

FUTURE OF COMPUTING

Computer Scientists Rethink Their Discipline's Foundations

"Suddenly, I didn't know what a computer was anymore," Richard Lipton of Princeton University recalls thinking recently. He's not the only one feeling that way. As Lipton and others seek the path to computers thousands or millions of times faster than today's—devices that will drive scientific and industrial research in scores of fields—they have strayed far from transistors, resistors, and wires. They are searching for the future of computation in a realm of new media, from optical materials to quantum circuits to DNA, and even new computing principles, far removed from the kind of sequential logic today's computers have inherited from the era of gear-and-lever machines.

Like any venture into alien territory, this effort can be disorienting for computer scientists reared on silicon microcircuits. "It's like an old joke," says Stuart Kurtz of the University of Chicago. "Two weeks ago I couldn't spell engineer and now I am one." That's how I feel about biochemistry," says Kurtz, who is studying DNA computers.

Those who are working in this new realm are determined to persist through their disorientation, because they're convinced that the rewards will be worth the rigor. "Inconceivable" is how Donald Beaver, a computer scientist at Pennsylvania State University, describes the theoretical degree of parallelism—the ability to perform many tasks simultaneously—of a computer consisting of a soup of DNA molecules. "Revolutionary" is the word James Merz of the University of Notre Dame uses for a "quantum dot" scheme that would replace transistors with spots of semiconductor so small that they would accommodate just one electron each. Similar optimism surrounds many other concepts.

Just how much these schemes—ranging from DNA computing to optical circuitry—will improve on today's most powerful hardware is a matter of guesswork, with guesses ranging from a factor of 100 all the way to 10^{12} . The reason for the huge range of uncertainty is that many of these new principles haven't advanced beyond theoretical papers or "toy" demonstrations, with practical applications years or even decades away. There's no doubt, however, about the demand for vastly greater computing power. "In almost any [technical] area," says Rick Stevens of Argonne National Laboratory, who has studied the need for faster computing, "you can identify problems where people are limited by ... computing capability."

Present supercomputers, says Stevens, top


out at tens to hundreds of gigaflops, or billion mathematical operations per second. But there's no shortage of problems that call for petaflops—million billion operations per second. Among them are detailed modeling of global climate, embryonic development, and elementary-particle interactions, and engineering problems such as simultaneously optimizing a jet's structural mechanics, acoustics, manufacturability, and cost.

To speed up conventional computer circuits, engineers make them smaller, shortening the distance the electrons have to travel. But that's an effort that "is eventually going to run out of gas," says Notre Dame's Merz. "Physical limits are being hit," agrees Doug Matzke of Texas Instruments, who has chaired several workshops on physics and computation. As the size of the features etched on silicon chips drops to tenths of a micron, they become harder to interconnect, and the challenge of dissipating the heat generated by electrical resistance grows. For cooling, today's most powerful workstation chip "has fins on it like an old Chevrolet," says Matzke. At some point, microcircuitry will no longer be able to take the heat.

Quantum gains. Even before that point, the effort to shrink microcircuits may run into another barrier, posed by quantum mechanics. At very small scales, electrons be-

have not as point particles but as waves. And that makes them hard to handle when circuit elements themselves dwindle to those scales. As electrons move, their wave functions spread out, making them apt to "tunnel" through the ever-thinner walls between circuit elements and cause the circuits to malfunction.

