

Think like a bee

Are silicon circuits that mimic the nervous systems of insects and other animals the future of computing? Jim Giles considers the merits of neuromorphic engineering.

Rodney Douglas is very impressed with the cognitive powers of the honey bee. “Just look at what it does. It learns how to fly, it can navigate, it has rather sophisticated pattern recognition and some level of communication,” he says. “But it has very grubby sensors. Have you seen the world through a bee’s eye? It’s horrible.”


If Douglas were a zoologist, such enthusiasm would not be surprising. But his team at the Institute of Neuroinformatics in Zurich does not study animals — it makes silicon circuits. Douglas is a leading figure in the field of neuromorphic engineering, which attempts to create microelectronic circuitry that mimics the way animal brains work.

Digital computers are extremely good at precise number-crunching. But animals, argue neuromorphic engineers, have evolved brains and senses that let them interact efficiently with the messiness of the real world. So why not develop devices that combine the best aspects of both? “Stupid-looking organisms are doing amazing computation,” says Rahul Sarpeshkar, an electronics engineer at the Massachusetts Institute of Technology. “We need to replicate this in electronics.”

Mind mimics

Attempts to mimic animal brains using artificial neural networks date back to the 1940s. The basic idea is to link together individual processing units — the artificial neurons — that can integrate signals from other units in the network and send signals on to another group of units. Each of the units can only process a limited amount of information, but combined into a network they should, in theory, mimic some of the functions of a real brain.

Neural networks really began to take off in




Fresh view: studies of insect eyes have helped Rodney Douglas to produce artificial retinas.

the early 1980s, when researchers developed sophisticated learning ‘rules’ that allowed them to ‘train’ the networks to recognize specific patterns^{1,2}. At that time, programmers were finding it difficult to make digital computers perform pattern recognition.

Signatures, for example, can be a problem. Every time we sign our names, we create patterns that are slightly different from each other. For an automated signature recognition system to work, it must recognize these different patterns as examples of the same signature. This is difficult to do using digital computers, but artificial neural networks can be taught the technique. Different examples of someone’s signature are presented to the input neurons, while the output neurons are set to keep giving the same output. By selectively strengthening connections between the output units and the various intermediary neurons triggered in response to different versions of the same person’s signature, networks can learn to identify signatures.

But building such a neural network using silicon is not easy. Strengthening the connections between individual neurons in the network is done by changing and maintaining the amount of charge stored at the junctions between them. For the pioneers of neuromorphic engineering, this was a big problem: every time the networks were switched off, the charges were lost and the networks ‘forgot’ the tasks they had learned.

Conventional computers do not suffer from such a problem, so instead of building networks from silicon, the researchers simulated them using software, and saved their trained states in a computer’s digital memo-




ry. Over the past three decades, simulated neural networks — some more brain-like than others — have been studied extensively. They have shed light on how real brains work and spawned a host of useful applications.

Faking it

But simulated networks deviate from the true spirit of neuromorphic engineering. In biology, brains and sensory organs work quickly and at low power. Simulations, however, are slow and need power-hungry computers. But thanks to research conducted from the late 1980s by Carver Mead and his team at the California Institute of Technology in Pasadena, genuine neuromorphic circuits are now being worked on by about a dozen research groups worldwide.

Mead was interested in creating silicon circuits based on analog electronics. In digital computer chips, problems are solved using strict algorithms. Numbers are represented in binary code — with ‘0’ and ‘1’ corresponding to distinct voltages — and a central ‘clock’ regulates how information is sent around the chip.

Analog circuitry seems chaotic by comparison. A range of voltages is used to represent different numbers, and signals flow



Thought provoking: can the adaptable attributes of networks of nerve cells (green) be married to the precise number-crunching of digital chips to produce a new generation of silicon circuits?

between different parts of the circuit without any central control. Precise algorithms are impossible to implement. Instead, the circuit's architecture is designed so that the 'natural' flow of signals produces useful processing — in much the same way that animal brains work.

One of Mead's biggest contributions was the invention of floating-gate analog structures³, devices that could reliably store charge for long periods of time. This helped to solve the memory problem, and paved the way for neuromorphic devices that simulate the function of the retina, the light-receptive layer at the back of our eyes.

The retina is far more than just a collection of photoreceptor cells. It performs computations, processing information to accentuate the edges of objects, called 'edge extraction', and adjusting the 'gain', or amplification of the signal, to compensate for bright or dark conditions. Powerful digital machines can replicate this 'preprocessing', but nervous systems do so using simple, low-power analog circuitry.

