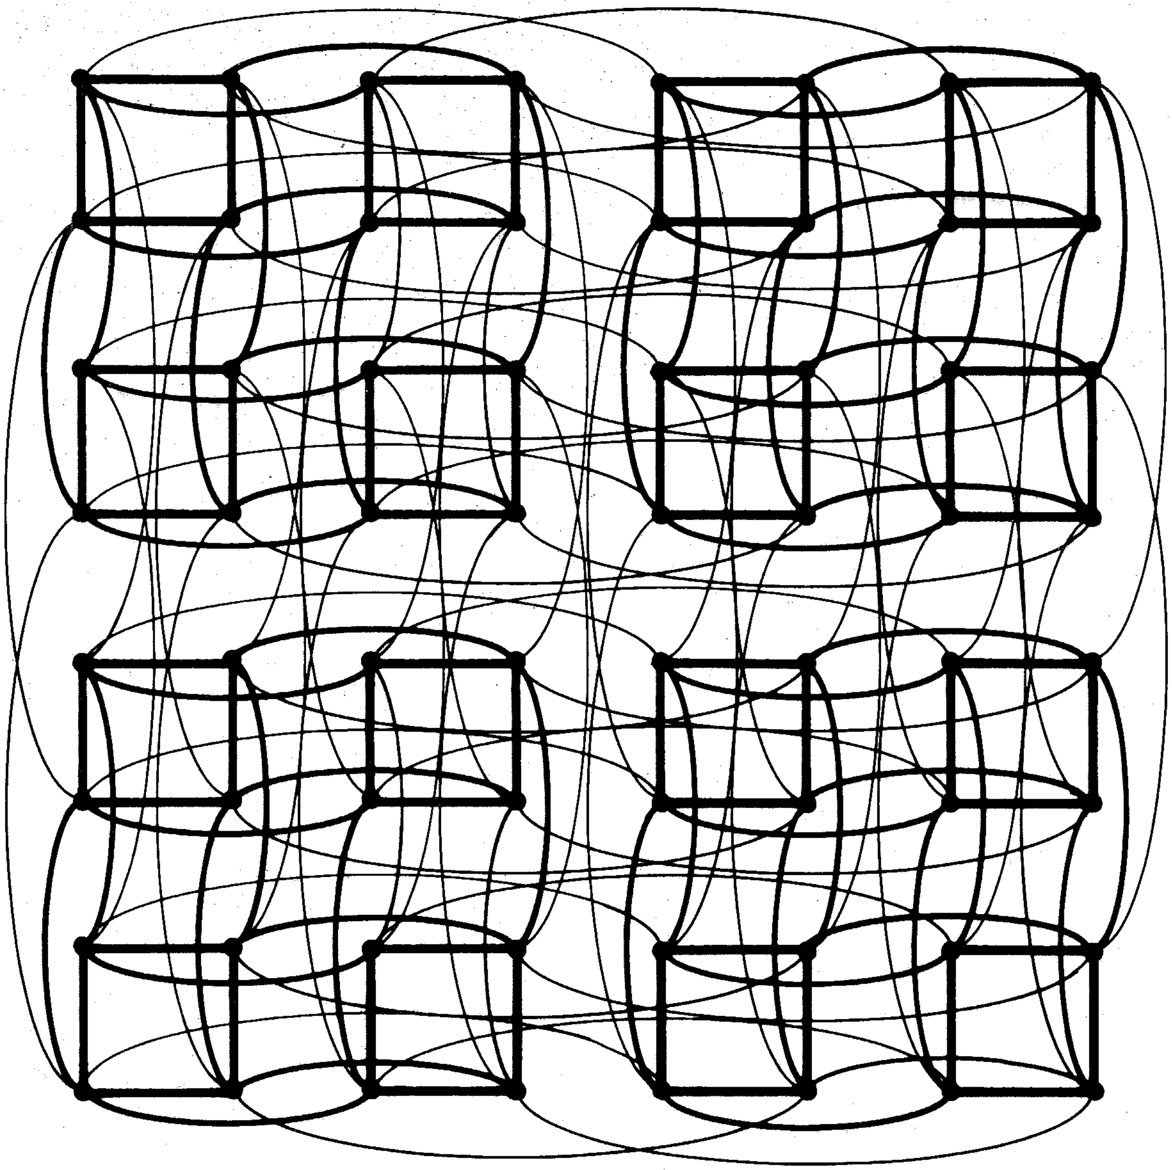

Caltech's concurrent computer is connected as a Boolean 6-cube, drawn schematically here. Each square is a 2-cube; with corresponding corners connected to a neighbor, it's a 3-cube; and with these two repeated and corners connected, you have a 4-cube. Repeating twice again and joining corresponding corners (lightest lines) you end up with a 6-cube.

Cosmic Cubism

Caltech's concurrent computer, resembling a horizontal stack of microprocessor boards, sits unassumingly on a desk top in Jorgensen Laboratory's basement. Aside from its nickname, the "Cosmic Cube," which offers a hint of its vast potential, there's nothing about it to give you a clue that it's the fastest computer on campus. It's hardly a cube; the name refers to its wiring. About 5 feet long, 8 inches high, and 14 inches deep, and consuming only 700 watts, the Cosmic Cube is comparatively tiny for its computing power, which is about a tenth of that of the Cray-1, the largest and most powerful computer in use today.

The Cosmic Cube is just what it looks like — a bunch of microprocessors (64 to be exact) harnessed together. It's called a "homogeneous ensemble machine" because it's a computer built as a group of identical computers working together. Instead of working on a problem sequentially, one calculation after another, as conventional, Von Neumann computers do, each of the concurrent processor's 64 small computers (nodes) works on a piece of the problem at the same time as all the others. Its homogeneous, modular architecture makes possible open-ended expansion (a 128-node machine is under construction and a 1,024-node model is planned), and its creators are talking about building computers 1,000 times the power of the Cray-1 in just a few years — at a fraction of the cost and size.

This revolutionary computer is the result of a collaboration between Charles Seitz, associate professor of computer science, and Geoffrey Fox, professor of theoretical physics and dean for educational computing. Seitz and his students developed the architecture and design of the machine, while Fox has applied the computer to numerically intensive scientific and engineering computation on the campus and at JPL.

