SUPERCOMPUTERS

Off-the-Shelf Chips Conquer The Heights of Computing

It's been a rocky year for the supercomputer industry. Late last summer Thinking Machines Corp., the company that pioneered the use of massively parallel computers, filed for Chapter 11 bankruptcy. In December another leader in the massively parallel field, Kendall Square Research, followed suit. And in March Seymour Cray, the father of much of today's supercomputer industry, announced that his latest progeny, Cray Computer Corp., was also bankrupt. (His first company, Cray Research, remains in business.) The series of failures led many in the media and elsewhere to speak of the decline—or even the coming demise—of the entire industry.

The reality is somewhat different. True, the industry has been shaken to its core, but those closest to it see the turmoil as signifying not sickness but a dramatic metamorphosis. "We're at a real turning point in the history of supercomputing," says Larry Smarr, director of the National Center for Supercomputing Applications at the University of Illinois.

Throughout that history supercomputers have been built with specialized processors—pricey, custom-made computer chips that provided performance that couldn't be obtained any other way. But in just the past year or so that situation has changed. The mass-produced microprocessors used in computer workstations have become so fast that companies like Silicon Graphics and IBM are building state-of-the-art supercomputers by stringing together these somewhat less powerful but far less expensive components. And as fierce competition for the vast workstation market drives companies to invest heavily in R&D, commodity microprocessors are continuing to improve much faster than their custom-built counterparts.

The result, predict Smarr and others, will be a complete transformation in how supercomputers are designed and made. "Many people think that in 10 years the Cray philosophy [of expensive, specialized processors] is bound to lose out," says Malvin Kalos, director of the Cornell Theory Center. When it does, the world will have changed for scientists who depend on supercomputers for their work—among them researchers trying to build computer models of circulating fluids in earth science and astronomy. (See Articles beginning on p. 1365.)

Supercomputers powered by commodity microprocessors promise not only to give much more computing bang for the buck, but also to affect the sorts of questions that researchers ask. "Applications reflect what problems are amenable to the particular type of supercomputer," says Sid Karin, director of the San Diego Supercomputing Center. For instance, says Wayne Pfeiffer, deputy director for research at the San Diego center, the new generation can deal with large databases more easily than traditional supercomputers can. As a result, he says, "data mining"—looking for relations in large data sets—will likely grow in popularity in coming years.

The rapid rise of this new type of supercomputer has caught many people by surprise. Since the late 1980s, industry observers have been transfixed by the battle between traditional "big iron" machines and massively parallel computers. The vector supercomputer, the specialty of Cray Research, was for 2 decades the most powerful type of computer and the only option for researchers with truly large problems. Originally developed to handle the fluid dynamics calculations of nuclear weapons design, the vector computer derives its power from expensive, custom-built processors that perform calculations simultaneously on long strings of numbers—vectors—instead of adding, subtracting, multiplying, and dividing numbers two at a time.

But in the 1980s a new idea began to take shape: Instead of relying on a few high-power, high-cost processors, why not amass hundreds or even thousands of cheaper, low-power processors and provide the same computing power at a fraction of the cost? The chief proponents of massively parallel processing, or MPP, were a young MIT graduate student named Danny Hillis and, across the country, Charles Seitz at Caltech. Hillis formed a company, Thinking Machines, to develop and sell MPP machines, while Seitz worked with Intel on a competing line. Several other MPP firms followed.

It wasn't smooth sailing. In particular, dividing the memory among hundreds of processors creates a programming nightmare. A programmer must not only figure out how to allocate the data to get the best performance but also how to sew together the results from the many processors to get a complete answer. Gradually, however, a combination of hardware and software fixes has smoothed the rough edges, and for many users the contest between the MPPs and vector supercomputers is now a tossup. Albert Semtner of the Naval Postgraduate School in Monterey, California, for instance, has done his ocean modeling calculations on both types of machines and is happy to work on either. "Vector versus MPP—I don't really care," he says. "Whichever one will deliver the best performance."

