

SpinNaker

**a universal
Spiking Neural Network
architecture**

version 0.4 - DRAFT
23 April 2007

SpiNNaker - a chip multiprocessor for neural network simulation

Features

- 20 ARM968 processors, each with:
 - 64 Kbytes of tightly-coupled data memory;
 - 32 Kbytes of tightly-coupled instruction memory;
 - DMA controller;
 - communications controller;
 - interrupt controller;
 - low-power ‘wait for interrupt’ mode.
- Multicast communications router
 - 6 serial inter-chip receive interfaces;
 - 6 serial inter-chip transmit interfaces;
 - 1024 associative routing entries.
- Interface to external SDRAM
 - over 1 Gbyte/s sustained block transfer rate.
- Fault-tolerant architecture
 - defect detection, isolation, and function migration.
- Boot, test and debug interfaces (to be determined).

Introduction

SpiNNaker is a chip multiprocessor designed specifically for the real-time simulation of large-scale spiking neural networks. Each chip (along with its associated SDRAM chip) forms one node in a scalable parallel system, interconnected to the other nodes through self-timed links.

The processing power is provided through the multiple ARM cores on each chip. Each ARM models multiple (up to 1,000) neurons, with each neuron being a coupled pair of differential equations modelled in continuous ‘real’ time. Neurons communicate through atomic ‘spike’ events, and these are communicated as discrete packets through the on- and inter-chip communications fabric. The packet contains a routing key that is defined at its source and is used to implement multicast routing through an associative router in each chip.

One processor on each SpiNNaker chip will perform system management functions; the communications fabric supports point-to-point packets to enable co-ordinated system management across local regions and across the entire system, and nearest-neighbour packets are used for system flood-fill boot operations and for chip debug.

Background

SpiNNaker was designed at the University of Manchester within an EPSRC-funded project in collaboration with the University of Southampton, ARM Limited and Silistix Limited. The work would not have been possible without EPSRC funding, and the support of the EPSRC and the industrial partners is gratefully acknowledged.

Intellectual Property rights

All rights to the SpiNNaker design are the property of the University of Manchester with the exception of those rights that accrue to the project partners in accordance with the contract terms.

Disclaimer

The details in this datasheet are presented in good faith but no liability can be accepted for errors or inaccuracies. The design of a complex chip multiprocessor is a research activity where there are many uncertainties to be faced, and there is no guarantee that a SpiNNaker system will perform in accordance with the specifications presented here.

The APT group in the School of Computer Science at the University of Manchester was responsible for all of the architectural and logic design of the SpiNNaker chip, with the exception of synthesizable components supplied by ARM Limited. All design verification was also carried out by the APT group. As such the industrial project partners bear no responsibility for the correct functioning of the device.

Change history

version	date	changes
0.0	27/12/05	first draft
0.1	16/8/06	sundry - document still developing
0.2	18/11/06	Comms controller and NN protocol details modified
0.3	19/02/07	Added ARM968 memory map, updated router pseudo-code, expanded system controller spec
0.4	23/04/07	Added DMA controller, updated system controller to use ADK watchdog, updated comms controller.

Contents

1. Chip organization	5
1.1 Block diagram	5
1.2 System-on-Chip hierarchy	6
2. System architecture	7
2.1 Routing	7
2.2 System-level address spaces	8
3. ARM968 processing subsystem	9
3.1 Features	9
3.2 ARM968 subsystem organisation	9
3.3 Fault-tolerance	9
3.4 Memory Map	9
3.5 Test	10
4. ARM 968	11
4.1 Features	11
4.2 Organization	11
4.3 Fault-tolerance	11
4.4 Test	11
5. Counter/timer and interrupt controller	12
5.1 Features	12
5.2 Register summary	12
5.3 Register details	13
5.4 Fault-tolerance	13
5.5 Test	13
5.6 Notes	13
6. DMA controller	14
6.1 Features	14
6.2 Using the DMA Controller	14
6.3 Register summary	16
6.4 Register details	17
6.5 Fault-tolerance	21
6.6 Test	21
7. Communications controller	22
7.1 Features	22
7.2 Packet formats	22
7.3 Control byte summary	23
7.4 Register summary	25
7.5 Register details	25
7.6 Fault-tolerance	27
7.7 Test	28
7.8 Notes	28
8. Communications NoC	29
8.1 Features	29
8.2 Block diagram	29
8.3 Arbiter structure	29
8.4 Fault-tolerance	30
8.5 Test	30
8.6 Notes	30

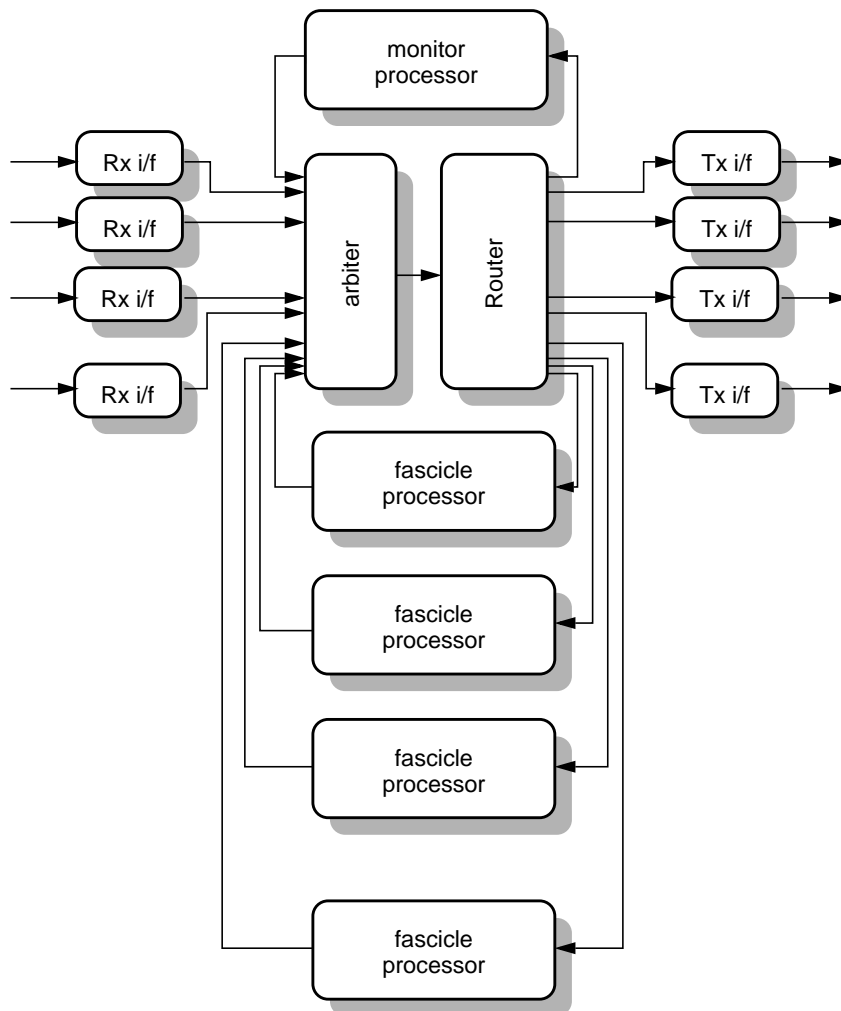
9. Communications Router	31
9.1 Features	31
9.2 Description	31
9.3 Internal organization	31
9.4 Multicast (MC) router	32
9.5 The point-to-point (p2p) router	35
9.6 The algorithmic (ALG) router	35
9.7 Time phase handling	35
9.8 Packet error handler	35
9.9 Emergency routing	36
9.10 Errant packets	36
9.11 Pseudo-code description	36
9.12 Registers	39
9.13 Fault-tolerance	39
9.14 Test	40
9.15 Notes	40
10. Inter-chip transmit and receive interfaces	41
10.1 Features	41
10.2 Programmer view	41
10.3 Fault-tolerance	41
10.4 Test	42
11. System NoC	43
11.1 Features	43
11.2 Fault-tolerance	43
11.3 Test	43
12. SDRAM interface	44
12.1 Features	44
12.2 Fault-tolerance	44
12.3 Test	44
13. System Controller	45
13.1 Register summary	45
13.2 Register details	45
13.3 Fault-tolerance	46
13.4 Test	46
14. Router configuration registers	47
14.1 Features	47
14.2 Register summary	47
14.3 Register details	47
14.4 Fault tolerance	49
14.5 Test	49
14.6 Notes	49
15. System RAM	50
15.1 Features	50
15.2 Fault-tolerance	50
15.3 Test	50
16. Boot ROM	51
16.1 Features	51
16.2 Fault-tolerance	51
16.3 Test	51

17. Boot, test and debug support	52
17.1 Features	52
17.2 Issues	52
17.3 Boot algorithm	52
17.4 Fault-tolerance	52
17.5 Test	52
18. Input and Output signals	53
18.1 External SDRAM interface	53
18.2	53
19. Area estimates	54
20. Power estimates	55

1. Chip organization

1.1 Block diagram

The primary functional components of SpiNNaker are illustrated in the figure below, which shows the Communications NoC and its clients.



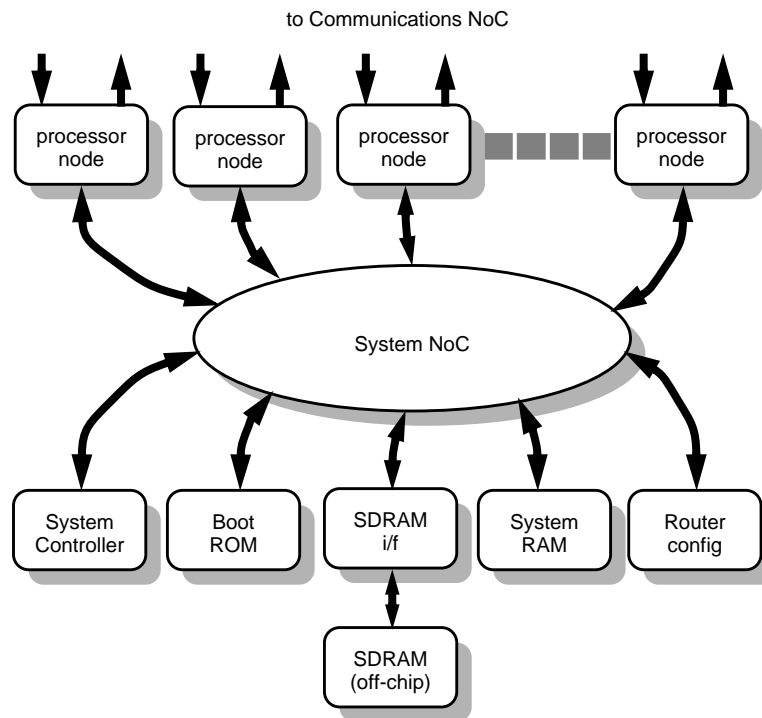
Each chip contains 20 identical processing subsystems each of which is responsible for modelling a number of neurons with associated inputs and outputs - a fascicle.

Following self-test, at start-up one of the processors is nominated as the Monitor Processor and thereafter performs system management tasks.

The router is responsible for routing neural event packets both between the on-chip fascicle processors and from and to other SpiNNaker chips. The Tx and Rx interface components are used to extend the on-chip communications NoC across to other SpiNNaker chips. The arbiter assembles inputs from the various on- and off-chip sources into a single serial stream which is then passed to the Router.

