SpiNNaker: A Multi-Core System-on-Chip for Massively-Parallel Neural Net Simulation

Eustace Painkras, Luis A. Plana, Jim Garside, Steve Temple, Simon Davidson,
Jeffrey Pepper, David Clark, Cameron Patterson and Steve Furber
The University of Manchester, United Kingdom
Email: painkras@cs.man.ac.uk

***Abstract*—The modelling of large systems of spiking neurons is computationally very demanding in terms of processing power and communication. SpiNNaker is a massively-parallel computer system designed to model up to a billion spiking neurons in real time. The basic block of the machine is the SpiNNaker multicore System-on-Chip, a Globally Asynchronous Locally Synchronous (GALS) system with 18 ARM968 processor nodes residing in synchronous islands, surrounded by a light-weight, packet-switched asynchronous communications infrastructure. The MPSoC contains 100 million transistors in a 102 mm^2 die, provides a peak performance of 3.96 GIPS and has a power consumption of 1W at 1.2V when all processor cores operate at nominal frequency. SpiNNaker chips were delivered in May 2011, were fully operational, and met power and performance requirements.**

I. Introduction

SpiNNaker [1] is a biologically-inspired, massively parallel computing architecture designed to facilitate the modelling and simulation of large-scale spiking neural networks of up to a billion neurons and a trillion synapses in biological real-time. It is a general-purpose, programmable platform for neuroscientists, psychologists and brain researchers to explore brain functions with software neuronal models.

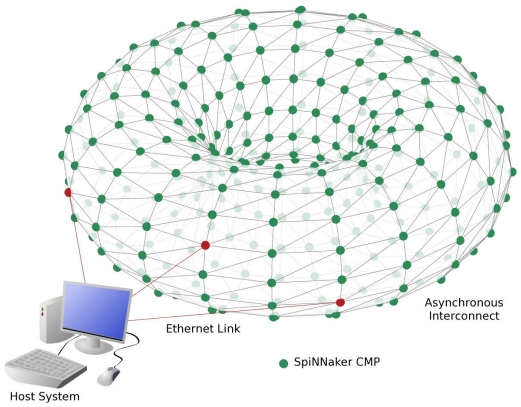

Fig. 1. SpiNNaker Machine

The SpiNNaker machine is designed as a large array of up to 2^{16} nodes, each node containing a Multi-Processor System-on-Chip (MPSoC) die and a 128 MB SDRAM die stacked and stitch-bonded together and housed in a single 300-pin BGA package (see Fig. 4(b)). Each CMP contains eighteen ARM968ES processing cores, each capable of simulating up to a thousand spiking neurons. The architecture scales from a single chip in its smallest configuration to a system of 65,536 chips with 1,179,648 processors in a fully-fledged machine, delivering peak processing power of over 233 Dhrystone TeraIPS. Figure 1 shows the connection of multiple chips to form a SpiNNaker machine and also the manner in which this system connects to the outside world.

The MPSoC design assumes that processors are free: the real cost of computing is energy. This is why we use energy-efficient ARM9 embedded processors and Mobile DDR (Double Data Rate) SDRAM, in both cases sacrificing some performance for greatly enhanced power efficiency. Additionally, inter-chip communication uses self-timed channels, which, although costly in wires, are significantly more power efficient than synchronous links of similar bandwidth.

Inter-processor communication is based on an efficient multicast infrastructure inspired by neurobiology. It uses a packet-switched network to emulate the very high connectivity of biological systems. The packets are source-routed, *i.e.*, they only carry information about the packet issuer and the network is responsible for delivering them to their destinations. The heart of the communications infrastructure is a bespoke multicast router that is able to replicate packets where necessary to implement the multicast function associated with sending the same packet to several different destinations.

The rest of this paper is structured as follows. Related work is presented in Section II. Section III reviews the design of the SpiNNaker MPSoC and its components. The following sections focus on the major design considerations for the SpiNNaker CMP (Section IV) and the experimental results (Section V). Section VI concludes the paper.

II. Related Work

Most large-scale neural simulations [2], [3] utilize supercomputers. The Blue Brain project [2] uses the BlueGene/P supercomputer [4] to simulate cortical columns. The BlueGene/P is not a custom architecture, but a general purpose massively parallel system. Ananthanarayanan et al. [3] also report using the Blue Gene/P machine for cat cortical column simulations as part of DARPA's Systems of Neuromorphic Adaptive Plastic Scalable Electronics (SyNAPSE) program.

