SCIENCE AND ENGINEERING RESEARCH COUNCIL
RUTHERFORD & APPLETON LABORATORIES

COMPUTING DIVISION

D I S T R I B U T E D C O M P U T I N G N O T E 5 0 3

CAMBRIDGE RING STANDARDS
Report on the Ring Standards Meeting at RAL on 28 Oct 81

Issued by
D A Duce

9 November 1981

DISTRIBUTION:
F R A Hopgood
R W Witty
D A Duce
C P Wadsworth
W P Sharpe
Technical Meetings/Ring Standards file
RL Support/Cambridge Ring/Working papers file

(new page)

IN CONFIDENCE

SCIENCE AND ENGINEERING RESEARCH COUNCIL

RUTHERFORD APPLETON LABORATORY
Computing Division

Notes on the Ring Standards Meeting held at Rutherford Appleton Laboratory on 28 October 1981.

Present

W P Sharpe, RAL (Chairman)
R Battell, GEC Computers
G Wilson, GEC Computers
M Cole, Logica VTS Ltd
S Wilbur, Logica VTS Ltd
Dr A Hopper, Orbis Computers Ltd
A W K Erasmuson, SEEL
G Strachan, SEEL
Dr K Heard, JNT
A R Cash, RAL
D Drury, Camtec Electronics
R D Morley, Toltec Computers Ltd
C Burns , Toltec Computers Ltd

Dr D A Duce, RAL

Apologies were received from the Computer Laboratory, University of Cambridge. The Laboratory do wish to be involved in the Standards Meetings but unfortunately could not send a representative to this meeting. It is hoped that Professor Wheeler will be involved when he returns from a period of sabbatical leave after Christmas.

1. INTRODUCTION

Mr Sharpe welcomed everyone to the meeting, and welcomed the level of interest and desire for cooperation shown. He stressed the urgency of completing standardisation work.

Mr Sharpe indicated that the Department of Industry had welcomed the establishment of this group and would be happy to give whatever stamp of approval necessary at the end of the exercise. The Department had also indicated that they would investigate the possibility of making available any necessary funds for the exercise.

Dr Heard said that the FOCUS Steering Committee for Local Area Networks at their first meeting recently had expressed pleasure about this initiative.

2. SUMMARY

Four Areas for work were identified:

1. Specification of area of overlap between existing systems at the repeater/wire level.

2. User requirements at repeater level (long hops, safety requirements etc).

3. Specification of existing (a) 100 and (b) 50 way interfaces.

4. An adequate specification for a 50 way interface and station functionality.

Individuals agreed to progress these areas as follows:

1. Tony Cash

2. SEEL and Logica to write to Bill Sharpe expressing their considered views in undertaking this work.

3(a).Andy Hopper,

3(b).Logica

4. Logica VTS are pursuing this area under funding from DoI. Mike Cole to circulate members of the group with a summary of the objectives of this work.

Timescales were agreed as follows:

1. Questionnaire circulated by Tony Cash at the meeting to be returned by FRIDAY 6 NOVEMBER. Replies to be collated and draft discussion paper circulated (about 2 weeks). Replies to that to be received within one week.

   Notes:

   i) Replies to questionnaire to be sent to Mr W P Sharpe, RAL

   ii) indicate clearly whether information is Confidential. Default is that information will be treated as confidential and NOT circulated to other members.

3. Document to be produced in 10 days, then circulated by Bill Sharpe. Comments to be submitted 1 week after that.

Mr Sharpe said there were two main levels of interface in the ring, the wire level and station/access logic. He suggested that the way to proceed was to agree on the main issues to be tackled at each level, to establish a timetable for the work to be done and then delegate responsibility to individuals to undertake the work. There was agreement on this procedure.

