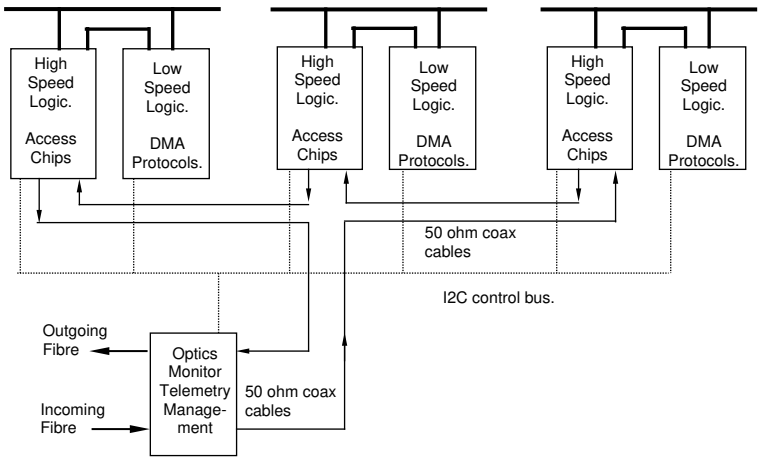

Figure 7: Example configuration where one optic board provides fibre connection and control for several Backbone Ring stations.

most common reason for station failure has been room mains power failure, and so groups of stations tend to need isolating from the ring in one step.

Reconfiguration time after switching in or out a group of stations depends on whether the active ring monitor is removed from the ring as a result of the switching. If the monitor is deleted, re-operation of the ring is delayed by the time required for a new monitor station to be elected. Owing to the 300 baud nature of the telemetry network, this can currently take over 15 seconds for a typical 10 station ring. It is hoped that such reconfiguration only rarely occurs.

Since the telemetry information is carried on the same fibres as the data, remote control and status information for stations isolated on the other side of the switch are lost. This problem can be overcome in a more sophisticated fibre switching module which has access to the telemetry data.

5 Current status and future plans.

Twenty five stations have been constructed at the 512 MHz rate and these are now carrying some of the Pandora traffic and a few other applications (e.g. XNS bridging of Ethernets).

The University of Cambridge has installed optical fibre between every College and University building. We plan to deploy the Backbone Ring using parts of this fibre base.

A future project will use the hardware experience of the Backbone Ring to build an ATM distribution network using 48 byte cells with a line rate of several GHz. The stations will include hardware structures commonly found in ATM fast packet switches so that ATM local area drop cables can be directly connected. The ring design will thus be more strongly oriented to its current use, which is as a distributed ATM switch.

Acknowledgements

Many thanks are due to the following people who have worked towards the Backbone Ring project: Dimitris Lioupis, Brian Robertson, Krzysztof Zielinski.

Reference List

[1] 'The Cambridge Fast Ring Networking System.' A Hopper and RM Needham. IEEE transactions on computers, Vol 37 number 10. October 1988.

[2] 'Pandora - An experimental distributed system for multimedia applications' by Andy Hopper. ACM Operating Systems Review, April 1990.

[3] 'The Cambridge Backbone Ring' by David J. Greaves, Andy Hopper and Dimitris Lioupis. In proceedings of IEEE Infocom 90, San Francisco 1990.

[4] 'The Cambridge Backbone Ring' by David J. Greaves and Andy Hopper In proceedings of EFOC/LAN 88, Amsterdam, 1988.

[5] 'Protocol Design For High Speed Networks.' University of Cambridge technical report 186. DR McAuley. December 1989.

[6] 'Preliminary performance results of the Backbone Ring v1s stations.' by David J Greaves, Krzysztof Zielinski. In proceedings of IFIP WG 6.1 conference, Palma Spain, June 91.

