S.S. Thomas.

Computer Laboratory
University of Cambridge, UK.
27 4 1978

GENERAL OUTLINE OF RING

Introduction

Various computers and devices are connected by a communication ring, arranged so that each pair of computers or devices can communicate with each other. The ring consists of a number of links going from one station to the next, and the messages are transmitted in blocks of packets of 37 bits each of which contains 16 or 8 message bits.

As each link is short, the data rate is relatively high but the system is vulnerable to errors in the repeaters at each station. Each packet passes through all the repeaters so that a single faulty repeater can affect all transmissions.

A packet is transmitted from the transmitting station to the receiving station and returned to the transmitter with a mark to indicate rejected, accepted, ignored or busy. The logic at the transmitter then takes proper action.

The ring itself conveys the packets in a serial manner, and as each packet has a source and destination number the ring is at the same time carrying packets from a number of sources to

their various destinations.

The bit speed of an ring is about 10 megabaud and the empty packet scheme is used for transmitting packets. In this scheme the complete packet is loaded into a shift register and the shift register is used to fill the first empty packet that passes the station. The packet is marked when it passes the destination and as it returns, it is compared with original packet and the passing packet is marked empty.

There are a number of consequences from this design. Firstly, for any two communicating stations the number of packets in flight around the ring is one. Secondly, the packet will be delayed before starting until an empty becomes available, but then the ring transit is swift. Thirdly the minimum delay at a station is a fraction of a bit, so that the delay around the ring is minimal. Forthly, the percentage of time a station is interacting with the ring is minimal so that reliability should be reasonable. Fiftnly, the traffic is self regulating and needs no central control. The passage of the empties causes a "round robin" system to be in force, so that no station can block the ring and each station affects all the other stations equally.

The monitor station is only needed during turn on and to monitor the ring and take action after certain types of error.

Design considerations

Experimental evidence indicates that the standard teletype cables which crisscross the laboratory will allow about 10

megabauds without particular precautions against noise, differential attenuation etc. as long as the signal is regenerated fairly frequently, say every 100 metres.

The cable system has to transmit clock, data, and power. It seems simpler to have a clock in a small system rather than an assynchronous system, which would require highly stable time sources at each repeater, and careful design to prevent one packet catching another up around the ring. The power is needed to keep each repeater working when the local power is switched off. It merely has to energise enough of the station to ensure that data is regenerated as it passes the station.

The laboratory teletype circuits usually are used in pairs so the choice of 1,2,4 pairs for the ring has to be taken. The four pair choice could, for example, send data, clock, power, and errors or control along separate pairs. The single pair choice would require that power, clock, and data were unraveled at each repeater. We have chosen to use two pairs and to use them with a special symmetrical self clocking scheme which uses both pairs for clock and data and sends a dc of about 50 volts around the ring as the power source, one pair being negative the other positive and both halves of each pair taking the current in parallel.

The modulation system amounts to using phase modulation along each pair with their digits half out of step with each other. However, the rules for both pairs are easier than phase modulation and the system cannot be half a bit out of synchronism. At each clock instant a change on one of the

pairs indicates a zero and a change on both the pairs indicates a one. The single changes use the pairs alternately. Thus the clock is recovered readily as there is always at least one change at each clock instant.

The system is completely balanced so that transformers are used at each repeater to isolate the dc power and reject noise and common mode interference from each pair. The repeater is designed to demodulate and modulate the ring signals, so that the interface between a repeater and station is very straight forward, and independent of the actual modulation system. Thus logically, different parts of the ring can use different modulation systems. In fact one section of the ring uses two optical fibres in a cable supplied by GEC. The attenuation and dispersion of the optical fibres is so low that this link could be up to some kilometres in length, witnout needing any intermediate repeaters.

The mode of operations between stations is for a transmitter to send a packet to a receiver, which copies it and marks it 'accepted'. When it returns to the transmitter the packet is marked empty.

