

Tackling Turbulence with Supercomputers

Computers only recently became powerful enough to illuminate simple examples of this great classical problem. In some cases, they will let engineers control it

by Parviz Moin and John Kim

We all pass through life surrounded—and even sustained—by the flow of fluids. Blood moves through the vessels in our bodies, and air (a fluid, properly speaking) flows into our lungs. Our vehicles move through our planet's blanket of air or across its lakes and seas, powered by still other fluids, such as fuel and oxidizer, that mix in the combustion chambers of engines. Indeed, many of the environmental or energy-related issues we face today cannot possibly be confronted without detailed knowledge of the mechanics of fluids.

Practically all the fluid flows that interest scientists and engineers are turbulent ones; turbulence is the rule, not the exception, in fluid dynamics. A solid grasp of turbulence, for example, can

allow engineers to reduce the aerodynamic drag on an automobile or a commercial airliner, increase the maneuverability of a jet fighter or improve the fuel efficiency of an engine. An understanding of turbulence is also necessary to comprehend the flow of blood in the heart, especially in the left ventricle, where the movement is particularly swift.

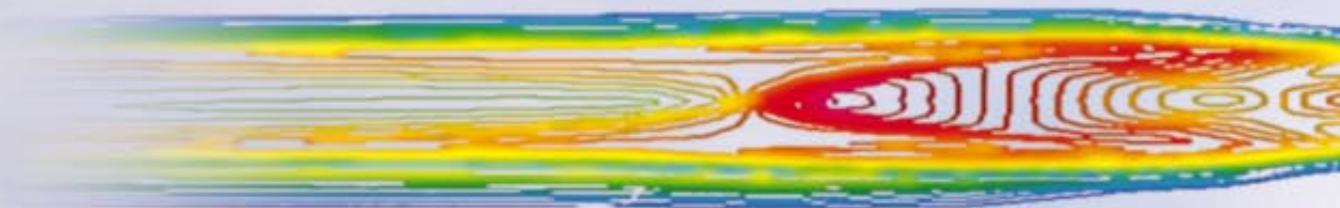
But what exactly is turbulence? A few everyday examples may be illuminating. Open a kitchen tap only a bit, and the water that flows from the faucet will be smooth and glassy. This flow is known as laminar. Open the tap a little further, and the flow becomes more roiled and sinuous—turbulent, in other words. The same phenomenon can be seen in the smoke streaming upward into still air from a burning cigarette. Immediately

above the cigarette, the flow is laminar. A little higher up, it becomes rippled and diffusive.

Turbulence is composed of eddies: patches of zigzagging, often swirling fluid, moving randomly around and about the overall direction of motion. Technically, the chaotic state of fluid motion arises when the speed of the fluid exceeds a specific threshold, below which viscous forces damp out the chaotic behavior.

Turbulence, however, is not simply an unfortunate phenomenon to be eliminated at every opportunity. Far from it: many engineers work hard trying to increase it. In the cylinders of an internal-combustion engine, for example, turbulence enhances the mixing of fuel and oxidizer and produces cleaner, more ef-

SPACE SHUTTLE SIMULATION was combined with a photograph of the shuttle for reference. In the bottom half of the image, different colors indicate air-pressure values at the vehicle's surface, from blue (low pressure) to red (high).



ficient combustion. And only turbulence can explain why a golf ball's dimples enable a skilled golfer to drive the ball 250 meters, rather than 100 at most.

Turbulence may have gotten its bad reputation because dealing with it mathematically is one of the most notoriously thorny problems of classical physics. For a phenomenon that is literally ubiquitous, remarkably little of a quantitative nature is known about it. Richard Feynman, the great Nobel Prize-winning physicist, called turbulence "the most important unsolved problem of classical physics." Its difficulty was wittily expressed in 1932 by the British physicist Horace Lamb, who, in an address to the British Association for the Advancement of Science, reportedly said, "I am an old man now, and when I die and go to heaven there are two matters on which I hope for enlightenment. One is quantum electrodynamics, and the other is the turbulent motion of fluids. And about the former I am rather optimistic."

Of course, Lamb could not have foreseen the development of the modern supercomputer. These technological marvels are at last making it possible for engineers and scientists to gain fleeting but valuable insights into turbulence. Already this work has led to technology, now in development, that may someday be employed on airplane wings to reduce drag by several percent—enough to save untold billions of dollars in fuel costs. At the same time, these insights are guiding the design of jet engines to im-

prove both efficiency and performance.

As recondite as it is, the study of turbulence is a major component of the larger field of fluid dynamics, which deals with the motion of all liquids and gases. Similarly, the application of powerful computers to simulate and study fluid flows that happen to be turbulent is a large part of the burgeoning field of computational fluid dynamics (CFD). In recent years, fluid dynamicists have used supercomputers to simulate flows in such diverse cases as the America's Cup racing yachts and blood movement through an artificial heart.