In addition to the primary function, there are additional resources accessible from the processor systems via the System NoC. Each of the fascicle processors has access to the shared off-chip SDRAM, and various system components also connect through the System NoC in order that, whichever processor is Monitor Processor, it will have access to these components.

The sharing of the SDRAM is an implementation convenience rather than a functional requirement, although it may facilitate function migration in support of fault-tolerant operation.




1.2 System-on-Chip hierarchy

The SpiNNaker chip is viewed as having the following structural hierarchy, which is reflected throughout the organisation of this datasheet:

- ARM968 processor subsystem
 - the ARM968, with its tightly-coupled instruction and data memories
 - Timer/counter and interrupt controller
 - DMA controller / System NoC interface
 - communications controller, including communications NoC interface
- Communications NoC
 - Router, including multicast, algorithmic, nearest-neighbour, default and emergency routing functions
 - 6 inter-chip transmit interfaces
 - 6 inter-chip receive interfaces
 - communications NoC arbiter and fabric
- System NoC
 - SDRAM interface
 - System Controller
 - Router configuration registers
 - Boot ROM
 - System RAM
 - System NoC arbiter and fabric
- Boot, test and debug
 - central controller for ARM968 JTAG functions

2. System architecture

SpiNNaker is designed to form (with its associated SDRAM chip) a node of a massively parallel system. The system architecture is illustrated below:



2.1 Routing

The nodes are arranged in a *hexagonal* mesh with bidirectional links to 6 neighbours. The system supports multicast packets (to carry neural event information, routed by the associative Multicast Router), point-to-point packets (to carry system management and control information, routed algorithmically) and nearest-neighbour packets (to support boot-time flood-fill and chip debug).

Emergency routing

In the event of a link failing or congesting, traffic that would normally use that link is redirected in hardware around two adjacent links that form a triangle with the failed link. This “emergency routing” is intended to be temporary, and the operating system will identify a more permanent resolution of the problem. The local Monitor Processor is informed of all uses of emergency routing.

Deadlock avoidance

The communications system has potential deadlock scenarios because of the possibility of circular dependencies between links. The policy used here to prevent deadlocks occurring is:

- *no Router can ever be prevented from issuing its output.*

The mechanisms used to ensure this are the following:

- outputs have sufficient buffering and capacity detection so that the Router knows whether or not an output has the capacity to accept a packet;
- emergency routing is used, where possible, to avoid overloading a blocked output;

- where emergency routing fails (because, for example, the alternative output is also blocked) the packet is ‘dropped’ to the local Monitor Processor;
- the local Monitor Processor is guaranteed to accept the dropped packet (eventually).

The expectation is that the communications fabric will be lightly-loaded so that blocked links are very rare. Where the operating system detects that this is not the case it will take measures to correct the problem by modifying routing tables or migrating functionality to a different part of the system.

Errant packet trap

Packets that get mis-routed could continue in the system for ever, following cyclic paths. To prevent this all packets are time stamped and a coarse global time phase signal is used to trap old packets. To minimize overhead the time stamp is 2 bits, cycling 00 -> 01 -> 11 -> 10, and when the packet is two time phases old (time sent XOR time now = 0b11) it is dropped to the local Monitor Processor and an error flagged. The length of a time phase can be adapted dynamically to the state of the system; normally timed-out packets should be very rare so the time phase can be conservatively long to minimise the risk of packets being dropped due to congestion.

2.2 System-level address spaces

The system incorporates a number of different levels of component that must be enumerated in some way:

- Each Node (where a Node is an SpiNNaker chip plus SDRAM) must have a unique, fixed address which is used as the destination ID for a point-to-point packet, and the addresses must be organised logically for algorithmic routing to function efficiently.
- Processors will be addressed relative to their host Node address, but this mapping will not be fixed as an individual Processor’s role can change over time. Point-to-point packets addressed to a Node will be delivered to the local Monitor Processor, whichever Processor is serving that function. Internal to a Node there will be some hard-wired addressing of each Processor for system diagnosis purposes, but this mapping will be hidden outside the Node.
- Neurons occupy an address space that identifies each Neuron uniquely within the domain of its multicast routing path (where this domain must include alternative links that may be taken during emergency routing). Where these domains do not overlap it is possible to reuse the same address, though this must be done with considerable care. Neuron addresses can be assigned arbitrarily, and this flexibility can be exploited to optimize Router utilization (for example by giving Neurons with the same routing requirements related addresses so that they can all be routed by the same Router entries).

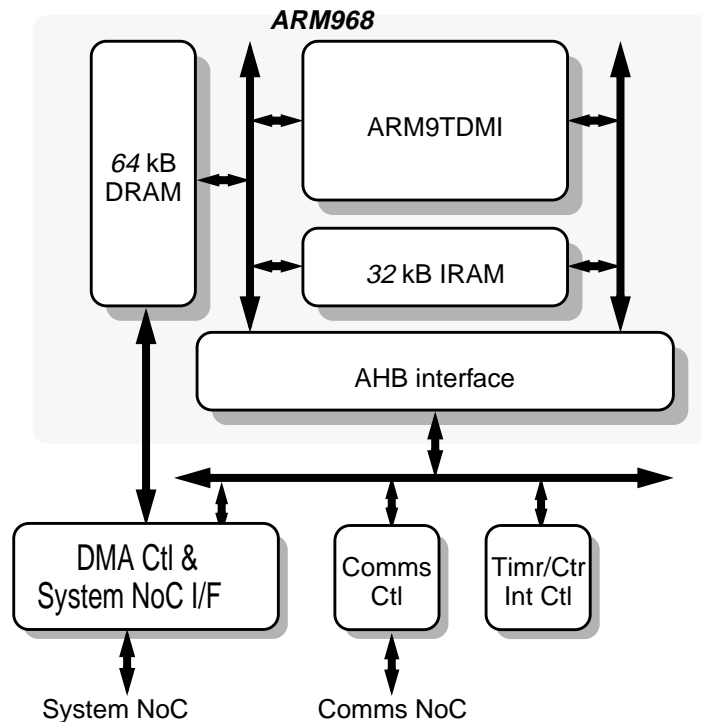
3. ARM968 processing subsystem

SpiNNaker incorporates 20 ARM968 processing subsystems which provide the computational capability of the device. Each of these subsystems is capable of generating and processing neural events communicated via the Communications NoC and, alternatively, of fulfilling the role of Monitor Processor.

3.1 Features

- a synthesized ARM968 module with
 - a 200 MIPS ARM9 processor
 - 32 kB tightly-coupled instruction memory
 - 64 kB tightly-coupled data memory
- a local AHB with
 - communications controller connected to Communications NoC
 - DMA controller & interface to the System NoC
 - timer/counter and interrupt controller

3.2 ARM968 subsystem organisation



3.3 Fault-tolerance

The fault-tolerance of the ARM968 subsystem is defined in terms of its component parts, described below.

3.4 Memory Map

The memory map of the ARM968 spans a number of devices and buses. The tightly coupled memories are directly connected to the processor and accessible at the processor clock speed. All other parts of the memory map are visible via the AHB master interface. This gives direct access to

the registers of the DMA controller, communications controller and the timer/interrupt controller. In addition, a path is available through the DMA controller onto the System NoC which provides processor access to all memory resources on the System NoC. The memory map is defined as follows:

```
// Local to ARM968 nodes
#define ITCM_STARTING_ADDRESS      0x00000000
#define DTCM_STARTING_ADDRESS      0x00400000

// 8MB address area (NNOP) reserved for
// "operation mapping" in the NN protocol
// Can identify this area by
// top 9 bits=000011111
#define NNOP_STARTING_ADDRESS      0x0f800000
#define NNOP_END_ADDRESS           0x0fffffff

// Local peripherals
#define COMM_CTL_STARTING_ADDRESS  0x10000000 // communications controller
#define INT_CTL_STARTING_ADDRESS   0x10400000 // counter-interrupt
controller
#define DMA_CTL_STARTING_ADDRESS    0x10800000 // DMA controller

// System NoC
#define SYSTEMNOC_STARTING_ADDRESS 0x20000000
#define SDRAM_STARTING_ADDRESS     0x20000000 // buffered write access
#define SDRAM_END_ADDRESS          0x2fffffff // buffered write access
#define RESERVED_SDRAM_UNBUF_START_ADDRESS 0x30000000 // unbuffered write access
#define RESERVED_SDRAM_UNBUF_END_ADDRESS  0x30000000 // unbuffered write access

#define UNALLOC_SPACE_START        0x40000000 // address space not
allocated
#define UNALLOC_SPACE_END          0xfeffffff

#define RTRCONFIG_STARTING_ADDRESS 0xff000000 // router
#define SYSCTL_STARTING_ADDRESS     0xff400000 // system controller
#define SYSRAM_STARTING_ADDRESS     0xff800000 // system RAM
#define BOOTROM_STARTING_ADDRESS    0xffff0000 // system ROM
```

3.5 Test

The test strategies for the ARM968 subsystem are likewise defined in terms of its component parts.

4. ARM 968

The ARM968 (with its associated tightly-coupled instruction and data memories) forms the core processing resource in SpiNNaker. It is a standard synthesizable IP component from ARM Ltd, and as such there is limited scope for customizing it for this application.

4.1 Features

- 200 MIPS ARM9TDMI processor.
- 32 kB tightly-coupled instruction memory (I-RAM).
- 64 kB tightly-coupled data memory (D-RAM).
- AHB interface to external system.

4.2 Organization

See ARM DDI 0311C – the ARM968E-S datasheet.

4.3 Fault-tolerance

Fault insertion

- ARM9TDMI can be disabled.
- Software can corrupt I-RAM and D-RAM to model soft errors. (Can these be detected?)

Fault detection

- The I-RAM and D-RAM are protected by parity bits?
- A chip-wide watchdog timer to catch runaway software.
- Self-test routines, run at start-up and during normal operation, can detect faults.

Fault isolation

- The ARM968 unit can be disabled from the System Controller.
- Defective locations in the I-RAM and D-RAM can be mapped out of use by software.

Reconfiguration

- Software will avoid using defective I-RAM and D-RAM locations.
- Functionality will migrate to an alternative Processor in the case of permanent faults that go beyond the failure of one or two memory locations.

4.4 Test

production test

start-up test

run-time test

5. Counter/timer and interrupt controller

Each processor node on an SpiNNaker chip has a local counter/timer and interrupt controller that is used to enable and disable interrupts from various sources, and to wake the processor from sleep mode when required. The interrupt controller provides centralised management of IRQ and FIQ sources, and offers an efficient indication of the active sources for vectoring purposes.


5.1 Features

- manages the various interrupt sources to each local processor:
 - arriving multicast packet with/without payload
 - arriving point-to-point packet with/without payload
 - arriving nearest-neighbour packet with/without payload
 - DMA complete/error/timeout
 - timer 1 & timer 2 interrupts
 - interrupt from another processor on the chip (usually the Monitor processor)
 - packet-error interrupt from the Router
 - system fault interrupt
- the counter/timer unit provides two independent counters, for example for:
 - millisecond interrupts for real-time dynamics

5.2 Register summary

Name	Offset	R/W	Function
r0: FIQ status	0x0	R	active FIQ interrupt sources
r1: FIQ enable	0x4	W	enable FIQ interrupt sources
r2: FIQ inputs	0x8	R	state of FIQ request inputs
r3:			
r4: IRQ status	0x0	R	active IRQ interrupt sources
r5: IRQ enable	0x4	W	enable IRQ interrupt sources
r6: IRQ inputs	0x8	R	state of IRQ request inputs
r7:			
r8:			

5.3 Register details



5.4 Fault-tolerance

Fault insertion

Fault detection

Fault isolation

Reconfiguration

5.5 Test

production test

start-up test

run-time test

5.6 Notes

- millisecond interrupt could be provided from a centralised C/T unit? But will all processors want to receive the same time interrupts?
- extra interrupt(s) for packet parity failure?