Mead wanted to replicate this simplicity,

and teamed up with the late Caltech biologist Misha Mahowald. The pair's original retina⁴ has since been improved by the invention of better photoreceptors⁵ and more sophisticated circuit design⁶. But the basic philosophy remains the same. Silicon retinas consist of photoreceptor arrays, in which each receptor is connected to its neighbours. A network of resistors, amplifiers and other devices allows the signals to flow between the receptors in real time. Designing this circuitry so that it physically mimics the networks in real retinas is impossible — the number of cells and connections between them is far too great. Instead, neuromorphics engineers work out how retinal networks do their pre-processing and then they design simpler analog networks to do the same job. Modern versions, such as the devices produced by Douglas and his Zurich colleagues⁷, now perform edge extraction and gain adjustment much as biological retinas do.

In 1992, Mahowald added another item to the tool-box of neuromorphic engineering in the shape of address-event-representation (AER)⁸ — a method that allows chips to com-



Carver Mead (seated, centre) and his Caltech team have pioneered chip-based neural networks.

municate with each other. Although neuromorphic chips often contain hundreds of interlinked artificial neurons, there is a limit to the number of connections a chip can make to the outside world. When two neuromorphic chips are connected together it is impossible to link all the individual neurons directly. In AER, incoming and outgoing signals associated with specific units are sent via a central 'bus'. Connections within chips are much easier to manage, so the bus can talk directly to individual units using their 'address' — a number that uniquely identifies every unit.

All-seeing mouse

Neuromorphics engineers can even point to a commercially successful product: the optical computer mouse, introduced in 1994 by Logitech of Fremont, California. Unlike a conventional mouse, which tracks movement using a ball in its base, the Logitech device detects changes in its position using a visual sensor that monitors movement by 'looking' at the desk below. The digital approach to such a problem involves comparing sequential snapshots of the scene as the sensor moves along — much too computationally demanding for a cheap device. Instead, Logitech commercialized a neuromorphic chip developed by André van Schaik, then at the Swiss Center for Electronics and Microtechnology in Zurich. Van Schaik had used his knowledge of a fly's visual system to design a cheap, low-power device⁹.

"The fly's brain compares the change in intensity at one photoreceptor with delayed versions from neighbouring receptors," says van Schaik, now at the University of Sydney in Australia. By analysing the output from neighbouring photoreceptors, cells in a fly's eye make estimates of the speed and direction of the moving objects it sees. But individual 'motion-detecting' cells give only crude estimates of these quantities. More accurate estimates of movement are generated when cells in the later stages of processing compare the output of the different motion-detecting cells. The optical mouse takes a similar approach.

Although relatively simple, Logitech's

mouse shows how analog neuromorphic devices can be combined with conventional digital computers, melding the latter's high-precision number-crunching with the fast, lower-power pattern analysis that animal brains have evolved to do. But if neuromorphic devices have so much potential, why are there only hundreds, rather than thousands, of researchers working in the field?

One reason is that analog computing has been a victim of the relentless improvements in digital-chip technology. Chip manufacturers are reluctant to invest in an analog speech-processing device, for example, when constantly improving computing power offers new ways of approaching the problem digitally.

Nervous niches

Given this dynamic, neuromorphics engineers are concentrating on niche applications in which the advantages of biologically inspired computing cannot be ignored. For researchers working on mobile devices that must function autonomously, for example, the energy efficiency of neuromorphic devices is appealing.

This is the logic behind using neuromorphic technology to produce 'bionic' implants. Current cochlear implants, used to restore hearing to some congenitally deaf people, contain mechanical versions of the hair cells that sense incoming sound waves. Artificial cochleae are relatively crude, using between 10 and 20 electrodes to simulate the input of 30,000 or so hair cells into the auditory nerve, but the results can be impressive. In the best cases, the implants allow their wearers to conduct telephone conversations.

But current devices use a bulky and power-hungry digital-signal processor that has to be worn externally. And the implant itself also needs recharging every few weeks, when the wearer is forced to sit next to a charging station for several hours.

Toumaz Technologies, a spin-off from Imperial College in London, now aims to produce a smaller, lower-power analog version of the digital-signal processor. Chris Toumazou, the electronics engineer behind the project, says the device, including both new, low-power electrodes and his processor, will fit within the ear and will only need recharging once a year. He hopes to begin clinical trials before the end of the year.

