# An Improved Interconnection Network for the Next Generation of Brain Simulators

Jonathan Heathcote

Short Report

## 1 Introduction

The human brain is the subject of intense research, a significant portion of which relies on simulations of large-scale models of the brain. Due to the unusual computational requirements of these simulations numerous specialised hardware platforms have been developed to cope [1, 2, 3, 4, 5]. This project will focus on the interconnection networks within such systems, in particular on network topology and power management.

Neural models of the brain typically consist of a directed graph of 'neurons' with simple, well defined behaviour and which communicate by sending 'spikes' to their neighbours. The only significant features of a spike are its origin and time of arrival. Such models have been demonstrated to reproduce high-level cognitive behaviours such as problem solving and memory [6].

Conventional super computer interconnection networks are poorly suited to human scale models containing tens of billions of neurons [7] each with thousands of connections [8]. Super computer networks are designed to handle the transmission of large blocks of data to a limited set of destinations. This is in sharp contrast with neural simulation where a spike may be sent to thousands of destinations and consists only of an origin ID and timestamp. Further, in realtime simulations timestamps may be omitted though the maximum delivery latency must be strictly limited. Due to this mismatch, alternative technologies are popular in neural simulators. Recent systems, such as SpiNNaker [1] and Bluehive [2], have begun to incorporate super computer style interconnect to benefit from increased bandwidth at the expense of power. It is here that this project hopes to contribute to a new generation of power efficient interconnection networks.

## 2 SpiNNaker

SpiNNaker is a hardware platform which combines low-power mobile-phone grade CPU cores using a custom, high-performance interconnection network to simulate large networks of neurons in realtime. The largest planned machine, which is in the final stages of manufacturing, is designed to simulate one billion ($10^9$) neurons in biological realtime [9] and will fill ten server-room cabinets. Figure 1 shows how over one million CPU cores are spread over thousands of chips and hundreds of circuit boards.