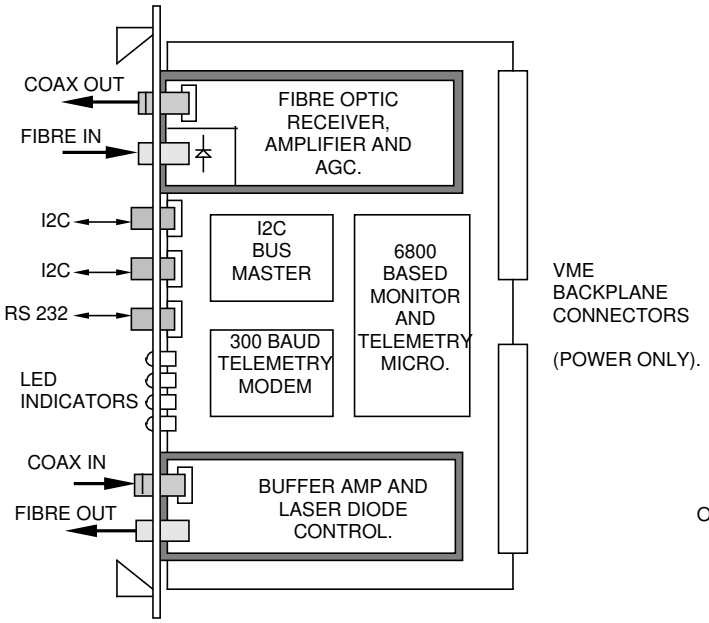

Figure 5: Layout diagram of the optical-electric and control board. This takes only power from the VME backplane, with no data connection.

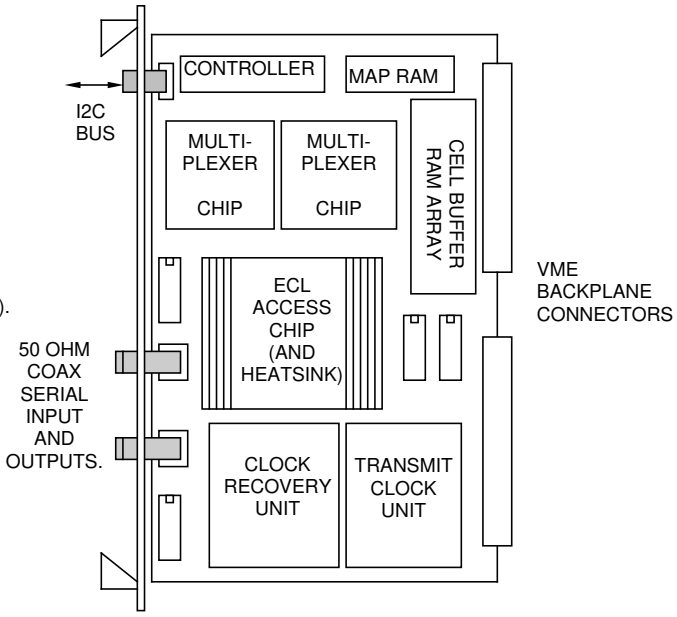

Figure 6: Layout diagram of the high-speed logic board.

dom logic. Its functions are to generate the addresses in the main cell buffer RAM for cell transfer, and to perform the VME bus protocol, providing DMA and user status and control registers. This board communicates with its high-speed logic partner via the 64 undedicated pins on VME J2.

Figure 7 shows how several pairs of high-speed and low-speed boards may be connected in a local cluster, controlled by a single optical board. Local interconnection is via the copper coax and the I2C cable. Separate VME segments are shown in the figure, but more than one station can be plugged into a single VME bus if desired. Using more than one station provides higher access bandwidth to the ring, although we have not yet investigated the VME bandwidth achieved in practice with this configuration.

Figure 8 shows the use of a fourth board-level component. This board contains an electromechanical, optical fibre, by-pass switch. There are two possible configurations of the (bi-directional) lightpaths through the switch, as shown. The switches are controlled via the I2C bus of the local station, but always revert to one configuration under conditions of power failure.

