Minutes of T2UK Meeting at RAL

on 7th January 1994

(Produced by Robin Middleton)

7th January 1994

1 Present

J. Strong (RHBNC)	S. Fisher (RAL)	F. Harris (Oxon)
S. Hunt (Oxon)	N. Gee (RAL)	F. Wickens (RAL) - chairman
S. Kolya (Manc)	I. Newman (RAL)	P. Clarke (UCL)
R. Cranfield (UCL)	B. Green (RHBNC)	R. Middleton (RAL) - secretary
J. Leake (RAL)		

2 Status Reports

2.1 Report back from Utrecht Workshop on C40s

John Strong reported on a meeting he attended in Utrecht in early December. A full set of slides are available for perusal from John.

- Zeuthen group have DBV44 equipped with two TIM modules and interfaced to a SPARC-10 Sbus. This is very similar to the system in place at NIKHEF.

  They have a design for a HIPPI destination with a C40 doing buffer management and control for use at HERA-B. It uses a dual port memory organised as 8k x 32 bits and is supposed to go at 30 Mbytes/s. Similarities with work done at RHBNC were noted !

  They have also developed an optical C40 link extender using Xilinx technology. This permits C40 links to be run at near full speed over long distances (e.g. between racks). This sounds potentially interesting.
- Josi Legrand presented the latest updates on his C40 data-driven architecture. Exactly what he's doing has yet to be unscrambled (see slides for details).
- John Renner-Hansen is looking at C40 to T9000 links (still waiting for the latter). He commented that there is now (from Analogue Devices) another candidate DSP sporting 40 Mbytes/s and 6 or more links per chip in addition to many other bells and whistles.

- Jos Vermeulen is planning a C40 system for use with MSGCs and using the 3L software with some custom extensions. He will also continue to simulate and ultimately test his ideas on a C40 funnel for local event building.
- Utrecht group seem to be more oriented towards computer science in that they are aiming for a large processing array employing switchable C40 links.

There was the suggestion the all DBV44 owners could pool resources for testing out any one (or more) idea more fully (e.g. Josi's system). So far no more on this has been heard.

2.2 C40 work

2.2.1 at RHBNC

A T2 buffer prototype now exists and plugs into a TIM site on a DBV44 board. As of the day before the meeting the system was proven to work electrically. Test software will be tried over the next couple of weeks. Input to the system is from the SLATE board. Barry showed a picture of the prototype.

2.2.2 at UCL

UCL now have the software development system for C40 and have borrowed a C40 from RHBNC. A DBV44 with 3 TIM sites populated has been ordered, though the delivery date is as yet unknown. Gordon Crone has written test software and will follow this with buffer manager code soon.

Peter Clarke mentioned that a company called MIZAR also now market a C40 board, although NOT based on the TIM standard. However, they offer a board with 8 C40s on it which might be interesting at a later date.

2.3 Choice of Global Processors

The original plan was to use an Alpha processor on a Futurebus+ board from DEC. This project appears to have been moth-balled within DEC. Instead two options have been reviewed, both based on new Alpha evaluation boards from DEC. The first, EB64+, is based on the existing 21064 chip with 128 bit path to memory. It is different from its predecessor, the EB64, largely in its support of PCI through a custom chipset. The second option, EB66, is based on the new Alpha chip 21066 (= approx. 80% of performance of 21064) which has PCI support on chip. However, the EB66 has only 64 bit path to main memory. Both boards are made available at near cost to DEC and include circuit schematics.

The EB66 will come initially with a cross-software environment based on GNU-C. It will later be followed by DECelx (UNIX real-time kernel based on VXWORKS). Indeed DECelx is expected to arrive on EB66 before EB64+. The availability of both EB64+ and EB66 is not before March / April although it may be possible to get access to the DECelx environment and circuit schematics very soon.

The preferred option is the EB66, which RAL have plans to purchase. Manchester also will purchase an EB66 and possibly an EB64+ as well.

