SpiNNaker
Biologically Inspired Massively Parallel Architectures

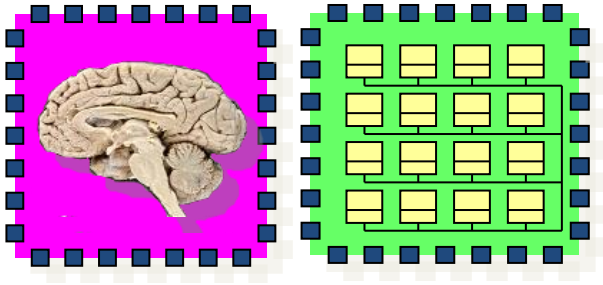

Self-Tuning Bio-Inspired Massively-Parallel Computing

Steve Furber
The University of Manchester
steve.furber@manchester.ac.uk

EPSRC UNIVERSITY OF CAMBRIDGE MANCHESTER 1824 The University Of Sheffield. UNIVERSITY OF Southampton ARM Silistix THALES

SpiNNaker
Biologically Inspired Massively Parallel Architectures

Outline

- 63 years of progress
- Many cores make light work
- Building brains
- The *SpiNNaker* project
- The networking challenge
- A generic neural modelling platform
- Plans & conclusions

SpiNNaker
Biologically Inspired Massively Parallel Architectures

Manchester Baby (1948)

SpiNNaker
Biologically Inspired Massively Parallel Architectures

SpiNNaker CPU (2011)

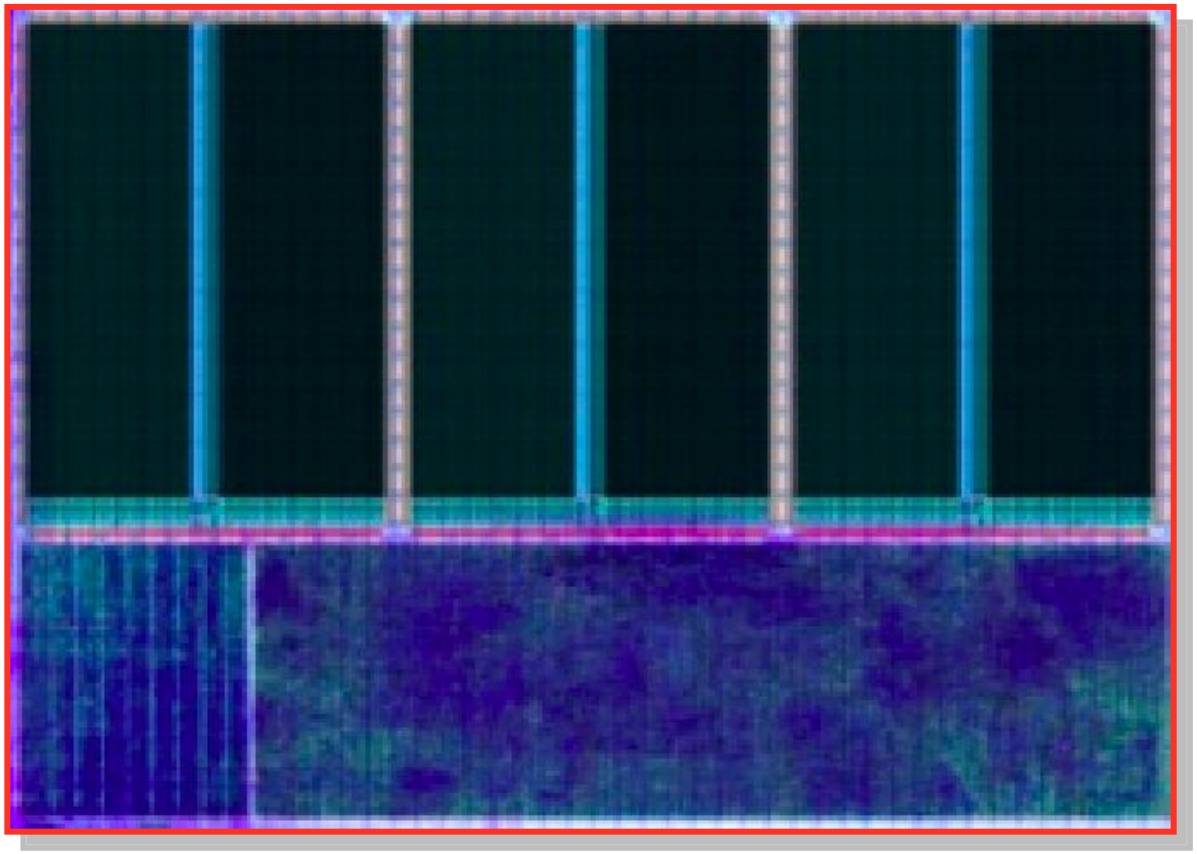

SpiNNaker
Biologically Inspired Massively Parallel Architectures

63 years of progress

- ***Baby:***
  - filled a medium-sized room
  - used 3.5 kW of electrical power
  - executed 700 instructions per second

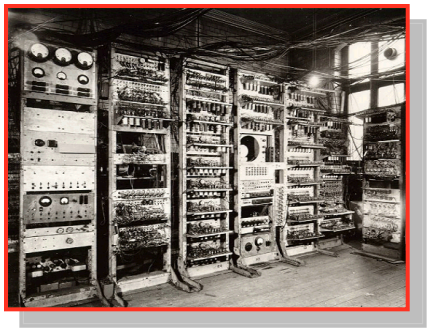

- ***SpiNNaker ARM968 CPU node:***
  - fills ~3.5mm^2 of silicon (130nm)
  - uses 40 mW of electrical power
  - executes 200,000,000 instructions per second

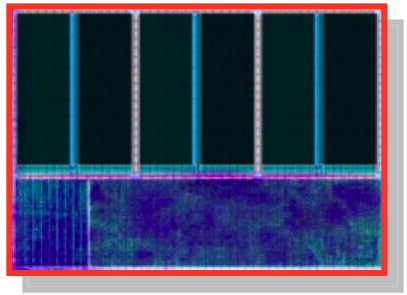

EPSRC UNIVERSITY OF CAMBRIDGE MANCHESTER 1824 The University Of Sheffield. UNIVERSITY OF Southampton ARM Silistix THALES

SpiNNaker
Biologically Inspired Massively Parallel Architectures

Energy efficiency

- Baby:
  - 5 Joules per instruction
- SpiNNaker ARM968:
  - 0.000 000 000 2 Joules per instruction

25,000,000,000 times better than Baby!

(James Prescott Joule born Salford, 1818)

SpiNNaker
Biologically Inspired Massively Parallel Architectures

Moore's Law

Transistors per Intel chip

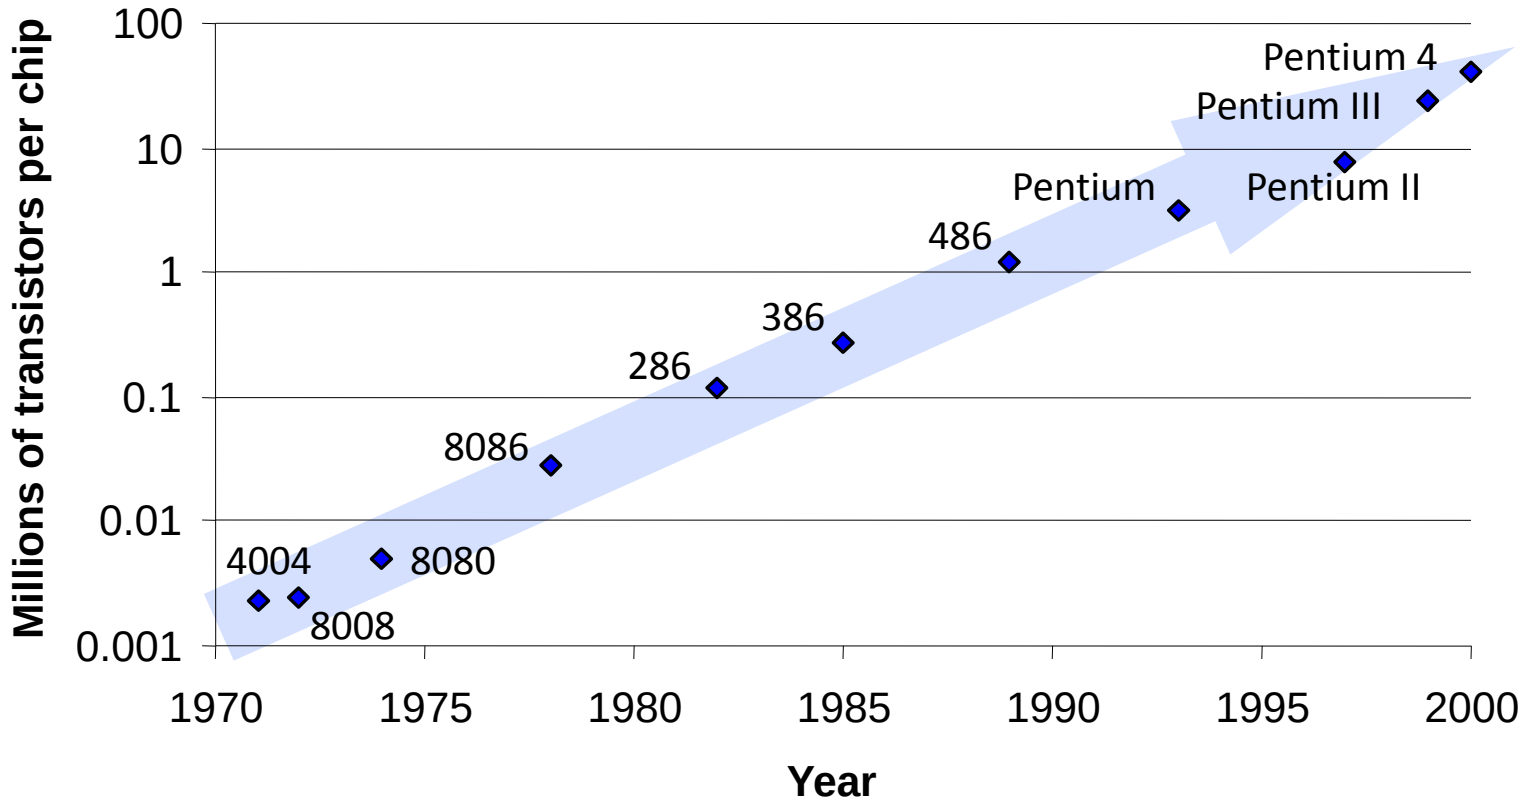

EPSRC UNIVERSITY OF CAMBRIDGE MANCHESTER 1824 The University Of Sheffield. UNIVERSITY OF Southampton ARM Silistix THALES

SpiNNaker
Biologically Inspired Massively Parallel Architectures

...the Bad News

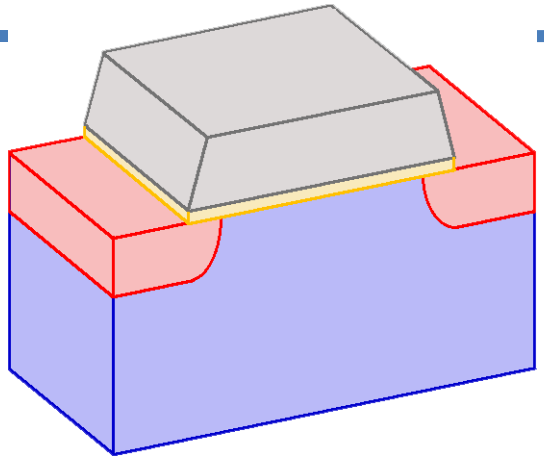

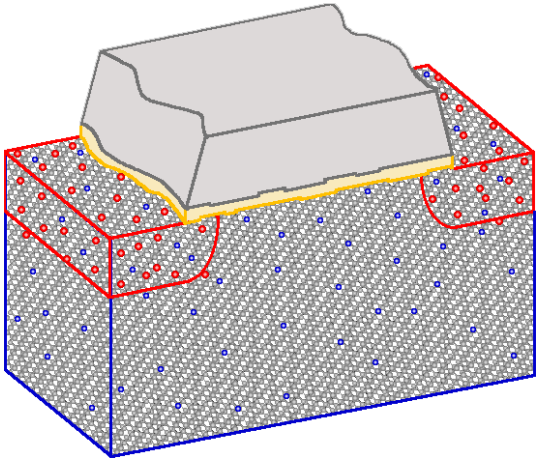

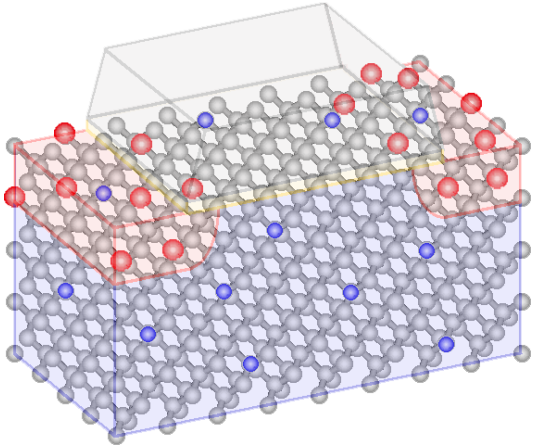

DEVICE MODELLING GROUP

UNIVERSITY of GLASGOW

- atomic scales
  - less predictable
  - less reliable

SpiNNaker
Biologically Inspired Massively Parallel Architectures

Outline

- 63 years of progress
- Many cores make light work
- Building brains
- The *SpiNNaker* project
- The networking challenge
- A generic neural modelling platform
- Plans & conclusions

SpiNNaker
Biologically Inspired Massively Parallel Architectures

Multi-core CPUs

- High-end uniprocessors
  - diminishing returns from complexity
  - wire vs transistor delays
- Multi-core processors
  - cut-and-paste
  - *simple* way to deliver more MIPS
- Moore's Law
  - more transistors
  - more cores

… but what about the software?

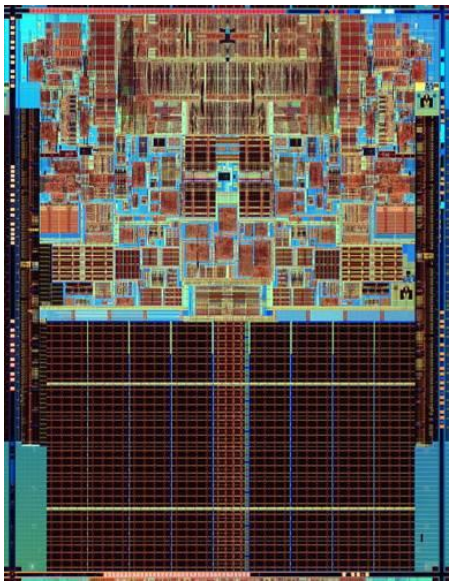

SpiNNaker
Biologically Inspired Massively Parallel Architectures

Multi-core CPUS

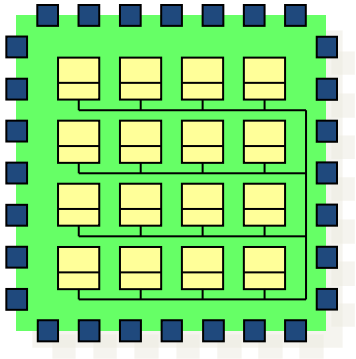

- General-purpose parallelization
  - an unsolved problem
  - the 'Holy Grail' of computer science for half a century?
  - but imperative in the many-core world
- Once solved
  - few complex cores, or many simple cores?
  - simple cores win hands-down on power-efficiency!

SpiNNaker
Biologically Inspired Massively Parallel Architectures

Back to the future

- Imagine…
  - a limitless supply of (free) processors
  - load-balancing is irrelevant
  - all that matters is:
    - the energy used to perform a computation
    - formulating the problem to avoid synchronisation
    - abandoning determinism
- How might such systems work?