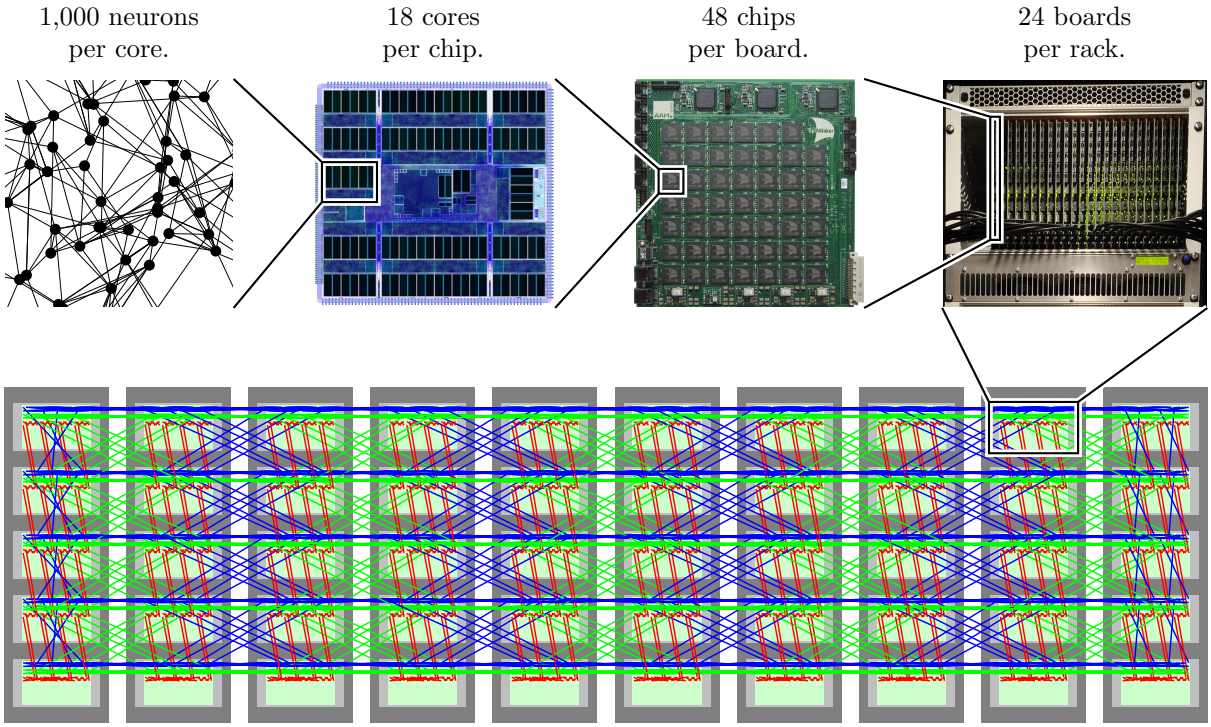


Figure 1: Construction of the largest planned SpiNNaker machine.

Between chips on the same board asynchronous, parallel signalling is used with groups of 16 wires required to communicate data between neighbouring chips [10]. This technique only consumes energy when the link is in use. Connecting the 768 wires required between boards would be prohibitively expensive so a different technology, High-Speed Serial (HSS), is used [11] reducing the number of wires to 24. HSS links can support high bandwidths and in SpiNNaker, one board-to-board HSS link carries the load of eight asynchronous links. Unfortunately, HSS links consume energy at the same rate when loaded and idle; to reduce power the links must be operated at a lower speed.

In conventional HSS applications, links change speed dynamically depending on load to cope with unpredictable traffic patterns [12]. Neural simulation traffic is also unpredictable but unfortunately, during link speed changes, HSS links become non-operational for around 1 ms [13]. To maintain realtime performance in SpiNNaker, the maximum latency for a spike traversing the network is also around 1 ms making delays during link speed changes unacceptable. As a result HSS links in SpiNNaker, as in other systems, are operated at full speed at all times despite average network loads between 5% and 10% of link capacity [2, 9].

This project will develop an improved HSS interconnect for neural simulations. Since SpiNNaker's HSS board-to-board links are implemented by reconfigurable FPGAs the platform lends itself to a prototype implementation.

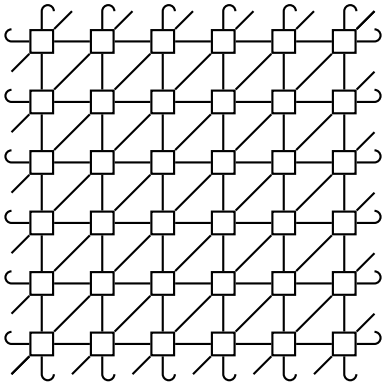

Figure 2: 2D hexagonal torus topology used in SpiNNaker. Each box represents an eighteen core SpiNNaker chip. Links at the edges of the network wrap-around to connect to those on the other side. Actual networks are much larger.

# 3 Preliminary work

Work has initially focused on developing improvements to existing neural interconnection networks and building infrastructure to enable their implementation within SpiNNaker.

## 3.1 SpiNNaker interconnect modelling

A detailed simulation model of SpiNNaker's interconnect was developed and compared with SpiNNaker hardware as part of a collaborative work exploring novel hardware-accelerated modelling of computer architectures. The result demonstrated a strong correspondence between the simulator and SpiNNaker indicating its suitability for use in the development of improvements to SpiNNaker's interconnect. This work is being prepared for journal submission in early September. This simulator will be exploited in upcoming work.

## 3.2 Wiring large systems with small-world connectivity

HSS technology imposes constraints on the lengths of wires used in large systems; for example in SpiNNaker, wires are limited to 1 metre or less [14]. SpiNNaker's network uses a 2D hexagonal torus topology (illustrated in figure 2) which, if laid out naïvely, would require cables long enough to span the whole system. For the large machine in figure 1, this would be over 6 metres. Preliminary work has developed a tool to aid the design of alternative physical organisations of SpiNNaker machines (such as in the figure) which maintain the same logical connectivity but only use wires less than a metre long.