Toumazou is also starting work on turning the analog retinas pioneered by Mead and Mahowald into practical medical implants. This will be more difficult than developing cochlear implants because pre-processing in the retina is much more complex, and many more nerves are involved — there are 1 million fibres in each of our two optic nerves, compared with 30,000 in each auditory nerve.

Meanwhile, Ralph Etienne-Cummings, an electronics engineer based at Johns Hop-


Chips offer insights into vision

Neuromorphics engineers are professional plagiarists, freely borrowing their ideas from nature's bag of computational tricks. But the flow is not one way. By building silicon models of animal brains, researchers can also learn about biology.

Many neuromorphics engineers pay lip service to this idea, but researchers at the Institute of Neuroinformatics in Zurich are putting it to work. Neuroscientist Kevan Martin and electronics engineer Shih-Chii Liu are collaborating on a project to investigate how our brains process visual images.

The project centres on a device made by Liu called the cortical chip, an analog silicon circuit based on part of the human visual system. Our brains analyse images using a series of different areas of the cerebral cortex. The first of these, known as 'V1', performs low-level processing such as detecting motion and the orientation of lines.

Liu's chip attempts to model two of V1's six layers. It uses a network of artificial neurons connected to mimic the links between layers 4 and 6. Layer 4 is interesting because it receives many of the inputs to V1. To make the chip as realistic as possible, the input to the chip is also taken from biology. Martin has recorded the activity of neurons in part of the cat brain involved in sending signals to



Visual clues: Shih-Chii Liu has designed an analog chip that models some of the brain processes involved in image analysis.

layer 4 of V1. Liu feeds these recordings straight into her chip.

Liu can adjust the connections and properties of the artificial neurons to see how they affect the performance of the whole network. By tweaking specific parameters — such as the number of links between layers 4 and 6 — Liu generates insights that Martin can relate to neurophysiological studies. "The number of questions we have to answer is overwhelming," says Martin.

For example, Martin is interested in how different

neurons in layer 4 form specific 'receptive fields' — in other words, why each responds to a different area of space within the eye's field of view. Each neuron's receptive field is thought to be controlled by inputs from another brain area known as the thalamus. But Martin suspects that there is more to it than that. "Fifty per cent of all connections to layer 4 neurons come from layer 6," he points out. "So what does layer 6 do? With this circuit, we can adjust the layer 6 synapses and use real data to find out."

kins University in Baltimore, is working on a circuit that will mimic the way the human spinal cord regulates muscle contraction in the legs during walking. Although we are not consciously aware of it, walking requires sophisticated and continuous real-time computation. Our spinal cord integrates information about our balance and leg positions to calculate the right set of muscle contractions. Etienne-Cummings has teamed up with Iguana Robotics of Mahomet, Illinois, to create a chip that might one day be implanted into the spinal cords of paraplegics to help them walk again.

Even the strongest enthusiasts for neuromorphic engineering accept that the current number of products — one computer mouse — is not that impressive. But a range of successful biological implants would be a differ-

ent matter. "It's been a bumpy road," says Andreas Andreou, a collaborator of Toumazou's based at Johns Hopkins, "but things are now looking very exciting."

Jim Giles is Nature's assistant News and Features editor.

1. Hopfield, J. J. *Proc. Natl Acad. Sci. USA* **79**, 2554–2558 (1982).
2. Rumelhart, D. E., Hinton, G. E. & Williams, R. J. *Nature* **323**, 533–536 (1986).
3. Mead, C. *Analog VLSI and Neural Systems* (Addison-Wesley, Reading, MA, 1989).
4. Mahowald, M. *An Analog VLSI System for Stereoscopic Vision* (Kluwer Academic, Boston, 1994).
5. Delbruck, T. & Mead, C. A. *CNS Memo #30* (Computation and Neural Systems Department, California Institute of Technology, Pasadena, 1994).
6. Boahen, K. A. *Analog Integr. Circ. S.* **13**, 53–68 (1997).
7. Liu, S.-C. in *Advances in Neural Information Processing Systems* Vol. 10 (eds Jordan, M. I., Kearns, M. J. & Solla, S. A.) 712–718 (MIT Press, 1998).
8. Mahowald, M. PhD thesis, California Institute of Technology (1992).
9. Arreguit, X., van Schaik, F. A., Bauduin, F., Bidiville, M. & Raeber, E. J. *Solid-State Circ.* **31**, 1916–1921 (1996).