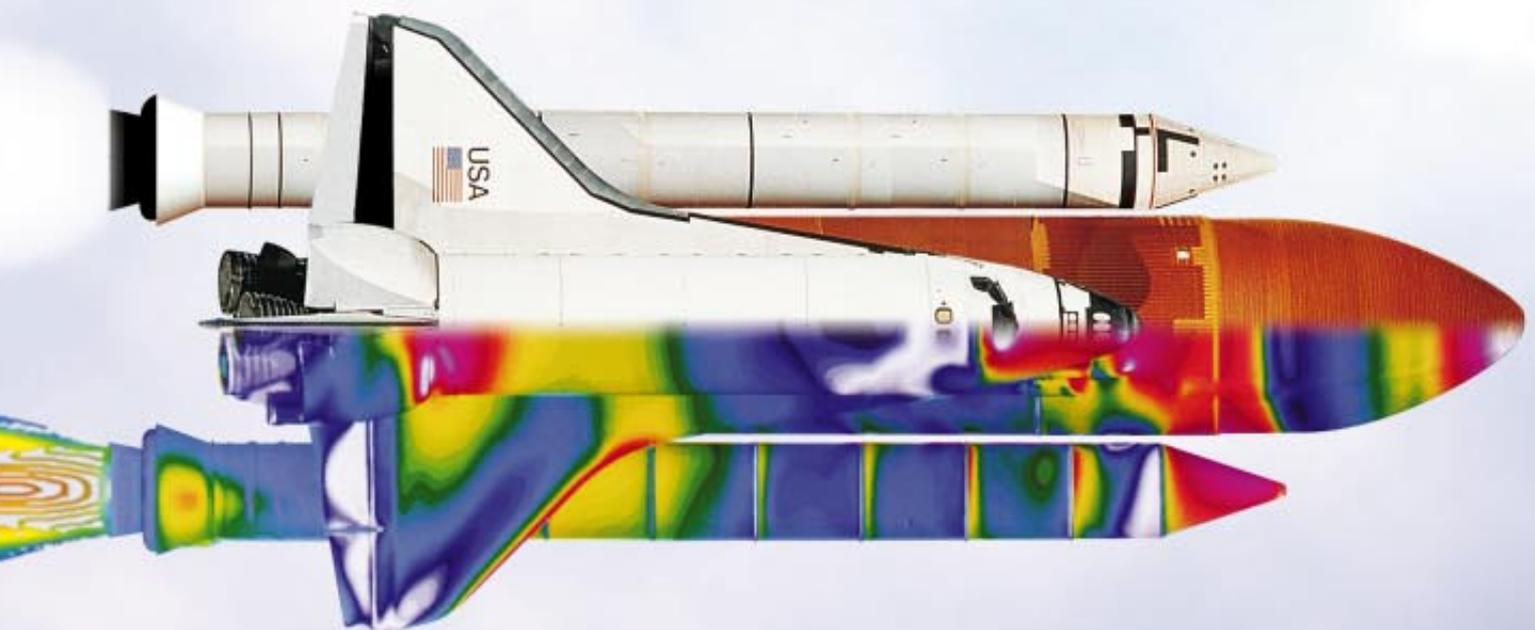
CFD: 150 Years in the Making

What do we mean when we speak of simulating a fluid flow on a computer? In simplest terms, the computer solves a series of well-known equations that are used to compute, for any point in space near an object, the velocity and pressure of the fluid flowing around that object. These equations were discovered independently more than a century and a half ago by the French engineer Claude Navier and the Irish mathematician George Stokes. The equations, which derive directly from Newton's laws of motion, are known as the Navier-Stokes equations. It was the application of supercomputers to these equations that gave rise to the field of computational fluid dynamics; this marriage has been one of the greatest achievements in fluid dynamics since the equations themselves were formulated.

Although the marriage has been successful, the courtship was a rather long one. Not until the late 1960s did supercomputers begin achieving processing rates fast enough to solve the Navier-Stokes equations for some fairly straightforward cases, such as two-dimensional, slowly moving flows about an obstacle. Before then, wind tunnels were essentially the only way of testing the aerodynamics of new aircraft designs. Even today the limits of the most powerful supercomputers still make it necessary to resort to wind tunnels to verify the design for a new airplane.

Although both computational fluid dynamics and wind tunnels are now used for aircraft development, continued advances in computer technology and algorithms are giving CFD a bigger share of the process. This is particularly true in the early design stages, when engineers are establishing key dimensions and other basic parameters of the aircraft. Trial and error dominate this process, and wind-tunnel testing is very expensive, requiring designers to build and test each successive model. Because of the increased role of computational fluid dynamics, a typical design cycle now involves between two and four wind-tunnel tests of wing models instead of the 10 to 15 that were once the norm.

