Minutes from the ATLAS-UK Level-1 Calorimeter Trigger Meeting

Monday 31st January 2000 at 10 am

RAL Conference Room 4 R27 (Atlas Centre)

Present: Eric Eisenhandler, Reg Gibson, Murrough Landon Ed Moyse (QMW), George Anagnostou, Paul Bright-Thomas, Scott Talbot, Alan Watson, Peter Watkins (Birmingham), Bruce Barnett, Ian Brawn, James Edwards, Norman Gee, Tony Gillman, Bob Hatley, Viraj Perera, Tara Shah (RAL).

Minutes generated by B.M.B.

General

How should we distribute the minutes and transparencies? All

A discussion took place concerning a proposal to move exclusively to electronic collation and distribution of the contributions to the level-1 UK meetings via the www. Norman expressed concerns about the systematic archiving of the information in such a scenario, in particular noting bad experience with the poor maintenance of old links.

In terms of work overhead, it was suggested that the web approach was probably more efficient in the collection stage, but that the subsequent generation of paper copies by multiple individuals at different institutes was apt to be less efficient.

It was decided to move to the proposed methodology, with individuals asked to provide copies of their transparencies in one of the following formats:

- postscript
- pdf
- or (exceptionally) scanned versions of hand written transparencies.

Individuals will be responsible to provide these by e-mail, either before the meetings (preferable) or shortly thereafter.

Data transmission

LVDS test results and lessons learned Richard/Paul

Paul presented transparencies in Richard's absence, summarizing the current status of the DSS LVDS tests underway at Birmingham.

The tests involved simultaneous transmission over 8 channels connected with two of the 4-pair cables from Datwyler incorporating matching pre-compensation LR circuits (tolerance +/- 2 m). The error rates were lower than 10-12 in all channels, and better than 10-13 in most, with cable lengths of both 10 and 15 meters. Furthermore, operation with a single 20 meter cable also runs at a BER of 10-12.

Very little change in the BERs was noticed, even with the source and sink in separate crates (albeit common grounding), or with intense VME activity in the sink crate.

It was noted again that these results satisfy the criteria set down at the last joint level-1 trigger meeting (Stockholm) upon which the choice of LVDS was contingent.

It was mentioned that the residual error rate appears related to poor jitter tolerance in the LVDS receiver, with rates not particularly dependent upon cable lengths.

Tests will proceed using the TTC, and with a new LVDS sink card from RAL. Tests to determine the effects of long PCB tracks, to determine the advantages of a (soon to be available?) faster LVDS chip-set, and eventually a redesigned LVDS source card are anticipated.

Discussion on DSS test plans (LVDS, G-link, S-link) All

Anticipating the LVDS tests utilizing TTC, Norman commented that Sam Silverstein has observed large timing jitter from the TTC clock, of the order of 400 ps. ((peak-to-peak). The TTC spec states an rms jitter of 70 ps. rms.) The speculation is that the problem is with the test clock, so that one should drive the system with the best clock possible. [Sam has posted some comments regarding their measurements on ATLAST1@LISTSERV.RL.AC.UK, Re: "Clarification of TTCrc measurement".]

The question arose as to what additional tests using the diagnostic power of the DSS are anticipated. Viraj commented that to date the only S-link tests that have been made were simple transmit/receive types of tests, so certainly yes, in this regime.

Hardware status summaries

Preprocessor MCM PDR summary Tony

Tony presented a summary of the PPr MCM PDR which took place on Monday 24 January. Comments had been received in advance by e-mail, and the major points were summarized and addressed through a list of recommendations presented to Ullrich at the end of the review.

Nine points were presented. Several issues have, in particular, possible serious longer term consequences:

- PHOS4 Timing: Momentary interruption of the clock to the chip corrupts stored data in a way which does not offer an obvious path for recovery. Ullrich is investigating the failure modes and feedback to CERN will be generated.
- Clock jitter on LVDS timing: TTC specifications show a maximum jitter larger than LVDS requirements. This requires a repeat of the LVDS link tests using the TTC system. In addition, adequate PS decoupling and pre-compensation requirements need to be included formally in the specification (as they are crucial to reliable LVDS operation).
- Fall back scenarios if LVDS has serious unforeseen flaws:
  If we must, for any reason, fall back to G-link technology, there will be major implications for scheduling.
- Scheduling: MCM scheduling should be meshed, in detail, with PPr ASIC and PPM scheduling.