3. WIRE LEVEL COMPATIBILITY - A R Cash

3.1 Presentation

Mr Sharpe introduced Mr Cash. Mr Cash was responsible for the construction of 6 Cambridge Ring systems for SERC's Distributed Computing Systems Programme (to Mr Wilbur's UCL designs) and subsequently was heavily involved in the preparation of a tender for ten further Rings purchased by the DCS Programme and acceptance testing of the chosen product.

A Standard should enshrine the best that can currently be offered, this implies that information on all current products is required at the outset. The Rutherford Appleton Laboratory will be happy to coordinate information to be used for the specification of a standard.

Product information might be of a commercially sensitive nature. All information gathered will be treated as Commercial in Confidence if suppliers so require, though it was not intended to encourage suppliers to regard product information in this way.

A questionnaire has been prepared which was distributed to participants.

The questionnaire concentrated on the ring and repeater level specification. The station access logic problem is to be considered separately.

Some questions may not be applicable to a given system, if this is the case then the supplier is asked to state why.

The questionnaire will be used as a means of gathering information and will result in the circulation of a draft standard discussion paper to all members for consideration firstly in-house and then at a joint meeting of the group.

It was agreed that the major contributors should be (in alphabetical order):

Cambridge University
Logica VTS
Orbis
RAL - Coordinators
SEEL
Toltec

3.2 Discussion

Wilbur: One should look at the wider issues of what the exercise is trying to achieve. There is a policy issue which needs to be decided as there are two distinct issues at the line level:

1. Compatibility of existing implementations - soluble in some sense.

2. Specification of a line level interface that will carry into the future and deal with long hops, noisy environments, standards that have to be complied with.

Issue 2 should be considered as a separate stage which will need research or consultancy work to establish a working engineering standard.

Cole: It is prudent to consider goals and non-goals first, as is done in the Xerox Blue Book.

Wilbur: Different companies will have different goals, for example Logica look for installations in a wide variety of environments, a flexible interconnection structure, ease of maintenance and installation etc.

One of the problems is constructing a specification at the wire level which is appropriate to different environments and technologies, for example fibre optics. I would like to see a standard, but worry whether it should be a single standard or family of standards.

Drury: One could have an electrical to fibre optic convector, invisible to the electrical standard.

Wilbur: Yes, but one should recognise optical fibres are a great asset to the ring.

Heard: One could have a specification at the functional level, mapping onto physical media. One could define one or many such mappings.

Sharpe: I do not like defining inputs and outputs and leaving something in-between undefined.

Drury: Aiming for six standards in parallel will result in the achievement of nothing.

Cash: There is a problem exemplified by, say, the attempt to define the performance of the transmitter and receive modulation schemes in that one might very well change the modulation scheme if using fibre optics.

Wilbur: One worry at present concerns radiated noise from twisted pairs. In the present system this is far in excess of what could be exported to the USA or Germany. Installation practice is also a worry in that say sophisticated test equipment should not need to be used to install the cables.

If one breaks down what the standard has to do, it will have functional, electrical and mechanical attributes. Electrical attributed will differ in different environments.

The functional interface is fairly clear, but it does impact on the phase lock loop design, whether one has an open or closed loop. Fairly extensive technical work needs doing to resolve these issues.

Sharpe: Please elaborate.

Hopper: I have done a lot of work on phase lock loops recently! It is not understood how the old design for the phase lock loop works, though one has one's latest theories. It is very important to specify the phase lock loop in terms of the filter used, or even better the rate at which pulses get fatter and thinner, and time constants in terms of the period of the whole ring.

There is a real puzzle surrounding the S124 chips used in the Mark II design. It is not understood what is happening inside these.

IBM have recently done some work with phase locked loops; in their ring it is possible to cut our repeaters and the ring will still lock, though the locking mechanism is fairly non-homogeneous.

Sharpe: Logica changed the design of the phase locked loop; could they explain what was changed and why.

Wilbur: This was approached from engineering considerations, supply of components etc.