But by adopting a new computing scheme, Notre Dame theorists Craig Lent and Wolfgang Porod found that they could take advantage of the electron's quantum nature to design switches and wires that are far smaller than present ones and should generate little heat. Their scheme traps individual electrons within quantum dots—blobs of semiconductor so tiny that, like an electron around an atomic nucleus, the electron wave in a dot is forced to occupy a specific energy state, narrowing the wave and holding it tight. Lent and Porod realized, however, that if they trapped a pair of electrons in an array of dots laid out at corners of a square, the electrons' mutual repulsion would force them to burrow into opposite corners of the pattern. The two possible configurations could represent a 0 or a 1, creating a simple data register:



Even better, the researchers noticed, because the mutual repulsion of electrons in adjacent arrays, or cells, is lowest when the pairs have the same orientation, a signal from a control cell would quickly cascade down a long row of cells. In a kind of domino effect, the electron pairs would all align themselves with the first one, transferring data—a 0 or a

COMPUTING CONCEPTS

Quantum Dots

Possible Speed Increase: 100–10,000, depending on the size of the smallest dots that can be produced in regular arrays. **Current Status:** Experimental tests are just beginning. **Negatives:** Early systems will operate at low temperatures.

Quantum Computers

Possible Speed Increase: A trillion or more in the most optimistic scenarios. **Current Status:** Experiments to create single logic gates are under way. **Negatives:** Limited range of problems that such computers can solve; exponential accumulation of errors.

Holographic Association

Possible Speed Increase: As high as 100,000. **Current Status:** Experiments are taking place in a variety of holographic media. **Negatives:** Could be limited to specific applications such as pattern recognition and artificial intelligence.

Optical Computers

Possible Speed Increase: 1000–100,000 without taking into account extra "parallelism" because photons can pass through one another. **Current Status:** Individual components exist. **Negatives:** Better switches need to be developed.

DNA Computers

Possible Speed Increase: A trillion in the most optimistic scenarios. **Current Status:** Proof-of-principle experiment has been carried out, and more sophisticated tests are beginning now. **Negatives:** Limited range of problems, errors, potential practical barriers.

1—without any heat-producing flow of current. Crossing these “wires” in various ways could in turn produce the logic gates of a computer, on a scale orders of magnitude smaller than conventional circuits.

Last year Merz resigned his post as the director of a quantum-electronics center at the University of California, Santa Barbara, and moved to Notre Dame to help Lent and Porod turn their idea into a working system. So far, he says, the results are “still pretty researchy.” When the first quantum-dot systems do appear, they will need to be cooled to a few kelvins to keep the electrons from jumping randomly from dot to dot. But still smaller dots, if Merz can make them, would bind the electrons more tightly and so could be operated at room temperature.

That’s not the only scheme for harnessing quantum mechanics to speed up computation. Some researchers, embracing the strange laws of the quantum world even more closely, are now building components of quantum computers that would take advantage of a single electron wave’s ability to exist in many different states at once. Each state could represent a piece of information, and all could be processed in parallel, speeding up certain computations by many orders of magnitude (*Science*, 7 July, p. 28).

The light fantastic. Instead of descending to the scale of individual electrons in search of new computational power, other researchers are abandoning electrons altogether and adopting a different information carrier: light. Not only are pulses of light the fastest messengers in nature, but they also pass through one another without effect. That should allow any number of activities to take place simultaneously in an optical circuit.

To make the elements of such a circuit, some optics researchers are learning how to make light guides that efficiently carry photons around sharp bends on a chip (*Science*, 26 May, p. 1131). Others, like Seng-Tiong Ho of Northwestern University, are designing the tiny, superefficient lasers that would generate the pulses of light at play in an optical circuit. At the Quantum Electronics and Laser Science meeting in Baltimore last May, for example, Ho and colleagues announced that they had induced lasing in a ring of semiconductor as small as 4.5 microns across. The ring acts as a “resonating cavity,” intensifying the laser light. But because its circumference is only several wavelengths long, other wavelengths can’t resonate within it and siphon off energy.

The design boosts the efficiency of the laser and reduces its output of waste heat—a feature vital for a component that might be used by the thousands in a single optical computer. “It’s a real breakthrough,” says Tony Campillo, head of the optical physics branch of the Naval Research Laboratory. And the photons generated by the laser eas-

ily escape from the ring to a surrounding waveguide, says Ho, where they could drive an optical circuit.

