THE RISE and Fall of Thinking Machines

THE BRILLIANT START-UP THAT IGNITED AN INDUSTRY NEVER GRASPED THE BASICS

by Gary Taubes

Some day we will build a thinking machine. It will be a truly intelligent machine. One that can see and hear and speak. A machine that will be proud of us.
—From a Thinking Machines brochure

In 1990 seven years after its founding, Thinking Machines was the market leader in parallel supercomputers, with sales of about $65 million. Not only was the company profitable; it also, in the words of one IBM computer scientist, had cornered the market "on sex appeal in high-performance computing." Several giants in the computer industry were seeking a merger or a partnership with the company. Wall Street was sniffing around for an initial public offering. Even Hollywood was interested. Steven Spielberg was so taken with Thinking Machines and its technology that he would soon cast the company's gleaming black Connection Machine in the role of the supercomputer in the film *Jurassic Park,* even though the Michael Crichton novel to which the movie was otherwise faithful specified a Cray.

In August of last year Thinking Machines filed for Chapter 11. It had gone through three CEOs in two years and was losing money at a considerably faster rate than it had ever made it.

What caused this high-flying company to come crashing to earth? The standard explanation is that Thinking Machines was a great company victimized by the sudden cutbacks in science funding brought about by the end of the cold war.

SETH RESNICK

The truth is very different. This is the story of how Thinking Machines got the jump on a hot new market—and then screwed up, big time.

Until W. Daniel Hillis came along, computers more or less had been designed along the lines of ENIAC. In that machine a single processor completes instructions one at a time, in sequence. "Sequential" computers are good at adding long strings of numbers and at other feats of arithmetic. But they're seriously deficient at the kinds of pattern-recognition tasks that a two-week-old puppy can master effortlessly—identifying faces or figuring out where it is in a room. Puppies can do that because their brains—like those of all animals, including humans—are "massively parallel" computers. Instead of looking at information one jigsaw-puzzle piece at a time, a brain processes millions, even billions, of pieces of data at once, allowing images and other patterns to leap out.

While a graduate student at MIT's Artificial Intelligence (AI) Lab, Hillis, whom everyone knows as Danny, had conceived of a computer architecture for his thesis that would mimic that massively parallel process in silicon. Hillis called the device a "connection machine": it had 64,000 simple processors, all of them completing a single instruction at the same time. To get more speed, more processors would be added. Eventually, so the theory went, with enough processors (perhaps billions) and the right software, a massively parallel computer might start acting vaguely human. Whether it would take pride in its creators would remain to be seen.

Hillis is what good scientists call a very bright guy—creative, imaginative, but not quite a genius. He is also an inveterate tinkerer, whose work has always been more fascinating than practical. On the fifth floor of Boston's Computer Museum, for instance, is a minimalist computer constructed of fishing line and 10,000 Tinkertoy parts. Hillis built it to play and win at tic-tac-toe, which it invariably does. His other work includes a robot finger that can differentiate between a washer and a screw but is flummoxed by a piece of gum; a propeller-driven jumpsuit that allows its wearer literally to walk on water; and a home robot constructed of paint cans, lightbulbs, and a rotisserie motor.

At the AI Lab, Hillis had become a disciple of legendary AI guru Marvin Minsky. The two were determined to build a connection machine as a tool with which to develop software programs for artificial intelligence. Because the cost would be prohibitive for a university laboratory, they decided to form a company. They went looking for help and found Sheryl Handler.

Handler had participated in the start-up of the Genetics Institute, a Harvard-based genetic-engineering firm. Her background was eclectic: she had studied interior design, held a master's degree in landscape architecture from Harvard, and at the time was pursuing a doctorate in city planning at MIT. She was also running her own nonprofit consulting firm, specializing in third-world resource planning. She had a taste for classical music and a fine appreciation for style. She'd even been the subject of a Dewars Profile that ran with the quote "My feminine instinct to shelter and nurture contributes to my professional perspective."

Handler also had a talent for cultivating friendships with brilliant and famous people. One of her Genetics Institute colleagues later called her a "professional schmoozer." She quickly proved her usefulness by connecting the people who would build the Connection Machine with CBS founder William Paley. Hillis, Minsky, and Handler pitched the idea to Paley and CBS president Fred Stanton in a meeting to which Hillis wore his customary jeans and T-shirt. Still, he managed to impress the television moguls, who with others eventually agreed to kick in a total of $16 million to the venture.

In May 1983, despite the lack of a business plan, the company was founded and took up shop in a dilapidated mansion outside Boston that once was owned by Thomas Paine, the author of the Revolutionary War pamphlet *Common Sense*. Hillis and Handler called their new company Thinking Machines because, says Hillis, "we wanted a dream we weren't going to outgrow." As it turned out, there was never much danger of that.

Brain man: Danny Hillis wanted to build machines that could achieve emotions—not machines that would solve companies' problems.

The new company's managers immediately got into a disagreement over the market for supercomputers. Hillis and Handler (Minsky quickly became a figurehead at the company) wanted to design a machine strictly along the lines of Hillis's thesis, a machine that would have its maximum impact as a research tool for scientists studying artificial intelligence. (Hillis envisioned his machine eventually becoming a sort of public-intelligence utility into which people would tap their home PCs, thereby bringing artificial intelligence to the world.) Howard Resnikov, a research director recruited by Minsky, on the other hand, argued for a more flexible architecture that could support whatever style of computing was needed to solve real-world problems. After all, the more problems the machine could solve, the more sales prospects there would be.

For a year, while the argument went on, the company did nothing. Finally, Handler and Hillis won out. "We had all sorts of reasoned discussions," says Resnikov, "and then emotional decisions were fundamentally made by Sheryl and Danny." Resnikov lasted another two years before he quit. Emotional decision making would last almost until the company fell.

In the first few years it didn't seem to matter. Thinking Machines didn't need to make good business decisions because it had the Defense Advanced Research Projects Agency. A research arm of the Defense Department, DARPA was looking for computer architectures that would enable tanks, missiles, and other weapons to recognize enemy targets and understand spoken orders. In 1984 Hillis and his colleagues at Thinking Machines repackaged Hillis's thesis and pitched it to DARPA. The agency responded by offering the company a multiyear $4.5-million contract. Now all Thinking Machines had to do was build one of the world's fastest computers in two years' time.

The company promptly went on a hiring binge. Its prime hunting grounds were the computer-science departments of MIT, Carnegie-Mellon, Yale, and Stanford—which happened to house four of the world's leading AI labs. Everyone, from programmers to administrative assistants, had to be interviewed by Handler, who had a very specific, if mysterious, idea of who would be good enough to work for Thinking Machines. (Many researchers later reported that once they were hired, they never got to speak to Handler again—even when they were alone with her in an elevator.)