6. DMA controller

Each ARM968 processing subsystem includes a DMA controller. The DMA controller is primarily used for transferring inter-neural connection data from the SDRAM in large blocks in response to an input event arriving at a fascicle processor, and for returning updated connection data during learning. In addition, the DMA controller provides access to other targets on the System NoC such as the System RAM and Router configuration target.

6.1 Features

- Multithreaded DMA engine supporting parallel operations
- Support for both DMA requests and direct pass-through CHAIN requests from ARM
- Dual buffers supporting simultaneous direct and DMA transfers
- Dual clock-domain FIFOs
- Support for error control in transferred blocks
- Interrupt-driven or polled DMA completion notification
- Parameterisable buffer size
- Direct and DMA request queueing
- Support for SDRAM address paging

6.2 Using the DMA Controller

There are 2 types of requests for DMA controller services. DMA requests use the full engine, are initiated by writing to control registers in the controller, execute in background, and issue an interrupt while complete - so that the ARM processor may set up a request and then return to other processing while the requests complete. Non-DMA requests go through the passthrough channel and are initiated by a request directly to the needed device or service. The DMA controller fulfills these requests transparently, the host processor retaining full control of the transfer. Invisible to the user, the controller may buffer the data from write requests for more efficient bus management. However, the ARM processor must control the entire data transfer throughout the process.

The controller is transparent to non-DMA requests, acting as a bridge between the AHB-Lite bus on the ARM port and the {AXI} bus on the Chain port. We will briefly outline the request latencies for various types of transfer on this bus. No operation will take less than 2 bus cycles, because of the 2 interfaces the request must pass through. Added to this 2 clock latency are the following additional clock cycles:

DMA controller register read or write	0
Buffered atomic data (ARM-Chain) write	0
Unbuffered atomic data write	1
Atomic data write, channel busy, buffer empty	1
Atomic data write, channel busy, buffer full, no DMA transfer	2
Atomic data write, channel busy, buffer full, DMA transfer in progress	indeterminate (min 2, max. burst_size +2)
Atomic data (ARM-Chain) read	1

Atomic data read, channel busy, no DMA transfer	indeterminate (min 2, max request_buffer_size +2)
Atomic data read, channel busy DMA transfer in progress	indeterminate (min 2, max burst_size+2)
Invalid request on the ARM (AHB slave) interface	0
Invalid request on the TCM(AHB master) interface:	1
Requests ending in an error on the Chain bus	2

Read requests while a DMA transfer is in progress require special handling. Such a request may have large latencies. Worst-case latency if the only outstanding request is the DMA transfer is the burst size plus 2 cycles. If, meanwhile, buffered write requests exist in the non-DMA queue, the read must wait until all these requests have completed (to ensure data coherency) before it can complete. The recommended procedure is for the ARM processor to interrogate the transfer_started bit in the DMA_Status register before requesting a non-DMA read.

To initiate a DMA transfer, the ARM must write to the following registers in the DMA controller: Control (CTRL), CHAIN Address (ADRC), TCM Address (ADRT), and Length (LEN). The processor may also optionally write the following register to set up specific transfer parameters: Interrupt Control (IRQC) Options (OPTS), and SDRAM Memory Segment Base Address (BASE). The processor may read from any register at any time. Once a transfer has started (as the transfer_started bit in the Status (STAT) register indicates), the processor may queue another request. The processor may have a maximum of 2 outstanding requests of which only one will be active. For the active request, the processor may only write to certain specific bits in the control and interrupt control register. An attempt to write any other register during an active transfer will result in an error. The controller does not clear the registers after a transfer so that the processor may, if desired, modify only those fields that change between transfers in order to initiate a new DMA transfer. Writing the Start bit in the control register commits the currently set up DMA request. There will be at minimum 4 clocks latency from the point when the Start bit is set to the point where the DMA transfer physically starts. Maximum latency occurs if there is a current non-DMA transfer in progress, in which case it could be as high as 20 (for a 16-word non-DMA burst started just as the Start bit was set). Data received on the source bus will appear on the destination bus 2 clock cycles later.


6.3 Register summary

Name	Offset	R/W	Function
CHAIN Address (ADRC)	0x00	R/W	DMA address on the CHAIN bus
TCM Address (ADRT)	0x04	R/W	DMA address on the TCM interface
Length (LEN)	0x08	R/W	Length of the transfer in bytes
Control (CTRL)	0x0C	R/W	Control DMA transfer
Status (STAT)	0x10	R	Status of DMA and other transfers
Interrupt Control (IRQC)	0x14	W	Set interrupt configuration
Interrupt Status (IRQS)	0x18	R	Current interrupt status
Options (OPTS)	0x1C	R/W	Set less-frequently changed options
SDRAM Segment Base (BASE)	0x20	R/W	Base address of Fascicle page in SDRAM
Bytes Transferred	0x24	R*	Number of bytes transferred, current DMA
CHAIN Address 2 (AD2C)	0x100	R*	Double-buffered CHAIN bus address
TCM Address 2 (AD2T)	0x104	R*	Double-buffered TCM bus address
Length 2 (LN2)	0x108	R*	Double-buffered length
Control 2 (CTL2)	0x10C	R*	Double-buffered control
Options 2 (OPT2)	0x11C	R*	Double-buffered options

* These double-buffer registers are automatically written to if there is an active DMA transfer by writing to the addresses of their corresponding primary registers


6.4 Register details

Offset 0x00: CHAIN Address.




Bit 31:0 - Starting address on the CHAIN interface (AHB Master). Note that a read is considered a data movement from a source on the CHAIN bus to a destination on the TCM bus. If the `RELATIVE_ADDRESS` bit is set in the Options register, then this address is an offset into the SDRAM address whose base value is in `BASE`. This is useful if the ARM (probably acting as a fascicle) will be performing multiple transfers to the same area of SDRAM.

Offset 0x08 Length.



Bit 15:0 - Length of the DMA transfer. The TCM as currently envisioned has a maximum size of 64k (for the data TCM), thus, the length is a 16-bit register. A DMA transfer must of necessity either take as a source or a destination the TCM, justifying this restriction. See also the Request Size field in `CTRL`.

Offset 0x0C: ControlRegister



Bit 0 - 0 Enable. When set to 1 this enables the DMA interface. Pass-through functionality is always available.

Bit 1 - Start. Setting this bit starts a DMA transfer. The CPU should have written all necessary control registers before this bit is set. Once set, the DMA transfer cannot be modified save by an abort, error, or restart command.

Bit 2 - Abort. End current transfer and clear the buffer. This command effectively wipes out the current transfer. The controller will wait for any pending bursts to complete but read data from such transactions will be discarded. This, the Enable bit, and the Restart bit, are the only bits of the control register the CPU can set when Start is 1.

Bit 3 - Restart. Resume a transfer that halted on an error or time-out. This is equivalent to setting the start bit with the transfer counter advanced to the position where the error occurred. Transfer will resume from that point. Note that the DMA does no checking to verify that any error-causing condition has been resolved: it is up to the CPU to correct the problem before resuming the transfer.

Bit 4 - Lock. Indicate to the DMA controller that it should initiate a locked transaction on the AHB Master interface. Normally this will be reserved for control-oriented transfers to the System RAM from the Monitor Processor, if at all.

Bit 5 - Direction. Indicate the direction of the transfer: 0 => read from SDRAM (CHAIN interface) 1 => write to SDRAM (CHAIN interface).

Bit 7:6 - Burst size in words. 00 = 1 word (single transfer) 01 = 4 words, 10 = 8 words, 11 = 16

words.

Bit 9:8 - Word width. 00 = 8-bit, 01 = 16-bit, 10 = 32-bit, 11 = 64-bit.

Bit 11:10 - Priority. Indicates how forcefully the DMA controller should request access to system resources.

00: No priority: the DMA controller accepts bus grants whenever they occur and does not schedule DMA transfers ahead of non-DMA transfers.

01: Low priority. The DMA controller indicates to the bus grant mechanism that it should favour this transaction over any pending 0-priority requests but otherwise does not reschedule requests.

10: Medium priority. The DMA controller does not place the transfer ahead of pending non-DMA transfers but any new non-DMA transfer will be queued behind it. It also requests the CHAIN arbiter to put this request ahead of all lower requests.

11: High priority. The controller schedules the request ahead of all pending transfers on its interface and requests the CHAIN arbiter to service it first.

Bit 15:12 - Timeout period (clocks) Specifies how long a pending transfer may wait for a bus response before timing out. Timeout period is = 2^(TO +1) clk.

Bit 31:16 - Request Size. Indicates the maximum size, in bytes that the DMA controller should treat as a single request. The controller will truncate longer requests and the second part will have to be resubmitted. This field is useful e.g. for protecting ITCM writes against wraparound that overwrites the bottom area of the ITCM (usually the boot area)

Offset 0x10: Status Register.

31	30	29	28	27	26	25	24	23	22	21	20	19	18	17	16	15	14	13	12	11	10	9	8	7	6	5	4	3	2	1	0
(Reserved)												Error code					Queue pos	A	T	W	S	C									

Bit 0 - Transfer Completed. Indicates that the current DMA transfer has finished and that the CPU may submit another. The CPU may need to examine the Error field for error messages.

Bit 1 - Transfer Started. Indicates that a transfer is underway. No DMA requests may be submitted. Non-DMA requests will be subject to wait states until the current transfer completes. As noted, the recommendation is that the CPU always interrogate this bit before a read operation.

Bit 2 - Waiting. Indicates a DMA transfer has been started - i.e. the Start bit in the Control Register has been set, but that the CHAIN arbiter has not yet granted access to the bus to the DMA controller.

Bit 3 - Timed out. Indicates the DMA controller timed out (after Timeout_period clocks) on a request.

Bit 4 - Token acquired. Indicates the DMA controller has been granted access to the CHAIN bus. If Transfer Started is not set then either the current transfer will begin immediately or will end immediately (This bit is mostly for internal use)


Bit 8:5 - Queue Position. Indicates where in the CHAIN request queue the current DMA request resides. The CPU can use this to estimate when the controller will be able to service its request.

Bit 19:9 - Error Code as follows

- 0x001 Illegal register modify attempt while transfer underway
- 0x002 Non-DMA buffer full
- 0x004 Non-DMA request buffer corrupted
- 0x008 Invalid address

0x010 Checksum Error

Offset 0x14 Interrupt Control.



Bit 0 - Enable Fast Interrupt on Completion. If this bit is set, the interrupt controller will signal the ARM via a Fast Interrupt that the data is ready for access.

Bit 1 - Enable Fast Interrupt on Error. If this bit is set, the interrupt controller will signal the ARM via a Fast Interrupt if an error occurs during a transfer (of any type)


Bit 2 - Enable Fast Interrupt on Timeout. If this bit is set, the interrupt controller will signal the ARM via a Fast Interrupt if a timeout occurs during an attempted transfer.

Bit 3 - Enable Interrupt on Completion. If this bit is set, the interrupt controller will signal the ARM via a Normal Interrupt that the data is ready for access.

Bit 4 - Enable Interrupt on Error. If this bit is set, the interrupt controller will signal the ARM via a Normal Interrupt if an error occurs during a transfer (of any type)

Bit 5 - Enable Interrupt on Timeout. If this bit is set, the interrupt controller will signal the ARM via a Normal Interrupt if a timeout occurs during an attempted transfer.