Watts and Strogatz demonstrated that adding a small number of random links to a torus topology can convert it into a 'small-world' network [15]. An example of a small-world network is the human social network where the idea is popularly known as the theory of 'six degrees of separation'. The theory states that any two people are connected by a chain of no more than six acquaintances. This property is also desirable in computer networks since it reduces the maximum and average number of 'hops' messages must make when travelling through a network. This both reduces the latency of messages and also

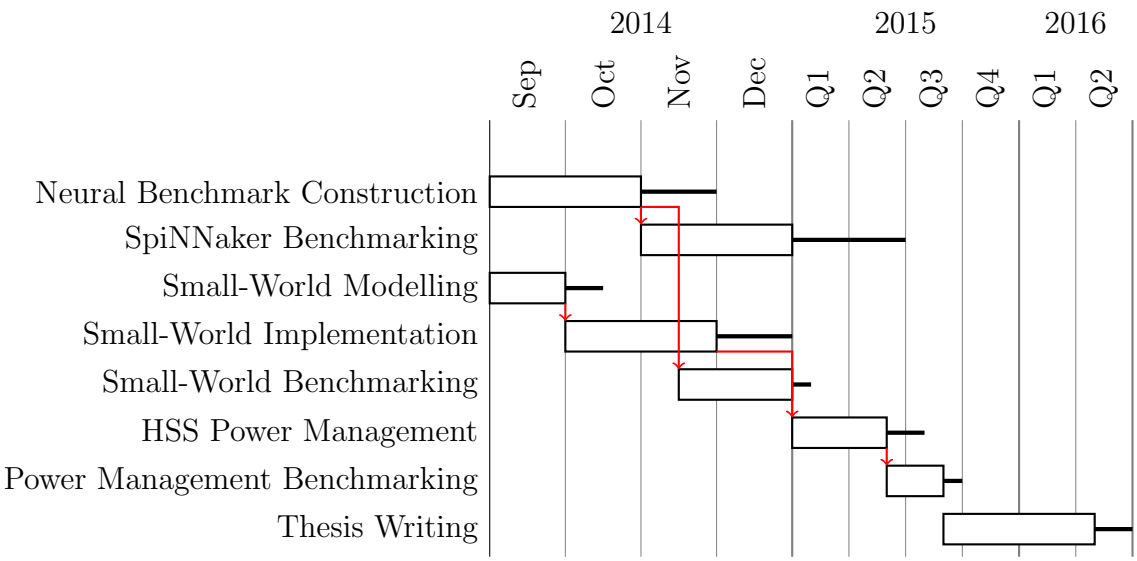


Figure 3: Research plan Gantt chart. Boxes indicate the expected duration of a task, thick lines indicate slack and red arrows show dependencies between tasks. Note the non-linear scale.

frees up resources on links [16]. This freeing of resources enables the network to operate at lower speeds and thus lower power.

Random links added to a network can easily require wires longer than a metre and allowing only short random connections was found to prevent small-world networks from forming. However, when systems were laid out to reduce wire lengths logically distant boards become physically close. This meant that even short cables could create logically long connections, once again allowing small-world networks to form.

## 3.3 High-speed serial (HSS)

Preliminary work has also resulted in the development of foundational hardware components for implementing new connectivity schemes within SpiNNaker's board-to-board interconnect. These components include a crossbar switch [17] which allows spikes to be selectively redirected and interfaces for connecting links operating at different speeds. These components will be built upon during upcoming work.

