chevron nozzle shown in figure 17 and to predict the far-field noise. This nozzle is the exact geometry of a nozzle under experimental investigation at NASA Glenn Research Center. Comparison to experiment has provided new scientific insight [63, 64, 65, 66]; confidence in the predictive simulations is so strong that new experiments are underway to re-test the few angles where there is some disagreement [67]. For the first time, we are able to match the far-field acoustic spectra from such a complex nozzle at all angles including the broadband shock-associated noise (see figure 17). Note especially the shocks created by the small chevron tips and gaps that act as a source for the observed high-frequency broadband shock-associated noise and peak at downstream angles.

While LES is an efficient tool for resolving the turbulent, highly vortical flow of the jet containing acoustic sources, it is not, however, efficient at propagating the resulting acoustic waves to the far field (see §3.3). Fortunately, in the irrotational flow surrounding the jet, the acoustic propagation reduces to a linear problem for which an analytic Green’s function exists. Therefore, a Ffowcs Williams-Hawkins (FWH) surface integral method whereby acoustic information is collected along a surface surrounding the jet from the unsteady, 3D numerical simulation database and then efficiently projected to the far-field as shown in figure 17. Figure 18 shows how aggressive grid refinement and coarsening near the FWH interface with unstructured grids has helped reduce the cost and improve the accuracy of




Figure 18: Mesh design and refinement for aeroacoustics using traditional grid stretching (left) and the advanced isotropic grid refinement capabilities of `charles` (right). The 45M CV mesh (right) has isotropic cells in the jet interior where grid-stretching is a particular problem for the 86M CV mesh (left). From the far-field acoustic spectra extracted from both cases (not shown), we find that the 45M CV mesh outperforms the 86M CV mesh because of the reduced grid-stretching. The unstructured mesh adaptation and refinement technology of `charles` enables higher-fidelity calculations of supersonic jet noise at substantially reduced cost compared to traditional methods. The FWH surface is shown in red.

such hybrid simulation methods.

Grid refinement is often used to verify that the simulation results converge to a physical solution, which can then be used to validate the flow predictions. In practice, a sequence of solutions is obtained on successively finer meshes until a sufficient range of scales in the flow field has been resolved and the statistical quantities of interest are invariant with respect to the mesh. This process is often referred to as grid convergence. For implicitly filtered (or traditional) LES, the grid-converged solution is not the true solution of the LES equations. In the limit when the mesh size is sufficiently small to capture the smallest scales of motion, an implicitly filtered LES will converge to a direct numerical simulation because the filter width also approaches the size of the smallest eddy. This limit, of course, is unreachable in most practical situations where LES is applied. Ideally, the true LES solution should correspond to the filtered velocity field, given a particular closure model and a well-defined spatial filter. Since the solution is grid-dependent, convergence is inconsistent and any attempt to verify, validate and quantify the uncertainty of sub-grid models is severely limited.

Grid-independent solutions of the governing equations for LES are needed for a consistent verification/validation study. A truly grid-independent solution is one such that the numerical errors, both truncation and commutation errors, are sufficiently small that the turbulent statistics demonstrate that a grid-converged solution has been reached [51]. If the flow statistics of the grid-independent solution of the explicitly filtered LES equations are not in good agreement with filtered direct numerical simulation (DNS) statistics, this failure can be attributed solely to the capability of the subfilter stress model employed. Without this framework, evaluating the fidelity of the closure model is ambiguous owing to the observed sensitivity of the subgrid model to numerical errors.

Bose and Moin [51] have recently performed an explicitly filtered LES of flow through a three-dimensional stalled diffuser flow with the same operating conditions and geometry




Figure 19: Left: instantaneous streamwise velocity in the yz -plane in a rectangular duct with cross-section $3.33H \times H$; contour levels from 0 to $1.4U_b$ [51]. Right: mean streamwise velocity profiles at different spanwise locations for the explicitly filtered LES (solid lines) and experimental measurements of Kolade (2010) (symbols).




Figure 20: Instantaneous spanwise velocity from LES of the 3D diffuser of Kolade (2010) at $z/H \approx 1.67$. Contours from $-0.2U_b$ to $0.2U_b$ [51].

(including side-walls) from a recent experimental campaign at Stanford. The experimental setup is designed to provide a challenging test case for numerical models at realistic Reynolds numbers involving strong adverse pressure gradients, rapid boundary layer growth, flow separation and associated unsteadiness. This is a coarse (50 million CV), explicitly filtered LES of a moderately high Reynolds number rectangular duct flow (figure 19) issuing into a three-dimensional stalled diffuser (figure 20). Since the flow is highly three-dimensional and anisotropic, the grid refinement procedure is done by using the subgrid kinetic energy surrogate $k^* = \frac{1}{2}\overline{u'_k u'_k}$ as a threshold for local grid refinement in regions of strong turbulent activity. The use of this criteria for grid refinement with fixed filter width is justified because it is likely that the influence of numerical errors will be most severe where the true solution of the Navier-Stokes equation is least resolved with respect to the filter width [51].

3.3 Numerical Methods

Many physical phenomena possess a broad range of length and time scales: turbulent fluid flows are a common example (recall table 2 and figure 4). Direct simulation of these phenomena require that all relevant scales are properly represented in the numerical model. Furthermore, large-eddy simulation also requires that all resolved scales are properly represented in the numerical model so that the physical, subfilter-scale turbulence model is not

overshadowed by unphysical numerical truncation, aliasing, or other discretization errors. Since the maximum resolvable range of scales in each direction for LES is always limited by the available computational resources (of order $\mathcal{O}(10^3)$ on today's most capable computers, $\mathcal{O}(10^4)$ with exascale resources), the numerical grid may only be fine enough to resolve the important large-eddy structures in most high-Reynolds number industrial or geophysical applications [68, 69]. Predictive turbulence models and simulation science therefore demand that any influence due to unphysical numerical discretization errors is minimized.

When we represent a continuous function $f(x)$ on a finite numerical grid (e.g. $N + 1$ grid points at $x_j = jh$ for $j = 0, 1, \dots, N$ with uniform grid spacing h), we may introduce numerical errors when we attempt to differentiate, integrate, interpolate, iterate or otherwise operate on the truncated and discretely sampled function $f_j = f(x_j)$. When solving differential equations like the Navier-Stokes (equation 2b), we must compute derivatives and compute non-linear products on a finite grid, which can introduce unphysical numerical truncation and aliasing errors (accumulation of round-off errors, transient flipping of bits and related fault-intolerant hardware errors may also become more important on future computing systems). For example, the formal truncation error due to the discrete approximation of the first derivative of a function $f(x)$ is of the form

$$\frac{\delta f}{\delta x} = \frac{df}{dx} + \text{truncation errors} , \quad (7)$$

for which a fourth-order differencing scheme (derived by Taylor Series expansion) is

$$\left. \frac{\delta f}{\delta x} \right|_j = \frac{f_{j-2} - 8f_{j-1} + 8f_{j+1} - f_{j+2}}{12h} = \left. \frac{df}{dx} \right|_{x=jh} + \mathcal{O}(h^4) , \quad (8)$$

where the formal truncation error, in the asymptotic limit of grid refinement (and $f(x)$ being smooth enough), would decrease as the fourth power of the grid spacing, i.e. mesh refinement by a factor of 2 improves the accuracy of a fourth-order scheme by a factor of 16.

Because turbulence typically contains a broad range of eddies, the smallest of which change very rapidly, an appropriate differencing method would ideally not dissipate or otherwise distort these small-scale, high-wavenumber components. The formal truncation error described by equation 7 does not, however, provide any information as to how the chosen difference approximation represents the exact result over the entire range of length scales a given mesh (or filter) can resolve. The formal truncation error (e.g. fourth-order) provides only the local error integrated over all wavenumbers in the asymptotic limit of grid-refinement and, hence, does not distinguish between different scales of motion. Fourier (or modified wavenumber) analysis provides an effective way to quantify such resolution characteristics of a given differencing method for all resolvable scales.

To illustrate this, consider how different numerical differencing schemes differentiate a simple Fourier wave $f(x) = e^{ikx}$ on a domain of length $L = Nh = 2\pi$. The exact first derivative of f is $f' = ikf$ but, in fact, most numerical discretizations, and finite-difference methods in particular, will have some finite error at all wavenumbers for this simple function (note that Fourier transforms differentiate such functions exactly up to the grid resolution,




Figure 21: Left: modified wavenumbers for three non-dissipative finite-difference schemes (second-order central, fourth-order central and fourth-order compact Padé) [70, 71]. Symmetric FD schemes such as these have zero imaginary component (i.e. $k'_i = 0$ and $k' = k'_r i$) and no-built-in dissipation. For the semi-discrete (exact time advancement) linear advection equation with constant phase speed c (equation 10), the phase speed for a wave of wavenumber k given by the finite difference scheme is $c' = (k'(k)/k)c$. Middle: modified phase speed c'/c as a function of wavenumber k . For these difference schemes, the modified phase speed lags the actual phase speed (i.e. $c'/c < 1$ for $0 < k < \pi$) so that computed waves travel at a lower velocity than the physical ones. Discrete-time-advancement can, however, introduce dissipation in the solution of equation 10. Right: amplitude ratio $|G| = |u^{n+1}/u^n|$ as a function of resolution. The amplitude ratio of the exact solution of equation 10 is unity (i.e. no change from timestep n to $n + 1$). Note how the fully discretized equations now have a dissipative character although the spatial discretization is non-dissipative. The figure shows the amplitude error for the solution of equation 10 using a fourth-order compact (Padé) scheme for spatial differences while a fourth-order Runge-Kutta scheme is used for time advancement. Three different CFL numbers ($\text{CFL} = c\Delta t/h$) are shown.

i.e. $f'_j = ikf_j$ for all wavenumbers $k_n = (2\pi/L)n$ and $n = 0, 1, \dots, N/2$ giving $k_{\max} = \pi/h$. For example, one can show that the fourth-order central difference approximation given by equation 8 yields $f'_j = ik'f_j$ where

$$k' = \frac{1}{6h} [8 \sin(hk) - \sin(2hk)] , \quad (9)$$

is the modified wavenumber as shown in figure 21. The distance from the exact $k'(k) = k$ solution provides a measure of the relative dispersive error introduced by each difference scheme. Note that modified wavenumbers for central difference approximations on uniform grids (e.g. equation 8) have zero imaginary components; this means that such schemes are non-dissipative for all wavenumbers, which is an important property for turbulence simulations (since we don't want the small amplitude, high-frequency wiggles of the smallest eddies to be damped by unphysical numerical dissipation). The reason for this becomes clear when




Figure 22: Left: modified wavenumber (real-part, k'_r) for three dissipative finite-difference schemes (first-order, third-order and fifth-order upwind) [70]. These non-symmetric (upwind) schemes have finite imaginary component (i.e. $k'_i \neq 0$), which creates amplification error. The first/third-order upwind schemes have the same k_r (and therefore dispersive error) as the second/fourth-order central differences shown in figure 21; however, both upwind schemes introduce dissipation whereas the central schemes have zero amplitude error. Middle: amplitude error measured by $\exp(k'_i h)$ shown as a function of wavenumber k . Right: amplitude ratio $|G| = |u^{n+1}/u^n|$ as a function of resolution for the solution of equation 10 using a fifth-order upwind scheme for spatial differences and a fourth-order Runge-Kutta scheme for time advancement. Three different CFL numbers ($\text{CFL} = c\Delta t/h$) are shown.

we discretize the linear advection equation

$$\frac{\partial u}{\partial t} + c \frac{\partial u}{\partial x} = 0, \quad (10)$$

where c is a constant wave speed. A particular (exact) solution of equation 10 is $u(x, t) = e^{-ickt} e^{ikx} = u(x, 0) e^{-ickt}$ where $u(x, 0) = f(x)$ is the initial condition of the waveform. The semi-discrete approximation (where we discretize only the spatial derivative in equation 10) is then $u(x, t) = u(x, 0) e^{-ick't} = u(x, 0) e^{-ick'_r t} e^{ck'_i t}$ where we have split the modified wavenumber $k' = k'_r + ik'_i$ into real and imaginary components. In this form, we immediately see that the imaginary part of the modified wavenumber contributes to numerical dissipation while the real part contributes to numerical phase/dispersion errors. Imaginary parts of the modified wavenumber introduce damping in this otherwise purely oscillatory system. As figure 22 shows, upwind biasing introduces numerical dissipation (i.e. negative k'_i) at intermediate and higher wavenumbers.

The smallest scales (large k) are poorly represented by finite-difference schemes; for upwind schemes, the amplitude decays faster and, for both upwind and central schemes, the waveform convects more slowly at the smallest scales. One can construct more accurate finite-difference schemes to provide better resolution at high wavenumbers, but the accuracy is best at low wavenumbers. While the dispersive error for upwind and central schemes are comparable, symmetric schemes make no contribution to numerical dissipation on uni-




Figure 23: One-dimensional frequency spectra downstream of the flow over a cylinder ($Re_D = 3900$) measured along the centerline at $x/D = 7.0$ [72]. A second-order central scheme, fifth-order upwind scheme, and the experimental data of Ong & Wallace (1996) are compared. The experimental spectra shows about a half decade of inertial range scaling. The second-order central difference scheme matches the experimental spectra much better than the fifth-order upwind scheme at intermediate to high frequencies. The smaller scales are more energetic in the simulation with central differences whereas the upwind numerics rapidly dissipate the small-scale scale fluctuations. The lower-left inset shows the separating shear layers and development of Karman vortex street in the flow over a circular cylinder at $Re_D = 3900$. Shown are 16 contours of instantaneous vorticity magnitude from $\omega D/U_\infty = 0.5$ to $\omega D/U_\infty = 10.0$ [73]. Measurement location $x/D = 7.0$ along centerline indicated.

form grids. In contrast, even very high-order upwind methods can quickly dissipate a large fraction of the resolvable turbulent eddies, which is not desirable for simulating turbulence. Figure 23 compares the one-dimensional energy spectra in the cylinder wake computed using fifth-order, upwind-biased finite-difference calculations and second-order, central-difference calculations with similar grid resolution [72, 73]. Numerical dissipation of the upwind scheme tends to suppress the medium to small scales. The low-order, but non-dissipative central finite-difference scheme matches the experimental data much better than the high-order, dissipative upwind scheme. In particular, the simulations are able to reproduce the inertial subrange observed in the experiment and not dissipate the small-scale turbulent eddies. In applications, such as flow generated noise and reactive flows, small-scale fluctuations are dynamically significant and therefore one must retain this energy in the small scales without unphysical dissipation. Similarly, the dynamic modeling procedure in LES utilizes information from the small scales of the resolved flowfield to estimate the sub-filter stress, which may be inaccurate if numerical truncation error dissipates energy at the small scales. In such applications, energy conservative schemes and discretizations with low dissipation and dispersion are preferable to upwind schemes.

Low dissipation schemes are, however, susceptible to non-linear numerical instabilities due




Figure 24: Breakup of a round liquid jet by a coaxial flow of gas (air/water: density ratio $\rho_l/\rho_g \approx 832$). (a) Instantaneous snapshot of the phase interface from simulations shown. Ax-symmetric and transverse azimuthal modulations on the liquid surface in the initial breakup stage are visible. (b) Transverse azimuthal modulations grow in amplitude producing small ligaments at the wave crests. A level set method is used to track the ligaments until they pinch off near the grid resolution. (c) A capillary breakup model is then used to split the filament into several drops of varying sizes whose motion is then tracked with Lagrangian methods. Comparison with experiment for the predicted instability wavelengths, ligament growth and drop size distribution are presented in [74].

to aliasing or resolving thin features. Aliasing errors, for example, appear when non-linear products are computed discretely (e.g. the advective form $u\partial_x u$ or division by variable density in compressible flow calculations) [69, 68, 70]. The product of two continuous and smooth functions, e.g. $u(x)$ and $v(x)$, can generate scales smaller than either term individually, e.g. if $u(x) = \sin(x)$ and $v(x) = \partial_x u = \cos(x)$, then $uv = u\partial_x u = \sin(x)\cos(x) = \frac{1}{2}\sin(2x)$. This presents a problem for marginally resolved turbulence simulations. If there is substantial energy at high-wavenumbers, then non-linear products may alias wavenumbers beyond the grid resolution to lower wavenumbers (e.g. $k_i + k_j = k_{i+j+lN}$ for integer l due to the periodicity of the exponential basis). Aliasing errors resulting from discretization of the non-linear terms can become significant and corrupt the solution.