In fact, Thinking Machines was becoming Handler's aesthetic creation as much as the Connection Machine was Hillis's. In the summer of 1984 the company moved into its new home—the top two floors of the old Carter Ink Building in Cambridge, Mass., a few blocks from MIT. Handler personally oversaw the design of the office space, insisting that each office be painted a different and specific color. Huge open spaces were created to stimulate idea sharing and creativity. A plush cafeteria was put in, complete with a gourmet chef. Couches were scattered throughout the offices so that researchers could take naps or even sleep there overnight, which many of them did. And

SETH RESNICK

the soft-drink machine was wired to a terminal. Researchers who wanted a drink simply typed in their choice.

In short, Thinking Machines was becoming a hacker's paradise. The thinking, says Lew Tucker, one of the company's research directors, was that "if they were fed, they'd practically live at Thinking Machines." If Hillis disapproved, he didn't make it known. Having taken to commuting in an antique fire engine, he could hardly play the pragmatist to Handler's stylist.

In May 1985, Thinking Machines announced the impending completion of the first Connection Machine, the CM-1. The announcement would be made on the third floor of the Carter Ink Building. Handler had every surface on the new floor repainted a slightly different shade of mauve. When it was done, she wasn't satisfied. So she had her researchers and scientists paint it again.

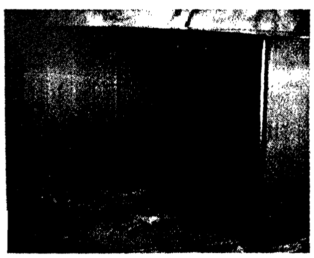

Dream machine: The CM-1 contains thousands of tiny processors instead of one big one.

The CM-1 was an AI researcher's dream. Unfortunately, few AI labs could afford a $5-million computer, and, as Resnikov had predicted, hardly anyone else was interested. When it came to general scientific computing, the CM-1 was "a dog," in the words of Gordon Bell, a computer guru and architect of the famous VAX computer at Digital Equipment Corp. It had no facility for running FORTRAN, the de facto standard computer language of science; nor could it do what are known as "floating-point operations," the operations that manipulate numbers in scientific computation.

Thinking Machines sold seven CM-1s, but only because DARPA brokered and subsidized most of the deals. If the company was going to stay in business, it would need a machine that could pull its weight outside AI research. Unfortunately, according to Resnikov, the decision to tailor the CM-1 to the AI "nonmarket" cost Thinking Machines three years in the real-world marketplace.

In April 1986, Thinking Machines announced the arrival of the CM-2, a machine the scientific community actually could use. The CM-2 was able to run FORTRAN and to do floating-point operations. It was also a piece of work artistically: a five-foot cube of cubes—done up in what Thinking Machines employees called "Darth Vader black"—in whose innards red lights flickered mysteriously. But the machine's exotic massively parallel technology still needed special software, which meant its users had to learn new programming techniques. The CM-2 might be more like the human brain than a sequential computer like the Cray was, but scientists knew how to write programs for the Cray. Many of Thinking Machines' first customers, says Dave Waltz, who ran the company's AI group, did most of their computing on the floating-point processors, ignoring the 64,000 single-bit processors.

As a result, there still wasn't much of a market for Connection Machines. But thanks to the support of DARPA, which continued to broker deals, Thinking Machines didn't have to seriously contemplate building a machine that had a natural market. "Our charter," says Tucker, "wasn't to look at a machine and figure out the commercial profit. Our charter was to build an interesting machine." But the definition of *interesting* would soon change.

In the late 1980s, DARPA and the Bush administration, having accepted the fact that the end of the cold war had reduced the urgency for military supercomputing, came up with a new challenge for parallel computing. They began to talk about solving what D. Allan Bromley, the president's science adviser, dubbed "grand challenge" scientific problems: modeling the global climate, analyzing the folding of proteins, mapping the human genome, predicting earthquakes, revealing the nuances of quantum mechanics. The problems didn't require artificial intelligence, just enormous computing power.

The official name of the new project was the High Performance Computing and Communication (HPCC) program, and DARPA was the lead agency, with a projected budget of several billion dollars through 1996 to accomplish its goals. At the top of the list: building a computer capable of a teraflop—a trillion floating-point operations per second.

Connection machine: CEO Sheryl Handler had lots of contacts. But could she run a company?

SETH RESNICK

Not surprisingly, Thinking Machines

had an inside track on getting a chunk of the projected budget. While other computer companies were out wooing customers, Handler had been cultivating a friendship with Bromley. As soon as Thinking Machines promised it would have a scaled-down version of a teraflop machine ready by 1992, the agency awarded the company an initial contract of $12 million.

In the meantime, several computer companies were exploring a new technology—a compromise between the comfort of sequential computing and the performance of massively parallel machines. A sort of "moderately parallel" design, the technology entailed stringing together a smaller number of the powerful, cheap off-the-shelf microprocessors used in PCs and workstations—rather than the thousands of highly customized but less powerful processors used in the Connection Machines—into a single supercomputer that would work with existing software.

The cost advantages of using off-the-shelf chips, as well as the functional advantage of running existing software, seemed overwhelming—especially considering the fact that few customers outside the tiny AI community had much interest in Thinking Machines' massively parallel design. Even Hillis eventually came around and chose the moderately parallel design for the company's next generation of machine. Unfortunately, the old dream died hard: the decision came only after 18 months of internal bickering. Once again, the company was off to a late start.

What's more, there were signs that the company was still chasing the wrong market. Industry analysts in 1992 were projecting that the growth in supercomputers was not in science but in business applications—in particular in what's known as "database mining," an area that could well become, as IBM parallel-computing expert Art Williams put it, "the killer application" for parallel computers. With the country in a recession, businesses needed every competitive advantage they could get, which meant knowing their customers' preferences and buying habits in intimate detail. They had begun to collect all conceivable data and were feeding them into their mainframes, looking for any insight that would help them maximize profits. But it sometimes took mainframes hours, even days, to churn out the answer to a single question. So large companies were beginning to check out parallel computers.

'Vendors handed money by the government have no interest in solving customers' problems,' growled KSR's Burkhardt.

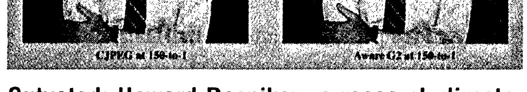

Outvoted: Howard Resnikov, a research director recruited by Minsky, argued for a more flexible architecture that could support whatever style of computing was needed to solve real-world problems.

PHOTO COURTESY OF AWARE, INC.

In fact, Thinking Machines had sold two Connection Machines to American Express. That got management at Thinking Machines talking about starting a business supercomputer group, an idea that appears at first to be a no-brainer. But at Thinking Machines the idea got stuck in endless discussions. Hillis and Handler already were bitter about having to target general scientific computing rather than artificial intelligence; they weren't about to jump on the idea of servicing mere merchants. Hillis later complained about the injustice of a world where "the real money is in handling Wal-Mart's inventory rather than searching for the origins of the universe."