SpiNNaker
Biologically Inspired Massively Parallel Architectures

Outline

- 63 years of progress
- Many cores make light work
- Building brains
- The *SpiNNaker* project
- The networking challenge
- A generic neural modelling platform
- Plans & conclusions

SpiNNaker
Biologically Inspired Massively Parallel Architectures

Building brains

- Brains demonstrate
  - massive parallelism (10^{11} neurons)
  - massive connectivity (10^{15} synapses)
  - excellent power-efficiency
    - much better than today's microchips
  - low-performance components (~ 100 Hz)
  - low-speed communication (~ metres/sec)
  - adaptivity – tolerant of component failure
  - autonomous learning

SpiNNaker
Biologically Inspired Massively Parallel Architectures

Bio-inspiration

- How can massively parallel computing resources accelerate our understanding of brain function?
- How can our growing understanding of brain function point the way to more efficient parallel, fault-tolerant computation?

SpiNNaker
Biologically Inspired Massively Parallel Architectures

Building brains

- Neurons
  - multiple inputs, single output (c.f. logic gate)
  - useful across multiple scales (10^2 to 10^{11})
- Brain structure
  - regularity
  - e.g. 6-layer cortical 'microarchitecture'

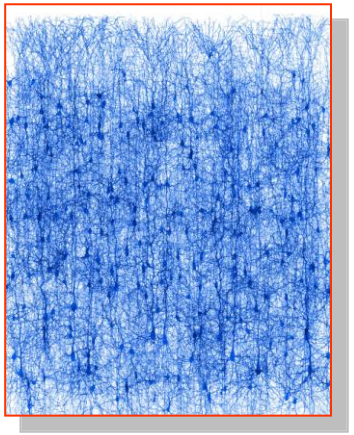

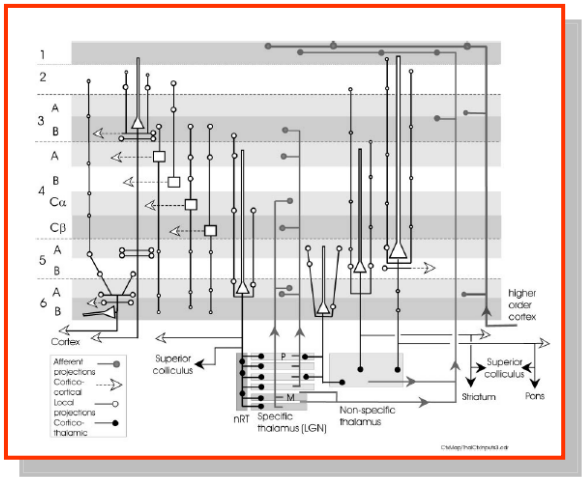

EPSRC UNIVERSITY OF CAMBRIDGE MANCHESTER 1824 The University Of Sheffield. UNIVERSITY OF Southampton ARM Silistix THALES

Spike Timing Dependent Plasticity

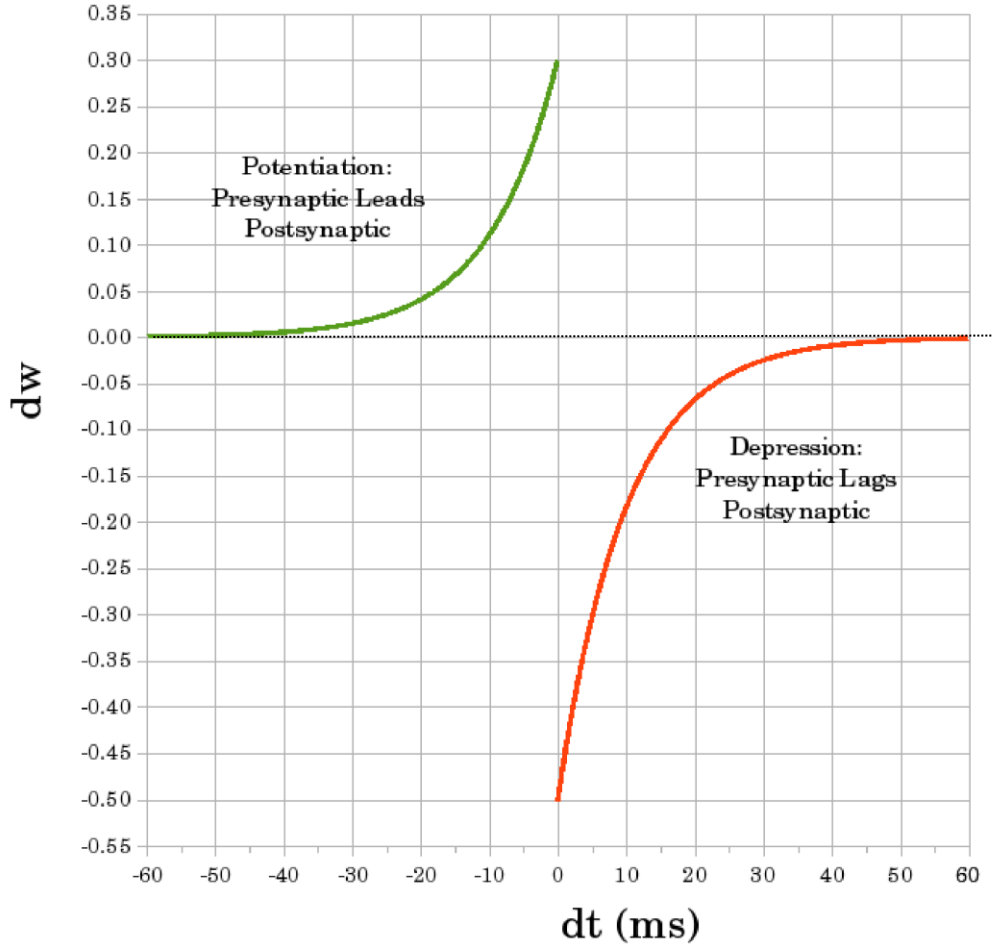

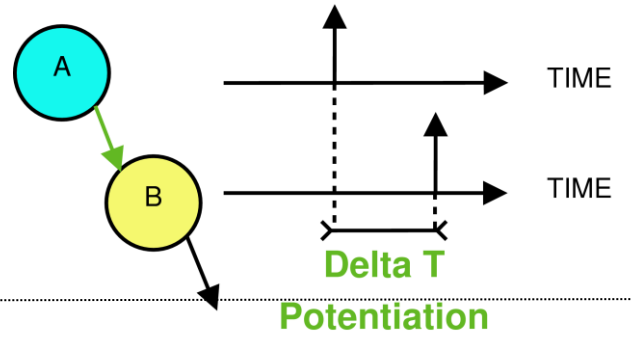

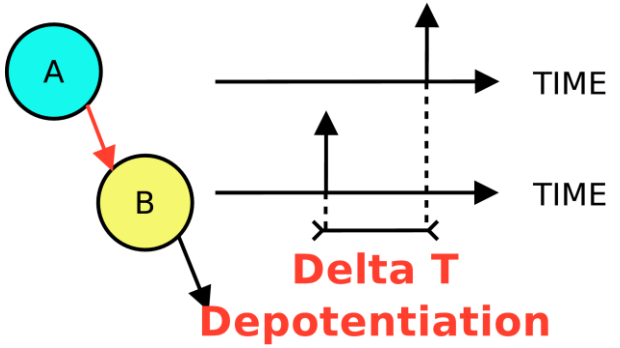

SpiNNaker
Biologically Inspired Massively Parallel Architectures

Learning patterns

- Spot the pattern?

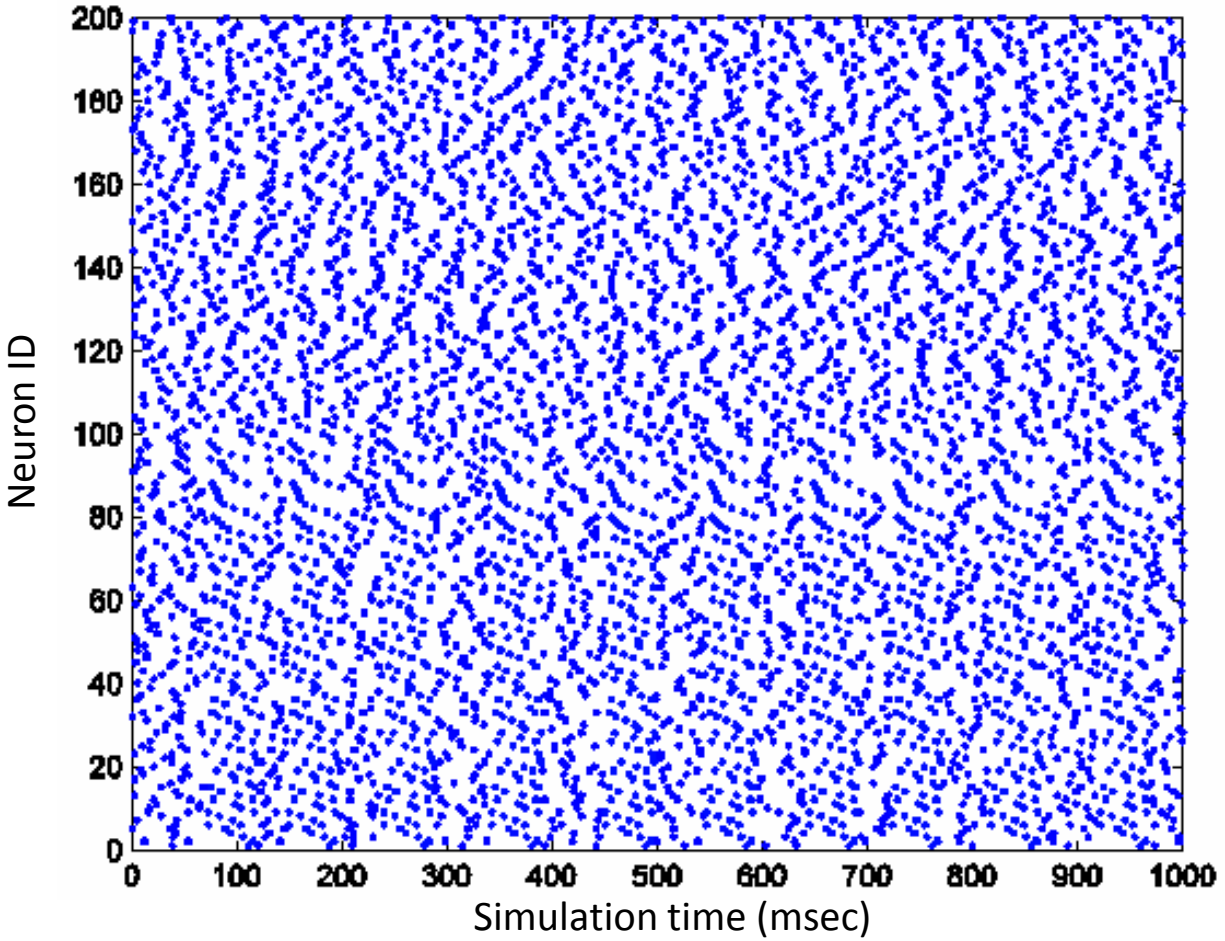

EPSRC UNIVERSITY OF CAMBRIDGE MANCHESTER 1824 The University Of Sheffield. UNIVERSITY OF Southampton ARM Silistix THALES

SpiNNaker
Biologically Inspired Massively Parallel Architectures

Learning patterns

- Now you see it!

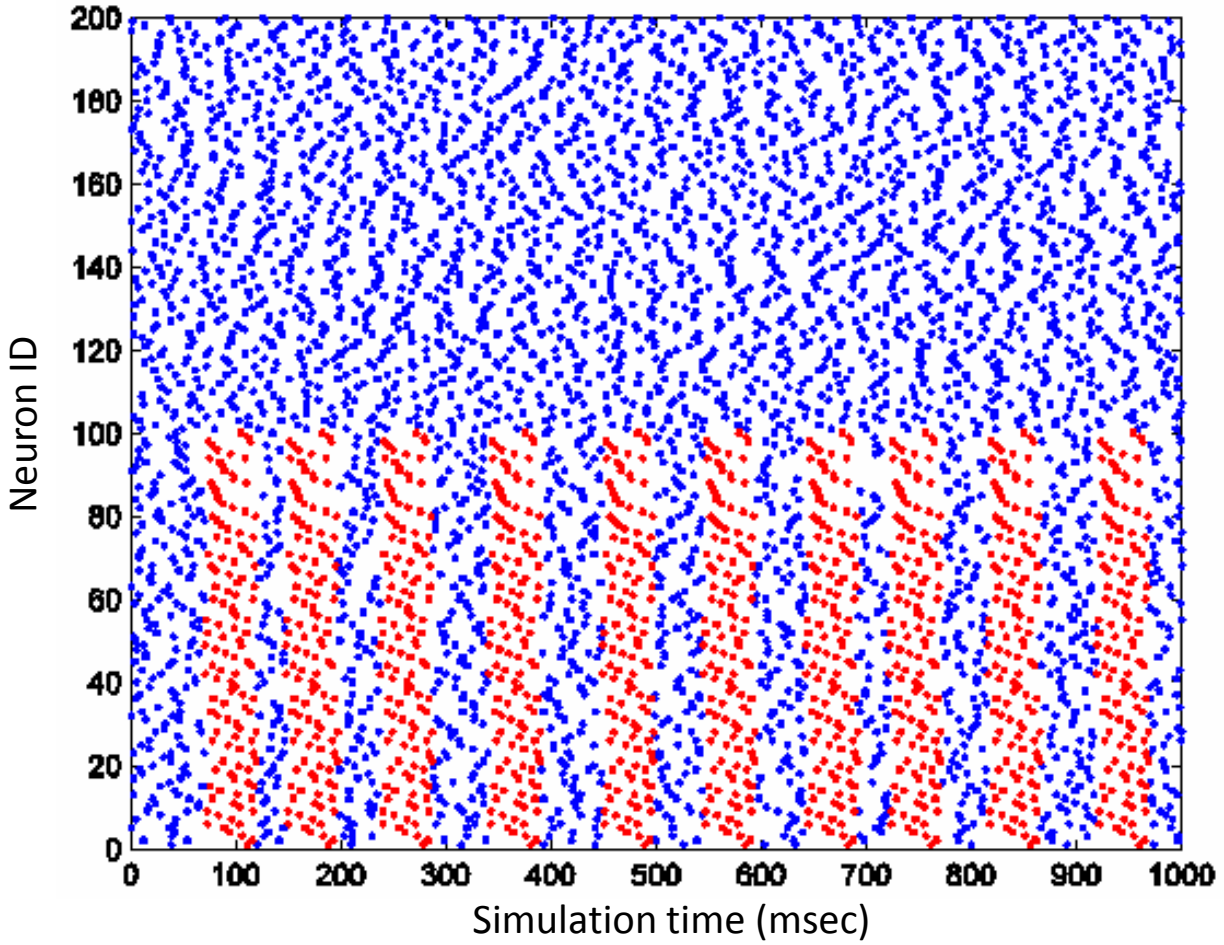

EPSRC UNIVERSITY OF CAMBRIDGE MANCHESTER 1824 The University Of Sheffield. UNIVERSITY OF Southampton ARM Silistix THALES

