EUROPEAN ORGANIZATION FOR NUCLEAR RESEARCH

CERN–ECP/95–16
30 August 1995

REALIZATION OF A 1000-NODE HIGH-SPEED PACKET SWITCHING NETWORK

R.W. Dobinson, B. Martin
CERN, Geneva, Switzerland

S. Haas, R. Heeley, M. Zhu
CERN, Geneva, Switzerland and University of Liverpool, Liverpool, England

J. Renner Hansen
Niels Bohr Institute, Copenhagen, Denmark

Abstract

Large data-switching networks are prominantly featured in future experimental proposals such as the High-Energy Physics programme for the Large Hadron Collider at CERN.
The technology to support this is now becoming available from industry in a variety of forms, one of which is the proposed IEEE P1355 standard. We describe a project to construct a large network testbed using P1355 point-to-point links and STC104 packet routers. A variety of nodes will be connected to this network, and its performance in terms of latency and throughput will be measured and compared with the results of simulation.
A large number of cheap, so-called simple nodes will be used to generate known traffic patterns. Against the background of this predefined traffic, test packets will be injected by more intelligent nodes and will be used to study network performance. Measurements have been done on engineering samples of the communication chips and these are compared with expected performance.
Finally, the future developments and measurements to be made will be presented.
This work is mainly supported by the MACRAMÉ project, ESPRIT 8603.

1 INTRODUCTION

High-speed packet switching networks connecting a large number of intelligent nodes over moderate distances are a feature of the data-acquisition and triggering systems of next-generation High-Energy Physics experiments [1], for example at the CERN Large Hadron Collider (LHC), Geneva, and at HERA-B at DESY, Hamburg.

The technology to support this is now becoming available from industry in a variety of forms, one of which is the proposed IEEE P1355 [2] standard. We describe a project to construct a large network testbed using P1355 point-to-point links and STC104 [3] packet routers. A variety of nodes will be connected to this network, and its performance in terms of latency and throughput will be measured and compared with the results of simulation.

A large number of cheap, so-called simple nodes will be used to generate known traffic patterns. Against the background of this predefined traffic, test packets will be injected by more intelligent nodes and will be used to study network performance.

It is also planned to connect to the network a variety of commercial processors including DSPs, T9000 transputers, PowerPCs, etc., running user applications.

2 A BRIEF DESCRIPTION OF THE P1355 STANDARD

The P1355 standard covers the physical connectors, cables, and electrical and logical protocols for implementing point-to-point serial interconnects operating at speeds of 100 Mbit/s and at 1 Gbit/s over copper and optical fibres. This technology has been developed in the Open Microprocessor Systems Initiative/Heterogeneous InterConnect Project (OMI/HIC) [4].

The goals of the standard are:

- to enable high-performance, scalable, modular, parallel systems to be constructed with low system-integration cost;
- to support high-performance communications infrastructures;
- to provide a transparent implementation of a range of high-level protocols and to support links between heterogeneous systems.

This project uses so-called Data/Strobe (DS) links running at 100 Mbits/s over printed-circuit and cable interconnections. A DS link is a four-wire protocol with two wires, data and strobe in each direction. The data line carries binary data and the strobe line changes state when the data line does not, as is shown in Fig. 1.

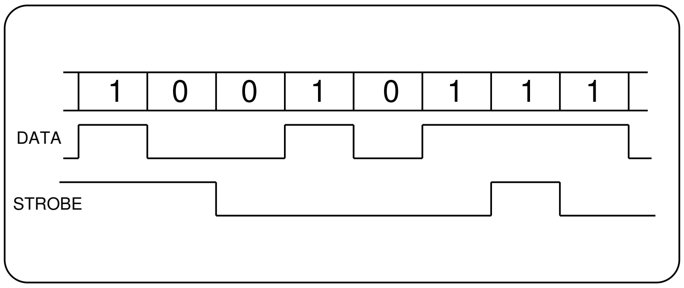

Figure 1: DS link protocol

As this protocol encodes the clock it allows for autobaud synchronization at the receiver and the possibility that, in any bidirectional pair, each direction may run at different speeds. In differential operation each of the four wires is transmitted and received as a differential pair (DS-DE) resulting in an eight-wire protocol. Studies on the transmission of DS links over copper have been carried out at CERN [5] and show reliable operation at distances up to 20 metres.

The serial bit stream is transmitted in groups of characters. Characters may be used to express either data or control information. The character format is shown in Fig. 2.

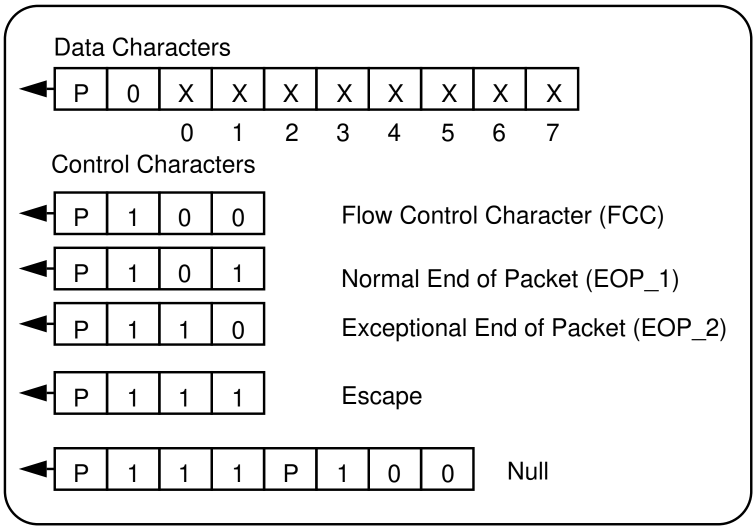

Figure 2: DS character encoding

The protocol exchanged between two DS link engines takes care of flow control, parity generation and checking, and the generation of null characters in the absence of other data to transmit. A packet is a sequence of characters consisting of a destination followed by a payload and delimited with an end-of-packet marker. There is no maximum packet size specified. For efficient operation the length should be restricted to avoid long packets blocking the network.

