Neuromorphic computing

The machine of a new soul

Computers will help people to understand brains better. And understanding brains will help people to build better computers

Aug 3rd 2013 | From the print edition

ANALOGIES change. Once, it was fashionable to describe the brain as being like the hydraulic systems employed to create pleasing fountains for 17th-century aristocrats' gardens. As technology moved on, first the telegraph network and then the telephone exchange became the metaphor of choice. Now it is the turn of the computer. But though the brain-as-computer is, indeed, only a metaphor, one group of scientists would like to stand that metaphor on its head. Instead of thinking of brains as being like computers, they wish to make computers more like brains. This way, they believe, humanity will end up not only with a better understanding of how the brain works, but also with better, smarter computers.
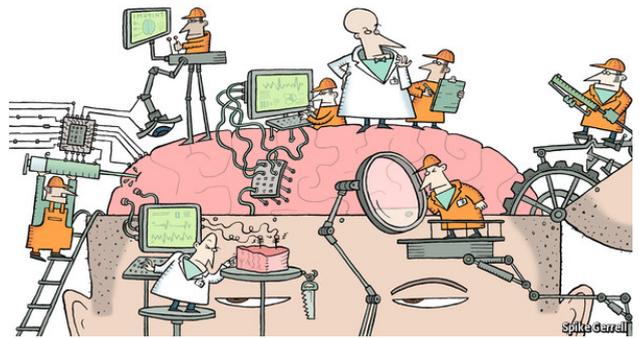
These visionaries describe themselves as neuromorphic engineers. Their goal, according to Karlheinz Meier, a physicist at the University of Heidelberg who is one of their leaders, is to design a computer that has some—and preferably all—of three characteristics that brains have and computers do not. These are: low power consumption (human brains use about 20 watts, whereas the supercomputers currently used to try to simulate them need megawatts); fault tolerance (losing just one transistor can wreck a microprocessor, but brains lose neurons all the time); and a lack of need to be programmed (brains learn and change spontaneously as they interact with the world, instead of following the fixed paths and branches of a predetermined algorithm).

To achieve these goals, however, neuromorphic engineers will have to make the computer-brain analogy real. And since no one knows how brains actually work, they may have to solve that problem for themselves, as well. This means filling in the gaps in neuroscientists' understanding of the organ. In particular, it means building artificial brain cells and connecting them up in various ways, to try to mimic what happens naturally in the brain.

Analogous analogues

The yawning gap in neuroscientists' understanding of their topic is in the intermediate scale of the brain's anatomy. Science has a passable knowledge of how individual nerve cells, known as neurons, work. It also knows which visible lobes and ganglia of the brain do what. But how the neurons are organised in these lobes and ganglia remains obscure. Yet this is the level of organisation that does the actual thinking—and is, presumably, the seat of consciousness. That is why mapping and understanding it is to be one of the main objectives of America's BRAIN initiative, announced with great fanfare by Barack Obama in April. It may be, though, that the only way to understand what the map shows is to model it on computers. It may even be that the models will come first, and thus guide the mappers. Neuromorphic engineering might, in other words, discover the fundamental principles of thinking before neuroscience does.

Two of the most advanced neuromorphic programmes are being conducted under the auspices of the Human Brain Project (HBP), an ambitious attempt by a confederation of European scientific institutions to build a simulacrum of the brain by 2023. The computers under development in these programmes use fundamentally different approaches. One, called SpiNNaker, is being built by Steven Furber of the University of Manchester. SpiNNaker is a digital computer—ie, the sort familiar in the everyday world, which process information as a series of ones and zeros represented by the presence or absence of a voltage. It thus has at its core a network of bespoke microprocessors.

The other machine, Spikey, is being built by Dr Meier's group. Spikey harks back to an earlier age of computing. Several of the first computers were analogue machines. These represent numbers as points on a continuously varying voltage range—so 0.5 volts would have a different meaning to 1 volt and 1.5 volts would have a different meaning again. In part, Spikey works like that. Analogue computers lost out to digital ones because the lack of ambiguity a digital system brings makes errors less likely. But Dr Meier thinks that because they operate in a way closer to some features of a real nervous system, analogue computers are a better way of modelling such features.