Nonetheless, thanks to DARPA, Thinking Machines went into the black for the first time. In 1989 the company reported a profit of $700,000 on revenues of $45 million. Handler promptly signed a 10-year lease with the Carter Ink Building for a whopping $6 million a year—about $37 a square foot. (Lotus Development Corp., which was virtually across the street from Thinking Machines, was paying $8 a square foot.) Thinking Machines also hired another 120 employees, bringing the total to over 400. Meanwhile, the company had developed an image as one of the leading high-tech companies in the country. It was, says Stephen Wolfram, who founded the highly successful software company Mathematica, "the place that foreign trade delegations would come to visit to see where American business was at these days."

Yet competition was looming. Cray Research launched a crash program in 1990 to get a massively parallel machine on the market within two years. IBM was doing the same. Even Fujitsu Limited, one of Japan's major supercomputer manufacturers, was in the process of opening a parallel-computing lab, looking toward marketing a 1,000-processor machine.

If there was ever a time that Thinking Machines could, and needed to, put itself on a solid financial and competitive foundation by merging with a deep-pocketed company or by going public, it was now. But Handler nixed all deal making. She felt the company could get a wildly successful teraflop machine out on its own.

As the company forged ahead with its frantic effort to bring the new machine out on time, the corporate culture started to shift from openness to paranoia. Employees weren't allowed to discuss the machine with one another in the cafeteria. Customers were kept in the dark. The new machine was dubbed the CM-5, to foil hackers acting as corporate spies who presumably would be rummaging through the company's files looking for a nonexistent CM-3.

Thinking Machines announced the CM-5 in October 1991. Hillis claimed it had the highest "theoretical" peak performance of any supercomputer ever, if you added enough processors to it. The reality: at the time completion of the CM-5 was announced, the machine was slower than its predecessor, the CM-2. Among other problems, the standard chips the company had chosen weren't ready, so some machines had to ship with slower, earlier-generation chips. Meanwhile, competitors like Intel, Kendall Square Research (KSR), MasPar Computer, and nCube were starting to ship faster supercomputers. More than ever, Thinking Machines was depending on its DARPA edge to move its products.

Then, in August 1991, as DARPA was about to start the process of determining which supercomputer vendors would win the lion's share of its planned spending spree, the *Wall Street Journal* broke the story that the agency had been playing favorites. It turned out that DARPA had subsidized—sometimes to the tune of the entire purchase price—the sale of some 24 Connection Machines in recent years. The subsidies added up to a gift to Thinking Machines of $55 million—20% of the company's lifetime revenues to that point.

DARPA had greased Intel's supercomputing wheels too but had left the rest of the supercomputer industry to fend for itself. And now the other players were howling. Perhaps the clearest and most damning criticism came from KSR founder Henry Burkhardt: "Vendors handed money by the government have no interest in solving customers' problems," he growled.

An embarrassed Bush administration put an end to Thinking Machines' DARPA gravy train. For the first time the company had to sell its machines on their merits in an open market. At the end of 1992, Thinking Machines reported a loss for the year of $17 million. The CM-5 wasn't selling, and the company was hemorrhaging money. Hillis was no longer spending much time in the office. The first round of layoffs had started. Salaries were frozen. Requests for new laptop computers were being denied.

Meanwhile, Handler had an enormous marble archway installed in the atrium of the Carter Ink Building. When a national supercomputer conference was held in Seattle, she decided to stay in San Francisco and commute to Seattle from the swank Stanford Court Hotel. She commissioned a $40,000 logo design for a CM-5 sweatshirt and then rejected it. While the company was sinking, she focused her attention on putting out a cookbook with recipes from the company's now-infamous cafeteria. Increasingly paranoid, she had a video camera aimed at her personal parking spot and, by some accounts, made people take meetings with her in her parked car. She hired a bodyguard, telling her colleagues that she had received death threats.

Some members of Thinking Machines' board suddenly seemed to realize that the person who had been running the company all those years had no business skills. The board discussed dumping Handler, but she managed to get her biggest enemies there kicked off.

In early 1993 a new president was brought in, but Handler, who remained CEO, quickly got rid of him. Later in the year a lawyer named Richard Fishman was hired as president. Fishman was a longtime friend of Handler, but when he realized that no outsider would fund the sinking company while Handler remained at its helm, he engineered her ouster.

Fishman focused the company on the business market and began looking for a partner. Sun and IBM were interested, says Tucker, but weren't willing to take on Thinking Machines' mounting debt, which included six more years of rent at the Carter Ink Building, a $36-million commitment.

In mid-August, Thinking Machines filed for bankruptcy protection, and Fishman resigned. Soon Hillis himself left the company that had been founded around his thesis. Thinking Machines would reemerge as a small software firm selling programs for its former competitors' parallel computers.

As late as 1989, says Fishman, Thinking Machines was still three years ahead of the rest of the world in parallel-processing technology. "While others caught up," he says, "Thinking Machines was losing time, losing customers, and not moving on to the next generation." Had the CM-5 been built without the miscues and the wasted time, the company might have gone on to live up to its considerable promise. But, as one of the company's senior scientists would later put it, what if pigs could fly? ❍

Gary Taubes is a New York–based science and technology writer. His most recent book is Bad Science: The Short Life and Weird Times of Cold Fusion

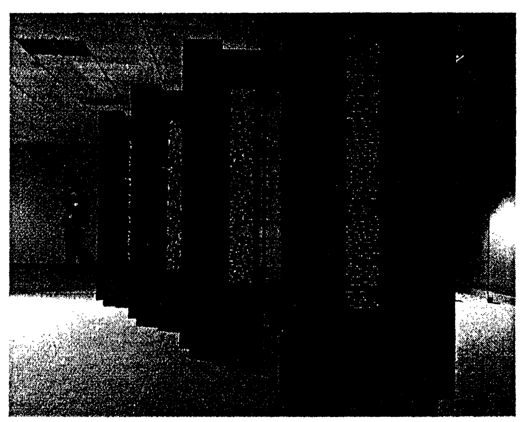

SETH RESNICK

Swan song: Without DARPA's push, the CM-5 didn't sell.

LAYING A BET TO REST

Gordon Bell claims to be the winner of the famous bet on supercomputers that he made with Danny Hillis five years ago. But the bigger issue in the dying supercomputer market is not who has the fastest machine, but whether anyone will still be making Big Iron in the future.

BY WILLIE SCHATZ

AS THE BELL tolls at the end of 1995, it tolls for Danny Hillis.