Upstart chips. While vector and massively parallel computers have been contending for the supercomputing market, however, a new factor appeared on the scene that changed everything: the availability of extremely powerful commodity microprocessors, the mass-produced chips at the heart of computer workstations. Ten years ago, workstation microprocessors were far slower than the processors in supercomputers. The fastest microprocessor was rated at 1 megaflop—1 million floating-point operations per second—while Cray's processors hit 200 megaflops. (A floating-point operation is the equivalent of multiplying two 15-digit numbers together.) But now, while Cray's processors have improved by a factor of 10—to 2 gigaflops in the brand-new T90—the fastest microprocessor runs at 600 megaflops, with an 800-megaflop chip rumored for release later in the year by Hewlett Packard. And, predicts supercomputer expert Jack Dongarra of the University of Tennessee and Oak Ridge National Laboratory, "Next year we'll have a chip that's 1 gigaflop."

And that power is available for a fraction of the cost of a traditional vector processor. Unlike vector processors, which consist of complex collections of chips, the micropro-

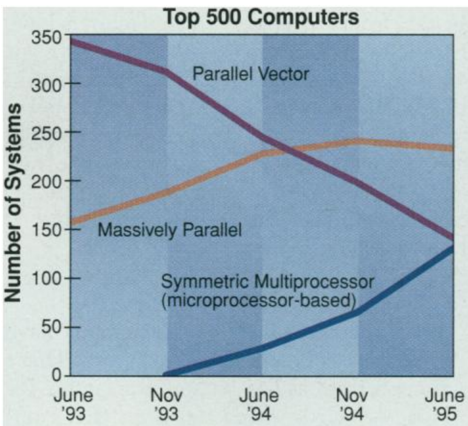

SOURCE: DONGARRA *ET AL.*

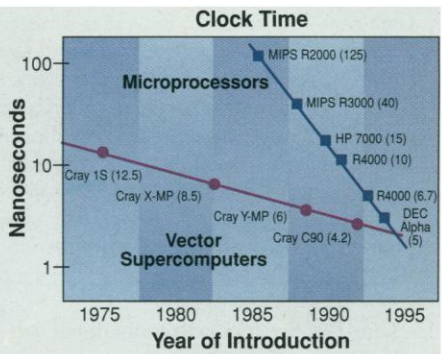

Chips are up. Microprocessors now rival vector processors in speed and number of top supercomputers.

SOURCE: F. BASKETT/SILICON GRAPHICS AND J. HENNESSY/STANFORD

Swimming Against a Supercomputing Tide

Burton Smith can recognize a trend as well as the next guy. He just doesn't believe it necessarily applies to him.

Now that most of the supercomputer world has decided the future lies with machines built with off-the-shelf microprocessors (see main text), only a few holdouts are still trying to compete using processors that are custom-designed. There are giant companies with established customer bases, like Cray Research. And then there is tiny Seattle-based Tera Computer Co., founded 7 years ago by Smith and Jim Rottsolk and now boasting 54 employees. It is testimony to the power of Smith's concept for a new parallel supercomputer that some observers give his still-unbuilt machine a fighting chance.

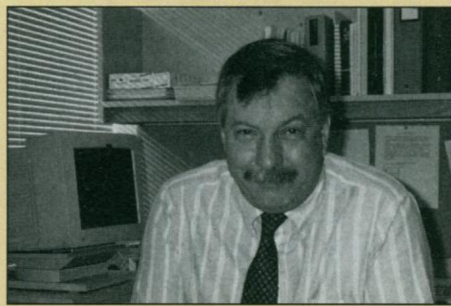

TERA COMPUTER

All in the timing? Burton Smith, Tera designer.

"It's a very clever architecture," says Bob Borchers, head of the National Science Foundation's supercomputing program. "It should be tested." And Sid Karin, director of the San Diego Supercomputing Center, says he hopes to be the first customer for Smith's supercomputer, which promises a combination of blistering speed and ease of programming.

What Smith has done is to design a solution to the problem of "memory latency" in massively parallel computers—the down time that afflicts a processor as it waits to get data from memory. As processors have gotten faster and faster, the time it takes to transfer data from memory to the processor has increasingly become a limiting factor on a computer's speed. And the problem is particularly acute for massively parallel computers, in which each processor must request data from the memories of other processors scattered around the computer.