By providing a universal by-pass module, rather than incorporating a single switch into each optical board, cheaper and more flexible arrangements are possible. Our ring networks have always been star-wired, so one possible use of the switch is to remove a failed star spur of the network, as shown. The

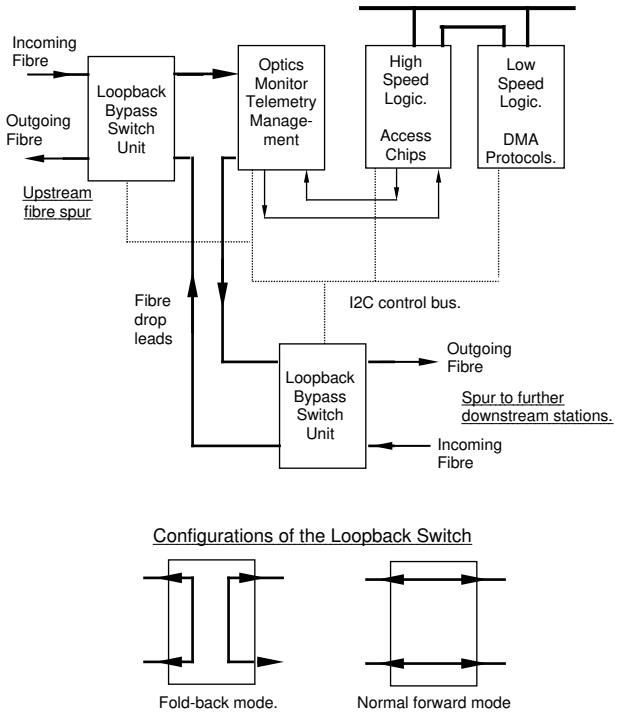

Figure 8: Examples of use of the fibre optic switch module to provided fault tolerance and branch isolation.

3 Backbone Ring VLSI chip set.

Figure 3 shows the simplest hardware configuration possible using the Backbone Ring chip set. The received data from the optical fibre is passed to an ECL access chip which performs two types of function. The first concerns the transformation of the serial data to eight bits wide and provision of clock and byte synchronisation. This requires the following functions:

- Decide whether the current input voltage is a one or zero using a D-type flip-flop,
- Convert from NRZI to NRZ,
- Decode 5 bit blocks into 4 bit data word or special code 'syn' used for synchronisation,
- Gain bit synchronisation using a slip method,
- Signal a code error on codebook violation,
- Gain frame synchronisation on frame header,
- Encode output stream into 4B5B and thence NRZI.

The ECL chip processes slots in *frames* of four consecutive slots. The ring contains an integral number of frames and a gap, filled with synchonisation symbols. The second type of function provided by the access chip is at the frame level and includes:

- Check CRC of received frame,
- Examine and update full flags to decide which slot is to be read or written,
- Generate data strobes for reading or writing the desired channel,
- Write new valid CRC to outgoing frame.

The 8 bit data streams from the ECL access chip are connected to the RAM packet buffer through the semi-custom multiplexer devices. Two similar devices are used in the basic station. Their purpose is to widen the data from 8 to 32 bits. This data widening results in a clock rate reduction of 4 to 1 and there is a further 4 to 1 reduction within these devices since they give access to only one of four slots in a frame at a time. The rate of the 32 bit words is then 6.25 MHz for a 1 GHz implementation, a rate and bus width which are fairly convenient for direct connection to computer backplane buses.

Figure 3 also shows a monitor and telemetry unit. This controls a maintenance sub-carrier which is carried by the physical fibre medium in addition to the main data channel. The sub-carrier uses mean-power modulation of the optical sources and so can operate even if the ring frame structure may have failed. The subcarrier carries data between 6800 style processors at each station. These processors perform ring monitoring functions, electing an active master station (for clocking and ring frame control) and providing a general telemetry service for bit-error rates, optical powers, etc.

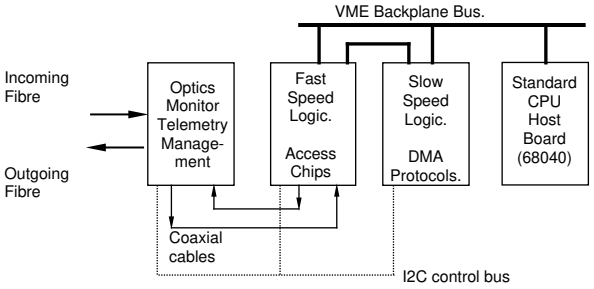

Figure 4: The current Backbone Ring stations, using three circuit boards.

