



THE ECONOMY

COMPUTERS SIMULATE PHYSICAL SYSTEMS

BOSTON

■ One of the most popular applications of supercomputers today involves using the machines to simulate physical systems. Often such simulations involve many millions of repeated calculations.

A program running on a conventional computer that calculates the motion of a protein molecule, for example, would have to consider each atom in the molecule one-by-one, calculate the force on each atom exerted by the other atoms, and then compute the new position of each one.

Solved by a Connection Machine, each of the molecule's atoms might be assigned to a different processor. Working at the same time, they would all calculate forces and determine the new position of the entire molecule. Each of the CM's physical processors can be programmed to simulate any number of "virtual processors," so even the company's 4,096-processor CM-2a can handle large problems.

Many scientists feel that this "data parallelism" approach is "a natural way to think about their problems," says Dr. Jill Mesirov, Thinking Machine's director of mathematical sciences research. Part of her job is to help show scientists how the machine can be used to solve their problems.

At the MIT Media Laboratory, graduate student Henry Holtzman is using a CM to increase the detail in television images: Each of the CM's processors calculates the color of a single point of the image.

"It's great to program. Once you get used to thinking the way [the machine] operates, it makes working with parallel data very simple," says Mr. Holtzman. "The real problem with programming it is trying to debug your program. It's kind of hard to watch what your program does" when it is operating on millions of numbers at a time.

- S. L. G.

SUPERCOMPUTERS

'Electronic Brain' Worthy of Name

By **Simson L. Garfinkel**
Staff writer of The Christian Science Monitor

BOSTON

COMPUTERS have long been thought of as "electronic brains." But the Connection Machine (CM) comes the closest yet to working like one.

Most computers have a central processing unit that executes one instruction at a time. The CM has thousands of simple processors, all computing in parallel like neurons in the human brain. In this way the CM breaks super-sized problems down into millions of pieces and solves them simultaneously.

The CM was invented by W. Daniel Hillis. Working at Massachusetts Institute of Technology in the early 1980s, Dr. Hillis realized that conventional computers were organized the wrong way for solving problems that required large amounts of information—the sort of problems he was encountering in his work on artificial intelligence.

"If you look at a very simple task that we can get a computer to do—recognizing an object—it might take a computer 15 minutes to do what an infant can do in a fraction of a second," Hillis says.

Yet neurons compute a thousand times slower than most computers. The problem, Hillis wrote in his doctoral thesis, was one of utilization: most integrated circuits in a typical computer simply store numbers and play little role in performing calculations. A much better approach, he reasoned, would be for the computer to use the majority of its silicon for "thinking."

Hillis founded Thinking Machines Corp. seven years ago to put his ideas to the test. Together with Sheryl Handler, a management and financing wizard who had recently put together a biotechnology venture in Boston, Hillis raised \$65 million in funding. The money came not from venture capitalists, who would be interested in short-term profits and the bottom line, but from investors who could afford to take a long-term view, says John Mucci, Thinking Machine's vice president for research and marketing.

In 1983 the company started building the machine that Hillis had described in his thesis. Three years later it was ready. By many accounts, Dr. Hillis had built the fastest computer in the world.

Called the CM-1, the computer had 65,536 microprocessors; Hillis hoped to eventually build a machine with over a million.

But it turned out that even with the initial number of microprocessors, the CM-1 could do a lot of things that no conventional computer could, Hillis recalls triumphantly. The machine excelled in solving scientific and engineering problems that involve vast amounts of information. Although not every problem could be phrased in such a way to take advantage of the CM's power, for those that could, the payoff was dramatic: the CM can run circles around a Cray, a conventional supercomputer costing four times as much.

In April 1987 the company announced an improved machine, the CM-2, that ran the same software and overcame many of the

"They've finally moved off the perimeter and are beginning to be viewed as mainstream."

By the end of 1988, the dozen CM-1s had been upgraded and eight more CM-2s had been sold. Since then, the company's sales have grown by 50 percent each year. Today there are more than 50 CM-2s in customer's hands—including one recently sold to a Japanese laboratory.


"When new machines are introduced, they are introduced in the research environment. They rely on the feedback loop of universities and the research community to assist in commercializing," says Mr. Smaby.

The real payoff for Thinking Machines, says Dick Shaefer,

inside of most desktop computers. Most of Intel's supercomputers average only 32 processors, but each one is many times faster than the processors inside the Connection Machine. "Because of the use of stock microprocessor technology, which will become increasingly more powerful, you will see the cost go down as the performance goes up," says Ken Harper, a spokesperson for the company.

Hillis doubts that Intel's approach is workable in the long-run. "The interesting thing about starting with a very large number of processors and improving the speed of the processor is that the programs stay the same. . . . Writing a program for 100 pro-

CURT CLEMENS



SEISMIC WAVE SIMULATION: Each processor of a CM-2a computes wave motion at a particular point inside the earth.

CM-1's problems. Each CM-2 processor has 16 times more memory, which lets the computer solve even larger problems. The new computer also boasts a color screen that can compute and display television images at full speed, allowing scientists to watch the results of their calculations as they occur.

With the CM-2 came the Data Vault, a mass-storage system which uses an array of 39 disk drives to hold 10,000 megabytes of information (compared with 20 megabytes of a typical desktop computer). A Data Vault can transfer data into the CM-2 at the rate of 25 megabytes per second—equivalent to 35 average-sized novels.

Even priced from \$2 million to \$10 million, the machines were in demand. "Parallel [processing] machines have always been viewed as being at the radical edge" of the supercomputer world, says Gary Smaby, a supercomputer analyst in Minneapolis.

founder of Technologic Partners, a Manhattan research firm, will come as applications outside of sciences and engineering are discovered for the machine. One such application is DowQuest, a service offered by Dow Jones Information Services. It uses the power of a Connection Machine to take questions in English and find articles that answer them, from a database of over 175 publications.

"That's a daring experiment in many ways," says Mr. Shaefer. "Thinking Machines needs more such sales."

DowQuest, says Mucci, shows the Connection Machine's ability to handle problems that are "simply beyond any computers capabilities because they require computing on so incredibly much data."

But other companies are also exploring parallel approaches. One is Intel Scientific Computers, a subsidiary of the company that makes the microprocessor that is

processors is not the same as writing a program for 10,000 processors. People invest a lot in learning how to program a certain kind of machine; if next year's machine is 10 times larger, they have wasted that investment."

Last November, the United States Defense Advanced Research Projects Agency awarded Thinking Machines a \$12 million contract to develop a computer capable of 1 trillion operations per second—a machine that is a hundred to a thousand times faster than the CM-2. (Intel was awarded a similar contract, Harper notes.)

A computer that fast, says Hillis, should be able to solve some of the "grand challenges" of supercomputing: "Things like actually predicting weather in real time, global climate modeling, and figuring out how proteins will fold. . . . Things like that require bigger computers than what we would call a supercomputer today."