The Cambridge Ring

Acorn Computers Ltd

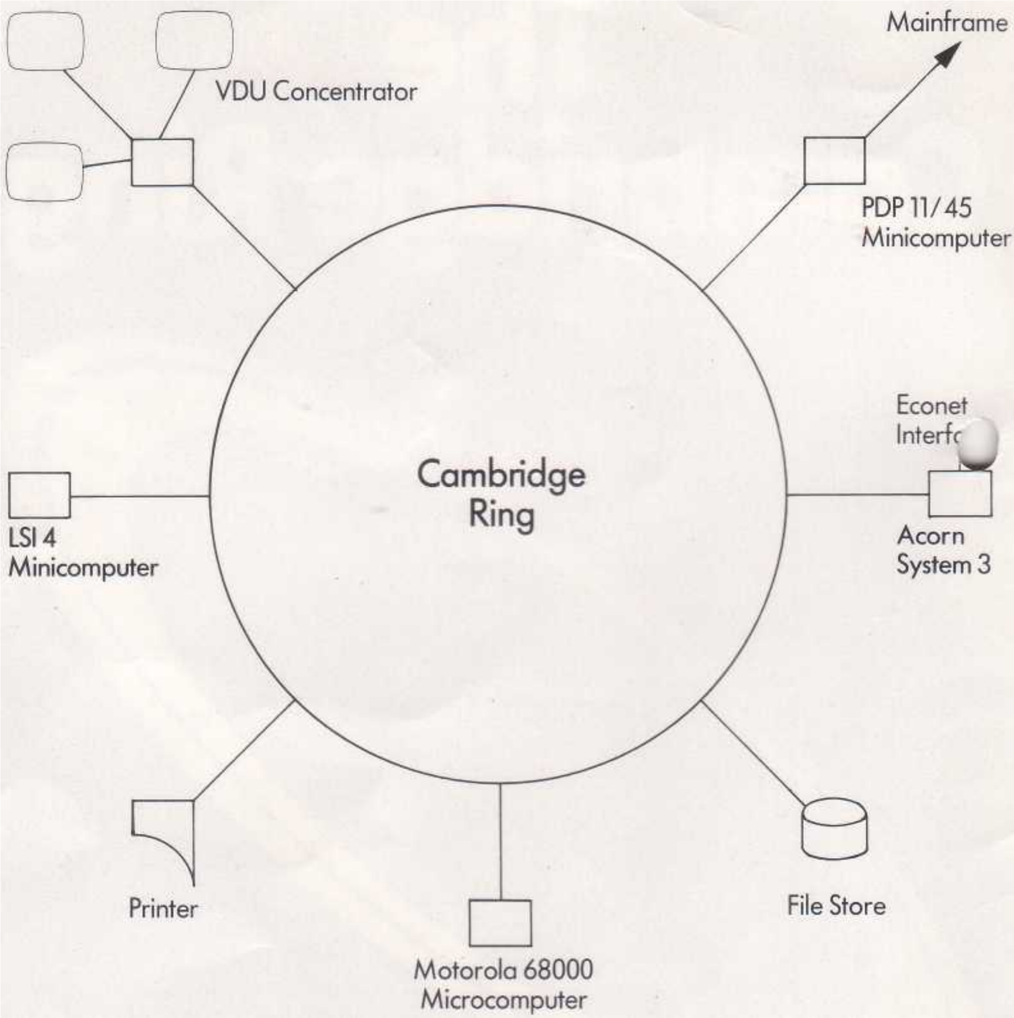

A typical Ring application.

To optimise the cost effectiveness of any computer system it is necessary to link computers with each other and their peripherals. The schematic diagram above illustrates a typical application.

For effective computer communications, four criteria must be satisfied — high speed, low cost, high reliability and versatility.

High speed so that the effective operation rate of the machines in communication is not reduced because of their shared data path.

Low cost so that it is economical to use the system to link even small microcomputers to mainframes and peripherals.

High reliability for precise data transfer without a heavy software overhead for error correction.

Versatility so that many different types of machine may be introduced.

The speed of the Cambridge Ring is in the order of a data rate of 10 million bits per second (10 Megabaud) which, when compared to typical data throughput rates of 9.6 kilobaud, 500 kilobaud,1 megabaud and 5-10 megabaud, for peripherals, microcomputers, minicomputers and mainframes respectively, illustrates the feasibility of networking without losing speed.

