Report of the Joint Final Design/Production Readiness Review

PRE-PROCESSOR MODULE

Abstract

The combined Final Design Review and Production Readiness Review of the Level-1 Calorimeter Trigger PRE-PROCESSOR MODULE (PPM) was held on 12 October 2006 at the Kirchhoff-Institut für Physik (KIP), Universität Heidelberg.

The Heidelberg group presented the work that they had carried out to prepare the Pre-Processor Module design for production. Following an internal Interim Design Review in December 2005, several important design changes were implemented and have been evaluated in an extensive test programme.

The reviewers were satisfied that the design and implementation of the PPM are now sound, and that the module achieves the required performance in all areas that have been tested so far. They requested that a number of further checks should be carried out to conclude the test programme.

Once this stage has been satisfactorily attained, full module production should commence as rapidly as possible. The reviewers made a number of suggestions which might accelerate the production and test schedule.

When the production and test programme has been completed, a major revision of the PPM documentation should be carried out, to provide a definitive specifications document for the final production module.

Prepared by :	*Checked by :*	*Approved by :*
Tony Gillman, RAL	**Dave Charlton, University of Birmingham** **Eric Eisenhandler, Queen Mary, Univ. of London** **Philippe Farthouat, CERN** **Norman Gee, RAL** **Stefan Haas, CERN** **Stephen Hillier, University of Birmingham** **Weiming Qian, RAL**	

for information, you can contact:	Tony Gillman	Tel. +44 1235 445521	Fax. +44.1235 446733	E-Mail a.r.gillman@rl.ac.uk

Distribution: EB Members, TDAQ Members, all participants mentioned in the report.

PURPOSE OF REVIEW

The purpose of this review was to check that:

- **The design and implementation of the Pre-Processor Module had been demonstrated to be technically sound**
- **Production issues had been addressed satisfactorily and a viable schedule had been prepared**
- **A full Test Plan had been prepared**

MEMBERS OF THE REVIEW COMMITTEE

Review Committee	Pre-Processor Module team
Dave Charlton, University of Birmingham	Paul Hanke, KIP
Eric Eisenhandler, Queen Mary, University of London	Kambiz Mahboubi, KIP
Philippe Farthouat, CERN	Karlheinz Meier, KIP
Stefan Haas, CERN	Klaus Schmitt, KIP
Stephen Hillier, University of Birmingham	Hans-Christian Schultz-Coulon, KIP
Norman Gee, Rutherford Appleton Laboratory	Rainer Stamen, KIP
Tony Gillman, Rutherford Appleton Laboratory *(chair)*	
Weiming Qian, Rutherford Appleton Laboratory	

Ex officio	For Information

AGENDA

Review agenda, with links to all the talks:

http://indico.cern.ch/conferenceDisplay.py?confId=7002

REVIEW OUTCOME

1 INTRODUCTION

The combined Final Design Review and Production Readiness Review for the Pre-Processor Module (PPM) of the ATLAS Level-1 Calorimeter Trigger was held at the Kirchhoff-Institut für Physik (KIP), Universität Heidelberg on October 12th 2006. The reviewers were:

Dave Charlton, University of Birmingham *(DC)*

Eric Eisenhandler, Queen Mary, London *(EE)*

Philippe Farthouat. CERN *(PF)*

Norman Gee, RAL *(NG)*

Tony Gillman, RAL (*Chair*)

Stefan Haas, CERN *(SHa)*

Stephen Hillier, University of Birmingham *(SHi)*

Weiming Qian, RAL *(WQ)*

The Pre-Processor Module design team was represented by:

Paul Hanke (Kirchhoff-Institut für Physik)

Kambiz Mahboubi (Kirchhoff-Institut für Physik)

Karlheinz Meier (Kirchhoff-Institut für Physik)

Klaus Schmitt (Kirchhoff-Institut für Physik)

Hans-Christian Schultz-Coulon (Kirchhoff-Institut für Physik)

Rainer Stamen (Kirchhoff-Institut für Physik)

This review was the third and final assessment of the PPM, following the Preliminary Design Review (PDR), held on 8th March 2001, and the Interim Design Review (IDR), held on 13th December 2005. The IDR was an internal review and no formal report was issued.