SpiNNaker
Biologically
Inspired
Massively
Parallel
Architectures

Learning patterns

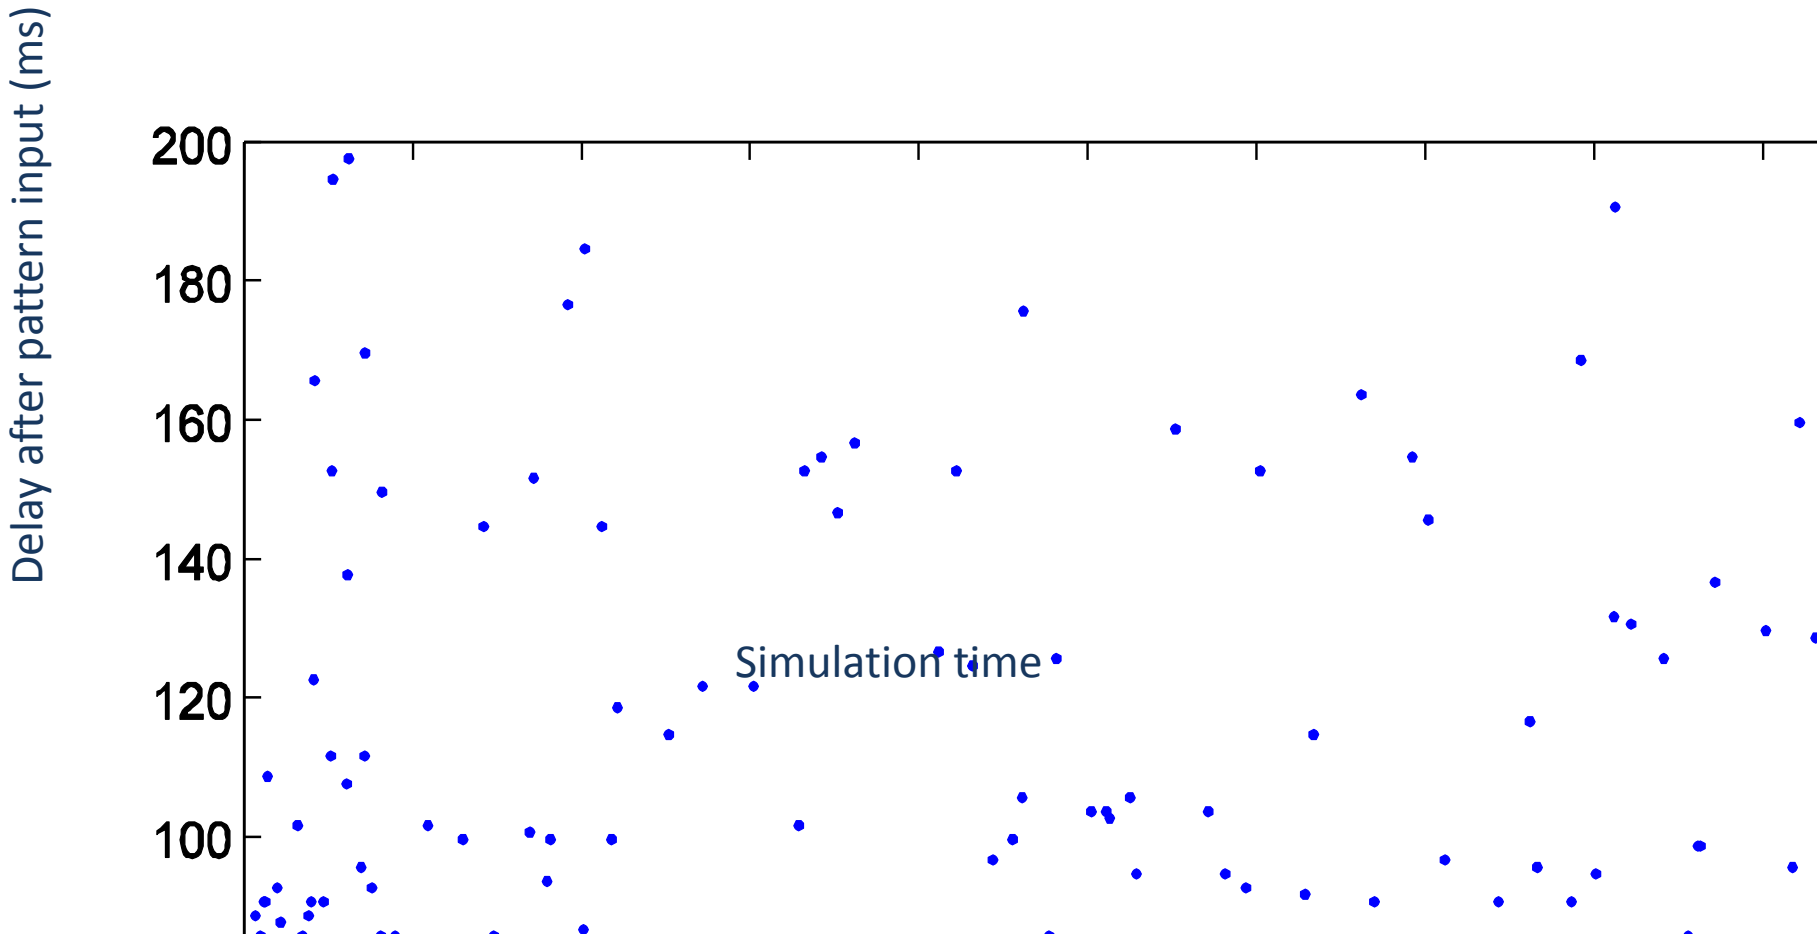

SpiNNaker
Biologically Inspired Massively Parallel Architectures

Self-tuning: in brains

- With STDP, and no other re-inforcement
  - neurons learn the statistics of their inputs
- and, with just a little mutual inhibition
  - populations distribute themselves across the range of presented inputs.
- New inputs are interpreted against these learnt statistics.
  - Bayes would be very proud!

Masquelier & Thorpe, 2007

SpiNNaker
Biologically Inspired Massively Parallel Architectures

Outline

- 63 years of progress
- Many cores make light work
- Building brains
- The *SpiNNaker* project
- The networking challenge
- A generic neural modelling platform
- Plans & conclusions

SpiNNaker
Biologically Inspired Massively Parallel Architectures

SpiNNaker project

- Multi-core CPU node
  - 18 ARM968 processors
  - to model large-scale systems of spiking neurons
- Scalable up to systems with 10,000s of nodes
  - over a million processors
  - >10^8 MIPS total

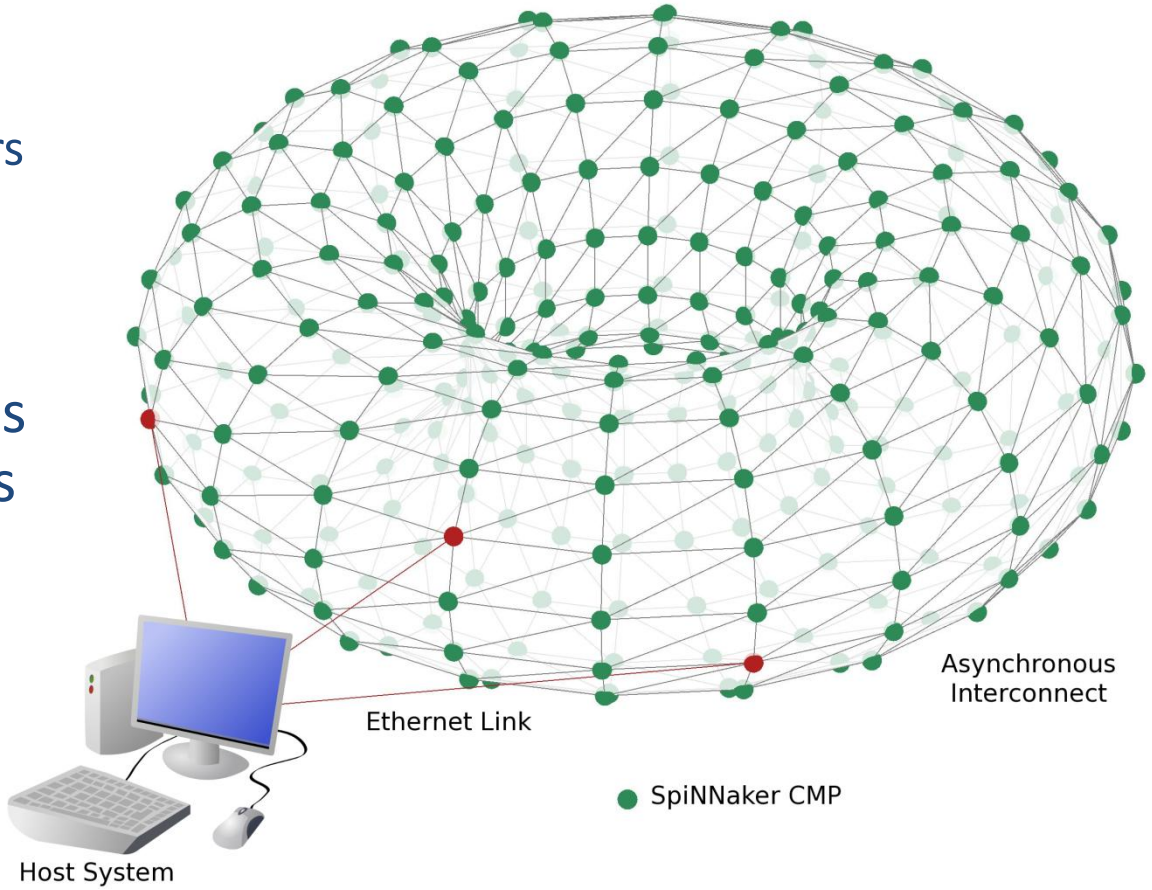

SpiNNaker
Biologically Inspired Massively Parallel Architectures

Design principles

- *Virtualised topology*
  - physical and logical connectivity are decoupled
- *Bounded asynchrony*
  - time models itself
- *Energy frugality*
  - processors are free
  - the real cost of computation is energy

SpiNNaker
Biologically
Inspired
Massively
Parallel
Architectures

SpiNNaker system

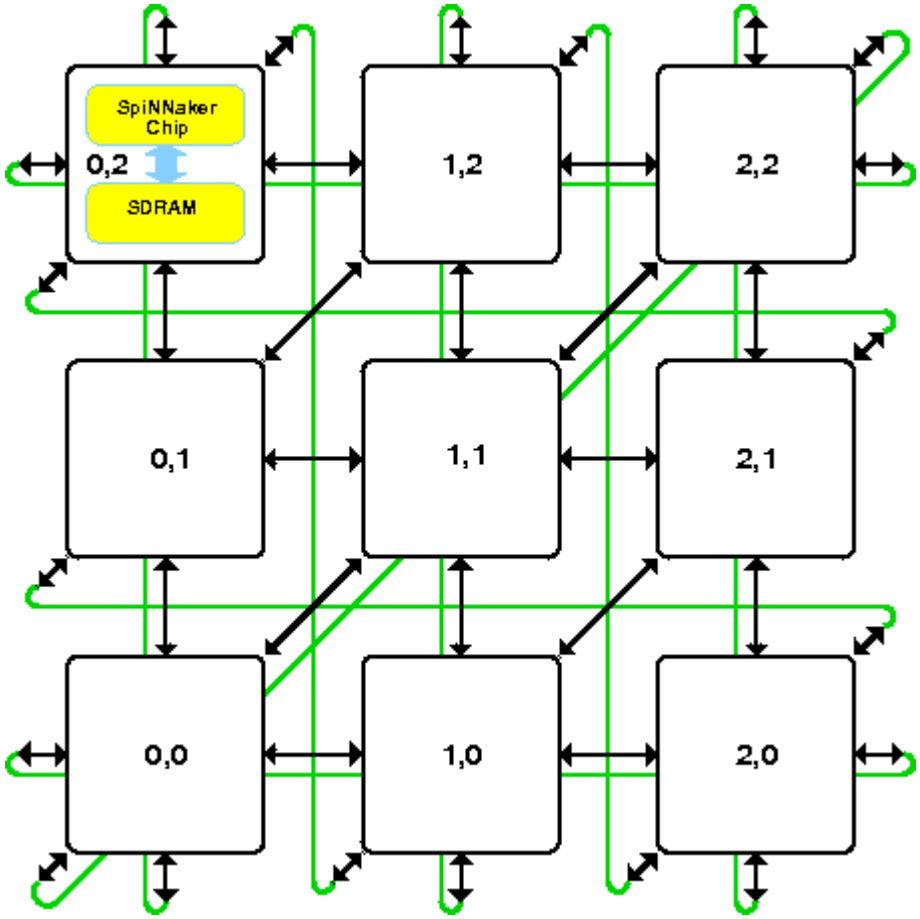

EPSRC UNIVERSITY OF CAMBRIDGE MANCHESTER 1824 The University Of Sheffield. UNIVERSITY OF Southampton ARM Silistix THALES

SpiNNaker
Biologically Inspired Massively Parallel Architectures

CMP node

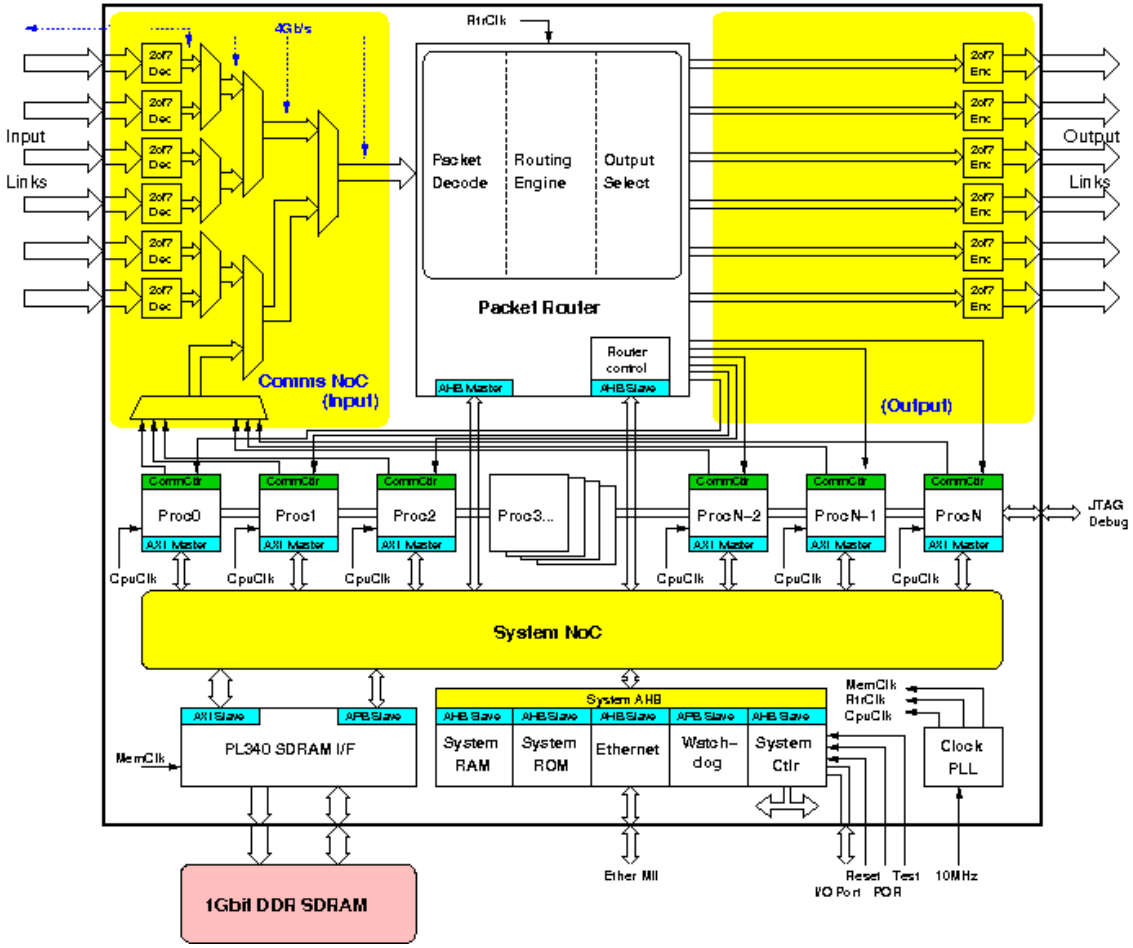

EPSRC UNIVERSITY OF CAMBRIDGE MANCHESTER 1824 The University Of Sheffield. UNIVERSITY OF Southampton ARM Silistix THALES

