Building an Artificial Brain

By Peter AJ van der Made, CTO, vWISP Pty Ltd. USC Innovation Centre, Sippy Downs.

The NeoCortex processor and its neural model has been granted a patent on the 27th of November 2008. This is a completely new generation of computer chips that are modeled on the functions of neurons in the human brain. The prototype device has been produced using a programmable gate array made by Actel. Each neuron consists of 3500 logic gates. It has taken 5 years of research and design work to come up with this building block for an artificial brain. Next steps include a full custom device that holds 15,000 neurons. While useful to build intelligent toys, safety devices for cars, eliminate false alarms in security systems and many other applications, this is still a long way for the 100 billion neurons and 100 trillion synapses in the human brain.

The design of the chip has been kept free of compromises, exactly mimicking the function of biological neurons. Over the years, I have found that if you are going to use a computer system to simulate any phenomenon that every compromise you make will cause the result to be skewed in some way, no matter how small the compromise. It's a lot like the old adage about setting out on a long journey with an error of 0.05 degrees and finishing up somewhere completely off target. Precision is crucial. This was true for my previous malicious-code detection system, which uses a virtual computer to analyze computer code, and it definitely true for a complex system like the human brain.

Our digital device accurately simulates the processing method of the human brain. If you are into Artificial Intelligence (A.I.), you may think that this has been done before but in every other case, compromises have obliterated most similarities between the device and the brain. A point in case is the so-called 'neural network', which may be useful as a technology but offends the name it carries; Neural Networks are about as neural as an abacus is a computer. On top of that there is a lot of misunderstanding, misinformation, and contradiction among neuro-scientists. There is disagreement about the way information is encoded in the brain, about the function of feedback and a thousand other things. I had to study neurology for four years to be able to understand the information that was relevant to building this device.

Compromises have been part of the A.I. industry for so long that they are taken for granted and necessary to get something working. I believe that is because the structure of modern computers is unsuitable to simulate a brain. Even though both are a kind of information processing system, the differences are just too great. A computer is a calculating machine, it does not think, it does not know what it is doing- it just runs programs. My point becomes more obvious when we turn it around: that the brain is unsuitable to simulate a computer in the same way that the brain model unsuitable to be executed by a computer. No one would even consider that possibility.

Most scientists approach the brain as a mathematical model and then create programs to simulate that model on a computer. Where a computer can be useful is in building a model of the lowest level building blocks of the brain called ‘neurons’. Neurons are specialized cells that can communicate with one another and are connected to each other through synapses. Recent models of neurons behave very similar to real neurons. One such model is Eugene Izhikevich’s ‘simple model’ – a pure mathematical model that is, as its name implies, simple and produces realistic looking results. What is the problem then?

The first problem is one of computing power. This is a case of ‘the first will be last, and the last will be first’ so bear with me. The brain is generally seen as a slow processor consisting of neurons. The processing time of a single neuron is about 0.001 seconds. In contrast, a simple desktop PC has a cycle time of 0.0000000003 seconds and a simple calculation tells us that is 3 million times faster. Too bad we are not comparing apples with apples here, or maybe that is a good thing or else the machines would have taken over the world by now. A neuron’s function is far more complex than the operations performed by a computer in a single cycle. Realistically, to simulate a ‘simple neuron’ as defined by Eugene takes at least 0.0007 seconds on a 3GHz dual-core desktop PC, and that is not too far from the response time of a real neuron. We need to add about .0004 seconds for synapse processing. Now this gets interesting, because a single neuron can have up to 10,000 synapses. To process 10,000 synapses correctly would take far longer than 0.001 seconds. Bring on the compromises! However, we are not there yet. The brain consists out of an estimated 100 billion intricately connected neurons. A computer is a serial execution device; it executes one step after the other no matter how many hyper-threading, pipelining and other tricks are used. Leaving synapse functions aside for a moment, our best case to simulate the entire brain takes 100 billion times 0.0007 seconds or 810 days – for a single cycle. That is a bit too long to wait for the thing to do something. After about 5 minutes I’d think the machine is frozen and reboot it. We are going to need some more compromises here or speed our computer up 116 million times to get anywhere near the response time of the brain. Expanding on Moore’s law such an enormous speed increase is going to take the next 11,000 years, if at all physically possible. The computer would also need at least 14 Terabyte of memory and would consume hundreds of times the amount of energy that the brain does. Our brain consumes about 10 Watts. Our clever computer does not look too good now, does it? Why is our ‘slow’ brain so fast?

The way the brain is constructed is very alien to computer engineering. It’s ‘architecture’ is as different as you may expect of a computer of extra-terrestrial origin. There is no Central Processing Unit, no Arithmetic Logic Unit and no Random Access Memory. It does not run programs, algorithms or Boolean algebra. A Computer Science degree is not going to do you any good here – toss all you have learned about computers and open your mind, or it is just going to confuse you.

The brain processes everything in parallel. Those 100 billion neurons and 100 trillion synapses are all working at the same time, not one after the other like in our computer model. The second reason is one of structure or, as computer scientist like to call it, system architecture. Our brain has a highly efficient and interconnected hierarchical structure. The third reason is that all data is entering the processor in parallel and as separate events. Each input pulse is an event, not a code.