Having done some experiments, we discovered some instability problems which do not seem to have been so apparent in the old Cambridge design. These problems always exist in second order phase lock loops and so we moved to a first order design which is unconditionally stable. We carefully tracked noise rejection against capture range.

Thought was also given to open versus closed loops and we too subsequently come across the IBM paper which describes a carefully designed third order open loop.

Hopper: The current chips use a second order design. There are a number of trade-offs to be considered - the lock range can be improved by making the oscillator more flexible. Many designs can be made to work, but fitting them together is a puzzle. Maybe one should fix on one design, but specifying the fitter alone is not sufficient, gain etc need to be specified also.

I think something can be aimed at, but it needs care and experimentation so that different systems can interact.

Cash: Mixing phase lock loops can cause trouble.

Heard: Are we concerned with a minimum subset allowing inter-working or a specification which each system adopts.

Wilbur: By sacrificing the phase lock loop specification, noise immunity of the system is worsened. First and second order loops inter-working will be a less than optimal design. To achieve good noise immunity requires a tight specification on all phase lock loops.

Cash: We are in a nasty situation having started this process late.

Drury: Are we talking about the future ring, or the 'now' ring and compatibility issues?

Sharpe: The 'now' ring definitely.

Wilbur: If we were designing a new low level interface we would want a much tighter specification. Noise rejection properties are good enough now for most environments.

Heard: We must address the 'now' ring whilst deliberately seeing where points of conflict and issues are.

Cash: Future depends on the technology used also. The actual specification has to wait until the technology is decided.

Drury: My own background is in CAMAC systems. CAMAC avoids modulation problems by specifying a logical interface in terms of bit handling and timing so that the modulation technique can be varied. We need agreement at the bit level.

Hopper: Taking that approach you then have to specify the clock which is equally difficult.

Wilbur: We started by looking for compatibility at this level, but then parameters such as lock range started to be used by others. We do not think these are the right parameters. To characterise a bit stream you have to specify things which have an effect at the modulation level.

Cash: One of the benefits of the closed loop system is temperature stability.

Sharpe: Are we saying it is or is not possible to specify what exists now?

Wilbur: Your compatibility requirement is plug as opposed to functional. Other customers may or may not have this requirement.

Wilson: There are a number of different standards at the wire level. We should select which is best as the results of a technical analysis of what exists now and everyone should fall in line with this, leading to an equivalent of the Ethernet Blue Book.

Sharpe: I get the impression we cannot do this easily.

Wilbur: There are really three different implementations. The ULA's are roughly compatible with Mark II.

There is a fair degree of compatibility as Tony Cash has established. I think it is possible to produce a specification, the problem is to make it sufficiently watertight so that someone else implementing a system to the specification can ensure that it will be compatible with existing systems.

Sharpe: We are wasting our time if we cannot do this.

Cash: Why not aim at a smaller target, defining a range of operation that everyone can fit. Maybe this is a naive view, but it should be possible with the adaptation of a small number of parameters in a limited number of cases.

The questionnaires should identify the problem areas.

Wilson: A hybrid of the best aspects of these systems will be the worst - why not select the best of the three in its entirety.

Hopper: I do not think this is a problem. One can mix these things as they are roughly similar.

Heard: We need to relate this discussion of technical matters to real world service impact.

Cash: Existing systems overlap, anyone bringing out a new system must at least cover the overlap, not make it narrower. It has been narrowed already, we do not want it to be more so.

Erasmuson: Equipment from each manufacturer must be mixable in one ring. We must present a united front.

We are looking at the problems of how to put rings in chemical plants, hospitals etc; we want to look ahead and develop new systems, but what? We do not want to create a similar problem in two years time.

Cole: You need to look at environments and requirements and construct hardware to meet these.

Sharpe: It is clear that everyone wants to solve the simple problem of the 'now ring'. Is there a general desire to cooperate on the 'future ring' also.

All: Yes.

Wilbur: Bill's [Sharpe] immediate goal seems to be to build a dam to stop more proliferation.