The review opened at 09:00, with two talks being presented in the morning session. Tony Gillman gave a short introduction to the review, outlining its structure and purpose:

http://indico.cern.ch/getFile.py/access?contribId=0&resId=1&materialId=slides&confId=7002

The main technical material for discussion was then presented by Kambiz Mahboubi in a composite talk, covering many different aspects of the current version (v2.0) of the module:

http://indico.cern.ch/getFile.py/access?contribId=1&resId=0&materialId=slides&confId=7002

He discussed the design and implementation changes from the v1.0 PPM, which had resulted from the Interim Design Review held ten months earlier. His talk also described the plans for manufacture and commissioning of the production modules.

The main points covered in this talk were:

The AnIn boards had been re-designed with single-channel operational amplifiers and improved layout to eliminate excessive inter-channel crosstalk, now measured to be ~0.25%.

Significant improvements had been made to the LVDS serial data signal paths, with careful re-design of the pre-emphasis circuitry both after the MCMs (LR network) and the LCD cards (CR network). The LCD cards had been subject to a very careful re-layout. In a variety of different configurations, four PPMs sourced error-free data over ten 11m LVDS cables in runs up to 20 minutes in length.

The ReMFPGA resource utilisation averaged between 30-50%, but was not yet fully optimised.

With four PPMs connected to a single ROD, an extensive series of readout tests had been carried out under a variety of extreme operating conditions, satisfying the maximum rate requirements of the level-1 trigger.

On-board clock distribution had been re-designed, with clock-trees improved by careful manual routing and the addition of suitable damping resistors. Considerable work had been carried out to ensure that the PHOS4 device would operate reliably, and to find a robust work-around of the known initialisation design bug. The reference clock, which must always be present at the PHOS4, is supplied from an on-board crystal oscillator, for which the most important parameter was found to be its duty cycle, which should be very close to 50%. The result of these improvements is reliable power-up initialisation of all 16 PHOS4 devices on a module.

The VME and Flash-Loader CPLDs are now two separate devices, and both have been extensively tested on the 20 pre-series PPMs.

Improvements had been made to the Hot-Swap Controller, by reducing its sensitivity to voltage/current deviations and by making allowances for power-up sequencing. These changes had been tested successfully.

A description was given of the rear Service Module, which will plug on to the backplane P0 connector and provide connector ports for JTAG access to the CPLDs and an RS232 interface to the CAN micro-controller.

A limited number of further tests are planned to take place at CERN immediately following the review (CAN controller, TTCdec long broadcast commands, *etc*).

Some details were presented of the proposed manufacturing plans and estimated schedule. All PCBs (sufficient for 160 PPMs) should be available by the end of 2006, with assembly taking place during Q1 of 2007. All SMD assembly would be in a single batch, with electro-mechanical (manual-intensive) assembly being performed in small (25%) batches. Comprehensive test plans existed for standalone and system tests of final production modules.

Several points were raised during this talk and during the subsequent discussion session *(see Section 3 – Discussion)*. The final working session before lunch was devoted to a presentation by Paul Hanke of a document giving more details of the proposed production and test plans for the PPM:

http://www.kip.uni-heidelberg.de/atlas/L1/DISCUSS/FDR_PRR_Planning.pdf

To conclude the morning session of the review, the reviewers were then invited to the Trigger Laboratory to see a demonstration of the full-crate PPr system in operation.

