

SMART CONNECTIONS

Computer chips inspired by human neurons can do more with less power.

Kwabena Boahen got his first computer in 1982, when he was a teenager living in Accra. “It was a really cool device,” he recalls. He just had to connect up a cassette player for storage and a television set for a monitor, and he could start writing programs.

But Boahen wasn’t so impressed when he found out how the guts of his computer worked. “I learned how the central processing unit is constantly shuffling data back and forth. And I thought to myself, ‘Man! It really has to work like crazy!’” He instinctively felt that computers needed a little more ‘Africa’ in their design, “something more distributed, more fluid and less rigid”.

Today, as a bioengineer at Stanford University in California, Boahen is among a small band of researchers trying to create this kind of

BY M. MITCHELL WALDROP

computing by reverse-engineering the brain.

The brain is remarkably energy efficient and can carry out computations that challenge the world’s largest supercomputers, even though it relies on decidedly imperfect components: neurons that are a slow, variable, organic mess. Comprehending language, conducting abstract reasoning, controlling movement — the brain does all this and more in a package that is smaller than a shoebox, consumes less power than a household light bulb, and contains

nothing remotely like a central processor.

To achieve similar feats in silicon, researchers are building systems of non-digital chips that function as much as possible like networks of real neurons. Just a few years ago, Boahen completed a device called Neurogrid that emulates a million neurons — about as many as there are in a honeybee’s brain. And now, after a quarter-century of development, applications for ‘neuromorphic technology’ are finally in sight. The technique holds promise for anything that needs to be small and run on low power, from smartphones and robots to artificial eyes and ears. That prospect has attracted many investigators to the field during the past five years, along with hundreds of millions of dollars in research funding from agencies in both the United States and Europe.



NEW ANGLES ON THE BRAIN

A Nature special issue

www.nature.com/neuroscience2013

RAMIN RAHIMIAN

Kwabena Boahen holds a 'neuromorphic' circuit board from his Neurogrid device.

Neuromorphic devices are also providing neuroscientists with a powerful research tool, says

Giacomo Indiveri at the Institute of Neuroinformatics (INI) in Zurich, Switzerland. By seeing which models of neural function do or do not work as expected in real physical systems, he says, "you get insight into why the brain is built the way it is".

And, says Boahen, the neuromorphic approach should help to circumvent a looming limitation to Moore's law — the longstanding trend of computer-chip manufacturers managing to double the number of transistors they can fit into a given space every two years or so. This relentless shrinkage will soon lead to the creation of silicon circuits so small and tightly packed that they no longer generate clean signals: electrons will leak through the components, making them as messy as neurons. Some researchers are aiming to solve this problem with software fixes, for example by using statistical error-correction techniques similar to those that help the Internet to run smoothly. But ultimately, argues Boahen, the most effective solution is the same one the brain arrived at millions of years ago.

"My goal is a new computing paradigm," Boahen says, "something that will compute even when the components are too small to be reliable."

SILICON CELLS

The neuromorphic idea goes back to the 1980s and Carver Mead: a world-renowned pioneer in microchip design at the California Institute of Technology in Pasadena. He coined the term and was one of the first to emphasize the brain's huge energy-efficiency advantage. "That's been the fascination for me," he says, "how in the heck can the brain do what it does?"

Mead's strategy for answering that question was to mimic the brain's low-power processing with 'sub-threshold' silicon: circuitry that operates at voltages too small to flip a standard computer bit from a 0 to a 1. At those voltages, there is still a tiny, irregular trickle of electrons running through the transistors — a spontaneous ebb and flow of current that is remarkably similar in size and variability to that carried by ions flowing through a channel in a neuron. With the addition of microscopic capacitors, resistors and other components to control these currents, Mead reasoned, it should be possible to make tiny circuits that exhibit the same electrical behaviour as real neurons. They could be linked up in decentralized networks that function much like real neural circuits in the brain, with communication lines running between components rather than through a central processor^{1,2}.

By the 1990s, Mead and his colleagues had shown it was possible to build a realistic silicon neuron³ (see 'Biological inspiration'). That

device could accept outside electrical input through junctions that performed the role of synapses, the tiny structures through which nerve impulses jump from one neuron to the next. It allowed the incoming signals to build up voltage in the circuit's interior, much as they do in real neurons. And if the accumulating voltage passed a certain threshold, the silicon neuron 'fired', producing a series of voltage spikes that travelled along a wire playing the part of an axon, the neuron's communication cable. Although the spikes were 'digital' in the sense that they were either on or off, the body of the silicon neuron operated — like real neurons — in a non-digital way, meaning that the voltages and currents weren't restricted to a few dis-

"WE ENVISION BUILDING FULLY AUTONOMOUS ROBOTS THAT INTERACT WITH THEIR ENVIRONMENTS IN A MEANINGFUL WAY."

crete values as they are in conventional chips.