Another advantage of supercomputer simulations is, ironically, their ability to simulate more realistic flight conditions. Wind-tunnel tests can be contaminated by the influence of the tunnel's walls and the structure that holds the model in

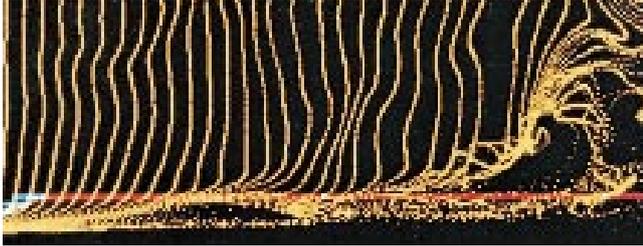


ROGER RESSMEYER/Corbis (photograph); PETER BUNING (simulation); SUMFILMS (digital montage)

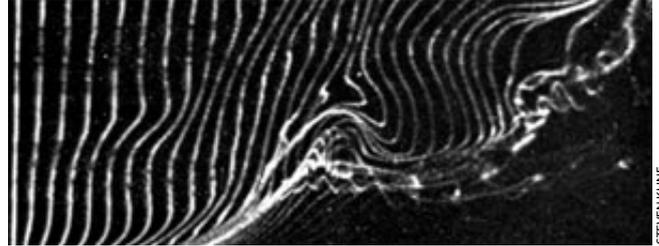
A Simulation Milestone

Until around 1980, few researchers attempted to simulate even very simple turbulent flows in their entirety. That year we and our co-workers at the NASA Ames Research Center used a pioneering parallel computer, the ILLIAC-IV, to perform the largest turbulence simulations achieved until then. The work was well received; soberingly enough, however, it was not the

quality of the data that won over many of our colleagues but rather a five-minute motion picture of the simulated flow. The movie showed trajectories of marker particles in a turbulent flow between parallel plates (*left*); remarkably, it resembled similar visualizations made two decades earlier, by filming actual water flow in a laboratory at Stanford University (*right*). —P.M. and J.K.



PARVIZ MOIN AND JOHN KIM



STEVEN KLINE

place. Some of the flight vehicles of the future will fly at many times the speed of sound and under conditions too extreme for wind-tunnel testing. For hypersonic aircraft (those that will fly at up to 20 times the speed of sound) and spacecraft that fly both within and beyond the atmosphere, computational fluid dynamics is the only viable tool for design. For these vehicles, which pass through the thin, uppermost levels of the atmosphere, nonequilibrium air chemistry and molecular physics must be taken into account.

Engine designers also rely extensively on computational techniques, particularly in the development of jet engines. A program called Integrated High Performance Turbine Engine Technology is seeking a 100 percent improvement in the thrust-to-weight ratio of jet engines and a 40 percent improvement in fuel efficiency by 2003. The project is supported by the U.S. Department of Defense, the National Aeronautics and Space Administration and various makers of jet engines.

The flow of air and fuel through a jet engine's various sections and passages is complex. A fan draws air into an internal chamber called a compressor. There multiple rotating and stationary stages increase the pressure about 20-fold. This high-pressure air is fed into a combustor, where it mixes with fuel and is ignited. Finally, the hot, very expanded exhaust drives a turbine. This turbine powers the fan and the compressor and, more important, generates thrust by directing the exhaust out of the rear of the engine at high velocity. Currently engineers use computational fluid dynamics

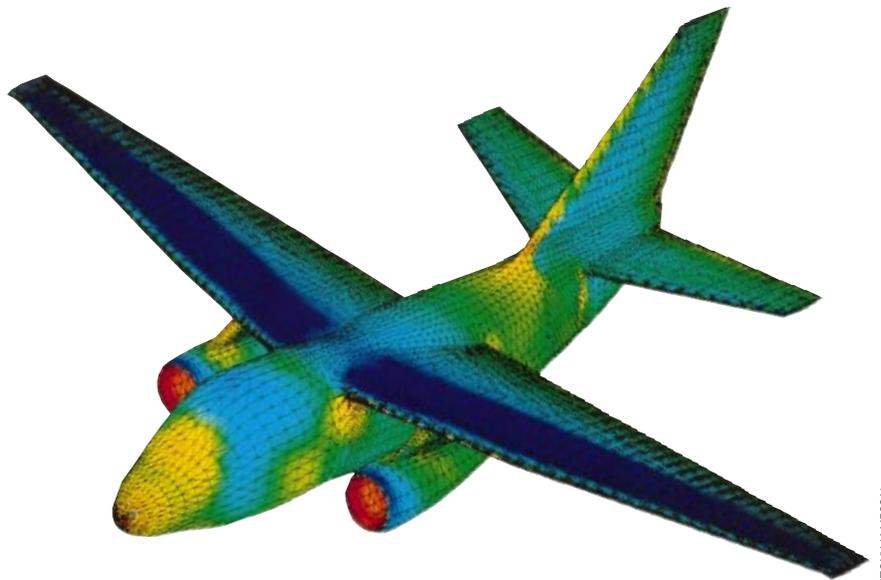
to design turbine blades, inlet passages and the geometry of combustors. Simulations also help engineers shape the afterburner mixers that, in military aircraft, provide additional thrust for greater maneuverability. And they play a role in designing nacelles, the bulbous, cylindrical engine casings that typically hang below the wings.