Offset 0x18 Interrupt Status.



Bit 0 - Fast Interrupt signalled on Completion. If this bit is set, a Fast Interrupt has been issued following successful completion of a DMA transfer. Setting the START bit in CTRL clears this value.

Bit 1 - Fast Interrupt signalled on Error. If this bit is set, an error occurred and the controller issued a Fast Interrupt to inform the ARM. Setting either the START or RESTART bits in CTRL clears this value.


Bit 2 - Fast Interrupt signalled on Timeout. If this bit is set, a timeout occurred during an attempted transfer and the controller issued a Fast Interrupt to inform the ARM. Setting either the START or RESTART bits as above clears this value.

Bit 3 - Interrupt signalled on Completion. If this bit is set, an interrupt has been issued following successful completion of a DMA transfer. Setting the START bit in CTRL clears this value.

Bit 4 - Interrupt signalled on Error. If this bit is set, an error occurred and the controller issued an interrupt to inform the ARM. Setting either the START or RESTART bits as above clears this value.

Bit 5 - Interrupt signalled on Timeout. If this bit is set, a timeout occurred during an attempted transfer and the controller issued an interrupt to inform the ARM. Setting either the START or RESTART bits as above clears this value.

Offset 0x1C Options.



Bit 0 - PREFERRED_BUFFER Indicates which of the 2 buffers the controller should select for a DMA transfer if both buffers are available and DOUBLE_BUFFER is not set. Passthrough requests will use the other buffer.

Bit 1 - DOUBLE_BUFFER Setting this option doubles the size of the data FIFO buffer for a DMA transfer. If this bit is set, no other transfers may happen in background. Thus, with DOUBLE_BUFFER set all transfers on the DMA controller are blocking transfers with respect to the local ARM.

Bit 2 - RELATIVE_ADDRESS Instructs the DMA controller to use the address in BASE as a starting address to which ADRC is added as an offset, in order to calculate the starting address for the CHAIN side of the interface. Recommended only for Fascicle SDRAM access.


Bit 4:3 - ERROR_DETECT[1:0] Enables detection of errors via checksum according to the following coding:

- 00 - no error detection
- 01 - generate checksums
- 10 - examine checksums
- 11 - both generate and examine checksums.

The controller evaluates the checksum at the end of a DMA transfer as the 2's complement sum of the data words in the transfer.


Bit 6:5 - DISABLE_BUFFER[1:0] This pair of bits disable each of the 2 buffers, the MSB disabling Buffer 1 and the LSB Buffer 0. If a buffer is disabled all requests will pass through the other buffer regardless of the setting of DOUBLE_BUFFER and PREFERRED_BUFFER. Disabling a buffer means that each access is blocking with respect to other requests. Disabling both buffers effectively pauses the DMA controller: the ARM can still issue new requests, but these will not be serviced until at least one buffer is enabled. The primary purpose of this field is for fault tolerance.

Offset 0x20 SDRAM Segment Base.



Bit 31:0 - Base address in memory of the SDRAM segment allocated to the local processor. This assumes a memory- management scheme where the Monitor processor allocates to each Fascicle a private area of SDRAM for its own use. Using this value with relative addressing and request size constraints ensures that no Fascicle can corrupt the other's memory. This is useful for fault tolerance as well as performance because it can isolate faulty processors and prevent them from corrupting the whole SDRAM's memory map.

Offset 0x24 Bytes Transferred



Bit 15:0 - Number of bytes transferred in the current DMA request. This register is the actual counter itself so that reading it will always give an accurate count of the progress of the DMA transfer.

6.5 Fault-tolerance

Fault insertion

Fault detection

Fault isolation

Reconfiguration

6.6 Test

production test

start-up test

run-time test

7. Communications controller

Each processor node on SpiNNaker includes a communications controller which is responsible for generating and receiving packets to and from the communications network.


7.1 Features

- Support for 3 packet types:
 - multicast neural event packets routed by a key provided at the source;
 - point-to-point packets routed algorithmically by destination address;
 - nearest-neighbour packets routed algorithmically by arrival port.
- Packets are either 40 or 72 bits long. The longer packets carry a 32-bit payload.
- 2-bit time stamp (used by Routers to trap errant packets).
- Parity (to detect corrupt packets).


7.2 Packet formats

Neural event multicast (mc) packets (type 0)


Neural event packets include a 32-bit routing key inserted by the source, and a control byte:



In addition they may include an optional (not normally used) 32-bit payload:




The 8-bit control field includes packet type (= 00 for multicast packets), emergency routing and time stamp information, a data payload indicator, and error detection (parity) information:




Point-to-point (p2p) packets (type 1)

Point-to-point packets include 16-bit source and destination chip IDs, plus a control byte and an optional (normally used) 32-bit payload:




Here the 8-bit control field includes packet type (=01 for p2p packets), a sequence code, time

stamp, a data payload indicator and error detection (parity) information:




Nearest-neighbour (nn) packets (type 2)

Nearest-neighbour packets include a 32-bit address or operation field, plus a control byte and an optional 32-bit payload:



Here the 8-bit control field includes packet type (= 10 for nn packets), routing information, a data payload indicator and error detection (parity) information:



7.3 Control byte summary

Field Name	bits	Function
parity	0	parity of complete packet (including payload when used)
data	1	data payload (1) or no data payload (0)
time stamp	3:2	phase marker indicating time packet was launched
seq code	5:4	p2p only: start, middle odd/even, end of payload
emergency routing	5:4	mc only: used to control routing around a failed link
route	4:2	nn only: information for the Router
T: nn packet type	5	nn only: packet type - normal (0) or direct (1)
packet type	7:6	= 00 for mc; = 01 for p2p; = 10 for nn

parity

The complete packet (including the data payload where used) will have odd parity.

data

Indicates whether the packet has a 32-bit data payload (=1) or not (=0).

time stamp

The system has a global time phase that cycles through 00 -> 01 -> 11 -> 10 -> 00. Global

synchronisation must be accurate to within less than one time phase (the duration of which is programmable and may be dynamically variable). A packet is launched with a time stamp equal to the current time phase, and if a Router finds a packet that is two time phases old (time now XOR time launched = 11) it will drop it to the local Monitor Processor. The time stamp is inserted by the local Router, so the Communication Controller need do nothing here.

seq code

p2p packets use these bits to indicate the sequence of data payloads:

- 11 -> start packet: the first packet in a sequence (of >1 packets)
- 10 -> middle even: the second, fourth, sixth, ... packet in a sequence
- 01 -> middle odd: the third, fifth, seventh, ... packet in a sequence
- 00 -> end: the last (or only) packet in a sequence

emergency routing

mc packets use these bits to control emergency routing around a failed or congested link:

- 00 -> normal mc packet;
- 01 -> the packet has been redirected by the previous Router through an emergency route along with a normal copy of the packet. The receiving Router should treat this as a combined normal plus emergency packet.
- 10 -> the packet has been redirected by the previous Router through an emergency route which would not be used for a normal packet.
- 11 -> this emergency packet is reverting to its normal route.

route

These bits are set at packet launch to the values defined in the control register. They enable a packet to be directed to a particular neighbour (0 - 5), to all neighbours (6), or to the local Monitor Processor (7).

T (nn packet type)

This bit specifies whether an nn packet is 'normal', so that it is delivered to the Monitor Processor on the neighbouring chip(s), or 'direct', so that performs a read or write access to the neighbouring chip's System NoC resource.

packet type

These bits indicate whether the packet is a multicast (00), point-to-point (01) or nearest-neighbour (10) packet. Packet type 11 is reserved for future use.

7.4 Register summary

Name	Offset	R/W	Function
r0: Tx control	0x0	R/W	Controls packet transmission
r1: Rx status	0x4	R/W	Indicates packet reception status
r2: send data	0x8	W	32-bit data for transmission
r3: send key	0xC	W	Send mc key/p2p dest ID & seq code
r4: receive data	0x10	R	32-bit received data
r5: receive key	0x14	R	Received mc key/p2p source ID & seq code
r6: reserved	0x18	-	-
r7: test	0x1C	R/W	Used for test purposes

A packet will contain data if r2 is written before r3; this can be performed using an ARM STM instruction.

7.5 Register details

r0: transmit control

31	30	29	28	27	26	25	24	23	22	21	20	19	18	17	16	15	14	13	12	11	10	9	8	7	6	5	4	3	2	1	0
I	E	T	F	V	Route	control byte										p2p source ID															

The functions of these fields are described in the table below:


Name	bits	R/W	Function
I: int. status	31	R	enabled Tx interrupt
E: int. enable	30	W	enable interrupt on Tx buffer full
T: Tx full	29	R	Tx buffer full
F: V int. enable	28	W	enable interrupt on Tx buffer overrun
V: Tx overrun	27	R/W	Tx buffer overrun (sticky)
Route	26:24	W	Set 'fake' route in packet
control byte	23:16	W	control byte of next sent packet
p2p source ID	15:0	W	16-bit chip source ID for p2p packets

The p2p source ID is expected to be configured once at start-up. The parity and sequence code fields of the control byte will be replaced by automatically-generated values when the packet is launched. The time stamp (where applicable) will be inserted by the local Router.

The transmit buffer full control is expected to be used, by polling or interrupt, to prevent buffer overrun. Transmit buffer overrun indicates packet loss and will remain set until explicitly cleared by writing 0 to bit 25.

The route field allows a packet to be sent by a processor to the router which appears to have come from one of the external links. Normally this field will be set to 7 (0b111) but can be set to a link number in the range 0 to 5 to achieve this.

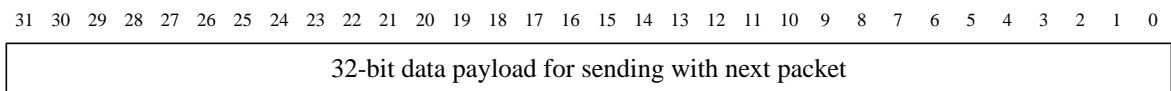
r1: receive status



The functions of these fields are described in the table below:

Name	bits	R/W	Function
I: int. status	31	R	interrupt caused by Rx packet
E: int. enable	30	W	enable interrupt on Rx packet
R: received	29	R	packet received
K: P int. enable	28	W	Rx packet parity error interrupt enable
P: parity	27	R/W	Rx packet parity error (sticky)
J: F int. enable	26	W	Rx packet framing error interrupt enable
F: framing error	25	R/W	Rx packet framing error (sticky)
U: unused	24	-	-
Control byte	23:16	R	Control byte of last Rx packet
U: unused	15:3	-	-
Route	2:0	R	Route field from packet

r2: send data




If data is written into r2 before a send key or dest ID is written into r3, the packet initiated by writing to r3 will include the contents of r2 as its data payload. If no data is written into r2 before a send key or dest ID is written into r3 the packet will carry no data payload.

r3: send mc key/p2p dest ID & sequence code


Writing to r3 will cause a packet to be issued (with a data payload if r2 was written previously).

If bits[23:22] of the control register are 00 the Communication Controller is set to send multicast packets and a 32-bit routing key should be written into r3. The 32-bit routing key is used by the associative multicast Routers to deliver the packet to the appropriate destination(s).




If bits[23:22] of the control register are 01 the Communication Controller is set to send point-to-

point packets and the value written into r3 should include the 16-bit address of the destination chip in bits[15:0] and a sequence code in bits[17:16]. (See 'seq code' on page 24.)