The question of the sensitive issue of FADC chip choice arose again (and had been an item in Tony's summary). The technical issues behind the choice need to be documented, but this is best to be done as a forward reference.

CP "chip" PDR summary — Tony

Tony presented a summary of the PDR of the CP ASIC/FPGA which took place on 25 January. Comments had been received in advance by e-mail, and these then formed the first basis for discussion. During their face-to-face discussion a number of additional points surfaced, proving the value of this type of PDR over that which would have comprised only e-mail and reaction. Future reviews will follow the same model.

Most of the pre-review proposals for change have already been incorporated in a new draft spec. Tony summarized the reviews in 8 points.

Modifications included:

- the inclusion of a 42 bit error register to allow the rapid identification of the channel generating parity errors.
- a specification of digital interpretation of saturation
- the inclusion of logic to ensure interpretation of a threshold of 255 as infinite (isolation off).
- inclusion of 'FIFO flush'
- ability to disable parity error generation from UN-driven (and hence forced to 0) serial inputs.
- inclusion of a summary of technical issues affecting choice of FPGA over ASIC technology.
- specification of consistent scheduling interplay relating CP ASIC/FPGA, Serializing FPGA and the CPM prototype.

A rather important issue was that of input data monitoring, and the need to explore techniques alternate to the scan path technique to provide monitoring of the incoming data.

ROD prototype — James

James commented briefly on the progress of the ROD prototype. There are no real problems to report. The ROD schematics should be complete, on schedule, by the end of March. From past experience, layout, manufacture and assembly will take about 3 months, so hopefully the hardware will be in hand at the end of the second quarter, which means end of June.

It was asked whether the same manufacturing problems experienced with the DSS due to limited available real-estate are anticipated in the manufacture of this prototype. It was stated no, in particular not with the use of large gate-count FPGAs.

Norman asked whether an intermediate inspection step might be helpful, but Tony commented he thought not. Additionally, no final design review will be needed on this project.

Serialiser — Ian

Ian presented a status report on the Serialiser.

The design of the serialiser itself matches its specification, and the timing simulation has been completed successfully. The pin-out has yet to be defined.

The test module, presented at the previous meeting, is a daughter board for the DSS. It incorporates a deserialiser which should look like the front end of the CP chip. Elements of that design have been taken over from the serialiser where possible: it is based on the RAL215 ASIC design.

The initial design of the test module is complete, as is its functional simulation. The next step is to synthesize the design, target it to a VirtexE FPGA and simulate the timing.

Paul raised the point that the Heidelberg people should check that the format Ian expects from the Preprocessor is what they plan to provide.

Timing Control Module — Bob

Bob described the specifications for the proposed Timing and Control module (TCM). This module is intended to distribute TTC signals (receiving optical and distributing electrical) and provide control and monitoring to the DCS. It must be in a form compatible with CP and JEP crates, and allow use with the preprocessor crate. Hence, it provides signals both to the backplane and the front panel through 18 buffered PECL outputs.

Monitoring is provided thorough a local crate bus to the relevant modules. A CANbus node, based on a programmable micro-controller and servicing either an intermediate or a global global CANbus, is anticipated.

Block diagrams of the system as proposed were shown.

A number of issues in the design remain outstanding. These include the choice of:

- physical format
- VME pin allocation
- fanout technology for the TTC signals
- VME version and supported signal subset

Murrough asked whether the suggestion was indeed that there should be only one control point to the DCS. Yes. Eric asked what would happen if individual crates were equipped with their own CAN interface, as is a suggestion for the ATLAS standard crates. This would need to be foreseen in the architecture.

Bob then described the TTCrx based decoder card (TTCdec) intended for use on all prototype trigger processor modules. It is a counterpart to the TCM, receiving the electrical signals (via bus or front panel) and decoding the bit-stream to generate local TTC signals. It consists of a small daughter card, designed with the flexibility to allow for various ASIC versions and various (3V or 5V) voltage levels as required by these. The TTCrx outputs are not buffered before use.