Applying the Equations

To understand how the Navier-Stokes equations work, consider the flow of air over an airplane in flight. In reality, it will probably be many decades

before computers are powerful enough to simulate in a detailed manner the fluid flows over an entire airplane. In theory, however, the Navier-Stokes equations can reveal the velocity and pressure of the air rushing by any point near the aircraft's surface. Engineers could then use these data to compute, for various flight conditions, all aerodynamic parameters of interest—namely, the lift, drag and moments (twisting forces) exerted on the airplane.

Drag is particularly important because it determines an aircraft's fuel efficiency. Fuel is one of the largest operating expenses for most airlines. Not surprising-



ANTHONY JAMESON

AIR PRESSURE over a Lockheed S-3A airplane in flight is highest near the craft's nose and inside the engine nacelles, which are below the wings. The grid visible on the surface of the image above is the computational mesh; a value of the air pressure was computed for each point at which grid lines intersect. Such simulations are critical means of

ly, aircraft companies have spent huge sums to reduce drag by even tiny increments. In general, though, lift is relatively easy to calculate, moments are harder, and drag is hardest of all.

Drag is difficult to compute mainly because it is the parameter most dependent on turbulence. Of course, in this context we are not referring to the bumpiness that provokes the pilot to remind passengers to fasten their seat belts. Even when a plane is flying smoothly, the flow of air within a few centimeters of its surface, in a volume known as the boundary layer, is turbulent. Because of turbulence, the high-speed air several millimeters above the surface of the wings is brought very close to the surface, where it undergoes a more abrupt—and momentum-robbing—deceleration. The equal and opposite reaction to this flow deceleration is drag on the aircraft. A great deal of the work of aerodynamicists involves understanding the mechanics of the generation and destruction of turbulence well enough to control it.

To solve the Navier-Stokes equations, engineers start by entering into the equations certain variables known as initial and boundary conditions. For an airplane in flight, the initial conditions include wind velocity and atmospheric disturbances, such as air currents. The boundary conditions include the precise shape of the aircraft, translated into mathematical coordinates.

Before the equations can be applied

to an aircraft, computer specialists must represent the aircraft's surface and the space around it in a form usable by the computer. So they represent the airplane and its surroundings as a series of regularly spaced points, known as a computational grid. They then supply the coordinates and related parameters of the grid to the software that applies the Navier-Stokes equations to the data. The computer calculates a value for the parameters of interest—air velocity and pressure—for each of the grid points.

In effect, the computational grid breaks up (the technical term is “discretizes”) the computational problem in space; the calculations are carried out at regular intervals to simulate the passage of time, so the simulation is temporally discrete as well. The closer together—and therefore more numerous—the points are in the computational grid, and the more often they are computed (the shorter the time interval), the more accurate and realistic the simulation is. In fact, for objects with complex shapes, even defining the surface and generating a computational grid can be a challenge.

Unfortunately, entering the initial and boundary conditions does not guarantee a solution, at least not with the computers available today or in the foreseeable future. The difficulty arises from the fact that the Navier-Stokes equations are nonlinear; in other words, the many variables in the equations vary with respect to one another by powers of two

or greater. Interaction of these nonlinear variables generates a broad range of scales, which can make solving the equations exceedingly difficult. Specifically, in turbulence, the range of the size of whirling eddies can vary 1,000-fold or even more. There are other complicating factors as well, such as global dependence: the nature of the equations is that the fluid pressure at one point depends on the flow at many other points. Because the different parts of the problem are so interrelated, solutions must be obtained at many points simultaneously.