2.4 SCI status

The first SCI node chips were expected in December '93. LSI, the manufacturer, supplied them to Dolphin in November last year. Initial tests showed problems with the by-pass fifo.

A pair of chips can be made to talk, but a ring can not be made to work. The fault appears to be with fabrication at LSI and a 3 month slippage is now expected, with first chips now due in March ’94.

On the bright side, SCI cables have been delivered from Dolphin; pity we don’t yet have anything to plug them into !

Scott reported that that the SCI to memory bus design was complete except for one or two cross checks to do against the SCI node chip. PC manufacture will be in a few weeks time and hence, assuming all goes well, will be on schedule for delivery of SCI chips in March.

2.5 Operating Systems for Global Processors

Fred Wickens briefly described work done on the Chorus operating system by Iain Newman and himself. A reasonably extensive evaluation of operating system features has been conducted so far and will result in a report. At present, Chorus does not run on Alpha processors (the GLVL2 target architecture) although it is known that there is commercial interest in such a port.

The Lynx/OS operating system has a significant following at CERN, particularly by the RD13 group who have run the TC/IX port (by CDC) of it on R3000 processors (RAID from CES) and are now porting their software to Lynx/OS proper on 68040 processors. Once again it does not yet run on Alpha processors, though there appears to be interest from outside HEP for it to do so.

Robin Middleton mentioned that the current favourite software environment for the Alpha processors is DECelx. This is based on the VxWorks real-time kernel, which also runs on C40s.

2.6 Physics Simulation Studies

No one was present from this sub-group (which is based at CERN). The principle workers are Lluisa Mir (RHBNC), Peter Sherwood (UCL) and Steve O’Neale (Bham). The work is directed towards understanding the implications for level-2 triggering of low luminosity running. Discussions on B-physics issues have taken place with Nick Ellis. More details of activities are expected at the Atlas meeting in February.

2.7 Discrete Event Simulation

Steve Hunt described work he had done in Oxford on a graphical interface written in MODSIM for use with simulation code from UCL. It allows to setup a run-time configuration of a C40 network prior to execution of a simulation.

Jason Leake stated that MODSIM had been purchased at RAL and was installed on the CSF farm of HP machines in the Central Computer Department. Jason has the SCI simulation code from Andre Bogearts (CERN) which needs pruning before use as the basis of a re-write. The immediate goal is to work on simulation of the global part of the system; Richard Hughes-Jones has also stated an interest in this.

Peter Clarke suggested that those active with MODSIM should get together more frequently.

Frank Harris introduced a workshop he is organising at the Coseners House for the weekend of 19th/20th February. There is likely to be participation from non-UK colleagues active in this area. There was a discussion on how best to advertise it, since only people ACTIVE in this area should participate. A provisional agenda is available from Frank.

Robin Middleton very briefly described some results from a "simple-minded" simulation written in C++. These were the result of a high statistics run over the Christmas period and show that fluctuations on some plots seen previously are purely statistical effects. Concerns were expressed on the use of C++, which had been adopted originally because of its superior performance over MODSIM. The latter is now supposed to have been speeded up significantly and it was felt appropriate to repeat the original benchmarking.

3 C40 / SCI Interface Definition

Over recent months both T2UK and RD24 have been considering the design and fabrication of a C40/SCI interface. Each group started thinking in this area independently and as of the end of last year informally agreed a degree of co-operation. To this end, T2UK (Fred & Robin) will specify our needs to RD24 prior to a meeting at CERN at the time of the next Atlas week.

There was much discussion of the needs of T2UK compared to the ideas within RD24 on this subject. The discussion was curtailed until a more complete session in the afternoon, the outcome of which is covered in note ATLAS/T2UK/94-1 (available from Robin or Fred).

4 Test Beam Work in 1994

T2UK has been encouraged to go for test beam studies during 1994 by both Nick Ellis and Daniel Froidevaux. So far only the systolic option based on a router and the ENABLE / MAXVIDEO systems have been tried in anger in beam. For us to have even a chance of a place in the technical proposal (due end '94) we should undergo a similar exercise to prove our asynchronous system.