The receiver must be able to mark a packet as it passes, and be able to recognise the destination address. The transmitter must recognise when the packet returns and mark it empty. If a receiver ignores a packet, perhaps because the station is switched off, then the transmitter will note the lack of a mark when it stops the packet.

Packet structure

At least one packet needs a front marker digit so that it is distinguished from gap digits. It needs one bit to indicate full or empty. It needs a source address and a destination address, and also space for data which has been fixed at 16 bits. It needs space for response bits to indicate accepted etc. and finally it is useful to have one extra bit which is changed as the packet passes the monitor station, so that the packet is deleted if it passes the monitor station again without the transmitter having picked it up. This allows simple deletion of unwanted packets by the monitor station.

It is best to place the source and destination address in front of the data so that the receiver has ample decoding time. The response bits are at the end of the packet to allow the maximum time for the receiver logic to decide whether to accept or reject the data. It seems best to put the monitor passed bit early in the packet as this enables a packet to be converted from a wrong packet to a null packet without requiring any long shift register in series with the ring at the monitor station. The full or empty bit has to be first to enable filling to be started without delay, which would otherwise need a delay in series with the ring. Thus the packet structure chosen is as follows:

Bit	Purpose
1	start of packet. used for packet counting
2	full/ empty

3	monitor passed digit
4-11	destination address
12-19	source address
20-36	data
36-37	response bits used by the receiver to indicate rejected, busy, accepted, or ignored.

The first bit of a packet is used to facilitate framing and to allow automatic sizing of the ring at each station. The packets are grouped into a train of packets terminated by a gap containing zeros. This allows each station to find out how many packets are circulating in the ring. The number of packets circulating is set by switches in the monitor station, which also detects, counts, and corrects framing errors and gap errors.

Traffic considerations

A transmitter can repeatedly send data to a receiver which ignores it. The fastest rate is attained with an immediate retry i.e. after a busy response which will take a ring transmission delay which of course depends on the current traffic. If we assume there are m active stations on the ring endeavouring to transmit as fast as possible, and then if a station endeavours to transmit a packet by ignoring the busy and continually sending the packet until it is accepted, then it will reduce the useful traffic by 1/m. If we consider that some devices may endeavour to drive teletypes in this way we can see an appreciable proportion of the total traffic may be

useless. The m stations active may of course represent perhaps 10m stations each transmitting at its own full rate but only 1/10 of that of the ring. Thus a few bad stations may well clog the ring with useless traffic.

The station logic is designed to mitigate this problem by delaying the response if it is busy, rejected, or lost, to about 2 ring delays, and to about 16 ring delays for the second and further tries. Thus, the jamming capacity of a badly designed access logic or computer program is reduced by a factor of about 8. This delay will not influence the useful traffic very much as the acceptance is not delayed at the station.

It should be noted that with immediate retry, the traffic depends on the number of transmitters and not their speed, so that capacity is drastically reduced by a large number of slow stations.

Errors and Delays.

There will be errors in the ring and it is desirable to be aware of their effect. A packet can become erroneous and the may affect any or all of the digits. The leader digit can be changed. This will upset the counts at all stations but can be corrected by the monitor station.

If the full or empty digit becomes changed, this can cause a full packet to circulate for ever or until the monitor stationremoves it. If full becomes empty, this may cause the receiver to ignore it , but in any case this will be detected

at the transmitter which will give an error response to the access logic. however it is possible that another station will fill the falsely emptied packet. This packet will be emptied as it passes the first station, so a third station may fill it before it reaches the second. This can continue indefinitely, and although the transmitters will detect the changed packets, unless the traffic is sufficiently low for the faulty packet to complete a circuit having been filled and falsely emptied. However, the incidence of faults at the transmitters affected should reduce the traffic. If empty becomes full, an eroneous packet will have been created which may be accepted by a receiver, but will not be emptied at a transmitter. However the monitor station can delete this as it passes for the second time.

An error in the source or destination, may cause the wrong receiver to collect or for it to assign it to the wrong source. A receiver which is accepting from a single source will not be affected. The transmitter will detect the returned packet as faulty, and give an error signal. An error in the response bits is more serious, as the transmitter may repeat a packet if for example accepted becomes busy, or unselected. The transmitter may not repeat a packet if busy becomes accepted.