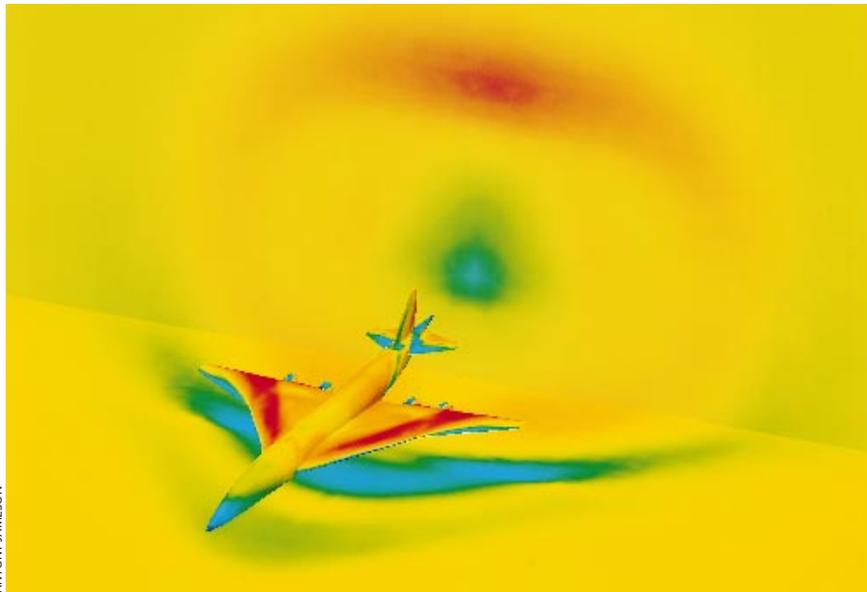
Computational Bête Noir

Although the preceding description conveys the basics of a fluid dynamics simulation, it leaves out turbulence, without which a realistic discussion of the capabilities—and limitations—of computational fluid dynamics would be futile. The complexities engendered by turbulence severely limit our ability to simulate fluid flow realistically.

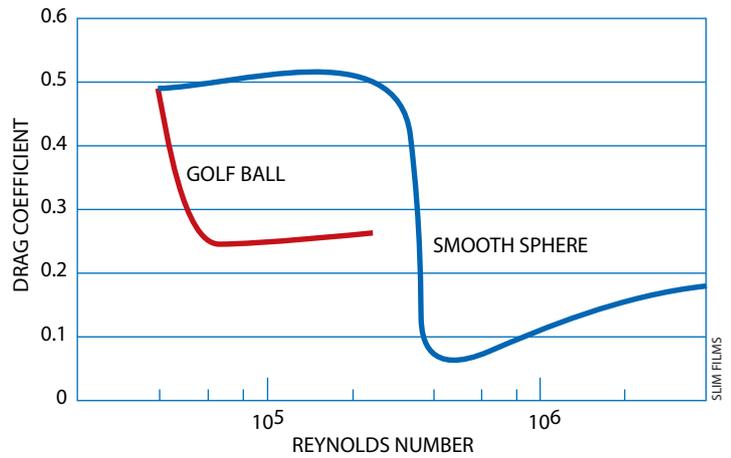
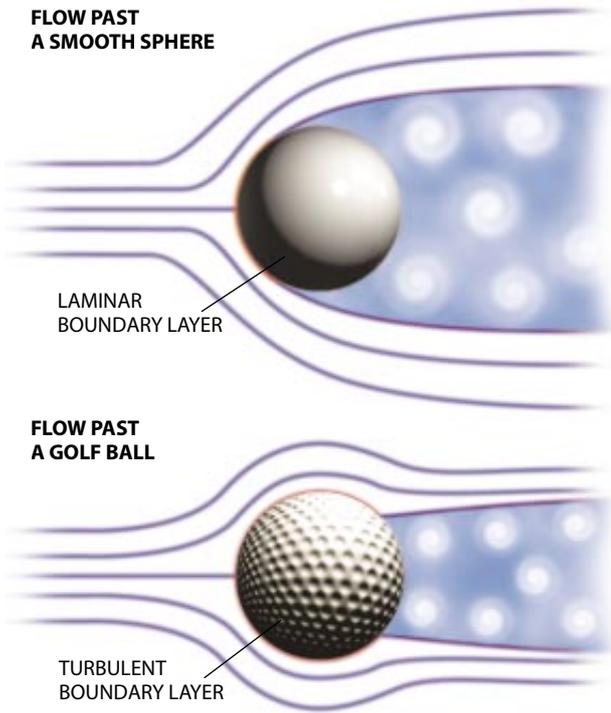
Perhaps the simplest way to define turbulence is by reference to the Reynolds number, a parameter that compactly characterizes a flow. Named after the British engineer Osborne Reynolds, this number indicates the ratio, or relative importance, of the flow's inertial forces to its viscous ones. (A flow's inertial force is calculated by multiplying together the fluid's density and the square of its velocity and dividing this product by a characteristic length of the flow, such as the width of the airfoil, if the flow is over a wing.)

Large inertial forces, relative to the viscous ones, tend to favor turbulence, whereas high viscosity staves it off. Put another way, turbulence occurs when the Reynolds number exceeds a certain value. The number is proportional to both the size of the object and the flow velocity. For example, the Reynolds number for air flowing over the fuselage of a cruising commercial aircraft is in the neighborhood of 100 million. For the air flowing past a good fastball, the Reynolds number is about 200,000. For blood flowing in a midsize artery, it is about 1,000.