The following points were noted :-

- RD6 are ready to buy HIPPI destination modules and RIO processors...this has significant implications for our T2 buffers.
- The last test beam work this year is in September (heavy ion work follows).

A full discussion of test beam possibilities took place in the afternoon. The end results are summarised in the Appendix A.

5 Admin & Future Meetings

Fred Wickens briefly mentioned the budget structure for Atlas work. 500 cost centres have been allocated to Atlas for its lifetime, of which 50 are designated for T2 capital expenditure. Currently just one is in use ! Travel will be organised within Atlas by institute and not sub-project.

Fred also showed the resource estimates for fiscal year 94/5 and a forward look up to the turn on of LHC (whenever that may be). They were based on consultations by email during

December. The sum requested from now to completion of Atlas is 3 million pounds with the most significant spend in the final few years before completion. A reasonable level of spend is expected in intervening years to cover test beam work and to improve technologies used. This information has now been submitted.

Only 10k pounds (approx.) have been spent from the PPESP funds so far (order confirmed after the meeting as having left RAL on 22/12/93). This has gone to purchase a C40 VME module (DBV44) equipped with 3 C40s. It was agreed that since the funds are extremely tight, ALL future spend on this budget should be agreed by colleagues in advance (as has happened so far); i.e. with LLVL2 colleagues for spend up to the local part of 17.5k, with GLVL2 colleagues for the global part of 17.5k and with T2UK as a whole for the remaining 5k.

Dates for your diary are :-

- Next T2UK meeting to be held at either RHBNC or UCL on Tuesday 22nd February (to start in the morning).
- Atlas week meetings of particular note :-
    - T2 meeting on Tuesday 1st February at 2pm in Bat 4-S-013.
    - FE/Trigger/DAQ meeting on Thursday 3rd February at 9am in LEP Auditorium.
    - Plenary meeting on Friday 4th February will include a special session on FE/Trigger/DAQ.
    - meeting with RD24 folk on C40/SCI interface -> to be fixed.
- East Collaboration meeting on Tuesday 8th February.
- MODSIM / simulation workshop on 19th/20th February at the Coseners House.

6 A.O.B.

None.

7 Conclusions

- T2UK will put test beam proof of system components as the primary goal for 1994.
- ALL T2UK spend to be agreed in advance by LLVL2 or GLVL2 colleagues, depending on nature of purchase.

8 Actions

- SK will think through consequences of putting a "FIFO-like" structure onto the SCI/PCI interface in addition to the standard memory bus interface.
- FJW/RPM will convey thoughts on C40/SCI interface to RD24.
- RPM will write a list of goals (first pass) to achieve beam tests and attempt to put them on a PERT chart.

- Steve Hunt, John Hunt and Jason Leake to get together on MODSIM issues more frequently.
- JLL to repeat benchmark comparisons of C++ and MODSIM done some 2 years ago.
- FH will finalise organisation of a MODSIM / Simulation workshop at Coseners House.

A Beam Tests in 1994

John Strong described current thinking on a mark II T2 buffer, now the concept is proven (well almost). It would be based around a three board set a) data input, b) T2 buffer itself and c) data output. In the initial proposal the input and output modules would be HIPPI based to fit in best with what seems to be in vogue today. This does not, however, preclude other data sources and sinks being designed as the interface between modules is kept very simple (just 32 bits + some control signals) and buffered through a fifo. In addition the T2 buffer module would support VME and have inputs from the LVL1 ROI builder and T2 supervisor and an output to local LVL2 processors. The VME interface can be substituted for the HIPPI output module if desired.