Such a circuit, however, also requires switches that enable one light signal to alter another, as a transistor does for electronic signals. Developing these transistor analogs has proved a major hurdle, because such switches require “nonlinear” materials that respond to light by changing their optical properties—and materials with a nonlinear response strong enough for information processing are scarce. As a result, not everyone favors riding the optical wave all the way to the beach. “I don’t believe in [purely] optical computing at all,” says Yoshihisa Yamamoto of Stanford University. Yamamoto thinks the actual information processing should be left to microelectronics. Light—and his own small, diode-based lasers—is best used only



Ring of light. A ring laser just 9 microns across sends light into the surrounding waveguide.

to speed communications between the electronic components, he says.

But some researchers see a way around the switching problem: Abandon conventional circuits and use a single large laser, bright enough to force a strong response even from existing nonlinear materials, to drive many computations at once. That’s the reasoning behind efforts to encode light beams with images or digital information and mix them in nonlinear media as diverse as gases, crystals, and bacterial proteins. The nonlinear medium allows information in one light beam to affect how the material “processes” a second beam. In effect, the medium performs many computations in parallel, enabling “holographic association,” as it is called, to compare two data sets tens of thousands of times faster than existing supercomputers can. “You can go through a database on the fly,” says Robert Birge, who is testing such systems at Syracuse University in New York.

The potential parallelism of holographic association is nothing, however, compared with that of DNA computing, in which up to 2^{70} DNA molecules (a few liters’ worth) act as individual “processors.” In essence, the nucleotide sequence of each molecule encodes a possible solution to the problem. By applying the techniques of molecular biology to clone, combine, and select subsets of the molecules, the operator of this biochemical

computer can force the system to sort through the entire astronomical range of possible solutions, leaving the correct sequence to be extracted and read out.

Each biochemical step can take minutes to hours, compared to a billionth of a second per operation for a supercomputer—but each step also acts on the whole panoply of molecules at once. The result, in the most optimistic scenarios, is a theoretical advantage of anywhere from 10^8 to 10^{12} in computing speed. The advantage holds, however, only for problems that can be solved efficiently by following many parallel computational paths, such as breaking cryptographic codes or finding routes to visit each of n cities exactly once. And DNA computing has already seen ups and downs in its brief history, which began a year ago with a paper published in *Science* by Leonard Adleman of the University of Southern California (11 November 1994, p. 1021).

In that paper, Adleman reported solving the path problem for seven cities using vanishingly small quantities of DNA. But a problem of that size can easily be solved with pencil and paper. Adleman now calculates that using the same algorithm to solve “even a problem of modest size”—involving, say, 50 or 100 cities—would require “tons of DNA.” Later, Penn State’s Beaver found that an apparently promising algorithm for factoring a 300-digit number actually called for “an ocean [of DNA] the size of the universe.”

Researchers in the field aren’t giving up on DNA as a computing medium. They think the answer lies in algorithms that “scale” better from simple situations (a few cities or small numbers, for example) to the more complicated and interesting ones. “I think we fairly understand what might and might not be possible,” says Dan Boneh, one of Lipton’s collaborators at Princeton. Boneh cautions, however, that other practical problems could crop up. In a year or two, says Adleman, the technique’s ultimate potential should come into focus.

Whatever the outcome of all these efforts to find computing’s future, says Adleman, “it seems as if the time is right for revisiting our notion of what a computer is.” For decades, he and his colleagues point out, computer scientists have regarded improved machines as nothing more than faster versions of the step-by-step mechanical devices of the 1930s—a notion that has shaped thinking about the conceptual “difficulty” of computing problems. Even parallel computing in its present form, with 100 or 1000 processors churning in tandem within a supercomputer, didn’t seem to challenge this mindset. But theorists fear it will finally crumble in the face of the outlandish parallelism of schemes such as DNA and quantum computing.

“It’s going to be a while,” says Lipton, “before we know what a computer is again.”

—James Glanz