The second half of the review began with a discussion session addressing the specific written comments and questions submitted in advance by the reviewers. The major areas on which they had focussed could be categorised as follows:

- Production and test plans
- LVDS tests with v2.0 PPMs
- VME functionality
- PHOS4 operation
- Power management *(crate PSUs, cooling, etc)*

The reviewers then held a closed session to discuss their conclusions and recommendations, which were finally presented to the KIP group before the review closed at 16:30.

2 DOCUMENTATION

- The Pre-Processor System in the ATLAS Level-1 Calorimeter trigger, Draft Version 0.1, 11/2005
  http://www.kip.uni-heidelberg.de/atlas/L1/PP_System/PPSys_Wrup.pdf
- The Pre-Processor Module (PPM) for the ATLAS Level-1 Calorimeter Trigger, Draft Version 0.1, 02/2006
  http://www.kip.uni-heidelberg.de/atlas/L1/PP_Module/PPMod_Wrup.pdf
- The FINAL Pre-Processor Module (PPM) for the ATLAS Level-1 Calorimeter Trigger, 09/2006
  http://www.kip.uni-heidelberg.de/atlas/L1/DISCUSS/FDRev_PPM_upd1.pdf
- The FINAL Pre-Processor Module (PPM) for the ATLAS Level-1 Calorimeter Trigger *(some more details on planning for production/testing)*, 09/2006
  http://www.kip.uni-heidelberg.de/atlas/L1/DISCUSS/FDR_PRR_Planning.pdf
- PPM Schematics, Revision 2.1, 09/2006
  http://www.kip.uni-heidelberg.de/atlas/L1/PP_Module/PPM21_schematics.pdf
- Channel Mapping on the Pre-Processor Module (PPM), 09/2006
  http://www.kip.uni-heidelberg.de/atlas/L1/PP_Module/ppm_channel_mapping.pdf

3 DISCUSSION

Following the formal presentations in the morning session of the review, the afternoon session was devoted to a detailed evaluation of the written points that had been submitted in advance by the reviewers. Many of the questions had been satisfactorily addressed during the morning session, but others required further discussion, as summarised below. *(It should be noted that only issues which could affect module production were considered relevant for this review; the focus was therefore very much on hardware.)*

The suggestion to improve further the rigidity of the PPM motherboard by extending the horizontal strengthening bars forward to the module front-panel was rejected as being unnecessary.

Concerns about the PHOS4 stability during data-taking were discussed, and the reviewers were satisfied that there is adequate real-time monitoring of the status of these devices, via the MCMs.

The LVDS data transmission tests on the latest v2.0 LCD cards used a variety of data patterns, long strings of zeroes/ones as well as pseudo-random binary sequence (PRBS) patterns.

There are still some checks of TTCdec functionality to be carried out; in particular, Broadcast operation and the integrity of all address bits.

In answer to a question about the status of firmware for the second CPLD, it was stated that is now complete and fully debugged.

The reason for not using System ACE for management of firmware configuration data files is that there is a lack of sufficient real-estate on the PPM motherboard.

The value of providing interrupt and block transfer capability in the VME CPLD was queried, but it was felt that there could be some future uses for these features (such as rate-metering?).

Operation of the enlarged (4 Mbyte) SRAM in PPM_2.0 has already been tested on a module.

In the absence of clear information about VME registers to provide CAN controller access, there had been concern that VME operations might interfere with the CAN controller operation, but the reviewers were satisfied that this had already been taken into account in the architecture.

After some discussion about the oscillation problems observed with the PPr crate power supplies, the proposed solution of reducing loop inductance for the supply leads was considered adequate. Tests had been carried out to verify that the oscillation frequency increased as the loop inductance was reduced, and that a sudden increase or decrease in supply current did not stimulate oscillation once stability had been attained.

It would be possible to add a Status Register flag to indicate the status of the RGTM on each PPM; the appropriate G-link bit is already tracked through the connector and the motherboard to a front-panel LED.