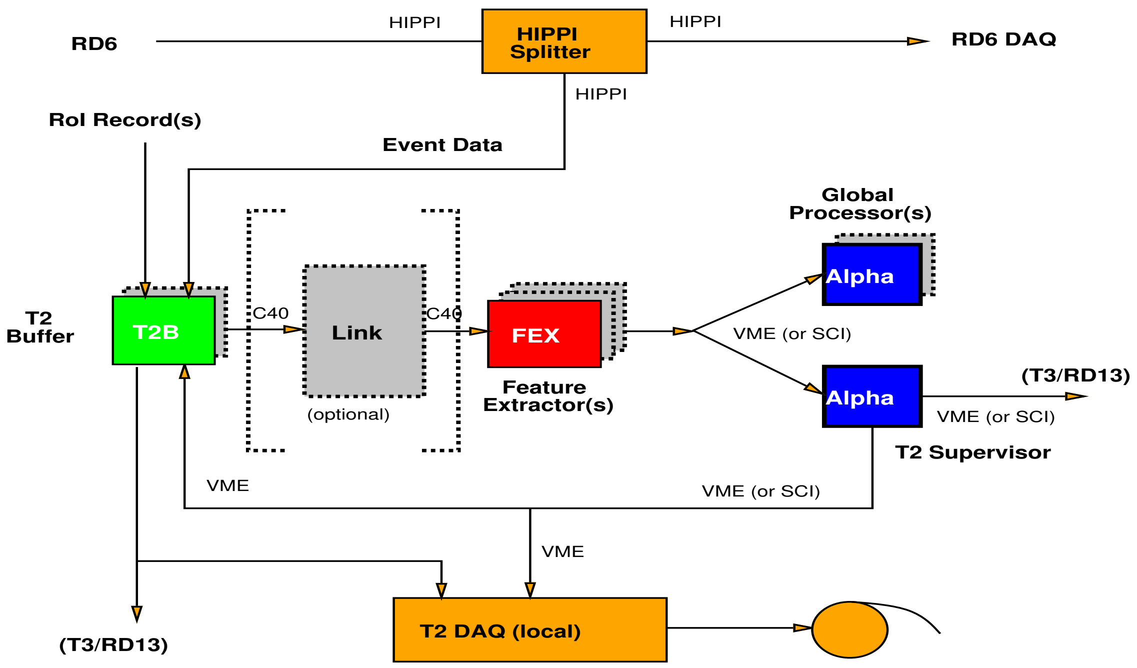

Figure 8-1 Proposed Test Beam Setup (1994)

There was some discussion whether the RD6 setup of HIPPI module and RIO processor could be used as a T2 buffer. However, it was felt that this would form a major diversion of effort away from pursuit of the optimum design currently being followed. Concerns were also expressed that T2 equipment should have minimal impact on RD6 operation and certainly should not get onto their critical path. To this end it was suggested that custom HIPPI spy hardware should be used to gather data in a "passive" mode from their system. Details of this have yet to be worked out.

The subject of seeking assistance from CERN was raised. Help with HIPPI work was mentioned as a possibility as was assistance with software. However, no concrete suggestion emerged.

Figure 8-1 shows a first guess at what the test beam setup could look like. It's strength appears to be that each major component can be demonstrated / tested to some extent independently of other parts, although we should naturally aim for a complete a system as possible. The preferred interconnect for the global part is SCI. However, given the slippage in SCI node chip delivery, it was felt that SCI may not be ready in time and that data could best be transported as a stop-gap between local and global parts over VME. The ensuing discussions centered around this option, although it was recognised that should a viable system incorporating SCI arrive on time, then it would be used.

B Progress Report on LLVL2 (from Manchester Meeting)

T2 Buffer

A pcb of the core unit has been delivered and is being assembled. We hope to have the unit operational by end of the year so that we can find out what changes are needed for the second version. This first board is intended to show that the central idea of using a C40 as a buffer manager works but the board does not have either a general purpose input channel or an adequate VME output. A discussion will take place on Friday 26th Nov. at RHBNC to try and set outline specs for the next iteration which should address both these areas. The input channel will be HIPPI and Tzvetomir Anguelov from CERN will stay at RHBNC for a week to work with Barry to put a HIPPI destination into the design. If the current design works well we could have a second board around Easter but more realistically we should not expect it until September 1994.