Smith's approach is deceptively simple: guarantee that the processors always have something to do while they wait. Like a secretary who uses down time to do some filing or to type letters, a processor in the Tera computer will switch to a second task—and a third, fourth, and so on—while it is waiting for the data it needs to complete previous tasks.

That strategy, together with a scheme for interconnecting the multiple memories and processors, will free programmers from having to worry about how to divvy up data most efficiently among all the memory locations. And it should open the way to stringing together an almost unlimited number of powerful processors—the initial generation will be rated at 1 gigaflop apiece—without the headaches and inefficiencies of today's massively parallel machines. For users, Smith says, this will mean "high performance with much easier programming and broader applicability." Those benefits won't come cheap, however: The Tera machines will cost from \$5 million to \$40 million, depending on the number of processors.

The computer's performance has been tested by simulations on other computers, but Smith has not yet built one. "We're hoping we can get the money together," he says. That may not be easy, Borchers notes. Failures of other specialty supercomputer manufacturers have "made it almost impossible to raise venture capital if you've got a good idea for a new computer."

But Smith thinks the rapid improvement in commodity microprocessors is slowing. Soon, he says, improved scheduling and communications within a supercomputer rather than raw processor speed will set the pace of computing. And then his custom processor may look less like a throwback than like the key to the future.

–R.P.

cessors are designed for mass production. Each is a self-contained miniature computer made on a single bit of silicon only a centimeter or so across, in a process honed by 2 decades of experience making integrated circuits. And the R&D costs for each microprocessor are spread over hundreds of thousands of chips instead of the hundreds of vector processors that Cray or any other traditional supercomputer manufacturer may make each year.

The inexpensive power of the microprocessor has made possible a new type of parallel supercomputer. Known as the symmetric multiprocessor (SMP), it connects as many as 20 commodity microprocessors to a single large memory instead of distributing the memory among its processors as MPPs do. In November 1993 not a single SMP could be found on a semiannual Top 500 list of the world's most powerful computers, compiled by Dongarra and Hans Meuer and Erich Strohmaier at the University of Mannheim in Germany. By June of this year 130 of the 500 were SMPs.

Silicon Graphics Inc. is leading the charge. Incorporating as many as 18 microprocessors, each rated at 300 megaflops or more, the company's supercomputers have peak speeds of more than 5 gigaflops. That's about a third as fast as Cray Research's 16-processor C90, which has been that company's flagship. And their modest price tags—\$1 million for a 16-processor machine, compared with \$30 million for the C90—have made them an instant hit with scientists, engineers, and businesses.

The SMPs are powerful enough for all but the most flop-hungry users, and for those who need something more, it's possible to create an SMP cluster, in which several machines are connected with fiber optic cable and their operations coordinated by software. The Illinois supercomputing center, for instance, has linked four eight-processor and two 16-processor machines from Silicon Graphics into a cluster with a peak speed of nearly 20 gigaflops. Smarr, the center's director, likes the price, flexibility, and future promise of SMPs so much that he plans to use them exclusively within 2 years. The center's older Crays were "cut up for scrap" at the beginning of this year, he says, and its massively parallel computers will be phased out by 1997.

Despite Smarr's move, however, MPPs are not on the way out. To the contrary, the availability of cheap, powerful microprocessors has given massively parallel computers a big boost in their competition with vector

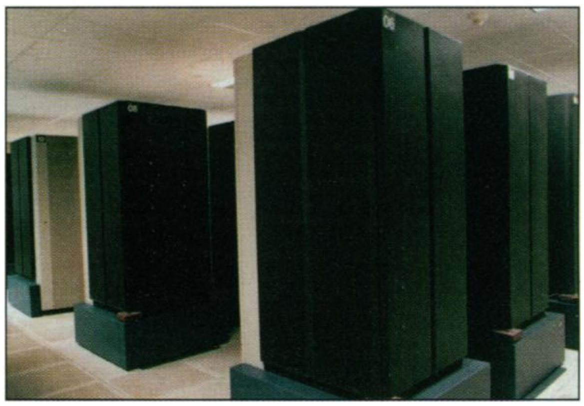

CORNELL THEORY CENTER

Towers of power. The IBM SP2, like this one at Cornell, is one of a new breed of massively parallel computers incorporating hundreds of workstation microprocessors.

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 cooccurrence 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