The destination address can be one or more bytes, although routing networks must know in advance the number to expect. In the simplest case of a point-to-point link the destination address can be a null.

3 SILICON COMPONENTS SUPPORTING P1355

This standard is currently implemented in the STC101 link driver chip [6] and in the STC104 thirty-two valent packet routing switch both developed by SGS-THOMSON Microelectronics.

3.1 The STC101 parallel DS-link adapter

The STC101 adapter drives a DS link at 100 Mbits/s full duplex and has a 16- or 32-bit parallel interface. The device operates in two modes depending on whether or not use is made of the internal packetizing function. Without internal packetizing the STC101 transmits input characters transparently and it is up to the application to provide explicit header information for routing.

If internal packetizing is exploited, the application writes routing and packet-size information to control registers and then supplies the data payload. The block diagram of the device is shown in Fig. 3. The presence of data FIFOs is essential for matching the speeds of external logic to the DS link flow-controlled protocol, but they add extra delay to the latency of data transmission and reception.

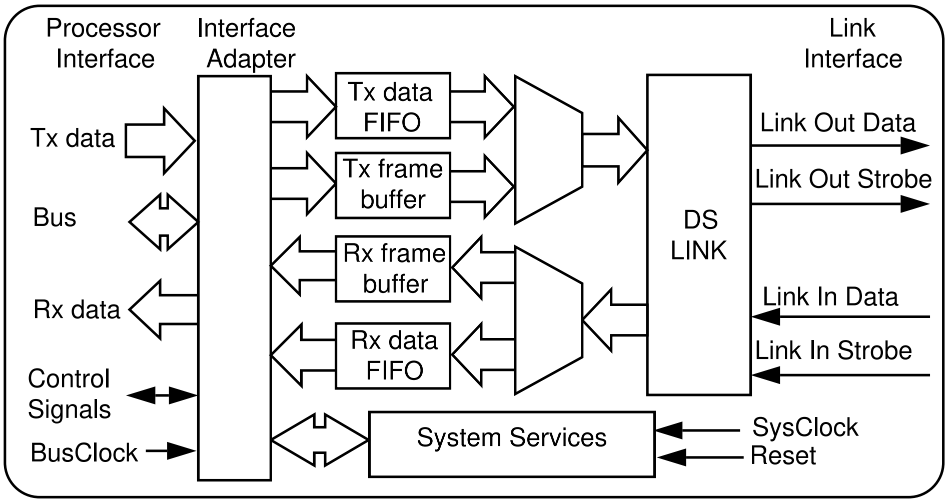

Figure 3: STC101 block diagram

3.2 STC104 asynchronous packet switch

This chip connects 32 DS bidirectional link ports via a 32×32 non-blocking crossbar switch thus enabling packets to be routed from any one of its links to any other link. The block diagram is shown in Fig. 4.

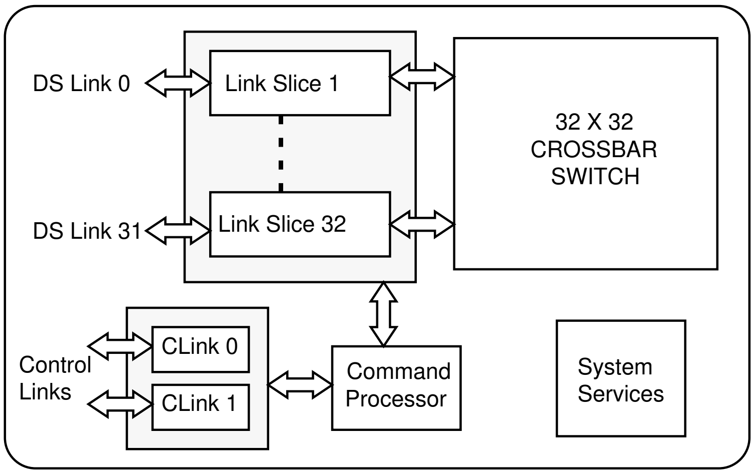

Figure 4: STC104 block diagram

It is a property of the switch that the links operate concurrently and traffic between any pair of links has no effect on the traffic between any other pair of links. The serial link speed is 100 Mbits/s with current silicon and the router processes up to 200 Mpackets/s with a maximum bandwidth of 300 Mbytes/s, and a latency of less than 1 μs.

To avoid store and forward buffering problems, the STC104 uses wormhole routing in which the routing decision is taken from a packet header as it arrives. The header destination address is compared against preloaded intervals which are used to determine the packet exit port. Packets of arbitrary length can be handled with full flow control.

In any network a major source of delay is caused by contention of several packets for a common resource. The STC104 offers two mechanisms to reduce this effect. Firstly, groups of links may be logically bundled together; concurrent incoming packets from different ports headed for the same output destination will be able to use any free link in the group. This of course requires link redundancy and the width of a group will be a function of the contention expected and the resources available. These are both functions not only of the traffic but also of the topology of the network.

However, it is not always practical or possible to assign more physical links per group than the peak number of packets wishing to concurrently pass through that group, and so local traffic contention, hotspots, may still occur. A second mechanism, universal routing [8], reduces this effect of hotspots by first sending a packet towards a random destination address. At an intermediate point in the network it is routed to its final destination.

