A CHORUS OF CPUs

Users of single and multiple instruction machines sing different tunes

The founder of Thinking Machines Corp., W. Daniel Hillis, irked a lot of people at last November's Supercomputing '90 conference in New York when he used the keynote address to belittle his competitors' machines.

In his address, Hillis knocked massively parallel architectures from rivals NCube Corp. and Intel Corp.'s Supercomputers Systems Group, saying they compared poorly with the architecture of Cambridge, Mass.-based Thinking Machines' own Connection Machine, which Hillis himself designed.

All supercomputers using massively parallel processing have certain things in common. Whereas an ordinary serial computer has only one CPU (central processing unit) to execute the instructions in a computation, a massively parallel machine has hundreds or even thousands of CPUs that can execute many instructions simultaneously. Each CPU has its own data memory, and the processors can share data when necessary. Operations performed by different processors are coordinated by a serial computer acting as a front end.

But the machines have differences as well. Hillis' Connection Machine is called an SIMD (single instruction, multiple data) machine, which means that each CPU has its own data memory but shares a common program memory. Like a choir singing in harmony, all CPUs simultaneously process different data according to the same instruction, which is broadcast to them by a special sequencer chip. Individual CPUs have the option of executing or not executing a given instruction, but they can't execute different instructions at the same time.

By contrast, in MIMD (multiple instruction, multiple data) machines such as those from Intel in Santa Clara, Calif., and NCube in Beaverton, Ore., each CPU has its own program memory and can execute completely independent programs. An MIMD application acts less like a choir and more like an orchestra: Each instrument plays something different but remains in concert with the others. In his keynote address, Hillis claimed that MIMD machines are inefficient because of the difficulty in synchronizing all the CPU operations.

Users of MIMD machines have found otherwise. David Audley, director of Prudential-Bache Securities Inc.'s financial strategies group in New York, has programmed Intel's iPSC/860, an MIMD machine, with 32 CPUs to run so-called "Monte Carlo" simulations 24 hours a day in support of the firm's securities analysis activities.

"We create a number of different econometric models and run each model on a different CPU," says Audley, who notes that the brokerage needed an MIMD architecture to run the application. "Each model assumes different economic scenarios, including different interest rates and other factors, and computes how a given investment would fare in that particular scenario. Since each of the computations runs separately, there's no problem synchronizing the CPUs."

SIMD machines, in comparison, are widely recognized as better suited for massive database searches. For instance, Dow Jones and Co. purchased a Connection Machine with 32,000 CPUs to support the Dowquest feature of its News/Retrieval Service. With Dowquest, an on-line user can type any phrase, such as "Persian Gulf," and the service uses its Connection Machine's CPUs to perform simultaneous text searches of the thousands of articles in its database to locate the ones containing the specified phrase.

Image processing is another common SIMD application. At MasPar Computer Corp. of Sunnyvale, Calif., special software has been developed for its SIMD machine to process video-screen images. Since the software assigns each pixel on the screen to a separate CPU, complex image changes can be performed very quickly in real time.

Most users, however, don't care whether they use a vocal chorus or a full orchestra, as long as they get the sound they want. "Right now, all massively parallel machines are difficult to program, although SIMD machines are a little easier," observes Prof. Kenneth Kennedy of Rice University in Houston and director of the National Science Foundation's Center for Research on Parallel Computation. But in the long run, he maintains, it is MIMD that will win out "because it's more flexible, and we'll figure out how to create programming languages that make synchronization of multiple processes easier."