The Blue Gene/Q chip, the basic processing element for IBM's latest offering - the Blue Gene/Q massively-parallel scientific computer [5] - employs 18 PowerA2 processor cores with floating-point units occupying a silicon real-estate of

359.5 mm^2 with 1.47 billion transistors fabricated in a 45nm SOI CMOS process. Peak performance for the chip was specified at 204.8 GFLOPS with 55W power dissipation when operated at 1.6GHz with a 0.8V supply.

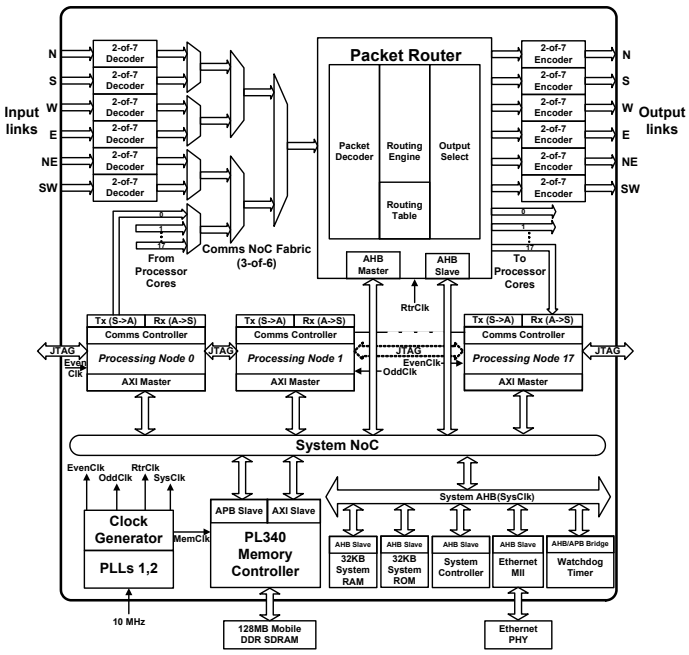

Fig. 2. SpiNNaker chip organization showing the CMP and the SDRAM

Alternatively, *neuromorphic* hardware has been used for large-scale neural net simulations. This type of hardware is extremely energy efficient but implements fixed neuronal and synaptic models, which impose too many restrictions for exhaustive model exploration.

The SpiNNaker MPSoC, with 100 million transistors in a 101.64 mm^2 die, peak performance of 3.96 GIPS and a power consumption of 1W at 1.2V when all cores operate at 180MHz, is a customized architecture which is much more energy efficient than general purpose machines while keeping the flexibility of software-implemented models.

III. The SpiNNaker MPSoC

The basic building block of the SpiNNaker machine is the SpiNNaker chip, shown in Fig. 2. The MPSoC is a GALS multi-processor SoC [6] with 18 ARM968 processor nodes residing in synchronous islands surrounded by a packet-switched asynchronous communications infrastructure. The GALS architecture simplifies timing closure in the SoC design and also facilitates isolation of faulty processor nodes [7]. Self-timed delay-insensitive on-chip interconnects based on CHAIN technology [8] are the backbone for on-chip and off-chip communications.

System-wide communications are handled by two separate hardware communications channels - the *Comms NoC* and the *System NoC*. The *Comms NoC* implements inter-processor communications. The communication is handled through a bespoke multicast router with six full-duplex links connecting to neighbouring chips, forming a 2D toroidal triangular mesh. The *System NoC* provisions the chip-wide sharing of system resources, viz. 32KB System RAM, 32KB System ROM, System Controller, Watchdog Timer and Ethernet interface. It also provides access through a memory controller to the 128 MB off-die SDRAM, private to each CMP but global to its processors. Designated CMPs communicate with the external world through 100-Mbit Ethernet interfaces. The System NoC implements Silistix's custom protocol with AMBA AXI [9] adapters whereas the Comms NoC has 2-of-7 NRZ for off-chip links and 3-of-6 RTZ for on-chip links [10]. As opposed to typical synchronous bus interconnect, the asynchronous NoC provides scalable, high-bandwidth, power-efficient communications.

A. SpiNNaker Processor Node

The SpiNNaker Processor Node is shown in Fig. 3. The processor is the power-efficient, small-footprint, 32-bit ARM968E-S processor designed for data-intensive applications with a Dhrystone performance of 1.1 DMIPS/MHz [11]. Each node has an ARM968 core, private, directly-connected 32KB Instruction Tightly Coupled Memory (ITCM) and dual-banked –for interleaved word access– 64KB Data TCM (DTCM) and peripherals such as a counter/timer, and controllers for vectored interrupts, communications, and direct memory accesses. An IEEE 1149.1-compliant JTAG port is also available for debugging purposes.