Costs are kept to a low level by the combination of elegant hardware design employing custom logic arrays and the use of standard interstation cabling of the sort used for teletypes. For example the cost of a node and 200 metres of connecting cable is approximately £300 (excluding computer interface module).

High reliability is achieved by providing for three powerful independent levels of error detection within the ring itself, and a node design that ensures correct operation of the Ring regardless of what happens to the stations connected to it. This system has been proved in use over the last six years at the Cambridge University Computer Laboratory.

Versatility is ensured by the use of intelligent work stations which provide data in a form readily understood by interface modules to most popular mini and micro computers.

The Ring system designed at the University of Cambridge Computer Laboratory provides a high speed, low error rate communications path between computers, microprocessors and other devices. The primary uses of the ring are for equipment sharing, file sharing and dumping, and for implementation of distributed systems.

The structure of a node is as shown in diagram 1 below. In addition to repeaters, *stations* and access boxes there is a unique station called the *monitor station.* This station is used for setting up the slot structure during turn on, for monitoring the ring and clearing lost packets, and for accumulating some error statistics. Immediately preceding the monitor station there is an error logging station which uses a normal station and receives packets containing error information sent from active ring stations and monitor station.

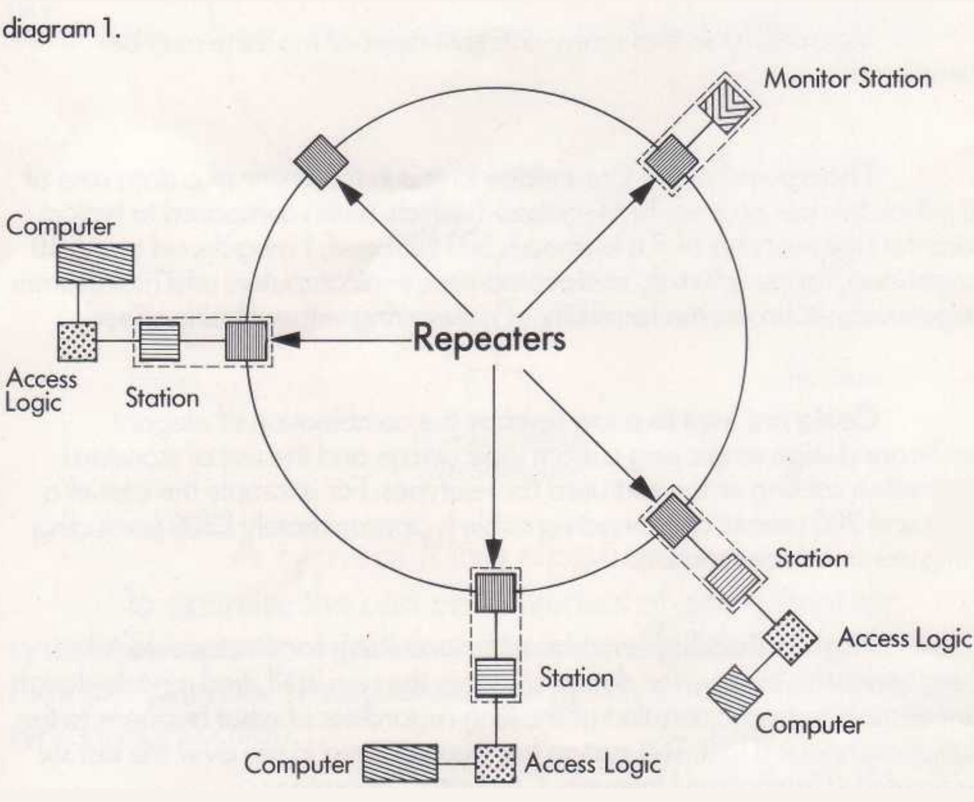

Cambridge Ring Structure.