The possibility of adding a simple form of air deflector in the central region of each PPr crate to try to eliminate the MCM hot-spots in slots 11 and 12 was discussed, but it was decided that it would probably be too difficult to make effective. The temperature differential was only ~10 degrees, and the MCMs had been operated at temperatures of 80 degrees without problems. The use of chilled (water-cooled) air in USA15 should also lower the MCM temperatures further.

It is not intended to use “home-brew” SBCs in the final PPr system; production module testing will be carried out with the ATLAS “standard” *Concurrent* SBCs. It was also confirmed that hot-swapping modules in a PPr crate does not cause the *Concurrent* SBC to re-boot.

It was confirmed that the modified pinning of the V8 and W8 LVDS fan-out signals has been carried out, as requested at the time of the Interim Design Review in December 2005.

The latency figure quoted is as measured for the outputs to the JEP sub-system, rounded up to the nearest integer, and it was confirmed that the equivalent figure for the outputs to the CP sub-system is smaller by 1 BC.

The reviewers suggested that it would be helpful to have a little more detail provided in the specifications about the analogue signals feeding the PPM; *e.g.* amplitude range, AC or DC coupling, source impedance, *etc*. The differential input impedance of the PPM is 86 ohms.

Although I2C access is disrupted in the absence of an external TTC signal, the system will recover automatically once the TTC signal is restored.

The CR component values in the LVDS pre-compensation networks on the LCD cards were designed for 15 metre cables. However, the tests that have also been carried out using the 11 metre production cables have shown no bit errors, although this will be re-checked.

There were a number of points relating to the VME operation of the PPM for which the reviewers suggested further clarification in the specifications document; these are listed in Appendix I.

Sending fill-frames is quite sufficient to establish link-lock between the PPM G-link transmitter and the ROD G-link receiver; no TTC use is required, as incorrectly stated in the specifications.

It is very important that the entire set of schematics and firmware source files are properly archived to ensure efficient maintenance and capability of future upgrades. It was agreed that this would be achieved by means of a CVS repository, which should also include all necessary text files.

Some time was spent in discussing the problems associated with initialising the PHOS4 chip reliably. As had been explained in the morning session presentation, extensive work had been carried out by the Heidelberg group to determine the critical operating conditions for this device. The most significant parameter was found to be the duty cycle of the 40 MHz reference clock, which should deviate only a few per cent from an ideal 50%. Routing a highly symmetric clock, generated from an on-board crystal clock feeding an FPGA DLL, directly to the PHOS4 guarantees the device a suitable reference clock even immediately after module power-up. A series of tests had been carried out on all of the 20 pre-series PPMs with satisfactory results.

One consequence of these changes to the PHOS4 configuration is that the readout data transfer to RODs is limited to a fixed frequency of 40 MHz. If a higher bandwidth were needed in future, a 25% increase could be obtained by operating the G-link transmitters on the RGTMs in 20-bit, rather than 16-bit, mode. After some further discussion, it was therefore agreed that the 40 MHz upper frequency limit would be acceptable. As successful tests of the RGTM operating with a pre-series PPM had already been carried out, approval could therefore now be given for production of the remaining 142 RGTMs.

There was some discussion on the production planning. It was noted that the proposed number of production modules was now 160, making a total (including the pre-series modules) of 180 PPMs, which is rather higher than originally expected. The assembly process for the production modules would be identical to that used for the pre-series modules. No Pb-free techniques will be employed. The current schedule for production of the CAN Controllers appeared rather late (~March 2007), so it was suggested that a second assembly company might be sought.

For production module tests, some form of G-link data sink device would be needed. It was suggested that the UK system of DSS with G-link receiver daughter-cards might be useful in this role. An LVDS data sink device was already in design in Heidelberg. It was noted that a 48V PSU would need to be added to the PPr crate for ROD operation. Also, some ancillary modules (TCM, TTCdecs, *etc*.) would need to be supplied from RAL.