Many practitioners introduce numerical dissipation to suppress these non-linear numerical instabilities. For example, upwind schemes naturally reduce aliasing errors by truncating the energy at high-wavenumbers as shown in figure 23. Unphysical dissipation such as this can provide some level of robustness but the resulting simulations may not be time-accurate. To limit these effects, one can apply dissipative (e.g. WENO/upwind) numerics in local, limited regions, e.g. near thin features such as the shocks seen in figure 17 and phase interfaces shown in figures 8 and 24, without dissipating the surrounding turbulent eddies too much [75, 76, 39]. Robust, time-accurate discretizations are, however, achieved by enforcing higher-order discrete conservation principles with non-dissipative differencing schemes.

As discussed earlier, artificial numerical dissipation leads to unphysical suppression of

turbulent eddies and should be avoided. Instead, a better way to achieve robustness is to construct energy conserving numerical discretizations [69, 77]. Failure to do so often results in computational instability and spurious conclusions [78, 79]. Many different discretely conserving discretizations for turbulence simulation have been developed [80, 77, 69, 56, 17, 81]. One common way to obtain these properties is to develop a discrete calculus to approximate the derivatives of the initial boundary value problem with accurate, non-dissipative operators that satisfy basic calculus identities, such as the product rule and integration-by-parts, at a discrete level. For example, in the continuous case, integration-by-parts on periodic domains gives $\int u(x) \frac{d}{dx} v(x) dx = - \int v(x) \frac{d}{dx} u(x) dx$. The discrete analogue is known as summation-by-parts, which takes the form

$$\sum_{j=1}^N u_j \frac{dv_j}{dx} = - \sum_{j=1}^N v_j \frac{du_j}{dx} , \quad (11)$$

for periodic domains. For explicit central difference schemes, e.g. the fourth-order scheme of equation 8, one can show that summation-by-parts is satisfied by simply substituting u'_j and v'_j from equation 8 and using the periodicity condition. For non-periodic boundary conditions, it is possible to construct finite-difference operators (central inside the domain and non-central near boundaries) such that

$$\sum_{j=1}^N h_j u_j \frac{dv_j}{dx} = u_j v_j \Big|_{j=1}^N - \sum_{j=1}^N h_j v_j \frac{du_j}{dx} , \quad (12)$$

is satisfied with weights h_j . Recall that, at the end of §3.1, we showed how, in the absence of external forces or viscous dissipation, momentum and kinetic energy are conserved by applying simple calculus relations for continuous functions (e.g. integration by parts) to the governing equations. By following these same steps using the above discrete calculus identities, we can develop numerical discretizations that also satisfy these general conservation properties [12].

Methods with low numerical dissipation and dispersion are essential for credible LES computations. As we have shown, numerical dissipation present in most RANS codes and upwind schemes is inadequate for LES, and therefore, special consideration must be given to construct low dissipation and dispersion error schemes for calculation of turbulent flows in complex geometries. A discrete calculus combined with non-dissipative discretization schemes provides a rational framework for the construction of accurate turbulence simulation algorithms on unstructured grids [77]. In the absence of numerical dissipation, robustness is achieved by enforcing higher order discrete conservation principles, which is a bit challenging in unstructured mesh settings. The `charles` codebase implements these desirable properties on unstructured grids.

3.4 Computational approach of `charles`

The flow solver developed at the Center for Turbulence Research, known as `charles`, implements these non-dissipative, energy-conserving numerical methods on unstructured grids




Figure 25: Parallel performance of `charles` measured on the Intrepid BlueGene/P at the Argonne Leadership Computing Facility (ALCF). Strong scalability of (a) compressible (40M CV grid) perfectly expanded supersonic jet, (b) incompressible (20M CV grid) scalar transport over urban city topography, and (c) parallel speedup for simulations of an imperfectly expanded supersonic jet with chevrons (528M CV grid, see figure 17) from 8 racks (4096 cores) to the entire Intrepid machine of 40 racks (163840 cores) shown.

and employs dynamic subfilter-scale LES turbulence models. This technology has enabled several large-scale, high-fidelity simulations of turbulent flows in complex geometry configurations including combustion, multiphase flows and aeroacoustics [14]. Many of the turbulent flow simulations presented in this article were produced using the `charles` codebase. For high-fidelity, predictive simulation, one must also effectively utilize available computing resources. We now describe the parallel performance of `charles` and outline some of the implementation details that enable `charles` to utilize leadership-class computing systems for production simulations.

For the compressible flow module of `charles`, the numerical method is purely hyperbolic and fully explicit. As a result, `charles` requires relatively little communication at each timestep. The interprocessor communication and parallel I/O scales as $O(p \log p)$ at worst, where p indicates the number of processors. `charles` uses MPI-2 for communication and MPI-IO for parallel file I/O. For the supersonic jet noise calculations shown in figure 17, the I/O module provided a throughput of 16GiB/s (average) when writing on 163840 cores of Intrepid producing about 25TB of data after running for just one day. To collect statistics for the FWH projection described earlier, several statistics along the FWH surface were outputted every few timesteps taxing the I/O system. For this run, the total time was split between about 20% I/O and 80% computation.

The analysis presented in §3.3 showed that central difference schemes are non-dissipative

on uniform grids and linearly stable. Highly skewed meshes, however, can have a strong effect on the numerical stability properties, accuracy and diffusion error of a given discretization scheme [70]. Mesh skewing is largely unavoidable in complex geometry configurations. To help eliminate these skewing-induced errors, `charles` uses widened numerical stencils with essentially zero dissipation in regions where the grid quality is good and applies some local dissipation in regions of less-than-perfect grid quality, e.g. near sharp grid transitions. These widened stencils require only slightly more communication than just simple nearest-neighbor data exchanges. In this way, `charles` achieves fourth-order accuracy on well-behaved meshes, reverting to second-order accuracy inherent to all unstructured solvers only in regions where the mesh becomes highly distorted. This is balanced by the ability to efficiently cluster grid points smoothly around specific areas of interest by isotropic, unstructured grid refinement (e.g. see figure 18 for an example). Standard domain decomposition packages such as Parmetis are used to partition these unstructured grids. Dual-constrained partitioning is often used for load-balancing to make sure each processor has roughly an equivalent computational workload. Load-balancing is required in such problems involving discretization operators of different complexity, Lagrangian particle transport, or code-coupling to make the computational workload for each processor more uniform.

For example, shock waves and turbulence appear together in many flows, including supersonic jet noise and supersonic mixing layers (e.g. figures 9 and 10). Simulations of both shocks and turbulence at the same time are made difficult by the often contradictory numerical methods used to capture such phenomena. `charles` uses a hybrid Central-WENO scheme to simulate flows involving shocks. The scheme has three pieces: (1) a central scheme, described previously, (2) a scheme appropriate for computing a flux across a shock, and (3) a hybrid switch, which detects where shocks are present in the flow, and activates the shock-appropriate scheme [75, 76]. The WENO reconstruction is often much more expensive than the standard central discretization. In many compressible turbulent flows, shocks or shocklets are sparsely distributed in the computational domain and fluctuate position within the unsteady flow. As a result, a single processor may end up with too many cells where the shock sensor is activated (e.g. near the chevron-tip nozzle of figure 17) and thus degrade the parallel performance if not properly load-balanced. Occasionally, the grid partition is occasionally updated or reset to account for these fluctuating eddy shocklets.

To demonstrate the scalability of the algorithms and the ability of this advanced unstructured LES technology to utilize leadership-class computing resources for a wide variety of problems, figure 25 shows the strong scaling of the `charles` code on the Intrepid Blue-Gene/P system at the Argonne Leadership Computing Facility (ALCF) for several different production simulations. Figure 25(a) shows scaling of the compressible `charles` module for a perfectly expanded supersonic turbulent jet calculation (40M CV grid) while figure 25(b) shows the scaling of the incompressible `charles` module for scalar transport over complex urban topography (20M CV grid). The primary parallel bottleneck for the incompressible `charles` solver is the linear solvers (in this case, the algebraic multigrid Hypre solver) for the pressure-Poisson equation; however, the scalability and parallel performance is quite favorable even in comparison to the fully explicit compressible flow solver. Both compressible

and incompressible `charles` modules show good strong scaling on Intrepid at ALCF from 128 cores to 16384 cores, i.e. down to about 1000 control volumes (CVs) per core, with a parallel efficiency greater than 70%.

Jets operating at perfectly expanded conditions do not generate screech tones or broadband shock-associated noise; as a result, the scaling study shown in figure 25(a) does not fully test the shock load-balancing capability of `charles`. Underexpanded supersonic jets do generate strong shocks and unsteady shocklet regions that require load-balancing to scale to high core counts. Figure 17 shows several regions where such shocks appear, e.g. the tiny shocklets generated in the chevron tips and troughs that break up the oscillating shock cell train into many small shock cells causing the chevron jet exhaust plume to expand in many small steps [65]. For the compressible module, load imbalance from the more expensive WENO reconstructions applied in about 1% of the flow domain in these regions where the shock sensors are activated presents a serious bottleneck to improved scaling. Improved adaptive load-balancing algorithms were developed to account for these more expensive reconstructions by weighting regions of intense dilatation more heavily when the grid is partitioned. Figure 25(c) shows how load-balancing enabled `charles` to effectively utilize all 40 racks of Intrepid for the underexpanded supersonic jet calculation with chevrons shown in figure 17. In total, this simulation used about 16 million core-hours (i.e. about 4 days on 40 racks of the Intrepid BlueGene/P at ALCF). The 80% parallel efficiency sustained on the full 40 racks (163,840 core) of the Intrepid BlueGene/P at ALCF depended critically upon properly load-balancing the shock-capturing faces.


4 Validation, verification, uncertainty quantification

Procedures to establish the quality of numerical simulations have been organized within the framework of verification and validation (V&V) activities [15, 16, 82]. Verification is a mathematical process that aims at answering the question: “are we solving the equations correctly?”. The objective is to quantify the errors associated to the algorithms employed to obtain the solution of the governing equations. Validation, on the other hand, aims at answering the question “are we solving the correct equations?”. The goal is to identify the appropriateness of the selected mathematical/physical formulation to represent the device to be analyzed. Validation always involves comparisons of the numerical predictions to reality, whereas verification *only* involves numerical analysis and tests.



An iterative process involving both verification and validation is necessary for credible LES computations. Figure 26 shows verification and validation applied to LES of flow past a cylinder. This example demonstrates the necessity of both verification and validation. Validation alone is insufficient; a poorly under-resolved simulation can fortuitously match an experiment. As figure 26 shows, early transition to turbulence due to unconsidered vibration in a physical experiment and numerical truncation noise in an under-resolved simulation ($N_z = 4$) can combine to give the appearance of agreement and predictive capability. In this case, verification by grid-refinement isolated the numerical errors from the physical modeling uncertainties such that the refined simulations predicted much longer

LES of Cylinder $Re_D = 3900$

(a) $N_z = 4$



(b) $N_z = 48$



(d) PIV Data




Figure 26: Validation/verification of flow over a cylinder at $Re_D = 3900$ and effects of grid resolution. Simulations with (a) $N_z = 4$ ($L_z = \pi D/2$) and (b) $N_z = 48$ ($L_z = \pi D$) grid points in the spanwise direction shown (instantaneous total vorticity contours shown). (c) Profiles of mean streamwise velocity at three locations downstream of the cylinder comparing the coarse/fine grids and experimental data of Lourenco & Shih (1993). (d) Instantaneous near-wake vorticity field at three Reynolds numbers from the experiments of Chyu & Rockwell (1996). The agreement between the under-resolved simulation and the experiment of Lourenco & Shih (1993) is fortuitous. Unconsidered external vibrations in the experimental facility led to early transition; numerical truncation error/noise in the $N_z = 4$ simulation also prompted early transition to turbulence in the separating shear layers and, consequently, a shorter vortex formation region, typical of much higher Reynolds numbers. The refined simulation with $N_z = 48$ exhibited long separated shear layers confirmed by (c) controlled PIV experiments, simulations and other experiments (e.g. see figure 23). Grid refinement helped verify that numerical truncation error was not affecting the results, and validation against other data helped establish the physical fidelity and credibility of these simulations.

shear layer development before breakdown to turbulence, which was later confirmed by other experiments and well-resolved simulations. In this way, one can get the right results for the right reasons.

There is a growing recognition of the fact that validation cannot be carried out without explicitly accounting for the uncertainties present in both the measurements and the computations. Experimentalists are typically required to report *uncertainty bars* to clearly identify the repeatability and the errors associated to the measurements. Validation must be carried out acknowledging the nature of the experimental uncertainties and by providing a similar indication of the computational *error bars*. One of the objective of uncertainty quantification methods is to construct a framework to estimate the error bars associated to given predictions. Another objective is to evaluate the likelihood of a certain outcome; this leads to better understanding of risks and improves the decision making process.

4.1 Errors vs. uncertainties

The American Institute of Aeronautics and Astronautics (AIAA) “Guide for the Verification and Validation of CFD Simulations” defines *errors* as recognizable deficiencies of the models or the algorithms employed and *uncertainties* as a potential deficiency that is due to lack of knowledge. This definition is not completely satisfactory because does not precisely distinguish between the mathematics and the physics. It is more useful to define *errors* as associated to the *translation* of a mathematical formulation into a numerical algorithm (and a computational code).

Errors are typically classified in two additional categories: acknowledged errors are known to be present but their effect on the results is deemed negligible. Examples are round-off errors and limited convergence of certain iterative algorithms. On the other end, unacknowledged errors are not recognizable¹ but might be present; implementation mistakes (bugs) or usage errors can only be characterized by comprehensive verification tests and procedures.

Using the present definition of errors, the *uncertainties* are naturally associated to the choice of the physical models and to the specification of the input parameters required for performing the analysis. As an example, numerical simulations require the precise specification of boundary conditions and typically only limited information are available from corresponding experiments and observations. Therefore variability, vagueness, ambiguity and confusion are all factors that introduce uncertainties in the simulations. A more precise characterization is based on the distinction between aleatory and epistemic uncertainties.

4.2 Aleatory uncertainty

Aleatory uncertainty² is the physical variability present in the system being analyzed or its environment. It is not strictly due to a lack of knowledge and cannot be reduced. The determination of material properties or operating conditions of a physical system typically leads to aleatory uncertainties; additional experimental characterization might provide more conclusive evidence of the variability but cannot eliminate it completely. Aleatory uncertainty is normally characterized using probabilistic approaches.

4.3 Epistemic uncertainty

Epistemic uncertainty³ is what is indicated in the AIAA Guide (AIAA 1998) as “uncertainty”⁴, i.e. a potential deficiency that is due to a lack of knowledge. It can arise from assumptions introduced in the derivation of the mathematical model used or simplifications related to the correlation or dependence between physical processes. It is obviously possible

¹In principle, using the AIAA definition, unacknowledged errors could be considered uncertainties because they are associated to lack of knowledge

²Aleatory uncertainty is also referred to as variability, stochastic uncertainty or irreducible uncertainty.

³Epistemic uncertainty is also called reducible uncertainty or incertitude

⁴Aleatory uncertainty is not mentioned in the AIAA Guidelines

to *reduce* the epistemic uncertainty by using, for example, a combination of calibration, inference from experimental observations and improvement of the physical models. Epistemic uncertainty is not well characterized by probabilistic approaches because it might be difficult to infer any statistical information due to the *nominal* lack of knowledge. A variety of approaches have been introduced to provide a more suitable framework for these analysis. Typical examples of sources of epistemic uncertainties are turbulence modeling assumptions and surrogate chemical kinetics models.

4.4 Sensitivity vs. uncertainty analysis

Sensitivity analysis (SA) investigates the connection between inputs and outputs of a (computational) model; more specifically, it allows to identify how the variability in an output quantity of interest is connected to an input in the model and which input sources will dominate the response of the system. On the other hand, uncertainty analysis aims at identifying the overall output uncertainty in a given system. The main difference is that sensitivity analysis does not require input data uncertainty characterization from a real device; it can be conducted purely based on the mathematical form of the model. Large output sensitivities (identified by SA) do not necessarily translate into important uncertainties because the input uncertainty might be very small in a device of interest.

4.5 Predictions under uncertainty

Computer simulations of an engineering device are performed following a sequence of steps. Initially the system of interest and desired performance measures are defined. The geometrical characterization of the device, its operating conditions, the physical processes involved are identified and their relative importance must be quantified. It is worthwhile to point out that the definition of the system response of interest is a fundamental aspect of this phase. The next step is the formulation of a mathematical representation of the system. The governing equations and the phenomenological models required to capture the relevant physical processes need to be defined. In addition, the precise geometrical definition of the device is introduced. This step introduces simplification with respect to the real system; for example small geometrical components are eliminated, or artificial boundaries are introduced to reduce the scope of the analysis. With a well defined mathematical representation of the system, the next step is to formulate a discretized representation. Numerical methods are devised to convert the continuous form of the governing equations into an algorithm that produces the solution. This step typically requires, for example, the generation computational grid, which is a tessellation of the physical domain. Finally the analysis can be carried out. To demonstrate how these elements of uncertainty quantification can be applied to complex engineering systems, we now describe an uncertainty quantification methodology applied to a complete aeroacoustic problem involving trailing-edge noise shown in figure 27.