That behaviour mimics one key to the brain's low-power usage: just like their biological counterparts, the silicon neurons simply integrated inputs, using very little energy, until they fired. By contrast, a conventional computer needs a constant flow of energy to run an internal clock, whether or not the chips are computing anything.

Mead's group also demonstrated decentralized neural circuits — most notably in a silicon version of the eye's retina. That device captured light using a 50-by-50 grid of detectors. When their activity was displayed on a computer screen, these silicon cells showed much the same response as their real counterparts to light, shadow and motion⁴. Like the brain, this device saves energy by sending only the data that matters: most of the cells in the retina don't fire until the light level changes. This has the effect of highlighting the edges of moving objects, while minimizing the amount of data that has to be transmitted and processed.

CODING CHALLENGE

In those early days, researchers had their hands full mastering single-chip devices such as the silicon retina, says Boahen, who joined Mead's lab in 1990. But by the end of the 1990s, he says, "we wanted to build a brain, and for that we needed large-scale communication". That was a huge challenge: the standard coding algorithms for chip-to-chip communication had been devised for precisely coordinated digital signals, and wouldn't work for the more-random spikes created by neuromorphic systems. Only in the 2000s did Boahen and others devise circuitry and algorithms that would work in this messier system, opening the way for a flurry of development in large-scale neuromorphic systems.

Among the first applications were large-scale emulators to give neuroscientists an easy way

to test models of brain function. In September 2006, for example, Boahen launched the Neurogrid project: an effort to emulate a million neurons. That is only a tiny chunk of the 86 billion neurons in the human brain, but enough to model several of the densely interconnected columns of neurons thought to form the computational units of the human cortex. Neuroscientists can program Neurogrid to emulate almost any model of the cortex, says Boahen. They can then watch their model run at the same speed as the brain — hundreds to thousands of times faster than a conventional digital simulation. Graduate students and researchers have used it to test theoretical models of neural function for processes such as working

memory, decision-making and visual attention.

"In terms of real efficiency, in terms of fidelity to the brain's neuronal networks, Kwabena's Neurogrid is well in advance of other large-scale neuromorphic systems," says Rodney Douglas, co-founder of the INI and co-developer of the silicon neuron.

But no system is perfect, as Boahen himself is quick to point out. One of Neurogrid's biggest shortcomings is that its synapses — of which there is an average of 5,000 per neuron — are simplified connections that cannot be modified individually. This means that the system cannot be used to model learning, which occurs in the brain when synapses are modified by experience. Given the limited space available on the chip, squeezing in the complex circuitry needed to make each synapse behave in a more realistic manner would require circuit elements about a thousand times smaller in area than they are at present — in the realm of nanotechnology. This is currently impossible, although a newly developed class of nanometre-scale memory devices called 'memristors' could someday solve the problem.

Another issue stems from inevitable variations in the fabrication process, which mean that every neuromorphic chip performs slightly differently. "The variability is still much less than what is observed in the brain," says Boahen — but it does mean that programs for Neurogrid have to allow for substantial variations in the silicon neurons' firing rates.

This issue has led some researchers to abandon Mead's original idea of using sub-threshold chips. Instead, they are using more conventional digital systems that are still neuromorphic in the sense that they mimic the electrical behaviour of individual neurons, but are more predictable and much easier to program — at the cost of using more power.

A leading example is the SpiNNaker Project,

BIOLOGICAL INSPIRATION

Neuromorphic technology is based on neurons and neural circuits in the brain. Like the brain, it uses much less power than standard computer chips.

Synapses

Each neuron has about 10,000 of these tiny junctions, which receive signals in the form of voltage spikes coming in from other neurons.

Neuron cell body

In both real neurons and neuromorphic versions, voltages and currents vary smoothly rather than jump in digital fashion from one discrete value to another.

Axon

This fibre, which can be up to 1 metre long, transmits the voltage spikes to other neurons.

Connections


Each emulated neuron receives signals through several thousand of these links, which are often much simpler than biological synapses.

Emulated neuron

Both real and emulated neurons add up, or integrate, incoming signals until they pass a threshold and 'fire', producing an outgoing series of voltage spikes.

Wire

This mimics the axon, and carries the voltage spikes to other emulated neurons.



led since 2005 by computer engineer Steve Furber at the University of Manchester, UK. This system uses a version of the very-low-power digital chips — which Furber helped to develop — that are found in many smartphones. SpiNNaker can currently emulate up to 5 million neurons. These neurons are simpler than those in Neurogrid and burn more power, says Furber, but the system's purpose is similar: "running large-scale brain models in biological real time".