Exploiting microelectronic technology and concurrency to multiply the power and speed of computing while lowering the cost is not unique to Caltech. In fact, there is, according to Fox, a "staggering amount" of work being done in the field of concurrent, or parallel, processing around the country, staggering at least in comparison to the amount of progress made so far. (Although the terms "parallel" and "concurrent" are often used interchangeably, the Caltech group prefers "concurrent" as a truer description of how the parts of the Cosmic Cube operate — independently, rather than in lockstep with each other.) But Seitz and Fox's work is far ahead of the rest because "we have real hardware and real software solving a real problem." Caltech's size and interdisciplinary atmosphere encourages the sort of collaboration that can create a prototype like the Cosmic Cube and put it to work immediately on state-of-the-art problems.

Seitz has been working on high-performance architecture suitable for very large scale integration (VLSI) since 1978. It was not until 1982, after he and his graduate students (in particular Dick Lang, Bart Locanthi, Erik DeBenedictis, Chris Lutz, and Bill Athas) had developed their ideas about the programming and technology of this family of machines and had studied them extensively through computer simulations, that they built the 4-node predecessor of the Cosmic Cube. The machine proved so successful in running Fox's programs that Seitz and Fox got together to build the 64-node version.

Both the 4-node and 64-node models use Intel 8086-8087 processors, the same family of chips used in the IBM personal computer. The technology of microcomputers and video games has developed rapidly out of intense competition, and its enormous popularity in the marketplace (home computers are selling in the millions) brings economies of scale. Seitz considers that Pac Man technology, which is in many respects more advanced than that of the Cray-1, has subsidized development of the Cosmic Cube. Even though this amazing technology has gone into some frivolous applications, he says, it's still "an elegant, expressive, and beautiful medium for switching systems of extraordinary complexity." The cost-effective technology is not, however, a good fit for the Cray-1. While the Cosmic Cube offers a tenth of the Cray's power, its cost, at $80,000, is more like a hundredth.

