supercomputers for the very top end of the market. IBM, for example, more than doubled the number of its computers in the Top 500 list between last November and June by introducing the SP2, which strings together up to 512 of the company's RS/6000 workstation microprocessors. Even Cray Research has introduced a new line of massively parallel computers that rely on microprocessors made by DEC.

The big losers in the supercomputer market, many believe, will be vector supercomputers and any other machines—including MPPs—that rely on custom-designed processors. That's the common thread connecting the past year's three major failures in the industry, after all. Cray Computer, for example, foundered in part on the technical difficulty of making chips from gallium arsenide rather than the usual silicon. But more important, industry observers say, the company was attempting to sell expensive, custom supercomputers to a market being invaded by fast, cheap microprocessors.

Kendall Square Research's problems were similar. The company had an imaginative plan for creating massively parallel computers with a scaleable shared memory—one that, although physically scattered through the computer, would simplify programming by acting as a single memory. But, explains Kalos at the Cornell Theory Center, that type of memory was impossible to build with commodity chips. Kendall Square had to create its own processors, and it was the victim of competition from commodity microprocessors.

Thinking Machines succumbed to much the same affliction—ironically, as Hillis had seen collections of cheap commodity microprocessors as the wave of the future. But as Smarr puts it, "Danny was ahead of his time." The commodity processors available when Hillis was developing his Connection Machine were not fast enough to dethrone the vector supercomputers, and Hillis was forced to beef up his microprocessors with special vector processors. Too late for Hillis, the chips needed to realize his vision are now available, and they promise to shape the supercomputing industry as much over the next couple of decades as vector computers did during the past two.

New talents. They'll shape the uses of supercomputers as well. Karin of the San Diego center notes that MPPs—which are coming to dominate the top end of the market thanks to commodity microprocessors—have much more memory than traditional supercomputers because each processor is equipped with its own memory. The extra memory in turn suits these machines to data-intensive applications, such as imaging or comparing observational data with the predictions of models.

And at the Illinois center, Smarr reports that he has already seen new kinds of users taking advantage of the SMP machines. The vector computer, he notes, excels at analyzing large, organized systems because its speed derives from performing the same calculation on an entire string of numbers at once. In contrast, says Smarr, "The new machines like lots of disorganized stuff."

One researcher, for instance, used the center's Silicon Graphics supercomputer to model the electrical activity of the heart and its surroundings in the chest cavity, taking into account the density, conductivity, and other properties of the various tissues. The model helped him simulate heart attacks in order to design better pacemakers. Another scientist, creating an innovative index for a digital library, analyzed 400,000 abstracts of electrical engineering articles looking for co-occurrence patterns between key words that would identify commonly linked ideas.

Besides opening supercomputing to new users, the microprocessor-based machines are breaking down barriers between supercomputing and the rest of the computing world. The new commercial order, Smarr says, may not be as much fun for supercomputer designers, who have always been able to "take their time and charge what they want" for their specialized machines. But for users, it should open the way to the steep performance increases and price drops that users of desktop machines have long been accustomed to. At that point supercomputers, like personal computers and workstations, will have become commodities.

–**Robert Pool**

FLUID DYNAMICS

Mathematicians Open the Black Box of Turbulence

When the mathematician John von Neumann was trying to stir up enthusiasm for electronic computing in the late 1940s, he pointed to several important problems the newfangled machines could help solve. Among them were two that had come to the forefront of scientific interest during World War II: calculating the behavior of shock waves like those generated by atomic explosions, and generating long-range weather predictions. Both problems are rooted in the physics of fluid dynamics, and both ultimately demand an understanding of the knotty problem at the heart of fluid flows: the seemingly random motions known as turbulence, which mix material and energy through a moving fluid.

It's hard nowadays to imagine any need to drum up support for computers. But it's also hard to imagine that someone as smart as von Neumann could have been serious in suggesting that the kind of computer then being built—in essence, an oversized pocket calculator run on vacuum tubes—was capable of solving any problem involving turbulence, much less two of the hardest. Like chaos, to which it is closely related, turbulence defies easy understanding because it amplifies infinitesimal fluctuations into major effects. Indeed, 2 decades after von Neumann's prediction, Richard Feynman remarked that scientists had yet to understand what goes on in one of the easiest of such problems, that of turbulent fluid flow in a simple, cylindrical pipe.

