NEUROMORPHIC ENGINEERING

Why Can't a Computer Be More Like a Brain?

Computer scientists may tout their machines' abilities to perform millions or billions of operations a second, but neuroscientist Christof Koch of the California Institute of Technology (Caltech) in Pasadena thinks they need a lesson in humility. "Any creature vastly outperforms any machine today," he says. "Not in logical thought, but in sensing the environment, smelling, seeing, moving about." The lesson, as Caltech's Carver Mead recognized in the early 1980s, is that biological systems are fantastically efficient at certain types of computation. Inspired by Mead's vision, some computer scientists are taking this lesson to heart—and they're using it to try to build a new kind of computer.

In an effort they call neuromorphics, researchers are capturing in silicon what Ralph Etienne-Cummings of Southern Illinois University (SIU) in Carbondale likes to call the "essence" of biological subsystems. Neuromorphic engineers, adds Koch, are "essentially adapting those features, those algorithms, those tricks that the nervous system came up with through the last 600 million years."

Those tricks include neurons' ability to change their behavior based on experience and to work as tiny, autonomous computers in their own right, performing operations that might take thousands of transistors in a conventional computer. By mimicking these abilities in silicon, says Koch, neuromorphic engineers are producing "systems [that] are able to do many useful tasks that have so far proved impossible." Research teams worldwide are using conventional Very Large Scale Integrated (VLSI) circuit technology to build these new kinds of chips, trying to reproduce elements of sensing and sense processing from the animal kingdom.

"Silicon retinas"—eyes on a chip that sense a scene or pattern and interpret it—are among the best developed products of this approach. These systems implement in silicon the kinds of neural circuits that control vision in biological systems, enabling them to perceive such features of a scene as brightness contrasts and the polarization, or orientation, of light. "We can now go out there and see the world with eyes that we did not have before," says Andreas Andreou of Johns Hopkins University. Neuromorphic engineers are also developing brainlike hardware to detect drugs and explosives, to generate music, and to allow vehicles to drive themselves, to mention just a few ongoing efforts.

Conventional digital computers speak only in ones and zeros, their millions of transistors linked into vast arrays of logic gates that require huge numbers of switchings to perform the most modest tasks. Even the simplest operation inside a computer, such as multiplication, requires at least 10,000 switchings. All this happens serially, in a strict sequence controlled by a systemwide clock that synchronizes the activities of every component. Even so-called parallel supercomputers are, in reality, modest collections of smaller, ordinary computers joined together.

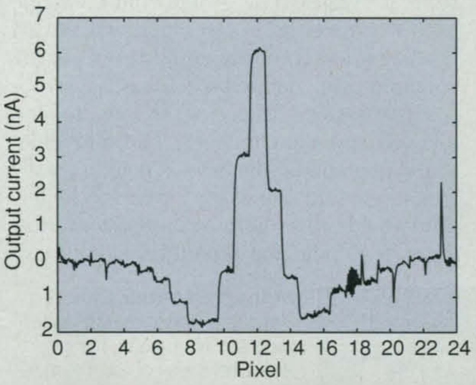

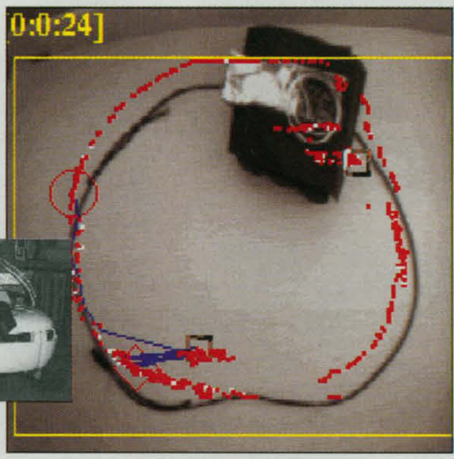

G. INDIVERI AND P. VERSCHURE

Holding a line. Koala, a robot guided by a silicon retina, faithfully follows a 2.4-meter loop independent of lighting and surface texture. A graph shows how a silicon retina responds to the image of a bar.

The brain, on the other hand, is totally different. As far as anybody knows, there is no systemwide clock in the brain: A neuron simply signals its neighbors when it is ready. What's more, "the individual components are very, very slow in a brain," says Koch. Yet the brain can perform 10^{16} operations per second, while consuming less power than an electric light bulb. To do the same amount of computation using a conventional digital chip would consume the output of an entire power station.