As we have seen, a distinguishing characteristic of a turbulent flow is that it is composed of eddies, also known as vortices, in a broad range of sizes. These vortices are continually forming and breaking down. Large eddies break down into smaller ones, which break down into yet smaller eddies, and so on.



verifying designs for supersonic (*above*) and hypersonic aircraft, for which wind-tunnel testing is impossible. The simulation of this supersonic aircraft design also shows the sonic boom, visible as a circle behind the vehicle. Such booms are a major issue in ongoing studies of whether the public will accept these aircraft.



DRAG ON A GOLF BALL comes mainly from air-pressure forces. This drag arises when the pressure in front of the ball is significantly higher than that behind the ball. The only practical way of reducing this differential is to design the ball so that the main stream of air flowing by it is as close to the surface as possible. This situation is achieved by a golf ball's dimples, which augment the turbulence very close to the surface, bringing the high-speed airstream closer and increasing the pressure behind the ball. The effect is plotted in the chart, which shows that for Reynolds numbers achievable by hitting the ball with a club, the coefficient of drag is much lower for the dimpled ball.

When eddies become small enough, they simply dissipate viscously into heat. The British meteorologist Lewis F. Richardson described this process in verse:

Big whorls have little whorls,
Which feed on their velocity,
And little whorls have lesser whorls,
And so on to viscosity.

To solve the Navier-Stokes equations for, say, the flow over an airplane requires a finely spaced computational grid to resolve the smallest eddies. On the other hand, the grid must be large enough to encompass the entire airplane and some of the space around it. The disparity of length scales in a turbulent flow—the ratio of largest to smallest eddy size—can be calculated by raising the flow's Reynolds number to the $3/4$ power. This ratio can be used to estimate the number of grid points that are needed for a reasonably accurate simulation: because there are three dimensions, the number is proportional to the cube of this ratio of length scales. Thus, the required number of grid points for a numerical simulation is proportional to the Reynolds number raised to the $9/4$ power. In other words, doubling the Reynolds number results in almost a five-fold increase in the number of points required in the grid to simulate the flow.

Consider a transport airplane with a 50-meter-long fuselage and wings with

a chord length (the distance from the leading to the trailing edge) of about five meters. If the craft is cruising at 250 meters per second at an altitude of 10,000 meters, about 10 quadrillion (10^{16}) grid points are required to simulate the turbulence near the surface with reasonable detail.

What kind of computational demands does this number of points impose? A rough estimate, based on current algorithms and software, indicates that even with a supercomputer capable of performing a trillion (10^{12}) floating-point operations per second, it would take several thousand years to compute the flow for one second of flight time! Such a “teraflops” computer does not yet exist, although researchers are now attempting to build one at Sandia National Laboratories. It will be about 10 times faster than the most powerful systems available today.

Simulation Shortcuts

Fortunately, researchers need not simulate the flow over an entire aircraft to produce useful information. Indeed, doing so would probably generate much more data than we would know what to do with. Typically, fluid dynamicists care only about the effects of turbulence on quantities of engineering significance, such as the mean flow of a fluid or, in the case of an aircraft, the drag and lift

forces and the transfer of heat. In the case of an engine, designers may be interested in the effects of turbulence on the rates at which fuel and oxidizer mix.

The Navier-Stokes equations are therefore often averaged over the scales of the turbulence fluctuations. What this means is that, in practice, researchers rarely calculate the motion of each and every small eddy. Instead they compute the large eddies and then use ad hoc modeling practices to estimate the effects of the small eddies on the larger ones. This practice gives rise to a simulated averaged flow field that is smoother than the actual flow—and thus drastically reduces the number of grid points necessary to simulate the field.

The ad hoc models that this averaging process demands range in complexity from simple enhanced coefficients of viscosity to entire additional systems of equations. All these models require some assumptions and contain adjustable coefficients that are derived from experiments. Therefore, at present, simulations of averaged turbulent flows are only as good as the models they contain.

As computers become more powerful, however, fluid dynamicists are finding that they can directly simulate greater proportions of turbulent eddies, enabling them to reduce the range of scales that are modeled. These approaches are a compromise between a direct numerical simulation of turbulence, in which

all scales of motion are resolved, and the turbulence-averaged computations.

For years, meteorologists have used a form of this strategy called large-eddy simulation for weather prediction. In meteorology, the large-scale turbulent motions are of particular interest, so in meteorological applications the relatively large eddies are generally simulated in their entirety. Smaller-scale eddies are important only inasmuch as they may affect the larger-scale turbulence, so they are merely modeled. Recently engineers have begun using these techniques for simulating complex fluid flows, such as the gases inside a cylinder of an internal-combustion engine.

Another current trend in computational fluid dynamics, also made possible by increasing computational speed, is the direct, complete simulation of relatively simple flows, such as the flow in a pipe. Simple as they are, simulations of some of these flows, which have low Reynolds numbers, offer invaluable insights into the nature of turbulence. They have revealed the basic structure of turbulent eddies near a wall and subtleties of their influence on drag. They have also generated useful data that have enabled engineers to validate or fine-tune the ad hoc models they use in practical simulations of complex flows.