3.3 Questionnaire

Tony Cash then continued his presentation.

The questionnaire assumes a number of pre-requisites:

1. A 2 channel communications system with modulation as described in Andy Hopper's thesis.

2. Phase lock oscillator, detection and remodulation of data. Repeater taps into ring.

3. Ring supplies DC power.

4. Monitor station to establish and maintain packet structure.

These were thought to be understood and accepted as axioms by all ring builders, but there now seem to be doubts about this.

The question was posed, 'Should RAL collect information on existing products and produce a draft specification/discussion paper for comment?'

Erasmuson, Hopper: Yes.

Wilson: In parallel with this start work on how compatible the three existing schemes are.

Cash: I have done some work on this already involving Logica and Cambridge Mark II components.

Wilbur: Are you intending to produce a report on these compatibility studies?

Cash: That will be included in the draft standard.

Drury: What is the impact on the user of the differences between these systems?

Cash: Firstly new DC voltage on the ring. Logica rings are powered from a 28 volt regulated power supply whereas the Cambridge Mark II design uses a 50 volt unregulated supply. A change to 28 volts is necessary, Mark II repeaters will operate at 28 volts; there is a suggested modification to operate Logica repeaters at 50 volts, but over a restricted voltage range, which is less than is normally found in a Mark II ring nominally operating at 50 volts. This change is necessary to meet IEC safety specifications.

A second restriction concerns the start-up frequency. The component selection tolerances on the Mark II design are very wide and the ring only works because of the terrific capture range and huge selection tolerance associated with the S124 chips. Occasionally (1 in 50) one has to change some components to get a Mark II repeater to function.

The restriction that applies when Mark II equipment is to inter-work with Logica equipment is that the monitor must free run at a frequency close to IOMHz (+-100 KHz). It is easy to add a trimmer to the circuit and adjust this such that this

condition is satisfied.

Compatibility is also influenced by ring size - in a large ring the frequency ends up close to 10MHz so there is no problem. In very small research rings these can be a problem as the operating frequency might be far removed from 10MHz.

Expected temperature and voltage ranges are also factors which influence how compatible equipment will be. Wide operating temperature ranges need to be accommodated if repeaters are to be mounted on external walls or in cable trunking.

The questionnaire refers to the number of repeaters that may be run from a single power supply and repeater spacing range.

The questionnaire also asks for experience with different types of cable, what is recommended for interbuilding or underground runs etc.

Question to Andy Hopper: 'In the ULA write-up it is not clear what hysteresis you have, could you clarify please?`

Hopper: I need to check, but I think it is a change of one-third of the most sensitive region before anything happens to the receiver.

Cash: Some questions may not apply to particular systems, if so please indicate why.

The questionnaire asks for the data handling (low error rate) frequency range (1.g). This is a difficult measurement. It was done here using a very large ring with large built-in time delays. Using a monitor with a variable frequency oscillator, steps of 50KHz are possible which enables a reasonable estimate to be made.

Hopper: As a side issue, what is the largest number of repeaters put together in a single loop?

Cash: I only went up to 16 and 300m cable.

Hopper: Has no one tried say, 100?

Wilbur: We have tried 40 nodes with 400m cable.

Cash: I have run a single 1Km leg successfully using BICC TR37 cable, standard drivers and receivers.

Sharpe: Can chips be tested at 10MHz?

Hopper: Yes, most of what we have been discussing can be tested, but the phase locking is a different question!

Ferranti are making 1000 chip sets, expected price in the region of 50 per set and expected to be available very soon. Digitally the chips will operate at 12MHz. The latest versions can be made to lock at 10 MHz, but some ranges are restricted.

It is also possible that one could ask Ferranti for n of thin best chips, and they would probably sell specially selected sets - at a price!

Drury: What is the guaranteed maximum limit of working of the digital circuitry?

Hopper: This is not clear as these parts of the test programme are not yet complete.