Finally, it was noted that the PPM design team have already developed a good database for maintaining up-to date records of the production/test status of all MCMs, which will be used also to record the status of the production PPMs and associated daughter-cards. (This database may also be very useful for maintaining records of production modules and components from other trigger sub-systems.)

Since the reviewers had not so far had the opportunity to study the production plans in detail, it was agreed that they should make their assessment as soon as possible after the review.

4 MAIN CONCLUSIONS AND RECOMMENDATIONS

The reviewers were satisfied that the Heidelberg group had thoroughly demonstrated that the design and implementation of the PPM is now technically sound, and that the tests that had been carried out to date on the current pre-series version of the module (v2.0) had demonstrated satisfactory performance.

However, the review revealed a number of areas where testing had not yet been completed, and it was therefore recommended that the following tests be carried out as soon as possible:

- *All features of the PPM CANbus system should be verified to work correctly.*
- *A test should be carried out to ensure that the PPM responds correctly to TTC Broadcast commands.*
- *A definitive and systematic check that the channel mapping is correct should be made by configuring one PPM to feed its JEP LVDS outputs to one JEM (or to one CPM).*
- *A single PPM should be configured to operate with as many LVDS outputs actively sourcing data as possible. i.e. by feeding 24 LVDS cables into a single JEM, or 20 LVDS cables into a single CPM.*

- *Correct operation of PPMs in a multiple-crate configuration should be verified by configuring four PPMs (two in each of two crates) to each drive data on several LVDS cables into a single CPM.*
- *The integrity of data sourced from a PPM to a ROD should be checked by operating the PPM in the presence of maximal VME and CANbus traffic in the PPr crate.*

Subject to the above tests being satisfactorily completed, approval is given for the full-scale production of the Pre-Processor Modules to proceed as rapidly as possible,

The following recommendations related to organisation of the production processes:

- *Following the hand-assembly and testing of the first pre-series CAN Controller daughter-card, a commercial company should be contracted to assemble the remaining 19 pre-series modules.*
- *Manufacture of the 160 production CAN Controller daughter-cards should proceed in parallel with manufacture of other modules/daughter-cards, employing another company if necessary.*
- *To avoid potential delays at the PPM testing stage, the production schedule for the TTCdecs and the Auxiliary Backplanes should be arranged to match the production schedule for the PPM motherboard.*
- *All PPr crates should be shipped to CERN as soon as they are fully-equipped, and then installed in USA15 to allow the complete LVDS cabling plant to be installed.*
- *It will be very important to populate the installed PPr crates in USA15 as quickly as possible, to allow calorimeter signal calibration studies to proceed efficiently. The production schedules for the PPMs and associated daughter-cards should therefore be streamlined as much as possible to accelerate the shipping of modules from Heidelberg to CERN.*

The draft version of the Production Test Plan presented at the review will be assessed by the reviewers, and their conclusions will be distributed within one month. *(Appendix – Section H)*

The final category of reviewers' recommendations related to the Pre-Processor Module documentation. It was generally felt essential for there to be a major update of the main specifications document, which currently contains some material which is either obsolete or requires revision, making it difficult to determine what is current.

The reviewers therefore recommend that a thorough revision of the documentation is undertaken in order to arrive at a full specification in a single document of the current production version of the module:

- *All of the firmware and schematics files (including explanatory text files) for the PPM should be archived at CERN, and appropriate links to the files should be created in the PPM specification document.*
- *Final revisions of the main PPM specification document must be completed and approved before the PPM is signed off (although module production is not contingent upon this).*
- *The documentation – in particular the main specifications document – must be updated as soon as possible to represent the final status of the module design and implementation.*
- *The reviewers recommended that additional material be added to the main specifications document in various places in order to clarify some details, and they also noted a number of minor typographical errors that should be corrected. Details of these points can be found in Section I of the Appendix.*
- *Fully-revised final versions of all documents should then be submitted to EDMS.*

APPENDIX REVIEWERS' WRITTEN COMMENTS ON THE REVIEW DOCUMENTATION

With the exception of Appendix H, which resulted from the Review itself, the following comments were received before the Review, and have been edited for clarity.