It can be seen that the effects of most errors are limited to one or more packets, and that in many cases the error is detected at the transmitter, so that unnoticed errors may be rare.

It will be seen that the monitor station can help in

limiting the propagation of errors and keeping a count of the errors it detects.

It is desirable that the monitor station should force the leader bits into the ring at each circulation and count the number that disappear and that it should convert any packet which passes for the second time, or is empty into a packet of all zeros apart from the leader digit. A count should bekept of those packets which passed for the second time.

It may be well to include somewhat more logic, so that faults can be analysed more easily. For example, perhaps the last one or two erroneous packets should be stored, so that intermittant errors can be more readily localised.

The design has split the logic into a number of separate parts; at each station there is a repeater which is powered from the ring and modulates and demodulates the data around the ring. This is isolated logic and sends data and clock to the station logic. It will be isolated when the station is turned off so that the integrity of the ring depends only on the switched on station units in addition to all the repeaters. The station unit modulates and demodulates the packet into a parallel interface for the access logic which depends on the partic ular device connected. It is possible to use different schemes for modulation and demodulation merely by changing the type of repeater for a particular section.

The station

The station is connected to the ring by a repeater, and to

the local device by the specialised access logic. It is designed on the basis of a 37 bit packet, which is copied into a passing empty packet and goes to the receiver which marks it, and then returns it to the transmitter, which frees the packet by marking it empty. The packet has a destination address which can be set by the access logic and a source address which is inserted at the transmitter. The receiver has to compare the destination address with its own, check and see if it is from a selected source, and mark the packet as it passes.

Packet structure.

A packet has a leader bit which is 1 and 36 usable bits.

The first bit is 1 when the packet is full

The second bit is 1 when the packet is filled, but changes into zero the first time past the monitor station, and is used by the MS to delete all bits of a packet if it is zero i.e. the second time past the MS since it was filled.

The next 8 bits give the destination address.

The next 8 bits give the source address

The next 16 bits are data.

The last two bits are set to 1 at transmission and used by the receiver. They are set to 00,10,01 according as the receiver is busy, unselected, or accepted if the packet is copied.

The addresses 0 and 255 are reserved for control purposes, but only 0 is now used..

Bus connections to the access logic.

The station is designed as a duplex station. That is, the receive and the transmit parts are independent and may be used simultaneously and independently.

The transmit part has a 16 bit two way bus which is controlled entirely from the access logic by gating signals. These allow data and destination to be set, and replies sensed. As there are independent gates which control each 8 bit half of the bus, these may be commoned to reduce the number of wires needed.

The receive part is similar, but also enables a selection register to be set so that packets can be limited to those received from one source, no source, or all sources.

The two buses can be commoned together, but again subject to the constraint that no two reading gates should be active simultaneously.

Station to access logic interface

There are two 16 digit 3 state buses. The gates onto the buses are either controlled directly in the access logic or indirectly by gating signals which are are sent along the interface cables. The gate signals are uncoded so that each signal directly controls one gate in the station. This does allow the buses to be grouped so that fewer wires are needed at the expense more gating steps.

The control signals are all negative going signals such that the active state is the zero state, this convention has the advantage that an absent cable will not generate spurious

signals.

The transmit bus is named TB0-15 and the receive bus RB0-15 numbered such that signal 0 is the least significant and corresponds to the earliest serial signal along the ring.

Control signals	Action
CS1	Sets the 8 bit source acceptable from RB0-7
CS2	Gates received data 0-7 to RB0-7
CS3	Gates received data 8-15 to RB8-15
CS4	Gates received source address to RB0-7
CS5	Gates source acceptable register to RB0-7
CS6	Gates read status to RB6,7 . RB6=1 if station has rejected a packet from an unselected source since the selector register was last set.
CS7	Echo response to access logic. The or of CS1-6,16 used for handshake gate timing control
CS8	Set data TB0-7 in transmission packet
CS9	Set data TB8-15 in transmission packet
CS10	Set destination address in packet from TB0-7
CS11	Gate status to TB0-4

TB0=0 if status is busy

TB1=0 if status is unselected

TB2=0 if status is accepted

TB3=0 if status is ignored

TB4=0 if status is erroneous packet returned

These responses are valid when CS14 has gone down.