Figure 27: Uncertainty quantification methodology for trailing-edge noise from blades. Dashed-green boxes correspond to the stochastic inputs and outputs from the corresponding “black-box” operations shown as solid-black boxes. Suggested inputs where uncertainty quantification (UQ) is applicable are denoted by red ellipses; we discuss the stochastic response from uncertain inflow velocity inputs in §4.6. LES and RANS of a fan blade are performed on a restricted domain using different inflow velocity inputs with 2.5% and 10% variation in streamwise and cross-wise velocity, respectively, about a mean profile obtained from a RANS computation of the full wind tunnel domain (see inset). After propagating this uncertainty through the LES and RANS, the computed trailing-edge wall-pressure spectra Φ_{pp} are input to a far-field noise model to evaluate the stochastic response of the far-field sound spectrum S_{pp} . Note that the RANS approach requires an additional modeling step to evaluate the wall-pressure spectra. Two such models (denoted by YR and PL) were evaluated, which require as input different global and local mean flow parameters available in steady RANS such as the trailing-edge velocity profile.

4.6 Uncertainty quantification for fan blade trailing-edge noise

The noise generated by modern rotating machinery includes multiple discrete tones and broadband components distributed over a very wide range of scales. Tonal or narrow-band noise is often the result of periodic forcing, e.g. tuned to the fan rotation frequency or due to coherent vortex shedding typical of flow past blunt trailing edges or at high angles-of-attack. Advances in passive and active noise control have helped reduce annoying tonal noise to levels indistinguishable from the background (broadband) noise [62, 61]. In many applications, broadband noise now limits the minimum achievable noise level. Reducing broadband noise at the source is then necessary to decrease the overall noise level further and comply with environmental noise regulations.

A key source of broadband noise is the interaction of turbulent eddies with the trailing-edge of lifting blades, especially at low Mach numbers [83, 62]. When turbulent boundary




Figure 28: Fan blade LES and RANS computations were performed on a restricted subset domain of the full experimental wind tunnel facility with installed fan blade (left). The boundary conditions for this smaller domain were parameterized by prior experiments and simulation results with specified aleatoric uncertainty bounds of 2.5% and 10% for (b) streamwise U and (c) cross-wise V velocity components, respectively. (a) Parameterization of inlet boundary condition $s \in [0, 1]$ (solid line) and velocity magnitude contours (dashed line). Note that standard outflow boundary conditions were applied along the vertical boundary connecting $s = 0$ to $s = 1$. Parametric inlet velocity profiles: (b) streamwise velocity U and (c) crosswise velocity V . Experimental measurements (square). The numerical wind tunnel inlet velocity profile was used as reference for stochastic collocation and not the experimental measurements. Uncertainty bounds around numerical inlet profile (dash-dot) reflect this variation in the nine LES collocation points.

layer eddies convect past the sharp edge of a solid body, the phase cancellation that otherwise occurs between the acoustic fields excited by turbulent fluctuations is destroyed. As a result, the trailing-edge makes the near-wall turbulence a more efficient noise source by preferentially scattering the near-field turbulent pressure fluctuations into strongly propagating sound waves. This trailing-edge noise often represents the dominant source of noise produced by rotating machines such as automotive and computer cooling fans, turboengines, propellers and wind turbines.

Recently, we have assessed some of the aleatory uncertainties associated with the prediction of trailing-edge noise [84]. The uncertainty quantification methodology and propagation path is shown in figure 27. Uncertainty related to the inlet velocity is propagated through both steady RANS and unsteady LES to predict the trailing-edge wall-pressure fluctuations. The uncertainty in the trailing-edge wall-pressure fluctuations is then propagated through two semi-empirical models for computing wall-pressure fluctuations from steady RANS and compared to unsteady LES predictions of the trailing-edge noise computed from an analytic far-field noise model under uncertainty related to the inflow velocity.

Input uncertainties are propagated through the differential equations and trailing-edge noise models using both classical Monte Carlo (MC) and stochastic collocation (SC) methods. Both approaches prescribe an ensemble of inputs for which the corresponding solutions (realizations) are collected to obtain statistics of various output quantities (such as the mean




Figure 29: Mesh configuration in x - y plane (left) for the fan blade 3D, unsteady LES. The 2D, steady RANS computations used a similar grid. In the LES, an unstructured grid was used so that grid coarsening in the spanwise direction was applied in proportion to the distance from the fan blade as well as grid refinement near the blade surface. To prevent poor mesh quality due to this coarsening, a smooth transition is applied by inserting a layer of pyramids or wedges between adjacent levels of the hexahedral grid (middle). This reduced the number of control volumes to 1.5 million. (b) Cut of the 3D unstructured mesh through the airfoil surface. (c) Close-up view of the grid expansion near the wall in the y - z plane.

and variance). Both schemes are *non-intrusive* in that deterministic flow solvers are used as black boxes where the effect of input uncertainties (e.g. initial conditions, boundary conditions and model parameters) on the output response quantity of interest (e.g. trailing-edge wall-pressure fluctuations) are evaluated without modifying the internal workings or source code to explicitly include stochastic expansions. In contrast, intrusive methods, such as stochastic Galerkin approaches, represent the uncertain solution in terms of a stochastic expansion explicitly incorporated into a stochastic flow solver.

Most statistics of interest, such as the expected value or variance, are then computed by integrating over the domain spanned by the input uncertainties. To do this, Monte Carlo methods randomly sample the stochastic space, but are slow to converge and require a relatively large number of samples to compute statistics of interest, which is prohibitive for turbulent flow realizations (and, in particular, the more high-fidelity and costly LES). Stochastic collocation methods instead precisely specify the uncertain input variables so as to take advantage of accurate numerical integration methods such as Gauss quadrature, which allows one to reduce the number of simulations needed for convergence as long as the number of uncertain input parameters is small enough.

In this example, the inlet velocity profile (U, V) along a boundary surrounding a two-dimensional fan blade is taken as the only uncertain input parameters to the LES and RANS flow solvers as shown in figure 28. Wind tunnel cross-wire velocity measurements and companion simulations of the entire wind tunnel facility were used to select uncertainty bounds of 2.5% for the streamwise velocity U and 10% for the cross-wise velocity V about the mean inflow velocity profiles for a smaller, restricted domain about which an ensemble of simulations were performed to evaluate the stochastic response. This aleatoric uncertainty is then propagated through an ensemble of both RANS and LES calculations.




Figure 30: Q-criterion isocountors ($Qc^2/U_\infty^2 = 2000$) colored by velocity magnitude for four of the nine LES collocation points (#1, #5, #7 and #9). The variation in angle-of-attack from case #1 to case #9 is about two degrees (i.e. one degree from #5 to #9). This small change in the input flow velocity has a large effect on the laminar-turbulent transition, separation bubbles, trailing-edge wall-pressure fluctuations and far-field noise for flow over this fan blade, especially at lower angles-of-attack, e.g. #7-#9. Q-criterion isocontours highlight this change by showing re-laminarization across the span, reduced leading-edge separation bubble and growth of turbulent length scales near the trailing-edge.

With the available computational resources, we performed an ensemble of 9 independent large-eddy simulations at the stochastic collocation points. These collocation points correspond to the 9 different inflow conditions (or angles of attack) shown in figure 28. The reference fan blade (case #5) is at a moderate angle-of-attack of 8° with the other cases within two degrees of the reference case. The change in flow topology between the different outputs due to this modest variation in inputs is shown in figure 29. In most cases, a small laminar separation bubble appears at the leading edge that triggers instabilities near the reattachment point that quickly transition to turbulence. The flow tends to re-laminarize toward mid-chord as the favorable pressure gradient accelerates the flow. When this gradient becomes adverse, the turbulent boundary layer thickens and larger turbulent structures develop near the leading edge. In LES #7, the lower angle-of-attack generates a weak separation bubble not strong enough to trigger transition across the entire span until the adverse pressure gradient at mid-chord triggers full transition and turbulence de-




Figure 31: Uncertainty quantification of mean integrated inputs and order convergence studies. (a) Mean wall-pressure coefficient C_p and corresponding uncertainty bars for RANS and LES computed by stochastic collocation compared with experimental measurements. Convergence of stochastic collocation and Monte Carlo sampling of RANS and LES computations was performed by increasing the number of collocation points N (i.e. order refinement not grid refinement). We compare the results for two different integrated mean statistics: (b) wall-pressure coefficient C_p shown in (a) for $N = 9$ and (c) the trailing-edge velocity profile $U(y)$ shown in figure 32. With only a few stochastic inputs, stochastic collocation is able to converge faster than Monte Carlo (e.g. the same level of convergence is achieved with 17 RANS computations as with 1000 Monte Carlo samples). Note that stochastic collocation (SC) should converge exponentially in the asymptotic limit for smooth functions. This asymptotic limit is not achieved for the LES stochastic collocation with only nine points due to the rapid change in statistics at lower angles-of-attack where the stochastic response functions become relatively non-smooth. Just as one performs grid convergence studies by increasing the number of grid points, one can increase the number of stochastic samples to perform an order convergence study of a given uncertainty quantification methodology. This is an important part of code verification.

velopment. At a slightly smaller angle-of-attack, LES #9 shows how the weak acceleration around the leading-edge is insufficient to generate the leading-edge separation bubble so that no transition to turbulence occurs before mid-chord.

The LES predicted a large and rapid change in flow character at lower angles of attack. Transition to turbulence in the LES occurs suddenly for a very small variation of incidence as in LES runs #8 and #9. The large uncertainty in the leading-edge separation bubble and the second laminar recirculation bubble at mid-chord for low incidence angles seen in the visualizations of figure 29 also appear in several statistics of interest shown in figure 30 such as the wall skin-friction drag and the wall-pressure coefficient. The RANS model predicted a monotonic variation in skin-friction and was not able to reproduce the more complex laminar-turbulent transition seen in LES at low angles-of-attack. This rapid change in statistics provided marginal but still sufficient convergence for stochastic collocation with only nine LES samples. Developing methods to handle such unsteadiness and discontinuous responses (e.g. laminar-turbulent transition, shocks, airfoil flutter, etc.) to determine the




Figure 32: Trailing-edge statistics and stochastic response (at $x_c/C = -0.02$). (a) Mean streamwise boundary-layer velocity-profile $U(y)$ and uncertainty bars (i.e. the stochastic mean profile about which the variation in the stochastic outputs appear as uncertainty bars is shown). Wall-pressure frequency spectra directly computed by (b) LES and (c) the YR model output from the RANS computations. The wall-pressure spectra outputs from the nine collocation points are shown (the corresponding stochastic mean and uncertainty bars are shown in figure 33). Note the increased energy of the wall-pressure fluctuations at lower angles-of-attack from the LES (about 6 dB for most frequencies). Experimental measurements at the same location are shown.

stochastic response of a system of equations to uncertainty is an active research area [82].

Stochastic collocation methods are well suited to problems where only a few stochastic parameters are of interest such as this and available computational resources are limited. For the RANS calculations, we found similar statistical convergence with 1000 Monte Carlo simulations as for 17 stochastic collocation computations. Over 1000 RANS were performed using desktop workstations. Performing several thousand LES, however, was not feasible given the available computing resources and time constraints (especially since we integrated each sample for an unusually long time to provide excellent statistical convergence at low-frequencies). The ability of stochastic collocation methods to reduce the required number of samples enabled the evaluation of the stochastic response of the flow past and noise produced by fan blades with LES. RANS was unable to predict the associated laminar-turbulent transition and recirculation bubbles near the leading-edge and mid-chord at low angles of attack.

The new physics accessible by LES were found to significantly affect the resulting stochastic noise response. Both LES and RANS agree with the experimental measurements (in that most single-point experimental measurements are included within the LES and RANS uncertainty bars). The trailing-edge boundary-layer profile and wall-pressure fluctuations for the LES and RANS computed using stochastic collocation are compared in figure 31. The LES trailing-edge velocity profiles had substantially smaller variation than RANS (which are input to the RANS wall-pressure fluctuation models). In both LES and RANS, wall-pressure fluctuations were used as input to an analytic trailing-edge noise model. Since steady RANS does not intrinsically account for wall-pressure fluctuations, the mean velocity statistics from




Figure 33: Comparison of (a) wall-pressure spectra and (b) far-field acoustic spectra in terms of mean and uncertainty bars. A far-field noise model (Amiet’s model) for trailing-edge noise is used to convert the input wall-pressure fluctuations to output far-field acoustic spectra. This is similar to how the FWH method used for supersonic jet noise shown in figure 17 is used to project the near-field turbulence information into far-field acoustics except that additional modeling assumptions are made in the formulation of Amiet’s model specific to trailing-edge noise and reflections from the leading-edge.

RANS were used to compute the inputs (e.g. trailing-edge boundary layer thickness, wall-shear stress, etc.) to two models for boundary layer wall-pressure fluctuations (denoted as the YR and PL models). Unsteady LES, on the other hand, does directly compute wall-pressure fluctuation statistics. In particular, the LES computations developed a recirculation bubble close to the trailing edge at low angles-of-attack that triggered much larger wall-pressure fluctuations at the trailing-edge for all frequencies as shown in figure 31. These computed wall-pressure fluctuations were then used as input to a far-field noise model. From this, the stochastic response of the noise spectrum produced by a fan blade under uncertainty related to the inflow velocity profile was computed using stochastic collocation and is shown in figure 32. Several statistics of interest agree with experiment, but the LES runs showed a larger spread (or standard deviation) in the output wall-pressure fluctuations and noise than RANS at most frequencies due to the additional flow separation and transition found at lower incidence angles. Extending this methodology to rotating machines with additional input parameters, such as mass flow rate and rotational frequency, would be valuable exercises (possibly, requiring different stochastic algorithms to address a much larger number of uncertain parameters). This example demonstrates the complexities of applying uncertainty quantification to a complete aeroacoustic problem with inlet velocity profiles as uncertain parameters.

5 Summary

Computations enable us to investigate phenomena where economics or physical and environmental constraints preclude experimentation. The last thirty years have seen the rise of computer-aided engineering in almost every industrial sector. Today, many aspects of product development, design, optimization, performance analysis and certification rely heavily on the use of computations. Computers also provide new avenues for scientific discovery. Today computational modeling is an integral part of aircraft and engine design and is responsible for dramatic reductions in the required expensive wind tunnel and engine tests. In spite of its successes, computational engineering is far from being predictive for complex engineering systems. Use of high fidelity methods, physics-based modeling, computer science, and validation, verification and uncertainty quantification tools in industrial settings, e.g. for high fidelity integrated simulations of complex engineering systems, are required before achieving the predictive status necessary to move past the current computational plateau.

The development of high-fidelity simulation techniques for prediction of multiple-physics, complex turbulent flows was examined in this article. Several applications ranging from the prediction of supersonic jet noise, combustion in realistic gas turbine engines, air-layer drag reduction for marine vehicles and separation control on airfoils were presented. The multi-scale nature of turbulence creates unique challenges for numerical simulations. Discretization methods must preserve the physical processes, reducing or eliminating artificial dissipation and dispersion. Moreover, the extreme computational effort required to capture all the temporal and spatial scales of motion leads to the introduction of physical models for unresolved flow features such as large-eddy simulation. In this regard, development of wall models for LES of high Reynolds number boundary layers, and subgrid scale models for multi-physics phenomena where the dynamics of thin interfaces are critical to the development of the large scale flow remain pacing items for LES. We have emphasized how validation, verification and uncertainty quantification can establish credibility and confidence in numerical simulations of turbulent flows. This ultimately leads to the design and certification of complex engineering devices.

References

- [1] P. Moin and J. Kim. Tackling turbulence with supercomputers. *Scientific American*, 276(1), 1997.

Highly recommended reading. What is turbulence? Why does turbulence matter? Why do golf balls have dimples? Why does your heart beat? How can we improve fuel efficiency? Why do planes, trains and automobiles work? Why industrial competitiveness is enhanced by tackling turbulence with supercomputers? How do we study turbulence? How do we predict turbulence? How can we control turbulence? Jam packed with everything you need to know about turbulence and supercomputers. Read it here: <http://www.stanford.edu/group/ctr/articles/tackle.html>

- [2] D. Chapman. Computational aerodynamics development and outlook. *AIAA Journal*, 17(12):1293–1313, 1979.