SpiNNaker
Biologically Inspired Massively Parallel Architectures

SpiNNaker chip

Mobile DDR SDRAM interface

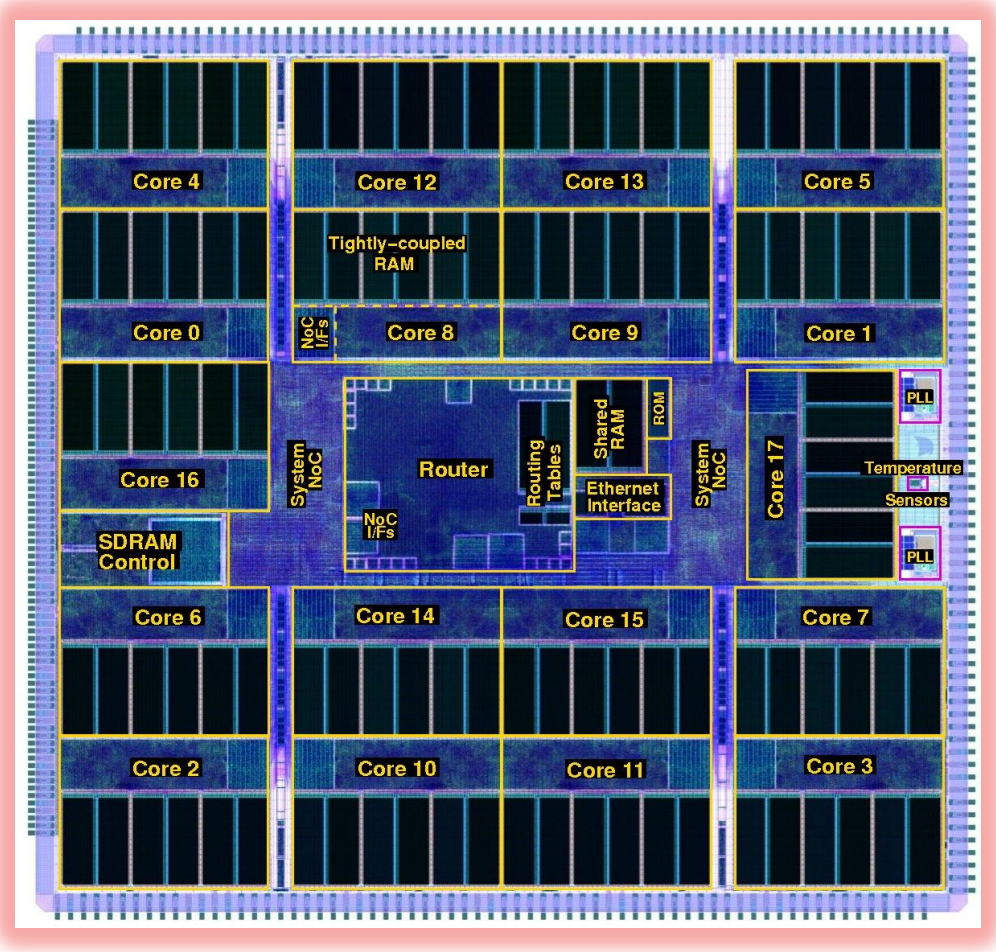

EPSRC UNIVERSITY OF CAMBRIDGE MANCHESTER 1824 The University Of Sheffield. UNIVERSITY OF Southampton ARM Silistix THALES

SpiNNaker
Biologically Inspired Massively Parallel Architectures

SpiNNaker SiP

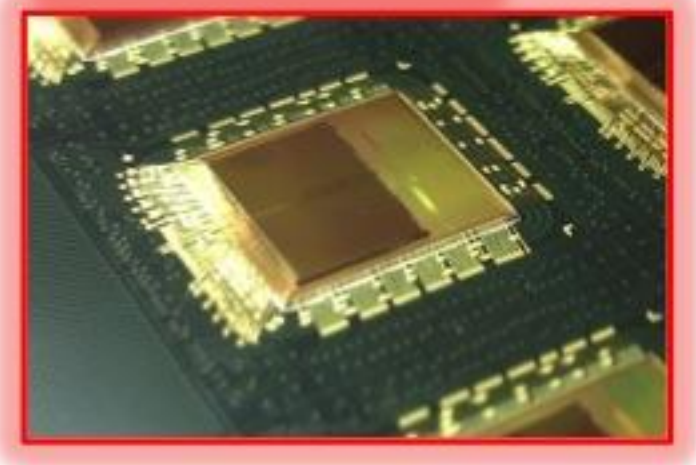

Multi-chip packaging by UNISEM Europe

EPSRC UNIVERSITY OF CAMBRIDGE MANCHESTER 1824 The University Of Sheffield. UNIVERSITY OF Southampton ARM Silistix THALES

SpiNNaker
Biologically Inspired Massively Parallel Architectures

Self-tuning: fault-tolerance

- Strategy: for all components consider:
  - fault insertion – how do we test the FT feature?
  - fault detection – we have a problem!
  - fault isolation – contain the damage
  - reconfiguration – repair the damage
- Goal: minimize performance deficit x time
  - real-time system, so checkpoint & restart inapplicable

SpiNNaker
Biologically Inspired Massively Parallel Architectures

Circuit-level fault-tolerance

- Delay-insensitive comms
  - 3-of-6 RTZ on chip
  - 2-of-7 NRZ off chip
- Deadlock resistance
  - Tx & Rx circuits have high deadlock immunity
  - Tx & Rx can be reset independently
    - each injects a token at reset
    - true transition detector filters surplus token

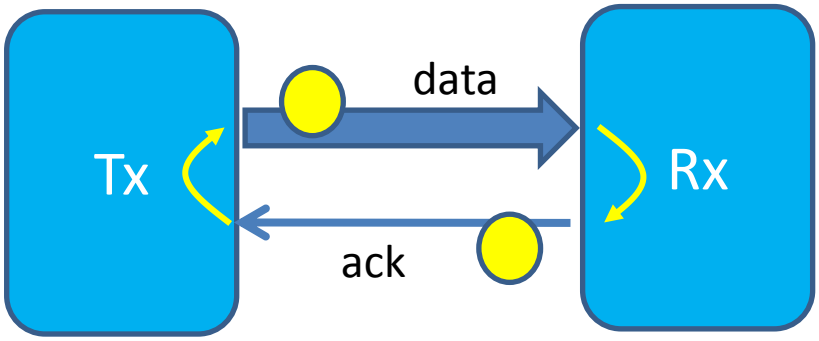

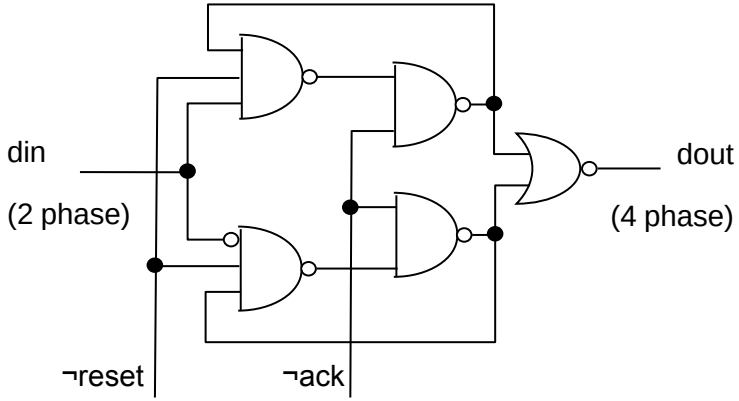

SpiNNaker
Biologically Inspired Massively Parallel Architectures

System-level fault-tolerance

- Breaking symmetry
  - any processor can be Monitor Processor
    - local 'election' on each chip, after self-test
  - all nodes are identical at start-up
    - addresses are computed relative to node with host connection (0,0)
  - system initialised using flood-fill
    - nearest-neighbour packet type
    - boot time (almost) independent of system scale

SpiNNaker
Biologically Inspired Massively Parallel Architectures

Application-level fault-tolerance

- Cross-system delay << 1ms
  - hardware routing
  - 'emergency' routing
    - failed links
    - congestion
  - permanent fault
    - reroute (s/w)

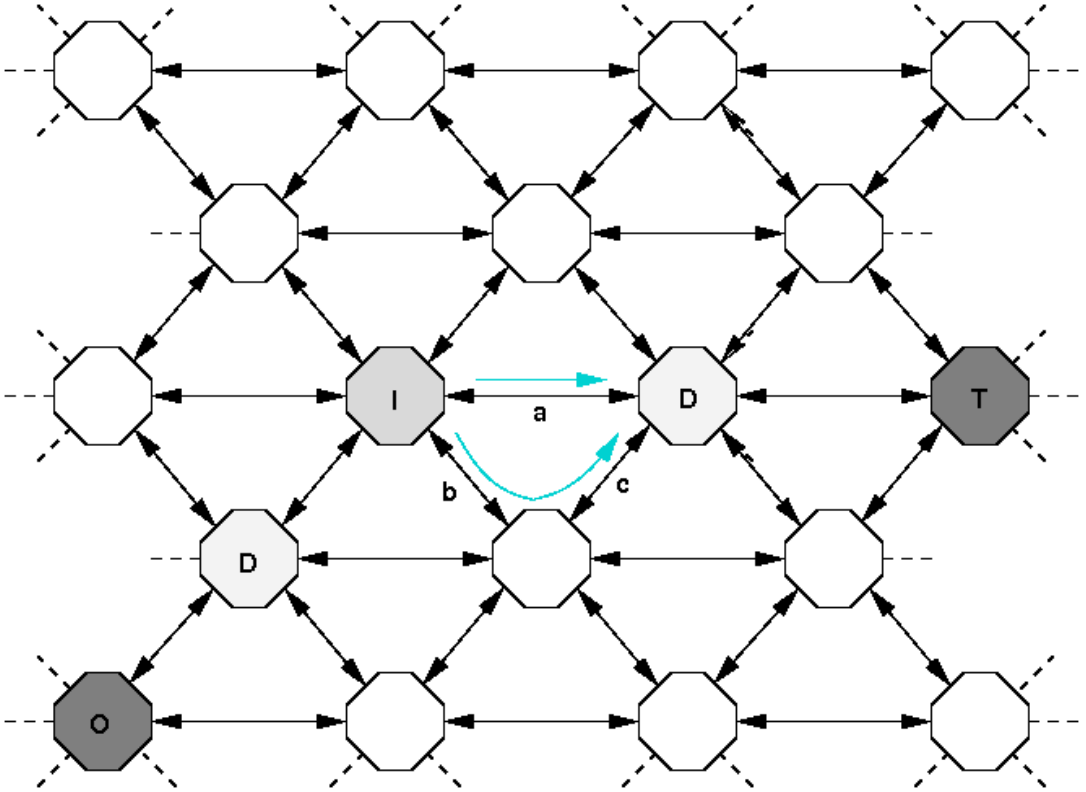

SpiNNaker
Biologically Inspired Massively Parallel Architectures

Outline

- 63 years of progress
- Many cores make light work
- Building brains
- The *SpiNNaker* project
- The networking challenge
- A generic neural modelling platform
- Plans & conclusions

SpiNNaker
Biologically Inspired Massively Parallel Architectures

The networking challenge

- Emulate the very high connectivity of real neurons
- A spike generated by a neuron firing must be conveyed efficiently to >1,000 inputs
- On-chip and inter-chip spike communication should use the same delivery mechanism

Network – packets

- Four packet types
  - MC (multicast): source routed; carry events (spikes)
  - P2P (point-to-point): used for bootstrap, debug, monitoring, etc
  - NN (nearest neighbour): build address map, flood-fill code
  - FR (fixed route): carry 64-bit debug data to host
- Timestamp mechanism removes errant packets
  - which could otherwise circulate forever

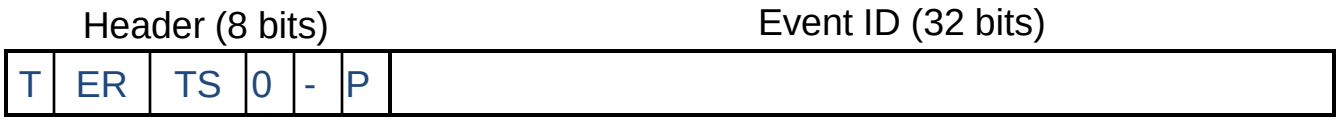

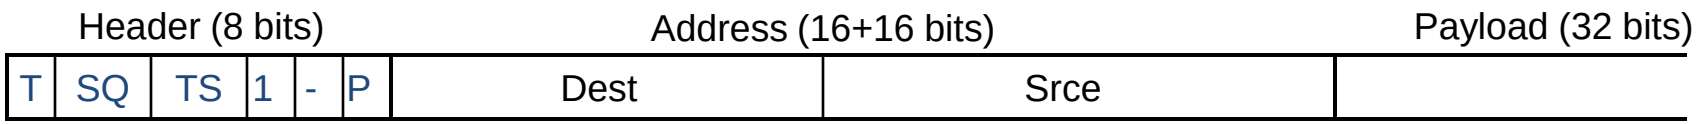

Network – MC Router

- All MC spike event packets are sent to a router
- Ternary CAM keeps router size manageable at 1024 entries (but careful network mapping also essential)
- CAM 'hit' yields a set of destinations for this spike event
  - automatic multicasting
- CAM 'miss' routes event to a 'default' output link