But von Neumann may have just been farsighted. Computers and computer algorithms have now gotten to the point where shock waves and flows perturbed by mild amounts of turbulence can be accurately computed. Meteorologists now produce surprisingly accurate regional forecasts a week or more ahead of time and are coming to grips with details of local severe weather. Aeronautical engineers rely as much on workstations as they do on wind tunnels to design new airplane parts and surfaces. Many researchers have now set their sights on extending the capabilities of computational fluid dynamics, or CFD as the field is called, to flows where turbulence plays a larger, often dominant role. Even Feynman's pipe-flow problem is beginning to yield its secrets.

The progress has come partly from bigger, faster computers, but also from insights into the nature of turbulence, which have opened the way to computer algorithms that do a better job of calculating complicated fluid flows. Theorists have learned, for example, that keeping track of a fluid's tendency to rotate and form eddies—its vorticity—allows them to infer the kind of turbulence likely to roil it. In another shortcut, they're learning how to treat turbulence statistically, for example by studying the spatial correlations of turbulence's energy-dissipating effects. Says Paul Dimotakis, a professor of aeronautics and applied physics at Caltech, "[People] don't treat turbulence as a black box anymore."

On one level, to be sure, turbulence hasn't been a black box since the 19th century. That was when fluid dynamicists developed a complete mathematical description of fluid flow—including turbulence—in the form of a system of partial differential equations known as the Navier-Stokes equations. These equations recast Newton's laws of motion, ordinarily applied to solid objects, in terms suitable for a fluid, where velocity and density can change at every point in space. Ideally, understanding a turbulent flow is just a matter of plugging specific properties of the fluid into the Navier-Stokes equations and then cranking out solutions.

But that's easier said than done. Solving the equations of fluid flow exactly is out of the question, except for the very simplest examples. That's because the equations are nonlinear, which means that a change in one variable can lead to a disproportionate change in another. Consequently, researchers wanting to predict fluid flows in, say, aerodynamics, weather prediction, or astrophysics have to search for approximate solutions using a computer.

The most direct way to do so is by means of a "full" or "direct" numerical simulation, which in effect immerses a three-dimensional grid of points in the fluid and then keeps track of conditions such as pressure, density, and velocity at each grid point. The computer applies the equations to recalculate the conditions at each point in incremental time steps, leading to an approximate picture of the fluid's behavior. The finer the grid, the better the approximation.

Because they work directly from the Navier-Stokes equations, full numerical simulations are considered the gold standard in computational fluid dynamics. But full simulations still aren't feasible for the turbulent flows that researchers are most interested in. The problem is that keeping track of the minute eddies that carry most of the energy in a turbulent flow requires more grid points and tinier time steps than even the fastest supercomputer can handle.

In general, the amount of computation required increases with the cube of an entity known as the Reynolds number, a dimensionless combination of overall properties of the flow. (For pipe flow, it is the diameter of the pipe times the average velocity of the fluid divided by the viscosity.) In other words, doubling the Reynolds number, say by doubling the speed of the flow in a pipe, increases the computational workload by a factor of 8. Computers currently balk at doing such simulations for Reynolds numbers much above 10,000—but many of the important fluid-flow problems, such as airflow over an airplane wing or the flow of coolant through a nuclear reactor, involve Reynolds numbers in the tens of millions.

For such flows, notes Steven Orszag, an applied mathematician at Princeton University, "the number of grid points [needed] is of order Avogadro's number," or comparable to the total number of molecules in the fluid. "You're not going to fit that on any conventional digital computer in any way whatsoever."