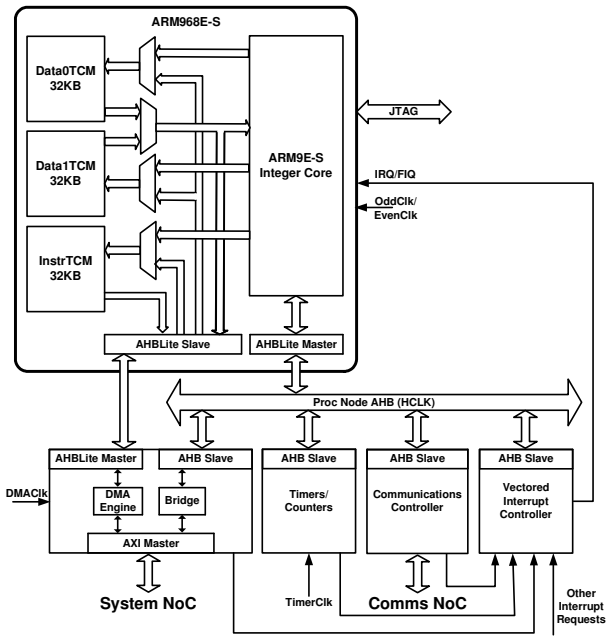

Fig. 3. Details of a SpiNNaker Processor Node

B. Neural Simulation and Communications

At boot time, one of the 18 cores is elected as monitor processor to carry out management and house-keeping tasks and the rest run application software. Each processing core runs independent, event-driven software simulating approximately 1,000 biologically-plausible neurons and their associated synaptic behaviour in real time. The TCMs hold

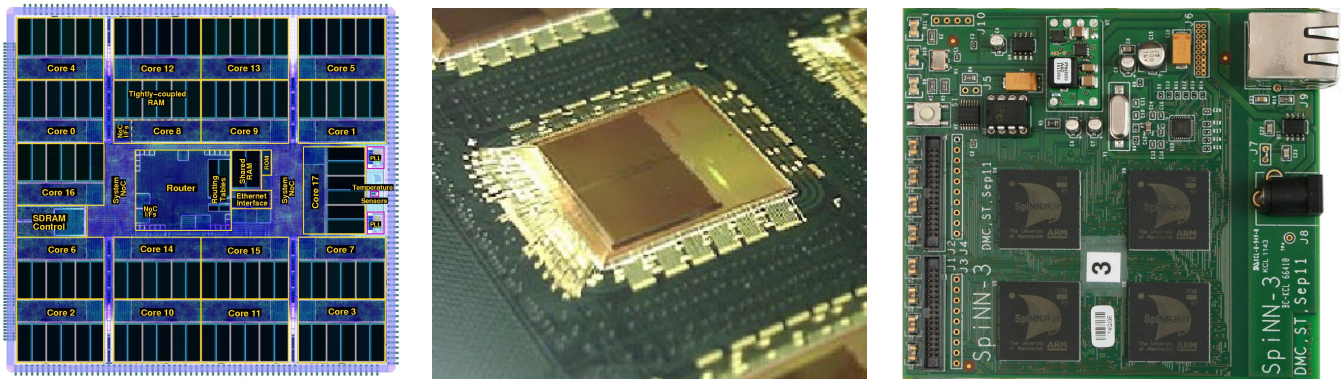

(a) SpiNNaker MPSoC die plot (b) Package substrate with SpiNNaker and SDRAM dies (c) A SpiNNaker test PCB with 4 chips

Fig. 4. SpiNNaker chip and a 3^{rd} generation SpiNNaker PCB

the neural processing model and neural states. The SDRAM contains synaptic connectivity, weights and axonal delays for the neurons simulated by all the application processors in the chip.

The massive, highly-complex interconnectivity and communication channels of the brain are modelled digitally as packet-based spike communication. When a neuron spikes, it generates a packet. The packets, which contain the address of the source neuron, are routed to target neurons distributed around the system using a source-routing algorithm. The arrival of a spike causes the associated synaptic data to be paged in from the SDRAM, using optimized DMA transfers. This can be thought of as a software-managed cache and has to be fairly and efficiently implemented for large-scale simulations to succeed.