Wilbur: The question of guarantees is worrying. Jitter on the clock for example could make it necessary instantaneously that the logic run at say 15 MHz and that could introduce errors.

Cash: What is scatter like in the product?

Hopper: Resistor variations can be quite large with the Ferranti process. The BBC machine runs at 16 MHz, but one has to specify very carefully what one wants.

4. THE ULA STATION INTERFACE - Dr A Hopper

4.1 Presentation

Ferranti are currently preparing a document to be circulated with the chips.

Some of the apparent options are not real, for example where to put the two control bits. Four extra clocks are provided over and above the data and address bits on the assumption that a 4 bit shift register will be used for the control bits. Two are dummy. Control bits could be at the beginning middle or end of the data bytes, but it is suggested they come after the data and before the response bits.

There is also the problem of what to do with the control lines CS19 and 20 which load and read these bits. Options include use of separate lines or existing control signals, the bits themselves being put on the bus. There are no strong views on a preferred practice.

Multibyte rings appear very nice on paper, but it is not clear in the longer term that the specification should be such that the customer can choose the size. A fairly serious incompatibility problem (which has not been addressed) could arise if the user had free choice. A sequential scheme (load next etc) for loading data such as is found in the Logica station interface, is appropriate.

Standardisation activity with regard to the chips is a multi-level activity. There is a base level which is compatible with the ring as now, a second level blending in some of the features of the Logica interface and higher levels incorporating multibyte rings etc.

Transmit on accepted. In the Mark II ring one cannot use an empty slot immediately following the one just marked empty because of the 4 bit settling time of the asynchronous command logic. Transmit on accepted allows one to preload the transmit shift register while the previous minipacket is still making its way around the ring.The new minipacket is launched if the previous one was marked accepted. There is a good chance that the slot following the one marked empty can be used. This is an important option for small rings, and DMA access logics.

The Broadcast address option may be ignored.

Slot retention is not sufficiently well explored for the repercussions on monitor station design to be understood. This option may be ignored.

Most of the differences are fine. At the lowest level there are just a few choices to be made which require no great debate. At the next level choices include multi data bytes and then transmit on accepted command.

4.2 Discussion

Sharpe: The Mark II design uses a 100 wire cable. This interface has more functionality, how do you cope with this?

Hopper: This is a problem, should you use a ribbon cable or backplane. The systems in whose development I am involved, use a backplane but the ring station interface is also available on a connector.

Heard: We need a specification for this. I would have thought the only feasible option was a plug and socket.

Wilbur: There is an argument which says that the use of chips will so change the way nodes are designed that there will not be a need for an interface at this level. I would not defend that view strongly myself.

Heard: The station interface must be accessible and visible as now.

I would like to see it resolved whether there is merit in maintaining separately the 100 wire and Logica interfaces. I get

the impression that the ULA specification is not complete - one has to decide to eliminate options.

5. POLYNET STATION INTERFACE - S Wilbur

5.1 Presentation

There were several reasons why Logica decided to go away from the Cambridge Mark II 100 wire interface. The Logica interface was designed in mid 1980 before the ULA chip specification was available, though there was a general idea of what it would contain. The flexibility to access the ULA chips when available was desired, also some 18 months experience of designing access logics to the old specification had repeatedly produced a number of standard problems whose solutions were desired.

1. Factorise out common interface functions into station. (Note a glossary of terms is long overdue!)
2. Remove variability by providing useful functions in a standardised way.
3. ULA compatibility.
4. Ease of interfacing to microprocessors.
5. High data rates for microprocessors - electrical ways of controlling data flow.

There are a number of improvements incorporated in the Logica interface:

1. node address readable
2. state of internal jumpers readable
3. retry on busy - important when using simple DMA chips so do not have to poll for a complex condition TX & done
4. cancel receive transmit requests - believed to be difficult on ULA's
5. 1-8 data bytes
6. type bits usable in 40 bit slots.