Hillis has been on one long, strange trip in the five years since Gordon Bell bet him a crow-eating essay about the future of supercomputing. Hillis, the founder and former chief scientist of parallel processing pioneer Thinking Machines Corp. (TMC), Bedford, Mass., no longer holds that position. He has returned to his roots at the Massachusetts Institute of Technology in Cambridge, Mass., this time at the younger and more hip Media Lab rather than the Artificial Intelligence Lab, dominated by Marvin Minsky, where Hillis first cobbled together a design for a massively parallel computer.

TMC itself is just emerging from the dark side of a Chapter 11 bankruptcy pleading. But to do so it had to die as a supercomputer maker and resurrect itself as a software company. The twin forced march of Hillis and TMC symbolizes the seismic shift that has rocked the traditional supercomputing market.

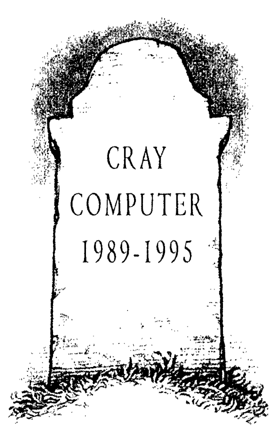

Bell didn't exactly predict that shift, but he did bet on Hillis' comeuppance. Bell, the legendary computer designer who created the PDP and VAX series of machines at Digital Equipment Corp., Maynard, Mass., and who is now dispensing wisdom for Microsoft Corp. from his office in Los Altos, Calif., bet Hillis that massively parallel supercomputers might not actually be the greatest revolution to hit the scene since the invention of rock and roll. Specifically, in 1990, Bell bet Hillis that in the last quarter of 1995, the highest number of sustained MFLOPS (millions of floating point operations per second) would be generated by a machine with fewer than 100 processors, rather than by a machine with many processors (more than 1,000). The wager concerns only supercomputers, or Big Iron machines costing more than $1 million and used for scientific purposes. That bet was chronicled in the January 1992 issue of UPSIDE.

The judge, jury and keepers of the bet are John Hennessey, a professor of electrical engineering and computer science at Stanford University, Stanford, Calif., and David Patterson, a professor and chair of the computer science division at the University of California, Berkeley. The loser has to 'fess up his humbling defeat to the world in writing.

The prohibitive favorite when the bet was made, Hillis is now such an underdog that Las Vegas bookies wouldn't touch him. "The only reason I might lose is if people get a few huge machines only doing floating point operations and lock them in a room and they just sit there and grind out MFLOPS," says Bell. "They'd be grinding out a shitload of numbers just to win the bet." A Cray T3E is still probably the most powerful machine not on wheels. "The [massively parallel] stuff is a bust," says Bell. "There's no market for it because no one can get applications for the machines."

If Bell, whose volume and frequency of pontification make the late Howard Cosell seem mute by

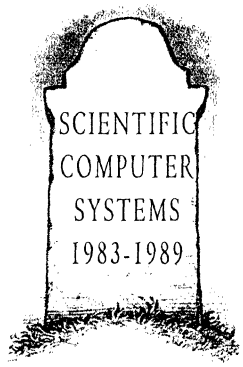

last year or so, three more Big Iron companies joined the list of roadkill: TMC (at least as a Big Iron company), Kendall Square Research and Cray Computer.

The high-end market, in fact, never ended up as anything more than a tiny blip on serious financial radar screens. The Smaby Group, Minneapolis, estimates last year's entire market for high-performance (scientific/technical/engineering) computing at $2.05 billion. The company projects that in 1999 the market will barely creep over $3 billion. But the true, high-end supercomputer segment of that market is beginning to crumble away. Smaby Group projects that the top end of that market—machines costing more than $5 million apiece—will decline by 6 percent in that time. Similarly, Chris Willard, manager, high-performance technology, in the Mountain View, Calif., office of Framingham, Mass.-based International Data Corp., predicts that by 1999 the revenue for the traditional supercomputer market will drop to $767 million from the 1994 total of $877 million, a negative 2.6 percent compound annual growth rate.

Market watchers, however, are not yet ready to bury the entire Big Iron market. Willard concurs that the high-end market "is not really dead, but it's certainly not growing."

But if it's not a dinosaur, it's at least a white rhinoceros at the top of the Endangered Species

"FROM THAT CHANGED MARKET STANDPOINT, I THINK WE BOTH LOST," SAYS HILLIS. "THE BET MAY FALL IN THE DEAD ZONE."

comparison, had said that three years ago, even his most ardent admirers would have questioned his sanity. Until recently, many people considered massively parallel machines to be the savior of high-performance computing.

But the trip to the future turned out to be even stranger than Bell might have predicted. It's not just the massively parallel machines that have disappointed their creators and investors. We're talking upheavals in the high-performance computing world—alias Big Iron—that no Richter scale could measure. And the aftershocks have only just begun. "There is no future for the Big Iron systems," declares Michael Burwen, director of the Palo Alto Management Group, a high-performance computing market research company in Palo Alto.

Now even Hillis is humbled. "The big surprise is [that] the supercomputing market basically doesn't exist anymore as a definable market," he said last October at the Media Lab's 10th anniversary party. "It's very clear that the dinosaurs are dying."

Many of the carcasses, in fact, already line the Information Highway: Alliant, Scientific Computing Systems, Multiflow, Floating Point Systems and Supercomputing Systems Inc. And in just the

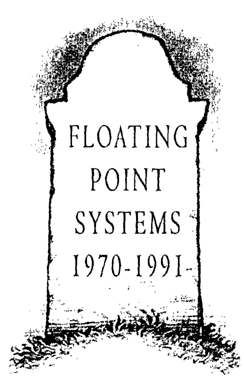

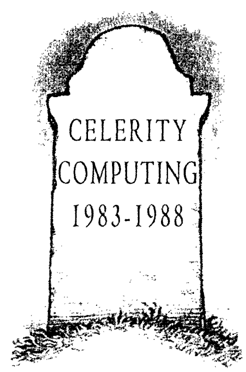

List. That's a huge contrast from the days when Hillis and Bell were betting on which type of supercomputer would get bragging rights to the title of fastest in the world. After all the hype and promise, most of the racehorses collapsed before reaching the finish line, and Hillis and Bell have found themselves betting over tombstones. "From that changed market standpoint, I think we both lost," says Hillis. "The bet may fall in the dead zone."

THE BIGGER THEY ARE...

It was one hell of a ride, though. A few short years ago, supercomputers and minisupercomputers were as hot as the World Wide Web is today. Venture capitalists were pouring in money and new companies were sprouting up every few weeks.

But today, Cray Research Inc., Eagan, Minn., still towers over the traditional supercomputing business. And perhaps it should. Cray Research has defined the excitement over supercomputing since it was founded in 1972. As soon as the first Big Iron box sprang from Seymour Cray's head, it was crystal clear that the thing had *attitude*—simultaneously mysterious, ethereal and fascinating. There was a unique majesty to the word *supercomputer.*

$40 MILLION VAPORWARE, ANYONE? How's this for a hot prospect?