Chapman outlines three motivations for development of computational aerodynamics: (1) access to new technological capabilities, (2) energy conservation, (3) economics. In 1979, energy costs associated with wind tunnel testing imposed significant restrictions on testing time whereas computation was comparably less energy intensive. For example, Los Alamos purchased the first Cray-1 system in 1976 for 8.8 million USD needing about 300 kW to operate; a 40 ft \times 80 ft low-speed ($M \sim 0.1$) air wind tunnel would cost more than 100 million USD to procure and require about $(1/2)\rho U^3 A \approx 20$ MW to operate. Today, however, the leadership computing facilities are increasingly limited by power constraints facing many of the same budget and time constraints seen both for procurement, construction and day-to-day operation of large wind tunnel facilities. For comparison, the first sustained petaflop system, the Roadrunner at Los Alamos, was procured for about 130 million USD requiring about 4 MW to operate continuously. Current DOE literature projects a power envelope of 20 MW for future exascale architectures.

Chapman also details the motivation and developmental history of the major approximations to the governing Návier-Stokes equations: linearized inviscid, nonlinear inviscid, Reynolds-Averaged Návier-Stokes, large-eddy simulation (LES). Chapman provides a detailed analysis of grid-resolution requirements for the LES of turbulent boundary layers relevant to aircraft [58]. To compute the flow over an airfoil, Chapman estimated that the required number of grid points for wall-resolved LES would scale as $N \sim \text{Re}_c^{1.8}$, where Re_c is the chord Reynolds number while wall-modeled LES would scale as $N \sim \text{Re}_c^{0.4}$. The nearly quadratic dependence on Reynolds number for wall-resolved LES is a consequence of the turbulent eddies becoming progressively smaller close to the wall. As a result, a wall-resolved LES must resolve down to the very small scales extremely close to a wall. Alternatively, one may model these fine-scale structures near a wall as a part of a general boundary condition for the resolved LES in the region only slightly away from walls. Wall-modeled LES attempts to do this by, for example, using unsteady RANS in the near-wall region to supply a boundary condition for LES in the outer layers. Chapman also emphasized the importance of continued communication and co-design between the turbulence community, computational scientists and hardware designers is essential. Without such attention to co-design, the ILLIAC-IV machine at NASA-Ames would not have produced so many landmark turbulence simulations.

- [3] H. Tennekes and J. L. Lumley. *A first course in turbulence*. MIT Press, 1972.

A classic text on turbulence. Highly recommended.

- [4] M. Van Dyke. *An Album of Fluid Motion*. The Parabolic Press, 10th edition, 2005.

Collection of black-and-white photographs of flow visualizations illustrating the great diversity of fluid motion compiled by Milton Van Dyke. Highly recommended.

- [5] R.W. Stewart. Turbulence. In A. Shapiro, editor, *National Committee for Fluid Mechanics Films*, 1969.

A great instructional video about turbulence. Highly recommended. You can view this video and many others about fluid motions here: <http://web.mit.edu/hml/ncfmf.html>

- [6] D.J. Mavriplis, D. Darmofal, D. Keyes, and M. Turner. Petaflops opportunities for the NASA fundamental aeronautics program. *AIAA 2007-4804*.

They observe that the engineering community is largely not actively participating in the renewed interest in high performance computing at the national level in recent years. Advocacy for high-performance computing has increasingly been championed solely by the science community with the oft-heard argument that computation is a “third-pillar” of scientific discovery alongside theory and experiment. They argue that, in recent years, computational engineering is not complex enough to warrant the deployment of state-of-the-art hardware for industrial design, research and development. They argue that an aggressive computational engineering program is at least (if not more) important as any science problem, particularly with respect to enhancing national competitiveness. They contrast this with the historical evidence of strategic computing targeted towards Grand Challenge problems of “integrated, multi-disciplinary simulations and design optimization of aerospace vehicles throughout their mission profiles” driven by engineering needs. This is well illustrated by the developments at NASA Ames (e.g. ILLIAC IV), the Center for Turbulence Research at Stanford University, and the original DOE Accelerated Strategic Computing Initiative.

They argue that the engineering community should strive to demonstrate the benefit of high-performance computing (as made clear in aerodynamic [9] and gas-turbine combustor design [7]). To do so, several grand challenge problems are proposed including (a) complete flight-envelope characterization (digital flight), (b) aerospace propulsion (e.g. the simulation of a Pratt & Whitney engine carried out at the Center for Turbulence Research at Stanford University supported by the DOE ASC program), and (c) computational design and optimization. They identify six major barriers to progress: (1) massive parallelism, (2) multi-scale resolution, (3) multiple physics, (4) simulation fidelity, robustness and automation, (5) software complexity and (6) technology transfer. They also suggest six areas of investment: (1) hardware

availability and scalable solvers, (2) algorithmic development, (3) physical modeling, (4) supporting experiments, (5) software infrastructure and (6) education.

- [7] W.W. Kim and S. Syed. Large-eddy simulation needs for gas turbine combustor design. *AIAA 2004-331*, 2004.

Reviews the 2004 capability of computational fluid dynamics for gas turbine combustor design and identifies when and where such methods are applicable. The potential of Large-Eddy Simulation is demonstrated through application to gas turbine combustor problems (e.g. swirl vanes, bluff bodies, recirculation zones, fuel/air mixing, turbulent combustion, jet-in-crossflow, liquid-fuel breakup, fuel injector spray modeling, turbulence-chemistry interaction, turbulent mixing, flamelets, breakup, coalescence, dispersion, vaporization, combustion instabilities, and most gas turbine applications covering the entire flight envelope of engines). LES provides improved fidelity and is a very attractive tool for numerical simulation of turbulent flows. They say that the variation among LES models is much smaller than RANS models in reacting flows so that LES is a clear “winner.” They argue that LES is required to enhance the capability and expand the applicable design space of CFD for combustor design. The two most critical needs to achieve “design tool” status are (1) continuous improvement of turbulent combustion models and (2) extensive data acquisition to validate reacting-flow LES sub-models. As modern combustor design evolve, the inevitable need to “push the envelope” on operability, durability, combustion, and emissions requires tools to support these more aggressive designs to resolve unsteady multi-phase flow, turbulent mixing and chemical reaction. They suggest LES represents a flexible modeling approach that provides continuously-improved fidelity as the available computational resources increase. This paper emphasizes the pressing industrial need for increased computing capabilities to enable the use of higher-fidelity design tools like LES with great potential to enhance industrial competitiveness and off-set expensive rig testing.

- [8] Energy and water development appropriations bill, 2012. *112th U.S. Congress*, 1st Session(Report #112-75), 2012.

The introductory paragraph of 2012 Senate Energy and Water Development Appropriations Bill in the section entitled *Title III, Department of Energy* states:

The Committee’s highest priority is accelerating breakthroughs in clean energy technologies to reduce the Nations dependence on foreign oil and developing carbon-free sources of energy that will change the way the United States produces and consumes energy.

The following paragraph, titled “Exascale Initiative”, makes similar remarks concerning the development of exascale computing by 2018:

The Committee supports the Departments initiative to develop exascale computing—1,000 times more powerful than todays most powerful computer. The Committee understands that the path to exascale computing will be extremely challenging and will require significant research and development breakthroughs. For example, an exaflop system made entirely out of todays technology would probably cost \$100,000,000,000, require \$1,000,000,000 a year to operate, need its own dedicated power plant to power the computing system, and be very unreliable. Despite these challenges, the Department has set an ambitious goal of 2018 to deploy the first exascale system.

The Committee also reiterates three “clear priorities” for basic scientific research, one of which is “the development and deployment of more powerful computing capabilities to take advantage of modeling and simulation to advance energy technologies and maintain U.S. economic competitiveness.” The Committee also “understands that exascale computing will help maintain U.S. industrial competitiveness. In particular, high-tech industries such as transportation, aerospace, nuclear energy, and petroleum will increasingly rely on high-performance computing, especially when traditional experiments would be impossible, dangerous, or inordinately costly to perform.”

- [9] F.T. Johnson, E.N. Tinoco, and N. Jong Yu. Thirty years of development and application of CFD at Boeing Commercial Airplanes, Seattle. *Computers & Fluids*, 34(10):1115–1151.

Describes the evolution of CFD use in the product development cycle at Boeing Co. Airplanes in Seattle, Washington. Emphasis on the costs, scheduling, trade-offs, and the phases of CFD development in industry. In particular, the paper addresses the translation of CFD as an academic curiosity to a full partner with established tools in the design of cost-effective and high-performing commercial transports. Practical aspects of the balance between design and optimization, coupling multiple engineering codes with different levels of approximation, and some motivating factors behind industrial involvement in computing. Figure 3 in this paper was used in the first few slides of the presentation to show the effect of CFD on the design cycle by reducing wind tunnel testing of multiple prototype wings. The same figure has also appeared in congressional testimony [10] and many unpublished company presentations.

- [10] M. Garrett. Testimony of Michael Garrett Director, Airplane Performance Boeing Commercial Airplanes. *Before the United States Senate Committee on Commerce, Science and Transportation Subcommittee on Technology, Innovation, and Competitiveness*, July 19, 2006.

Applications of CFD and HPC at Boeing including aircraft noise reduction, wing design, and cabin ventilation are discussed. In the future, Boeing would prefer to “test only two or three wings in wind tunnels versus the 11 for the 787” (see figure 1). Instead of developing a new airplane once a decade, Boeing expects to develop aircraft in significantly less time using HPC. Computation of the acoustic signature of an airplane during takeoff and engine emissions are suggested grand challenge applications. Key goals of HPC are to reduce overall design cycle times and validate new technology at lower overall cost. The primary contribution of HPC at Boeing has been in technology validation (e.g. chevrons, see figure 17) and its application in new product development.

- [11] C.D. Pierce and P. Moin. Progress-variable approach for large-eddy simulation of non-premixed turbulent combustion. *Journal of Fluid Mechanics*, 504:73–97, 2004.

This paper describes the first large-eddy simulation of turbulent combustion in a realistic combustor configuration. Rather than solving transport equations for each individual species in a typical chemical mechanism and modelling the unclosed chemical source terms, the authors instead use an indirect mapping approach whereby all of the detailed chemical processes are mapped to a reduced system of tracking scalars: a mixture fraction variable, which tracks the mixing of fuel and oxidizer, and a progress variable, which tracks the global extent of reaction of the local mixture. The mapping functions, which describe all of the detailed chemical processes with respect to the tracking variables, are determined by solving quasi-steady diffusion-reaction equations with complex chemical kinetics and multicomponent mass diffusion. The performance of the new model is compared to fast-chemistry and steady-flamelet models for predicting velocity, species concentration, and temperature fields in a methane-fueled coaxial jet combustor. Comparisons with available experimental data are presented. The authors show how the progress-variable approach captures the unsteady, lifted flame dynamics observed in good agreement with the experiment; in contrast, fast-chemistry and steady-flamelet models predict an attached flame. Recent extensions of the flamelet progress variable approach to supersonic combustion regimes is described in [50].

- [12] P. Moin. Advances in large eddy simulation methodology for complex flows. *International Journal of Heat and Fluid Flow*, 23(5):710–720, 2002.

This review paper focuses on the derivation of the constitutive equations for large-eddy simulation, subgrid scale modeling, and wall modeling with application of LES applied to turbulent combustion. Simulations on an unstructured mesh of a sector of a combustor in an operational gas turbine are presented. Immersed boundary methods and applications to vehicle-drag reduction and tip-clearance flow in a stator-rotor combination are highlighted. Accurate and robust on LES on unstructured grids derives from the use of

non-dissipative differencing schemes and ensuring conservation of higher order conservation principles. The importance of discretely conserving kinetic energy is illustrated by plotting the total kinetic energy as a function of time for the simulations of the Taylor problem at high Reynolds numbers. The energy conserving scheme is robust while a non-dissipative scheme that only conserves momentum blows up after some time. Even in the inviscid limit, or at very high Reynolds numbers where the dissipative scales are not resolved, the numerical simulation conserves energy.

- [13] G. Iaccarino. Predictions of a turbulent separated flow using commercial CFD codes. *Journal of Fluids Engineering*, 123(4):819–828, 2001.

Numerical simulations of the turbulent flow in an asymmetric diffuser are compared with three commercial CFD codes each using two different turbulence models. Barriers and challenges to using commercial CFD codes are presented. In particular, the challenge of choosing between the many different physical and numerical models is addressed. Uncertainties associated with (a) different computational grids, (b) boundary conditions, (c) convergence and (d) numerical schemes suggests, even for relatively simple, controlled problems, commodity CFD leaves much to be desired. This paper is a brief but succinct description of some of the common pitfalls and advantages of commodity CFD.

- [14] P. Moin and G. Iaccarino. Complex effects in large eddy simulations. In S.C. Kassinos, C.A. Langer, G. Iaccarino, P. Moin, T.J. Barth, M. Griebel, D.E. Keyes, R.M. Nieminen, D. Roose, and T. Schlick, editors, *Complex Effects in Large Eddy Simulations*, volume 56 of *Lecture Notes in Computational Science and Engineering*, pages 1–14. Springer, 2007.

Several examples of advances in LES methodology applied to complex-geometry, multi-physics applications are presented, e.g. the integrated jet engine simulations shown in figure 2. Numerical methods for non-dissipative, energy-conserving methods on unstructured grids are discussed. Wall modeling approaches for LES of complex engineering flows are documented. Novel applications of LES to aero-optics, aero-acoustics, two-phase flow, optimization and control are presented.

- [15] L. Shunn, F. Ham, and P. Moin. Verification of variable-density flow solvers using manufactured solutions. *Journal of Computational Physics*, 231(9):3801–3827, 2012.

Code verification is a procedure by which a particular implementation of a numerical algorithm is confirmed to have all the correct properties of the desired algorithm. A properly verified code should then be free of programming errors (bugs) that affect the theoretical order-of-accuracy of the numerical algorithm. State-of-the-art simulation codes are complex (e.g. millions of lines

of code and complex, interacting physics modules involving different algorithms and different grid topologies). Verification of codes and solutions are equally complex for such problems, but are essential steps toward building confidence in the predictive capabilities of simulation software. The method of manufactured solutions (MMS) provides a procedure by which we “manufacture” (or assume/generate) an analytical solution (or ansatz), without concern for physical realism (e.g. unphysical slip boundary conditions may be used on walls). We then substitute the ansatz into the governing equations to formally derive a general source term $F(\mathbf{x}, t)$ that would reproduce the chosen ansatz (i.e. our educated guess). If we then run our simulation code with the added forcing term $F(\mathbf{x}, t)$ (and compatible initial/boundary conditions), we should reproduce our ansatz and the theoretical order-of-accuracy by performing a systematic grid refinement study. Essentially, such an approach ensures that exact solutions to the governing equations can be reproduced by the code with the same inputs.

This paper is unique in that it applies this method of manufactured solutions to combustion applications on unstructured grids and addresses several challenging issues including the use of tabulated state properties (i.e. density), mesh type and the effect of sub-iterations in the time-advancement scheme on the convergence and accuracy of a multi-physics reactive turbulent flow. Linear interpolation of the equation-of-state led to numerical fluctuations that were found to impede convergence and reduce accuracy and many outer iterations were necessary to eliminate splitting errors in highly nonlinear combustion problems. A complex balance exists between factors such as the size of the tabulated equation of state table, time step, grid size and outer iterations for a given problem. Any of these effects can potentially degrade the theoretical accuracy of the solver and the optimal choice of settings is a problem-dependent task.

- [16] G. Iaccarino, R. Pecnik, J. Glimm, and D. Sharp. A QMU approach for characterizing the operability limits of air-breathing hypersonic vehicles. *Reliability Engineering and System Safety*, 96(9):1150 – 1160, 2011.

Quantifying uncertainty, margin to failure, risk or risk mitigation is important in many industrial problems to establish confidence that what you’ve built will function as expected, be safe, be reliable, among other qualities. Many approaches and perspectives on how best to accomplish this exist. This paper applies on such framework to investigate the operability limits of a supersonic combustion engine for an air-breathing hypersonic vehicle. The time-dependent compressible flow equations with heat release are solved in a simplified configuration. Verification, calibration and validation are carried out to assess the ability of the model to reproduce the flow/thermal interactions that occur when the engine unstarts due to thermal choking. Quan-

tification of margins and uncertainty (QMU) is used to determine the safe operation region for a range of fuel flow rates and combustor geometries.

- [17] P. Moin, K. Squires, W. Cabot, and S. Lee. A dynamic subgrid-scale model for compressible turbulence and scalar transport. *Physics of Fluids*, 3(11), 1991.

The dynamic subgrid-scale model of Germano *et al.* [55] is generalized for the large eddy simulation (LES) of compressible flows and transport of a scalar. The mixing of scalars is a classical turbulence problem with applications in turbulent combustion, multiphase flows and environment flows dealing with stable and unstable stratification as well as pollutant transport in the wind. The application of large-eddy simulation to chemically reacting flows was facilitated by this work. In many turbulent flows, the transport of mass, momentum and energy is often driven by the large, resolved scales so that the eddy diffusivity (Smagorinsky) model coupled to a dynamic procedure as used in this paper often works well. In some flows, however, scalar transport may be dominated by sub-Kolmogorov transport physics (e.g. high-Schmidt number mixing characteristic of heat transport in the ocean) and additional subgrid-model fidelity may be required. The dynamic procedure applied to compressible turbulence and scalar transport outlined in this paper would remain largely the same but further subgrid model development is needed.

