MANCHESTER 1824
The University of Manchester

Author
Jonathan Heathcote (jdh@cs.man.ac.uk)

Supervisor
Dr Jim Garside

Co-Supervisor
Prof Steve Furber

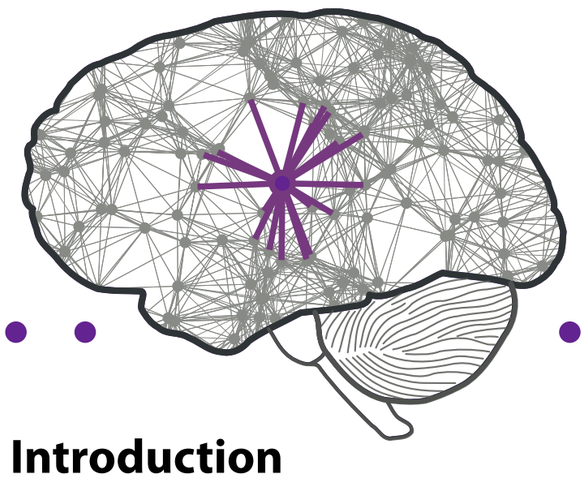

Improving the Interconnection Network of a Brain Simulator

Introduction

The current generation of spiking neural models exhibits a range of basic cognitive functions and realistic biological properties [1]. Neurons in the brain connect to around 10,000 others near-by (illustrated in the brain above) and communicate by sending 'spikes' to their neighbours. The timing of these spikes is the significant factor requiring many tiny, latency-sensitive messages to be sent within a neural simulator's interconnection network.

Despite the network's impact on performance, simulator architectures typically focus attention on novel neuron-modelling hardware [2,3]. This work focuses on improving networks and their topologies.

Practical Network Topologies: A Case Study

The SpiNNaker brain simulator may contain up to 57,600 chips arranged in a "torus" topology (figure 1).

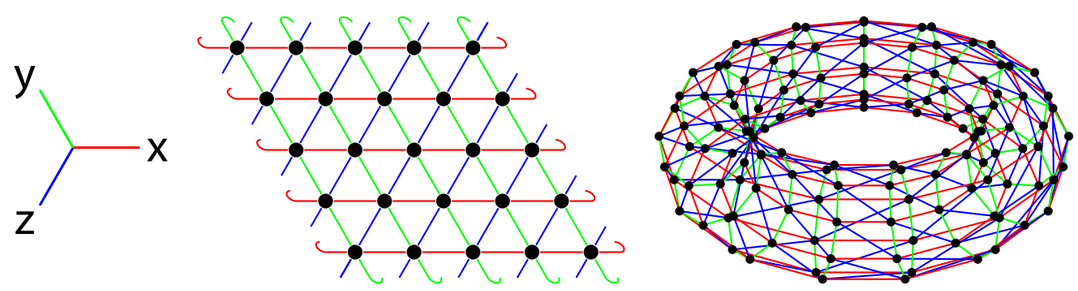

Figure 1: SpiNNaker's topology shown flat and in 3D

Practical systems must be assembled from many circuit boards housed in cabinets (figures 2 & 3) using only physically short wires to minimise latency. A tool was created to design wiring schemes for networks using only short cables and only a few repeating patterns.