This stunning speed and efficiency results in part from "the massive, massive parallelism" of the brain's hundred million neurons, all working at the same time, says Koch. It also reflects the style of computing that goes on in neurons: analog computing, an approach that computer scientists have traditionally shunned. The digital computers of today solve a problem by imposing a computational recipe, or algorithm, on general purpose hardware. So a PC can, in principle and with enough time, solve any problem that a supercomputer can. Analog computers, by contrast, embody a specific computational problem in the actual physics of the hardware. Where digital computers traffic in ones and zeros, analog computers use continuously varying quantities. For example, a set of rods, springs, and weights can be turned into an analog computer—a mechanical model—to rapidly evaluate a bridge design.

Similarly, brains use electrical signals and varying membrane properties, instead of stretchy springs and weights, to do their job. Like the springs in the bridge design, which can stretch to many different lengths, a neuron has maybe a hundred internal electrical levels, giving it far more information content than the binary on-off of a digital switch. And instead of treating all inputs alike, as a digital circuit does, neurons can give added weight to pulses coming from certain favorite neighbors.

What's more, synapses—the junctions where one neuron receives input from another—act as little memory elements, aware of their previous inputs. "The circuitry you compute with is also the circuitry that remembers, in a sense," says Koch. That property dramatically reduces the need for data to be swapped during a computation, increasing efficiency. And like any analog computer, neurons are much faster than their digital counterparts.

Even though the transistors in standard integrated circuits are restricted to flipping between "on" and "off" states, they are capable of mimicking some of this behavior. Simple circuits already exist in which transistors are used in analog mode—not as switches, but as amplifiers that can operate at many different voltages—to perform operations such as multiplication, division, and subtraction.

Neuromorphic engineering aims to go much further, by transforming microcircuitry into an analog computing medium resembling neural tissue. "If we use that [circuitry] in the peculiar way we do, we can generate physical processes that are similar to neurons," says Rodney Douglas, who heads the Institute of Neuroinformatics in Zurich, Switzerland.

Computer scientists have already tried to

A Subtler Silicon Cell for Neural Networks

Nature is the model for artificial neural networks. These networks of processors—either real or simulated on a conventional computer—"learn" from experience by adjusting the strength of their connections, much like networks of real neurons. Small neural nets have become commonplace, doing tasks such as predicting how stock prices may fluctuate and recognizing handwritten characters. But Nabil Farhat of the University of Pennsylvania, Philadelphia, thinks he can build a better neural net by making its constituent neurons even closer to biology.

Neural nets traditionally consist of so-called sigmoidal neurons, circuits that add up incoming signals until they reach a fixed threshold and then fire themselves. Farhat's so-called bifurcation neurons, in contrast, switch between different modes of operation—between regular and chaotic firing, for example—depending on subtler factors. These include not just the value of a particular train of incoming signals, but also the interaction between many incoming signals and the neuron's recent history. So far, Farhat has made only single neurons, and he hasn't linked them together into a complete network. But his latest simulations suggest, he says, that they could yield neural nets with more lifelike behavior than has been seen in networks to date, such as the ability to see, recognize, and even react to the world in real time.

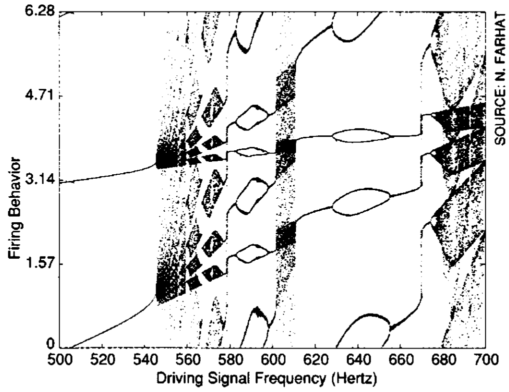

Nervous behavior. The bifurcation neuron switches among many different firing modes, depending on the frequency of the signal it receives.

Whether he will succeed is still an open question, says Daniel Collobert, a neural net expert at France Telecom. But he notes that Farhat's bifurcating neuron provides a level of behavioral complexity "that [artificial] neural networks could not previously [show], and I guess never will, because of the functional simplicity of [their] neurons."

The neurons in traditional nets sacrifice important information, because they know only how many spikes reached them in a given period, but not when each spike arrived. To an ordinary neuron, the periodic signal 110110110 (where a 1 is a spike) would be exactly the same as 101101101.

Yet real neuronal nets, in the brain, capture this timing information. A neuron fires because incoming signals cumulatively depolarize the excitable membrane of the neuron's output device, the axon. Afterward, there follows a period when the membrane cannot respond at all to an incoming signal, which gives way to another slow buildup. A pulse arriving immediately after firing will have a completely different effect on the output of the neuron from one that arrives immediately before.