Lately the number of engineers and scientists seeking access to these data has swelled to the point that immense data sets have been archived and made available by the NASA Ames Research Center. Although most researchers do not have the computing resources to perform direct simulations of turbulence, they do have sufficient resources, such as powerful workstations, to probe the archived data.

From Prediction to Control

As supercomputers become faster and faster, fluid dynamicists are increasingly able to move beyond predicting the effects of turbulence to actually controlling them. Such control can have enormous financial benefits; a 10 percent reduction in the drag of civilian aircraft, for example, could yield a 40 percent increase in the profit margin of an airline. In a recent project, researchers at the NASA Langley Research Center demonstrated that placing longitudinal V-shaped grooves, called riblets, on the surface of an aircraft's wing or fuselage leads to a 5 to 6 percent reduction in viscous drag. Drag is reduced despite the

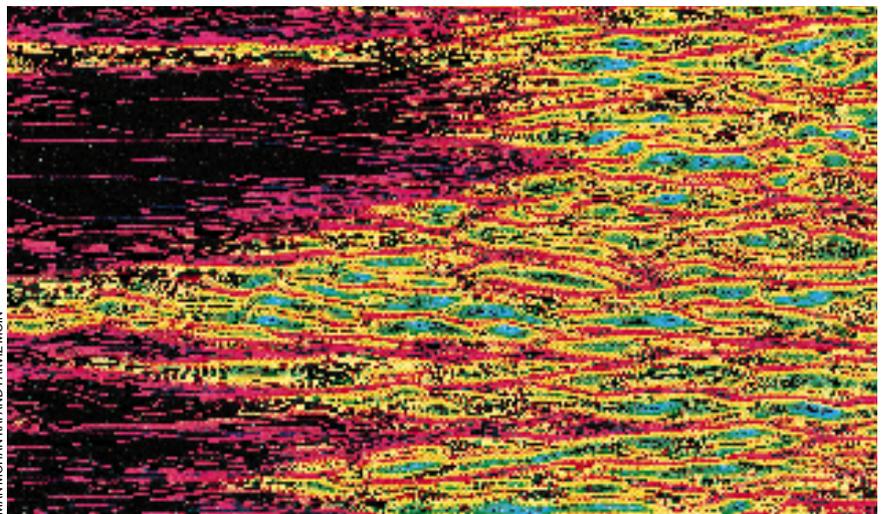
increase in the surface area exposed to the flow. For typical transport airplane speeds, the riblets must be very finely spaced, about 40 microns apart, like phonograph grooves; larger riblets tend to increase drag.

During this work, the researchers came across a Soviet study on toothlike structures, called denticles, on the skin of sharks. These denticles strikingly re-

sembled riblets, a fact that has been interpreted as nature's endorsement of the riblet concept. Ultimately, however, it was the direct numerical simulation of turbulent flow along riblets that showed how they work. The riblets appear to inhibit the motion of eddies by preventing them from coming very close to the surface (within about 50 microns). By keeping the eddies this tiny distance

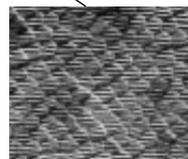
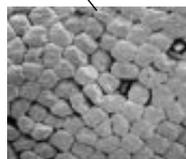
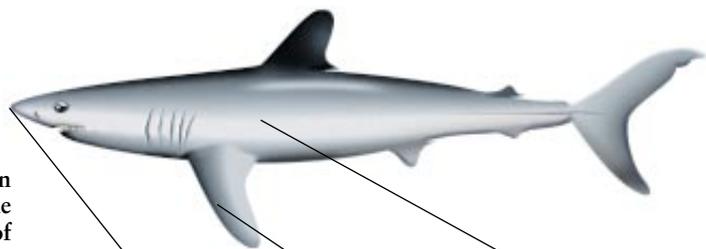
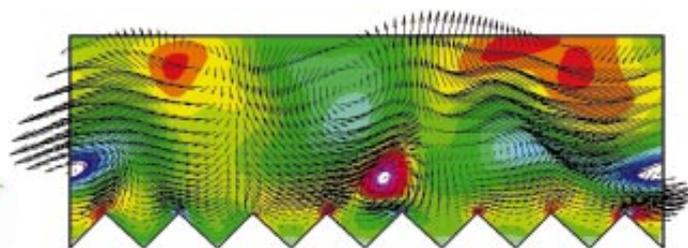
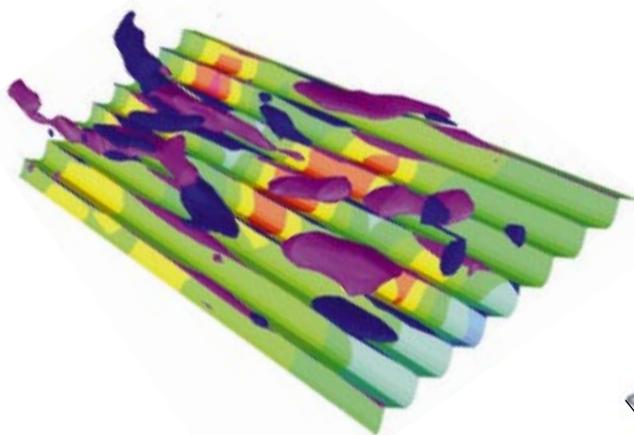