- [18] G. Tryggvason, R. Scardovelli, and S. Zaleski. *Direct numerical simulations of gas-liquid multiphase flows*. Cambridge University Press, 2011.

Many industrial applications involve gas-liquid multiphase flow. Heat transfer by boiling is commonly used in both conventional, solar and nuclear power plants (as well as the common kitchen stove). Liquid gasoline must evaporate in an oxygen-rich air mixture before it burns in car engines. For efficient combustion, the liquid is atomized into many small droplets that increase the net surface area thereby increasing the evaporation rate. Several other applications with more complex physics, e.g. mass transfer, chemical reactions, and solidification, are also discussed and numerical methods able to directly resolve such flows are discussed in this book. Particular emphasis is placed on the “one-fluid” formulation wherein a single set of equations involving the balance of mass, momenta and energy are used to describe the entire flow field with variable material properties and material interface physics, such as surface tension, treated as localized source terms in these equations.

- [19] O. Pauluis and J. Schumacher. Self-aggregation of clouds in conditionally unstable moist convection. *Proceedings of the National Academy of Sciences*, 2011.

Moist Rayleigh-Bénard convection is investigated using a Boussinesq model with simplified thermodynamics for phase transitions. High-resolution, three-dimensional direct numerical simulations were performed for extended layers

with aspect ratios up to 64. The authors find that the transition to convective turbulence depends not only on the amplitude of the initial perturbation of the quiescent equilibrium but also the aspect ratio of the convective domain. In contrast to classical dry Rayleigh-Bénard case, moist convection is highly asymmetric with respect to the vertical direction. Moist upwelling air inside turbulent cloud aggregates is surrounded by ambient regions of slowly descending unsaturated air. It is also found that conditionally unstable moist convection is inefficient at transporting energy. The authors find that the net heat transfer (Nusselt number) in moist convection is lower than that of classical dry convection. The influence of periodic boundary conditions and the need to simulate very large aspect ratio domains to resolve these anisotropic, low-wavenumber flow features is highlighted.

- [20] T. Sayadi, C.W. Hamman, and P. Moin. Fundamental and subharmonic transition to turbulence in zero-pressure-gradient flat-plate boundary layers. *arXiv:1110.3986v1 [physics.flu-dyn]*, pages 1–9, 2011.

Simulations of transition to turbulence in compressible (free-stream Mach number $M_\infty = 0.2$), zero-pressure-gradient flat-plate boundary layers triggered by fundamental (Klebanoff K-type) and subharmonic (Herbert H-type) secondary instabilities of Tollmien-Schlichting waves are highlighted in this paper. The accompanying fluid dynamics video won a 2011 Milton Van Dyke Award for its scientific merit, originality and artistry/aesthetic appeal in the video visualization category (click here [high-resolution] or here [low-resolution] to watch the video). Over a billion grid points (i.e. $4096 \times 550 \times 512$ points) were used in each simulation. Simulations were performed on both the BlueGene/L system at Lawrence Livermore National Laboratory and the BlueGene/P system at the Argonne Leadership Computing Facility on 32k cores for most production runs.

- [21] T. Sayadi, C.W. Hamman, and P. Moin. Direct numerical simulation of complete transition to turbulence via H-type and K-type secondary mechanisms. *Journal of Fluid Mechanics (Under Consideration)*, 2012.

Transition in wall-bounded flows is sensitive to the type of disturbance. In this paper, we examine direct numerical simulation databases of two boundary layers from laminar (Blasius $Re_\theta = 210$) inflow that transitions into fully turbulent flow via two different physical disturbance transition scenarios:

1. Fundamental (Klebanoff K-Type) transition ($Re_{\theta_{\max}} = 1410$), and
2. Subharmonic (Herbert H-Type) transition ($Re_{\theta_{\max}} = 1250$).

Visualizations of K-Type transition are shown in figure 3. These DNS databases provided validation data used in the development of a dynamic subgrid scale LES modeling procedure that predicts the location of transition from laminar to turbulent flow irrespective of the path to transition; the

RANS approach and constant coefficient eddy viscosity models are sensitive to the type of transition, but dynamic models are able to predict the particular transition scenario without tuning or specification. This simulation database has also enabled further studies of turbulence flow structures and development of reduced-order models of the transition mechanisms. Simulations were performed on both the BlueGene/L system at Lawrence Livermore National Laboratory and the BlueGene/P system at the Argonne Leadership Computing Facility on 32k cores for most production runs.

- [22] P.R. Spalart and J.D. McLean. Drag reduction: enticing turbulence, and then an industry. *Phil. Trans. R. Soc. A*, 369(1940):1556–1569, 2011.

Drag-reduction methods are reviewed with an emphasis on how to bridge the gap between pure science through engineering to what makes these inventions go into service in industry. The practical factors that limit deployment of technology concepts that have shown significant drag reduction potential in computations for many years are discussed.

- [23] X. Wu and P. Moin. Direct numerical simulation of turbulence in a nominally zero-pressure-gradient flat-plate boundary layer. *Journal of Fluid Mechanics*, 630:5–41, 2009.

First direct evidence (in the form of a direct numerical solution of the Navier-Stokes equations) where hairpin packets are seen to appear. Turbulent flows near walls such as this are a major focus of engineering research for their importance in determining the lift of an aircraft wing or drag on a car. This paper contributed significantly towards our current understanding of such flows and is also highlighted in the Focus on Fluids article by Marusic [40] and appeared on the front cover of Volume 630 of the Journal of Fluid Mechanics.

- [24] P. Moin and K. Mahesh. Direct numerical simulation: A tool in turbulence research. *Annual Review of Fluid Mechanics*, 30:539–578, 1998.

Review of direct numerical simulation that emphasizes the use of DNS as a research tool and not a brute-force solution to the Navier-Stokes equations for engineering problems. The wide range of scales in turbulent flows requires careful attention to their numerical solution, boundary conditions, and spatial-temporal discretization. Scientific insight into turbulence physics obtained from DNS of certain idealized flows are highlighted and the important achievements for modeling, control and evaluation of new concepts in turbulence are highlighted. Parallel computing is found to be “the enabling technology for the next generation of archivable simulations.”

- [25] X. Wu and P. Durbin. Evidence of longitudinal vortices evolved from distorted wakes in a turbine passage. *Journal of Fluid Mechanics*, 446:199–228, 2001.

Direct numerical simulations of turbulent wakes swept past the inlet of a low-pressure turbine cascade were investigated. About 57 million grid points were used and simulated on the one teraflop Nirvana Blue machine at Los Alamos c. 1999. Nirvana Blue was composed of several shared memory multiprocessor SGI Origin 2000 machines with 128 processors on each node (similar to the LANL ASC Blue Mountain machine). Only one node was used for this particular simulation. As a result, OpenMP was used for parallelization (decomposed along the homogeneous spanwise dimension). Intense streamwise vortices, descending from the passing wakes were observed on the pressure surface. In high-pressure turbines, such vortices are known to enhance heat transfer causing the turbine pressure surface to run hotter by several hundred degrees Fahrenheit than the suction surface. While the DNS of low-pressure turbines is feasible, this simulation provided much needed insight into this phenomenon in higher Reynolds number turbomachinery flows. More recently, the lead author has carried out a more comprehensive study of the effects of passing wakes in triggering transition to turbulence in boundary layers [23].

- [26] S.A. Orszag and G.S. Patterson Jr. Numerical simulation of three-dimensional isotropic turbulence. *Physical Review Letters*, 28:76–79, 1972.

DNS begins with this paper and much of the foundation developed when the authors were at the National Center for Atmospheric Research (an exceptionally strong turbulence research and simulation group composed of people like J.W. Deardorff [54], D.K. Lilly [53], J. Smagorinsky [79], A. Arakawa [78], and W. Washington who were at NCAR around this time and the preceding decade(s), see Fox & Lilly [41] for a review). Since that time, computer capacity has increased by several orders of magnitude and the number of computers (i.e. workers) in the field has also grown. Much progress has been made in the efficiency and accuracy of computational algorithms for turbulence, but the primary pacing item remains the speed and memory size of computing hardware [2]. This early DNS paper discusses simulations of homogeneous, isotropic turbulence at Reynolds number $R_\lambda = 35$ with 32^3 grid points.

- [27] R.S. Rogallo. Numerical experiments in homogeneous turbulence. *NASA-TM*, 81315, 1981.

Simulations of homogeneous turbulence in an incompressible fluid subject to uniform deformation or rotation are considered. The simulations (which actually began around 1977, e.g. see the earlier NASA Technical Memorandum titled “An ILLIAC program for the numerical simulation of homogeneous turbulence” by R.S. Rogallo Nov. 1997) were quite comprehensive and covered a wide range of different turbulent flow configurations. The simulations used 128^3 grid points. This may not seem like much today, but, at the time, every word of available memory was needed for such a simulation. Rogallo made

very meticulous use of the what the ILLIAC-IV computing platform had to offer. Constrained by computer memory size, he had to reduce the number of words per node to a minimum. Memory constraints are expected to become increasingly tight on future exascale computing platforms; algorithms that reduce memory usage, as Rogallo showed, are helpful here.

Several algorithm changes and storage patterns to this effect are documented in this paper and the earlier 1977 NASA Technical Memorandum (which, in turn, had to be entered on punch card decks and fed into a reader). In order to use periodic boundary conditions with Fourier-spectral approximations in each direction, Rogallo solved the equations in a time-varying, transformed coordinate system that moved with the mean velocity. From time to time (at a frequency $\sim \frac{1}{2}dU/dz$), the code would re-mesh the flow field to avoid numerical deterioration; to prevent unwanted aliasing when re-meshing, Rogallo had to develop novel dealiasing techniques. In addition, he showed that, for a two-step time integrator, one can apply a phase shift to the velocity field at negligible extra cost to reduce aliasing errors (the more costly 3/2-rule, where the number of modes in each direction are increased by 3/2, requires about three times as much computation and more memory storage, when computing non-linear products is then avoided). This is particularly important in DNS that are only partially well-resolved (since we are often willing to accept some error in order to obtain a higher Reynolds number).

- [28] M. Rogers and P. Moin. The structure of the vorticity field in homogeneous turbulent flows. *Journal of Fluid Mechanics*, 176:33–66, 1997.

The vorticity fields are examined in homogeneous turbulent shear flow by direct numerical simulation of the unsteady, incompressible Navier-Stokes equations with up to 128^3 grid points; $Re_\lambda = 72.6$. They found that vortex filaments in homogeneous shear flow are strongly anisotropic and aligned with the major principal axis of the stress tensor. This is in contrast to the vorticity field in homogeneous turbulence without shear, which exhibits no such preferential alignment. This fundamental direct numerical simulation motivated further development into the structure of vorticity in shear flows such as the boundary layer simulations of Wu & Moin [23], which also observed similar phenomena. The authors found soon discovered that this anisotropy appears in a wider class of turbulent shear flows, which has influenced the development of turbulence models. The authors also used the code originally developed by Bob Rogallo [27] for the simulations with essentially the same grid. All simulations were done on the Cray X-MP at NASA-Ames Research Center. At the time, the Cray X-MP had a Solid-State Storage System (SSD) that had about 256 MB of storage space in addition to about 16 MB of high-speed, random-access main memory. For the 128^3 simulations with three velocity components, pressure and three others scalars (Rogers & Moin were

also interested in developing models for the turbulent scalar flux for arbitrarily imposed mean scalar gradient), one snapshot of the flow field (about 60 MB in single-precision) would simply not fit into main memory; in fact, only about 6% of the data could be loaded into main memory at any given time! As a result, whenever they wanted to compute a derivative, the program could only operate on few lines of data at a time as it had to be repetitively moved from the SSD to main memory, operations done on the subset, and then sent back to the SSD. In the end, the simulations took about 45 seconds wall-clock time per timestep in total (20 seconds to compute the hydrodynamic field and 25 seconds to simulate the three passive scalar fields). Integrations were carried out for about 900 to 1400 timesteps for each simulation before the computational mesh became too distorted to resolve the turbulence (since the grid was adaptively deformed as time advanced in order to account for the mean shear in a numerically convenient way). In recent years, it is not unusual for large direct numerical simulations to integrate for a few hundred thousand timesteps (for improved statistical sampling and as a consequence of the smaller timescales one must resolve at higher Reynolds).

- [29] J. Kim, P. Moin, and R. Moser. Turbulence statistics in fully developed channel flow at low Reynolds number. *Journal of Fluid Mechanics*, 177:133–166, 1987.

Landmark direct numerical simulation of fully developed turbulent channel flow. The numerical method used is described by P. Moin and J. Kim [85]. A mesh with $192 \times 129 \times 160$ ($\sim 4 \times 10^6$) mesh points was used to achieve a Reynolds number of 3300, based on the centerline velocity and channel half-width, and friction Reynolds number $Re_\tau = 180$. About 250 CPU hours on a Cray X-MP machine were used to simulate about 10 non-dimensional times with each timestep taking about 40 seconds. A fully spectral method, with Fourier series in the homogeneous direction and Chebyshev polynomials in the wall-normal direction to calculate spatial derivatives, and a second-order time advancement scheme are used. Comparisons with experiments are presented. A direct numerical simulation of a turbulent channel flow where all essential scales of motion are resolved was performed. Many of the statistical correlations for turbulent channel flow, which were found to complement the existing experimental data, were reported for the first time. More recently, Hoyas and Jiménez [36] have simulated an $Re_\tau = 2003$ channel.

- [30] P.R. Spalart. Direct simulation of a turbulent boundary layer up to $R_\theta = 1410$. *Journal of Fluid Mechanics*, 187:61–98, 1988.

The turbulent zero-pressure-gradient flat-plate boundary layer is simulated numerically between $R_\theta = 225$ and $R_\theta = 1410$. The author used a spectral method to solve the three-dimensional, time-dependent Navier-Stokes equations with up to about ten million grid points. A scaling procedure is used

to approximate the slow streamwise growth of the boundary layer and avoid the difficulty of inflow conditions. Several more recent direct numerical simulations of zero-pressure-gradient flat-plate boundary layers have addressed many of these difficulties and are free of such approximations, e.g. [23, 20, 21]

- [31] J. Jiménez, A. Wray, P.G. Saffman, and R.S. Rogallo. The structure of intense vorticity in isotropic turbulence. *Journal of Fluid Mechanics*, 255:65–90, 1992.

A comprehensive study of the celebrated tube-like intense vortical structures in homogeneous turbulence at several Reynolds numbers was carried out as part of the 1992 Center for Turbulence Research Summer Program. The diameter of the tubes scale with the Kolmogorov scale η and their lengths scale with the integral scale ℓ . These structures, also known as “worms”, had prior to this work only been observed in forced isotropic turbulence calculations, there was some concern that they may be artifacts of the forcing. The authors presented evidence that the worms are robust and occur without forcing. A novel numerical experiment was performed to investigate the influence of such worm-like structures. The authors artificially removed the vortex tubes from their solution and switched off the large-scale forcing. The subsequent evolution of the flow without the “worms” was compared to a similar decaying turbulence simulation where the tubes were retained. The decay rate of the kinetic energy was found to be identical in both computations. The worms occupy a smaller fraction of the volume of the flow with increasing Reynolds number and are primarily associated with intense events found in the tails of the probability distribution functions of velocity gradients that become more intense with increasing Reynolds number. The simulations were carried out on the 128-compute-node Intel iPSC/860 hypercube at NASA-Ames and the newly installed (and one-of-its-kind prototype) 512-compute-node Intel Touchstone DELTA machine at Caltech. In particular, the Delta machine was fully utilized for the very large simulations of forced isotropic turbulence with 512^3 degrees of freedom described in this paper (i.e. each compute node applied a Fast Fourier Transform on its $1 \times 512 \times 512$ plane of data and then transposed data amongst all other nodes to compute along the other direction). An earlier version of this paper is available in the proceedings of the 1992 CTR Summer Program pp. 21–45.

- [32] S. Lee, S.K. Lele, and P. Moin. Direct numerical simulation of isotropic turbulence interacting with a weak shock wave. *Journal of Fluid Mechanics*, 251:533–562, 1993.