"The Company is a development stage enterprise that had an accumulated loss of approximately $9.6 million as of June 30, 1995. The Company has experienced net losses in each year of operations and expects to incur substantial further losses while it builds its MTA (Multithreaded Architecture System) system prototype and commences production, and possibly thereafter. The Company has had no revenue or earnings and does not expect to recognize revenue from the sale of its MTA system sooner than the second half of 1996, if ever."

Sounds like the kind of statement that just builds confidence, doesn't it?

Well, that's the wording from the September 25 IPO statement of Seattle-based Tera Computer Co. The company went public in order to raise money for its nonexistent supercomputer, whose selling price it expects to set between $5 million and $40 million.

Since its December 1987 inception, Tera has spent $27 million to develop the MTA system. More than $18 million of that was a gift from the Advanced Research Projects Agency (ARPA). Still, the prospectus notes that the MTA system "has been subject only to computer simulation and the Company has not yet built its initial prototype."

So who would buy stock in such a company? Tera put 850,000 units up for sale, each consisting of two shares and one warrant. The stock price was $12 per unit. The warrant entitles the holder to purchase, at any time over a five-year period starting September 25, 1995, one share of common stock at $7.20 per share through March 24, 1998, and at $8.40 per share through September 24, 2000, when the warrants expire.

A lot of people bought. The offering was a monster success, raking in $8.55 million ($9.9 million if the underwriter's over-allotment is exercised).

"We didn't have much difficulty selling it because we've got a very exciting and contrarian story to tell," says Jim Rottsolk, Tera's president, CEO and CFO. "It's unusual to go public when you have no revenue. But it's the promise of the product that interests investors and the government."

Indeed, the government seems especially interested. Aside from helping to fund the company's research, ARPA is first in line to buy its products—once they're done. Last January the agency signed a contract with Tera that provides ARPA with options to purchase up to $20 million of MTA systems for early evaluation over the next three years. Tera is also negotiating a contract with ARPA to jointly fund the development of certain components of its next-generation MTA system.

In addition, the San Diego Supercomputer Center (SDSC) has submitted a proposal to ARPA to purchase Tera's first production MTA system. ARPA has told Tera that it plans to exercise an option under the January 1995 contract to buy an MTA system to place at SDSC. (Calls to ARPA officials seeking comment about this expenditure of taxpayers' funds were not returned.)

This kind of backing indicates there's a new guru in town. "This is absolutely startling," says Bob Stern, a Washington-based IT consultant. "There seems to be a serious national investment in Burt Smith [Tera's founder, chairman and chief scientist]. Seymour Cray was in the same position six months ago—he needed $25 million but couldn't get it."

Stern also notes that this seems to be government supercomputer business as usual. "This obviously shows that the days of the 'state computer' [Gordon Bell's term for companies kept alive by the federal government] aren't dead. It's amazing that Smith pulled this off."

Get real, counters Patrick W. Grady, senior vice president in the San Francisco office of Rochester, N.Y.-based H.J. Meyers & Co. Inc., the offering's underwriter. Grady claims that he has been approached by many of Tera's competitors several times and hasn't given them a damn thing. "I'm completely unconvinced that Cray Research holds the future of high-performance computing," he adds.

Tera's management insists it does have the answer. According to Vice President Gerald Loe, customers are desperately seeking general purpose scalable parallel machines with large-scale memories. Tera promises to create such machines because it claims it can solve the memory latency problem that slows down other architectures. Grady contends that cure extends all the way to the desktop.

"We know the machine's not built," Rottsolk concedes. "We know no one knows if the hardware works. We know that everyone could lose all their money. But we strongly believe that we're going to release our prototype in the first quarter of 1996 and deliver it by June 30.

"Right now, though, you've got to take it on faith because the product doesn't exist." Amen.—*W.S.*

It crunched numbers and solved problems that previously had been the stuff of dreams.

The genius of supercomputers, Seymour Cray, changed the world—and made one hell of a lot of money doing it—as the master of the bipolar-logic superprocessor. When the supercomputing universe centered around single-processor systems, he who made the fastest CPU ruled. In that domain, Seymour Cray was unbeatable.

But in recent years, even the mighty Seymour has begun to look outmoded. While most people recognized that killer microprocessors are taking over the field, Cray refused to yield. He tenaciously—some say stubbornly—clung to his mission of making the fastest, most powerful superprocessor. In 1989, the pioneer left the company he founded to start another venture, Cray Computer Corp., to create even more powerful superprocessors from gallium arsenide.

But the superprocessor crusade, which had worked so well for so long, had become an anachronism. In 1990, Lawrence Livermore Laboratory agreed to pay $42 million to be the proving ground for an eight-processor Cray-2 and Seymour's newest creation, a 16-processor Cray-3. But the Cray-3 was a stillborn prototype.

Because of Seymour's reputation, the mere possibility of completing the Cray-4 kept the company alive long after less tolerant—or hero-worshipping—creditors would have ripped out its life support systems. Death came last April, not with a bang but a whimper when financing finally dried up.

Steve Chen, another supercomputer genius from Cray Research, also ran away to fight another day. Reveling in his X-MP and Y-MP glory, he formed Supercomputer Systems Inc. (SSI) after convincing IBM that he had the cure for its supercomputing sickness. After four years and possibly as much as $250 million from Big Blue, however, Chen proved incapable of walking the walk. He allegedly produced a prototype based on superprocessors, but he never sold a machine. So a sadder but wiser IBM, now a self-made supercomputing heavyweight, pulled the plug in 1992.

Unlike his mentor, however, Chen recently demonstrated that he may not still be crazy after all these years. Funded entirely by MCSB Systems PTE Ltd., a Singapore-based technology conglomerate, Chen Systems Corp. (formerly SuperComputers International) in Eau Claire, Wis., began beta testing its Pentium Pro-based Chen 1000 server line last April. Its eight-processor machine was released on September 18, mostly to yawns. The company at the time claimed it had 20 orders for the new machine, and if so, that's an impressive debut. But it's hardly a supercomputer. As one source who declined to be identified asked, "What's the big deal about another Intel microprocessor-based machine?"

Unable to match Cray at the high end, most potential rivals went low. They gradually realized that more microprocessors meant more power to the people. They coupled a few, then tens, then dozens, then hundreds in a single system. They didn't fare any better, however.

No company rose higher quicker and fell lower faster than the classic of parallel processing, Thinking Machines. TMC at the beginning of the decade had sales of $65 million and thoroughly dominated the burgeoning parallel processing market it had created. The future seemed limitless.