T2 Buffer .vs. RIO

A RIO exists with a HIPPI destination, software memory management and VME. The performance of this unit and its projected successor does not reach the 100MB/s of HIPPI or the T2B design (which is HIPPI limited). We would need to build an interface card for the RIO to obtain C40 link compatible communications (for T1/T2 RoIs, data to T2 local and T2 yes/no). With the unit managing all functions by software what would be the sustainable input data rate? (for the current unit Rudy Bock tells me that they get about 10MB/s) We should prepare a short paper setting out our case for building a T2 Buffer and compare its spec against the current and projected RIO. We need agreement with ATLAS that this is a worthwhile project to provide LHC like capability in the near future. If ATLAS says that the RIO is good enough we will have to consider our position.

SCI Local

If SCI chips ever reach a low price for a low performance version they could become an interesting technology for a local architecture of the type we propose. SCI nodes could be used to connect together the buffers for a local area and bridges used to provide the nearest neighbour links. At present we learn about both C40s and SCI as network elements and we should not therefore change our current plans, merely bear the possibility in mind. We need to find out if any one is working on this (in RD-24?) but without provoking anyone to start the activity (i.e. ask everyone what they are doing, not ask if they are doing this.)

Embedded C40 board

In future we want to be able to embed C40s both for link units and for T2 buffers in order to reduce the cost by not employing TIMs. If/when manpower is available we should design a board to see how easy/hard it is to embed.

C40 Link Units

An order for a DBV44 with 4 TIMs for just over 10k with VAT is with RAL. We must check that it does not get lost/delayed in the system. UCL have now received there XDS510 system and RHBNC will lend them a TIM plus base board to get started on the software testing while they wait for the LSI board.

Utrecht Meeting

A meeting of C40 interested parties will be held in Utrecht on 14th Dec. John Strong will go, Peter Clarke and Robin Middleton have been invited. Do we want to try to set up a link test with all available DBV44 boards from all parties (Amsterdam, RHBNC, UCL, Utrecht and Zeuthen)?

Other Points

- PPESP has accepted revised milestones and ignored our moans.
- All institutes should send Peter Norton their ATLAS expenditure estimates for next year. As our equipment budget is the balance of the PPESP funds left by next April we need to know what we will spend this year and inform Peter of the balance. To avoid double counting we should ensure only one equipment estimate goes in from T2UK. On travel we should make sure that our travel is included in the institute bid so that a) we are not left out b) we are not double counted. It should be from the institute as Peter will provide a travel budget for each institute and not each sub-project if the procedure is the same as last year. If it is felt necessary, we can inform each other of our respective needs. For the forward look, we should adopt a similar approach but we should take into account test beam running from spring 1995. We need to discuss our strategy for the forward look equipment claim based on the total we believe T2 needs from the UK (~3M) and a spend profile to our best estimate of LHC turn on date. The equipment forward look bid should be a single T2 return.
- We should establish contacts with people looking at the data acquisition for the UK prototype work to see if we can produce a compatible system. One contact is Maurice Goodrich at Cambridge who is looking at one type of read out for the silicon. We need to know who is working on other systems.
- Moscow State University has a group of engineers and technicians interested in contributing to T2 work. Have we any projects to give them? Contact John Strong for more details.
- Although the Technical Proposal is likely to be delayed, we should assume that the original schedule holds (May) and prepare a proposal based on our ideas (we should leave room for a router approach to some detectors to avoid appearing to grab all the action). The proposal should include costings and John Strong will circulate current T2 estimates to interested parties on request.
- A meeting must be arranged in the near future to discuss our requirements for a C40/SCI interface. This can be compared with the RD24/Valencia proposal to see if they are compatible and for us to decide our programme on this. (My suggestions:- RAL 3rd or 10th December)

C Distribution

Attendees +
Peter Norton
Tony Weidberg