4 PROJECT GOALS

- To construct very large networks with different topologies: Clos, grid, mesh, tree.
- For each topology, to measure the performance of the network, latencies and throughput, as a function of packet length and traffic patterns.
- To investigate traffic patterns corresponding to those found in a variety of application areas, both commercial and in the field of High-Energy Physics.
- To investigate the network behaviour under different packet routing algorithms.
- To determine the behaviour of a network under conditions ranging from sparse to heavy traffic (near saturation and saturation) and to establish a performance envelope.
- To compare measured performance with the results of simulation and to thereby calibrate simulation models.
- To study price–performance trade-offs for particular applications.

5 NETWORK SATURATION ISSUES

In any network there is a trade-off between bandwidth and latency. Unlike some networks that may discard data to meet fixed latency requirements, the P1355 protocol will deliver all packets but with a non-deterministic latency. As already mentioned hotspots may develop and propagate. Studies have shown [7] that this effect, known as tree saturation, is common to any multistage network with distributed routing and non-uniform traffic patterns. The effect is largely independent of the topology, switching mode (packet or circuit), and whether the network is being used for memory access or message passing. Simulations [8] based on STC104 technology for a message size of 32 bytes show that for a fully loaded network a few microseconds of contention is sufficient to reduce total throughput by a factor of two or more and that milliseconds may be required to re-establish a steady state. With the test network being constructed we can map this effect more thoroughly and provide guidelines to users as to just where the limits are of linear behaviour for any given combination of traffic and connectivity.

6 NETWORK COMPONENTS

The construction of this network is essentially a packaging exercise. There are many parameters and constraints to take into consideration. It is obviously desirable to produce as few boards as possible and as few different boards as possible. This is balanced by the difficulty and cost of having to place too many components on one board and the fact that different topologies have different interconnectivity needs. It is preferable to have as few racks as possible to simplify power supply distribution and cooling but handling such a large number of cables imposes limits of pacticality for front panel design. The control network needs to handle traffic quickly and with maximum inherent redundancy yet without occupying too large a share of connector and board space. After studying many possibilities it was decided to proceed with the design and construction of the following components:

- a test board to verify the simple source design and performance
- a switch board with 32 external links
- a switched-node motherboard carrying the STC104 chip, sixteen external links and carriers for piggy-back daughter boards
- switched node daughter boards each carrying four simple sources
- an intelligent node for time-stamped packet handling
- a link diagnostic module for traffic probing
- a root and control structure.

6.1 The simple node

6.1.1 Requirements and compromises

The traffic profile is described in terms of statistical distributions of the following parameters:

- payload length
- destination address
- time of despatch

The design should be capable of maintaining maximum bandwidth for the whole range of traffic profiles and cope with the effects of traffic congestion. To meet these requirements we define a profile in terms of a set of packet parameters which are stored in the simple node. During run time these parameters are extracted sequentially, interpreted, and a corresponding packet is transmitted. The cost and size issue resulted in four compromises.

1. The number of entries in the packet set is a trade-off between statistical distribution and component cost. We estimate that a store of 4000 to 8000 events (dependent on header size) is sufficient given that the sequencer can loop indefinitely on the same set.
2. It becomes difficult to meet the maximum bandwidth requirement for low packet lengths without excessive complexity and we estimate that bandwidth will be limited for packets of less than 10 bytes. The behaviour of packets shorter than this can be studied by the intelligent nodes.
3. Coping with traffic congestion is described in detail in Section 6.1.4. Simply put, there is a need for a counter. The length and time resolution of this counter determines the maximum time that a port can be stalled and still recover normal operation once the congestion eases. This design can cope with delays of up to 32 ms with a resolution of 500 ns. If greater delay is required this can be done, albeit with less resolution.
4. The requirements for time of despatch refers to one of two possible despatch algorithms. In the first, a packet is sent some known time after the previous one was sent. In the second, a packet is sent some known time after an incoming packet has been received. Having both possibilities available in the state machine resulted in excessive logic. By choosing RAM-based field-programmable gate arrays (FPGA) it is possible to load, at start-up time, either of these algorithms into any number

of the simple source nodes. The actual data in a packet is of little interest except as an identifier for diagnostic purposes or as a probe for data sensitivity. The simple nodes provide only a single data byte which is repeated for the whole packet length.

6.1.2 Principles of operation

Figure 5 shows the main elements of the simple node design.

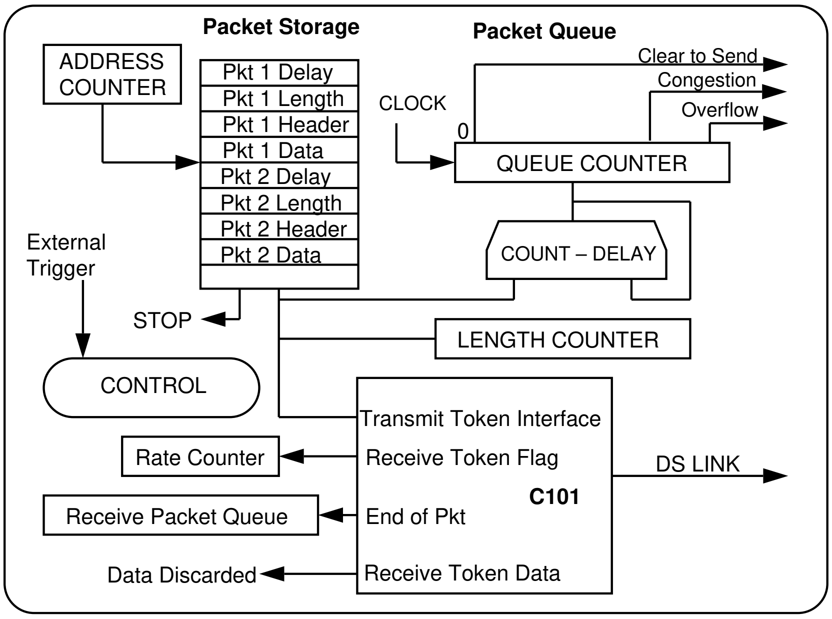

Figure 5: Simple node functional block diagram

6.1.3 Packet storage

The packet descriptors in the store are shown from a functional point of view. In practice the descriptor fields are of different lengths and packed for storage efficiency