Real neurons also respond differently to signals when they are correlated than when they arrive separately. Farhat explains that signals arriving simultaneously through different dendrites produce a periodic modulation of the neuron's electrically excitable membrane, just as two beams of light from the same source produce a periodic pattern of light and dark patches when they interfere with each other. The oscillation modulates the neuron's response to later inputs.

Conventional digital circuitry, with its arrays of on-off switches, can't efficiently mimic this kind of behavior. So Farhat has been combining resistors, capacitors, and other components into so-called analog circuits, which can adopt any intermediate state between "on" and "off." One proof-of-principle design incorporated two capacitors, which charge up in parallel as incoming signals build up. Eventually, one capacitor "breaks down," allowing current to flow, which switches on a light-emitting diode. The diode discharges both capacitors, reversing the breakdown and allowing the charging to begin again. This behavior makes the circuit time-sensitive: Signals arriving when the capacitors have just been discharged have a different effect from signals arriving earlier or later.

The latest incarnations of Farhat's neurons display more complex behavior (see diagram). When many neural inputs (spike trains) arrive at the same time, they interact to generate an electrical oscillation, which affects the neuron's firing. Small changes in the frequency of this oscillation, caused by varying input signals, produce huge shifts in the output behavior. For instance, oscillation frequencies below about 550 hertz produce periodic firing with two spikes per cycle. Just above this frequency, the output rapidly changes to chaotic firing.

Farhat and his colleagues are now planning to link such neurons into a full array using optical signals, which should allow them to create the dense thicket of interconnections needed for a large neural net. In a paper to be published in the *Journal of Intelligent and Robotic Systems* early next year, Farhat describes a computer simulation that offers a glimpse of how such a network would behave. The bifurcation neurons seem to form "netlets"—subsets of the neurons that work together. In an even more recent simulation, Farhat found that the netlets formed a kind of neuroanatomy, with different clusters of netlets responding to stimuli from different sources in the environment. "That's exactly the same as people observe when they look at functional MRI [magnetic resonance imaging] and PET [positron emission tomography] scans of the brain," Farhat says. "Depending on the inputs, the stimulus from the outside, or the cognitive task that the person is engaged in, we see different parts of the brain firing." It's that kind of complexity, Farhat thinks, that could make his networks of bifurcation neurons capable of simple abilities that we take for granted. **–Sunny Bains**

Sunny Bains is a science and engineering writer in Edinburgh, U.K.

mimic some aspects of the brain in so-called neural nets, networks of "processors" linked by "synapses"—connections that strengthen or weaken depending on activity, enabling the net to learn from experience. But these neural nets generally aren't real physical devices—instead, they are simulated ones, running as software on conventional computers. What's more, their neurons are, with few exceptions (see sidebar), generally much simplified versions of the real thing. Neuromorphics, on the other hand, is an effort to capture some of the richness of actual neurons in hardware—transistors, capacitors, and resistors, all fabricated onto silicon chips—in what is called analog VLSI, or simply AVLSI.

Besides allowing transistors to operate at many different voltage levels, neuromorphic engineers are designing them to serve as both calculation and memory elements. Work by Lance Glasser at the Massachusetts Institute of Technology, and by Mead and his team at Caltech, has led to the design of a new type of transistor, the floating gate transistor, which can reliably store analog information as electrical charge, enabling it to keep track of previous signals fed to it. These new transistors should open the way to building systems that can learn from experience, as neural-net software does, but more efficiently, because the hardware itself is learning.

Labs around the world are already exploiting AVLSI to build silicon noses, ears, and especially eyes. At the University of Adelaide in Australia, Alireza Moini and his team have built a succession of "bug-eye" chips, using a design abstracted from insect eyes. One variant senses motion by tracking regions of changing light intensity—an ability that could lead to collision sensors for cars. Andreou, together with Kwabena Boahen at Caltech, has built the most advanced silicon retina yet, with a resolution of 210 by 230 pixels. And Tamás Roska and his team at the Computer and Automation Institute in Budapest, Hungary, have produced a programmable "visual microprocessor" that can analyze a scene and swiftly pick out patterns for applications such as medical diagnosis.

In Zurich, Paul Verschure and Giacomo Indiveri are working on a silicon-retina-controlled robot that can follow a line across the floor of their laboratory. "We showed that we can reliably track an edge, independent of colors and textures of the surface," says Verschure. Their robot can stay on course for more than 100 meters—"the time we got fed up looking at the device doing the same thing."


ETU/UNIV. OF ADELAIDE

Collision avoider. Toy cars, one equipped with a silicon retina that detects a second car approaching from behind, and designer Alireza Moini.