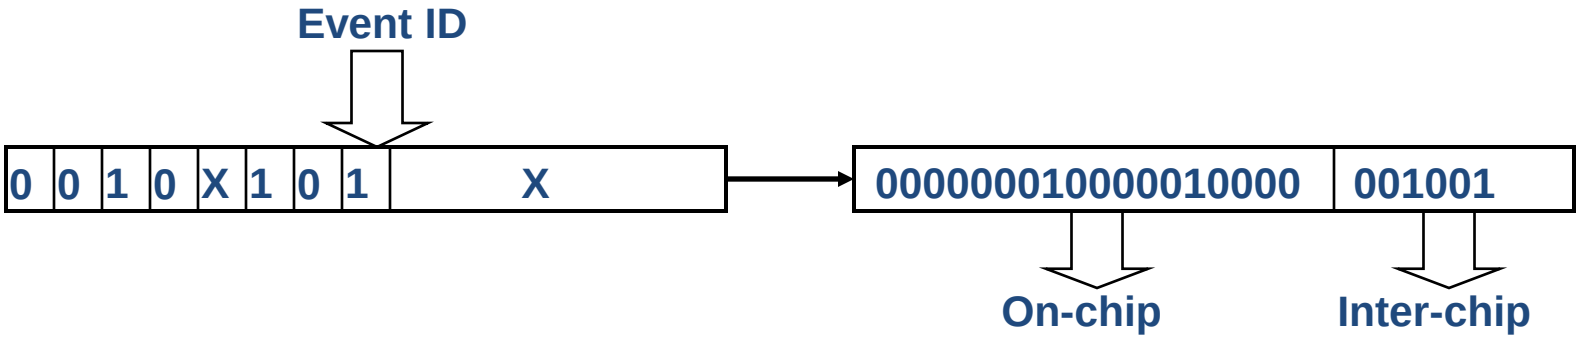

SpiNNaker
Biologically Inspired Massively Parallel Architectures

Topology mapping

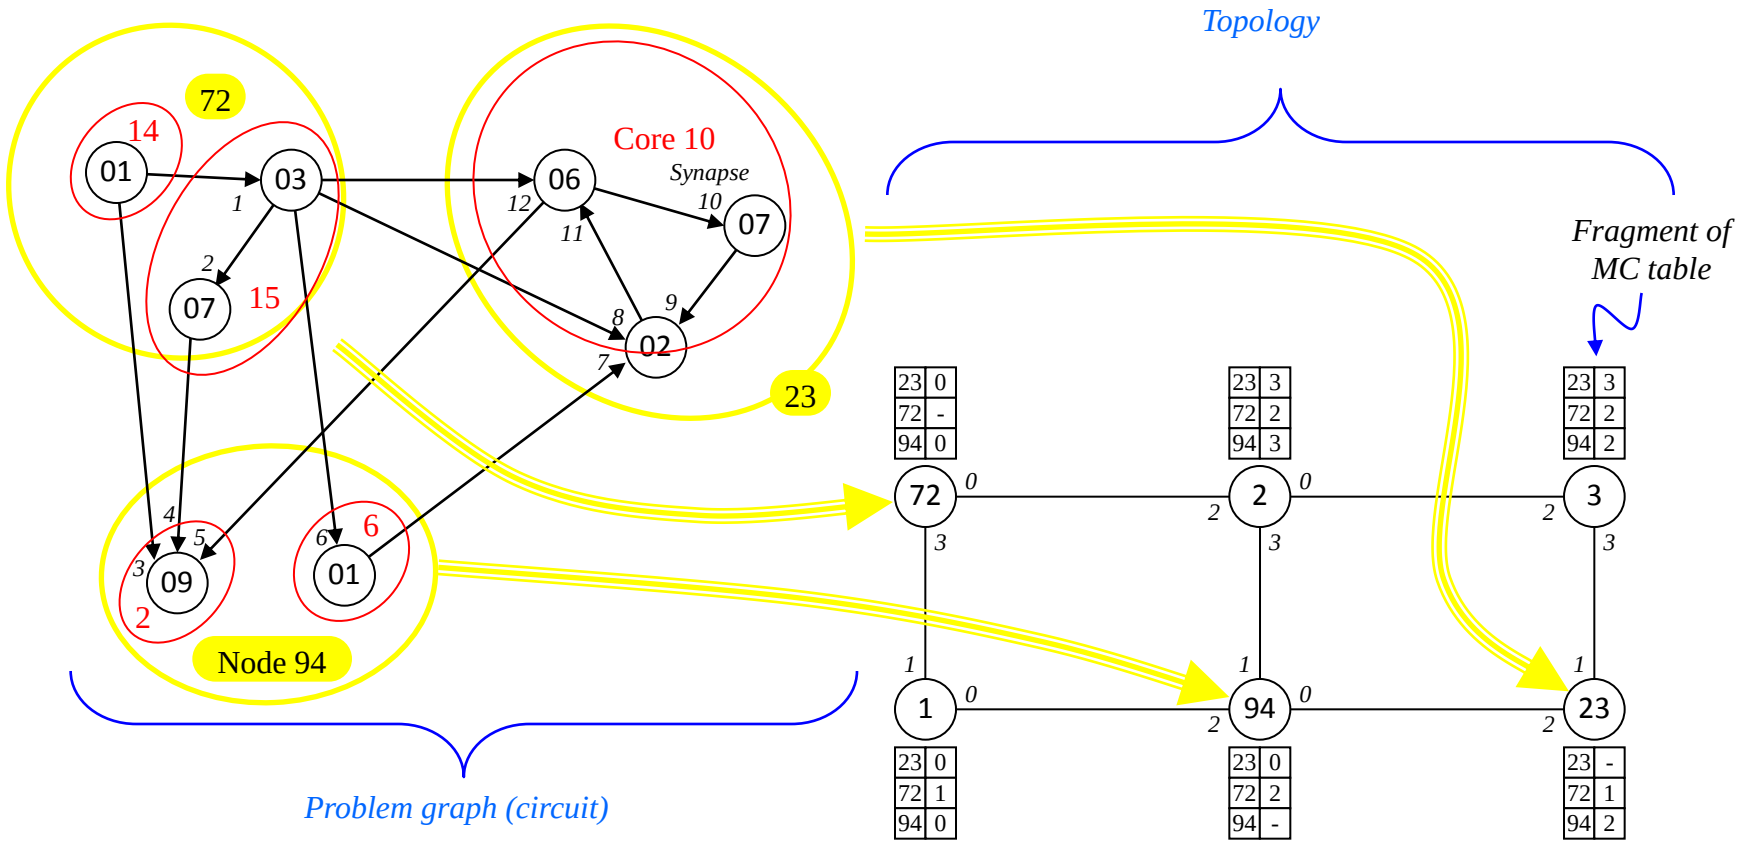

SpiNNaker
Biologically Inspired Massively Parallel Architectures

Problem mapping

SpiNNaker:

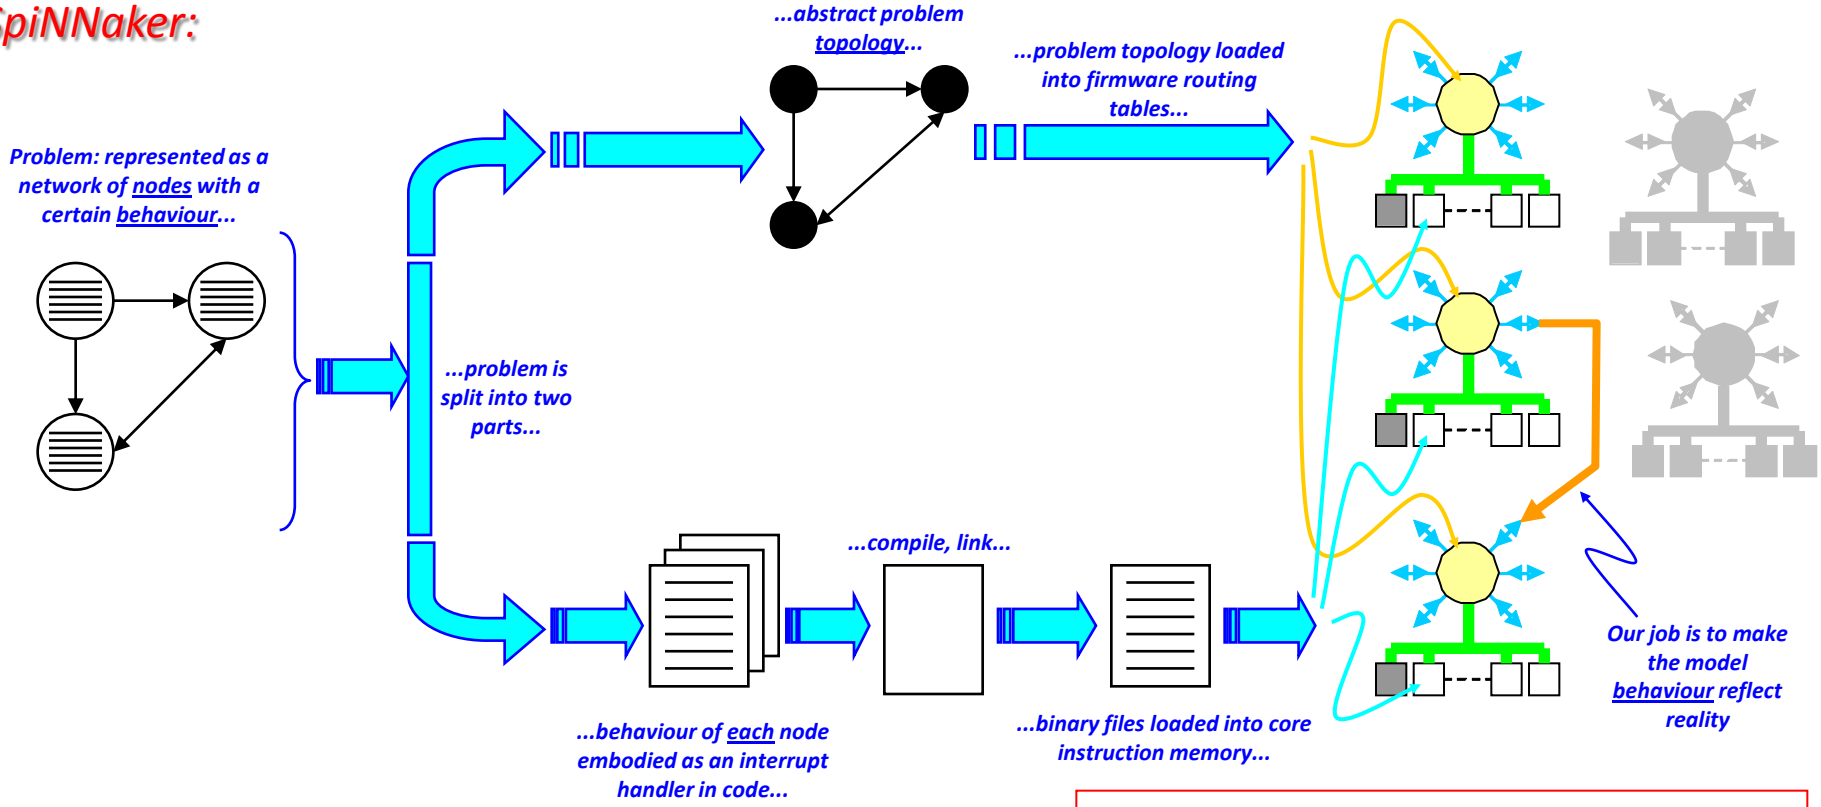

The code says "send message" but has no control where the output message goes

EPSRC UNIVERSITY OF CAMBRIDGE MANCHESTER 1824 The University Of Sheffield. UNIVERSITY OF Southampton ARM Silistix THALES

SpiNNaker
Biologically Inspired Massively Parallel Architectures

Bisection performance

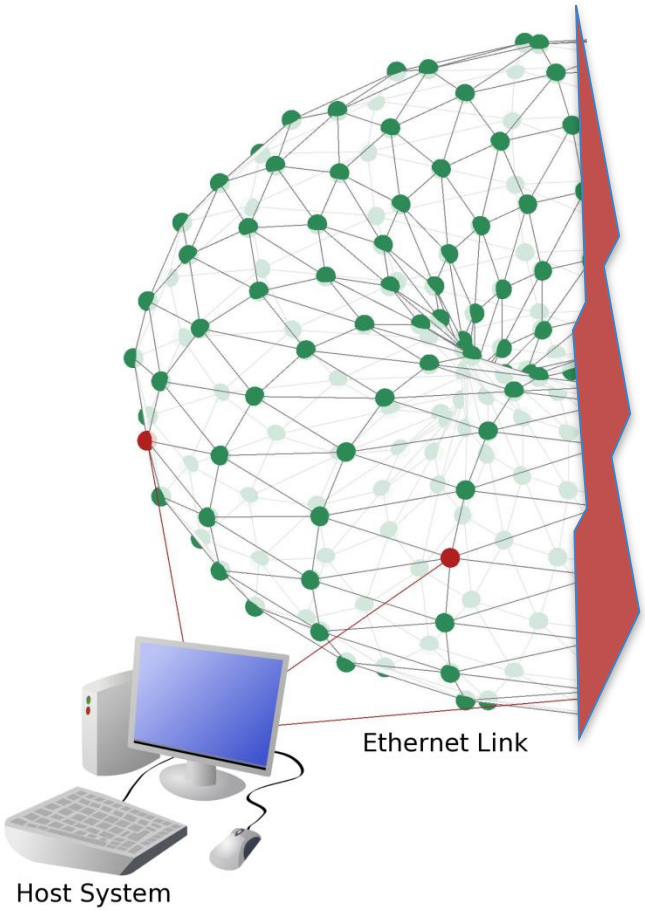

- 1,024 links
  - in each direction
- ~10 billion packets/s
- 10Hz mean firing rate
- 250 Gbps bisection bandwidth

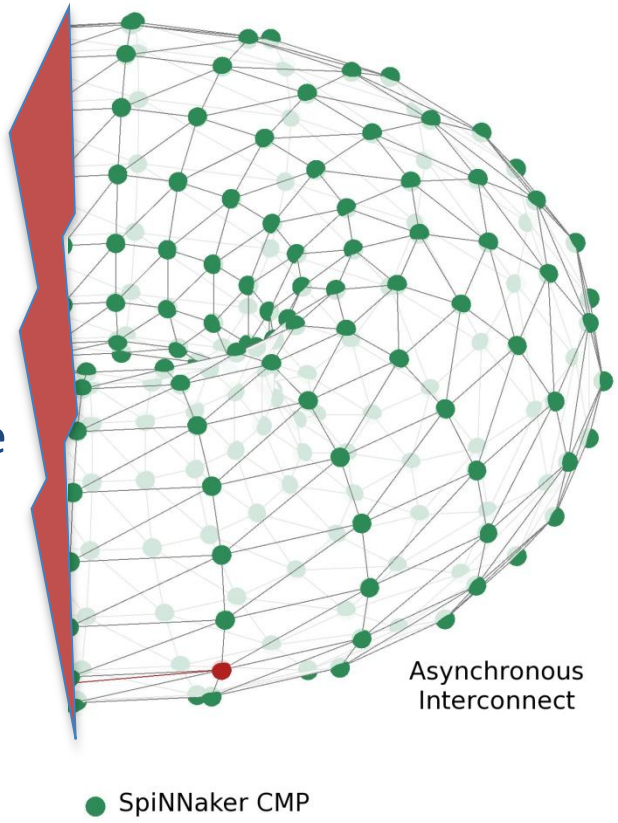

SpiNNaker
Biologically Inspired Massively Parallel Architectures

Outline

- 63 years of progress
- Many cores make light work
- Building brains
- The *SpiNNaker* project
- The networking challenge
- A generic neural modelling platform
- Plans & conclusions