4 Current Implementation

Figure 4 shows the three-board set used in the current VME bus implementation of the Backbone Ring. These operate at a clock rate of 512 MHz. A 68040 processor card is also shown. Not shown are further VME cards providing access to ATM point-to-point links or CFR. These provide conectivity to the Backbone Ring from the local devices. The traffic performance results have been presented in [6].

The optics and telemetry board, shown pictorially in Figure 5, contains a bidirectional pair of optical to coax converters and the 6800 microprocessor, together with a telemetry modem and an RS-232 interface. At this point, a terminal may be connected to any station to read off the status of the whole network. The optical link will extend in excess of 20 kilometers. The coax cables are limited to a few meters, but they need only be very short, since they connect directly to the high-speed logic board, which is normally in the same rack.

The high-speed logic board, shown pictorially in Figure 6, contains the clock recovery unit, the local transmit clock synthesiser, the access chips, VCI map ram and main cell buffer array. It also has the data buffers for data transfer to and from the VME host bus. A local, single chip controller receives supervisory instructions via the I2C bus from the 6800 processor.

A third board, called the low speed logic board, was shown in Figure 4. This contains only PALs and ran-

Full flag (1)	monitor passed flag (1)	Virtual circuit etc (32)	Data Payload (256)	CRC (16/4)

Figure 2: Significant fields of a cell as it appears on the Backbone Ring with their sizes in bits. For efficiency, the CRC field is shared between, four cells as explained in [3].

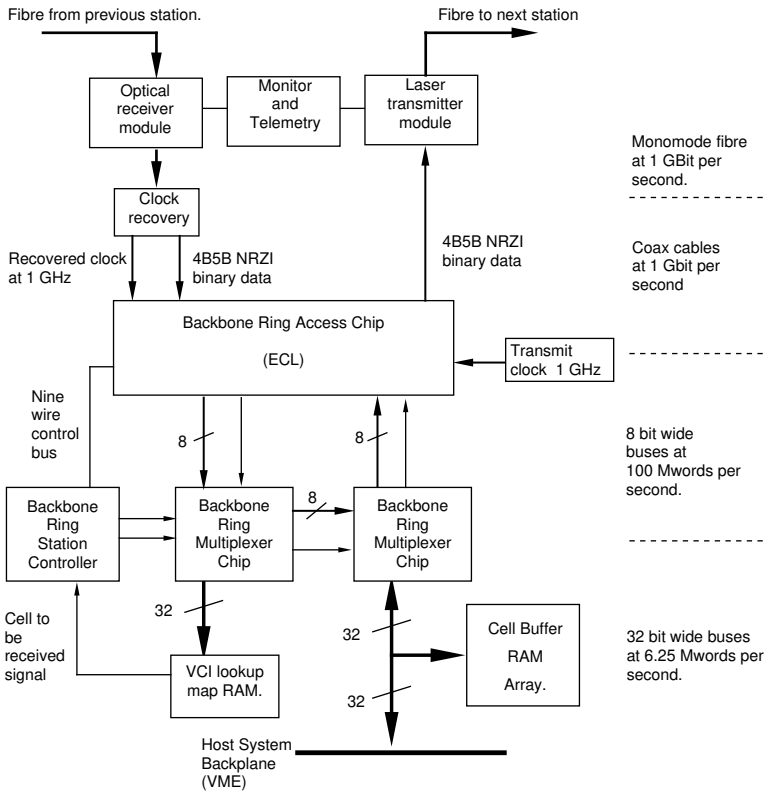

Figure 3: Configuration of the most simple type of station that can be constructed with the Backbone Ring chipset. The stated speeds assume a clock rate of 1 GHz, although the current implementation of the multiplexer chip limits the operation of the network to 512 MHz.

