

Thinking BIG!

*How electronics technology is supporting attempts to model the brain. By **David Boothroyd**.*





Simulating complex objects and systems on computers is now a vast industry, giving us an understanding of things, such as the global climate, that would otherwise be impossible. While the weather may be complex, it is simple compared with the object now being modelled by projects worldwide: the human brain.

Generally recognised as the most complex known object in the universe, the human brain contains around 100 billion nerve cells or neurons, each linked by around 10,000 synapses. Every cubic millimetre of the cerebral cortex contains roughly 1bn synapses. Even by the standards of today's semiconductor chips, this is an impressive piece of packaging.

Despite the huge technological challenges of modelling such an object, serious progress is being made and it may not be that long before we have created elements of the brain that work very much as the real thing does. Major benefits could follow – for example, studying the effects of drugs, or understanding why brain diseases occur.

Some truly novel techniques are being developed to achieve brain modelling, two of them involving the use of analogue processing. One is the Europe wide project called FACETS (Fast Analogue Computing with Emergent Transient States). The aim is to build a brain like computer using data from studies of rat brains and, to do that, FACETS is adopting the technique of wafer scale integration (WSI), first suggested decades ago (but for different reasons) by the likes of Sir Clive Sinclair.

WSI does not slice up wafers to create individual chips, instead it links all the chips on the wafer. WSI holds out the prospect of creating an electronic system that has the potential to achieve something like the massive interconnection capability of the brain. Project coordinator Professor Karlheinz Meier of Heidelberg University explains.

"We started with single chips, each with 384 neurons and 100,000 synapses, which can be combined on a backplane holding 16 chips. The system now under development is going from chips to 8in wafers, each wafer holding up to 200,000 neurons (depending on complexity) and 50m synapses. Since the neurons are so small, the system runs 100,000 times faster than the biological equivalent and 10m times faster than a software simulation. We can simulate a day in one second."

The FACETS architecture uses analogue processing to emulate the performance of neurons and digital processing to carry out the communications between them. The basic processing unit is called HICANN (High Input Analogue Neural Network), which contains the mixed signal neuron and synapse circuits, together with the necessary support circuits and host interface logic.

Two crucial aspects of the FACETS system are fault tolerance and low power consumption.

"Already on the current wafer system it is essential not to lose the

entire device because of local manufacturing faults," Prof Meier says. "This will become even more important if one goes to nanoscale components, as other projects are looking to do."

"Power consumption is the major issue in achieving WFI for the analogue neural network. To limit the power consumption of event transmission, we have developed a novel asynchronous, differential low voltage signalling scheme. Also, the static power consumption of all circuits is minimised, and especially the synapse, which uses no static power at all. These techniques keep average power consumption to less than 1kW for a single wafer. If this is distributed uniformly across a 20cm silicon wafer, it equates to a power density of 1.6W/cm², well within the capabilities of standard air cooling methods, enabling us to mount the wafer systems densely in industry standard racks."

FACETS could lead to practical devices in less than five years, Prof Meier says. A possible application would be improved search tools, because the system should be able to make predictions based on previous experience.


Another brain modelling project underway at Stanford University has two similarities with FACETS – it exploits analogue processing and it also harks back to a past innovation; neuromorphic chips, pioneered in the 1980s by computer science legend Carver Mead.

"Mead developed the first silicon retina at Caltech and, in 1990, correctly predicted that present day computers would use 10m times more energy per instruction than the brain uses per synaptic activation," says Kwabena Boahen, associate Professor at Stanford's Bioengineering Department. "He sought to close this efficiency gap by building microelectronic circuits based on the brain, and succeeded in mimicking ion flow across a neuron's membrane with electron flow through a transistor's channel."

In 2004, the Stanford team developed a more sophisticated version of Mead's device, based on extensive analysis of the visual centres in various animal brains, resulting in a more detailed model of the mammalian retina and primary visual cortex than any previous ones.

This work has since been developed, resulting in a device called the Neurogrid, designed to model the cortex. As with FACETS, Neurogrid uses both analogue and digital signals. The analogue circuits simulate the ion channels in a brain cell's membrane, which allow the passage of a signal's charged particles into and out of the cell body. Digital signals then run to neighbouring cells and also to a memory location – a ram chip – which redirects the signal to other cells. Since this rerouting is variable, it functions like a plastic brain, changing connections whenever needed.

"Instead of designing different electronic circuits to emulate each of a variety of ion selective protein pores that stud neurons' membranes, as Mead did, we have produced a more versatile device that emulates the range of behaviour they display," Boahen says.



"Despite the huge technological challenges of modelling [the brain], serious progress is being made and it may not be that long before we have created elements of the brain that work very much as the real thing does."