The Cosmic Cube achieves its huge calculating capacity through the teamwork of its 64 computers wired together according to a fairly simple mathematical structure — a Boolean lattice, or six-dimensional hypercube. (Actually, the string of microprocessors could still solve a fair number of difficult problems even when randomly wired together, according to Seitz, and Fox has longed to get his hands on the rows of personal computers that sit alone and unused at night in, say, a bank, and team them up for some giant physics problem.) Each node communicates only with its six nearest-neighbor nodes in the hypercube topology; there is no shared memory with the rest of the computer.

Problems for the Cosmic Cube to solve must be capable of being decomposed into approximately equal segments for each node. The operation of the Cube is particularly simple when the variables at one site are affected only by the variables near it, and you only have to know local conditions to predict the next value. For example, says Fox, to predict the weather in

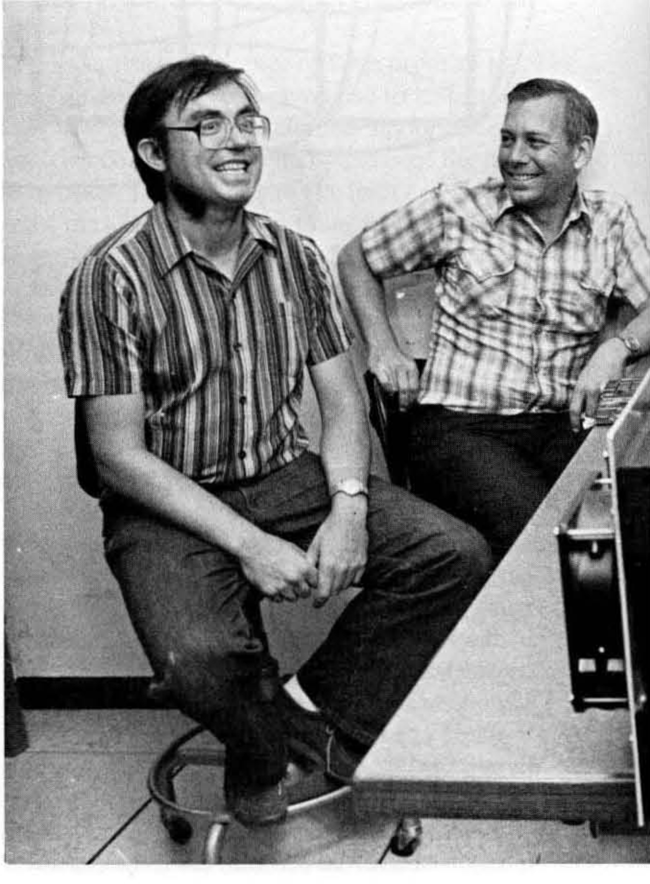

Geoffrey Fox (left) and Charles Seitz display their 64-node Cosmic Cube. One of the nodes, a complete microprocessor in itself, is visible at the front end of the concurrent computer.

Los Angeles in the immediate future, you don't need to know the weather in New York, even though they might be related over a long time.

Meteorology is, indeed, one of the fields with problems well suited to the concurrent processor's capacities — problems that depend on relatively simple algorithms, or equations, but in very large volume. Efficient algorithms for concurrent processing must be such that the time spent communicating is very low compared to the time spent calculating. It turns out that there's a wide range of scientific and engineering problems that fit this approach, as discovered by Fox and the Concurrent Computation Group, which surveyed the heavy computer users on the campus and at JPL for possible efficient application of concurrent processing to their problems.

A number of problems in high energy physics, astrophysics, geophysics, chemistry, and applied mechanics have been seriously hampered by lack of computing power for their vast calculations, since conventional computer technology has leveled off, and current large computers, although they are getting cheaper, are no faster than they were 10 years ago.