They should *not* be regarded as final conclusions, which are listed clearly in Section 4.

A GENERAL COMMENTS

(EE) The PPM is very late, but the state of the module does look good now. A lot of very thorough work has at last been rewarded.

(NG) There's a lot of work gone into the documents - thanks to those involved.

(SHi) Thanks for all the effort and new information put into these documents. However, I'm afraid I found some of the documents a bit of a confusing mixture of status, history and plans. We do need at some stage a clear and full specification of the final PPM implementation with properly updated details (many of tables and details in the current specification are out of date). I would also have liked to see a test-plan for the modules. I know most of the functionality of the PPr system is already well tested in the MCM tests, but it would also be good to see how the connectivity of the module itself will be systematically tested after production.

(DC) Overall, the module looks in pretty good shape - very good news.

B THE PRE-PROCESSOR SYSTEM (INTRODUCTION)

http://www.kip.uni-heidelberg.de/atlas/L1/PP_System/PPSys_Wrup.pdf

(SHi) I've only really skim-read this, but I think it does nicely fill a hole in the description of the system, and so is a very good addition to the documentation, bringing together all sorts of diverse general information. However, it does have some unfortunate misleading and incorrect information in the details of the crate layout (section 4.2). For example on page 17, the grouping of EM and Hadronic PPMs in alternate slots is completely untrue! On page 19, the CP algorithms only extend to 2.5 not 3.2. Just below, the figure of 6 PPMs connected with the forward jets is a very strange number. There are 4 PPMs dealing with the FCAL (PPM 8a), and another 8 dealing with the high eta end-cap region (PPM 8). I'm not at all sure where this number 6 comes from.

C THE PRE-PROCESSOR MODULE (PPM) (SPECIFICATIONS)

http://www.kip.uni-heidelberg.de/atlas/L1/PP_Module/PPMod_Wrup.pdf

(EE) It's far too late, of course, to change the FPGA situation, and we understand the history, but the REM_FPGA is doing too many different things (with limited resources) and should ideally be split into two or three separate chips.

(EE) Section 3.4.7, italics just before "Local Power Control" - Now that you have a full crate, has this now been done?

(EE) "Local Control" section - We hear about the possibility of detailed error codes using the ATMEGA, but how much is actually implemented and is there a list? (i.e. is this alphanumeric display actually doing anything useful?)

(EE) Section 3.5 - Latency shown is for the JEP; please confirm that it's one b.c. less for CP.

(SHi) Page 5 - What is the expected threshold for external BCID, and is this only ever to be used as a back-up option?

(NG) Page 9, Section 2.4, para 2 - Is a separate 40MHz test crystal provided, or is it the same one that drives the G-Link? How does the PPM behave if there is no TTC or if the TTC stops?

(NG) Page 16, section 3.4.4 para in box - Does this mean that the I2C BUS stops working if the TTC clock disappears? How can it be recovered?

(SHi) Page 16 - Here, and in other places, there's an explicit mention of 15m cables, and that the LCD is specifically tuned to this value currently. However, we know the final length of the cables to be used, and it's not 15m. Are adjustments needed, or have they already been done?

(NG) Page 17, just before section 3.4.6 - Does this readout have any impact on the event data transfer rate?

(NG) Page 19, greyed-out para - Has this power check been done? Does it depend on the incoming pulses? During a calibration run, there may be many large calo signals at the same time. Is the power consumed OK?

(NG) Page 19 - Has it been verified that hot-swapping doesn't cause a concurrent CPU to reboot?

(WQ) Page 20, Section 3.5 - Are these latency figures measured? I wonder why all of these are integers?

(EE) Page 20, section 3.5 - Latency shown is for the JEP; please confirm that it's one b.c. less for CP.