Dr Furber and his team have been working on SpiNNaker since 2006. To test the idea they built, two years ago, a version that had a mere 18 processors. They are now working on a bigger one. Much bigger. Their 1m-processor machine is due for completion in 2014. With that number of chips, Dr Furber reckons, he will be able to model about 1% of the human brain—and, crucially, he will be able to do so in real time. At the moment, even those supercomputers that can imitate much smaller fractions of what a brain gets up to have to do this imitation more slowly than the real thing can manage. Nor does Dr Furber plan to stop there. By 2020 he hopes to have developed a version of SpiNNaker that will have ten times the performance of the 1m-processor machine.

SpiNNaker achieves its speed by chasing Dr Meier's third desideratum—lack of a need to be programmed. Instead of shuttling relatively few large blocks of data around under the control of a central clock in the way that most modern computers work, its processors spit out lots of tiny spikes of information as and when it suits them. This is similar (deliberately so) to the way neurons work. Signals pass through neurons in the form of electrical spikes called action potentials that carry little information in themselves, other than that they have happened.
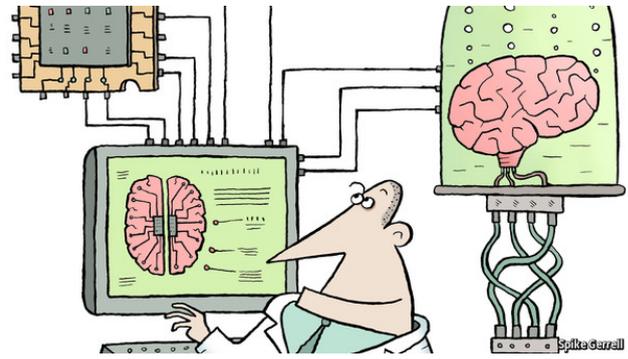
Such asynchronous signalling (so called because of the lack of a synchronising central clock) can process data more quickly than the synchronous sort, since no time is wasted waiting for the clock to tick. It also uses less energy, thus fulfilling Dr Meier's first desideratum. And if a processor fails, the system will re-route around it, thus fulfilling his second. Precisely because it cannot easily be programmed, most computer engineers ignore asynchronous signalling. As a way of mimicking brains, however, it is perfect.

But not, perhaps, as perfect as an analogue approach. Dr Meier has not abandoned the digital route completely. But he has been discriminating in its use. He uses digital components to mimic messages transmitted across synapses—the junctions between neurons. Such messages, carried by chemicals called neurotransmitters, are all-or-nothing. In other words, they are digital.

The release of neurotransmitters is, in turn, a response to the arrival of an action potential. Neurons do not, however, fire further action potentials as soon as they receive one of these neurotransmitter signals. Rather, they build up to a threshold. When they have received a certain number of signals and the threshold is crossed—basically an analogue process—they then fire an action potential and reset themselves. Which is what Spikey's ersatz neurons do, by building up charge in capacitors every time they are stimulated, until that threshold is reached and the capacitor discharges.

Does practice make perfect?

In Zurich, Giacomo Indiveri, a neuromorphic engineer at the Institute of Neuroinformatics (run jointly by the University of Zurich and ETH, an engineering university in the city) has also been going down the analogue path. Dr Indiveri is working independently of the HBP and with a different, more practical aim in mind. He is trying to build, using neuromorphic principles, what he calls "autonomous cognitive systems"—for example, cochlear implants that can tell whether the person they are fitted into is in a concert hall, in a car or at the beach, and adjust their output accordingly. His self-imposed constraints are that such things

should have the same weight, volume and power consumption as their natural neurological equivalents, as well as behaving in as naturalistic a way as possible.

Part of this naturalistic approach is that the transistors in his systems often operate in what is known technically as the "sub-threshold domain". This is a state in which a transistor is off (ie, is not supposed to be passing current, and thus represents a zero in the binary world), but is actually leaking a very tiny current (a few thousand-billionths of an amp) because electrons are diffusing through it.

Back in the 1980s Carver Mead, an engineer at the California Institute of Technology who is widely regarded as the father of neuromorphic computing (and certainly invented the word "neuromorphic" itself), demonstrated that sub-threshold domains behave in a similar way to the ion-channel proteins in cell membranes. Ion channels, which shuttle electrically charged sodium and potassium atoms into and out of cells, are responsible for, among other things, creating action potentials. Using sub-threshold domains is thus a good way of mimicking action potentials, and doing so with little consumption of power—again like a real biological system.