A port (register) structure has been defined. Four control lines are used for accessing functions and an 8 bit wide bus for data values. There was considerable debate over the choice between:

dual 16 bit wide bus
single 16 bit wide bus
dual 8 bit wide bus

single 8 bit wide bus

The minimum solution was chosen because it was believed many applications would be concerned with microprocessors. The only real requirement for a dual bus is with bridges and there are other ways to solve that problem.

Ports 0-8 follow the standard layout, 9 and 10 are reserved and 11-15 are available to the user/interface unit.

There are a number of facilities which differ from Cambridge Mark II.

The node address is visible to the user.

Difficulties have been found with some implementations because of different procedures for enabling and powering up the station which could cause ring errors. A simple technique was required. The Logica interface has only one enable bit which does not allow the user to manipulate the power relay. The programmer can cause at most one ring error, if he fails to programme power fail safe correctly.

There are 2 configuration options available; 38 or 40 bit mini- packets. IN the 40 bit option, the 2 type bits follow the trailing end of the data.

Broadcast addressing has been implemented.

The number of bytes per minipacket can also be selected.

The interface also contains auto retry logic. Simple DMA chips may use this facility to retry a transmission without host involvement. There is an enable auto-retry bit in the control logic which can be read and set. TX done is only set when the transmission has properly completed, not on each retry round the ring. Retry is continuous until either the packet is accepted or auto-retry mode is disabled.

A major concern was that the 8 bit bus should allow an 8 bit DMA controller to address the interface as a normal peripheral device and independent of the number of bytes per slot to first order. Essentially there is a single port with multi-function capability.

Writing to port 6 for example loads successive bytes into the minipacket transmit buffer and transmits the packet when it is full. Port 5 is similar except the packet is not transmitted when full. This is left - for symmetry with the receive case where successive bytes from the received minipacket are returned when the buffer is read. Port 7 is a control signal - reading or writing this port issues a receive or transmit minipacket command.

It was then necessary to decide whether to provide access to the counters stepping through the minipacket byte structure. It was decided this was unnecessary as the programmer can track this. Transmit and receive reset initialises the counter to zero. Enabling the node will also reset the counters to zero.

Logica wanted to be compatible with the ULA chip rings and also to co-

exist with existing products at the wire level hence an option to support type bits was included. The type bits can be written provided TX done is set (normal timing rules apply) and can be read if RX done is set. In 38 bit mode the type bits are read as zero and writing has no effect.

It turned out to be very simple to implement broadcast addressing though there are limitations if using basic block protocol. Broadcasting is likely to be useful for multicast addressing of voice packets etc. It was included as it involved negligible costs in logic. Broadcast enable and a status register to indicate when the packet has been received are provided. This is compatible with the ULA ring.

A ring-on derivative of CS17 has been included - AOK when ring switched on, node powered, no breaks upstream and address plug inserted. Effectively this is a software version of the front panel light!

The electrical interface has been reduced from 100 to 50 wires. The most important aspects are the encoded addressing structure and 8 bit bi-directional bus. These provide no real hardship for 95% of applications. Control lines include:

TDONE	action when minipacket successfully transmitted
TCLOCK	pulses every time minipacket transmitted round ring. Counting this can be used to limit retries
TERROR	as corresponding bit
RDONE	receive done
BPR	broadcast packet received
Type Bits	are brought out
TEST	for commissioning - external manual enable of node
GND	Limitations to powering node through interface
+5v	(voltage drop in cable etc).
	1 spare!

An attempt has been made to keep interface timing as sensible as possible so that if it is necessary to multiplex interface bus requests it can be done with reasonable spare timing.

5.2 Discussion

Sharpe: DoI have placed a contract with Logica to develop an interface standard, how is that work progressing?

Cole: This activity is concerned with specifying this interface in such a way that it represents a watertight description of the electrical characteristics and functionality of the interface and also incorporates as formal a definition as possible of the station lying behind the interface. When complete the definition will be sent to interested parties to solicit their views on its acceptability as a standard.