(NG) Page 23, bullet 3 - Does VME access slow down DAQ?

(SHi) Page 24 - Readout merger: Is the situation with needing completely different firmware loads for different slice-number configurations likely to persist? If so, how many different configurations can be stored locally for quick transitions?

(SHi) Page 24 – VME: What is the time needed currently to load individual, different values into each LUT in all channels of all PPMs in a crate? What mechanisms are foreseen to speed this up if it is too slow?

(NG) Page 25, Section 4.3.2, para 3 - If a value is loaded from VME into the TTCrx, then changed via a TTC command (e.g. a reset), what value will be read back from the PPM through VME?

(NG) Page 37, section 4.4.1 - How can we be certain that the required DAC load has been completed?

(NG) Page 37, section 4.4.2 - What will happen if I attempt a read on a phos4?

(NG) Page 38, section 4.5 - What does this data look like if one MCM channel is disabled?

(NG) Page 38, section 4.5 - I thought the LUT data could also read 0-127 slices?

(NG) Page 40, section 4.5.2 - The RemFPGA needs to check that the G-link Tx is phase-locked to the PPM 40.00 MHz clock. Sending fill frames (DAV false) should be sufficient to lock the ROD. There is no need for TTC involvement.

(NG) Page 40, section 4.5.2 - Need to verify if Rx G-Link on ROD can lock to 16 and 20-bit data.

(WQ) Page 40, section 4.5.2 – What is implemented eventually for GLINK?

(NG) Page 40, section 5 para 3 – Where is the complete design and all sources of the firmware archived?

(NG) Page 41 – How will the analogue performance of the production PPMs be checked?

(NG) Page 41 – We must confirm that a PPM is able to provide error-free LVDS streams from all outputs while VME, CAN and TTC traffic are underway.

(NG) Are there any spare pins on any of the FPGA devices? Where are they tracked to?

D THE FINAL PRE-PROCESSOR MODULE (PPM) (REVIEW DOCUMENT)

http://www.kip.uni-heidelberg.de/atlas/L1/DISCUSS/FDRev_PPM_upd1.pdf

(EE) Page 15, para 2 - "A test with CAN software might be carried out at CERN." A test (location not relevant) at minimum of the CAN interface and programming, is ESSENTIAL, not optional.

(EE) Page 2, 2nd reference - "The 'programming section' ... will be actualised after the review' - you mean added (rather than "invented")? If so, good, it's essential - but what about the tables in section 4 of that document? Isn't that what they are?

(EE) Page 2, last reference - Better to give URLs to the EDMS document, rather than local ones, to be sure of getting current version.

(DC) Page 3, Figure 1 - Mechanical stiffening: it appears that there is less stiffening in the "AnIn" region of the board, if I am interpreting figure 1 correctly. I guess this may be because the AnIns themselves provide additional rigidity, but is it impossible to continue the horizontal rods through to the front panel? (Would this actually help?).

(WQ) Page 4, "Crosstalk on ANIN" - What is the trigger requirement on the crosstalk level?

(DC) Page 4, bottom - You mention that there is no observable cross-talk. Can you give more information please? I.e what is the sensitivity (cross-talk limit)? How many AnIn_2.0's have been measured?

(NG) Page 5 - How linear is this? What are the residuals like?

(NG) Page 6, para "DAC Range" - Are these figures for one channel of one PPM, or averaged over several?

(NG) Page 7, plot - Again, what is the linearity like. Does the FADC become non-linear at the top or bottom?

(DC) Page 8 - Phos4 - Can we monitor in operation if this is working correctly? Is it obvious, or do we have to wait for a timing run?

(PF) Page 8 - The PHOS4 behaviour is really a pain and the proposal to use the DLL in the Rem_FPGA looks the right thing to do. I do not completely understand the question of reloading the firmware in mid-operation. If this happens, I am afraid that the clock seen by the Phos4 will be disturbed anyway and that a re-initialisation is necessary.