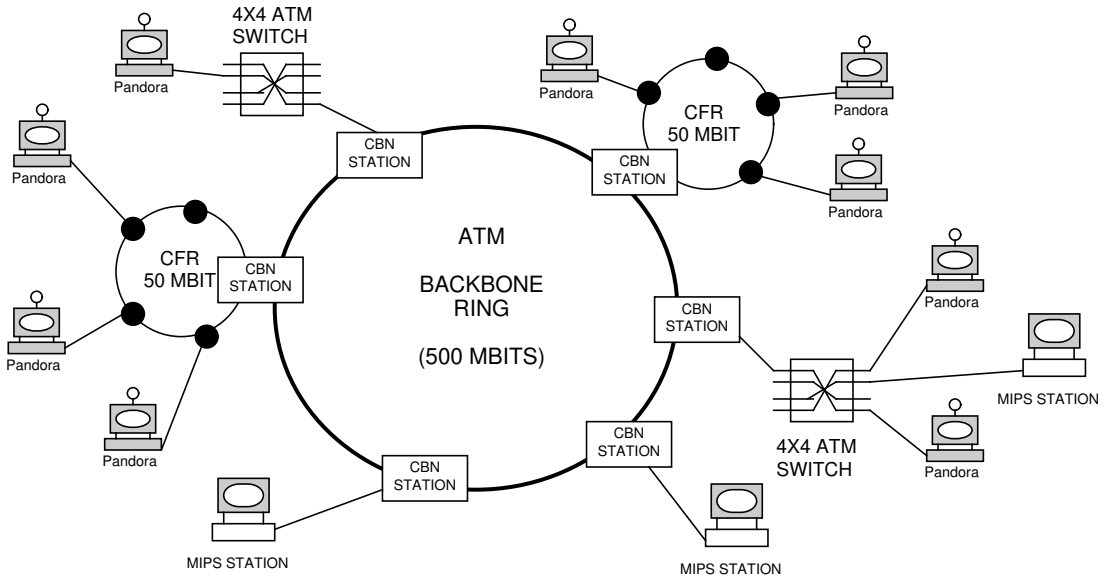

Figure 1: Part of the installed hardware in Cambridge. The Backbone Ring has run run between Olivetti Reaserch Ltd and the University of Cambridge Computer Laboratory since December 1990.

(switches and host interfaces) which uses 48 byte cells.

2 Access protocol for the Backbone Ring.

The Backbone Ring uses an access protocol know as MSR (multiple slotted ring). A slotted ring treats the electrical delay of the stations and fibre as a circular shift register which rotates constantly with a period equal to the total delay. This shift register is formatted into fixed length slots. One is shown in Figure 2. The frame format is set up and maintained by one of the stations which is the active monitor. Any station can adopt this role without interrupting its normal operation. The monitor functions include initially clearing the ring and ensuring that all stations have gained correct bit synchronisation, then writing out a frame structure of the correct length. The maintenance functions are clearing permanently full slots through the use of the monitor passed flag and passively monitoring the ring operation in case the format becomes corrupted.

A station with a cell to transmit waits for a slot with a clear full bit and then marks it full and fills in the header and payload fields of the slot. The CRC is appended by hardware to provide an integrity check. At any station, cells with faulty CRCs are marked bad and these are then discarded at the receiving station. For hardware and bandwidth efficiency, the CRC is shared over four cells, as described in [3] and [4].

Reception operates as follows: Stations monitor the VCI field of all[2] cells with the full flag set, looking up 16 bits of it in a 64K entry, directly indexed map-ram. If the map entry is appropriate, then the cell is copied into the reception buffers of the station. The map also contains futher information regarding the conditions for when to interrupt the host. In the MSDL datalink protocol of MSNA, receiving stations are able to allocate their own VCI's. Therefore for a new connection, a station can always allocate an entry in its map ram not currently in use.