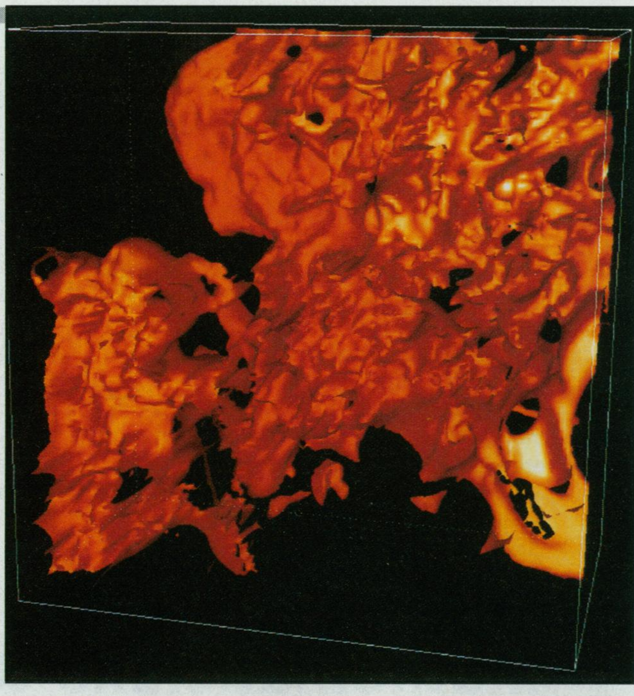

H. CATRAKIS AND P. DIMOTAKIS/CALTECH

Unruly fluid. A surface of constant concentration traces the turbulent mixing of a fluid jet into its surroundings. The image was reconstructed by computer from laser-induced fluorescence data.

Searching for short cuts. To predict the effects of turbulence in these kinds of settings, researchers have to look for shortcuts: models of fluid flow that combine theoretical insights and physical intuition to produce simplifying assumptions. "If you want to get to larger Reynolds numbers, you have to do some modeling and some theory," says Orszag. And that's where much of the progress since von Neumann's time has come, largely in recent years. As Dimotakis puts it, "There are new ideas that are being proposed all the time, and the complexity of these models is increasing tremendously."

Many of the ideas center on an approach known as large eddy simulation, or LES. In LES models, researchers use the behavior of eddies big enough to show up in the simulation to infer how eddies smaller than the grid are influencing the flow. Other ideas are based on a rival approach called transport modeling. Transport models replace the Navier-Stokes equations with equations derived from assumptions about how the fine-scale fluctuations of turbulence transport energy within the fluid.

Researchers gain or lose confidence in such models by comparing their predictions for low-Reynolds-number flows with the results of full numerical simulations—and even with observations of actual turbulent flows. Turbulence experts all have their favorite techniques, but they generally agree that theory-based modeling is finally realistic enough to take over where full numerical simulation leaves off. Indeed, says Alexandre Chorin, a mathematician at the University of California, Berkeley, "I think we are not many years away from a satisfactory solution" of how to compute turbulent flows.

Laboratory observations of turbulent flows are spurring theorists to refine their models even further. Orszag and colleagues Alexander Smits and Mark Zagarola in the mechanical and aerospace engineering department at Princeton, for example, are sharpening their understanding of turbulence in a cylindrical pipe—precisely the problem Feynman pointed to. "The goal is to tie experiment together with theory and computation," Orszag explains. Part of the project involves a huge piece of plumbing they fondly refer to as Superpipe, which circulates air compressed up to 240 times atmospheric pressure. The high pressure allows the Princeton researchers to generate flows with Reynolds numbers up to around 50 million.

They've found that at Reynolds numbers above 1 million—a realm important for practical problems—the nature of turbulence seems to change. Orszag explains that the relation between the friction generated as the air rushes past the wall of the pipe and the Reynolds number shows a new slope—a new scaling, it's called—at high Reynolds numbers. "The fact that we see this gorgeous scaling at Reynolds numbers between a million and 30 million says that there's something rich in the subject," says Orszag. "There's fundamental physics to understand." Orszag thinks the models he and his colleagues are developing, especially some recent analyses of energy dissipation, have a good shot at explaining the experimental results.

Dimotakis, however, suspects that many laboratory experiments will just add to the theorists' puzzlement. "The technology of measurement is probably undergoing a faster explosion and evolution than the technology of computation," he notes. He and colleagues, for example, are developing equipment that will be able to record fluid-flow images at the rate of a thousand frames per second. In a test run last year using a less capable, preliminary version of this system, they took pictures of a turbulent jet of liquid, hoping to capture details of the intricately wrinkled surfaces of constant temperature and pressure across which turbulent mixing takes place. "The experiment lasted a little over 10 seconds, and we have spent a year trying to understand what the three-dimensional surfaces look like," Dimotakis says.

At that rate, von Neumann's vision may have to wait a few years longer.

–**Barry Cipra**