6.1.4 Packet queue

Packets are to be despatched according to some predetermined schedule. The delay is defined as the time between the sending (or receiving) of packets. In the case of network congestion there may be packets stalled to the point that the STC101 transmit FIFO is unable to accept further data. A queue counter is maintained which increments with time starting from the first packet token being written to the STC101. For subsequent packets the state machine compares the elapsed time, the counters value, with the time it should wait, the delay descriptor. When equal, the counter is loaded with its actual value minus the delay value, nominally a zero, and the next packet is despatched. With congestion, the point is reached when the elapsed time is greater than the delay descriptor, therefore the comparison shows a positive result and so the packet is immediately despatched, but late. If the congestion eases then subsequent packets will also be flushed as fast as possible until the counters value is less than the intended delay and synchronism is restored. However if congestion worsens, the packet despatch will be further and further delayed until the counter overflows and synchronism is irretrievably lost.

6.1.5 Packet despatch

The STC101 is run in transparent mode. This was done partially because the external logic is simpler in terms of chip count and printed circuit board area, and partly because the test packet tokens are in any case pre-defined as a function of the test to be performed. Once able to send, the state machine fetches the next descriptor, which is the packet length, and writes it into the transmit counter after which the packet header is fetched and is written into the STC101 if the token port is free. Then the data descriptor byte is fetched and applied to the token input of the STC101 for the duration of the packet transmission. Transfer of data into the token port FIFO follows as rapidly as the STC101 can absorb it and continues until the length counter indicates the end of transfer.

6.1.6 Loop structure

For most purposes the simple node will be used in loop mode whereby once all the events have been despatched the address counter loops around and sends them again.

Even if none of the nodes has lost synchronism, there will typically still be an uncertainty as to just where any particular node is in its loop. Even if it were known, the action of the STC101's FIFO contributes timing uncertainty. It is desirable to have as close a control as possible over certain packet despatches in order to generate, on demand, a proportion of known, conflicting packets. To achieve this it is necessary to start, on command, in the middle of a stable traffic background, a set of nodes which output only a single set of packets and then stops. This is done by issuing a selective trigger from the host to start a chosen set of nodes which have a STOP flag set in the packet descriptor to inhibit the controller from looping once the packets have been despatched.

6.1.7 Packet reception

At the receive port of the STC101 the incoming tokens are flushed as fast as they arrive. Link error conditions can be flagged and cleared. The main requirement here is that a receiving node offers no extra latency or bandwidth constraints since it is the network properties that need to be measured, not the nodes.

The token-received trigger is counted by a divider which generates an interrupt every 32 000 tokens to permit an external microprocessor to calculate the received bandwidth. This division is done to reduce the processors' load especially since there are four such sources of interrupt per processor.

For diagnostic purposes, however, in a search for missing or misdirected packets it is necessary to know exactly how many packets have been received and so the packet-received trigger is made available to external counters.

6.1.8 Implementation

The link traffic source is implemented in three chips: a STC101, a memory chip and a FPGA. They require some controller capable of loading, starting, and monitoring the links' progress, handling errors and reporting status to the host computer. A cost-effective single chip solution to this is the T225 [9] transputer which has a clock speed sufficient to handle up to four link sources at a time at full bandwidth. Lastly there is a need for one electrically-programmable logic device to handle the decoding and interrupt combining. The complete simple node is shown in Fig. 6.

The OS links from the controller are used to connect nodes to the host, the State Machine's drive status information flags and the Packet Receive signals are taken to external counters for diagnostic purposes.

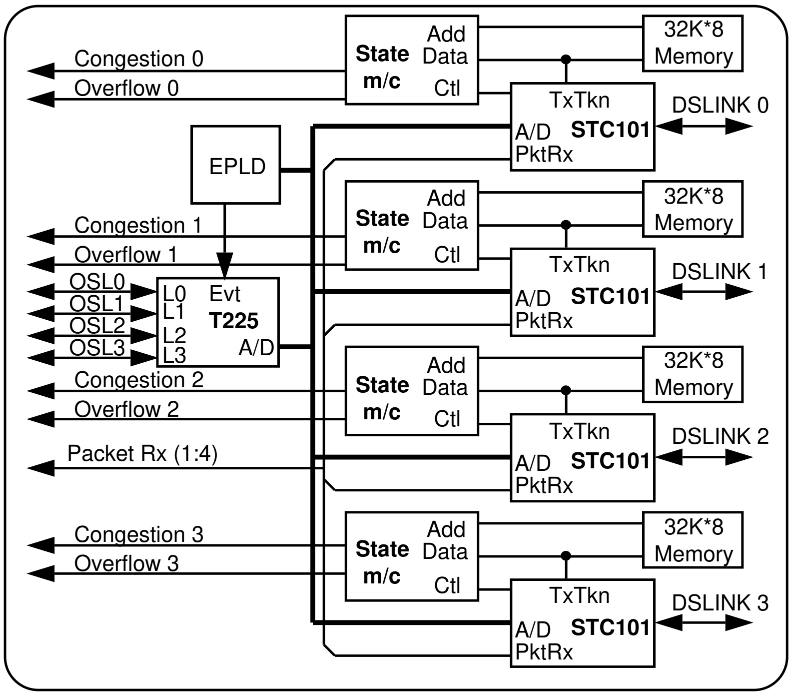

Figure 6: Simple nodes and controller

6.2 The intelligent node

It is from the definition of the simple node that the rest of the design stems. The intelligent node should adhere to the same control structure as the simple nodes, and provide additional functionality. Figure 7 shows the block diagram in which a T805 [10] controls a single STC101. This is a cost-effective solution meeting both the speed and communication requirements.