Sharpe: What about the extra functionality?

Cole: I think it does need to be included.

Morley: I think your proposal needs thought.

Erasmuson: We are supplying the unwritten community standard product - Logica are proposing something new.

Cole: We are offering this as something to get views on as something that may be useful.

Our work was motivated by intelligent interfaces and controllers, likely to contain embedded microprocessors.

Heard: I would like to see us discuss this further. I think we are not dealing with a black and white situation. What are the merits of the different approaches:

1. Encoded method so simpler to attach microprocessors above.
2. Lay bare the functionality of the underlying system at minimum cost with highest performance.

There may be merit in both; are they so fundamentally different?

Sharpe: I have got the impression from informal discussions that it is possible to map the 100 way interface to the 50 way interface apart from the extra functionality and repeat on busy.

Wilbur: Repeat on busy can be done with some messy logic on top of the ULA chips. I believe the only incompatibility is the ability to cancel outstanding requests.

The 100 wire interface exists but supports 2 bytes for a mini-packet only. For ULA's you have to move forward anyway. There is an existing 38 bit interface and one Logica claim is applicable to ULA's similar to that but embracing the ULA functionality.

Hopper: I am not clear about the 8 bit bus and its implications, when responses are available etc. It is necessarily worse than the 100 wire interface as it is one quarter of the width.

The Orbis Acorn view is to go for a simple general interface that will fit on a small card. For lightly loaded systems there are no objections to a simple interface which loads up the system. For sophisticated interfaces one would build something like a Gizmo which has enough cycles to look at the ring at the lowest level.

Repeat on busy is reflected in software somewhere along the line.

I am keen on a low level interface on price grounds. The width of the bus and its implications on larger machines is a difficult question to answer.

Sharpe: To describe what exists now would require less than the functionality of the Polynet interface. The Logica interface is a bid to satisfy user requirements.

Cole: We should carefully consider the notion of having 2 interfaces available. To select one of the other easily, both should be available in silicon which may answer some of Dr Hopper's points. Additional custom logic requirements may be less for one than the other in a given case. We should consider both in the light of putting them into silicon.

Hopper: Fine. At the higher level something needs to be said about DMA, the number of channels, where they lie, how they are read/loaded etc.

Sharpe What stage has the DoI contract reached?

Cole: A plan of work has been formulated, ground work is in progress but I haven't a clear idea how it related to this activity as I am not sure how this is progressing.

Wilbur: There is no overlap at the line level. It is only in the second area where there is commonality. Part of the brief is to document the Logica interface which implies documenting the functionality underneath.

Drury: I have not seen an equivalent document for this alternative 100 way interface.

Hopper: We are faced with a status quo - it is too late to decide on one or the other.

Morley: There are two today, we should move to one in the future.

Wilbur: The technological incentive to change comes with the ULA chips and silicon. Although I have an axe to grind, I worry about the use of the 100 way interface in future as it does not embrace the functionality of the ULA's.

Sharpe: Can the Logica interface be viewed from 2 angles:

(1) to which the existing interface can be mapped
(2) desirable given the advent of the ULA chips, but subject to discussion amongst those here.

It should not be contentious to say that there should be 100 and 50 wire interfaces.

Wilbur: There should be agreement on functionality, but maybe there should be 100 and 50 wire electrical interfaces.

Erasmuson: The question is given that you have 16 data bits and 16 bit machines, why break this down to 8 bits?

Battell: Data may be on 1 not 2 byts boundary so there is a need to go down to the byte level.

Erasmuson: Agreed. Should we keep it that way?

Wilbur: I do not believe the arguments are cut and dried. I notice no one is arguing strongly for a dual bus.

Battell: I see one of the problems as to whether to go to full duplex operation. The 100 way interface had simpler logic operating.