Hillis had his own worshippers, including Steven Squires, director of the Computer Systems

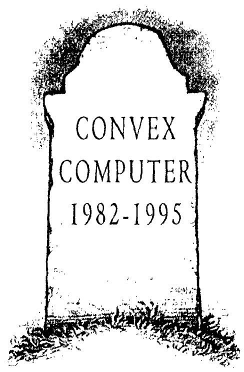

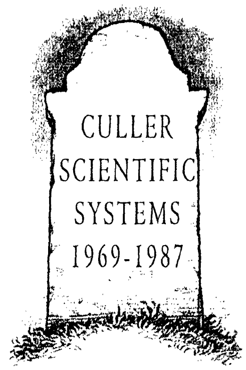

their big brothers. Convex Computer Corp. was started in 1982 as a venture capital-financed mini-supercomputer alternative to the expensive, massive Big Iron boxes that dominated the high-performance computing market. The plan was to build "affordable" supercomputers—in other words, smaller, cheaper and more efficient machines.

Everything proceeded smoothly for almost a decade. Convex went public in 1986 and its stock traded at more than $20 by 1990. It had an impressive string of consecutively profitable quarters. The company was all the rage on Wall Street.

But when the Cold War went down, it also brought Convex with it. The Defense Department stopped writing checks as freely as it did in 1983, when defense official Richard Perle told a Congressional hearing that an Apple II was capable of starting World War III. Convex's profits turned to losses—$140 million since 1993, including $15.3 million in the first half of this year. Money went out faster than it came in; revenue plummeted to $144.2 million last year from $231.8 million in 1992. "I hate to use the term, but this paradigm shift is changing everything that has been the norm for the last 10 years," says Steve Wallach, co-founder and senior vice president of technology at Richardson, Tex.-

WHAT HAPPENED TO THE MARKET? THE REAL PROBLEM WAS THAT THE MARKET NEVER REALLY HAPPENED.

Technology Office at the Advanced Research Projects Agency (ARPA). Through targeted grants and sweetheart deals, ARPA came very close to violating the spirit, if not the letter, of government contracting regulations, enabling Thinking Machines to lead its charmed life.

Since its May 1983 founding on little more than a wing and a prayer, TMC never had to worry about business plans, competitive strategies or the color of its balance sheet. Every time the company leaked, ARPA patched the hole.

But the days of the future passed in a nanosecond. With the Big Red Menace broken and no longer the justification for an unlimited military budget, ARPA pulled the plug.

The company's Connection Machine couldn't cut it in the real world. TMC was also wracked by management turmoil. In August 1994 it sought shelter in the Chapter 11 bankruptcy womb.

Kendall Square Research, another one-time high-flier in the massively parallel world, even drew Bell as investor and consultant. Founder Henry Burkhardt was a terrific technologist and a terrible accountant. After revealing accounting irregularities in 1995, it also filed for Chapter 11.

The fate of the minisupers paralleled that of

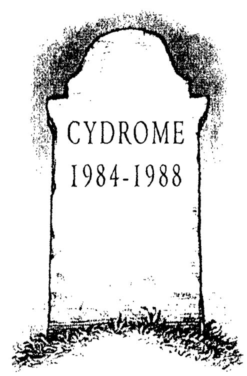

based Convex, now the Convex Technology Center of Hewlett-Packard Co., which bought the company last September for $150 million.

BACK FROM THE DEAD

What happened to the market? The real problem was that the market never really happened. The breathless anticipation of the military turned into a panting enthusiasm from entrepreneurs and venture capitalists, and none of it was really deserved. Thinking Machines and others discovered that corporations are not as free-spending as the military. Says Bell, "TMC got in so much trouble because in the beginning they extrapolated, from a few sales directly related to government placements, that there was a market for huge processing machines. There never was a real market there."

For a while, the military enthusiasm seemed infectious. There was a certain amount of prestige to owning a supercomputer. Apple Computer Inc. even bought a Cray for its research efforts. But as the chill of the Cold War thawed, million-dollar supercomputers became as popular as $500 toilet seats. Business process reengineering, a buzzword that mostly means cutting expenses and people, accelerated the trend. The CIOs and the MISers began buying solution cycles, not just MFLOPS. Users stopped genuflecting at the high-end altar and started asking what the machines could actually do for them.

"The ultimate in hot-vector technology were the machines from [the extinct] Cray Computer Corp.," says Stephen Brobst, managing partner at market researcher Strategic Technologies and Systems, Cambridge, Mass. "They were phenomenally fast but they weren't economically feasible." Adds market researcher Burwen, "There won't be any vector machines sold in another five years. There's no future for those systems."

At first it seemed as though massively parallel systems might in fact win out as much more cost-effective than the vector machines. However, it was much harder to create software for them.

The software steadily improved but not as fast as the microprocessors themselves. Suddenly, real-world problems that people thought only Big Iron could solve—with either vector or massively parallel machines—were being tackled by machines with fewer than 100 processors, and getting the job done. The knotty problem of making software run on a thousand microprocessors simultaneously is often simply not worth tackling. "Microprocessors are getting so powerful that even a small number of them can handle problems that traditional supercomputers used to do," says Burwen.

In a recent study, the Palo Alto Management Group predicted that the market for parallel processing systems would increase at better than a 40 percent annual rate to about $14.3 billion in 1999. Approximately three-quarters of that amount will come from commercial applications such as on-line transaction processing (OLTP), decision support systems (DSS) and multimedia. Science and engineering will still be important markets but will lose ground by century's end.

The market will not be fed by Big Iron machines. "There will be no 1,000-processor machines sold this quarter," predicts UC Berkeley's Patterson. "Even machines costing a few million dollars are very unlikely to have more than 64 processors. There's no demand for the traditional big supercomputers anymore."

Companies are surviving by adopting the new religion that says smaller can be better. Even Seymour Cray's original company, Cray Research, once synonymous with Big Iron, has gotten the drift. Its share of the traditional supercomputer market is actually rising as competitors die off, but the market is shrinking. A few years ago Cray realized that it could no longer live on supercomputers alone. In 1991 it created Cray Research Superservers after purchasing selected assets of Floating Point Systems.