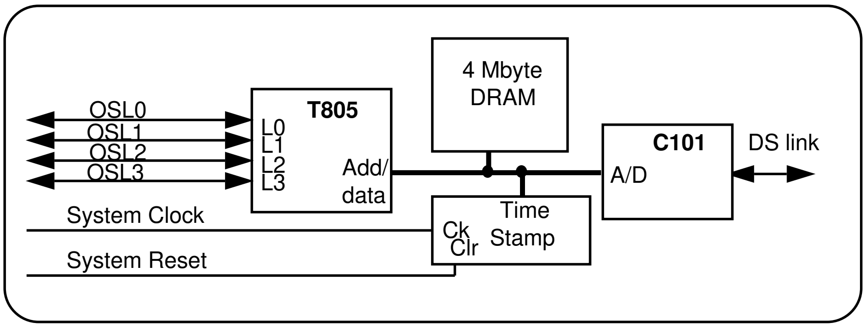

Figure 7: Intelligent node

The timing information for latency measurements is provided by a time stamp counter. All counters in the system are clocked with the same system clock and reset with the system reset to maintain synchronicity across the network. It is this value which is fed into the data stream of the outgoing packet. The receiving T805 must subtract its local time from the incoming time of despatch and maintain the statistics during the period of measurement.

6.3 The switch boards

Inspection of the networks under study reveals two modes of connecting switch chips:

- in the first mode some of the links go to nodes and the remainder to other switch links. This link distribution determines how rich in connectivity the net is compared to the maximum load it must carry
- the second mode of connection is where switch chip links go only to other switch chip links. This is true of the higher levels of a Clos or tree network.

Off-board links are carried in differential DS-DE form and the driver chip, connector, cable and share of printed circuit board area are a significant proportion of the total network cost and should be minimized. This implementation offers two solutions as shown in Fig. 8. In the first case 16 simple sources are packaged on a 6U VME board and consume half the links of a STC104 chip. The other half are brought out to the front panel. This ensures that there is at least as much bandwidth on the net side as on the source side of the switch. All of the simple nodes will be packaged this way. In the second case the STC104 is packaged on a 6U VME board and all its links are brought out to the front panel.

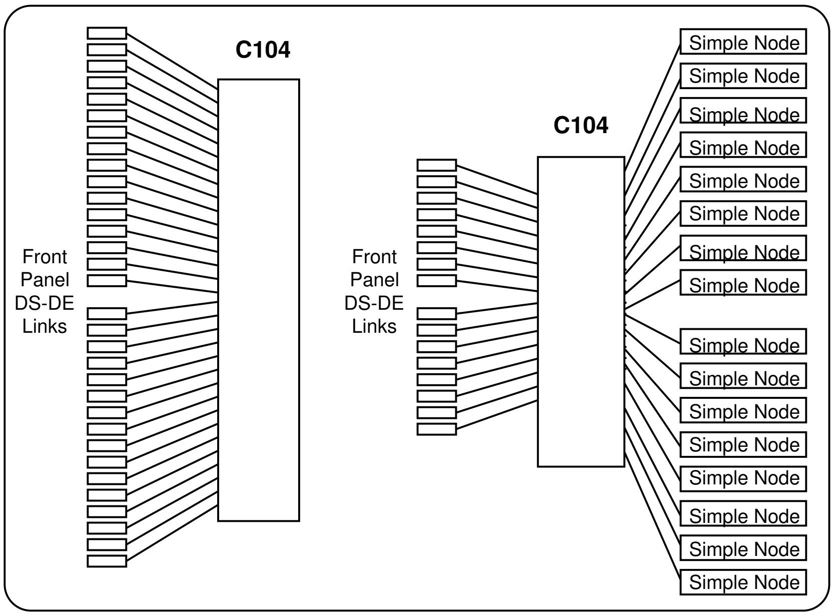

Figure 8: Switch board (left) and switched node board (right)

6.4 Hybrid nodes

As well as the nodes already described, development is ongoing to provide user application nodes to study the effect of different topologies and architectures on their specific traffic sets. One such development is being done in a study for the ATLAS second-level trigger [11]. In this model, shown in simplified form in Fig. 9, samples of data from a detector are buffered in memories and passed to a farm of micro-processors which extract the data concerning areas of interest known as features. These features are then passed on, through another switch to further processors which combine the results on an event basis and make a 'global' decision on whether or not the event is to be retained for further study, or discarded.

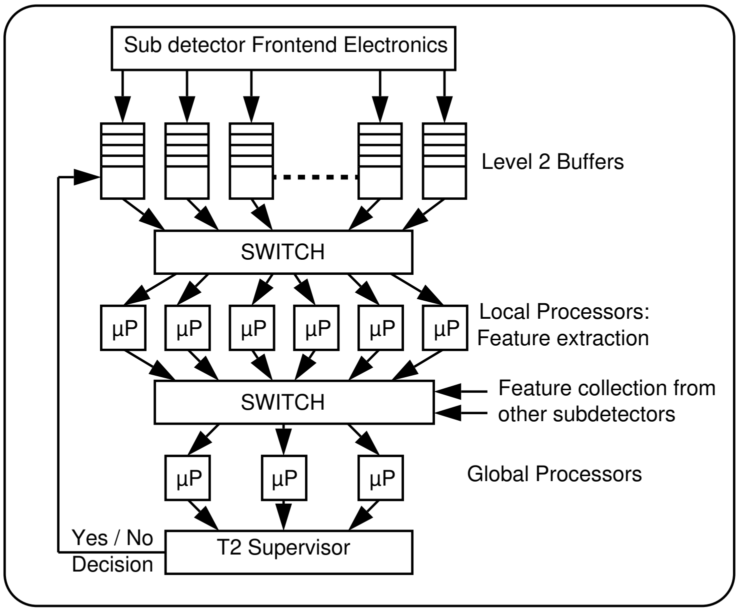

Figure 9: ATLAS level-two trigger farm

The speed with which this can be done is a function not only of the processors and the algorithms they run but also of the latency of the switch networks. The maximum total time for the entire process is a critical parameter since it determines how much memory buffering is needed in the total experiment and the efficiency of the trigger detection.

