

A real-time simulation of the human brain

Steve Furber, 11 November 2004

Understanding the human brain is, arguably, the grandest challenge of them all. It is the most powerful information processing system known to mankind, with capabilities we are nowhere near to emulating with our artificial computing machines. Using fairly low performance unreliable component technology, humans can understand the visual world and control a complex of muscle systems in ways that surpass our engineering knowledge. Then there is the emergent phenomenon of natural language...

The human brain consists of 10^{11} neurons each with, on average, around 10^4 inputs. Neurons typically fire at a few 10s of Hz, so the computing power required to model the brain at this level must be able to handle around 10^{16} connections per second. If we knew exactly how each neuron behaved, how it is connected and how to model it, this would represent a formidable computing challenge. But we don't even know that much. There are major challenges in understanding the abstract computational function of a neuron (that is, what aspects of its natural behaviour are essential to its computational function as opposed to keeping it alive, providing energy, or gratuitous complexity resulting from its evolutionary heritage), the coding of information in neural connections, the high-level architecture of brain structures, and so on.

Progress is, however, being made. The Foresight Project into Cognitive Systems...

<http://www.foresight.gov.uk/servlet/Controller/ver=575/userid=2/>

...looked into the prospects for establishing inter-disciplinary activity that brings together neuroscientists and computer scientists to see if there is potential for mutual benefit from sharing their respective understandings of natural and machine intelligence. Analytical tools employed by neuroscientists are providing insights into the functions of neural systems in more detail than ever before, and computer models of natural systems are increasingly successful in modelling their behaviours.

Even partial progress towards this goal offers great benefits in the understanding (and therefore potentially the treatment of) brain pathologies and the development of computer systems that can emulate human capabilities more readily, easing human-computer interactions. Achieving the full objective will yield machines with capabilities which up to now have existed only in the imaginations of writers of science fiction.

Feasibility

The feasibility of this “vision” depends on computing resources (processing speed, memory, communications) and the availability of biological data. It also depends on less tangible factors such as a breakthrough in our understanding of neural codes, new mathematical models, neurons not being a whole lot more complex than they appear to be, and so on. Here I just look at the modelling issues and do some rough calculations on the machine resources that might be required:

Memory: If we assume that each of the 10^{15} synaptic connections requires a byte to indicate its strength the memory needed is 10^{15} bytes (regularity in the structure may reduce this significantly), which with memory today costing around \$1 for 10Mbytes would cost \$100M, reducing by a factor 1,000 over the next 20 years. This is clearly within the scope of a research budget, but not a high-volume consumer product.

Processing: The processing power required can only be guessed, but let's assume that each connection requires 10 instructions. We therefore require 10^{17} IPS or 10^{11} MIPS. Current general-purpose processors deliver 10^3 MIPS, rising to 10^4 MIPS over the next 10 years. 10^7 processors will therefore be required, costing perhaps \$100 each, so \$1,000M in all. As another reference point, high-end supercomputers are aiming for a petaflop (10^{15} FLOPS) over the next 5 years and will cost around \$100M. A lot depends here on the complexity of the neural algorithms. If they can be cast into hardware the cost may be reduced by several orders of magnitude; the estimates above suggest this is a requirement for success.

Communication: A hardware-based system might require around 10^5 nodes, each node modelling 10^6 neurons with 10^9 bytes of RAM and 10^6 MIPS. Each neuron would generate 10 events/s, so the communication requirement would be around 1Gbit/s per node, which is quite manageable today. The total communication bandwidth is 10^5 Gbit/s, so the communication architecture would have to exploit locality and be carefully designed to avoid bottlenecks.