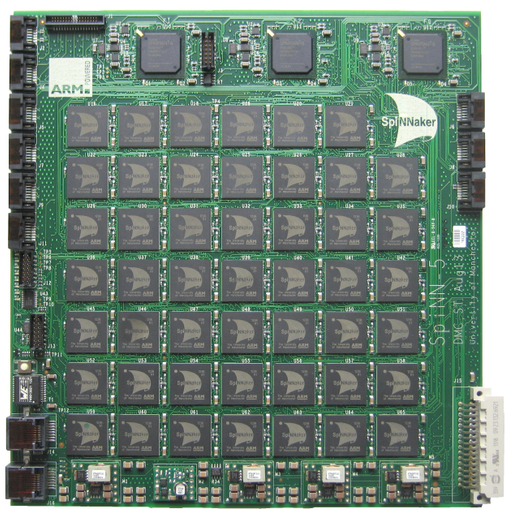

Figure 2: SpiNNaker Board

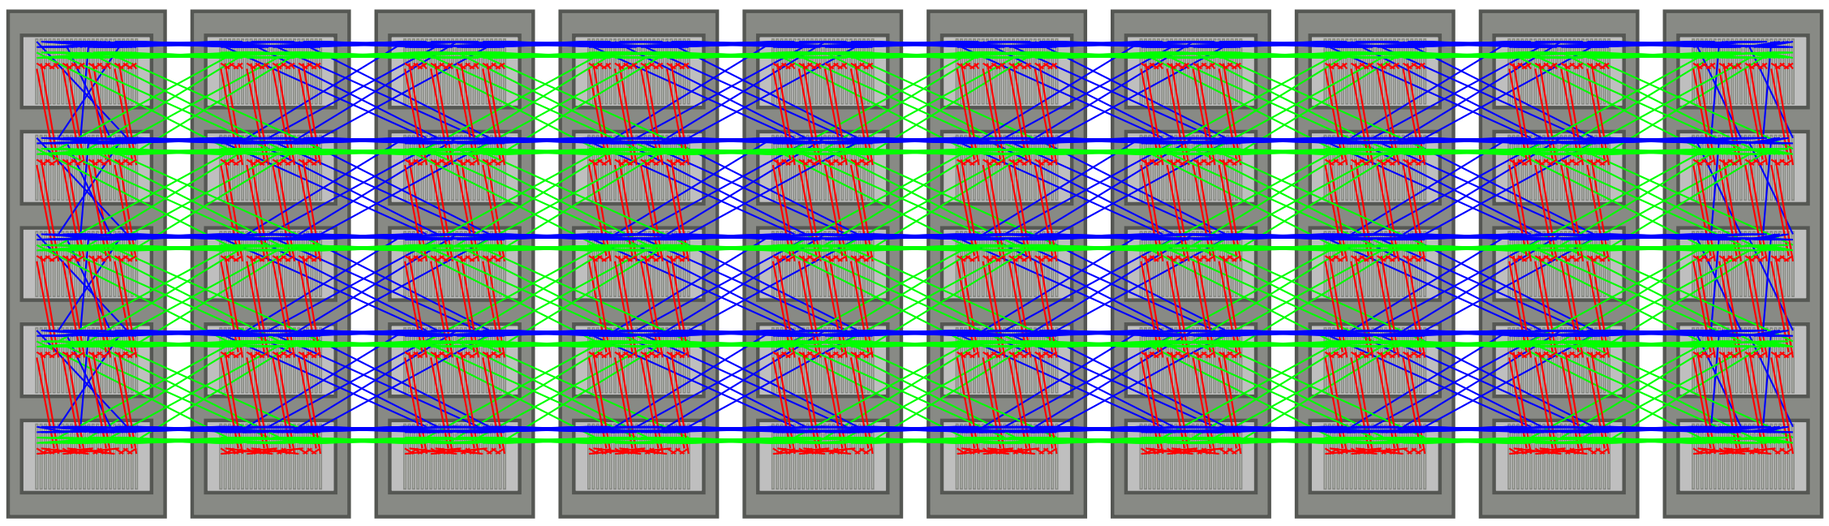

Figure 3: A 6.7m wide.1,200 board system in cabinets arranged such that all wires <1m

Small World Super Computers

Small-world networks, such as social networks, are large and sparsely connected yet the number of links separating any pair of people is small, estimated to be fewer than six [4]. These properties are desirable for interconnection networks and can be achieved by adding random connections to conventional super computer topologies.

