Research News

The Connection Machine Goes Commercial

The computer's massive parallelism allows programmers to tackle problems in ways they never would have contemplated before

THE Connection Machine, an ambitious experiment in achieving very high computing speeds by putting thousands of processors to work in parallel, is moving out of the laboratory. On 30 April a commercial version of the device was announced by Thinking Machines Corporation of Cambridge, Massachusetts, a spin-off of the artificial intelligence laboratory at the Massachusetts Institute of Technology (MIT). A model using 16,384 processors will cost roughly $1 million, and a model using 65,536 processors will cost roughly $3 million. Eventually the company hopes to have a version that uses as many as one million processors.

Originally conceived in 1981 by W. Daniel Hillis, then a graduate student at MIT and now chief designer at Thinking Machines, the commercial version of the Connection Machine can execute between one billion and seven billion instructions per second, depending on the kind of application being run. Such performance puts the Connection Machine in the same league as scientific supercomputers, says Hillis. However, its speed per se is not the point. "It's not just a faster computer," he says. "It's a different *kind* of computer."

Conventional, one-step-at-a-time computers are already beginning to approach some fundamental speed limits, he explains, largely because their microcircuitry can only be miniaturized so far. The best way to get around those limits is thus to develop parallel machines that perform many steps at a time. The Connection Machine is hardly the first such computer; in fact, dozens of projects are under way worldwide. The Connection Machine is not even the first commercial parallel processor. Last year, for example, the Intel Corporation introduced a line of machines that can incorporate up to 128 processors, each roughly equivalent to an IBM AT personal computer.

However, the Connection Machine *is* the first commercial machine to implement an approach known as data-level parallelism, or massive parallelism, says Hillis. Unlike the Intel machines and most other parallel designs, which work by linking together a relatively small number of powerful computers and then assigning a piece of the program to each, the Connection Machine works by linking lots of very small processors together and running the same program on all of them simultaneously. One can think of the Connection Machine as a single giant memory bank, in which each tiny chunk of memory has been given a little processing power of its own.

The crucial element in this approach is communications, says Hillis. For data-level parallelism to work, any given processor has to be free to pass data back and forth to any other processor on demand. Thus, the innards of the Connection Machine constitute a kind of electronic telephone network. Indeed, devising a means to handle the internal information flow was the biggest single challenge in developing the machine. But the payoff is that the processors of the Connection Machine can adapt and readapt their communication network into any configuration necessary, even while a program is running. And this means in turn that the machine can automatically concentrate its computational resources on the parts of the problem that need them the most. In general, says Hillis, the Connection Machine's data-level parallelism works best when relatively simple operations need to be performed on huge amounts of data. Some examples:

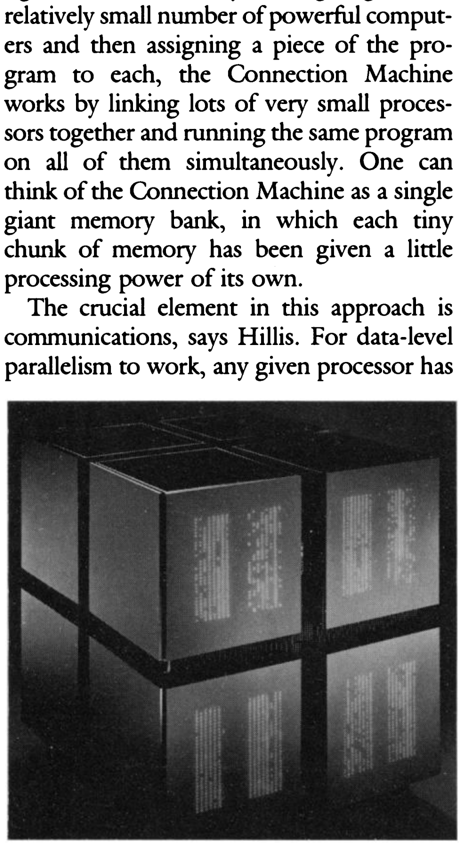

The Connection Machine

A translucent black cube roughly 1.5 meters on a side, the computer holds 65,536 individual processors.

- Image processing. A high resolution, digitized image can easily contain more than a million pixels, which means that image processing on conventional computers tends to be very slow. On the Connection Machine, however, each pixel, or each small group of pixels, can be assigned its own processor. Starting from stereo images looking down on mountainous terrain, the computer can produce a detailed contour map of that terrain in somewhat under 10 seconds.
- Document retrieval. Ordinarily, one compiles a computerized bibliography by typing in a few keywords such as "Halley's Comet"; the computer then searches for matching keywords by looking through all the documents in its database one by one. On the Connection Machine, however, each document can be stored in its own processor. A search of 3 months of stories from the Reuter's news service—15,000 documents—takes less than a second.
- Fluid dynamics. The usual method of calculating how air flows around, say, an airplane wing, is by numeric integration of a partial differential equation such as the Navier-Stokes equation. Using massive parallelism, however, programmers can model the fluid as a swarm of individual particles, and thus dispense with differential equations entirely. The space through which the fluid moves can be modeled as a regular grid of cells, each containing a handful of particles; the dynamics of the fluid can then be modeled by a simple rule that tells how the particles move from cell to cell. (In general such systems are known as *cellular automata*. One of the most active researchers in the theory of cellular automata, Stephen Wolfram of the Institute for Advanced Studies in Princeton, New Jersey, is a corporate associate of Thinking Machines.) Researchers at Thinking Machines have implemented such a fluid dynamics program that operates a two-dimensional grid of 4000 by 4000 cells,

typically containing some 32 million particles. With each processor in the Connection Machine taking care of several cells, the system can perform roughly 1 billion cell updates per second. A simulation of 100,000 time steps reportedly takes less than 30 minutes and produces results that compare favorably with those produced by traditional numerical integration.