Another effort sticks with neuron-like chips, but boosts their speed. Neurogrid's neurons operate at exactly the same rate as real ones. But the European BrainScaleS project, headed by former accelerator-physicist Karlheinz Meier at Heidelberg University in Germany, is developing a neuromorphic system that currently emulates 400,000 neurons running up to 10,000 times faster than real time. This means it consumes about 10,000 times more energy than equivalent processes in the brain. But the speed is a boon for some neuroscience researchers. "We can simulate a day of neural activity in 10 seconds," Meier says.

Furber and Meier now have the money to push for bigger and better. Together they constitute the neuromorphic arm of the European Union's ten-year, €1-billion (US\$1.3-billion) Human Brain Project, which was officially launched last month. The roughly €100 million devoted to neuromorphic research will allow Furber's group to scale up his system to 500 million digital neurons; Meier's group, meanwhile, is aiming for 4 million.

The success of these research-oriented projects has helped to stoke interest in the idea of using neuromorphic hardware for practical, ultra-low-power applications in devices from phones to robots. Until recently, that hadn't been a priority in the computer industry. Chip designers could usually minimize energy consumption by simplifying circuit design, or splitting computations over multiple processor 'cores' that can run in parallel or shut down when they are not needed.

But these approaches can only achieve so much. Since 2008, the US Defense Advanced Research Projects Agency has spent more than \$100 million on its SyNAPSE project to develop compact, low-power neuromorphic technology. One of the project's main contractors, the cognitive computing group at IBM's research centre in Almaden, California, has used its share of the money to develop digital, 256-neuron chips that can be used as building blocks for larger-scale systems.

BRAIN POWER

Boahen is pursuing his own approach to practical applications — most notably in an as-yet-unnamed initiative he started in April. The project is based on Spaun: a design for a computer model of the brain that includes the parts responsible for vision, movement and decision-making. Spaun relies on a programming language for neural circuitry developed a decade ago by Chris Eliasmith, a theoretical neuroscientist at the University of Waterloo in Ontario, Canada. A user just has to specify

a desired neural function — the generation of instructions to move an arm, for example — and Eliasmith's system will automatically design a network of spiking neurons to carry out that function.

To see if it would work, Eliasmith and his colleagues simulated Spaun on a conventional computer. They showed that, with 2.5 million simulated neurons plus a simulated retina and hand, it could copy handwritten digits, recall the items in a list, work out the next number in a given sequence and carry out several other cognitive tasks⁵. That's an unprecedented range of abilities by neural simulation standards, says Boahen. But the Spaun simulation ran about 9,000 times slower than real time, taking 2.5 hours to simulate 1 second of behaviour.

Boahen contacted Eliasmith with the obvious proposition: build a physical version of Spaun using real-time neuromorphic hardware. "I got very excited," says Eliasmith, for whom the match seemed perfect. "You've got the peanut butter, we've got the chocolate!"

With funding from the US Office of Naval Research, Boahen and Eliasmith have put together a team that plans to build a small-scale prototype in three years and a full-scale system in five. For sensory input they will use neuromorphic retinas and cochleas developed at the INI, says Boahen. For output, they have a robotic arm. But the cognitive hardware will be built from scratch. "This is not a new Neurogrid, but a whole new architecture," he says. It will trade a certain amount of realism for practicality, relying on "very simple, very efficient neurons so that we can scale to the millions".

The system is explicitly designed for real-world applications. On a five-year timescale, says Boahen, "we envision building fully autonomous robots that interact with their environments in a meaningful way, and operate in real-time while [their brains] consume as much electricity as a cell phone". Such devices would be much more flexible and adaptive than today's autonomous robots, and would consume considerably less power.

In the longer term, Boahen adds, the project could pave the way for compact, low-power processors in any computer system, not just robotics. If researchers really have managed to capture the essential ingredients that make the brain so efficient, compact and robust, then it could be the salvation of an industry about to run into a wall as chips get ever smaller.

"But we won't know for sure," Boahen says, "until we try." ■

M. Mitchell Waldrop is a features editor for *Nature* based in Washington DC.

1. Mead, C. *Analog VLSI and Neural Systems* (Addison-Wesley, 1989).
2. Mead, C. *Proc. IEEE* **78**, 1629–1636 (1990).
3. Mahowald, M. & Douglas, R. *Nature* **354**, 515–518 (1991).
4. Mahowald, M. A. & Mead, C. *Sci. Am.* **264**, 76–82 (May 1991).
5. Eliasmith, C. et al. *Science* **338**, 1202–1205 (2012).