If bits[23:22] of the control register are 10 the Communication Controller is set to send nearest neighbour packets and the 32-bit nn address field should be written in r3.

r4: received data



If a received packet carries a data payload the payload will be delivered here and will remain valid until r5 is read.

r5: received mc key/p2p source ID & sequence code

A received packet will deliver its mc routing key, nn address or p2p source ID and sequence code to r5. For an mc or nn packet this will be the exact value that the sender placed into its r3 for transmission; for a p2p packet the sequence number will be that placed by the sender into its r3, and the 16-bit source ID will be that in the sender's r0.

The register is read sensitive - once read it will change as soon as the next packet arrives.

r6: reserved

This register is reserved for future use.

r7: test

Setting bit 0 of this register makes all registers read/write for test purposes. Clearing bit 0 restricts write access to those register bits marked as read-only in this datasheet. All register bits may be read at any time.

7.6 Fault-tolerance

Fault insertion

Software can cause the Communications Controller to misbehave in several ways including inserting dodgy routing keys, source IDs, destination IDs.

Do we need to be able to force parity errors in transmit packets?

Fault detection

Parity of received packet; received packet framing error; transmit buffer overrun.

Fault isolation

The Communications Controller is mission-critical to the local processing subsystem, so if it fails the subsystem should be disabled and isolated.

Reconfiguration

The local processing subsystem is shut down and its functions migrated to another subsystem on this or another chip. It should be possible to recover all of the subsystem state and to migrate it, via the SDRAM, to a functional alternative.

7.7 Test

production test

start-up test

run-time test

7.8 Notes

- time phase accuracy: if we assume that the system time phase is F and the skew is K (that is, all parts of the system transition from one phase to its successor within a time K), then a packet has at least $F-K$ to reach its destination and will be killed after at most $2F+K$. Thus, if we want to allow for a maximum packet transit time of $F-K = T$ and can achieve a minimum phase skew of K , then T and K are both system constants and we should choose $F = T+K$. The longest packet life is then $2T+3K$.

8. Communications NoC

The communications NoC has the primary role of carrying neural event packets between Fascicle Processors on the same or different chips.

8.1 Features

- On- and inter-chip links
- Router which handles multicast, point-to-point and nearest neighbour packets.
- Arbiter to merge all sources into a sequential packet stream into the Router.
- Individual links can be reset to clear blockages and deadlocks.

8.2 Block diagram

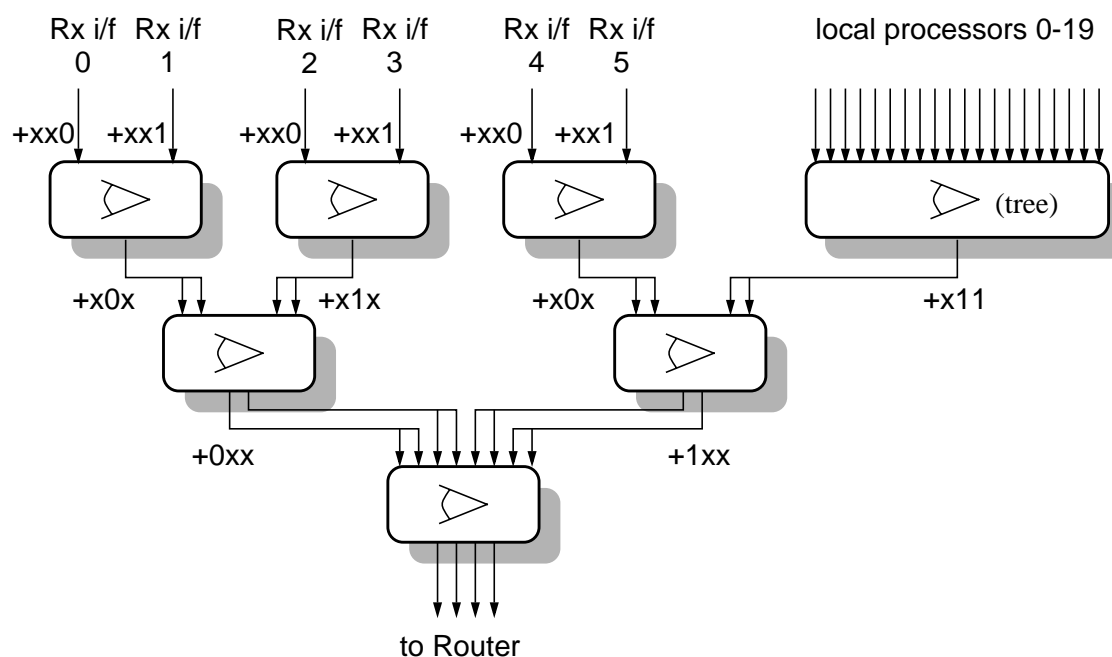
A block diagram of the Communications NoC was given in section 1.1 on page 5.

8.3 Arbiter structure

As the input links converge on the Router they must merge through 2-way CHAIN arbiters, and the link width must increase to absorb the bandwidth. The following hierarchy is proposed:

- the local processor links can all be merged through a single-link arbiter tree as the local bandwidth is low, e.g. at most 20 processors x 1,000 neurons x 100Hz x 40 bits = 80 Mbit/s.
- the Rx interfaces can each carry up to 1 Gbit/s, about half the on-chip single-link bandwidth, so the first layer of arbiters can be single-link, the 2nd layer dual-link and the 3rd layer quad-link (i.e. 8-bits or 48 wires wide).
- buffering is required wherever the link width increases to ensure that the full arbiter bandwidth is used. Each buffer must be at least half a packet long - 36 bits?
- at each arbiter merging Rx interfaces the packet must pick up 1 bit to indicate its source, for default routing [unless the source tagging is done by the Rx interface?]

The Arbiter structure is illustrated below. Each doubling of the wires represents a doubling of the CHAIN link width. The numbers indicate source tagging of the packets.



8.4 Fault-tolerance

Fault insertion

There is little direct control of the Communications NoC fabric except at the periphery as noted in the sections below.

Fault detection

Most failures will cause local asynchronous deadlock, which is readily detected at both the transmitting and receiving ends of the link.

Fault isolation

If links fail their clients will have to be disabled and their functions migrated.

Reconfiguration

Client functional migration is required.

8.5 Test

production test

start-up test

run-time test

8.6 Notes

- must decide whether to add source tags for default routing in arbiters or in Rx interfaces.

9. Communications Router

The Communications Router is responsible for routing all packets that arrive at its input to one or more of its outputs. Its primary function is to route multicast neural event packets, which it does through an associative multicast router subsystem. But it is also responsible for routing point-to-point packets (for which it uses algorithmic routing), for nearest-neighbour routing (which is a simple algorithmic process), for default routing (when a multicast packet does not match any entry in the multicast router) and for emergency routing (when an output link is blocked due to congestion or hardware failure).

Various error conditions are identified and handled by the Communications Router, for example packet parity errors, time-out, and output link failure.

9.1 Features

- 1024 programmable associative multicast routing entries.
 - associative routing based on source ‘key’.
 - with flexible “don’t care” masking.
- algorithmic routing of point-to-point and nearest-neighbour packets.
- support for 40- and 72-bit multicast, point-to-point and nearest neighbour packets.
- default routing of unmatched multicast packets.
- automatic re-routing around failed links.
- failure detection and handling:
 - packet parity error
 - time-expired packet
 - output link failure
 - corrupt (wrong length) packet

9.2 Description

We assume that messages arrive from other nodes via the link receiver interfaces and from internal clients and are presented to the router one-at-a-time. The Arbiter is responsible for determining the order of presentation of the messages, but as each message is handled independently the order is unimportant (though it is desirable for packets following the same route to stay in order).

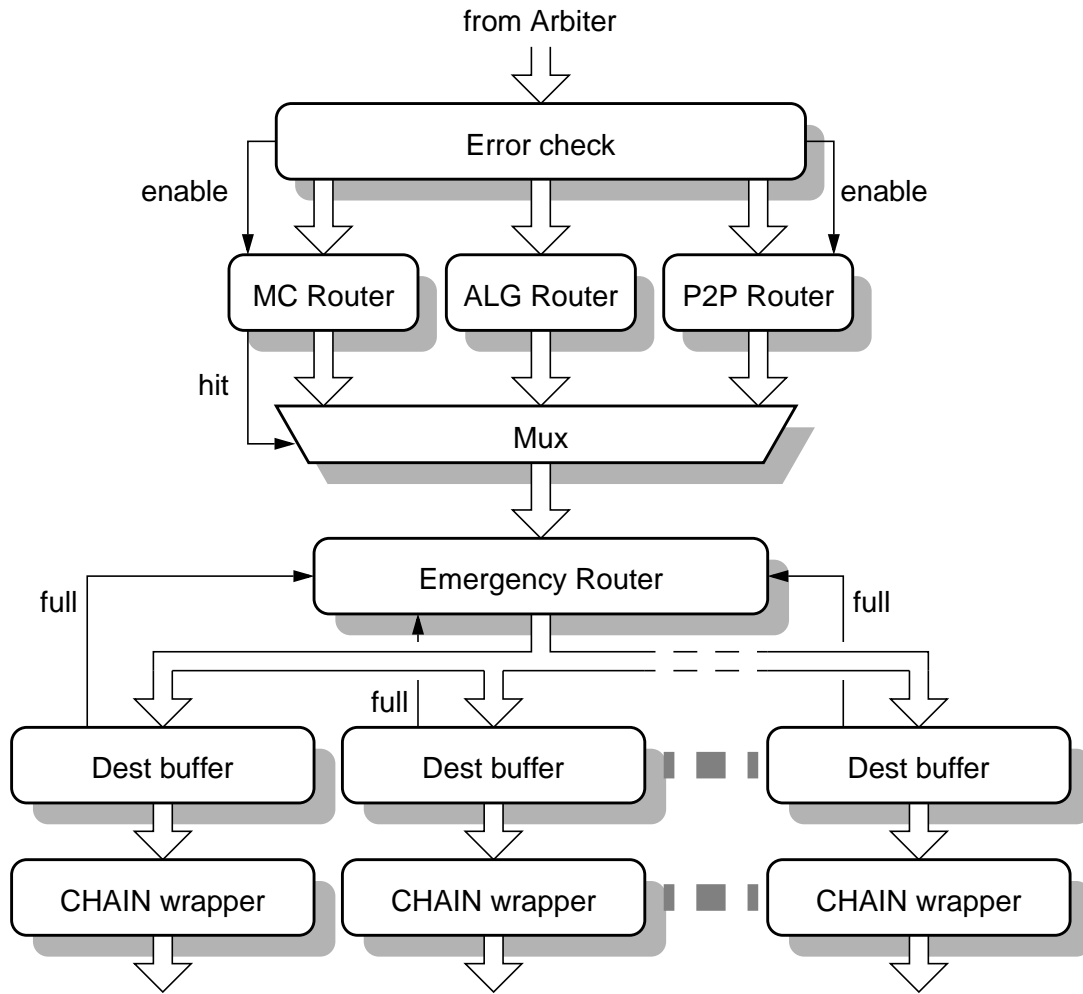
Each message contains an identifier that is used by the Communications Router to determine which of the outputs the message is sent to. These outputs may include any subset of the output links, where the message may be sent via the respective link transmitter interface, and/or any subset of the internal processor nodes, where the message is sent to the respective Communications Controller.

For the neural network application the identifier can be simply a number that uniquely identifies the source of the message – the neuron that generated the message by firing. This is ‘source address routing’. In this case the message need contain only this identifier, as a neural spike is an “event” where the only information is that the neuron has fired.