The draft specs are to appear on the web and will be agreed within the U.K, before being submitted to the rest of level-1 for comment.

These cards are intended, again, for the prototyping - not production - stages.

CP Module prototype/Backplane and cable connectors Richard

Paul presented the work of Richard on this subject as well. The PDR will take place in March 2000, with the specifications to be circulated internally at Birmingham before submitted to a broader audience. Recent progress has been in:

- the choice of the BERG Metral connector for common use in the common CP/JEP backplane.
- Readout Controller specification.
- LVDS link tests.
- The specification of the control bus. In this case a specification for a compact VMEA24/D16 bus (eliminating some redundant signals) has been made. The specification for this may be found on Richard's home page (URL: http://www.ep.ph.bham.ac.uk/user/staley/cpm_spec.ps)

The choice of which subset of VME64 to use was pursued. Norman suggested that a VME64 subset would be preferred and that effort should be made to understand the full ramifications of removing other signals. In particular modules should not misunderstand scans over VME which take place outside the address space implied by the available signals. The need for a decision to be made soon was underlined.

Common Merger Module Norman

Norman summarized the CMM brainstorming which took place January 21, just before the PDR sessions described by Tony.

The connector issue has been resolved, by default as the major contender to the Berg is not available at a reasonable price.

All backplanes, even at the prototype stage will be production backplanes. This means that the module height (9U) and pin-out will be fixed at that point. Regarding the depth of the cards, 40 cm is assumed. Some uncertainty had been expressed regarding the ease of finding a manufacturer capable of providing boards of that size. Sam and Uli will investigate the possibilities.

Crates will be required from the middle of the year onward. The discussion regarding choice of VME spec (also discussed above) is converging. Some issues involve what interrupt needs are anticipated, and how to insure against misinterpretation of VME cycles generated outside those supported by the bus.

The full final design is anticipated even at prototype stage, vis-a-vis the controller design (thereby reducing engineering effort) but the number of channels replicated on the board will be a reduced set.

Stockholm and RAL have agreed that RAL will assume responsibility for the design and its implementation.

Although the specs of input signals across the backplane are agreed, and issues such as transmission of information via a compact floating format resolved, some functional questions remain outstanding, in particular pertaining to energy summing and FCAL where physics simulation/trigger issues need to be resolved.

Norman will write a functional specification which will then be sent for engineering specification, when agreed.

More brainstorming is anticipated at a meeting to take place on 3 March, the location of which is to be determined.

Schedule and reviews update Tony

Tony summarized the status of reviews which have taken place in 1999 and are scheduled for 2000. In the former, 5 reviews:

- DSS, PPr ASIC, Serialising FPGA,
  ROD prototype, TCM prototype

took place.

In 2000, in addition to the 2 PDRs of January (PPr MCM and CP ASIC/FPGA, 2 more are imminent:

- March 2000 for the CPM prototype
- April 2000(expected) for the JEM prototype (incl. FPGAs)

Additionally, 2 PDRs (PPM and CMM) will occur at as yet unspecified dates later in this year, as will the final design review (FDR) for the PPr ASIC in March or April 2000, date to be confirmed.

The intention of holding frequent ‘brainstorming’ sessions in order to set the frame for and complete these reviews was again emphasized.

Trigger simulation

Forward-jet trigger thoughts Alan

Alan presented some thoughts on the forward jet trigger, a topic which has been re-emerging of late. Some of these thoughts were circulated previously on ATLAST1@listserv.rl.ac.uk (Jan 07, 2000), and related items published on the web:

- http://www.ep.ph.bham.ac.uk/publications/atw/lvl1/jetbits_addendum.html
- http://www.ep.ph.bham.ac.uk/user/pbt/lvl1/threshold_splitting.txt

There are a number of physics motivations for including such a trigger:

- improvement of signal/noise in Higgs/SUSY analyses (forward jets).
- triggers on QCD physics where the cross sections populate the regions at large eta.
- unforeseen physics interests.

but in addition, proper treatment of saturated regions (especially those which occur within the FCAL region) may demand some triggering abilities in any case.