—*John J. Xenakis*

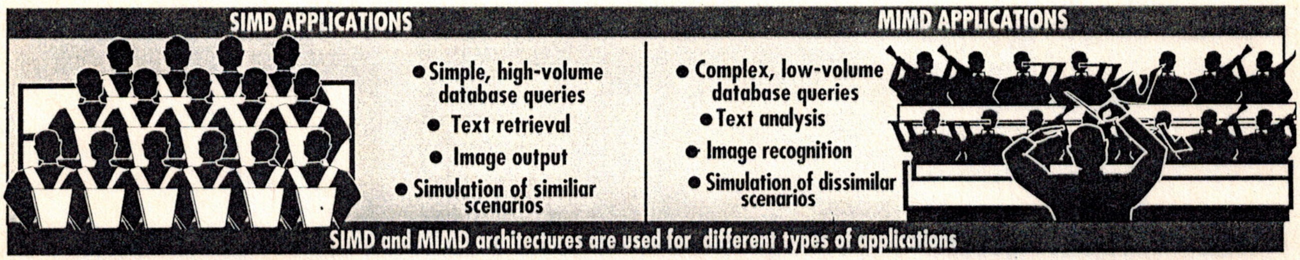

Thinking of Machines

Hillis & Company race toward a teraflops

In 1981 a graduate student in Marvin L. Minsky's Artificial Intelligence Laboratory at the Massachusetts Institute of Technology wrote a curious memo. First he complained that computers were too slow: try to instill even a twinkle of intelligence in them, and they bog down even more, he said. Then he proposed a solution: a novel computer built by connecting thousands of weak processors.

Ten years and 450 employees later W. Daniel Hillis and his Cambridge, Mass., company, called Thinking Machines, plan to unveil his third computer based on that concept. This one, Hillis believes, may finally muster enough speed and grace to let him tackle artificial intelligence. According to Hillis, the raw processing speed of the largest version of this series of Connection Machines, dubbed CM-5s, could be as much as a trillion floating-point operations per second (teraflops)—a trophy eagerly sought by more than half a dozen of the world's leading supercomputer designers.

The models of the machine that Hillis was slated to unveil in late October will not reach a teraflops. But based on that design, Hillis says he can plug together about 16,000 processor nodes and "comfortably" calculate at the teraflops level. Thinking Machines has not yet priced such a machine, which would likely occupy an area as large as a tennis court. Smaller CM-5s will range from $1.5 million for a 32-processor-node machine to $20 million for an 8,000-node machine.

Like the first computer introduced by Thinking Machines in 1987, the CM-5 is a parallel processor—one that sets many computing elements to work on a problem simultaneously. The new design should further push so-called parallel computing into the mainstream, partly with the assistance of IBM, which recently gave a nod of approval to Thinking Machines in the form of a cooperative research agreement.