The Router then functions simply as a look-up table where for each identifier it looks up a routing word, where each routing word contains 1 bit for each destination (each link transmitter interface and each local processor) to indicate whether or not the message should be passed to that destination.

9.3 Internal organization

The internal organization of the Communications Router is illustrated in the figure opposite.



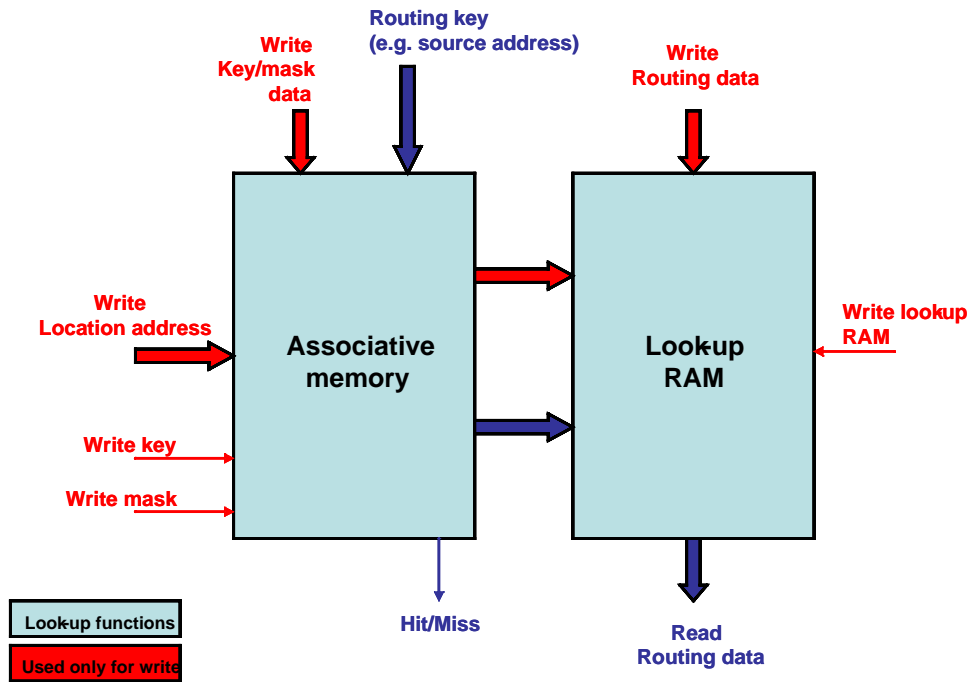
Packets are passed as complete 40- or 72-bit units from the Arbiter, together with an identifier of the Rx interface that the packet arrived through (for nearest-neighbour and default routing). The first stage of processing here is to identify errors. The second stage passes the packet to the appropriate routing engines – the multicast (MC) router is activated only if the packet is error-free and of multicast type, the point-to-point (P2P) handles point-to-point packets while the algorithmic (ALG) router handles nearest-neighbour packets and also deals with default and error routing. The output of the router stage is a vector of destinations to which the packet should be relayed. The third stage is the emergency routing mechanism for handling failed or congested links, which it detects using ‘full’ signals fed back from the individual destination output buffers.

Notes

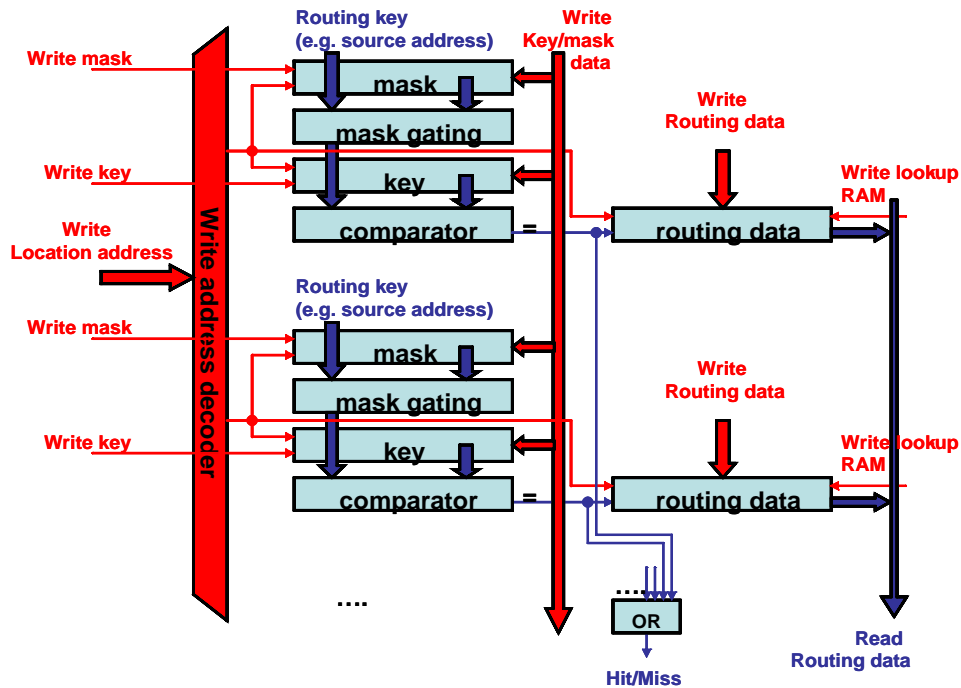
- the Router needs to know which is the Monitor Processor for routing terminating p2p, nn, error, and dropped packets
- how are details of errors communicated to the Monitor Processor?
- Emergency routing may cause two packets to be issued onto the same output link (one for the MC routed data and the second for the alternative route for the blocked link). These are merged into a single packet with a different emergency-routing type to indicate its dual purpose.

9.4 Multicast (MC) router

The internal organisation of the multicast router is illustrated in the figure below.



Implementation



Multicast router optimisations

The simple look-up table as described above works in principle but in practice would be too large to fit onto a chip. However, several optimisations address this problem, reducing the table to a practical size:

- The table within a particular node need contain entries only for message identifiers whose routes from their source to all of their destinations pass through, are generated in, or end in that node.

- Default routing can be supported that passes messages from a link receiver interface to the diametrically opposite link transmitter interface when the message identifier is not found in the look-up table.
- Groups of message identifiers can be routed using the same look-up table entry by making some of the identifier bits “don’t care” as far as the look-up process is concerned.

The logical structure of the Router is then an associative (content-addressed) memory, with programmable masking on a per-entry basis to support the “don’t care” optimization, connected to a conventional memory that holds the per-entry output routing word. The associative memory can be implemented using any of the usual techniques such as VLSI CAM cells or hash-addressed RAM.

Additional mechanisms can support the default routing mentioned above when there is no match for the message identifier in the associative memory and, through partitioning the identifier address space, provision can be made for conventional 1-to-1 destination address routing and broadcast mechanisms.

Illustrative example

By way of illustration, let us assume a neural modelling system where each neuron has a unique 14-bit number, and when it fires it transmits this number prepended by two zero bits as the message identifier (and the message has no other content). Further, we assume that the neurons are handled in groups of 256, where each group of 256 is assigned to a particular processor on a particular node and each neuron in a group has the same set of inputs and the same output destinations as every other neuron in the same group.

The message identifiers can then be processed by the Routers as the 16-bit binary number: 00nnnnnnXXXXXXXXXX, where:

- 00 indicates that this is a source address message identifier that should be routed according to the routing table;
- nnnnnn is the neuron group identifier, so at most $2^6 = 64$ routing look-up table entries are required;
- XXXXXXXX indicates that the bottom 8 bits of the message identifier can be treated as “don’t care” and play no role in the routing. They will be used only by the destination processor to identify the neuron that fired within the group.

Messages beginning with 01, 10, and 11 may be used for 1-to-1 destination address routing, broadcast, and some other purpose respectively, using conventional routing algorithms.

It can be seen from this example that the optimisations are a vital aspect of the invention, reducing the size of the look-up table in this very small example from 16,384 ($= 2^{14}$) to at most 64 entries. For larger systems the benefits of the optimisations are likely to be even more significant.

The “don’t care” bits are programmable independently for each look-up table entry on each node, and they can be distributed anywhere across the message identifier for maximum flexibility, so a single look-up table entry may route several groups together at one node, whereas at the next node they may be routed independently to different destinations.

Route set up

The routing look-up tables may be configured using external software that takes a neural “netlist”, describing the way the neurons interconnect, and then maps the neurons onto processors and determines the routing table values using algorithms similar to those used to configure an FPGA. As with FPGA configuration, resource constraints such as the routing table size and link bandwidth limitations must be taken into account during the mapping process.

The routing table configuration is then loaded into the local Router by a local processor that follows instructions from a control system using the 1-to-1 message routing mechanism.

For static neural network modelling the routing is fixed after initialisation. It is possible to allow

local processors to modify the routing tables while the system is running, if this is required to model developmental processes (for example), provided this is done with due care.

9.5 The point-to-point (p2p) router

The p2p router uses the 16-bit destination ID in a point-to-point packet to determine which output(s) the packet should be routed to. A 64K entry x 8-bit SRAM lookup table directs the p2p packet to:

- the local Monitor Processor, and/or
- adjacent chips via the appropriate links.

Each 8-bit entry has one bit which determines whether the packet is delivered to the local Monitor Processor, one bit for each of the six output links, plus a parity bit. Thus there is a form of broadcast capability available here.

9.6 The algorithmic (ALG) router

nn routing

Nearest-neighbour packets are used to initialise the system and to perform run-time flood-fill and debug functions. The routing function here is to send 'normal' nn packets that arrive from outside the node (i.e. via an Rx link) to the monitor processor and to send nn packets that are generated internally to the appropriate output (Tx) link(s). This is to support a flood-fill OS load process.

In addition, the 'direct' form of nn packet can be used by neighbouring systems to access System NoC resources. Here an nn 'write' packet (which is a direct type with a 32-bit payload) is used to write the 32-bit data defined in the payload to a 32-bit address defined in the address/operation field. An nn 'read' packet (which is a direct type without a 32-bit payload) uses the 32-bit address defined in the address/operation field to read from the System NoC and returns the result (as a 'normal' nn packet) to the neighbour that issued the original packet using the Rx link ID to identify that source. This 'direct' access to a neighbouring chip's principal resources can be used to investigate a non-functional chip, to re-assign the Monitor Processor from outside, and generally to get good visibility into a chip for test and debug purposes.

default and error routing

In addition, the algorithmic router performs default and error routing functions.

9.7 Time phase handling

The Router maintains a 2-bit time phase signal that is used to delete packets that are out-of date. The time phase logic operates as follows:

- locally-generated packets will have the current time phase inserted (where appropriate);
- a packet arriving from off-chip will have its time phase checked, and if it is two phases old it will be deleted (dropped to the local Monitor Processor).

9.8 Packet error handler

The packet error handler is a routing engine that simply flags the packet for routing to the local Monitor Processor if it detects any of the following:

- a packet parity error;
- a packet that is two time phases old;
- a packet that is the wrong length.

There must be a means for the Monitor Processor to recognise packets passed to it with errors. Rather than complicating the Communications Controller, this is probably better done by providing

error information via the Router configuration registers.

9.9 Emergency routing

If a link fails (temporarily, due to congestion, or permanently, due to component failure) action will be taken at two levels:

- the blocked link will be detected in hardware and subsequent packets rerouted via the other two sides of one of the routing triangles of which the suspect link was an edge.
- the Monitor Processor will be informed. It will assess the problem, and take appropriate action:
 - if the problem was due to transient congestion, it will note the congestion but do nothing further;
 - if the problem was due to recurring congestion, it will negotiate and establish a new route for some of the traffic using this link;
 - if the problem appears permanent, it will reset the link (incurring some packet loss) and then, if this does not clear the problem, negotiate and establish new routes for all of the traffic using this link.