A key question is that of FCAL signal granularity, and in particular whether this would need to be altered to achieve the trigger. If not, an implementation should be relatively simple. If granularity in eta (other than forward/backward) is required, then the anticipated 16 JEMs can be used if the bin-size in the PPr is changed.

A 17^{th} JEM may be avoided if the available threshold bits for some jet triggers are redefined so as to only count multiplicities using two bits, leaving the remaining bit for the FCAL trigger. A 17^{th} JEM may still be desirable, if azimuthal granularity above that anticipated (16 fold) is required, or the number of available jet thresholds is insufficient. Also the 17^{th} JEM would also improve the bandwidth available between modules. The implications for a total jet-Et trigger, such as has been discussed by Juergen but not yet included in our design, also needs to be explored.

Simulation studies are proposed (initially with ATLFAST, later with ATRIG) to investigate the alternatives, with code currently being written. A discussion of the technical questions (and possible solutions) is ongoing.

It was emphasized that the hardware schedule (particularly in the case of the CMM, which is not a prototype) requires fast convergence in this discussion.

Present and future trigger physics simulation software — Alan

Alan presented this tutorial on trigger physics simulation. Two tools exist:

ATLFAST is a fast simulation which doesn't track but rather models detector resolution, longitudinal (but not transverse) shower development, etc. It allows access to a simplified model of the (final) ATLAS reconstruction. There is a deep level-1 simulation available for ATLFAST, but this is not a standard part of the release. ATLFAST has an unofficial C++ rework, ATLFAST++, which however provides for no such simulation. Work on formal OO-ification of ATLFAST will likely commence after April when the Architecture Task Force's report, submitted last year, is approved.

ATRIG is the detailed (legacy) simulation which has been been the workhorse of full simulations to date. It is relatively well planned, but shows its heritage as 'grown' code. It is currently quite well documented and stable.

In operation, output from a physics Monte Carlo generator produces 4 vectors (KINE), which are then passed to DICE which produces DIGI output for input to ATRIG. The first two stages are normally run centrally. ATRIG then allows the study of level-1 and level-2 trigger algorithms based on this input (which can then filter and produce input for higher level studies in turn). The slides from Alan's talk contain additional details along with useful URLs which help the novice get started (or the experienced user debug).

In the context of level-1, moving to an OO framework for the trigger simulation code poses its own challenges. One will need to understand the trigger functionality and the existing code. A rough study in this direction exists at:

http://www.hep.ph.rhbnc.ac.uk/atlas/trigger_sw_req.txt

There were a number of comments. In particular, Norman commented that with availability of advanced prototype (and final) modules in the near future one needs to consider in what framework should one write online analysis code so that it moves effectively between [online/off-line/simulation] environments. One needs to understand how different are the environments, Also the degree to which one can use generated test vectors to test or certify the hardware.

Tara contrasted the nature of earlier approaches to software description of the level-1 hardware functionality. "The Model" was more a formal description of components and their interrelationships within hardware, the aggregate of which could generate any output state based on the input state. The direction of Mike Pentney's work was to define a simpler paradigm concerned more with input and output, each stage making a transformation, the inner details of which

need to be simulated but not necessarily reproduced in full detail. This latter direction is apt to be the more appropriate in the future, but the detail required in the simulation will need to be understood.

Concerning ATLFAST, it was commented that Juergen Thomas had done some work already on studies of the Sum Et trigger ... and that this needs to be kept on the agenda (... Sam/Uli).

Concerning organizational questions, and how to move on.

- Obtain a proper briefing from Simon George
- Talk about the issues amongst the level-1 UK people. (eg: those present)
- Norman emphasized that so far we have had a rather rigid division separating off-line and online environments, but here we have to move together.
- We have to determine where the crossover/overlap between off-line and online is. Probably the commonality does not extend all the way to the front end ... but algorithms should be framed to reflect the commonality which does clearly exist.

Online software

Work on DAQ software.and DAQ-1 Norman

Norman reviewed the status of the DAQ software, and the process of merging and evolution. In moving from the DAQ which has been in use until now, to one based on the "DAQ-1" environment, it is necessary to maintain a running system for ongoing tests, and in particular to have a working system available for the ROD prototype. As not all necessary elements are available yet within "DAQ-1", a mixed system is anticipated in the interim.