It must be said that there is still considerable room for skepticism about all this. Massive parallelism is a new and relatively untried concept, and it is not at all clear that the Connection Machine will ever occupy more than a specialized niche. Indeed, it was originally conceived as a device for doing artificial intelligence research; only recently have Hillis and his colleagues begun to explore its potential for more general-purpose scientific computing. Furthermore, as a purely practical matter, anyone who wants to use the Connection Machine for general scientific problems will have to do so without the help of the millions of lines of computer code that have already been written in FORTRAN: Thinking Machines currently offers only LISP and C as programming languages.

On the other hand, it is already clear that massive parallelism allows programmers to tackle problems in ways they never would have contemplated before. And as a purely practical matter, programming code developed for the Connection Machine tends to be short, straightforward, and conceptually transparent. "Not having FORTRAN is a short-term disadvantage," says Hillis, "We debated a lot about whether to offer it. But we think these new languages [LISP and C] offer so much more productivity that people will be won over."

Thinking Machines has accordingly targeted this first version of the Connection Machine at sophisticated users who can explore what its possibilities really are. The first machine has already been delivered to the Defense Advanced Research Projects Agency (DARPA) in return for the agency's $4.7 million of support during the development phase. DARPA has also ordered a second machine, and will be using both in its Strategic Computing project. Two other machines are going to MIT, and one apiece to the Perkin-Elmer Corporation and to Yale University.

"This is among the very best of our computer architecture projects," says Steven Squires, deputy director of the information technology office at DARPA. "It's going to push back the frontiers for years to come. People have had many very good ideas that they couldn't pursue just because the computers have been too slow. Now they can." ■ M. MITCHELL WALDROP

AIDS-Related Brain Damage Unexplained

Most AIDS patients develop a variety of neurological problems. The most common is a dementia that seems to be caused by the AIDS virus itself

NEW results reported at two recent scientific meetings indicate that the retrovirus that causes acquired immune deficiency syndrome (AIDS) not only gets into the brain, but also attacks specific brain regions and cell types. Richard Price and his colleagues have evidence that the AIDS virus is most often in white matter and in gray matter deep within the brain. Anthony Fauci, Scott Koenig, and their coworkers find that most of the virus in brain is in multinucleated giant cells that are derived from macrophages.

Price and Tomas Pumarola, of the Memorial Sloan-Kettering Cancer Center in New York, identify some "process-bearing cells," which may include astrocytes and neurons, as containing a protein made by the AIDS virus. But, like Fauci, they see most of the staining in multinucleated giant cells. Price described a new study of 70 AIDS patients at the recent American Academy of Neurology meeting in New Orleans.*

Fauci, Koenig, and Howard Gendelman, at the National Institute of Allergy and Infectious Diseases, identified cell types infected by the AIDS virus in the brains of three AIDS patients. "By in situ hybridization, the multinucleated giant cell was identified as the predominant cell type containing 95% of the viral RNA," according to Koenig. He described one of these patients at the recent meeting of the Association of American Physicians in Washington, D.C.†

"A very large percentage of AIDS patients have neurological problems," says Richard Johnson, of Johns Hopkins University School of Medicine. "The exact incidence is not known, but as many as 60% will eventually develop dementia. About 10% of AIDS patients present neurological symptoms first, including dementia, neuropathy, or opportunistic infections of the central nervous system. These facts are very important when you are talking about therapies. You could suppress viral replication extraneurally and the patient would continue to dement."

The consequences of AIDS in the brain are devastating. In early stages of the disease, many AIDS patients complain of "forgetfulness, a loss of ability to concentrate, mild confusion, and being mentally slow," says Price. In as little time as a few months later, they may be very confused, unable to speak or function independently, and seemingly unaware of how sick they are.

Accompanying these cognitive changes are motor problems such as leg weakness, an unsteady gait, poor coordination, and trouble with handwriting. Many AIDS patients also become apathetic, withdrawn, agitated, or depressed.

The time course over which AIDS patients manifest different aspects of their disease varies greatly among individuals. The usual sequence of events is as follows: infection by the AIDS virus, seropositivity (having circulating antibodies to the virus), AIDS-related complex (ARC), and full-blown AIDS.

Price estimates that out of all the AIDS patients who develop neurological problems, about 10% have neurological symptoms first, before any signs of ARC. Approximately the next 40% show their neurological symptoms after signs of ARC have appeared, but before they have full-blown AIDS. The remaining 50% develop neurological symptoms after they are diagnosed as having AIDS.

Price, Bradford Navia, also of Sloan-Kettering, Eun-Sook Cho of the University of Medicine and Dentistry of New Jersey, and Carol Petito of Cornell University Medical College, studied 70 AIDS patients, 46 of whom were demented. When Price's group examined the brains of these patients, they found only 5 of the 70 that were histologically normal, and those were from nondemented patients.

Typically, the brain of an AIDS patient is shrunken, a change that can be detected in a living patient by computerized tomography (CT) scans. The ventricles, spaces within the brain that contain cerebrospinal fluid (CSF),

*Annual meeting of the American Academy of Neurology, 27 April–3 May 1986, New Orleans, Louisiana.
†Annual meeting of the Association of American Physicians, 2–5 May 1986, Washington, D.C.