The hardware support for these processes include:

- default routing processes in adjacent nodes that are invoked by flagging the packet as an emergency type;
- mechanisms to inform the Monitor Processor of the problem;
- mechanisms the Monitor Processor can use to reset the link;
- means of inducing the various types of fault for testing purposes.

Emergency rerouting around the triangle requires additional emergency packet types for mc packets. p2p packets will find their own way to their destination following emergency routing.

9.10 Errant packets

In order to ensure that packets cannot circulate for ever within the system each packet includes a time phase field. This is set when the packet is launched, and if a packet arrives at a Router two time phases after it was launched it will be routed directly (and only) to the local Monitor Processor for error-handling purposes.

9.11 Pseudo-code description

The following pseudo-code describes the detailed operation of the Communications Router:

```
Pipeline stage 1: Error Checking
inputs:      72-bit Packet  p;
             3-bit  SourceID src;
local info:  2-bit TimePhase timePhase

Begin stage 1=====

% check error conditions

PPerr = (packetParity(p) == EVEN);           % parity error
TPerr = (src < 6) AND (p.timeStamp = timePhase EOR 0b11)
        AND (p.type == 0b0x);              % time phase error
LNerr = (p.lastSymbol != EOP);              % packet length error

error = PPerr OR TPerr OR LNerr;

% update counters

if (PPerr) incPacketParityErrorCounter();
if (TPerr) incPacketTimeStampErrorCounter();
if (LNerr) incPacketLengthErrorCounter();
```

```

incPacketCounter();

% insert Time Phase

if (src == 7) AND (p.type == 0b0x) {           % local p2p or mc packet
    p.timeStamp = timePhase;
    ParityFix(p.parity);
}

% engage appropriate Router

enMC = (not error) AND (p.type == MC) AND (p.emergencyRouting != 0b10);
enP2P = (not error) AND (p.type == P2P);

End stage 1=====

Pipeline stage 2: Routing
inputs:      72-bit Packet    p;
             3-bit SourceID src;
             Booleans PPerr, TPerr, LNerr, error;
local info: 5-bit MonitorProcessorID mpID;

Begin stage 2=====

% enable relevant Router

if (enMC) {hit, MCvect} = MCrouter(p.MCkey);
else hit = 0;
if (enP2P) {P2Pvect}    = P2Prouter(p.destID);

% default emergency routing vector

erVect = 0;

% send all errors to Monitor Processor

if (error) vect = 2^(mpID+6);
else {

% routing depends on packet type

case (p.type) {

MC: if (hit) vect = MCvect;
    else if (src == 7)
        vect = 2^(mpID+6)           % local: miss => error
    else if (p.emergencyRouting == 0b0x)
        vect = 2^[(src+3)mod6];    % normal default
    else if (p.emergencyRouting == 0b11)
        vect = 2^[(src+2)mod6];    % ER 2nd stage default
    else vect = 0;                  % ER only
    if (p.emergencyRouting == 0b01 or 0b10)
        erVect = 2^[(src-1)mod6]; % ER 1st stage

P2P:      vect = P2Pvect;

NN: if (src == 7) {                 % local source
    if (p.route == 7)               % local MP
        vect = 2^(mpID+6);
    else if (p.route == 6)          % all neighbours
        vect = 0b00000000000000000000111111;
    else vect = 2^(packet.route);  % one neighbour
} else {                             % external source
    if (p.T) {                       % direct NN
        if (p.data) write.SystemNoC(p.address,p.data);
        vect = 0;                    % packet goes nowhere
    } else {
        p.payload = read.SystemNoC(p.address);
        p.data    = TRUE;
    }
}
}

```

```

        p.T          = 0;                % change to normal
        ParityFix(p.parity);
        vect = 2^src;                    % return to sender!
    } else vect = 2^(mpID+6);           % normal NN
    }
} % end case
} % end else (Â-error)

End stage 2=====

Pipeline stage 3: Emergency Routing
inputs:    72-bit Packet  p;
           26-bit Vector  vect;
           6-bit  ERvector erVect;
local info: Booleans buffFull bFull[0..25];
           5-bit MonitorProcessorID mpID;

Begin stage 3=====

% check for output contention & wait fixed max time to resolve

clockCycles = 0;
do {
    blocked = FALSE;
    for (i = 0; i++; i<6) {
        if (bFull[i] AND (vect.bit[i] OR erVect.bit[i])) blocked = TRUE;
    }
    for (i = 6; i++; i<26) {
        if (bFull[i] AND vect.bit[i]) blocked = TRUE;
    }
    clockCycles++;
} while (blocked AND (clockCycles < MaxWaitBeforeER));

% now look into Emergency Routing options & wait fixed max time

clockCycles = 0;
do {
    blocked = FALSE;
    for (i = 0; i++; i<6) {
        if (bFull[i] AND ((vect.bit[i] AND (bFull[(i-1)mod6] OR (p.type ==
NN))) % check if these may be ER'd
                                OR erVect.bit[i])) %
these are not ER'd
        blocked = TRUE;
    }
    for (i = 6; i++; i<26) {
        if (bFull[i] AND vect.bit[i]) %
these are also not ER'd
        blocked = TRUE;
    }
    clockCycles++;
} while (blocked AND (clockCycles < MaxWaitForER));

% if Emergency Routing has failed...

if (blocked) {
    if (bFull[mpID+6]) {
        discardPacket();                % throw away packet
        incDiscardedPacketCounter();    % record packet loss
        doSomethingWithMonitorProc();   % try to revive monitor?
    } else {
        sendPacketTo(p, buff[mpID+6]);  % send to Monitor Proc
        incDroppedPacketCounter();      % record packet loss
        sendToSystemController(bFull);  % report blocked links to system
    }
}
}

% can now proceed

```



```

for (i = 0; i++; i<6) {
  if (NOT bFull[i]) {
    p2 = p;
    case (p.type) {
      MC: if (vect.bit[i] OR erVect.bit[i]
             OR (bFull[(i+1)mod6] AND vect.bit[(i+1)mod6])) {
          if (vect.bit[i]) {
            if (bFull[(i+1)mod6] AND vect.bit[(i+1)mod6])
              p2.emergencyRouting = 0b01;    % normal + ER 1st
              incERpacketCounter();          % record ER
            else p2.emergencyRouting = 0b00;    % normal
          } elseif (bFull[(i+1)mod6] AND vect.bit[(i+1)mod6]) {
            p2.emergencyRouting = 0b10;    % ER 1st stage
            incERpacketCounter();          % record ER
          } elseif (erVect.bit[i]) {
            p2.emergencyRouting = 0b11;    % ER 2nd stage
          }
          ParityFix(p2.parity);
          sendPacketTo(p2, buff[i]);
        }

      P2P: if (vect.bit[i]
              OR (bFull[(i+1)mod6] AND vect.bit[(i+1)mod6])) {
          sendPacketTo(p2, buff[i]);
        }

      NN: if (vect.bit[i]) {
          sendPacketTo(p2, buff[i]);
        }
    }
  }
}
for (i = 6; i++; i<26) {
  if (vect.bit[i] AND NOT bFull[i]) {
    p2 = p;
    if (p2.type == NN) {
      p2.route = src;
    }
    sendPacketTo(p2, buff[i]);
  }
}

End stage 3=====

```

9.12 Registers

The Router configuration and error-reporting registers are detailed in section 14. on page 47.

9.13 Fault-tolerance

The Communications Router has limited fault-tolerance capacity, mainly coming down to mapping out a failed multicast router entry. This is a useful mechanism as the multicast router dominates the silicon area of the Communications Router.

Fault insertion

- enable Router to flip packet parity bits?

Fault detection

- packet parity errors
- packet time-phase errors
- packet unroutable errors (e.g. a locally-sourced multicast packet which doesn't match any entry in the multicast router).

- wrong packet length.

Fault isolation

- a mechanism is required to disable a multicast router entry if it fails. Possible just an 'entry valid' bit?

Reconfiguration

- since all multicast router entries are identical the function of any entry can be relocated to a spare entry (within the same segment of the router if segmentation is used to save power).
- if a router (segment) becomes full a global reallocation of resources can move functionality to a different router (segment)

9.14 Test

production test

start-up test

run-time test

9.15 Notes

- The Router will require a number of traffic monitor features, e.g. packet counters, congestion indicators, count packet under match & mask, dropped packet count, emergency routing count, count on each output link, ...

10. Inter-chip transmit and receive interfaces

Inter-chip communication is implemented by extending CHAIN links from chip to chip. In order to sustain CHAIN link throughput, there is a protocol conversion at each chip boundary from standard CHAIN 1-of-5 (including EOP) return-to-zero to 2-of-7 non-return-to-zero. Each conversion maps two 2-bit CHAIN symbols to a single 4-bit 2-of-7 symbol.

10.1 Features

- transmit (Tx) interface:
 - converts two on-chip 1-of-5 RTZ symbols into one off-chip 2-of-7 NRZ symbol;
 - control input to induce a fault;
 - failure detection output.
 - fault reset input.
- receive (Rx) interface:
 - converts one off-chip 2-of-7 NRZ symbol into two on-chip 1-of-5 RTZ symbols;
 - control input to induce a fault;
 - failure detection output.
 - fault reset input.
 - adds source tag to packet for default routing [unless this is done in the Communications NoC arbiter?]

10.2 Programmer view

There are no programmer-accessible features implemented in these interfaces. In normal operation these interfaces provide transparent connectivity between the routing network on one chip and those on its neighbours.

10.3 Fault-tolerance

The fault inducing, detecting and resetting functions are controlled from the System Controller (see 'System Controller' on page 45).

Fault insertion

- an input controlled by the System Controller causes the interface to deadlock

Fault detection

- an output to the System Controller indicates deadlock

Fault isolation

- the interface can be disabled to isolate the chip-to-chip link. This may be the same input from the System Controller that is used to insert a fault.

Reconfiguration

- the link interface can be reset by the System Controller to attempt recovery from a fault
- the link interface can be isolated and an alternative route used

10.4 Test

production test

start-up test

run-time test

11. System NoC

The System NoC has a primary function of connecting the Fascicle Processors to the SDRAM interface. It is also used to connect the Monitor Processor to system control and test functions, and for a variety of other purposes.

11.1 Features

- supports full bandwidth block transfers between the SDRAM and the Fascicle Processors.
- the Router is an additional initiator for system debug purposes.
- can be reset (in subsections?) to clear deadlocks.

11.2 Fault-tolerance

Fault insertion

Fault detection

Fault isolation

Reconfiguration

11.3 Test

production test

start-up test

run-time test

12. SDRAM interface

The SDRAM interface connects the System NoC to an off-chip SDRAM device. It will be the ARM PL340.

12.1 Features

- lots of bandwidth, please!

12.2 Fault-tolerance

Fault insertion

Fault detection

Fault isolation

Reconfiguration

12.3 Test

production test

start-up test

run-time test

13. System Controller

The System Controller incorporates a number of functions used by the Monitor Processor for system start-up, fault-tolerance testing (invoking, detecting and resetting faults), general performance monitoring, and such like. At present, the exact layout of the registers has not been defined and bits and functionality may move between registers as the design progresses. The descriptions below are therefore given in general terms.