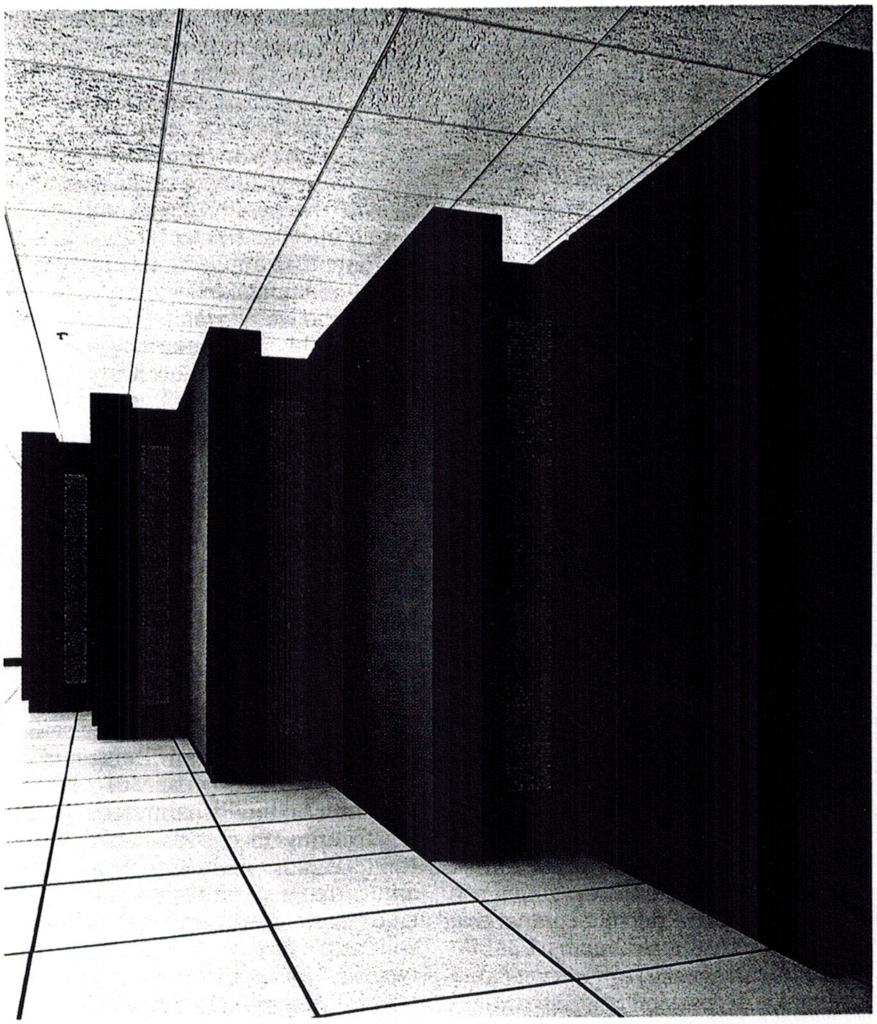

COMPUTERS should not look like refrigerators, Danny Hillis declares. This version of his new supercomputer uses 1,024 processor nodes. Photo: Steve Dunwell.

Yet whereas the early Connection Machines relied on thousands of relatively simple one-bit processing units, the CM-5 could use as few as 32 processing nodes. These nodes rely on a package of four accelerator chips, along with a sophisticated microprocessor, called a RISC chip, which runs fast by carrying out simplified instructions.

From a user's point of view, the CM-5 may also help bridge the two divergent schools of parallel programming that have developed over the past few years. The new architecture has been designed to let users take advantage of both SIMD (single instruction, multiple data) and MIMD (multiple instruction, multiple data) programming. Hillis had previously emphasized SIMD programming, in which one instruction is broadcast to all processors, which then carry out the operation on their own datum. Other designers have pursued MIMD, in which more complex processors use different instructions to manipulate different data sets. Although MIMD architectures can adapt more flexibly to a wider range of problems than can SIMD designs, they can be difficult to program.

During the past year or two, SIMD machines, such as those built by Thinking Machines and MasPar Computer in Santa Clara, Calif., have been gradually edging closer to MIMD designs, points out Jeffrey C. Kalb, president of MasPar. In the CM-5, every node can operate either on its own independent set of instructions (in an MIMD fashion) or on instructions broadcast to all processors (SIMD). Independently operating nodes are synchronized when necessary. Depending on the work at hand, processing nodes might undertake 10,000 operations without being synchronized—or communicate with one another at every step.

Physically connecting such processing nodes, however, is no trivial task. According to Hillis, researchers from Thinking Machines and from M.I.T. have created a network in which the amount of communication possible between processors increases at the same rate as that at which more processors are added. Hillis is still reluctant to provide precise details of the network that makes the machine so "scalable." Even though communication between nodes physically adjacent to one another will still be faster than between distant nodes, Hillis says that the precise geometry of the processors is far less important than it was in earlier Connection Machines.

The communications scheme is also

key to two other important features of the machine: its reliability and the speed with which it fetches or sends data to external memory banks (called I/O, for input/output). Almost a third of the circuits in the communications network are devoted to ensuring the machine is functioning properly. Should a processor fail, the machine will automatically reroute tasks and data, without significantly slowing down its speed, Hillis says.

Shuttling data in and out of an external memory is essentially an extension of the communications network used within the machine. As a result, multiple I/O ports, each of which can operate at about 20 megabytes per second, can be teamed to transfer data at hundreds of megabytes or even tens of gigabytes per second.

Hillis is betting that the SIMD/MIMD duality of his machine will open up new classes of problems that earlier models of Connection Machines did not handle well. By late November, when the first four CM-5s were scheduled to be working at customers' sites, users will have had a chance to answer that question for themselves. But first they will have to resolve some software problems. Because the architecture of the Connection Machine has changed, users of earlier models must recompile, or rearrange, their programs before the software can run on the new CM-5. "The single biggest problem the industry faces is developing software," declares Kalb of MasPar.

Thinking Machines' recently announced partnership with IBM might help boost its software reach. IBM has clients who need highly parallel architectures, says Irving Wladawsky-Berger, an assistant general manager for IBM's enterprise systems division. At the same time, "there is a class of customers for whom having it in a blue box with 'IBM' on the side is the solution. They need speed integrated to IBM," Hillis points out. Although the two computer makers do not plan to market their products jointly, they will develop techniques for moving customers' programs from IBM's 3090 series of mainframes to various Connection Machines.

Hillis is not planning to make any major changes in future architectures. "We've pretty much settled into the right general trend for parallel machines for the next decade," he says. Hillis himself plans to spend his time using the machine, particularly concentrating on artificial-intelligence applications. "Now," he adds, "we can't blame the power of the tools for not having a thinking machine."—*Elizabeth Corcoran*