A proof-of-concept demonstration system was produced which connects SpiNNaker to a host PC using spare HSS connections on SpiNNaker boards. This system replaces an existing Ethernet connection to SpiNNaker (out of the scope of this report) used for system management and interaction with neural models. Though the prototype link operates at a fraction of the maximum speed of an HSS link, it has achieved an order-of-magnitude improvement in host bandwidth for certain neural loads by bypassing bottlenecks in the Ethernet subsystem.

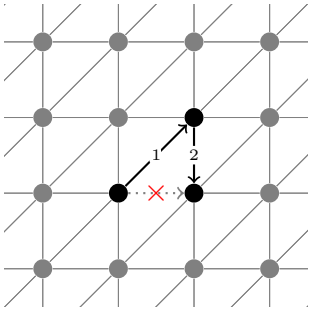


Figure 4: Connectivity between boards (circles) in part of a SpiNNaker system. A link undergoing a speed change is avoided by traversing two neighbouring links.

# 4 Research plan

An outline of planned research is given in figure 3. The major components of this plan are outlined in this section.

## 4.1 Benchmarking

To measure success, a suite of benchmarks is required. A method of automatically converting both models developed by neuroscientists and a number of synthetic tests for the Neural Engineering Framework (NEF) [18] into benchmarks has been proposed. This approach ensures that the benchmarks will represent realistic neural loads while also allowing systematic experimentation and exploration of corner cases.

## 4.2 Small-world

Building on preliminary work, models of small-world configurations in SpiNNaker's network will be modelled and, if found to be effective, implementation will follow. It is hoped that the associated reduction in traffic will allow a global reduction in link speed. The most significant open problem to be solved is the determination of the complexity of routing strategy required to exploit small-world connections. Simpler designs may be insufficient for realistic systems but complex routers may consume more power than they save.

## 4.3 HSS power management

As described earlier, HSS systems present an energy efficiency challenge for neural loads due to their load-independent energy consumption and slow link speed changes. This work focuses on reducing the time HSS links are unavailable during link speed changes making it possible to change link speeds dynamically in response to changes in system load.

Two complementary approaches are being considered: the first involves transmitting data redundantly while a link is changing speed. During speed changes, the link becomes unreliable and existing systems simply wait until the link stabilises at the expense of latency to avoid adding additional hardware for this special case. By using forward

Error Correction (FEC) [19] link errors can be corrected inexpensively using redundant encoding of data sent over the link.

The second approach is to redirect data via an alternative route during link speed changes. This is possible due to redundancy inherent in the toroidal mesh topology of board-to-board links in SpiNNaker; while a link changes speeds, data can be redirected via two neighbouring links as shown in figure 4.

Energy consumption will be measured during the execution of the benchmarks described above using power measurement techniques used in other SpiNNaker power usage studies [20, 21].

# 5 Conclusions

Large scale neural simulations represent a rich opportunity for neuroscientists to understand the brain but progress has been held back by their computational expense. While specialised architectures such as SpiNNaker have significantly advanced the state of the art, their interconnection networks are only beginning to adopt HSS interconnect technologies which enjoy heavy investment from the super computer industry [22].

This work aims to develop a new generation of HSS-based networks for neural simulators which improve both system performance and power efficiency. Preliminary work has developed network modelling tools and proposed a novel technique which adds small-world connectivity to HSS networks, reducing both latency and network load. Planned work will implement this small-world connectivity scheme within SpiNNaker's HSS board-to-board interconnect and complement this with the development of an HSS power management scheme appropriate for neural systems.

# References


[1] Steve Furber and Steve Temple. Neural systems engineering. *Journal of the Royal Society interface*, 4(13):193–206, 2007.

[2] Simon W Moore, Paul J Fox, Steven JT Marsh, A Theodore Markettos, and Alan Mujumdar. Bluehive-a field-programable custom computing machine for extreme-scale real-time neural network simulation. In *Field-Programmable Custom Computing Machines (FCCM), 2012 IEEE 20th Annual International Symposium on*, pages 133–140. IEEE, 2012.