The station which transmitted the cell is responsible for clearing the full flag when the slot again passes that station on the next ring revolution. This enables promiscuous multi-cast VCI's to be implemented, by simply setting the reception condition for the same VCI on several stations. When a station has freed a slot, it may not reuse it immediately, but passes it on free to the next station, thus providing fairness. A future enhancement would be to extend the map ram function to indicate whether to free a slot. This would enable spatial reuse of slots on non-promiscous VCI's, but at the expense of fairness.

[2] In fact, only one quarter of the VCI's need be examined owing to the four channel TDM system used, as described in [3].

Design and Implementation of an ATM Backbone Ring

David J Greaves, David R Milway and David Garnett
and Andy Hopper.

Olivetti Research Limited, Cambridge, UK.

Abstract

The Cambridge Backbone Ring (CBR) is a collaborative project between Olivetti Research Limited and the University of Cambridge Computer Laboratory. The project has designed and built an experimental ring communication network for computer data, and real-time applications. The network forms a backbone facility for interconnecting groups of CFR (Cambridge Fast Ring) networks and Pandora multi-media workstations in the Cambridge area. Access control uses the empty slot protocol, using short, fixed-length cells which contain 32 bytes of payload and a 4 byte header containing the virtual circuit identifier.

Although the architecture was designed for line rates of 1 Gbit/second and higher, the current implementation operates at 512 MHz. Twenty-five stations have been constructed, to date. A revision incorporating the now-standard, 48 byte cells, and operating at a higher line rate, is planned.

1 ATM local-area environment.

The Cambridge Fast Ring (CFR) is a previous project completed in collaboration between Olivetti Research Ltd, Acorn Computers and the University of Cambridge [1]. This was a slotted ring using 256 bit cells. It is still in daily use in Cambridge, and currently forms the interconnection network for groups of Pandora multi-media workstations [2]. The CFR is limited in the distance it can cover, both by the multi-mode technology used for its fibre links, and by its access protocol. The access protocol restricts each station to using only one slot per ring revolution. This places an upper limit on the point-to-point bandwidth available between stations inversely proportional to the number of slots on the ring. Clearly, the number of slots increases both if either the total fibre length and/or the line rate are increased.

This paper concerns the Backbone Ring, a new ring project designed to interconnect groups of CFR networks and for other applications. This paper mainly concentrates on the current hardware implementation of the Backbone Ring. The architecture has been described in full technical detail in [3], [4].

The Backbone Ring, like the CFR, is a slotted ring using 256 bit payload cells, but the Backone Ring access protocol removes the restriction of using only one slot per ring revolution. The Backbone Ring is designed as an ATM LAN or MAN, able to operate at line rates roughly between 0.5 to 5 Gbit/second, and with up to 200 kilometres of total fibre length. This implies a ring with hundreds of slots stored in the fibre, and hence this is the reason why the alternative access protocol is used.[1] The choice of 256 bit cells was made for upwards compatibility with the CFR and was fixed in the current version of silicon VLSI before the CCITT adopted 48 byte cells as the B-ISDN standard.

Figure 1 shows parts of the current infrastructure in Cambridge. The networks are split over the two sites (Olivetti and Computer Laboratory), but all equipment is part of a common ATM internet through the use of our multi-service network architecture (MSNA). MSNA is described fully in [5]. The equipment shown includes CFR, Backbone Ring, 4 by 4 ATM switches, Pandora multi-media workstations and MIPS workstations using the DEC Turbochannel interface.

MSNL, the network level protocol, is currently the most visable part of MSNA, since this provides transparent shipping of ATM cell payloads between MSNL-service access points, setting up the required concatenation of virtual circuits across the intermediate set of ATM switches and rings. MSNL will also indicate to the end points the size of cell to use on the circuit (32 or 48 bytes) so that the 32 byte cells can sensibly be integrated with the more recently designed equipment

[1] The new protocol cannot easily use a MAC-layer response mechanism for automatic retransmission. However, the Backbone Ring stations are provided with sufficient receiver buffering for retransmission due to receiver overloading not to be required. Other uses of the MAC-layer response mechanism are not appropriate when vusing ATM style higher-level protocols.