Understanding how turbulence interacts with a shock wave is of great importance in many complex flows. This paper describes one of the first simulations of such shock-turbulence interactions. The interaction of isotropic quasi-incompressible turbulence with a weak shock wave was studied by direct numerical simulation with as many as $193 \times 64^2 \approx 0.8 \times 10^6$ grid points

. Upon interacting with a weak shock wave, the turbulence was found to become stronger as evidenced by the amplification of turbulent kinetic energy and transverse vorticity component and smaller turbulent length-scale found downstream of the shock. The shock wave interface was found to become distorted by the upstream turbulence; at sufficiently strong turbulence levels, the shock front distortions would fluctuate rapidly and create regions (or small holes) of locally non-shocked flow. Such effects of shock strength and turbulence intensity were discussed in a follow-up paper by the same authors titled “Interaction of isotropic turbulence with shock wave: Effects of shock strength” in the *Journal of Fluid Mechanics* Volume 340, pp. 225 (1997)

- [33] H. Le, P. Moin, and J. Kim. Direct numerical simulation of turbulent flow over a backward-facing step. *Journal of Fluid Mechanics*, 330:349–374, 1997.

Turbulent flow over a backward-facing step with expansion ratio 1.20 and Reynolds number 5100 based on step height and inlet free-stream velocity was computed. A turbulent boundary layer enters from the left boundary, separates at the step ($x = 0$) and then reattaches downstream (at $x \approx 7$ step heights). About 8.3 million points (768 cells in the streamwise direction, 64 in the homogeneous spanwise direction and 192 in the vertical direction; grid points in the solid step are ignored) on a staggered grid were used. This simulation was one of the first DNS of a wall-bounded, turbulent separation region including flow reattachment. The results of the simulations are in excellent agreement with experimental data from Jovic & Driver (1994) and includes Reynolds-stress budgets. The simulations were done a few years before this paper; much of which is documented in Dr. H. Le’s thesis and in the TF-58 report of the Thermosciences Division, Dept. of Mech. Engineering, Stanford University by H. Le and P. Moin (1994). The simulations took about 54 days of time on a Cray C-90 to integrate for about 2.1×10^5 timesteps).

- [34] J.B. Freund, S.K. Lele, and P. Moin. Compressibility effects in a turbulent annular mixing layer. Part 1. Turbulence and growth rate. *Journal of Fluid Mechanics*, 421:229–267, 2000.

Compressible turbulent mixing occurs in many technological applications. In scramjet combustion, for example, it is desirable to inject the fuel stream parallel to the air stream so as to maximize thrust efficiency and minimize momentum losses that occur when fuel injection is made perpendicular to the free stream. Mixing between parallel streams of fluid at high Mach numbers, however, is significantly suppressed. Understanding the fundamental physics that of such mixing layers is important develop techniques to enhance mixing between parallel flows while incurring minimal losses. This paper investigates the fundamental properties of compressible turbulent free shear flow by direct numerical simulation of a turbulent annular mixing layer over a range of Mach

numbers. See also “Part 2. Mixing of a passive scale” (in the same JFM volume) and the follow-up direct numerical simulations (which used the same code as this one) of a Mach 1.92 jet and its sound field by the same authors in the AIAA Journal volume 38 number 11. For reference, the simulations were done on the IBM SP, the Cray T3D and the Cray T3E usually with 64 processors and all post-processing was done on a Cray C90 (because individual IBM SP processors did not have enough memory to load the entire flow fields). Parallel domain decomposition was done along a single coordinate: either the periodic streamwise or azimuthal directions where Fast Fourier Transforms were used to compute the corresponding derivatives.

- [35] M. Yokokawa, K. Itakura, A. Uno, T. Ishihara, and Y. Kaneda. 16.4-Tflops Direct Numerical Simulation of turbulence by a Fourier spectral method on the Earth simulator. *Supercomputing, ACM/IEEE Conference*, 2002.

High-resolution DNS of incompressible turbulence with as many as 4096^3 grid points using a Fourier spectral method (single-precision arithmetic). Most of the computation time (about 90%) was spent in calculating the large three-dimensional Fast Fourier Transforms involved in the simulation. Methods to accelerate the FFT computation and distributed parallelization are discussed. Simulations were performed on the Earth Simulator (fastest supercomputer in the world from 2002 to 2004).

- [36] S. Hoyas and J. Jiménez. Scaling of the velocity fluctuations in turbulent channels up to $Re_\tau = 2003$. *Physics of Fluids*, 18(1), 2006.

This paper was the first reported simulation of a turbulent channel flow beyond $Re_\tau = 2000$. At this point, the authors were just beginning to go about analyzing the copious quantities of data, some of which is available at <http://torroja.dmt.upm.es/ftp/channels>. After a year or two, several scientific publications by many different research groups concerning the fundamental physics of turbulence near walls made use of this data.

- [37] X. Wu and P. Moin. A direct numerical simulation study on the mean velocity characteristics in turbulent pipe flow. *Journal of Fluid Mechanics*, 608:81–112, 2008.

Fully developed incompressible turbulent pipe flow at bulk-velocity- and pipe-diameter-based Reynolds number $Re_D = 44,000$ was simulated with second-order finite-difference methods on 630 million grid points.

- [38] J. Larsson and S.K. Lele. Direct numerical simulation of canonical shock/turbulence interaction. *Physics of Fluids*, 21:126101, 2009.

Direct numerical simulation of isotropic turbulence passing through a nominally normal shock wave are described. This paper builds upon the studies of

[32] and, in particular, investigates the “wrinkled” and “broken” shock regime in greater detail.

- [39] S.K. Lele and J. Larsson. Shock-turbulence interaction: What we know and what we can learn from peta-scale simulations. *Journal of Physics: Conference Series*, 180(1):012032, 2009.

Many applications in engineering and physical sciences involve turbulent flows interacting with shock waves, e.g. high-speed aerodynamic flows and propulsion systems. Supernova explosions and implosion of cryogenic fuel pellets for inertial confinement fusion also involve the interaction of shock waves with turbulence and strong density variations. Numerical simulations of such physical phenomena impose conflicting demands on the numerical algorithms. Capturing broadband spatial and temporal variations in a turbulent flow suggests the use of high-bandwidth schemes with minimal dissipation and dispersion, while capturing the flow discontinuity at a shock wave requires numerical dissipation. Estimates for the computational resources necessary for studying this fundamental shock-turbulence interaction problem at higher Reynolds number on peta-scale computing systems are given [38].

- [40] I. Marusic. Unraveling turbulence near walls. *Journal of Fluid Mechanics*, 630:1–4, 2009.

Brief summary article on the direct numerical simulation of Wu & Moin [23] of turbulent flow near walls.

- [41] D.G. Fox and D.K. Lilly. Numerical simulation of turbulent flows. *Reviews of Geophysics*, 10(1):51–72, 1972.

Review paper that describes the early beginning of large-eddy simulation, especially at the National Center for Atmospheric Research. At the time, direct numerical simulation was only practical for two-dimensional or marginally turbulent three-dimensional flows. The passage of time has allowed for slightly more than marginal simulations of turbulent flow (see table 3), but computational resources are always a fundamental limitation for DNS as discussed in §3. They say “A more useful approach for practical applications is to directly simulate the larger scales of motion and only consider the small unresolved scales with respect to their gross statistical interactions with the larger scale.” Variable-eddy viscosity (Smagorinsky-type) approaches are discussed. The issue of truncation and aliasing errors when selecting a particular discretization or numerical integration technique are also addressed.

- [42] D. Kim and P. Moin. Direct numerical study of air layer drag reduction phenomenon over a backward-facing step. *Annual Research Briefs, Center for Turbulence Research*, 2010.

CTR Annual Research Briefs are located here:

<http://www.stanford.edu/group/ctr/publications.html>

- [43] S. Balachandar and J.K. Eaton. Turbulent dispersed multiphase flow. *Annual Review of Fluid Mechanics*, 42:111–133, 2009.

Turbulent dispersed multiphase flows occur in many engineering and environmental applications (e.g. coal gasifiers, bio-reactors, soot in combustion systems, and ice-crystal growth in clouds). The increased range of scales and stochastic nature of the dispersed-phase particle distribution embedded in a turbulent flow makes this problem significantly more complex for both experiment and simulation than single-phase turbulence. The strengths, limitations and opportunities for experimental and computational techniques are reviewed. In particular, this review article addresses preferential concentration of particles, droplets and bubbles; coupling mechanisms between turbulence and the dispersed phase; and how microscopic particle/bubble interactions change the behavior of the macroscopic turbulent flow.

- [44] J. Jiménez. Computing high-Reynolds-number turbulence: will simulations ever replace experiments? *Journal of Turbulence*, 4:111–133, 2003.

It is often claimed that experiments can be run at higher Reynolds numbers than simulations. The author argues that the Reynolds number (and cost) difference between laboratory and computation has been steadily eroding with the advances in computing technology, and that in many respects both are now comparable. This is particularly true when the quantity of interest is more complicated than turbulence intensities or the mean velocity profile statistics. In the case of turbulent channel flows, the early simulations at $R_\tau = 180$ [29] have been extended by a factor of ten to $R_\tau \approx 2000$ [36] and, in this case, the lack of well-documented laboratory channel data prevents validation of numerical simulation databases with physical experiment. In the near-future, simulations of wall-bounded flows including the logarithmic layer are expected to occur. Careful numerical simulations of the Navier-Stokes equations “are just a different kind of experiment” [24]. Advantages of simulation, such as lack of ambiguity, ability to measure most any quantity of interest, and the ease in which one can perform “thought” experiments, are discussed.

- [45] J.H. Chen. Petascale direct numerical simulation of turbulent combustion – fundamental insights towards predictive models. *Proc. of the Combustion Institute*, 33:99–123, 2011.

The author shows how petascale computing has enabled advances in the simulation of turbulent combustion and predicts what future exascale computing architectures may provide. The paper emphasizes how DNS combined with detailed chemistry models has provided substantial scientific insight into turbulence-chemistry interactions and provided a valuable tool for validation

and development of combustion model necessary for complex fuels and flow scenarios. Several impressive examples involving stabilization of autoignitive lifted turbulent jet flames, a reactive jet in crossflow, premixed flame propagation, pollutant formation, as well as extinction and reignition in turbulent jet flames. In particular, the examples highlight the need for development of improved finite-rate chemical mechanisms and turbulence models for mixed regimes of combustion. Methods to address the data deluge of large-scale simulations of combustion are also discussed (e.g. in-situ data processing to avoid writing full-field files to disk). The author says “simulation currently can play only a minor role in influencing engine design because the coupling of turbulence with detailed fuel chemistry over multiple engine cycles is beyond both our scientific understanding and current computational resources.” Rather than using the $1000\times$ computational capability from petascale to exascale to provide a $10\times$ increase in Reynolds number, the author appears to suggest that a $100\times$ increase in chemical transport and a $10\times$ increase in chemical mechanism complexity may be a more fruitful use of such resources.

- [46] T. Poinso and D. Veynante. *Theoretical and numerical combustion*. 3rd edition, 2012.

Most fluid flows are turbulent. Most reacting fluid flows, with the notable exception of the common candle flame, are also turbulent. This book provides a comprehensive introduction to the basic numerical methods and physics behind direct, large-eddy, and Reynolds-averaged simulation for a wide array of reacting flows from candle flames to complex geometry burners/combustors found in industrial applications. Numerical stability, boundary conditions, wall-flame interactions and combustion models are also discussed.

- [47] R.O. Fox. Large-eddy-simulations tools for multiphase flows. *Annual Review of Fluid Mechanics*, 44:47–76, 2012.

The large range of length and time scales in turbulent multiphase flows make direct numerical simulation of the microscale governing equations intractable for many applications. At the microscale, boundary layers on bubbles and turbulent bubble wakes can form behind bubbles that are much smaller than a numerical grid even when the macroscopic flow is laminar. This review article addresses different modeling approaches and, in particular, emphasizes a kinetic approach rather than an ensemble-averaging approach to modeling the dispersed phase dynamics.

- [48] C.M. White and M.G. Mungal. Mechanics and prediction of turbulent drag reduction with polymer additives. *Annual Review of Fluid Mechanics*, 40:235–256, 2008.

Review article on polymer drag reduction in wall-bounded turbulent shear flows. Addition of small quantities of high-molecular weight polymers into

turbulent boundary layers is known to reduce drag by as much a 80% compared to flow without polymer injection. Simulations provide direct evidence that polymers disrupt conversion of free-stream momentum into wall shear stress in turbulent boundary layers by directly interacting with and dampening quas-streamwise vortices. These near-wall vortices act to stretch the polymers thereby transferring Turbulent kinetic energy from the vortices to the polymers. In doing so, the self-sustaining cycle of wall turbulence is disrupted and turbulent skin friction (drag) is reduced. The need to optimize polymer drag reduction and develop models for higher Reynolds number flows in complex geometries are discussed. Matching polymer relaxation time and time-scale of quasi-streamwise vortices, accounting for polymer degradation and detailed numerical models of polymer dynamics are suggested as future research areas.

- [49] D. Richter, E.S.G. Shaqfeh, and G. Iaccarino. Numerical simulation of polymer injection in turbulent flow past a circular cylinder. *Journal of Fluids Engineering*, 133(10):104501, 2011.

Injection of polymers within realistic range found to stabilize the turbulent wake behind a circular cylinder. Instead of a detached turbulent shear layer and low-pressure chaotic vortex behind the cylinder, a more coherent primary vortex with increased core pressure and reduced turbulent kinetic energy is shed when a viscoelastic polymer solution is injected from the upstream face of the cylinder. Drag reduction was also observed. Applications of polymer injection beyond drag reduction include suppression of propeller tip vortex cavitation. The base cylinder flow and Reynolds number ($Re_D = 3900$) are the same as that studied in [86] where the effect of hydrophobic surfaces instead of polymer injection was investigated. The primary difference between these studies is that, in polymer injection, polymers are transported with the flow and help stabilize the near-wake region whereas hydrophobic surfaces were found to help delay flow separation but could not directly influence the subsequent wake dynamics.

- [50] A. Saghafian, V.E. Terrapon, F. Ham, and H. Pitsch. An efficient flamelet-based combustion model for supersonic flows. *AIAA-2267*, 2009.

A Flamelet/Progress Variable [11] combustion model suitable for supersonic flows is introduced and allows the use of complex chemical mechanisms. A direct simulation of a temporal mixing layer with a detailed chemical mechanisms for hydrogen/oxygen combustion (9 species and 29 reactions) was initiated by the lead author (see figures 9 and 10) so as to validate such flamelet models and evaluate the effects of heat release and mixture composition in such supersonic flows. See also “Direct numerical simulation of compressible

reacting flow” by A. Saghafian and H. Pitsch under consideration for publication in the Journal of Fluid Mechanics (2012) for more details of the detailed numerical simulation.

- [51] S.T. Bose, P. Moin, and D. You. Grid-independent large-eddy simulation using explicit filtering. *Physics of Fluids*, 22(10), 2010.

The governing equations for large-eddy simulation are derived from the application of a low-pass filter to the Navier-Stokes equations. Most common LES approaches assume that discrete operations performed on a particular grid act as an implicit filter, causing results to be sensitive to the mesh resolution. Explicit filtering instead provides a methodology to formally separate the filtering operation, and hence the resolved turbulence, from the underlying mesh distribution alleviating the grid dependent solutions upon further grid refinement. The authors apply explicit filtering in large-eddy simulation to obtain numerical solutions for turbulent channel flow that are grid independent. Recent simulations of a three-dimensional diffuser using this framework applied to unstructured grids can be found in S.T. Bose, P. Moin and F. Ham, “Explicitly filtered large eddy simulation on unstructured grids”, CTR Annual Research Briefs, 2011.

- [52] P.R. Spalart. Strategies for turbulence modelling and simulations. *International Journal of Heat and Fluid Flow*, 21(3):252 – 263, 2000.

Evaluates the possibilities for the numerical prediction of a turbulent flow and, specifically, what is needed to target a complete airplane, turbine or car. Evaluates the solution strategies range from Reynolds-Averaged Navier-Stokes (RANS) equations to Direct Numerical Simulation (DNS), with Large-Eddy Simulation in between. The many intermediate strategies such as VLES, URANS and DES are discussed and compared with conventional RANS. Spalart argues that, for many decades, practical methods such as RANS will be necessary, possibly unsteady, or RANS/LES hybrids (i.e. wall-modeled LES), pure LES being unaffordable. He claims that the empirical content of such solution strategies will remain high and standard turbulence modeling concepts (e.g. law of the wall, mixing-length models, Van-Driest, etc.) will remain substantially important. The author suggests that the role of grid refinement is now physical instead of numerical (i.e. meaning that grid refinement weakens the role of the modelled eddies thereby improving the fidelity of the simulation). The merits and drawbacks of grid-independent large-eddy simulation where one does not automatically link the width of the LES filter to the grid spacing are discussed. The author makes predicts that wall-modelled LES of an airliner will be ready in 2045 as a “Grand Challenge” problem requiring about $10^{11.5}$ grid points and $10^{6.7}$ timesteps with industrial use coming several years later.