Dr Indiveri's devices also run at the same speed as biological circuits (a few tens or hundreds of hertz, rather than the hyperactive gigahertz speeds of computer processors). That allows them to interact with real biological circuits, such as those of the ear in the case of a cochlear implant, and to process natural signals, such as human speech or gestures, efficiently.

Dr Indiveri is currently developing, using the sub-threshold-domain principle, neuromorphic chips that have hundreds of artificial neurons and thousands of synapses between those neurons. Though that might sound small beer compared with, say, Dr Furber's putative million-processor system, it does not require an entire room to fit in, which is important if your goal is a workable prosthetic body part.

Unusually, for a field of information technology, neuromorphic computing is dominated by European researchers rather than American ones. But how long that will remain the case is open to question, for those on the other side of the Atlantic are trying hard to catch up. In particular, America's equivalent of the neuromorphic part of the Human Brain Project, the Systems of Neuromorphic Adaptive Plastic Scalable Electronics, SyNAPSE, paid for by the Defence Advanced Research Projects Agency, is also sponsoring two neuromorphic computers.

The Yanks are coming

One of these machines is being designed at HRL Laboratories in Malibu, California—a facility owned jointly by Boeing and General Motors. Narayan Srinivasa, the project's leader, says his neuromorphic chip requires not a single line of programming code to function. Instead, it

learns by doing, in the way that real brains do.

An important property of a real brain is that it is what is referred to as a small-world network. Each neuron within it has tens of thousands of synaptic connections with other neurons. This means that, even though a human brain contains about 86 billion neurons, each is within two or three connections of all the others via myriad potential routes.

In both natural brains and many attempts to make artificial ones (Dr Srinivasa's included) memory-formation involves strengthening some of these synaptic connections and pruning others. And it is this that allows the network to process information without having to rely on a conventional computer program. One problem with building an artificial small-world network of this sort, though, is connecting all the neurons in a system that has a lot of them.

Many neuromorphic chips do this using what is called cross-bar architecture. A cross-bar is a dense grid of wires, each of which is connected to a neuron at the periphery of the grid. The synapses are at the junctions where wires cross. That works well for small circuits, but becomes progressively less wieldy as the number of neurons increases.

To get around this Dr Srinivasa employs "synaptic time multiplexing", in which each physical synapse takes on the role of up to 10,000 virtual synapses, pretending to be each, in turn, for 100 billionths of a second. Such a system requires a central clock, to co-ordinate everything. And that clock runs fast. A brain typically operates at between 10Hz and 100Hz. Dr Srinivasa's chip runs at a megahertz. But this allows every one of its 576 artificial neurons to talk to every other in the same amount of time that this would happen in a natural network of this size.

And natural networks of this size do exist. *C. elegans*, a tiny nematode worm, is one of the best-studied animals on the planet because its developmental pathway is completely prescriptive. Bar the sex cells, every individual has either 959 cells (if a hermaphrodite) or 1,031 (if male; *C. elegans* has no pure females). In hermaphrodites 302 of the cells are neurons. In males the number is 381. And the animal has about 5,000 synapses.

Despite this simplicity, no neuromorphic computer has been able to ape the nervous system of *C. elegans*. To build a machine that could do so would be to advance from journeyman to master in the neuromorphic engineers' guild. Dr Srinivasa hopes one of his chips will prove to be the necessary masterpiece.

In the meantime, and more practically, he and his team are working with AeroVironment, a firm that builds miniature drones that might, for example, fly around inside a building looking for trouble. One of the team's chips could provide such drones with a brain that would, say, learn to recognise which rooms the drone had already visited, and maybe whether anything had changed in them. More advanced versions might even take the

controls, and fly the drone by themselves.

The other SyNAPSE project is run by Dharmendra Modha at IBM's Almaden laboratory in San Jose. In collaboration with four American universities (Columbia, Cornell, the University of California, Merced and the University of Wisconsin-Madison), he and his team have built a prototype neuromorphic computer that has 256 "integrate-and-fire" neurons—so called because they add up (ie, integrate) their inputs until they reach a threshold, then spit out a signal and reset themselves. In this they are like the neurons in Spikey, though the electronic details are different because a digital memory is used instead of capacitors to record the incoming signals.