SpiNNaker
Biologically Inspired Massively Parallel Architectures

Event-driven software model

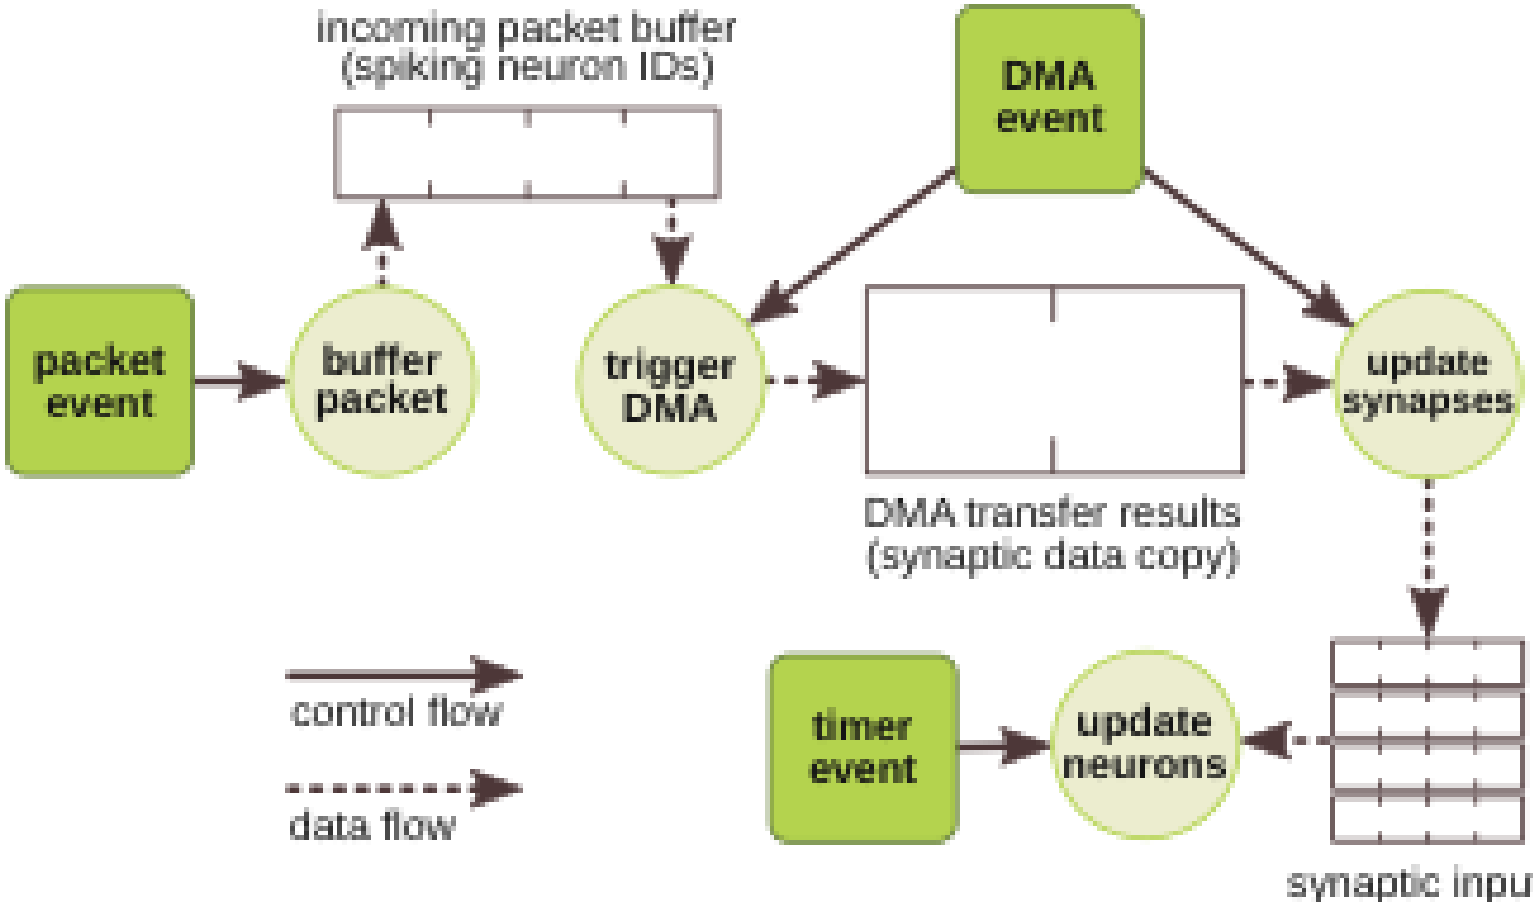

SpiNNaker
Biologically Inspired Massively Parallel Architectures

Event-driven software model

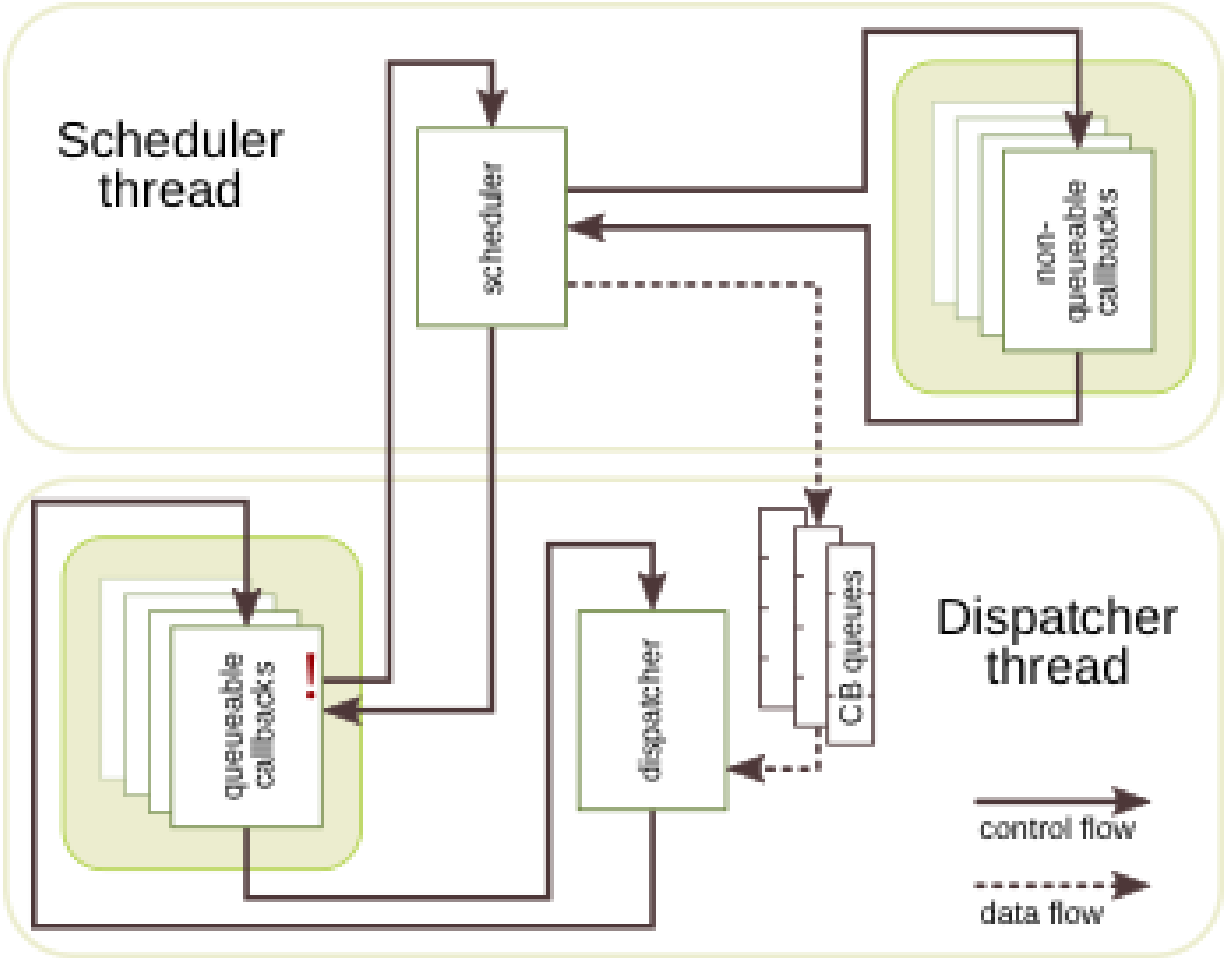

EPSRC UNIVERSITY OF CAMBRIDGE MANCHESTER 1824 The University Of Sheffield. UNIVERSITY OF Southampton ARM Silistix THALES

SpiNNaker
Biologically Inspired Massively Parallel Architectures

PACMAN
- Partitioning and Configuration Manager

High Level Interface
(PyNN, NEST, Lens, Damson etc.)

Splitting and Grouping
using SQL defined rules

Mapping
available cores + constraints

Binary File Generation

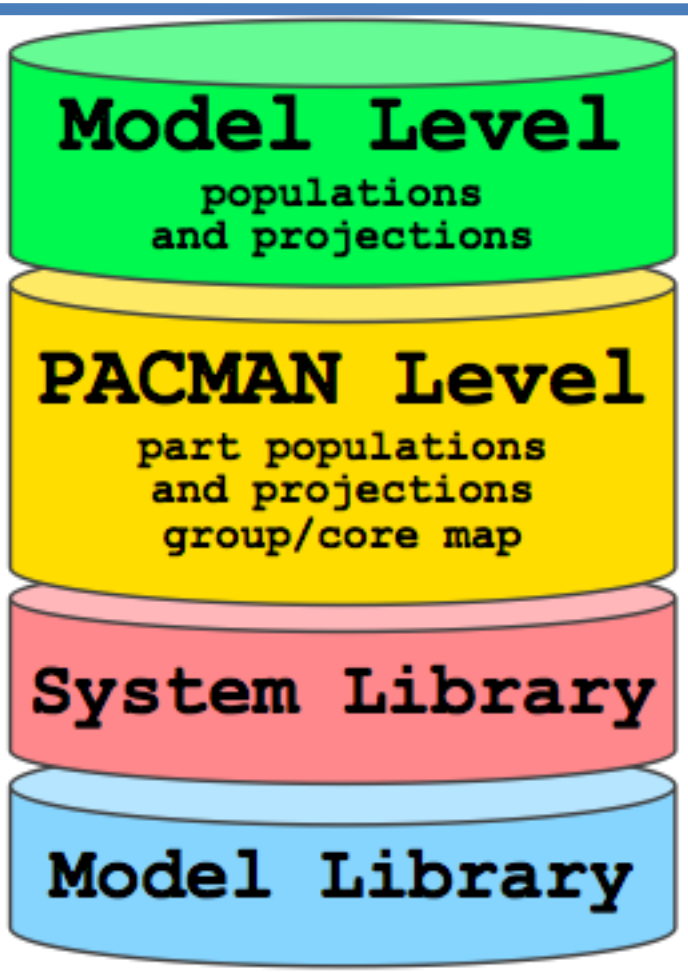

SpiNNaker
Biologically Inspired Massively Parallel Architectures

Self-tuning: software

- PACMAN: extrinsic configuration
  - good for small systems
- 1000-processor system
  - move table creation into *SpiNNaker*
- 10,000-100,000 processors
  - increasingly intrinsic configuration
- Million processor system
  - application loaded in one place
  - relax configuration across machine
  - continue relaxation at run-time to relax hot-spots

We don't know how to do this!

SpiNNaker
Biologically Inspired Massively Parallel Architectures

PyNN integration

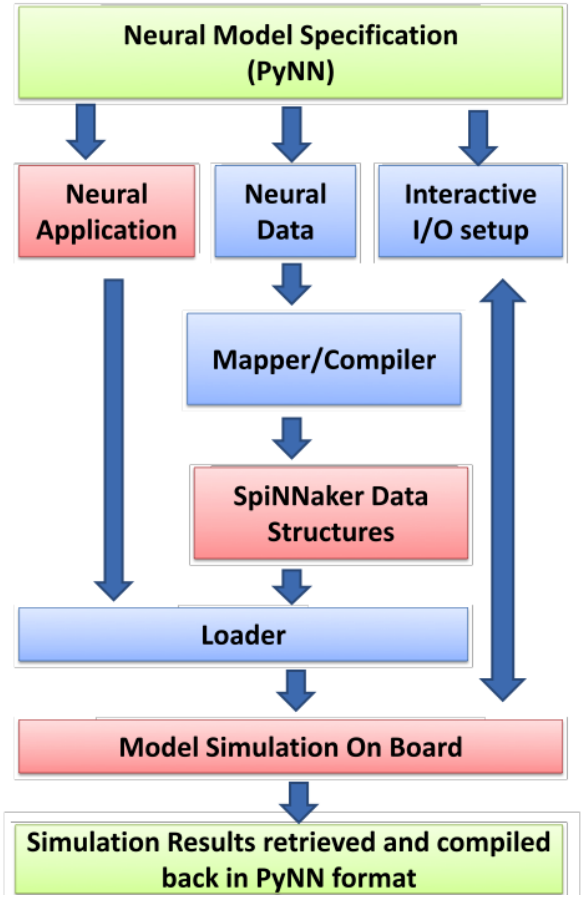

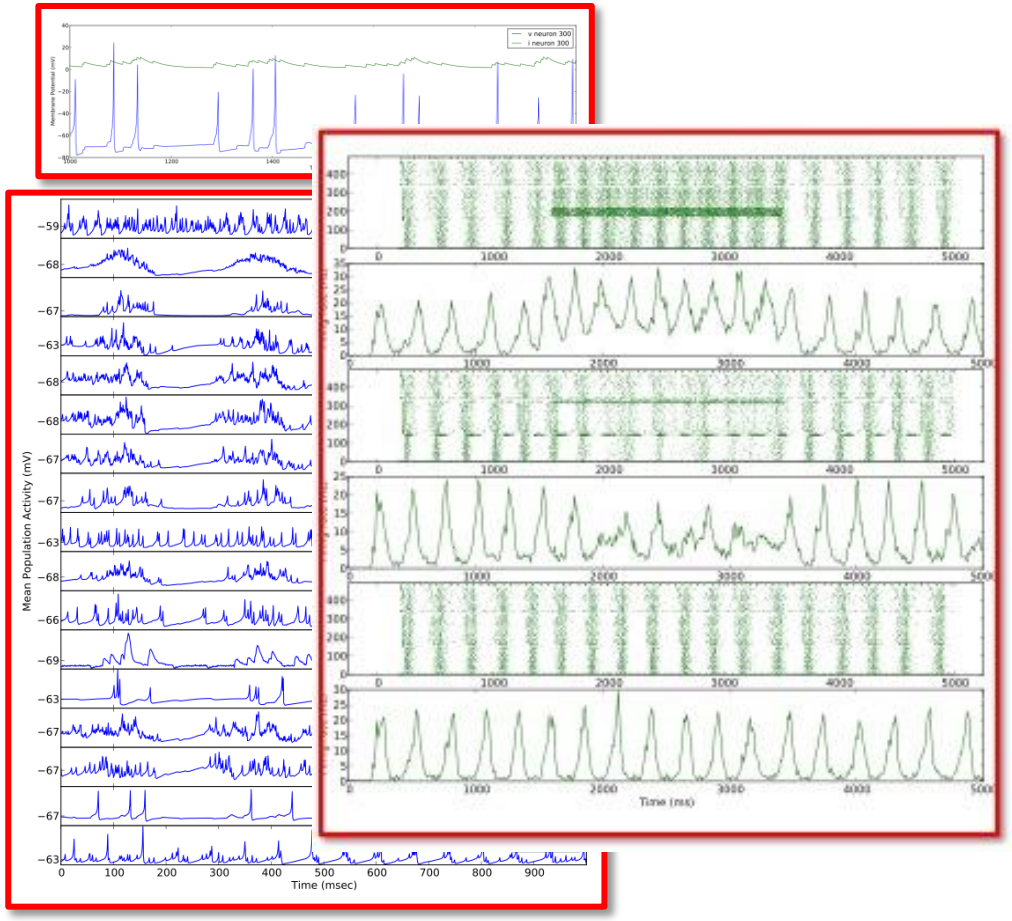

EPSRC UNIVERSITY OF CAMBRIDGE MANCHESTER 1824 The University Of Sheffield. UNIVERSITY OF Southampton ARM Silistix THALES