DAVID SCHARF



MANMOHAN RAI AND PARVIZ MOIN

CLOSE-UP VIEW of the simulated airflow over an airplane wing shows the transition from smooth, or laminar, flow (dark areas at left) to turbulence (rippled areas at right). Tiny actuators called microflaps (top photomicrograph) would tilt upward or remain flat in response to pressure variations to control small eddies on the wing's surface.



HAECHON CHOI, PARVIZ MOIN AND JOHN KIM

SLIM FILMS

WOLFF-ERINSTRUP

V-SHAPED GROOVES called riblets (*above*) have been found to inhibit the motion of eddies, thereby lessening the drag they can exert on an aircraft's wing. A cross section of the airflow over the grooves (*above, right*) shows the vorticity contours—warm colors are clockwise, and cool colors are counterclockwise. The arrows are velocity vectors for the flow over the riblets. To be effective, the riblets must be very closely spaced, like phonograph grooves. The skin of a shark has tiny, toothlike denticles (*photomicrographs at right*) that seem to serve the same function as the riblets, lessening the drag on the creatures as they move through the water.

away, the riblets prevent the eddies from transporting high-speed fluid close to the surface, where it decelerates and saps the aircraft's momentum.

Another recent and exciting application of direct numerical simulation is in the development of turbulence-control strategies that are active (as opposed to such passive strategies as riblets). With these techniques the surface of, say, a wing would be moved slightly in response to fluctuations in the turbulence of the fluid flowing over it. The wing's surface would be built with composites having millions of embedded sensors and actuators that respond to fluctuations in the fluid's pressure and speed in such a way as to control the small eddies that cause the turbulence drag.

Such technology appeared to be far-fetched as recently as three years ago, but the advent of microelectromechanical systems (MEMS), under the leadership of the U.S. Air Force, has brought such a scheme to the brink of implemen-

tation. MEMS technology can fabricate integrated circuits with the necessary microsensors, control logic and actuators. Active control of turbulence near the wing's surface also has an encouraging analogue in the form of a marine creature: dolphins achieve remarkable propulsive efficiency as they swim, and fluid dynamicists have long speculated that these creatures do it by moving their skins. It seems that smart aircraft skins, too, are endorsed by nature.

Getting back to those golf balls we mentioned earlier: they, too, present an intriguing example of how a surface texture can advantageously control airflow [see illustration on page 66]. The most important drag exerted on a golf ball derives from air-pressure forces. This phenomenon arises when the air pressure in front of the ball is significantly higher than the pressure behind the ball. Because of the turbulence generated by the dimples, a golf ball is able to fly about two and a half times farther

than an identical but undimpled ball.

The growing popularity of computational fluid dynamics to study turbulence reflects both its promise, which is at last starting to be realized, and the continued rapid increase in computational power. As supercomputer processing rates approach and surpass a trillion floating-point operations per second over the next few years, fluid dynamicists will begin taking on more complex turbulent flows, of higher Reynolds numbers. Over the next decade, perhaps, researchers will simulate the flow of air through key passages in a jet engine and obtain a realistic simulation of an operating piston-cylinder assembly in an internal-combustion engine, including the intake and combustion of fuel and the exhaust of gases through valves. Through simulations such as these, researchers will finally begin learning some of the deep secrets expressed by the equations uncovered by Navier and Stokes a century and a half ago. SA

The Authors

PARVIZ MOIN and JOHN KIM worked together in the early 1980s at the National Aeronautics and Space Administration Ames Research Center. Moin is now Franklin and Caroline Johnson Professor of Engineering at Stanford University and director of the Center for Turbulence Research there. His work lately has been on turbulence control, on the interaction of turbulence with shock waves and on noise generated by turbulent flows. Kim is Rockwell International Professor of Engineering at the University of California, Los Angeles. Much of his recent work has been in the application of neural networks to the control of turbulence.

Further Reading

PHYSICAL FLUID DYNAMICS. D. J. Tritton. Oxford Science Publications, 1988.
 TURBULENT TIMES FOR FLUIDS. Tom Mullin in *Exploring Chaos: A Guide to the New Science of Disorder*. Edited by Nina Hall. W. W. Norton, 1993.
 FEEDBACK CONTROL OF TURBULENCE. Parviz Moin and Thomas Bewley in *Applied Mechanics Reviews*, Vol. 47, No. 6, Part 2, pages S3–S13; June 1994.