IV. Design Considerations

Choice of Process Technology – SpiNNaker was fabricated in UMC 130nm 1.2V 1P8M Fusion process with a combination of Standard Performance (SP) and Low Leakage (LL) standard cell libraries. This choice of libraries resulted in CMPs that consume 70% less leakage power than they would if the High Speed (HS) library alone was used with only a marginal (3%) increase in die area at the desired operating frequency of 180 MHz. The development of SpiNNaker, being of academic nature, is not intended to match the cutting-edge consumer industry, which is shipping products at 32nm and below, as the chip development costs will be exorbitant. Cost considerations have led us to opt for a mature, competitively-priced process technology.

Power Optimization – To minimize the operating costs of the SpiNNaker machine, the main strategy employed is to reduce the power consumption. The SpiNNaker system is built out of energy-efficient embedded processors and mobile DDR SDRAMs. Rather than employing large high-performance power-hungry processing cores, small cores have been chosen for the system, amortizing the area cost across multiple cores. Processor nodes operate at a relatively low frequency of 180MHz. Therefore, they consume much less area and power and it is possible to pack 18 of them in a single MPSoC. The ARM968 processors implement 32-bit fixed-point arithmetic as opposed to the floating-point operations available in general-purpose processors, once again sacrificing performance for energy efficiency. In addition, mechanisms have been built-in to power-off idle nodes of the machine, putting the processors into *sleep* mode when they are not used for computation and *wake* them up when the need arises.

Balancing Area, Performance, and Power – The implementation employs architecture and logic-level clock gating. The design methodology has been fine-tuned so that the processor cores consume low power with special emphasis on the power-efficiency of the clock networks. Power-aware synthesis was used throughout the flow. The ARM processor cores were brought in as IP, their physical implementation fine-tuned through numerous iterations to achieve the most compact, power-efficient and best achievable performance design using the ARM-Synopsys Galaxy reference design methodology. Other blocks such as the Memory Controller, Router and System AHB were also implemented adopting the same design methodology. Special mention is to be made of the implementation of the asynchronous *System* and *Comms NoCs*, which is detailed in [7] but the rest of the chip was implemented as an SoC.

Monitoring and Fault Tolerance – A hallmark of the SpiNNaker system design is its built-in fault tolerance features. 18 processor cores, 6 I/O links and 2 PLLs engender redundancy in the main components of the MPSoC. On-chip temperature sensors have been placed in every CMP to track the thermal state during operation. The System Controller (Fig. 2) has built-in fault-monitoring mechanisms to gauge and report the health status of every core in the CMP. The monitored data is transported along with the neural application packets without impacting the application data traffic, thereby obviating the need for separate interconnect resources to communicate the monitored data. Additionally, the system software is capable of doing run-time diagnostics and reconfiguration with the assistance of hardware units such as neighbouring chips using the router as a portal into the resources. The Comms NoC can detect parity and framing errors while the DMA controller has built-in CRC encoder-decoder for communications through the System NoC. Finally, a hardware-implemented emergency routing is designed into the system

to cope with failures in the inter-chip communication links.

Clocking – The SpiNNaker CMP has two independent PLLs providing clock inputs to processor nodes and other system components. The processor nodes and router operate at a nominal frequency of 180MHz, the shared system resources at 100MHz and the SDRAM at 166MHz. In order to maintain uniform clock tree power consumption throughout the entire clock period, the clock nets have been purposely skewed evenly across the processing nodes. This deliberate introduction of skew reduces the peak supply current. However, the phase differences in the arrival of the clock inputs at the different cores do not have any impact on the operation of the CMP due to the GALS design methodology.

V. Results

The SpiNNaker MPSoC was designed by a small team of academic researchers and post-graduate students with the associated restrictions and constraints regarding fabrication cost, access to process technologies, standard cell libraries and IPs. Overall, a design effort of approximately 40 person-years has gone into the design, implementation and verification of the SpiNNaker CMP. A test chip with 2 cores was taped-out in August 2009 followed by the 18-core chip in December 2010. The SpiNNaker die area is 102 mm^2 (10.386 mm $\times$ 9.786 mm). The first batch of fully-functional packaged chips was delivered on May 20th, 2011. The MPSoC contains 100 million transistors, provides a peak performance of 3.96 GIPS and a peak power consumption of 1W at 1.2V when all the processor cores are operating at 180MHz.

As mentioned earlier, the optimal utilization of the SDRAM bandwidth is crucial to the performance of spiking neural simulations. Figure 5 illustrates the results of an experiment to evaluate this key indicator. The figure shows the aggregate SDRAM bandwidth utilization of a number of cores. The number of cores is increased progressively from 1 to 14. The experiment is based on DMA accesses to memory, with the first 7 cores reading from memory and the last 7 doing write operations. It is clear from the figure that the read channel from memory saturates at just over 600 MBytes/s and the write channel adds around 300 MBytes/s on top of that, for a total aggregate bandwidth of 900 MBytes/s. The figure also shows that, as expected, the simple mechanism of restricting each core to a single outstanding command results in a fair share of the bandwidth for all cores.