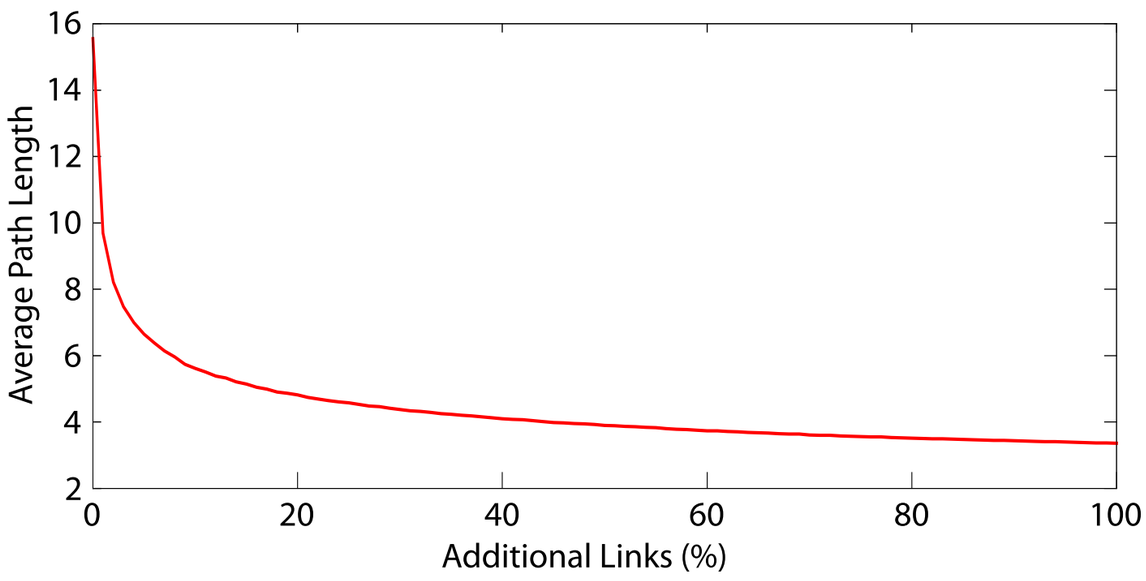

Figure 4: Rewired 40x40 network

Few additional links are required to yield a large drop in average path length and thus to reduce latency (figure 4). Unfortunately, these random connections may require long, and therefore high-latency, cables.

Practical Small World Super Computers

Limiting cable length for random connections in systems laid out naïvely (figure 3) has a negative impact on path length reduction. However, when laid out as in figure 3, this effect is reduced.

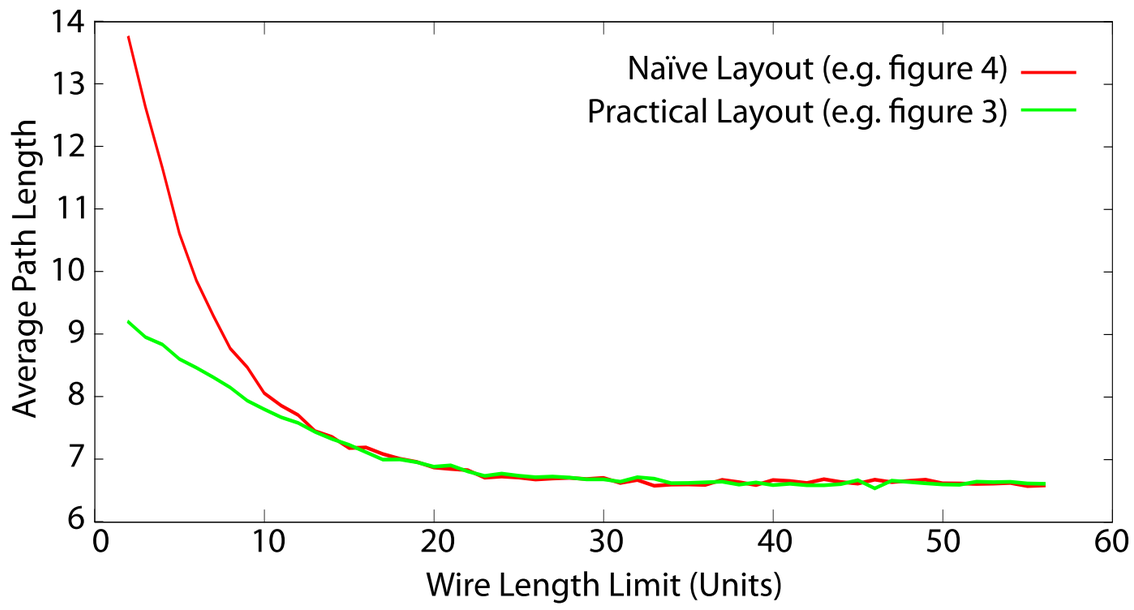

Figure 5: Node positioning effects on wire-length limited wire lengths and 5% additional links

References

[1] 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.

[2] SB Furber, Steve Temple, and AD Brown. High-performance computing for systems of spiking neurons. In *AISB06 workshop on GC5*: Architecture of Brain and Mind, volume 2, pages 29–36, 2006.

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