Last February, Superservers merged with two

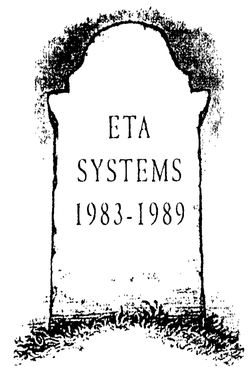

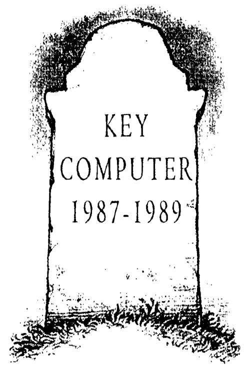

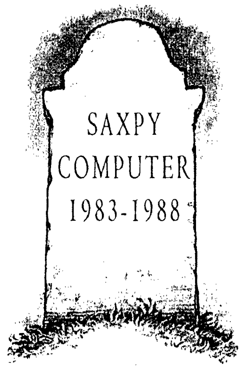

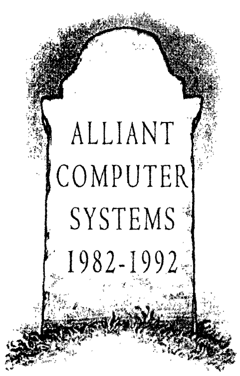

other groups and became the Business Systems Division. That entity's sole offering is the CS6400 enterprise database server, a multiprocessor system based on the 85MHz SuperSPARC II microprocessor from Mountain View, Calif.-based Sun Microsystems Inc. Cray recently demonstrated a 48-processor machine cranking away at a 1.6 terabyte database. Cray claims the CS6400 has penetrated more customer sites in a shorter time than any competing equipment. Cray's commercial customer base has grown so that it currently accounts for more than half of all Cray sales revenue.

"Even major parts of Cray believe [the traditional supercomputer market] is dead," says consultant Brobst. "But they've done a terrific job engineering the 6400 and moving into the commercial market. Only Cray could have done that."

Thinking Machines has dropped out of the Iron business altogether. In October 1994 it reorganized as a software-only company and promptly went on to four consecutive profitable quarters. It apparently has learned its lesson well.

"The market spoke loud and clear, and told this company that building some of the world's fastest computers was not, by itself, enough to sustain growth," TMC president and CEO Robert Doretti admitted recently. "We took a hard look at our core competencies and quickly realized that we had substantial expertise in the software that harnesses the power of multiprocessor computers. We believe there's a huge untapped market for this capability, and this is our strategic focus going forward."

STAYING ALIVE

There are, however, naysayers to all this negativity. "It's a big misperception that the high-end supercomputer market is dying," contends Cray Research President and COO Bob Ewald. Just last November, Cray sold a top-of-the-line 32-processor T90 supercomputer to Nippon Telegraph and Telephone Corp. (NTT). According to Cray Chairman and CEO J. Phillip Samper, demand at the end of the third quarter for the T90 line represented 45 percent of the company's backlog. Cray also claims to have $90 million in orders for its T3E supercomputer, which isn't scheduled for delivery until the first quarter of 1996.

And what does it say about the condition of a market when Tera Computer, which since its founding eight years ago has done absolutely nothing but lose money—a mere $9.5 million—goes public and oversubscribes the offering? "Big Iron is not stone-cold dead," contends Patrick Grady, senior vice president, corporate finance, at H.J. Meyers & Co., the Rochester, N.Y.-based underwriter of Tera's IPO. "It's contracting for the companies still able to play in it, but there's still plenty of business there. People are waiting for a new approach. Vector processing machines are long in the tooth. Parallel processing machines are good but not exceptional. The high-performance world needs a major breakthrough in programming." He believes that Tera is the company to do exactly that.

In fact, the technical supercomputing market could continue to stumble along for several more years. Bell paints a picture with several types of surviving supercomputers, including Cray-style evolutionary supercomputers with multiple vector processors; multicomputers formed from microprocessor-based workstations connected via high-bandwidth, low-latency switches; and "multis," or multiple microprocessors connected to large caches that access a common memory via a common bus.

Bell also believes that a few trends could keep the momentum for such machines going for a long time. For one, the government's "buy U.S." policy is still alive, although not as visible as it was a few years ago. For another, Cray, IBM and Silicon Graphics Inc. have large installed, loyal customer bases for their supercomputers and super servers.

In fact, Intel Corp. announced the sale of a massively parallel machine just a few months ago. The Department of Energy's ASCI (Advanced Scientific Computing Initiative) program is paying

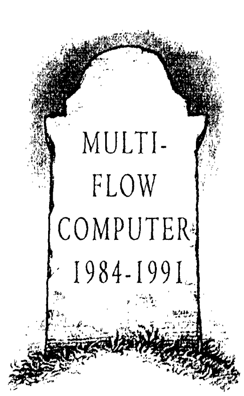

market dominated by nuclear weapons, weather forecasting and a few industries, such as petroleum and aerospace, that needed major horsepower.

"Now we've come completely around," the adviser says. "Big Iron supports a few niche markets, and the rest of the world—because of the amazing increase in computer power and distributed computing, now has on its desk machines more powerful than Big Iron was 20 years ago. You use software tools and couple these together over corporate LANs and public data networks and it beats the shit out of Big Iron in cost effectiveness."

But those niches may no longer be the breeding ground for entrepreneurs. "All the upstarts that entered the supercomputer business are either dead or part of another company," market researcher George Smaby says. And those that aren't—MasPar, NCube (which is nominally independent but would disappear if Oracle Corp. President Larry Ellison decided he had better things to do with his fortune), Meiko and Tera—seem one phone call away from being eaten.

John Toole, director of the National Coordinating Office for High Performance Computing and Communications, says that the huge amount of capital required to play the high-performance game will prevent any startups from joining the roster.

"NOW WE'VE COME COMPLETELY AROUND," A WHITE HOUSE ADVISER SAYS. "BIG IRON SUPPORTS A FEW NICHE MARKETS."

$46 million for a 9,000-processor machine, which it claims is the first teraflop box. It's scheduled to be installed at Sandia National Laboratory in New Mexico in 1996.

But don't look to it as a model for future supercomputer sales. "That's a one-shot, boutique computer deal," Patterson scoffs. "It's a research vehicle [to help develop a nuclear testing visualization program]. It's not going to affect the high-performance market."

Mostly, the supercomputer business has become a market for niche players. "Forget general purpose players," says John Harte, president of MasPar Computers, Sunnyvale, Calif. "In today's market you'd better have a specialty where you own a significant part of the niche you're playing in. You can't maintain a competitive advantage unless you've got a specialized enough architecture that the large players won't adopt." Some examples of such niches: Data mining (extracting info from huge databases) and gene sequencing.

In other words, supercomputing has returned to where it started: a very specialized niche market. "We're not strictly back to the future, but we're pretty close," says a White House technology policy advisor who asked not to be identified. "Big Parallel Iron, and maybe all Big Iron, is reverting to the niche market it was 20 years ago," when it was a small

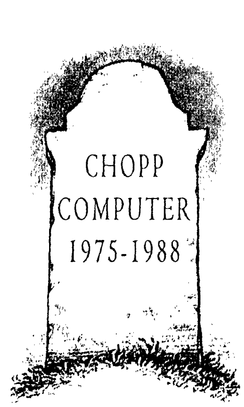

And Hillis adds that technology is changing so rapidly that buyers cut risk wherever they can.