Courtesy FACETS project (www.facets-project.org)



“For example, some open when the voltage across the membrane is high, others open when the voltage is low, and everything in between. As few as eight transistors sufficed to replicate this behaviour, allowing millions of distinct ion channel populations to be modelled with a single chip.”

The result is potentially dramatic. Neurogrid can simulate 1m neurons and 6bn synapses by using two subcellular compartments to minimise ion channel population count and by using local analogue communication to minimise bandwidth.

One way of assessing Neurogrid is to compare it with one of longest running and most ambitious brain modelling projects, Blue Brain, so called because it has been built on IBM’s Blue Gene supercomputer. At least in terms of power consumption, the result is dramatic: Neurogrid can simulate 1m neurons in real time, while consuming a million times less energy than Blue Brain: 1W instead of 1MW!

Despite its power consumption, Blue Brain – based at the Ecole Polytechnique Federale in Switzerland – is generally regarded as the most ambitious of all brain modelling projects. That is partly because, while other projects have aimed to model brain like computation or to simulate parts of animals’ brains at a relatively high level, Blue Brain aims to reverse engineer mammalian brains from laboratory data and to create an extraordinarily detailed model, right down to the level of the molecules that brain cells are built of.

Blue Brain began in 2005 and the first phase is now complete, with researchers having built a model of the neocortical column, which handles higher brain functions and thought. With around 8000 processors running in Blue Gene, this can simulate up to 50,000 fully complex neurons in real time, or about 100m simpler models of neurons.

In April, at a conference in Prague called Science Beyond Fiction, Blue Brain project leader Professor Henry Markram announced that the neocortical column is being integrated into a virtual reality agent, a model of an animal in a simulated environment, enabling researchers to observe detailed activities in the column as the animal moves around the space.

“It starts to learn and remember things. We can actually see when it retrieves a memory and where it retrieved it from because we can trace back every activity of every molecule, every cell, every connection and see how the memory was formed,” Prof Markram explained.

The next phase of the project will use a more advanced version of the IBM Blue Gene supercomputer and the aim now is to begin a ‘molecularisation’ process.

“We plan to add in all the molecules and biochemical pathways, which we couldn’t do on our first supercomputer,” Prof Markram says. “A very important reason for going to the molecular level is to link

gene activity with (the brain’s) electrical activity. Ultimately, that is what makes neurons become and work as neurons – an interaction between nature and nurture.”

The Blue Gene supercomputer has been used for another brain modelling exercise, by IBM’s own Almaden Research laboratory working with the University of Nevada. They built a model that they say was equivalent to around half a mouse’s brain, with some 55m neurons, each with more than 6000 synapses.

Recently, a new multidisciplinary project was announced, with the Almaden centre working with five US universities to integrate data from biological systems with the results of supercomputer simulations. This time, the aim is to develop a system with the level of complexity of a cat’s brain.

Leader Dr Dharmendra Modha claims a ‘perfect storm’ is occurring, with three different elements coming together – sufficiently powerful computing, detailed biological data and, now, electronics in the form of nanotechnology.

“Technology has only recently reached a stage in which structures can be produced that match the density of neurons and synapses from real brains – around 10bn per square centimetre. The real challenge is to manifest what will be learned from future simulations into real electronic devices using nanotechnology.”

Meanwhile, there are plenty of other brain modelling projects. Californian software company Numenta, founded by Palm Pilot inventor Jeff Hawkins,

is aiming to tackle problems using the structure and operation of the neocortex. Its system is called Hierarchical Temporal Memory and the company claims it will help in areas like machine vision, fraud detection, and semantic analysis of text.

In another project, Georgia Institute of Technology’s Paul Hasler is using field effect transistors that imitate the adaptable behaviour of synapses – ‘strengthening’ (that is, letting more current through) or ‘weakening’ (letting less through), depending on how often current flows through them.

Finally, a team led by neuroscientist Gerald Edelman at the Scripps Research Institute in San Diego has built a model of the circuitry in the thalamus and hippocampus, key elements in the brain. As the models complexity increases, they begin to generate their own activity, just as happens in the brain.

What’s more, oscillating waves of synchronous neural firings that are not explicitly built in to the models have emerged spontaneously.

The complexity of the brain is daunting and it will be many years before anything like human performance is approached. But the work on brain modelling that has begun in the first decade of the 21st Century will surely be looked back on as the start of something truly extraordinary.



The Blue Brain project at EPF Lausanne is using an 8192 processor version of IBM’s Blue Gene supercomputer in order to model the brain’s functionality in real time.



Image courtesy IBM