An implementation of the local system uses TMS320C40 DSPs [12] for combined buffer management and local processing. To interface this system, built around DSP-links, with a switching fabric using DS links, a translation module is being prepared as shown in Fig. 10. A TMS320C44 [13] processor receives data from up to four DSP-links and prepares them for despatch with the required destination information and passes them to the high-speed 64 kbyte dual-port memory. This memory is shared with a T9000 [14] processor which despatches the stream of features into the switching network. Communication between the two processors is done through message pointer queues established in the shared memory defining the size, location and direction of the data packets to be transmitted. Interrupts are used between processors to signal change of state.

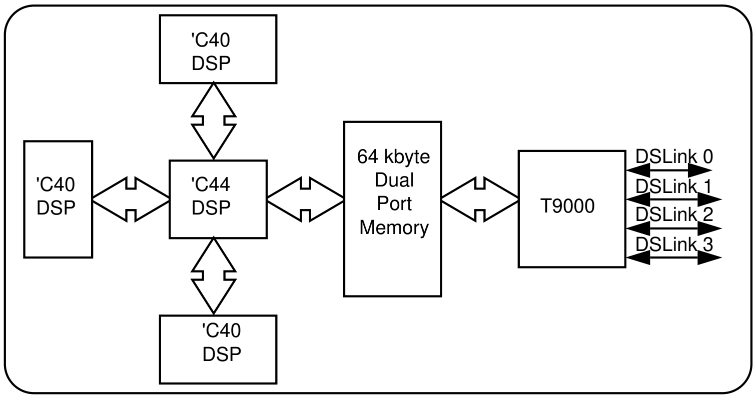

Figure 10: DSP link to DS link module

6.5 Root and control structure

The network will be controlled and monitored from a single root driving two branches as shown in Fig. 11. The OS link [15] branch fans out through a tree of T805 transputers to four racks of switched simple nodes. Each rack has sixteen switched source boards connected serially. Each board carries four T2 controllers connected as shown and each controller manages four simple sources.

The T8 fan-out tree also connects to a serial chain of T805 controlled intelligent nodes and to a single system module.

The DS link branch fans out through a STC104 switch chip to the four racks of 16, serially connected STC104 switches. The host computer communicates with the root node over Ethernet.

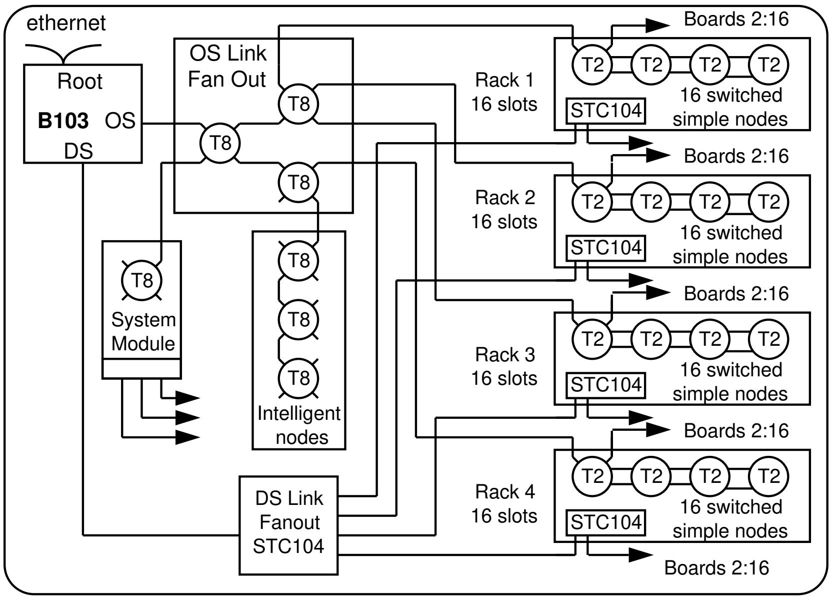

Figure 11: Root and control network

The root module is the IMSB103 [16] which has the capability to control not only the OS links of the control chain but also the DS links of the STC104 switch chain.

The root node channels all the run-time monitoring information, via Ethernet, to a host workstation.

The systems module is processor controlled and is linked to the control chain through its OS links. It is the source of the global clock and reset signals which are distributed throughout the network. It is also the source of the global trigger which starts a test run and the selective triggers used for controlled conflict generation.

6.6 Link diagnostic module

Statistics from the intelligent nodes will only give information on the latency of a given traffic between two given points. It will give no information on where the congestion was, nor how it was distributed, nor over which exact paths the packets passed. The link diagnostic modules will be used to probe points on the network during latency measurement tests to find out how packets are being routed and where delays are being experienced. A module may be inserted in series with any link cable. The link traffic is received in the module and immediately retransmitted causing only a few nanoseconds of extra delay. A single FPGA deserializes the received traffic and does the token detection and analysis. The token stream is encoded and written together with the local time-stamp into a memory shared with a T225 transputer. The received packet time-stamp is also extracted and saved in memory. The T225 transputer passes this information over the control network back to the root host. Only small samples of data flow can be handled this way and the module will be triggered externally under host control.