Dr Modha's chip has 262,000 synapses, which, crucially, the neurons can rewire in response to the inputs they receive, just like a real brain. And, also like those in a real brain, the neurons remember their recent activities (which synapses they triggered) and use that knowledge to prune some connections and enhance others during the process of rewiring.

So far, Dr Modha and his team have taught their computer to play Pong, one of the first (and simplest) arcade video games, and also to recognise the numbers zero to nine. In the number-recognition program, when someone writes a number freehand on a touchscreen the neuromorphic chip extracts essential features of the scribble and uses them to guess (usually correctly) what that number is.

This may seem pretty basic, but it is intended merely as a proof of principle. The next bit of the plan is to scale it up.

One thing that is already known about the intermediate structure of the brain is that it is modular. The neocortex, where most neurons reside and which accounts for three-quarters of the brain's volume, is made up of lots of columns, each of which contains about 70,000 neurons. Dr Modha plans something similar. He intends to use his chips as the equivalents of cortical columns, connecting them up to produce a computer that is, in this particular at least, truly brainlike. And he is getting there. Indeed, he has simulated a system that has a hundred trillion synapses—about the number in a real brain.

After such knowledge

Spike Gerrell
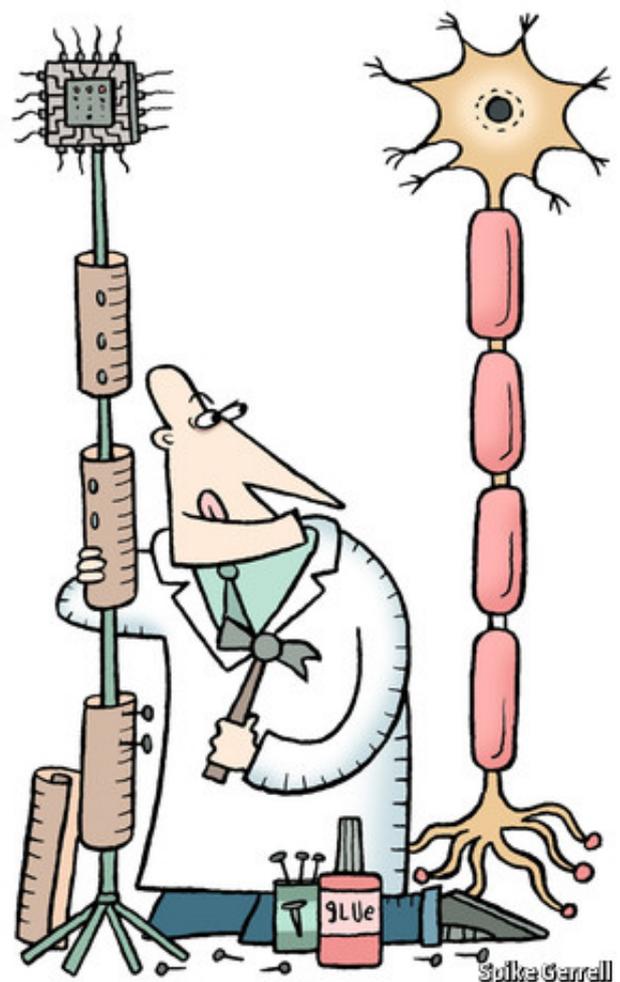
There remains, of course, the question of where neuromorphic computing might lead. At the moment, it is primitive. But if it succeeds, it may allow the construction of machines as intelligent as—or even more intelligent than—human beings. Science fiction may thus become science fact.

Moreover, matters may proceed faster than an outside observer, used to the idea that the brain is a black box impenetrable to science, might expect. Money is starting to be thrown at the question. The Human Brain Project has a €1 billion ($1.3 billion) budget over a decade. The BRAIN initiative's first-year budget is $100m, and neuromorphic computing should do well out of both. And if scale is all that matters, because it really is just a question of linking up enough silicon equivalents of cortical columns and seeing how they prune and strengthen their own internal connections, then an answer could come soon.

Human beings like to think of their brains as more complex than those of lesser beings—and they are. But the main difference known for sure between a human brain and that of an ape or monkey is that it is bigger. It really might, therefore, simply be a question of linking enough appropriate components up and letting them work it out for themselves. And if that works perhaps, as Marvin Minsky, a founder of the field of artificial intelligence put it, they will keep humanity as pets.