"I'm not suggesting that there won't be startups," IBM's Barnes says. "What people are really saying is that for a hardware vendor to get into the high-end space without fantastic technology is almost impossible. The capital costs are too huge."

For proof, just look at the Cemetery of the Innovators. "Everyone talks about how great the high-performance industry has been for innovation," Barnes says. "But I'm not sure there ever was a time like that. The half-dozen examples that we all talk about—Kendall Square, Thinking Machines, Cray Computer—never really made it. Pieces of their technology may show up in other products, but the companies aren't around."

So if all the innovators and tweakers and tinkerers and experimenters and revolutionaries have turned to dust, who's still standing when the high-performance sun sets?

Does anyone want to bet on the last "Big Iron company?" After all, there are still 30 white rhinos left. ■

Willie Schatz heads the Schatz Group, a Washington, D.C., firm specializing in technology, policy and communications. He is editor of HPCC Week. *His last article for* UPSIDE *was "Cutting the Gordian Knot" in January 1995.*


R.I.P.
THINKING MACHINES
1983-94
R.I.P.
KENDALL SQUARE RESEARCH
1986-95

Thinking Machines Returns

Remember the hardware company Thinking Machines Corp.? The erstwhile maker of Connection Machines has turned its focus to software and introduced GlobalWorks, which allows users to take networks of existing systems–particularly Suns, GlobalWorks runs on top of Solaris–and make them behave as though they were multiprocessor devices.

TMC describes GlobalWorks as a parallel-computing environment that runs under UNIX. What it runs on, meanwhile, is the GlobalWorks server. This is a small high-bandwidth network of Sun UltraSPARC-based systems–which themselves can be multiprocessor devices–housed in cabinets that can hold up to 10 systems.

"The GlobalWorks server is essentially a collection of Sun servers connected by ATM or Fibre Channel," says Jacek Myczkowski, TMC's vice president of development and technology.

At the moment, GlobalWorks requires a GlobalWorks server to run. But, says Myczkowski, there is no reason why you might not someday be able to use the technology to link ATM networks of Suns into similar clusters that don't reside in a TMC chassis. "In the–we hope–not *too* long term, you'll be able to choose what you buy from us," he says. "It could be hardware, or just software."

GlobalWorks marks a significant shift for TMC, not only in terms of its product, but also in the way it thinks about technology. "Parallelism takes on a different meaning," says Myczkowski. "In the past, it meant grand challenges. Today, it means a throughput problem." Where before, it meant supercomputing and complex design or research problems, now it means things like database searches and data mining for commercial users.

So what happened that brought TMC from the grand challenges of the past to the throughput problems of the present? Six years ago, TMC, then of Cambridge, MA, was one of the most glamorous of all the high-tech firms along what was then called "AI Alley," a stretch of office buildings not far from Massachusetts Institute of Technology. TMC had come into being to produce the Connection Machine, a supercomputer-like device based on thousands of individual

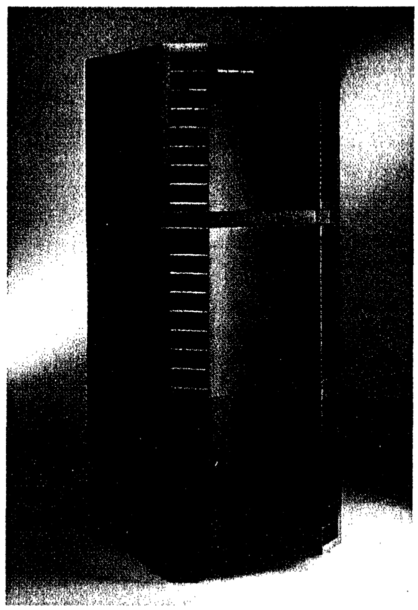

The GlobalWorks server from Thinking Machines houses multiple Sun systems. With the company's GlobalWorks environment, the device can also function as a parallel-processing supercomputer.

processors. Connection Machines were thought to be the devices most likely to give established high-performance computer vendors, such as Cray Research Inc., a run for their money.

Moreover, Connection Machines were supposed to be the hardware that might, someday, achieve the Holy Grail of computer science, the sentient machine. The company's motto was "someday, we shall build a machine which can think...and it will be proud of us."

As to what went wrong the first time, Myczkowski cites various technical and marketing reasons. "I think that technology and economics took a turn that left TMC with products that weren't competitive," he says.

Specifically, the company was busy selling big, expensive systems at a time when the industry was discovering that it could do almost as well with inexpensive little ones.

"And," admits Myczkowski, "there were also internal problems...with management style." In particular, he cites the company's long-standing belief that superior engineering alone was necessary to build a successful company, and that actually selling boxes was unnecessary and even vulgar. "We had espoused the philosophy that if we build it, they will come," he says. But, they didn't come. "We should have become a marketing driven company."

But management style problems went even deeper for TMC. The company not only failed to market its product, it was very successful at alienating potential customers. It gained the reputation for being among the most difficult companies in the industry to work with.

But, if pride goes before a fall, then sometimes falls are instructive. Myczkowski says that the TMC of 1996 is a new company with a new handle on things. The company, he says, has learned "that the marketplace is governed by laws that are not the same as those of science."

Thus, TMC's new direction. "We decided that our core competency was to leverage our software expertise and make it available on a wider set of platforms," says Myczkowski. "We decided to move it from the supercomputer domain to a broader marketplace."

In other words, TMC's hardware was never exotic. Connection Machines were remarkable in appearance, but underneath they were primarily a collection of reasonably standard components. What was unique about them was TMC's methods of making the components cooperate. The company now intends to take those methods and sell them for other devices, like workstations and servers. "The network," says Myczkowski, "is a perfect environment for this sort of technology."

But can TMC come back with a software-only solution? It's not, after all, alone in this. Products that allow networked devices to cooperate have been around for quite a while. Over a decade ago, Apollo Corp. (now a division of Hewlett-Packard Co.), for example, had an operating system called Domain that allowed networked workstations to perform as a single device.

During the days of lavish government funding on the Strategic Defense Initiative, there were several federally sponsored efforts in the same direction. And there were several Transputer-based approaches that likewise sought a market in the 1980s.

None of these came to much. Why, then, does TMC think it can do what others haven't? Myczkowski thinks that the situation has changed. "In those days, there was a big difference between the MIPS you had and the bandwidth of your network," he says. Now, high-performance communications technologies, like Fibre Channel, make it possible for networks to behave as if they were parallel-processing systems, says Myczkowski.

But if that's true, doesn't TMC run the risk of once again being outflanked by a commodity product? Won't every vendor have a similar capability in its own product line?

Myczkowski has an answer for that too. "I think we know how to do this. We have had the experience. We have had 12 years of dealing with parallel computing," he says.

In short, he says, "We've already made all the mistakes." And that alone, he believes, is a significant advantage.