7 PROJECT STATUS

A test board containing four simple nodes and their associated T225 controller has been constructed to validate the design and is undergoing extensive testing prior to the construction of the 1024 nodes for the full network.

The daughter board for the four simple nodes is being laid out and will go to construction shortly. This is a critical phase since, even though the node has been designed to a minimum chip count, the space constraints are severe. Once the geometry of the layout is accepted the switched node motherboard can then proceed to layout and construction.

The STC104 Switch board bringing out all 32 links to the front panel has been built and fully tested. The hybrid node is under construction and the intelligent node, system, and diagnostic modules are in the design phase.

We expect to have all components available by the end of 1995. The control and support software is under development using an ad hoc test rig of a B103 root module and a few target T225 processors. Performance measurements have been carried out on a Clos network of STC104 chips driven by T9000 transputers in the GPMIMD [17] machine.

8 EARLY RESULTS

8.1 STC101-based simple node

The testing of the simple node design of the test board yields data about both the node design and STC101 performance. The graph in Fig. 12 is based on test data in which two simple nodes are connected via a STC104. The data shown is for an unidirectional data transmission, source to sink, using two byte headers. The theoretical curve shows the maximum possible unidirectional bandwidth of a DS link in terms of packet size.

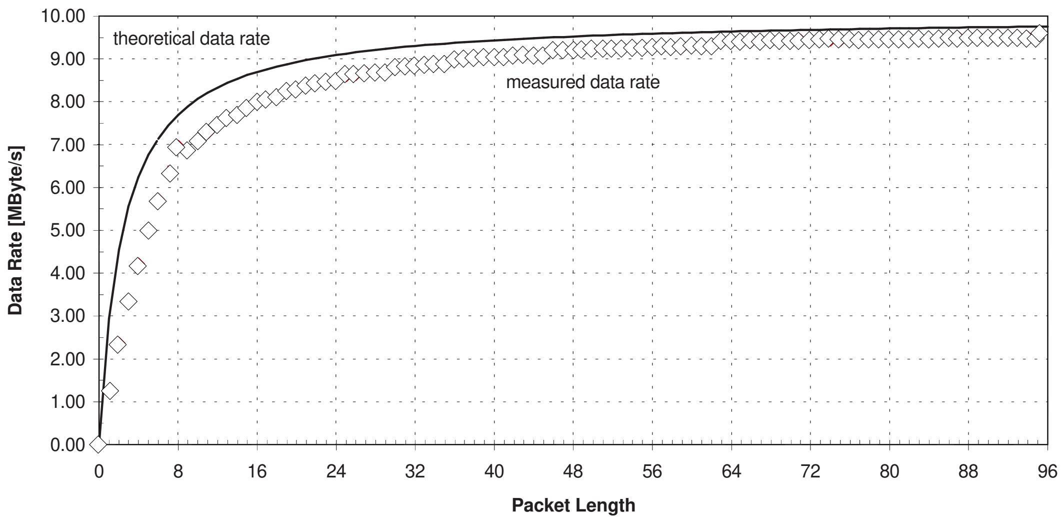

Figure 12: STC101 to STC101 unidirectional bandwidth through a STC104

The measured curve shows the effect of handling short packets already described which results in 95% of the maximum bandwidth being attained only for packets larger than 30 bytes in length. For long packets, the measured rate approaches the theoretical asymptote of 10 Mbytes/s.

The bidirectional data rate is slightly lower due to the overhead of the link flow control protocol. For long packets we have measured a data rate of 9.5 Mbytes/s corresponding to the theoretical asymptote of 9.52 Mbytes/s.

8.2 STC104 latency

The packet latency has been measured using a digital oscilloscope as follows: output header delay from input header = 1.00 μs. The measured jitter is of the order of ± 0.1 μs and is due to the passage of the packet across different clock domains.

8.3 STC104-based Clos network system bandwidth

Measurements have been carried out on the aggregate cross-sectional bandwidth of the GPMIMD machine. This machine has a Clos network architecture in which there are three stages of routing chips as seen in Fig. 13. For this test a reduced configuration was used in which the first and third stage of switching have 112 DS links connected to T9000 processors. The second, central, stage supports 64 concurrent circuits between the first and third stages. A varying number of these central links can be employed by declaring or omitting them in the routing tables. There are 28 processors programmed to communicate across the network with 28 corresponding processors in full duplex message passing. There is a one-to-one communication pattern with no destination address conflicts. Each processor uses seven virtual links per physical link yielding a measured maximum total bidirectional throughput of 33 Mbytes/s for long messages. This is the maximum that a 20 Mhz T9000 can sustain. Thus for 28 processors there is a bandwidth capability of: $28 \times 33 = 924$ Mbytes/s over $28 \times 4 = 112$ source links.

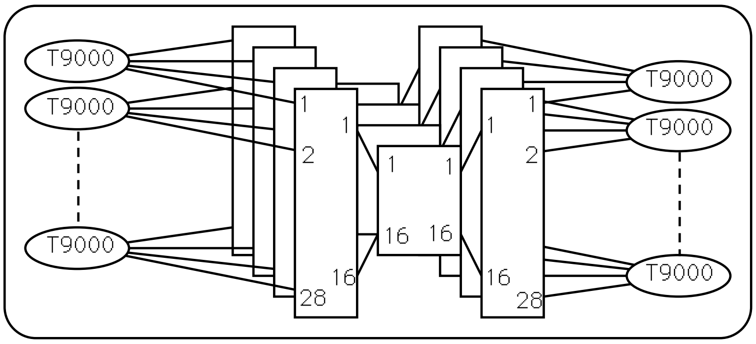

Figure 13: Three stage routing

The three plots in the graph shown in Fig. 14 shows the total throughput as a function of message length and the number of central links, 16, 32, or 64. It is clear that the bandwidth scales linearly between 16 and 32 links, a clear demonstration of the efficiency of grouped adaptive routing in exploiting all available resources. When 64 links are available for use, the T9000 bandwidth itself becomes the limiting factor.