- [53] D.K. Lilly. The representation of small-scale turbulence in numerical simulation experiments. *NCAR Manuscripts*, 281, 1966.

Emphasizes that the future practicality of turbulence simulations requires development of models for the transport of turbulent kinetic energy into and through the inertial range. This would allow one to simulate energy transfer to scales smaller than the resolving power of the mesh. Early examples of large-eddy simulation closure models applied to real-world problems are given. Later appears as an article in an IBM Scientific Computing Symposium on environmental science (1967), pp. 195–210.

- [54] J.W. Deardorff. A numerical study of three-dimensional turbulent channel flow at large reynolds number. *Journal of Fluid Mechanics*, 41(2), 1970.

For the first time, the three-dimensional Návier-Stokes equation were integrated numerically in time for the case of a turbulent, plane channel flow at very large Reynolds numbers. Only 6720 grid points were used! Sub-grid scale effects were simulated with an eddy viscosity proportional to the local velocity gradient. The constant of proportionality was uniformly set to $C_S \approx 0.1$ with filter width equal to the grid size in this mean shear flow. Note that, before this, Lilly [53] had simulated homogeneous isotropic turbulence but found $C_S \approx 0.23$ gave better results. A couple decades later, dynamic Smagorinsky models began to be developed to address this non-universal nature of the eddy viscosity coefficient [55, 17]. In common parlance, Deardorff’s simulation was a large-eddy simulation (LES) with constant-coefficient Smagorinsky model. Statistics were compared to existing experimental measurements. Good to marginal agreement was found. Deardorff performed many of the pioneering simulations of turbulence (this paper being only one of them). He concludes that “the problem of turbulence at large Reynolds numbers is already profitable, with increased accuracy to be expected with modest increase of numerical simulation.” With only 600 times as many grid points, Kim, Moin & Moser [29] did find significant profit in the simulation of turbulence by DNS of turbulent channel flow. With increased availability of computational resources, we expect this profit to science, industry and society to continue.

- [55] M. Germano, U. Piomelli, P. Moin, and W.H. Cabot. A dynamic subgrid-scale eddy viscosity model. *Physics of Fluids*, 3(7), 1991.

This paper introduces the dynamic procedure to the development of subgrid-scale (SGS) closure models. The same eddy viscosity formulation of Smagorinsky [79] is assumed where the small scales are assumed to be in “local-equilibrium” so that energy production balances viscous dissipation. The key difference of a dynamic SGS model is that the eddy viscosity coefficient is

allowed to vary in both space and time and is dynamically computed using information contained in the resolved turbulence scales based on first-principles physics and is much less sensitive to model parameters; the only adjustable coefficient being the ratio of filter widths, which specifies the scales of resolved motions that contribute most to subgrid-scale turbulent stresses. The dynamic procedure is particularly successful in addressing the limitations of the Smagorinsky model in flows undergoing transition to turbulence, relaminarization and in the viscous near-wall region (e.g. the dynamical model produces the correct limiting behavior of vanishing eddy viscosity near walls while a constant-coefficient Smagorinsky model artificially generates non-zero Reynolds stress in the viscous near-wall region).

- [56] D. You, F. Ham, and P. Moin. Discrete conservation principles in large-eddy simulation with application to separation control over an airfoil. *Physics of Fluids*, 20(10), 2008.

Unstructured LES is used to investigate the external turbulent flow separation over an airfoil with and without synthetic-jet control via a spanwise cavity slot inside the airfoil. The synthetic jet actuation is shown to stabilize the boundary layer by adding/removing momentum to/from the boundary layer and enhancing mixing between the inner and outer parts of the boundary so as to prevent flow separation and stall at high-angles of attack. By careful consideration of the discrete equations, a novel numerical approach that minimizes the non-conservation of kinetic energy due to the pressure term on irregular unstructured skewed meshes is presented. They verify that this approach improves numerical accuracy and stability on such irregular unstructured skewed meshes and validate the method by application to synthetic-jet actuation numerically confirming the experimental observation that such synthetic jets effectively delay the onset of flow separation with about 70% increase in the lift coefficient. Paying attention to such discrete conservation principles is particularly important for predicting subtle separation effects in turbulent boundary layers where energetics play a crucial role.

- [57] P. Moin and M. Wang. Wall modeling for large-eddy simulation of turbulent boundary layers. In G. Meier, K. Sreenivasan, H.-J. Heinemann, and G. M. L. Gladwell, editors, *IUTAM Symposium on One Hundred Years of Boundary Layer Research*, volume 129 of *Solid Mechanics and Its Applications*, pages 269–278. Springer Netherlands, 2006.

Near a no-slip wall, turbulent eddies scale with the distance from the wall and move increasingly closer to the wall as the Reynolds number increases. These eddies are dynamically important and their effects on the large eddies away from the wall must be accounted for in large-eddy simulation. This remains a pacing item for the use of LES in engineering applications. This article surveys methods to combine LES and wall modeling to alleviate the stringent

near-wall resolution requirement (e.g. as discussed in [58]). The authors emphasize methods whereby the dynamics of the near-wall eddies are modeled and coupled to the outer LES flow simulation by imposition of approximate boundary conditions off the wall. With wall modeling, the LES can use a relatively coarse grid such that the overall computational cost is only weakly dependent on the Reynolds number. In this way, simulations of high-Reynolds number turbulent flows typical of engineering application are feasible.

- [58] H. Choi and P. Moin. Grid-point requirements for large eddy simulation: Chapman’s estimates revisited. *Physics of Fluids*, 24(011702), 2002.

In near-wall regions, additional modeling is required for LES. Near-wall modeling has been a significant pacing item for the practical use of LES in high-Reynolds number, complex geometry, industrially-relevant flows. This paper revisits Chapman’s scaling estimates in light of new empirical data for flat plate boundary layers. The important point is that wall-modeled LES allows the computational cost to scale linearly with Reynolds number (i.e. a $1000\times$ increase in computational concurrency would provide a $\mathcal{O}(1000\times)$ increase in the accessible Reynolds number; wall-resolved LES is however not much different from direct numerical simulation in terms of resolution requirements. Resolution requirements for LES, estimated by Chapman [2], are updated using accurate formulae for high Reynolds number boundary layer flow. The new estimates indicate that the number of grid points N required for wall-modeled LES is proportional to R_x , but wall-resolving LES depends almost quadratically as $N \sim R_x^{13/7}$. Direct numerical simulation of wall-bounded turbulent shear flows, on the other hand, requires $N \sim R_x^{37/14}$. The author did not account for the possibility of using nested, anisotropic grids to resolve the turbulent boundary layer; as a result, the results for wall-resolved and direct simulation are conservative estimates.

- [59] R. Verzicco, M. Fatica, G. Iaccarino, P. Moin, and B. Khalighi. Large eddy simulation of a road vehicle with drag-reduction device. *AIAA Journal*, 40(12):2447–2455, 2002.

Flow around an idealized vehicle is simulated using large-eddy simulation with an immersed boundary numerical approach to capture the geometric complexity of the vehicle. Effects of Reynolds number and wake modifications produced by different geometrical models (a cavity and boat-tail attached to the base of the vehicle) are investigated. Overall drag reduction and modifications to the wake structure, including mean velocity profiles, low-frequency axial wake pumping and high-frequency shear-layer instabilities, were analyzed and found to be in good agreement with the measured value from companion experiments carried out by General Motors Corporation.

- [60] R. Mittal and G. Iaccarino. Immersed boundary methods. *Annual Review of Fluid Mechanics*, 37:239–261, 2005.

Meshing or grid generation has long been a pacing item and significant barrier to the widespread use of computational fluid dynamics. In complex geometry configurations such as the flow past a pick-up truck, past flexible filaments, inside moving piston engines, helicopter blades, industrial fan blades and other flows, meshing a body conforming grid can take several days or even months of tedious work. Geometry processing and surface grid generation accounts for a significant fraction of the total time require to run a CFD simulation. When geometries change dynamically during a simulation, re-gridding can be cost-ineffective. To avoid these time sinks, immersed boundary methods effectively carry out the simulation on a nominally Cartesian grid (which is simple to keep track of) and carefully handle the numerics at non-conformal boundary interfaces. This review paper discusses the mathematical formulation and provides several example applications of large-eddy simulation and immersed boundary methods, e.g. flow past a pickup truck.

- [61] P. Moin and T. Bewley. Feedback control of turbulence. *Applied Mechanics Reviews*, 47(6), 1994.

This review paper concentrates on active feedback control of turbulence rather than passive control mechanisms, e.g. the dimples on a golf ball or longitudinal V-shaped grooves, known as riblets, on aerodynamic surfaces. In active feedback control methods, external energy is added to the flow (active control), e.g. by synthetic jet actuators [56], and coordinated with real-time measurements of the flow itself (feedback control), e.g. by distributed micro-electro-mechanical system sensors that sense local wall-pressure fluctuations or other variations in the flow. Most promising applications of feedback control schemes involves flows regions that are most sensitive to such modification, e.g. areas of separation or transition. Adaptive schemes, dynamical systems, optimal control schemes, suboptimal control, cost function relationships and physical arguments for control are discussed. Implementation issues for quiet sensing of the flow and responding by several different actuation mechanisms is discussed. Significant improvements in drag reduction are suggested.

- [62] A.L. Marsden, M. Wang, J.E. Dennis, and P. Moin. Trailing-edge noise reduction using derivative-free optimization and large-eddy simulation. *Journal of Fluid Mechanics*, 572:13–36, 2007.

Shape optimization was applied to a time-accurate turbulent flow calculation using LES. Computation of the airfoil lift and drag for each new shape by with RANS for the entire airfoil. This was done to quickly evaluate whether or not the new airfoil shape would violate any constraints (e.g. reduced lift or increased drag) in advance of doing a full LES to evaluate the cost function. By separating the constraint and cost function evaluations in this way, large savings in computational cost were made by avoiding unnecessary LES

evaluations in cases when constraints were violated. Incompressible LES was used to simulate the flow over the trailing-edge of an airfoil. To calculate the far-field sound, a semi-infinite half-plane acoustics Green' function was used. Each LES used about 7 million mesh points. Mean-velocity boundary conditions were interpolated from the steady RANS results onto the smaller rectangular LES domain. The LES domain included only the trailing-edge portion of the airfoil. Turbulent-inflow was generated on the fly using the "rescale and recycle" technique. Note that the original airfoil was originally designed to study trailing-edge noise was not meant to be quiet; therefore, the large amount (89%) noise reduction achieved by optimization, mainly through suppression of vortex shedding, is expected. This work did demonstrate the successful coupling of shape optimization to a time-accurate turbulent-flow calculation using LES. The use of a novel derivative-free optimization technique provided a flexible optimization framework which can be applied to a variety of complex flow problems.

- [63] J.W. Nichols, S.K. Lele, P. Moin, F.E. Ham, G. Bres, and J.E. Bridges. Large-eddy simulation for supersonic rectangular jet noise prediction: effects of chevrons. *to appear 2012 AIAA aeroacoustics conference, June 2012, Colorado Springs, CO.*

Recent results from massively parallel simulations using 40 racks (163,840 cores) on the BlueGene/P Intrepid machine at the Argonne Leadership Computing Facility are presented. Flow and noise predictions from a supersonic rectangular jet with chevrons are documented and compared with companion experiments.

- [64] J.W. Nichols, F.E. Ham, S.K. Lele, and J.E. Bridges. Aeroacoustics of a supersonic rectangular jet: experiments and les predictions. *AIAA-678*, 2012.

The aeroacoustics of supersonic jets are intricately connected to turbulence in the jet shear layers and just downstream of the potential core. For pressure-mismatched jets, the interplay between this turbulence and shocks is also important. Despite significant scientific investigation into this complex multi-scale dynamic phenomenon, jet noise remains a large component of the overall noise generated by supersonic aircraft. This is of particular concern on aircraft carriers, where it is necessary for deck personnel to be in close proximity to aircraft at takeoff and landing. In such environments (over 150 dB), even the most advanced hearing protection cannot offer complete protection from long-term exposure. Experiments, however, have shown that it is possible to reduce jet noise at its source through modification of the nozzle geometry. Current jet noise reduction technologies include the adding chevrons to the nozzle lip, altering the nozzle's cross sectional shape, beveling the nozzle, and including micro-jets around the nozzle perimeter. In this article, noise from

isothermal and heated under-expanded supersonic turbulent jets are investigated using the high-fidelity, unstructured `charles` LES infrastructure and hybrid acoustic projection based on Ffowcs-Williams Hawkins (FWH) equations. The results of a grid resolution study for validation against laboratory measurements with the same nozzle geometry are compared and contrasted.

- [65] J.W. Nichols, F.E. Ham, S.K. Lele, and P. Moin. Prediction of supersonic jet noise from complex nozzles. *CTR Annual Research Briefs, Stanford, CA*, 2011.

This CTR Annual Research Brief (available here: <http://www.stanford.edu/group/ctr/publications.html>) documents the challenges and barriers that had to be overcome to enable the massively parallel several hundred million CV LES of supersonic jet noise from rectangular chevron nozzles [63]. Starting from a Computer-Aided Design (CAD) file (see upper-left inset of figure 17 provided by James Bridges at the NASA Glenn Research Center), a body-fitted (including the volume inside the nozzle and extending downstream) hexahedral mesh was made, which was progressively refined from 6 million to 262 million control volumes (CVs) using the isotropic grid refinement capability of `charles` to cluster grid points inside the FWH surface. High-resolution (263M CV) cases were found to scale near ideally on as many as 65,536 core on the Intrepid BlueGene/P machine at the Argonne National Laboratory Leadership Computing Facility. While this chevron jet simulation produced farfield noise spectra of similar shape and amplitude as in the experiment, large discrepancies were observed in the mid- to high-frequency range of the downstream directed sound. Removal of the near-nozzle portion of the FWH integration surface, however, corrected this discrepancy, suggesting more resolution is needed close to the nozzle lip to capture the complicated three-dimensional transition to turbulence in this region. The transition of isotropic grid refinement near the edge of the FWH from fine to coarse resolution was found to be too abrupt leading to some unphysical reflections at the FWH-fine-to-coarse-grid transition interface that may corrupt the FWH surface data. To remedy these issues, further adaptive refinement was performed near the nozzle lip and a strategy to smooth the fine-to-coarse mesh transition near the FWH surface was applied. The resulting mesh contains 528M CVs and at the time this paper was written was running on 131,072 cores on Intrepid. A full system dedicated run on all 40 racks was soon requested making the groundbreaking simulations presented in [63] possible.

- [66] S.K. Lele, S. Mendez, J. Ryu, J. Nichols, M. Shoeybi, and P. Moin. Sources of high-speed jet noise: Analysis of LES data and modeling. *Procedia Engineering*, 6:84 – 93, 2010.

Results from recent LES studies of high-speed jet flows and its near and far-field noise are reviewed with an emphasis on validation of the LES result.

- [67] J. Bridges. Validating LES for Jet Aeroacoustics. *AIAA-0017*, 2011.

From the abstract, “Engineers charged with making jet aircraft quieter have long dreamed of being able to see exactly how turbulent eddies produce sound – and this dream is now coming true with the advent of large eddy simulation (LES).” Recent advances in computing capability and optical diagnostics have provided aeroacoustic engineers with massive experimental and computational databases. How can we assess the accuracy of these increasingly sophisticated tools. What measures or metrics should be used in validation? Code validation and interpretation of these massive datasets are seen as the primary pacing items. The use of advanced experimental techniques (e.g. particle image velocimetry (PIV) and Raman and Rayleigh scattering) to provide the high-fidelity, detailed data needed to validate LES solutions is discussed. In particular, the author suggests that validation on higher-order statistics, critical to aeroacoustic phenomena, is needed. For non-experimentalists and simulation scientists, this paper addresses many common questions concerning validation methodologies for advanced experimental diagnostics and how this can inform companion LES studies of jet aeroacoustics (e.g. [63, 64, 65]).

- [68] F.K. Chow and P. Moin. A further study of numerical errors in large-eddy simulation. *Journal of Computational Physics*, 184:366–380, 2002.

Numerical errors in LES arise from aliasing and discretization errors, and errors in the subfilter-scale turbulence model. A systematic analysis of these errors is difficult given the non-linearities involved. The authors evaluate the effects of these errors by numerical tests and validation for a stably stratified shear flow and compare with DNS. Guidelines for choosing the proper grid size and subfilter-scale are given. This is particularly important to ensure the subfilter-scale terms are larger than the numerical errors due to aliasing and truncation.

- [69] A.G. Kravchenko and P. Moin. On the effect of numerical errors in Large Eddy simulations of turbulent flows. *Physics of Fluids*, 131(2):310–322, 1997.