CS12 Echo response to access logic. The or of CS8-11,13

used for handsake gate control.

CS13 Transmit command. Will transmit after setting source and control bits.

CS14 Transmit reply. This signal will go to one after the end of CS13. It becomes zero again when the transmitted packet returns and the response bits are again valid.

CS15 Received packet available signal. This signal will go to one at the end of CS16 and become zero again when the receiver register has again been loaded.

CS16 Receive command to the station. The current packet is discarded.

CS17 This signal indicates the gap preceeded by an empty packet so that on the average it indicates the average delay round the ring.

Note. If gates CS8-10,13 are given while the previoustransmission is not complete tney are ignored.

Station to repeater interface

The signals are listed below.

-DS Serial data to station

DR Serial data to repeater

-CS clock to station. Timing on ve edge.

GR Gate to repeater to cause the station to be bypassed.

EMS4 This signal changes polarity each time a special monitor signal is received.

The station may be considered as made of the following parts

1) Framing logic. This locates the passing packets and gaps, and counts the number of packets between gaps.

2) Receiver logic. This recognises packets sent to it marks them accepted, busy, etc. and copies the packet and signals the access logic.

3) Transmission logic. This accepts packets from the access logic.
and awaits a passing empty packet, fills it, awaits its return, compares it with the original packet, and marks the packet empty and will signal the access logic with the response.

4) Echo logic. This will mark special packets from the monitor station and signal by a low bandwidth ground signal, so that ring breaks can be located by the maintainance logic of the monitor station.

5) Delay logic. This delays responses other than accepted to the access logic so that unintelligent devices cannot swamp the ring by continual retries.

The monitor station

The purpose of the monitor station is to give means of starting the ring and limiting the propagation of some types of errors. There are two main modes of operation namely, starting and running. The starting mode is automatically

forced when the ring is turned on and the running mode is then entered. The starting mode sets up a packet structure, reads the switches giving the packet count and also counts and holds the number of gap digits. The transition to running mode will clear the error counters and indicaters, and will then maintain the standard framing structure.

The structure consists of a fixed number of packets each with a leading digit of 1 and 37 digits in total. If the second or third digits are zero, then the entire packet other than the leading digit is cleared to zero. The data returning from the ring is checked against allowable structures, and certain errors are detected. The loss of the leading digit, ones app pearing in the gap, and a full packet circulating past the monitor station for the second time, cause a indicator to be set and one added to a common counter.

The packet structure is refreshed as it passes the monitor station station.

A number of maintainance aids have also been included in the monitor station. It is possible to issue arbitary packets, or have them filled with random data and addresses and they will be checked as they return. The error count and indicaters will detect simple faults, and will allow long unattended tests. The station can also issue packets to fixed destinations with zero source and data, so that the echo signal can be stimulated from individual stations. This is independent of the returning data so that this can be done when the ring is broken. This is done by setting a control switch and the appropiate address on

the load data switches.

The repeater

This takes phase encoded signals from two pairs of twisted pairs and demodulates them into a clock and data signal which is supplied to the station logic. It accepts gate and data signals from the station (synchronised to the clock) and combines these with the incoming data and then modulates these onto the output twisted pairs.

It also contains a DC transformer which takes power supplied along the twisted pairs and generates its own supplies for its circuits. Thus it is independent of local power. The gate signal from the station (which is the only destructive signal) is protected by an open reed relay when the station power is off.

A phase controlled oscillator is used to increase working tolerances (the system is self clocking) and a special circuit can pass a low bandwidth signal back to the monitor station for maintainance purposes.

The repeater is isolated from the ring by four small transformers so that so that common mode induced voltages cause no problems.