The packet structure is shown in diagram 2 below and is chosen to allow the maximum timing tolerance and minimum delay at the transmitter and receiver. The leading bit is always a one and is followed by a bit to indicate whether the slot is full or empty. Now follows a control bit used by the monitor station to mark as empty packets which are circulating indefinitely due to an error in the full/empty bit. This is followed by a number of eight-bit bytes the first two of which are used for destination and source addresses and the rest (maximum 8) for data. Finally there are two user control bits, two response bits used for acknowledgement purposes and a parity bit for ring maintenance.

diagram 2.

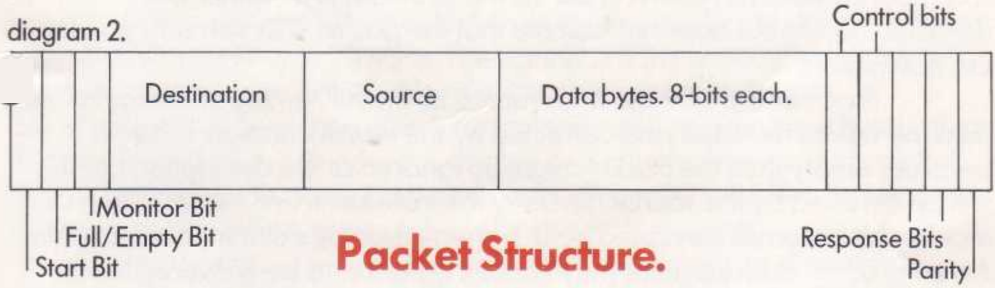

Packet Structure.

When a station has a packet ready for transmission in its shift register it waits until the beginning of the next slot. It now reads the full/ empty bit and at the same time writes a one at the output. If the full/empty bit was a zero it transmits the packet, if however the full/empty bit was a one the slot is already occupied and the algorithm is repeated for the next packet. This scheme minimises the delay at each node.

The transmitted packet makes its way to the destination where the control bits are set on the fly to indicate accepted, busy, or rejected. It now returns to the source where the slot is marked empty. If the packet returns with the control bits unchanged it was not recognised by any destination. Each station knows the total number of slots in the ring and can thus clear the full/empty bit immediately.

Thus on transmission the packet is delayed until an empty slot is found but then the transmission is rapid. As the destination does not explicitly signal when it is ready to receive the next packet the ring can easily become clogged with packets returning marked busy when devices with varying speed characteristics are being interconnected. In order to overcome this the following algorithm is incorporated in the hardware to reduce the number of busies. If a source transmits a packet and it returns marked busy then it is not allowed to retransmit it until some time later. This additional delay is dependent on ring loading and is approximately the time to aquire the next empty slot. If any further transmissions are attempted the extra delay is increased to about 16 times the original delay. Thus the number of busies is decreased and performance is improved.

Each station possesses a station select register which is initialised by the host. This register can be set to accept or to reject all packets, or to receive from one source only. When combined with a time out mechanism it can be used to allocate resources on the ring.

Maintenance and Error Recovery

There are no end-to-end CRC or parity checks on transmitted packets, however, a parity of the information is retained at the source and is compared with the returning parity. This provides a powerful error detection facility but does not indicate that the packet was correctly copied at the destination.

If one of the SOP bits is corrupted or the full/empty bit becomes full then this will be detected and corrected by the monitor station. If full becomes empty then the packet might be ignored at the destination but this will be detected by the source. Similarly the transmitter will detect if the monitor bit becomes corrupted in such a way that the slot is marked empty. An error in the address fields may cause the packet to be delivered incorrectly or be assigned to the wrong source. An error in the response bits might have a more serious effect as it will not be detected by the transmitter, which might repeat the packet or assume it was received correctly when this was not the case. Errors are generally detected by the source or monitor station within one ring delay.

The Cambridge Ring provides powerful maintenance facilities for localising transient faults and ring breaks. This is done by using the parity bit at the end of the packet. Each station continually computes the parity of every passing packet and empty slot. This parity is written into the parity field, which is simultaneously sensed, thus not increasing the delay. If the new parity does not match the old, a fault has occurred. Since the correct parity is inserted at each station, the fault must have occured since the last active station. Each station requires two parity circuits, one to check the incoming packets, and one to produce a parity for an out-going (and perhaps changed) packet.