Wilbur: This seems to be a major criticism of the single bus. If one goes to a dual bus, register layout needs careful consideration. Bus structure and to some extent width is an issue.

Battell: Interleaving two independent operations is a problem, because of time constraints on access to the host's store etc.

Drury: This is getting into the area of what is on the other side of the interface.

Battell: Switching times for the ring are not enough - you have to do something with the data!

Wilbur: I understand this and agree timings do not have immense margins, but a real system design would incorporate FIFO's anyway. I did not try to cover the situation when one is aiming for very high performance.

Hopper: This is an area for debate, there are no obvious answers. I would suggest this needs looking at and that formal/informal comments be made. It is certainly correct to tidy up what exists.

6. ORGANISATION OF FUTURE WORK

Sharpe: We have loosely identified four separate pieces of work, we should nominate four individuals to progress each.

1. Specification of area of overlap between existing systems which we want to keep no narrower than it is now. Tony Cash will coordinate this activity.

   It was agreed that replies to the questionnaire were to be returned by Friday 6 November 1981. Replies should be channelled through Bill Sharpe.

   Members of the group should make it perfectly clear what information can be circulated to other members and what cannot. The default is that information will NOT be circulated.

2. User requirements at repeater level; (long hops, safety regulations etc).

4. The exercise Logica are already pursuing - an adequate 50 wire interface and station functionality.

The third area is more difficult to define precisely but is something like:

3. Specification of 100 and 50 wire interfaces as they are today.

Hopper: There is room for tidying up - it is not clear if this should be implemented in hardware or software, nor how high it should go.

Wilbur: Timing rules and a functional description are needed for the 100 wire interface.

Cash: Maybe that needs a questionnaire also.

Hopper: I would be happy to take this on; documentation of the original 100 wire interface. A separate issue is how the ULA chips could map into this.

Sharpe: Please cover functionality as well as electrical specification.

Heard: Comments should be invited from SEEL and Toltec.

Hopper: I should be able to produce the document in 10 days.

Sharpe: Please send it to me and I will circulate it.

Comments should be received within a week after that.

If members undertaking work on behalf of the group will indicate formally how much effort the exercise will consume, we will approach DoI with a view to obtaining at least part funding.

What about the repeater wire level specification from the long term point of view?

Cole: It would be useful to have user input for this - the definition should be driven by use requirements. A poor standard in this area would be a disaster.

Wilbur: The design and drafting of an appropriate standard is not a task to be undertaken lightly.

Wilson: This is all very well, but what interface specification do we use now? We have to make that decision long before this exercise can be completed.

Cash: Now that the Polynet specification has been circulated, suppliers can consider what it would cost to incorporate it into an existing product.

Erasmuson: The two existing interfaces will continue - you have to consider support for kit already in the field.

Morley: A new specification could be rolled into a chip product.

Sharpe: A subset of the Logica interface can be mapped into the 100 wire inter- face. This at the moment is the only level which will insulate you from all the existing implementations.

Wilson: Are you drawing a line at the access logic to node interface and not considering higher interfaces?

Sharpe: Yes. The specification of the higher transport service interface will come out of the JNT ring exercise.

Wilbur: I am worried that that should be decided by a closed user community.

Sharpe: What I mean is that the transport service interface and high performance access logics are sufficiently difficult that they are research areas. The JNT will continue their exercise in the widest possible context and input of ideas.

Wilson It would help me if you would outline what other initiatives and committees exist in this area and how they are related.

Sharpe: The Focus steering committee under the DoI umbrella are considering standardisation of LANs, common products, interfaces etc.

A JNT programme for establishing the transport service and GIZMO is in preparation.

A JNT Protocols Working Party for the definition of basic block, TSBSP etc as a firm document is also being launched by Ken Heard.

The standing of the JNT and its relationship to SERC and the Computer Board was explained.

NEXT MEETING

It was agreed that the next meeting will be on Wednesday 9 December, commencing at 10.00 am.