But some problems, for their sheer size, require speed; otherwise they would need years or even lifetimes of costly computer time to solve. Anything that has a basic unit (an event, grid point, star, matrix element, pixel) that is manipulated by a basic algorithm and is computationally intensive because the unit is repeated an enormous number of times, is well suited to concurrent processing. This includes all problems involving simulation, modeling, or image processing. Problems depending on partial differential equations, as does Fox's own work in high energy physics, or on matrix inversion, as do some problems in chemical physics, are good candidates for this approach. Applied areas such as oil exploration, aircraft design, and microcircuit design also stand to benefit significantly from the computing power promised by concurrent processing. Seitz, for example, plans to make use of the Cosmic Cube for circuit simulation in designing new generations of chips for its descendants. "Of course, you always wish you had the machine you're building to design its parts," says Seitz.

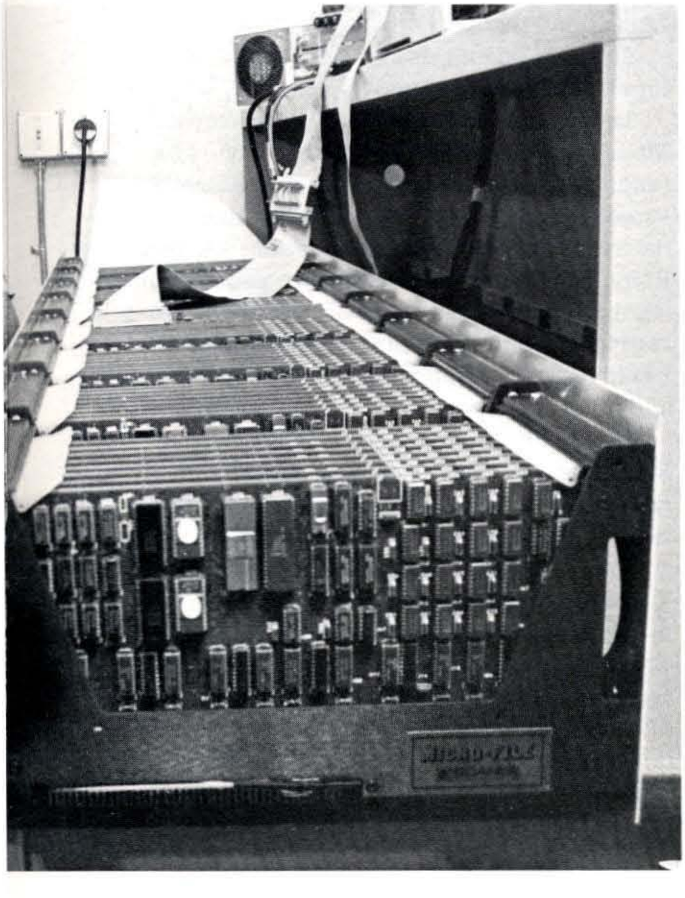

Fox's group in high energy physics (including research fellow Steve Otto and grad students Eugene Brooks and Mark Johnson) is the first to use the Cosmic Cube to work on a previously unsolvable problem — numerical calculation of the quantum field theory called Quantum Chromodynamics. To do a good job of calculating all the interactions in four dimensions, Fox figured he would need something like 100 to 1,000 times the power of a Cray (or 10,000 years on the VAX). The numerical techniques for solving the quantum field theory problem, which are identical to those in statistical mechanics, were pointed out several years ago. But trying to solve the problem in more than two dimensions was very difficult and inconclusive with current computers. "We believe we understand the fundamental physics; we know how to solve the basic equations; but we didn't have the computing power to calculate the consequences," says Fox.

Although the Cosmic Cube still isn't the answer to Fox's prayers, it will typically run about eight times faster than the VAX. It has already put in more than 1,500 hours toward solving Fox's problem as a three-dimensional grid, and with the potential of concurrent processing, he believes that all the fundamental questions describing elementary particles will be solved in the next few years. "It makes research much more meaningful," says Fox. He thinks there is every reason to believe that major scientific breakthroughs in a number of fields will result from the application of concurrent processing.