[3] Liam P Maguire, T Martin McGinnity, Brendan Glackin, Arfan Ghani, Ammar Belatreche, and Jim Harkin. Challenges for large-scale implementations of spiking neural networks on FPGAs. *Neurocomputing*, 71(1):13–29, 2007.

[4] Swadesh Choudhary, Steven Sloan, Sam Fok, Alexander Neckar, Eric Trautmann, Peiran Gao, Terry Stewart, Chris Eliasmith, and Kwabena Boahen. Silicon neurons that compute. In *Artificial Neural Networks and Machine Learning–ICANN 2012*, pages 121–128. Springer, 2012.

[5] Janardan Misra and Indranil Saha. Artificial neural networks in hardware: A survey of two decades of progress. *Neurocomputing*, 74(13):239 – 255, 2010.

[6] Chris Eliasmith, Terrence C Stewart, Xuan Choo, Trevor Bekolay, Travis DeWolf, Charlie Tang, and Daniel Rasmussen. A large-scale model of the functioning brain. *Science*, 338(6111):1202–1205, 2012.

[7] Robert W Williams and Karl Herrup. The control of neuron number. *Annual review of neuroscience*, 11(1):423–453, 1988.

[8] S.Y. Kung. Parallel architectures for artificial neural nets. In *Systolic Arrays, 1988., Proceedings of the International Conference on*, pages 163–174, 1988.

[9] J. Navaridas, M. Luján, J. Miguel-Alonso, L.A. Plana, and S. Furber. Understanding the interconnection network of SpiNNaker. In *Proceedings of the 23rd international conference on Supercomputing*, pages 286–295. ACM, 2009.

[10] John Bainbridge and Steve Furber. Chain: A delay-insensitive chip area interconnect. *IEEE Micro*, 22(5):16–23, 2002.

[11] Abhijit Athavale and Carl Christensen. *High-speed serial I/O made simple*. Xilinx Inc, 2005.

[12] InfiniBand Trade Association. *InfiniBand Trade Association Architecture Specification, Volume 1 General Specifications*. Relase 1.2.1 edition, 2008.

[13] Xilinx, Inc. *Spartan-6 FPGA Data Sheet: DC and Switching Characteristics (DS162 v3.1)*, June 2014.

[14] Serial ATA International Organisation. Serial ATA Revision 3.0. `https://www.sata-io.org/`. Accessed: 26 June 2014.

[15] Duncan J Watts and Steven H Strogatz. Collective dynamics of 'small-world' networks. *Nature*, 393(6684):440–442, 1998.

[16] Ji-Yong Shin, Bernard Wong, and Emin Gün Sirer. Small-world datacenters. In *Proceedings of the 2nd ACM Symposium on Cloud Computing*, pages 2:1–2:13. ACM, 2011.

[17] William James Dally and Brian Towles. *Principles and Practices of Interconnection Networks*. Morgan Kaufmann Publishers, 2004.

[18] Chris Eliasmith and Charles H Anderson. *Neural engineering: Computation, representation, and dynamics in neurobiological systems*. MIT Press, 2004.

[19] Richard W Hamming. Error detecting and error correcting codes. *Bell System technical journal*, 29(2):147–160, 1950.

[20] Thomas Sharp, Francesco Galluppi, Alexander Rast, and Steve Furber. Power-efficient simulation of detailed cortical microcircuits on SpiNNaker. *Journal of neuroscience methods*, 210(1):110–118, 2012.

[21] Evangelos Stromatias, Francesco Galluppi, Cameron Patterson, and Steve Furber. Power analysis of large-scale, real-time neural networks on spinnaker. In *Neural Networks (IJCNN), The 2013 International Joint Conference on*, pages 1–8. IEEE, 2013.

[22] InfiniBand Trade Association. InfiniBand Trade Association Homepage. `http://www.infinibandta.org/`. Accessed: 21 July 2013.