The authors compare spectral and finite difference methods when applied to large-eddy simulation. Spectral methods are found to be energy-conserving only if dealiasing is performed. While, for skew-symmetric and rotational form of the non-linear equations, both spectral and finite-difference methods are energy-conserving in the inviscid limit even with aliasing errors. The relative importance of aliasing and truncation errors as compared to subgrid scale model terms in large eddy simulations is analyzed and discussed. Since dynamic subgrid scale models compute its coefficients by sampling information in the high wavenumber part of the resolve spectrum, it is important that this part of the spectrum not be distorted by numerical truncation errors.

- [70] P. Moin. *Fundamentals of Engineering Numerical Analysis*. Cambridge University Press, 2nd edition, 2011.

Concise, self-contained textbook on numerical analysis. Numerical interpolation, differentiation, integration, and the key ideas and methods needed to develop, analyze and use computers to solve ordinary and partial differential equations. Highly recommended.

- [71] S.K. Lele. Compact finite difference schemes with spectral-like resolution. *Journal of Computational Physics*, 103(1):16–42, 1991.

Finite difference schemes for the evaluation of derivatives, interpolation and filtering are presented. Emphasis is made toward improving the representation of a range of wavenumbers rather than achieving the highest possible formal accuracy (recall the discussion concerning fifth-order upwind and second-order central difference schemes for LES in §3.3 and figure 23). In particular, the differencing schemes described by the author provide an improved resolution of the short length scales and have no built-in artificial dissipation. Fourier analysis (e.g. modified wavenumbers) is shown to provide a tool to optimize the resolving efficiency of discretization schemes suitable for particular applications (e.g. acoustics and combustion).

- [72] R. Mittal and P. Moin. Suitability of upwind-biased finite difference schemes for large-eddy simulation of turbulent flows. *AIAA Journal*, 35(8), 1997.

Simulations of flow past a circular cylinder at Reynolds number 3900 using an energy-conservative, second-order central difference scheme is compared to simulations by Beaudan and Moin with a fifth-order upwind difference scheme for spatial discretization and with experiments to assess the impact of numerical diffusion on the computed flowfield. Numerical dissipation inherent to the upwind schemes was found to remove substantial energy at the small scales while the smaller scales were more energetic in the central difference scheme. In applications such as flow generated noise and reactive flows where small-scale fluctuations play an active role, the use of energy conservative schemes would be preferable over upwind schemes.

- [73] A.G. Kravchenko and P. Moin. Numerical studies of flow over a circular cylinder at $Re_D = 3900$. *Physics of Fluids*, 12(2), 1999.

Flow over a circular cylinder at Reynolds number 3900 is studied numerically with large-eddy simulation. Computations performed using several different numerical methods are compared to hot-wire measurements in the far-field wake of the same flow. In the near-wake region, each simulation method are in agreement with each other. Further downstream, however, the schemes using non-dissipative numerics are in better agreement with the available

experimental data. Effects of grid resolution on shear layer transition is evaluated. Numerical transition at coarse resolutions leads to early transition and poor prediction. Grid refinement verified that increased spanwise resolution reduced the numerical truncation error responsible for the premature transition and thereby improving the agreement.

- [74] D. Kim and P. Moin. Numerical simulation of the breakup of a round liquid jet by a coaxial flow of gas with a subgrid lagrangian breakup model. *Annual Research Briefs, Center for Turbulence Research*, 2011.

CTR Annual Research Briefs are located here:

<http://www.stanford.edu/group/ctr/publications.html>

- [75] A. Mani, J. Larsson, and P. Moin. Suitability of artificial bulk viscosity for large-eddy simulation of turbulent flows with shocks. *Journal of Computational Physics*, 228(19):7368 – 7374, 2009.

The simultaneous prediction of shock waves and turbulence remains a challenge. Capturing an essentially discontinuous shock wave on a requires the addition of numerical dissipation. As discussed in §3.3, such dissipation should be avoided when computing turbulent flows. Using too much dissipation for shock capturing may severely damp turbulent fluctuations. This paper presents a new formulation of artificial bulk viscosity methods parameterized by the rate of dilatation rather than strain rate magnitude is shown to provide significant sound-prediction capability in turbulent flows with shocks. The dilatation more effectively localizes the dissipation around the shocks and provides a less intrusive shock-capturing treatment.

- [76] E. Johnsen, J. Larsson, A.V. Bhagatwala, W.H. Cabot, P. Moin, B.J. Olson, P.S. Rawat, S.K. Shankar, B. Sjögreen, H.C. Yee, X. Zhong, and S.K. Lele. Assessment of high-resolution methods for numerical simulations of compressible turbulence with shock waves. *Journal of Computational Physics*, 229(4):1213 – 1237, 2010.

Shock waves and turbulence appear together in many flow, e.g. scramjet propulsion, supernovae explosion and inertial confinement fusion. Simulations of both phenomena at the same time are made difficult by the often contradictory numerical methods used to capture such phenomena. Accurate numerical methods for turbulence seek to minimize numerical dissipation so that the physical dissipation due to small-scale eddies is not overwhelmed while shock-capturing schemes introduce artificial numerical dissipation to stabilize the solution in such flow regions. This paper provides a critical performance analysis and assessment of the prevailing numerical methods in use today capable of simultaneously handling shocks and turbulence for a comprehensive

suite of test cases. WENO methods are found to provide sharp shock profiles but overwhelm the physical dissipation. Hybrid WENO/central-difference methods (e.g. [75]) are found to provide significant improvements.

- [77] A.E. Honein and P. Moin. Higher entropy conservation and numerical stability of compressible turbulence simulations. *Journal of Computational Physics*, 201(2):531 – 545, 2004.

A numerical formulation for the treatment of nonlinear instabilities in shock-free compressible turbulence simulations. The formulation is high order and contains no artificial dissipation. Numerical stability is enhanced through semi-discrete satisfaction of global conservation properties stemming from the second law of thermodynamics and the entropy equation. The numerical implementation is achieved using a conservative skew-symmetric splitting of the nonlinear terms. The robustness of the method is demonstrated by performing unresolved numerical simulations and large eddy simulations of compressible isotropic turbulence at very high Reynolds number.

- [78] A. Arakawa. Computational design for long-term numerical integration of the equations of fluid motion: Two-dimensional incompressible flow. Part I. *Journal of Computational Physics*, 82, 1956.

Showed that, when central differences are used, conservation of certain integral parameters (enstrophy) in the absence of viscous dissipation is required for long time integration of the flow equations. Otherwise, such a code would blow-up. Conceptually, “blow-up” implies a spurious unphysically large increase in energy. Arakawa modified the numerical scheme so that, at the discrete level, the method conserves total energy exactly to machine precision. As a result, “infinite” velocities are numerically not possible (and they are, of course, not physically possible either). To stabilize calculations using methods that do not preserve such quantities as Arakawa did, artificial viscosity is often introduced either explicitly or implicitly by using dissipative finite-difference schemes, especially for high-Reynolds number flows. Such methods may be called more “robust” or forgiving but, when numerical error is left unchecked, bad things (like blow-up and unphysical solutions) happen. Applied to turbulent flows, such conservation is even more critical for a numerical method to have as demonstrated by several researchers [80, 81]. Artificial viscosity methods are often used in flows involving shocks (a classic example being the paper by Von Neumann and Richtmyer [87]). Such methods are, however, not suitable for turbulent flows with shocks since the unphysical dissipation will not only thicken the shock but kill the turbulent eddies. For compressible, turbulent flows problems with shocks or shocklets, improved methods exist that preserve the relevant conservation principle and avoid artificial dissipation by

restricting its use to limited regions near thin, unresolved flow features such as shocks and phase interfaces [75, 76].

- [79] J.W. Smagorinsky. General circulation experiments with the primitive equations I. The basic experiment. *Monthly Weather Review*, 91:99–164, 1963.

An ambitious simulation of the planetary-scale atmosphere is documented in this paper (the work was complete about five years prior to publication). The longest resolved wavelength was bounded by the circumference of the earth and the smallest resolvable length scale was a few thousand kilometers (with $\Delta \approx 500$ km). Smaller-scale motions, such as hurricanes, clouds and convection, were not resolved. The equations Smagorinsky solved were not fully three-dimensional, but he did develop a model for the small eddies, which Deardorff [54] later adapted for the first three-dimensional turbulence simulations. In this early atmospheric model, the lateral eddy transfer of heat and momentum was introduced as a non-linear lateral diffusion coefficient with a tuning parameter k_H . If k_H is made too small, the smallest resolvable scales accumulate potential and kinetic energy by the non-linear eddy cascade, which can quickly lead to an unphysical “non-linear computational instability.” This small-scale lateral diffusion is introduced to “simulate the *physically real* net cascade of energy from the larger than grid-size scale to the smaller scales which have been truncated by the discrete difference.”

This approach has hence become known as the Smagorinsky model with a constant coefficient. Dynamic smagorinsky models, on the other hand, essentially eliminate the tunable parameters and model uncertainty associated with the unspecified coefficient [55, 17]. Subgrid-scale models, such as the constant coefficient Smagorinsky model, in effect utilize the high wavenumber information from a large-eddy simulation to extrapolate the effect of the unresolved scales on the resolved scales, in principle, at each point in space or instant in time. In many engineering (and geophysical) flows, the inertial range is often sparsely resolved if at all by the simulation grid. As a result, additional model equations that account for the pertinent length scales of the residual field may be needed (e.g. sub-grid models for macroscopic cloud formations).

- [80] Y. Morinishi, T.S. Lund, O.V. Vasilyev, and P. Moin. Fully conservative higher order finite difference schemes for incompressible flow. *Journal of Computational Physics*, 143(1), 1998.

In the development of a finite-difference scheme, the authors begin first with the requirement that the method analytically satisfy conservation of mass, momentum and kinetic energy. This may seem obvious given they are solving the incompressible Návier-Stokes equations, but, when applied on a discrete grid with a finite number of points, it turns out that a more carefully and

refined analysis of the governing equations at the discrete level are required. A general family of fully conservative high order accurate finite difference schemes are developed for staggered grid systems in this paper.

- [81] J. Kim and P. Moin. Application of a fractional-step method to incompressible Navier-Stokes equations. *Journal of Computational Physics*, 59(2), 1985.

To compute three-dimensional and time-dependent incompressible flows, fractional-step (or “projection methods”) are often used to advance the momentum equation and enforce the continuity equation (i.e. conservation of mass). The choice of appropriate boundary conditions and the exact manner in which one “corrects” the pressure at each timestep, which involves solving a Poisson equation, had hitherto been somewhat arbitrary and different choices were used by different people, many of which resulted in a reduction in the numerical order of accuracy. The authors showed that by using a staggered grid, certain boundary conditions, and a proper pressure-correction scheme, one could simultaneously conserve mass, momentum and kinetic energy discretely to machine precision for incompressible flows. In particular, the numerical method satisfies the continuity equation at the end of each timestep. This numerical method was used to compute flow in a driven cavity and over a backward-facing step [33].

- [82] G. Iaccarino. Quantification of uncertainty in flow simulations using probabilistic methods. In *RTO-EN-AVT-162 (VKI Lecture Series)*, September 2011.

The growing interest in uncertainty quantification has motivated the development of a variety of mathematical approaches. In particular, probabilistic uncertainty propagation methods have received considerable attention. Several probabilistic UQ methods are outlined in this report with an emphasis on those useful for typical fluid mechanics simulations. Alternative approaches not based on probabilistic reasoning have been proposed and used with some success. It is not generally clear when probabilistic methods fail or are insufficient; the treatment of epistemic uncertainty remains difficult and possibly the greatest challenge in uncertainty quantification.

- [83] M. Wang, J.B. Freund, and S.K. Lele. Computational prediction of flow-generated sound. *Annual Review of Fluid Mechanics*, 38:483–512, 2006.

Review of computational techniques for flow-noise prediction and underlying theories. Hybrid approaches, in which the turbulent noise source field is computed and/or modeled separately from the far-field calculation, are emphasized. Properly validated hybrid simulation methods are expected to remove much of the empiricism in the prediction of flow-generated sound. These developments require careful verification that numerical errors are small and validation against laboratory or full-scale data, including high-quality data

on the unsteady flow and sound. Figure 17 provides an example of a hybrid approach to the computational prediction of flow-generated sound.”

- [84] J. Christophe, S. Moreau, C.W. Hamman, J.A.S. Witteveen, and G. Iaccarino. Uncertainty quantification for the trailing-edge noise of a controlled-diffusion airfoil. *Proceedings of the 2010 Summer Program, Center for Turbulence Research*, 2010.

The understanding and control of acoustic radiation (more commonly known as noise) remains at the forefront of investigations in science and engineering. Reducing acoustic emissions (without sacrificing performance) is a chief concern of the transportation and energy conversion sectors where issues of human comfort and noise regulations are important (e.g. 2MW size wind turbines, residential communities near airports and automobile radiator cooling fan noise). Designing efficient computer cooling systems with a minimum noise configuration is also very important for the consumer electronics industry and for the leadership-class computing facilities.

In this paper, the authors attempt to quantify the uncertainty in simulation of the trailing-edge noise of a controlled-diffusion airfoil embedded in the potential core of a jet in an anechoic wind tunnel experiment using RANS and LES. A non-intrusive stochastic collocation method is used to propagate several aleatoric, aerodynamic uncertainties (e.g. angle of attack and inflow conditions) through both deterministic incompressible flow solvers. For the RANS calculations two different wall pressure fluctuation models were used to evaluate the efficiency of a non-intrusive stochastic Galerkin method based on a stochastic collocation expansion which was found to be superior to a more traditional, and expensive, Monte Carlo approach. When applying the method to the LES calculations, which have increased physical fidelity and introduce sensitivities not seen in the RANS calculations, the authors find a considerable increase in uncertainty of the far-field radiation as a function of the unknown inflow velocity, relative to the RANS results.

- [85] P. Moin and J. Kim. On the numerical solution of time-dependent viscous incompressible fluid flows involving solid boundaries. *Journal of Computational Physics*, 35:381–392, 1979.

A basic numerical problem associated with fully explicit pseudospectral numerical simulations of turbulence near walls is described. A semi-implicit scheme which resolves this problem is presented. The formulation in terms of Fourier series and Chebyshev polynomial expansions is given. The paper notes that “A common objective of the large eddy and direct simulation techniques is to test and suggest statistical models of turbulence which can in turn be used in a simpler method for complex flows. In this case, it is imperative that errors introduced by the numerical scheme are minimized. This is necessary for an objective evaluation of the turbulence model.” The authors

outline such an accurate numerical method for the simulation wall-bounded turbulent flows using spectral methods. Only a few years later, the authors used this numerical method for their landmark calculation given in the 1987 J. Kim, P. Moin and R. Moser paper [29]. Careful numerical analysis on the drawing board surely saved the author's many frustrating experiences and tortuous hours debugging their code and reconciling the outputs spewed from their computers with the physics.

- [86] D. You and P. Moin. Effects of hydrophobic surfaces on the drag and lift of a circular cylinder. *Physics of Fluids*, 19(081701), 2007.

Flow over a cylinder with no-slip boundary conditions is compared to different cylinder configurations with regions of slip boundary condition. The slip boundary conditions are intended to model the effects of hydrophobic surfaces (e.g. micro/nano-textured surfaces that preferentially trap small gas bubbles in a liquid environment) that attempt to prevent the surface from becoming fully wetted, which would lead to increased skin-friction drag. This paper shows that strategic hydrophobic surface treatment can significantly reduce the drag and lift in cylinder flow (by up to 27% and 75%, respectively). Simulations that explicitly account for the detail micro/nano-scale surface texture characteristic of hydrophobic surfaces rather than using slip boundary conditions are under investigation. The base cylinder flow and Reynolds number ($Re_D = 3900$) are the same as that studied in [49] where the effect of polymer injection instead of hydrophobic surface was investigated.

- [87] J. Von Neumann and R.D. Richtmyer. A method for the numerical calculation of hydrodynamic shocks. *Journal of Applied Physics*, 21(3):232–237, 1950.

Artificial dissipative terms are introduced into the governing equations of motion so that a shock is smeared out to have a width at least the size of the grid space on the numerical mesh. Very strong shocks have physical dimensions that are very nearly a mathematical discontinuity (being only a few mean free paths thick). Such methods allow the shock to be “resolved” on the grid without worrying too much about the small-scale details. Such methods applied to turbulent flows with shocks are, however, not appropriate as the same terms responsible for thickening the shock dissipate the small-scale turbulent eddies (which are typically still very much larger than a few mean free paths in size). Today, turbulent flows with shocks should at least selectively apply such artificial viscosity or compressibility terms in the near-shock region and not where turbulence is present. This requires an appropriate criterion to separate shocked and unshocked regions (which may be difficult in flows with both strong turbulence, shocks and shocklets). Such methods are discussed in the following references [75, 76].