SpiNNaker
Biologically Inspired Massively Parallel Architectures

PyNN integration

- LIF

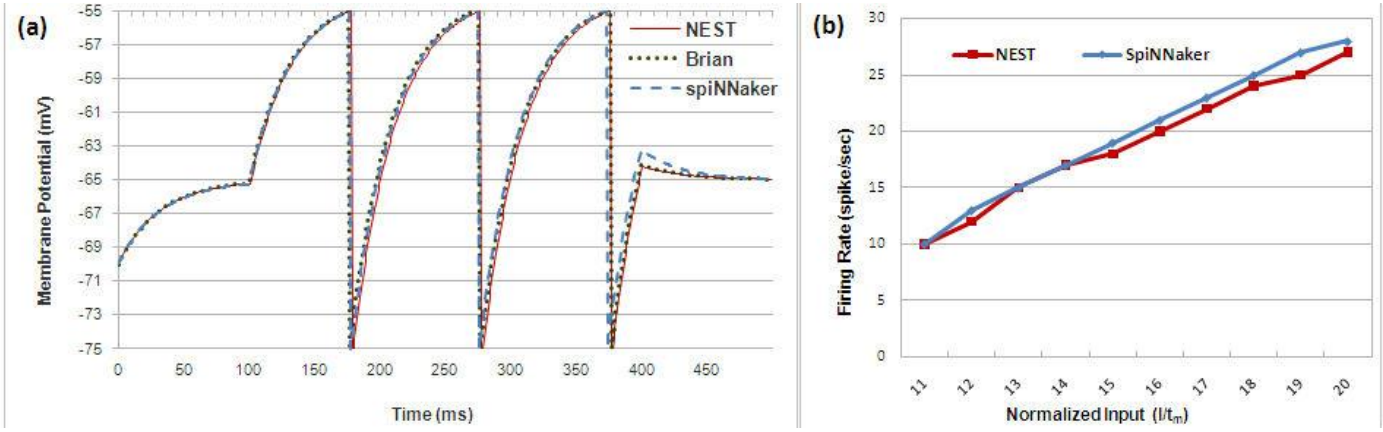

- Izhikevich

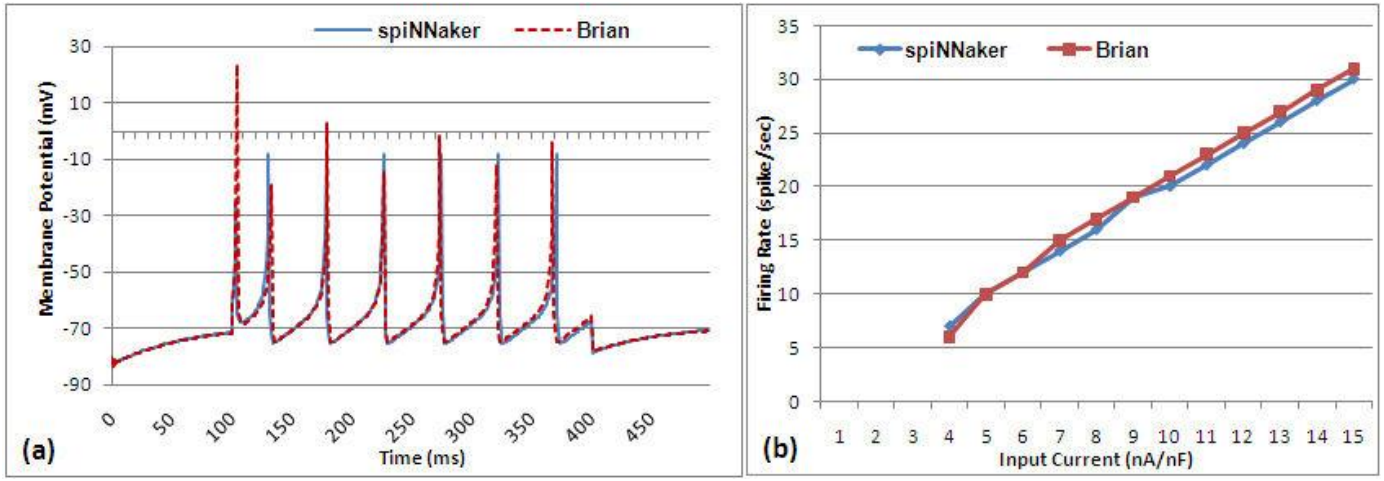

EPSRC UNIVERSITY OF CAMBRIDGE MANCHESTER 1824 The University Of Sheffield. UNIVERSITY OF Southampton ARM Silistix THALES

SpiNNaker
Biologically Inspired Massively Parallel Architectures

PyNN integration

- Vogels-Abbott benchmark
  - 500 LIF neurons

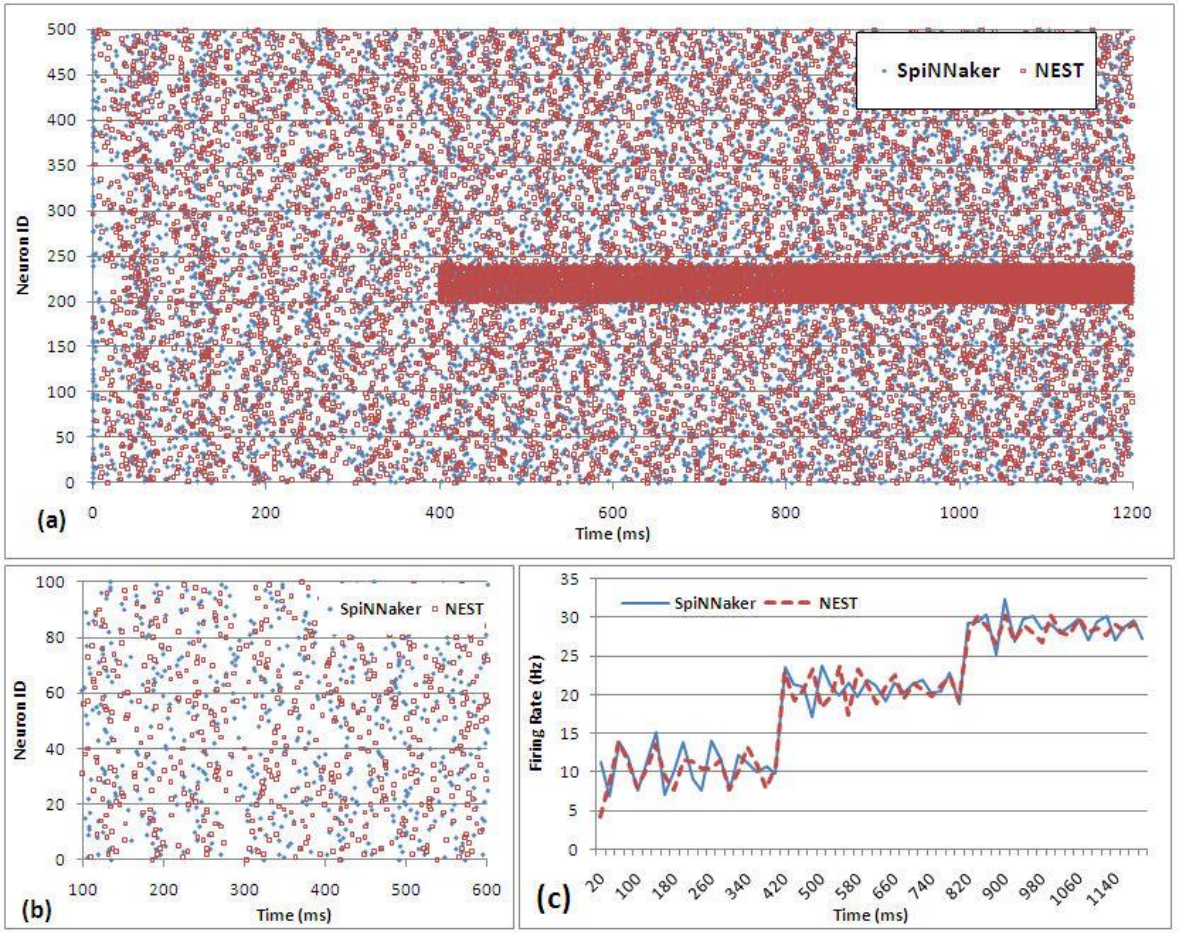

EPSRC UNIVERSITY OF CAMBRIDGE MANCHESTER 1824 The University Of Sheffield. UNIVERSITY OF Southampton ARM Silistix THALES

SpiNNaker
Biologically
Inspired
Massively
Parallel
Architectures

SpiNNaker robot control

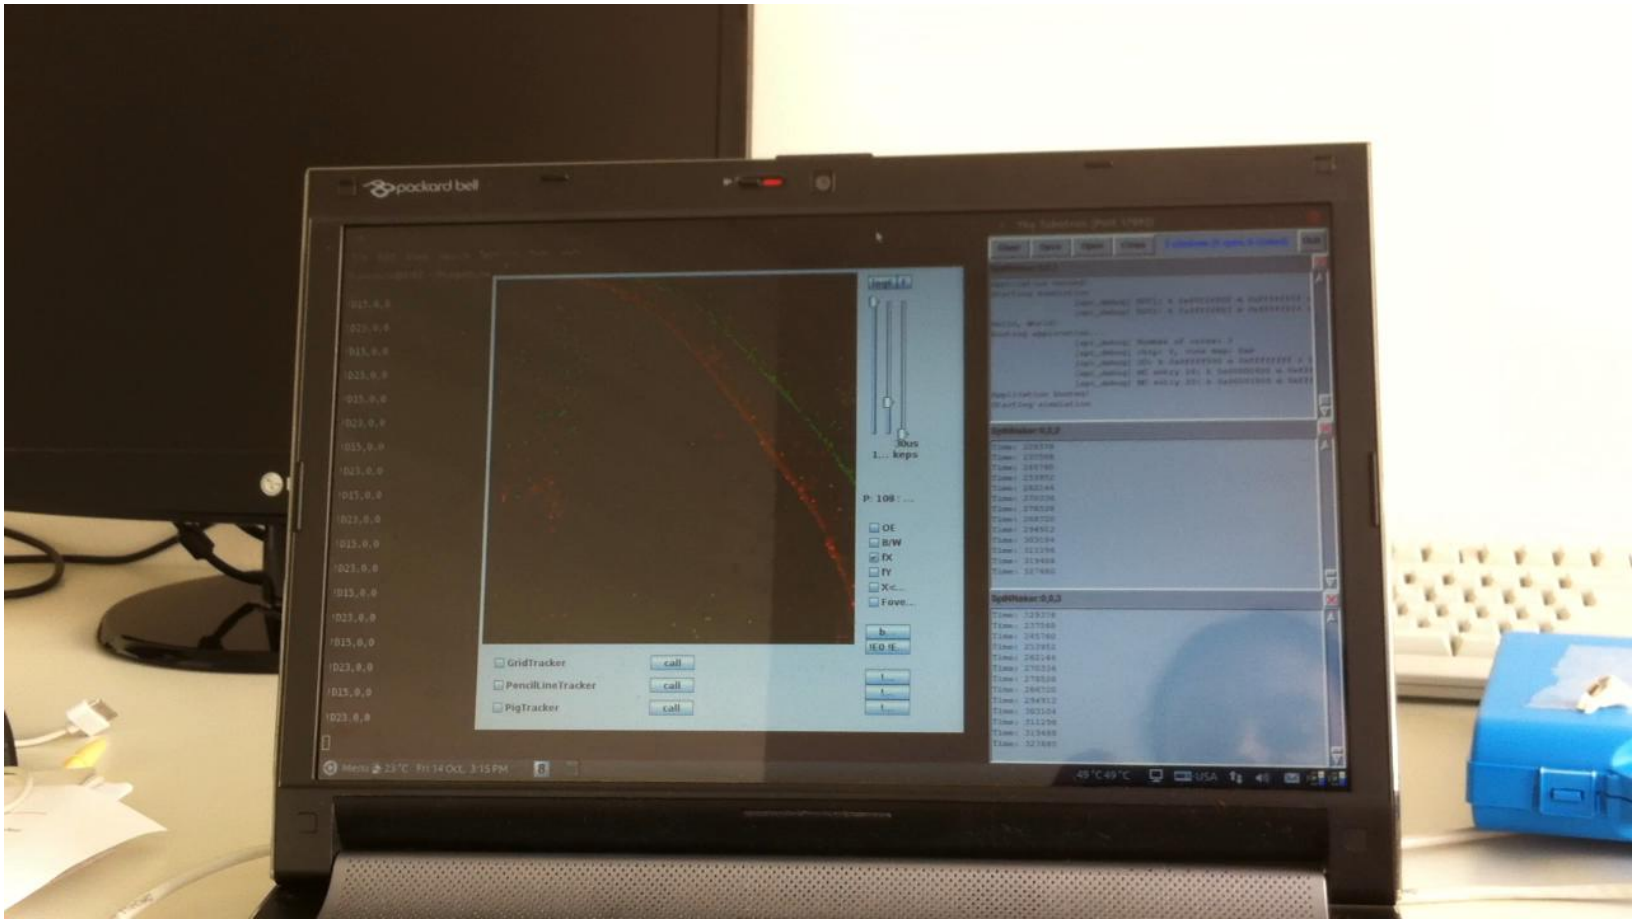

EPSRC UNIVERSITY OF CAMBRIDGE MANCHESTER 1824 The University Of Sheffield. UNIVERSITY OF Southampton ARM Silistix THALES

SpiNNaker
Biologically Inspired Massively Parallel Architectures

Outline

- 63 years of progress
- Many cores make light work
- Building brains
- The *SpiNNaker* project
- The networking challenge
- A generic neural modelling platform
- Plans & conclusions

SpiNNaker
Biologically Inspired Massively Parallel Architectures

Hexagonal PCB structure

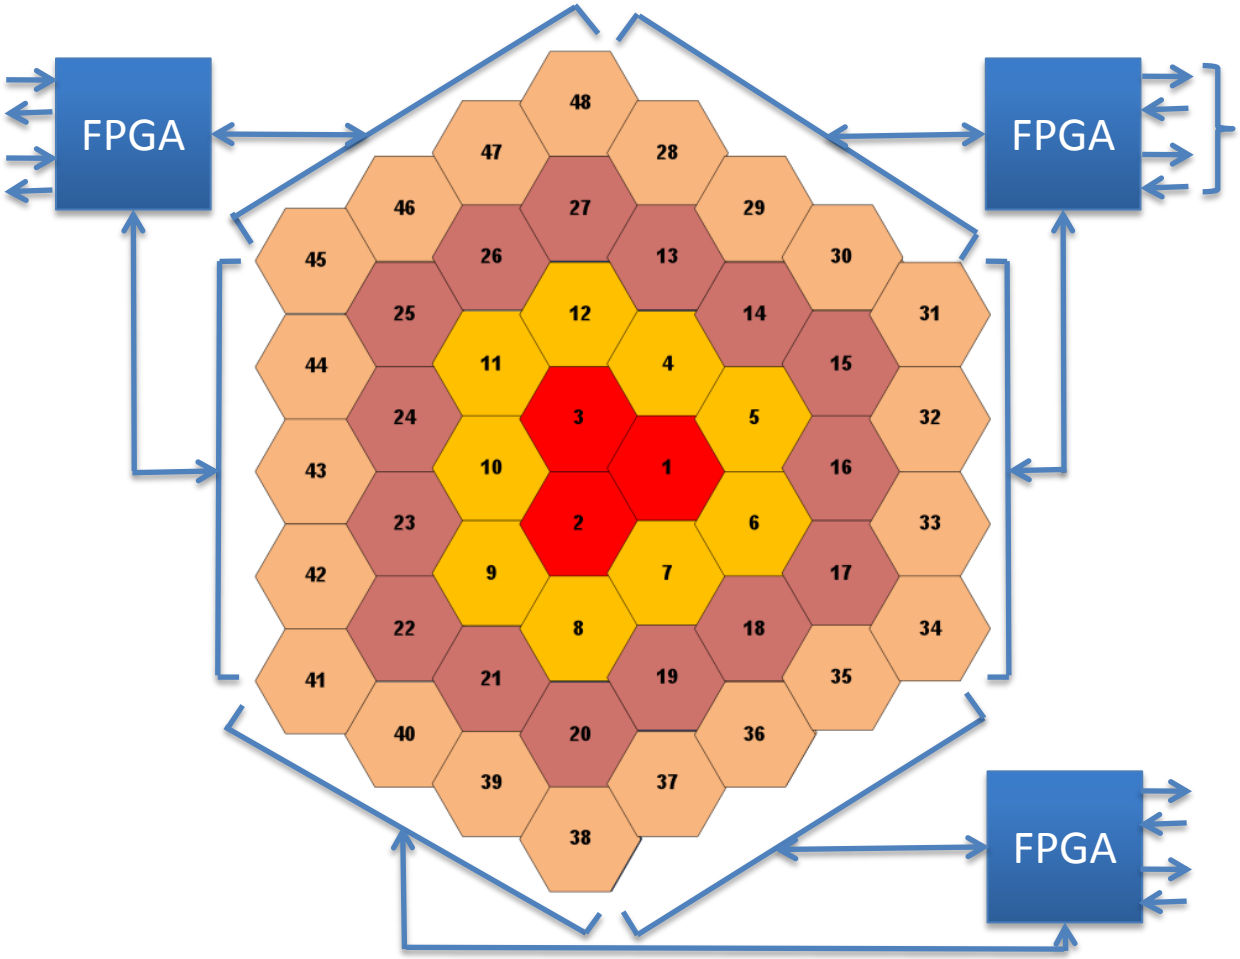

2x 3.1 Gbps SATA links

3-board basic unit:

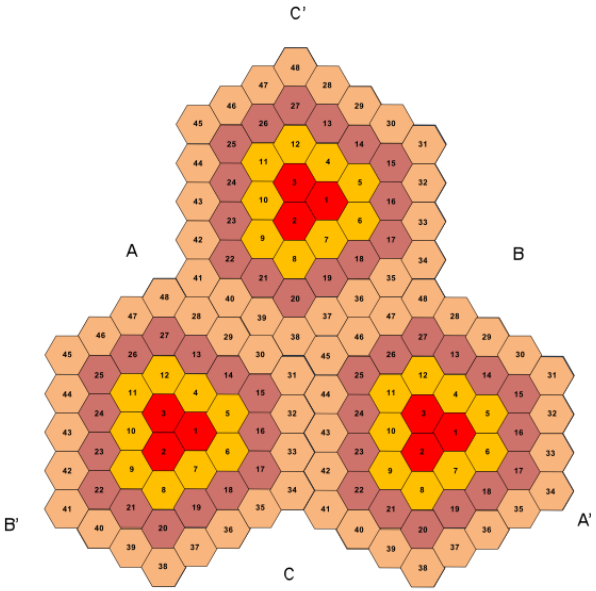

EPSRC UNIVERSITY OF CAMBRIDGE MANCHESTER 1824 The University Of Sheffield. UNIVERSITY OF Southampton ARM Silistix THALES

SpiNNaker
Biologically Inspired Massively Parallel Architectures

Hexagonal PCB structure

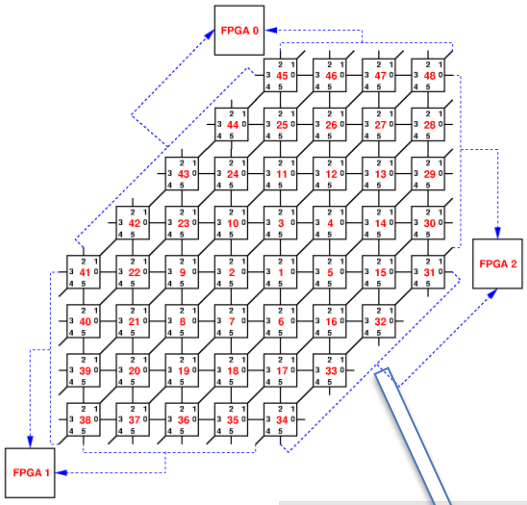

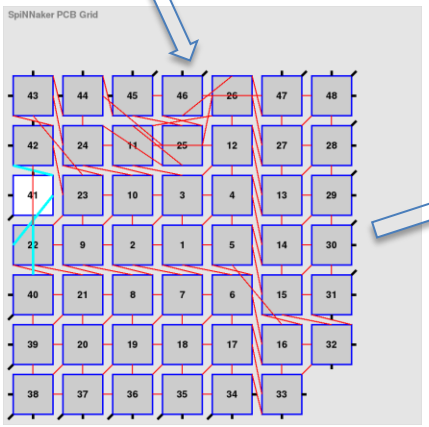

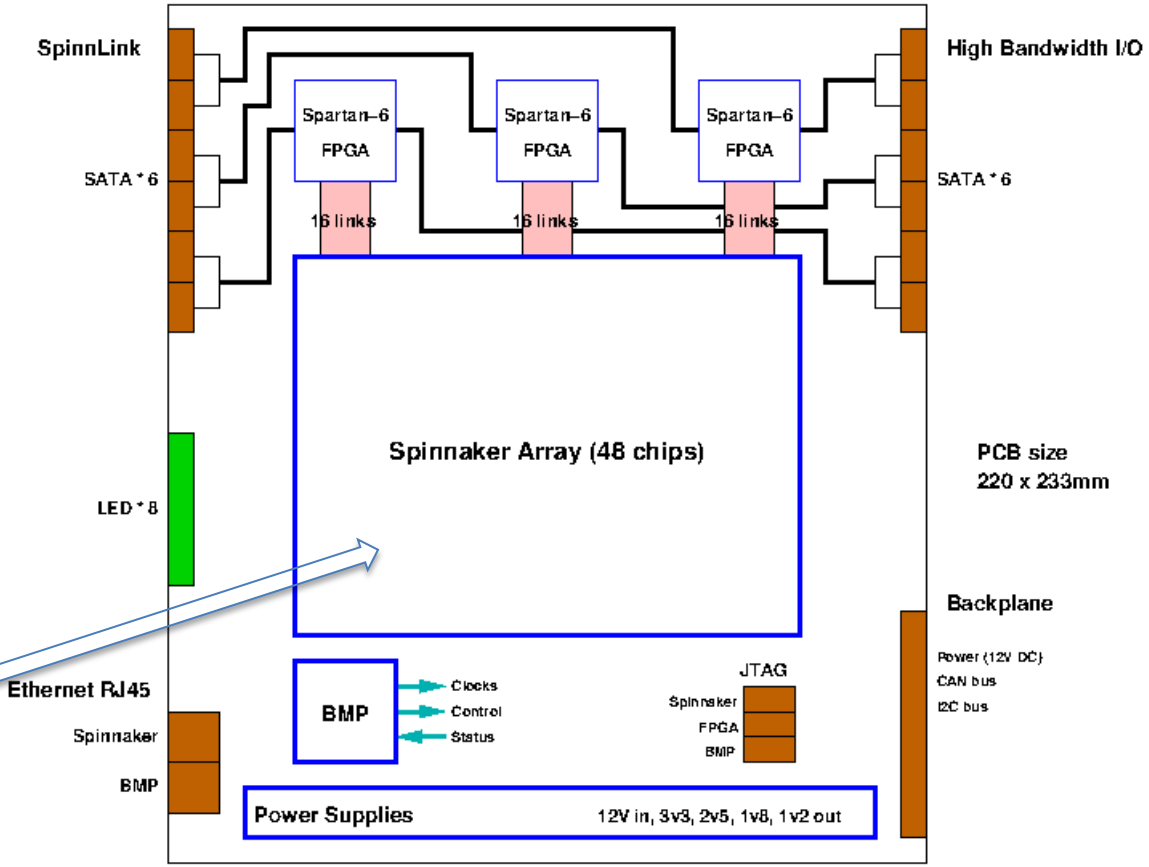

EPSRC UNIVERSITY OF CAMBRIDGE MANCHESTER 1824 The University Of Sheffield. UNIVERSITY OF Southampton ARM Silistix THALES

SpiNNaker
Biologically Inspired Massively Parallel Architectures

48-node PCB

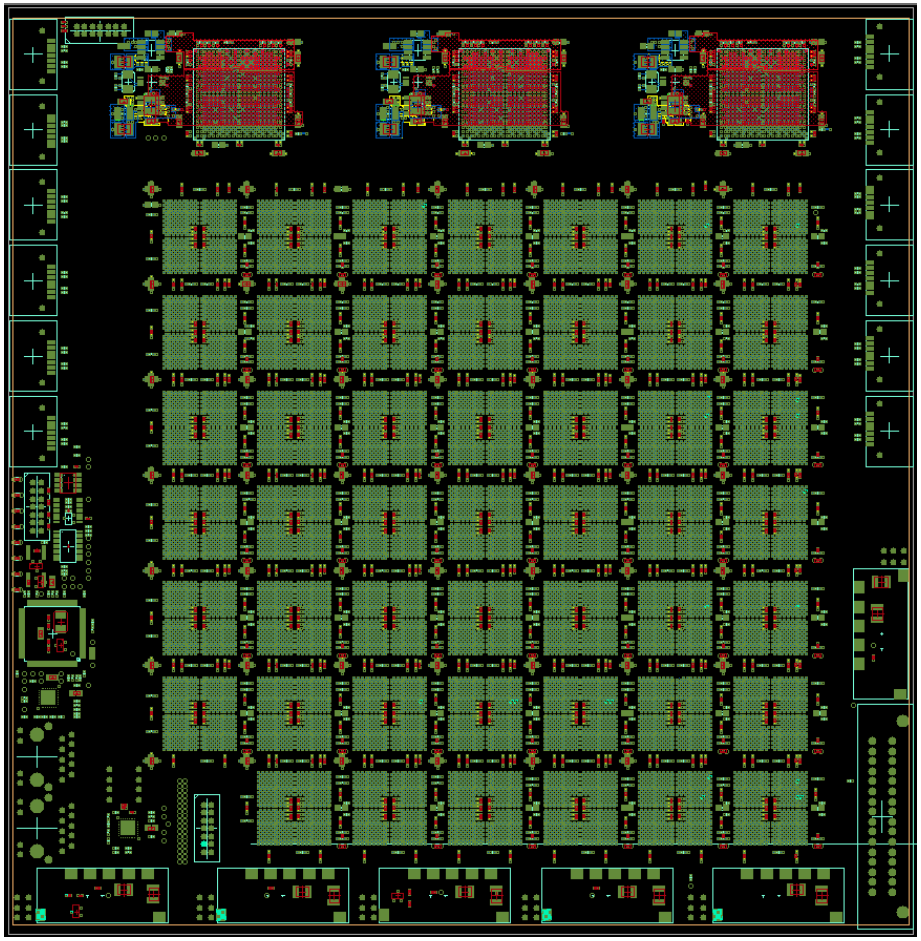

SpiNNaker
Biologically Inspired Massively Parallel Architectures

SpiNNaker machines

103 machine: 864 cores, 1 PCB, 75W

104 machine:10,368 cores, 1 rack, 900W
(NB 12 PCBs for operation without aircon)

105 m/c: 103,680 cores, 1 cabinet, 9kW

106 m/c: 1M cores, 10 cabs, 90kW

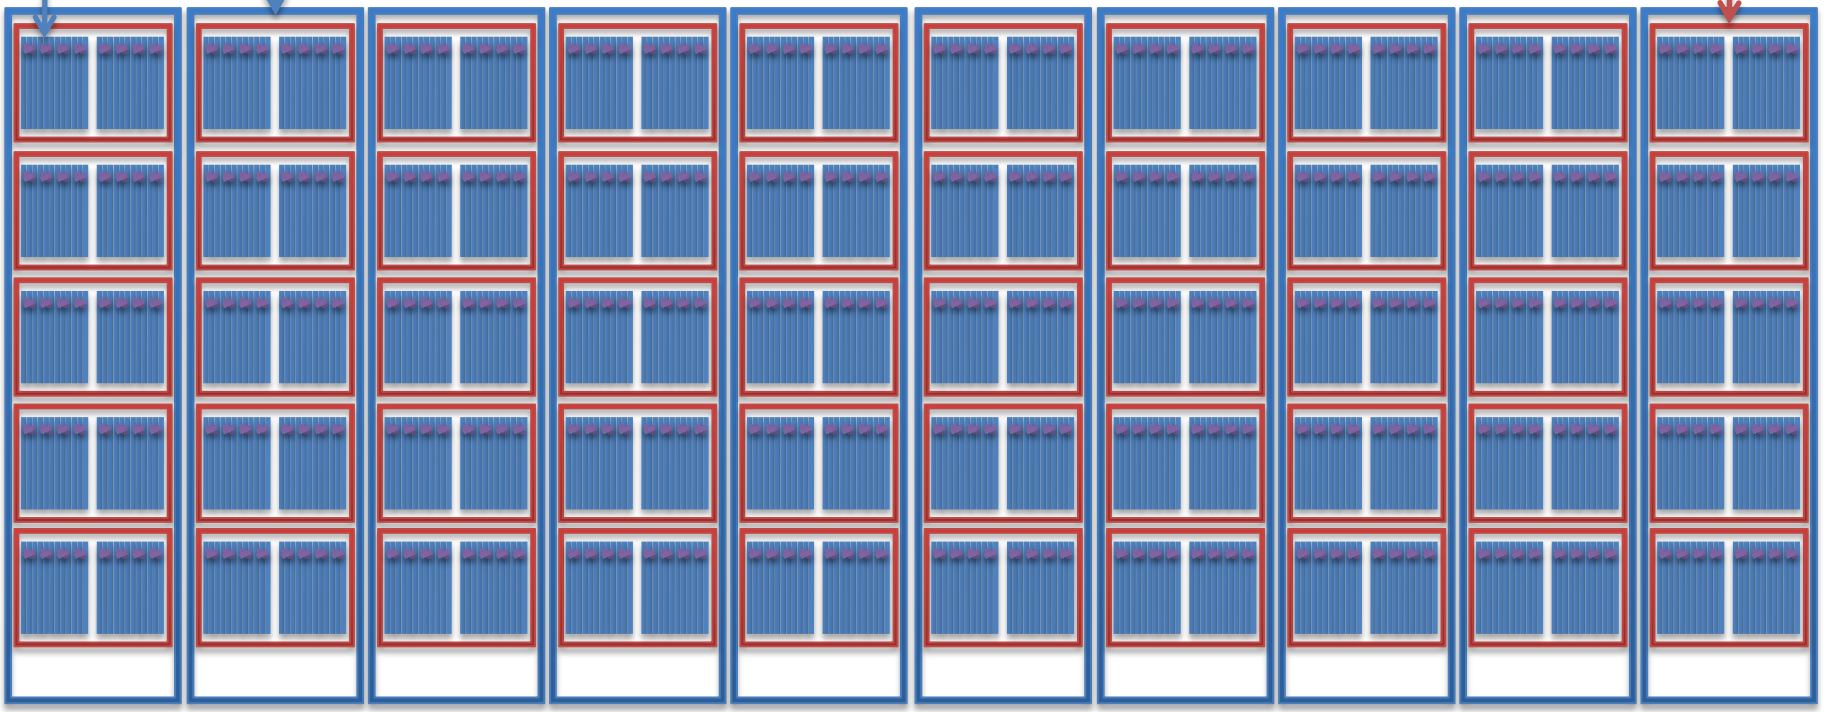

EPSRC UNIVERSITY OF CAMBRIDGE MANCHESTER 1824 The University Of Sheffield. UNIVERSITY OF Southampton ARM Silistix THALES

SpiNNaker
Biologically Inspired Massively Parallel Architectures

Current status…

- Full 18-core chip: arrived 20 May 2011
- Test card: 4 chips, 72 processors
  - Cards can be linked together
- Neuron models: LIF, Izhikevich, MLP
- Synapse models: STDP, NMDA
- Networks: PyNN -> SpiNNaker, various small tools to build Router tables, etc

…and the next steps:

- 48-chip 103 machine (Q1 2012),
  500-chip 104 machine (Q2 2012), 5,000-chip 105 machine (H2 2012), 50,000-chip 106 machine (end H2 2012).

SpiNNaker
Biologically Inspired Massively Parallel Architectures

Conclusions

- Brains represent a significant computational challenge
  - now coming within range?
- *SpiNNaker* is driven by the brain modelling objective
  - virtualised topology, bounded asynchrony, energy frugality
- The major architectural innovation is the multicast communications infrastructure
- Self-tuning at many levels
  - hardware (for fault-tolerance), software and, most effectively, in the neurons themselves!
- We have prototype working hardware!

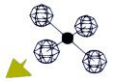

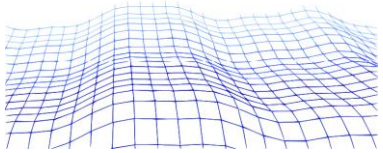

SpiNNaker
Biologically
Inspired
Massively
Parallel
Architectures

SpiNNaker team

Manchester

Southampton

EPSRC UNIVERSITY OF CAMBRIDGE MANCHESTER 1824 The University Of Sheffield. UNIVERSITY OF Southampton ARM Silistix THALES