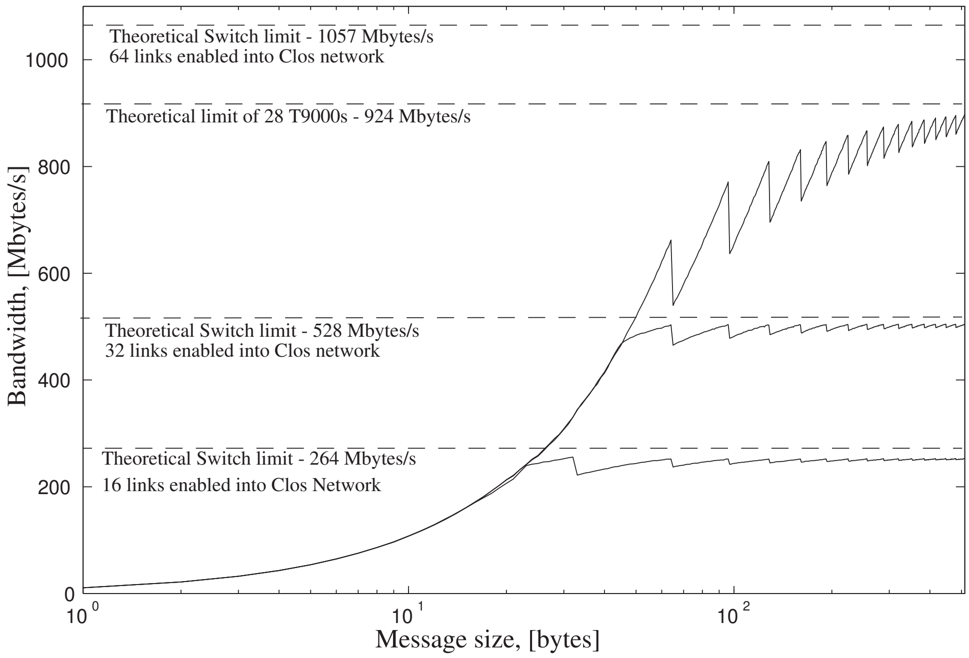

Figure 14: Cross-sectional bandwidth of a 256-node STC104 Clos network

9 CONCLUSIONS

The performance of the first-generation silicon supporting the P1355 standard has been demonstrated and measured and the use of this technology for the construction of large, flexible, switching networks has been shown to be feasible.

Early test board results show that the design parameters are being met now and with every expectation of full-scale performance figures for the 1000-node machine in the near future.

The GPMIMD machine results show agreement with earlier simulation [8] studies.

ACKNOWLEDGEMENTS

This study is being carried out to investigate a large switching network using European technology designed and built with the support of the Open Microprocessor Initiative, OMI, and prepared for market exploitation within the MACRAMÉ project, ESPRIT 8603.

We especially wish to thank Peter Thomson member of SGS Microelectronics Ltd., for his continued and valuable technical insight and members of the ATLAS experimental groups for their many suggestions and valuable collaboration.

References

[1] The ATLAS Technical Proposal, CERN/LHCC/94–43, LHCC/P2, ISBN:92-9083-067-0. The CMS Technical Proposal, CERN/LHCC/94–38, LHCC/P1.

[2] IEEE Draft Std P1355 Standard for Heterogeneous InterConnect (HIC) (Low Cost Low Latency Scalable Serial Interconnect for Parallel System Construction. IEEE Inc., 1995.

[3] STC104. Asynchronous Packet Switch, Preliminary Data Sheet, June 1994, SGS-THOMSON Microelectronics.

[4] ESPRIT Open Micropocessor systems Initiative 1992: High-Performance Heterogeneous Interprocessor Communications (OMI/HIC) project 7252, The Synopses, ISBN 92-826-4817-6.

[5] S. Haas, X. Liu and B. Martin, Long Distance Differential Transmission of DS Links over Copper Cable (CERN). http://www/hensa.ac.uk/parallel/vendors/inmos/ieee-hic/copper.ps.gz.

[6] STC101. Parallel Link Adaptor, Preliminary Data Sheet, June 1994, SGS-THOMSON Microelectronics.

[7] G.F. Pfister and V. A. Norton, Hot Spot Contention and Combining in Multistage Interconnection Networks, IEEE Trans. Comput. C-34 (1985) 943–948.

[8] A. Klein, Interconnection Networks for Universal Message-Passing System, Esprit '91 Conference Proceedings, pp. 336–351, Commission of the European Communities, 1991, ISBN 92-826-2905-8.

[9] The Transputer Data Book, 2nd ed., SGS-THOMSON Microelectronics, 1989, pp. 453–462.

[10] The Transputer Data Book, 2nd ed., SGS-THOMSON Microelectronics, 1989, pp. 85–126.

[11] The FEAST Collaboration. CERN-NBI-RHBNC-Krakow-UCL-Stockholm.

[12] TMS320C4x User's Guide, Texas Instruments.

[13] TMX320C44 Digital Signal Processor Data Sheet, Texas Instruments No. SPRS031A.

[14] Networks, Routers and Transputers, edited by M.D. May, P.W. Thomson and P.H.Welch. T9000 pp. 15–36 ISBN 90 5199 129 0.

[15] The Transputer Data Book, 2nd ed., SGS-THOMSON Microelectronics, 1989, p. 24.

[16] IMSB103 Ethernet to DS-Link Interface, Product Information August 1994, SGS-THOMSON Microelectronics.

[17] The Esprit GPMIMD1 project P5404, Commission of the European Communities, 1991.