

Making Connections

A cube of eight cubes, matte black, starred with scores of tiny red lights, each blinking on its own but, somehow, in concert.

This is the Connection Machine, the supercomputer at the heart of the Boston University Center for Computational Science, which is being officially inaugurated this term. The center, headed by Professor Claudio Rebbi of the Physics Department, is interdisciplinary; its purpose, already undertaken in its several years of informal operation, is to support advanced scientific computation by researchers throughout the University.

"What gets them in the door is that we have a raw computational power that can allow researchers to do things they never could do before," says Associate Professor of Electrical, Computer, and Systems Engineering Roscoe Giles, a member of the center's council. One of the largest supercomputers in the state, the machine's capacity is "by a factor of ten over what we had. Researchers can get ten times the performance. That's important."

The Connection Machine is a parallel processor, rather than the more traditional vector processor; it is among the first commercial versions of this new breed. "In the classical von Neumann model of computation," explains Giles, "you have a single processor, a memory, and some I/O [input/output] devices. The computation is about reaching into memory, bringing something into the CPU [central processing unit], doing something to it, and sending it back.

The idea here is to achieve high performance in the computer by having thousands of processors, each of which executes the same instruction at the same time. Not just one CPU that does everything, but thousands." Thirty-two thousand, in the case of the center's Connection Machine.

*"We have a raw computational power that
can allow researchers to do things they
never could do before."*

"While it's a very hard thing to do, to take a traditional single-processor CPU and give it ten times higher performance, it's not hard to imagine taking ten times as many of them and putting them together - if you can figure out the software and hardware and reliability issues. That's why it's important to be able to push in the direction of large-scale parallel computers: that's where you might be able to get factors of ten."

"Some problems," observes Giles, "are naturally parallel." His own work with magnetic materials, for instance. In this research, "there is a physical system to simulate, and the computer has an idea of what the state of that system is.
Continued on page 4

Connections

Continued from page 1

The computer's job is to go forward in time and figure out what the next state is. All the different atoms and molecules in the material respond to the focus on them. They don't call somebody up and say, 'Hey, I have no orders—what do I do next?' Each responds, to a degree, independently. At a certain basic level, it's parallel.

"What we've done in the past," he continues, "is, we've told people [with parallel problems] how to go to a serial program, in other words, how to translate what they know [to be parallel] into a series of instructions: take the first molecule, calculate the force on it, do something, and then the second, and then the third..."

"Now," Giles goes on, "we're presenting the same people with the parallel computer. The first reaction is, 'I have a serial program with a list of instructions. How do I make that go in parallel? Oh, my God!' It seems like a big job, and as stated that way, it is a big job. But really, the problem was parallel to begin with. What is required is that they take one step back, think to themselves not 'what program did I used to use to solve the problem?' but 'what problem do I want to solve?' And if they're able to take that step, and are lucky, and the problem really does suit, and lots of other footnotes, they can take a big step forward."


"There are other problems harder to make go in parallel than in serial, or which have some aspects easy to parallelize and some which are hard. From the computer research side, that's

interesting stuff: addressing those applications and building tools for doing things in parallel that hadn't existed before." In that sense, the work of the Center for Computation is extremely important in the development of parallel computing;

accordingly, much work is done cooperatively with the Cambridge company that made the Connection Machine, Thinking Machines, and the center is a test site for Thinking Machines software. According to Robert Putnam, resident applications engineer from Thinking Machines, the Center for Computation is ideal for such work. "What's [uncommon] is the breadth of applications running on the computer," says Putnam. Giles agrees: "The big users are in physics and engineering, but there's use by researchers in the [University's] Center for Adaptive Systems, Computer Science, Information Technology, Mathematics, and Chemistry."

Even more unusual than this variety of users is the contact that the center fosters among them. Usually, "supercomputer centers are focused on people from far away bringing in jobs, running them, and taking them away. This center is rare in that it really does bring people together to work and talk to each other," says Giles.

Since parallel computing techniques tend to be developed independently in different fields, this degree of contact among researchers in different fields is fruitful. "There are things that computer scientists know, that the applied math people know, that physicists know, that engineers know. The same information could plug in from one discipline to another," says Giles. Nevertheless, "during the normal course of business there's no cross fertilization."



With the Connection Machine, Roscoe Giles of the Electrical, Computer, and Systems Engineering Department; Glenn Bresnahan, director of the Computer Graphics Lab; and Claudio Rebbe, professor of physics and director of the Center for Computational Science.