To provide some tolerance to runaway code on any processor it is expected that access to these registers (and similar functionality elsewhere on the chip) will only be permitted in a privileged CPU mode. A control bit can be provided to disable this if necessary, so that user mode code can also access these resources. Access might also be restricted just to the monitor processor at certain times. A global watchdog timer is also provided here which the monitor processor is expected to refresh. Failure to do so results in the chip being reset (and a new monitor chosen? - how?)

13.1 Register summary

Name	Offset	R/W	Function
CPU Reset	0x0000	R/W	Each bit allows a processor to be reset
CPU Interrupt	0x0004	R/W	Each bit generates an interrupt at a processor
CPU Clk disable	0x0008	R/W	Each bit disables the clock of a processor
Monitor ID	0x0010	R/W	ID of monitor processor
Set CPU OK	0x0020	R/W	Writing a 1 sets a CPU OK bit
Clr CPU OK	0x0024	R/W	Writing a 1 clears a CPU OK bit
Reset Code	0x0030	RO	Indicates cause of last chip reset
I/O port	0x0040	R/W	Access to external I/O pins
Misc control	0x0050	R/W	Miscellaneous control bits
Misc status	0x0054	RO	Miscellaneous status bits
Misc test	0x0100	R/W	Miscellaneous chip test control bits
Watchdog	0x1000	Various	Base of ADK watchdog registers

13.2 Register details

CPU Reset, Interrupt, Clk disable

These three registers contain individual bit enables for each processor. They are all initialised to zero on reset. Setting a bit will either reset, interrupt or disable the clock of a particular processor. It is expected that only the monitor processor will access these registers.

Monitor ID

This register is written with the ID of the processor which has been chosen as the monitor processor. Its output is required by the router (via the Router Control Register) in order to route error packets to the monitor processor. It is initialised by power-on reset to an invalid value which does not refer to any processor. Other forms of reset do not change this register.

Set/Clr CPU OK

These two registers contain individual bits for each processor and are used to indicate which processors are operational. Writing a 1 to a bit in the Set register will set a bit and writing a 1 to a bit in the Clr register will clear a bit. This mechanism ensures that the state is updated atomically when several processors write these register 'simultaneously'. When read, the both registers return the same value, being a bit mask indicating the current set of operation processors. All bits are cleared when the chip is reset.

I/O port

The I/O port controls a small number of I/O pins on the chip. A minimum of two would allow a serial I/O interface such as I2C to be implemented. LEDs may also be driven by the I/O pins.

Misc control

This register contains a collection of bits which provide general chip control. The following functions may be provided - PLL control, bits to allow User mode access to various parts of the system, direction control for the I/O pins, etc, etc

Misc status

This register provides a collection of status bits. For example, reason for last reset, state of the off-chip links, 'one-shot' bit which reads as 1 only once following reset (used to assist in selecting monitor processor), PLL lock status, etc, etc

Test control

The test control register provides control for on-chip testing. For example, bits to simulate error conditions in various parts of the chip, bits to reconfigure pins to allow testing to proceed, etc, etc.

Watchdog

This is implemented using the watchdog component from the AMBA design kit. It must be periodically written (by the monitor processor) to prevent the watchdog timer from expiring and resetting the chip. NOTE: What is the timer period & do we want to try to force a different monitor processor when the watchdog expires.

13.3 Fault-tolerance

Fault insertion

Fault detection

Fault isolation

Reconfiguration

13.4 Test


production test

start-up test

run-time test

Name	bits	R/W	Function

register 1: Router error flags




The functions of these fields are described in the table below:

Name	bits	R/W	Function
LE[5:0]	5:0	R/W	Tx link transmit error (sticky)
FPE[19:0]	25:6	R/W	Fascicle Processor link error (sticky)
P	26	R/W	Rx packet parity error (sticky)
F	27	R/W	Rx packet framing error (sticky)
T	28	R/W	Rx packet time stamp error (sticky)
U: unused	29	-	-
E	30	W	Enable Router interrupt
I	31	R	Router interrupt active


The Tx and Fascicle Processor link error flags indicate that a packet has been dropped because it could not be sent through the respective link and, in the case of the Tx links, any attempt at emergency routing was unsuccessful. These flags, and the F and P flags, are ‘sticky’ - once set they remain set until explicitly cleared by the processor writing a 0 to the respective bit position in this register.

register 2: packet count



Register 2 maintains a count of the total number of packets that have passed through the Router. It may be initialized to any value (for example, zero) by the Monitor Processor by writing to the register, and it may be read at any time.

register 3: time stamp



14.4 Fault tolerance

Fault insertion

Fault detection

Fault isolation

Reconfiguration

14.5 Test

production test

start-up test

run-time test

14.6 Notes

- time stamp: writeable, updated on newer incoming packet and by internal counter?
- dropped packet counter, emergency routed counter, ...?

15. System RAM

The System RAM is an additional 128 kByte block of on-chip RAM used primarily by the Monitor Processor to enhance its program and data memory resources as it will be running more complex (though less time-critical) algorithms than the Fascicle Processors.

As the choice of Monitor Processor is made at start-up (and may change during run-time for fault-tolerance purposes) the System RAM is made available to whichever processor is Monitor Processor via the System NoC. It is probably important that accesses by the Monitor Processor to the System RAM are non-blocking as far as SDRAM accesses by the Fascicle Processors are concerned, so the System NoC should ensure this is the case.

The System RAM may also be used by the Fascicle Processors to communicate with the Monitor Processor and with each other, should the need arise.

15.1 Features

- 128 kB of SRAM, available via the System NoC.
- can be disabled to model complete failure for fault-tolerance testing.
- can we include parity or ECC to improve fault-tolerance?

15.2 Fault-tolerance

Fault insertion

- It is straightforward to corrupt the contents of the System RAM to model a soft error – any processor can do this. It is not clear how this would be detected.
- The System RAM can be disabled to model a total failure.

Fault detection

- The Monitor Processor may perform a System RAM test at start-up, and periodically thereafter.
- It is not clear how soft errors can be detected without some sort of parity or ECC system.

Fault isolation

- Faulty words in the System SRAM can be mapped out of use.

Reconfiguration

- For hard failure of a single bit, avoid using the word containing the failed bit.
- If the System RAM fails completely the only option is to use the SDRAM instead, which will probably result in compromised performance for the Fascicle Processors due to loss of SDRAM bandwidth. An option then would be to relocate some of the Fascicle Processors' workload to another chip.

15.3 Test

production test

- run standard memory test patterns from one of the processing subsystems.

start-up test

run-time test

16. Boot ROM

16.1 Features

- a small on-chip ROM to provide minimal support for:
- initial self-test, and Monitor Processor selection
- Router initialisation for bootstrapping
- system boot.

16.2 Fault-tolerance

Fault insertion

Fault detection

Fault isolation

Reconfiguration

16.3 Test

production test

start-up test

run-time test

17. Boot, test and debug support

17.1 Features

- means of booting system and flood-filling the distributed operating system efficiently at start-up.
- back-up boot mechanism in case Boot ROM fails.
- access to ARM968 EmbeddedICE features.
- sundry features to facilitate production, start-up and run-time testing.

17.2 Issues

At system power-up we can make few assumptions about what is and isn't working within the system. What is the minimum that must work for each chip to run internal self-tests, appoint a Monitor Processor, and then participate with its peers in an efficient bootstrap process that loads a distributed operating system into every node?

The inter-chip communication system is very soft and must be initialised before any mc or p2p communication can take place. But each node has no initial knowledge of where it is in the system, so how can it initialise the Router?

The ultimate system is large, so the bootstrap process must be efficient and employ flood-fill algorithms.

17.3 Boot algorithm

- Following power-on reset, each chip will perform internal self-tests and a Monitor Processor will be selected, probably as a result of asynchronous arbitration processes. The node will go into receptive mode, relying on the default boot routing process to communicate.
- The host system will begin sending OS load packets in nearest-neighbour format, tagged with sequence numbers, to the node to which it is directly connected. All nodes receive all incoming nn packets and, if they have not been seen before, retransmit them to all neighbours. Any packet which has been seen before will be dropped and not retransmitted.
- Once all sequence numbers have been received a node will perform a CRC check and, if this is correct, begin executing the loaded OS code.

17.4 Fault-tolerance

Fault insertion

Fault detection

Fault isolation

Reconfiguration

17.5 Test

production test

start-up test

run-time test

18. Input and Output signals

18.1 External SDRAM interface

Signal	Type	Function

18.2

19. Area estimates

We are targeting a UMC 130nm process 10mm x 10mm die. (Europractice runs on this process are multiples of 5mm x 5mm. The test chip will be 5mm x 5mm.)

Assumptions

- RAM is around $2\mu\text{m}^2/\text{bit}$ = 3M T/mm².
- logic is 0.2 x the density of RAM = 100k gates/mm².
- The pad ring occupies 0.25 mm all round the chip, so the core is 9.5 x 9.5 = 90.25 mm².

Using these assumptions we total up the core logic area as follows:

- The processor nodes = 20 x 3.8 = 76 mm².
 - An ARM968 with 32 kByte I-RAM and 64 kByte D-RAM is 3.5 mm².
 - DMA, interrupt, counter/timer, communications controllers: 20 k gates = 0.2 mm².
 - Communications and Systems NoC interfaces = 0.1 mm².
- The Communications NoC = 9.7 mm².
 - The associative router with 1024 associative entries is ~ 9mm².
 - The algorithmic router with 64k x 3 entries is 0.2 mm².
 - The Arbiter is small ~ 0.1 mm².
 - The Tx and Rx interfaces are small: altogether ~ 0.2 mm².
 - Communications network fabric ~ 0.2 mm².
- The Systems NoC = 4 mm².
 - The 128 kByte System RAM is 2 mm².
 - The Boot ROM is small ~ 0.2 mm².
 - The System Controller with 20k gates is 0.2 mm².
 - The SDRAM controller with 60k gates is 0.6 mm².
 - The network fabric is ~ 100kgates = 1 mm².
- Boot, test and debug = 0.5 mm².

Total area

The total core logic area is thus 76 + 9.7 + 4 + 0.5 = 90.2 mm².

Notes

Associative Router = 1024 x 96 latches + 96 gates = 500k gates = 5 mm²?

20. Power estimates

Processor

ARM968 (from ARM web site) consumes 0.12 to 0.23 mW/MHz on a 130 nm process, and delivers 1.1 dhrystone MIPS/MHz. Thus, to a good approximation, its power-efficiency is 5,000 to 10,000 MIPS/W and it uses 100-200 pJ/instruction.

neuron dynamics

30 instructions at 1 kHz = 30 kIPS = 3-6 μ W.

connection processing

1,000 inputs at 10 Hz (ave.) and 10 instructions/input = 100 kIPS = 10-20 μ W.

SDRAM access

assume SDRAM uses 250mW at 1 Gbyte/s; accessing 4 bytes costs 1 nJ.

1,000 inputs at 10 Hz (ave.) = 40 kByte/s = 10 μ W.

communications link

2.5V I/Os, 10 pF/wire = 30 pJ/transition

3 transitions/4 bits + EOP = 33 transitions/spike = 1nJ/spike/link.

Router

assume power budget at full throughput of 200 MHz is 200 mW, so 1 nJ/route.

neuron total

at 10 Hz (ave.), with H hops, power = 3-6 + 10-20 + 10 + (1 + 2H)10⁻³ μ W
= 23-36 μ W (routing & inter-chip hops are negligible).

Chip

20 processors x 1,000 neurons/processor x 13-26 μ W = 260-520 mW.

Node

chip + SDRAM = 460-720 mW.

System

1 billion neurons = 50,000 nodes = 23-36 kW.

S p i N N a k e r