VI. Conclusions

SpiNNaker is a massively-parallel computer system designed specifically for large-scale neural net simulations, a computation- and communication-intensive application. The main design and implementation strategy used to minimize SpiNNaker operating costs was to reduce power consumption. Amongst many power-saving measures, it is built out of energy-efficient embedded processors and mobile DDR SDRAMs.

Experimental results show that, for massively-parallel neural net simulations, a customized multi-core architecture can be more energy-efficient than general purpose machines while keeping the flexibility of software-implemented neuronal and synaptic models, absent in current neuromorphic hardware.

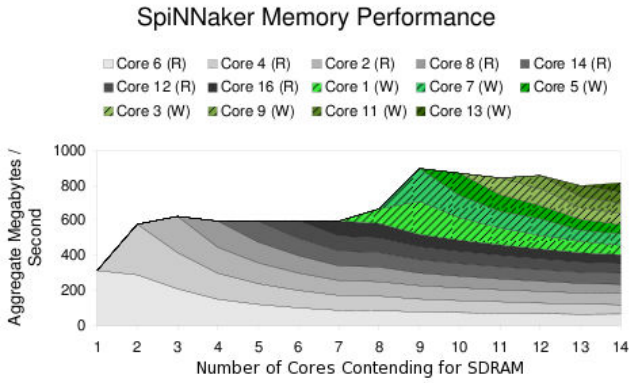

Fig. 5. SDRAM Bandwidth Utilization

Though SpiNNaker is an application-specific architecture, it can be used as a general-purpose machine to run parallel applications such as ray-tracing and protein folding, which are outside the purview of neuroscience applications. For the neural applications, a major advantage is gained in the flexibility afforded by the software implementation of neural models in the processor cores and the efficiency of the bio-inspired asynchronous communications infrastructure.

Acknowledgements

The SpiNNaker project is supported by the Engineering and Physical Sciences Research Council of the UK, through Grants EP/D07908X/1 and EP/G015740/1, and also by ARM and Silistix. The die photo in Fig. 4(b) is courtesy of Unisem Europe Ltd.

References

[1] S. Furber and S. Temple, "Neural Systems Engineering," *J. Roy. Soc. Interface*, vol. 4, no. 13, pp. 193–206, Apr. 2007.

[2] H. Markram, "The Blue Brain Project," *Nature Reviews Neuroscience*, vol. 7, no. 2, pp. 153–160, Feb. 2006.

[3] R. Ananthanarayanan *et al.*, "The Cat is Out of the Bag: Cortical Simulations with 10^9 Neurons, 10^{13} Synapses," in *Proc. Conf. High Performance Computing Networking, Storage and Analysis*, 2009, pp. 63:1–63:12.

[4] IBM Blue Gene Team, "Overview of the IBM Blue Gene/P Project," *IBM J. Research & Development*, vol. 52, no. 1/2, pp. 199–220, 2008.

[5] R. Haring *et al.*, "The IBM Blue Gene/Q Compute Chip," *Micro, IEEE*, vol. 32, no. 2, pp. 48–60, march-april 2012.

[6] L. Plana *et al.*, "A GALS Infrastructure for a Massively Parallel Multiprocessor," *IEEE D&T Computers*, vol. 24, no. 5, pp. 454–463, 2007.

[7] ——, "SpiNNaker: Design and Implementation of a GALS Multicore System-on-Chip," *ACM J. Emerging Technologies in Computing Systems*, vol. 7, no. 4, pp. 17:1–17:18, Dec. 2011.

[8] J. Bainbridge and S. Furber, "CHAIN: A Delay-Insensitive Chip Area Interconnect," *IEEE Micro*, vol. 22, no. 5, pp. 16–23, 2002.

[9] ARM, "AMBA AXI Protocol Specification, Rev.1.0," *http://www.arm.com/products/system-ip/amba/*, 2004.

[10] Y. Shi *et al.*, "Fault Tolerant Delay Insensitive Inter-chip Communication," in *Proceedings of the 2009 15th IEEE Symposium on Asynchronous Circuits and Systems*, ser. Async 2009, 2009, pp. 77–84.

[11] ARM, "ARM968E-S," *http://www.arm.com/products/processors/classic/arm9/arm968.php*, 2011.