We are going to need some background to understand this third reason. If you are a computer boffin you would know about interrupts – or ‘exceptions’ as they are called now. An exception is an event that stops normal program flow, goes off and does something that is time-critical and important, and then returns to continue normal program flow. Every one of millions of inputs and all the pulses at interconnections between neurons in the brain are such events; all are time critical and are handled immediately. Both the timing of each sequence of events and the timing between different events is critical. To do this correctly we would need a computer with a 100 trillion parallel interrupt lines and the processing power to respond to each one instantly.

Another misconception is the idea that the brain performs complicated pattern recognition. It never compares patterns to look for a match. Everything the brain does is by association of temporal and spatial pulse patterns. If you are adding 2 and 3, it is because your eyes have ‘preprocessed’ the edges of the numbers and sent multiple temporal pulse trains to your brain. Your eyes contain neural cells that associate dark and light areas and detect edges. Your brain associates these pulse-trains to form a belief that it sees certain simple shapes and at the next level associates the simple shapes etcetera, to eventually form a belief that represents the numerals 2 and 3. You know what they are because in primary school you have learned to associate those shapes with the numbers 2 and 3 in the higher regions of your cortex. You have also learned to associate the shapes 2, a plus sign and 3 with 5. In the same way you have learned to associate the frequency pattern of the spoken word “two” with the word ‘two’ and number 2. Because the brain is a hierarchy of neurons many of the neurons that are triggered by visual association are also triggered by the audio association of these numbers – neurons are extensively ‘reused’. At each higher level the association becomes more sophisticated. Massive feedback paths modify the association at lower levels to reduce association errors. So next time you have trouble understanding something, consider that all that is missing are some of the neural paths required to associate the basics of whatever it is that you are trying to understand. Once you fix those missing associations, you can understand anything. Your brain is amazing; the only thing that holds us back is our attitude and our belief systems.

We learn to associate patterns by repetition. Recently a brain process called ‘STDP’ or Synaptic Time Dependent Plasticity was documented that shows that repetition modifies the connection strength of synapses. The effectiveness of synapses is constantly modified, whenever they contribute or do not contribute to the firing of a neuron. When they contribute, they are strengthened and when they do not contribute, they are weakened. The memory gets stronger when events are repeated, but we also remember things we have experienced only once, at least for a while. Invariance, the recognition of shapes in any size or angle, is an intrinsic property of the brain. That is because we do not look at pixels, but look at shapes associated in a hierarchy, which consists of primitive shapes at the lower levels.

The patented vWISP neural processor does not run programs, has no mathematical model stored anywhere and has no central processor or random access memory. It has been designed from the 'gate' level up. It is an exact simulation in binary logic of the functions of a biological neuron, with the underpinning that a hierarchical array of such devices, with feedback, will behave in the same way as a section of the brain. Each neuron receives a multitude of pulse trains and associates those pulse trains through a synaptic circuit coupled to a leaky integrator circuit. The output is a pulse train that expresses an associated belief.

The synapses are constantly updated through STDP to reinforce existing knowledge and to learn new associations. The NeoCortex-1 is a FPGA chip that contains a humble 10 of such neurons and 140 synapses. FPGA technology has been designed to work in 'normal' microprocessor-type circuitry, is not well suited to the very different logic circuit of an artificial neuron, so efficiency of standard-cell usage is very low but the device can be demonstrated, and has taught us much about the next generation of neural processors that we are designing. The target is to put 15,000 neurons and 240,000 synapses on a single chip. Such as chip would have the same complexity as a standard desktop computer processor with 240 million transistors, but would execute a neural algorithm 60 times faster and does not slow down as more devices are added to form a larger matrix.

The world is not in danger of being taken over by machines based on these devices though – 15,000 neurons is still a long way from the 100 billion neurons that make up the human brain. Sorry S.F. fans, no 'terminator' in the near future. These devices form the core of intelligent machines that perform tasks in areas where current computers perform poorly. Existing applications that are improved by the use of this technology are the recognition of spoken words and sentences, invariant image recognition, fingerprint recognition and association of images with sound. Applications are found in security, exploration, mining, robotics and medical prosthesis. New applications could be safety devices in cars that will stop a driver from taking dangerous risks. "You can not do that Dave" may be the words you hear when you are trying to take that corner a bit too fast.

Once trained, the training image can be read back and replicated to other devices. To train the system to recognize words we initially trained a device with a signal generator and a spectrum analyzer. The device responded to spikes representing a single frequency. The spectrum analyzer performs the function of a cochlea. We trained the next level in the hierarchy to recognize phonemes by using the knowledge the device had already learned.

A few additional advantages and challenges flow from the NeoCortex architecture. If a single transistor fails in a microprocessor, the whole system is affected. In contrast, the brain can lose tens of thousands of neurons without any significant effects. Even though the NeoCortex processor is just as reliable as any other processing device, it too can loose many neurons without affecting its function. This adds another layer of robustness to any intelligent system based on these processors.

There are a few challenges that flow from an exact simulation of the brain – the same challenges that exist in biological brains; when hierarchical arrays get very large it becomes increasingly difficult to control what the system learns. There is no way to ‘program’ the ‘three laws of robotics’ into an intelligent system that does not run programs. In addition, things like forgetfulness, imperfect recall, uncertainty and boredom may become issues for large intelligent machines.

That is, unless we think that we can do a better job than God can by using increasingly complex mathematical models and faster computers to build a perfect brain. Given where that science is now I would say that it is a tower of Babel, and will suffer the same consequences.