Work on the old DAQ has included removal of old code supporting defunct hardware and re-compilation under g++, thereby allowing direct linkage with C++ code.

Movement or merging of database, buffer management, and data-producer functionality into the "DAQ-1" environment is the next step.

HMINI calls will be replaced with an interface to a network-based ROOT system, VME interfaces will be replaced with a variant of the VME-daemon supporting the RAL diagnostics, and error messaging previously provided by EMU will be moved to target MRS.

Comments on DAQ-1 installation at RAL Bruce

Bruce commented that the "DAQ-1" framework/ demonstration DAQ as generated by Murrough at QMW had been installed at Rutherford. Some preliminary work at a crude interface of that framework with the control (daqctl) code has been made.

Comments on ROOT implementation at RAL Tara

Tara showed a test plot generated by ROOT with code destined for the HMINI replacement. In this direction (and for eventual use in the DAQ-1 environment) he is adapting some software which originated with CDF and provides a client-server access to ROOT histogram objects. These objects can be saved locally, and then transmitted over a socket-based interface to a local machine for display.

Tara is planning to attend a ROOT workshop at CERN during the first week of February and intends to report at an upcoming software meeting.

Work with Heidelberg diagnostic software Scott

Scott described his experience with the installation and USE of HDMC at Birmingham. In spite of some compilation problems experienced with QWT, he has been able to test compilations generated at RAL and run at Birmingham.

Having added DSS and TTCvi interfaces, he finds HDMC crashes consistently when Memory components are added. There are a few bugs, with register content view and update. Scott commented that although new elements may be added with the software, direct text editing of the file describing the elements is in his experience a quicker alternative

Work on TTC diagnostics Bruce/Kithsiri

Bruce reported that Kithsiri completed a version of his diagnostic software for the TTC system before returning to Sri Lanka, which Bruce installed for him (from tar file) at RAL (hepvme2). Kithsiri will make some improvements to the code and will continue to interact with the group. He is available via e-mail at "kithsiri@sjp.ac.lk".

The code currently runs only under Linux/PC via the network server to the LynxOS system, although that problem (the inability to compile natively) is under investigation.

Other topics

Trigger/DAQ "Commission" status Eric

Eric reported that progress had been made on the T/DAQ reorganization whose preliminary description had been circulated in the December Birmingham meeting. There had been a lot of discussion at Beatenberg and subsequently, and a new draft is in preparation for discussion at the institutes board meeting during ATLAS week (Feb, 2000). [Eric and Fred have circulated copies of the new draft.] Level-1 is not at the hub of the discussion, but the key question is still the choice of project leader. The choice of the coordinator of data flow is also a difficult choice, with no obvious candidate. One possibility remains to bring in fresh blood to fill the position.

ATLAS notes, LEB and IEEE papers for 2000 All

It was pointed out that there are a number of notes whose submission (via the ATLAS web-pages) has been pending for some time now.

- Alan is nearing the completion of a rewrite on his, as is that of Paul, the latter incorporating comments from Stockholm. The authors are hoping to get these submitted as soon as possible.
- It was requested that those of Richard (backplane tests) which has had a draft available on the web for some time and Kithsiri (which was sent to Nick for comments (?)) should be submitted directly.

The topic of conferences was brought up. The IEEE meeting (Lyon) and the LEB (Krakov) have deadlines for summaries and abstracts, respectively, which are approaching - sometime around April.

Possible topics include:

- DSS experience and testing.
- LVDS test results.
- The role of/ experience with compensation in LVDS
- Experience in TTC integration into our systems.
- G-link experience.

It was requested that a more coordinated approach than that of 1999 (which left something to be desired) be taken!

AOB

Open Source Model All

Under AOB, the DAQ-1 Open source proposal was discussed. Bob Jones is actively investigating licensing implications and so on. Murrough was encouraged to send a message to Bob Jones indicating the support and enthusiasm of level-1 in such a direction. DAQ-1 representatives are proposing to hold a workshop to help people become involved and integrated in this aspect of DAQ-1.

Date of next meeting

Note: the date set at the meeting has been changed to:

- Monday 20th March at Birmingham, Watson room 310