Mimicking biology does have its disadvantages. Relying on continuously varying electrical states rather than the clearly defined ones and zeros of digital computing means that tiny variations in the components may cause them to respond differently to identical inputs. In silicon retinas, for example, pixels may vary in performance by as much as 20%. The brain can interpret sensory information reliably in spite of that kind of variability, thanks to the huge numbers of interconnections that allow it to smooth and correct data, says Koch.

Mimicking those dense interconnections is the field's other great challenge. "What the brain has that we do not have is connection technology," says Koch. Each cubic centimeter of the brain contains 100,000 cells and 2 kilometers of wiring, enabling each neuron to talk to 10,000 others. "We don't have that kind of technology right now." In the future, optical interconnections, relying on pulses of light that can crisscross freely, could solve the wiring problem. And nearly 10 years ago, the late Misha Mahowald at Zurich proposed a scheme to reduce the number of physical connections needed in a neuromorphic system. In her method, extended in a collaboration with Douglas, silicon neurons exchange addresses rather than actual pulses. A sender neuron communicates with its target over a common line linking many neurons, telling it to expect a signal from a particular address. The receiving neuron then recreates the signal, as if it had come over a dedicated line from the sender.

Even if they solve these problems, neuromorphic engineers are under no illusions about displacing conventional computing technology, which is unbeatable for number crunching. Ultimately, they hope to create neuromorphic sensors that can feed their readings of the world around them to digital electronics for subsequent processing, says Koch, who, for example, envisages cameras with a neuromorphic "seeing end" feeding a conventional data-processing unit. Says SUI's Etienne-Cummings: "If engineers can mimic the benefits of biological organisms while capitalizing on the speed of [digital] electronics, the resulting computational systems can be very powerful."

–Andrew Watson

Andrew Watson is a science writer in Norwich, U.K.

ARTIFICIAL LIFE

After 50 Years, Self-Replicating Silicon

The workings of living things are an inspiration to avant garde computer scientists, but so far the simple act of reproduction has them stymied. In fact, it's defeated them since the late 1940s, when the legendary computer scientist John Von Neumann first tried to see whether a computer could be made to reproduce. He managed to conceptualize a self-replicating computer using cellular automata—identical computing devices arranged in a checkerboard pattern that change their state based on the states of their nearest neighbors. But his scheme called for an enormously complicated device made of millions of 29-state cellular automata, if not more. "It was so big," says Stanford University's John Koza, "nobody has ever even done a simulation."

Now researchers at the Swiss Federal Institute of Technology in Lausanne are on the verge of achieving in practice what Von Neumann could only work out in theory—and they are doing so in a far smaller system. In the September issue of the journal *Robotics and Autonomous Systems*, Daniel Mange and his colleagues report that they have made a self-repairing, self-replicating version of a specialized computer. It's able to perform only one specific task, but they hope to do the same soon with a "universal" computer—a necessary step, says Koza, toward creating computers that truly mimic life by reproducing and evolving.

Like Von Neumann's scheme, the Swiss system is based on cells of identical processors, which they call "biodules." Each cell contains a random-access memory and a single field programmable gate array, which is a collection of circuits that can be rewired by software, allowing it to assume new functions (see p. 1931). The biodules are laid out in a two-dimensional array, with a "mother cell" at one corner. Each one is programmed with an artificial chromosome—a string of bits that encodes all the information necessary for all the cells to function together as a computer.

Mange explains that each cell uses the mother cell as a reference point to calculate its position in the array, extracts from the bit string the information that a cell at that position needs to carry out its particular functions, and wires itself accordingly. The resulting computer can perform just one task: checking a string of parentheses to see if every left parenthesis belongs to a closed pair.

The system is able to repair itself by enlisting spare cells that sit off to one side of the working array. When a cell is identified as faulty, its entire column is deactivated. Then the functions of each column are shifted one column over, so that a spare column takes over the function of what used to be the last working column of the computer. Mange suggests that such a system might have applications in avionics, for instance, for computers that require extraordinary fault tolerance, but he admits that there is a "rather high" price to pay in efficiency: the need to store the complete "genome" in every cell. "It's the same price biology agrees to pay with every living being to have a very safe architecture," he says.

Self-replication is an extension of the same idea. Mange and his colleagues have shown that with enough spare cells in the array, all of the working cells of the computer can simply copy themselves into a new set of cells. Moving on to a self-replicating universal machine should be relatively easy, says Mange. "We should be able to realize the original dream of Von Neumann in the very near future," he says.

–Gary Taubes