Several other projects (in astrophysics, chemical physics, geophysics, and, again, high energy physics) are being readied for the 64-node computer to tackle in the next six months, and Fox hopes to have perhaps a dozen working well within a year and a half. One of those coming up soon on the agenda is a study of the formation of galaxies, how they evolve from a system of stars and where the spiral arms come from. It's easy to see how the problem can be broken down into groups of stars, but on first thought, the problem of dealing with long-range gravitational forces might seem to spoil the minimal communication requirement of concurrent processing. But, says Fox, even though you do need to know the value of variables in the whole system rather than just the nearest neighbors, the Cosmic Cube actually should deal with long-range forces quite nicely. The essential criterion is not locality but rather the ratio of communication to calculation time, and although problems involving long-range forces do imply greater communication time, they also imply an even greater calculation time, and the ratio remains small.

Software is being developed at Caltech to implement the concurrent algorithms for numerous problems. Although critics have claimed that programming for concurrent computers will be virtually impossible, Seitz and Fox insist that it's not particularly difficult at all, and point to running application as evidence. Their programming tools make extensive use of software portability, and programs are written in the same languages used by conventional sequential computers, but with a library of procedures for dealing with communication between concurrent processes.

While Fox puts the Cosmic Cube to work, Seitz is designing future generations of machines. A 128-node Cosmic Cube Mark II is being built at JPL to provide additional cycles for scientific applications by next summer. It will have more than twice the performance and four times the storage capacity. Advanced technology nodes, with 10 times higher performance, are being designed using the latest single-chip instruction and floating point processors and will be in use in two to three years. Another machine with much smaller and faster node elements is already well under way, using the "Mosaic" chip Seitz's students have designed. Mosaic, a single, thumbnail-size chip, uses the same architecture, with the storage capacity scaled down, as the 78-chip Cosmic Cube node. It gains in processing what it sacrifices in storage. Many huge computational problems don't really need much storage.

The technology to manufacture Mosaic chips is at the state of the art of custom VLSI and allows a 1,000-node concurrent computer, which would outrun even the Cray-1 for many problems, to be packed into one cubic foot. The essential critical improvement is not in size but in performance for its cost. The Cray-1, which costs roughly $5 million, has a performance on typical problems of 50 million floating point operations per second, a performance/cost ratio of 10. One of the Cosmic Cube's 64 nodes, at a cost of $1,000 and performance of 50 thousand floating point operations per second, has a ratio of 50. And the Mosaic chip, which will cost about $50, will have a performance/cost ratio of more than 2,000.

Continued advances in microelectronics will make it possible within a few years to provide on a single chip the performance of Mosaic elements with increased storage, resulting in a machine similar to the Cosmic Cube with single-chip elements. This was in fact the original idea behind the Cosmic Cube; it's representative of a system whose mode elements could be integrated on a single chip using the technologies that are today in transition between the research laboratories and commercial use. In anticipation of this advanced microelectronic technology, the Cosmic Cube is a "hardware simulation" that allows experiments with the applications, algorithms, and programming of future concurrent supercomputers.

Seitz expects to complete a 1,000-element Mosaic machine by the summer of 1985, with funding from the Defense Advanced Research Projects Agency. The Department of Energy, the System Development Foundation, and the Ralph M. Parsons Foundation are also continuing to support development of concurrent processing, and Intel Corporation and Digital Equipment Corporation have contributed chips and other hardware.

Because the Cosmic Cube is open ended, Seitz believes he can keep using basically the same techniques to expand it until wiring limitations creep in at the 100,000-node range. His goal is to keep making open-ended computers "because these crazy guys have open-ended problems." There will always be problems big enough to be divided into 10^5 or 10^6 parts. Fox agrees that once scientists and engineers know one thing, they'll keep going on to even bigger things. After he solves his favorite problem of quantum field theories in four dimensions with a 1,000- or 10,000-node computer, some even larger problem will certainly turn up.□ — *JD*