Thus, the ring is continuously monitored for errors and each detected fault is located to the nearest active downstream station. The error information is then transmitted to the monitor station by inserting into the next empty slot a packet which is entirely zero apart from the source field and the full bit, and a bit in the data field to indicate the maintenance packet type. Thus the monitor station bit is also cleared so that this packet can be sent independently of the transmission shift register and is removed by the monitor station. The fault message itself may become corrupted, giving rise to further valid fault messages; never-the-less the indicators reaching the monitor station will at least be correct for the nearest fault.

The scheme is also used to detect ring breaks. This is done by arranging that the phase-locked loop at each repeater continues to operate in the centre of its frequency range when there is no data at the input. This is interpreted as a string of zeros, and the station is made to give a repeated fault message packet to which other repeaters downstream synchronise. Thus, when the ring is completely broken and unable to be used for data transmission the forward path is used to send fault localizing messages.

Hardware

The ring is built using TTL technology and operates at 10 MHz with a maximum distance of 300 metres between repeaters. An enhanced version is available operating at 20 MHz. The signals are transmitted along twisted pairs of the type normally used for duplex operation of teletypes. Transformers are used throughout for isolation and common mode rejection. As the repeaters have to operate reliably and whether they are connected to a station or not, they are powered directly from the ring. This power is injected into the system at a number of independent points.

Each station is fully duplex so that it can transmit, and receive, concurrently and independently. The number of bits delay at a repeater is three bits and the minimum ring delay is about 5 microseconds.

A number of different access boxes have been developed where complexity varies according to the speed of the host, and according to the level of buffering. Access boxes can be grouped as implementing three types of interfaces; polled, interrupt and direct memory access (DMA). A polled interface is the simplest and is suitable for attaching microprocessors to the ring. Each control signal at the hosts station interface is memory mapped in the microprocessor and data transmission and reception consists of writing and reading from memory. Such micro-processors have been developed on the Cambridge Ring as printer controllers, VDU concentrators or where a simple data processing function is required. An interrupt interface is the simplest way of connecting a minicomputer. If the interrupt is at the macro-program level the data transmission rate may be low. However, this may be remedied by servicing the interrupt in the microprogram. DMA operations can be performed by a simple micro-sequence or by a more sophisticated device which buffers a number of ring packets and writes them to store in one block. This is suitable where high performance is required, and systems of this kind have been developed based on the 8X300 and 6809 microprocessors.

SOME QUESTIONS ANSWERED

Q What happens in the event of workstation failure? Won't that affect the whole signalling system and clog the ring?

A The workstation and repeater are powered separately, with the repeater powered by the same source as the ring. No matter what happens to the workstation the repeater will continue to function independently.

Q What system is used for interfacing stations and devices?

A Interfacing is by means of custom logic. For the simplest systems this consists of about half a dozen chips. For more sophisticated systems this can be extended to provide buffering and protocol functions.

Q Are synchronous communications between nodes possible?

A The small packet size on the Cambridge Ring provides very low latency for certain types of communication. Synchronous communication between a pair of nodes can easily be interleaved with other transactions on the ring.

Q If a packet is returned marked 'busy,' what happens if further transmissions are attempted by that station?

A The delay is increased to around sixteen times the original delay, decreasing the number of 'busies' and improving performance.

Q The protocols appear to be very simple. Are they really capable of dealing with errors?

A A feature of the system is the high bandwidth and low error rate. The error rate is of the order of one error in 5 x10'^{1} bits. Protocols can thus be simple because faults are infrequent and can be handled by repeating at a high level.

Q What machine interfaces are possible?

A Interfaces (access boxes) to the following exist already—DEC PDPII, LSI II, CA LSI 4, NOVA, 6502, 6809, Z80 and 68000 microprocessors. Many others are under development.

Q What happens if the Ring is broken? A The repeater after the break reports the fact to the monitor.

Q Can I link two or more rings?

A Yes, a gateway can be constructed wiring a computer with two access boxes.

Q Can I use the ring in adverse electrical environments? A Yes, using screened cable the ring has a high immunity to noise.

Q Can I use fibre optic links? A Yes, as the Ring is a point to point system it can utilise most transmission media.

Q Can I transmit speech on the Ring? A A single Ring can support several hundred telephones in normal use.

ACORN COMPUTER

For prices and further information write to: Acorn Computers Limited, (Orbis) 4a, Market Hill, Cambridge CB2 3NJ.