(PF) Page 8 - Do you foresee a systematic check of the PHOS4 after initialisation and/or during the run?

(SHa) Page 9 - It is a bit surprising that this pre-compensation circuitry is necessary to transmit the serial LVDS signals at 480Mbps over 30cm of FR-4 when you can go over 15m of cable. Doesn't the pull-up resistor on the negative wire of the pair unbalance the differential pair (I guess this relies on the fact that most of the time zeros are transmitted)? It would be nice to see an eye-diagram of the signal at the LCD input.

(ML) Pages 9,10 - The LVDS cable length of 12-15m is mentioned. Actually it is 11m. Results are shown for tests with 15m cables. Has the new LCD been tested with 11m cable to CPM or JEM? Is this needed before production?

(SHa) Page 10 - There isn't much detail on the BER measurements, it would be interesting to know what pattern was used, how many outputs were active, if the test was done including the cables, how long the test run were and if any errors were observed at all.

(PF) Page 11, Figure 11 - The eye pattern shown (cable transmission) looks very nice. Do you use only random patterns or have you also tried pattern with long sequences of zero and/or one which then give problems?

(SHa) Page 11 – It would be nice to also see the eye-diagram at the receiver, i.e. after the cable.

(WQ) Page 11, Figure 11 - Are these measured with LCD_2.0 on PPM_2.0 with latest (final) TTCdec? I suggest to measure this eye diagram with several PPM_2.0 with final TTCdec. If all looks good, then we can start production for TTCdec.

(NG) Page 12, "Firmware Code" - Some work is still needed to check that all errors which can be detected in the RemFPGA are actually reported to DAQ. Also the behaviour to disable individual input channels needs to be updated.

(WQ) Page 12, bullet 1 - Has this new firmware loader function been tested? How about f/w loading upon power-up?

(SHa) Page 12 - Many of the other L1Calo modules seem to use the Xilinx SystemACE for configuring the FPGAs, was there a specific reason not to do so on the PPM?

(DC) Pages 12, 14 – Second CPLD: You stress that the CPLD code for VME is well debugged. What about the code for the second CPLD? Is this fully written and tested/debugged, too?

(WQ) Page 13, para 1 - Has this new RAM been tested?

(WQ) Page 14, Figure 14 - I'd like to see a scope shot of the clock after PPM on board distribution since this was problematic on PPM_1.0.

(NG) Page 14, "Can controller" - I didn't find any definition of the VME registers to access the CAN device, nor what the individual bits mean. We have to be very careful that VME does not interfere with the normal operation of the CAN device.

(SHa) Pages 14, 29 – Is the design of the CAN controller exactly the same as on the other L1Calo modules?

(WQ) Page 15, "Use of Hot-swap" - Does this mean that now one PPM can be changed while other PPMs in the same crate are running?

(SHa) Page 15 - How many layers does the PCB have? Are you also using lead-free components on the board? Having a mixture of leaded and RoHS compliant components can give problems if it is not properly handled at the assembly. Are you planning to use the JTAG ports on the board, e.g. for boundary-scan testing?

(WQ) Page 16, bullet 1 - There is no power supply oscillation any more now with the new arrangement of power cables. But the current consumption in the PPM crate was not at the peak during the test either.

(PF) The problems with power supplies seem to be understood and the proposed actions look OK. Have you checked the behaviour for several loads (i.e. different number of modules in the crate)? Are you sure that a change in current consumption can not initiate an oscillation?

(DC) Pages 16, 17 - Power oscillations: the inductive loops cannot be made of precisely zero size. While you have shown that the oscillations can be strongly suppressed, how confident are we that they will not recur? Could we be "close to the edge"?

(DC) Page 17 - Similar comment on 3.3V current overload. How close to the limit are you by increasing the current limit by 20%? Would you be happy to raise this further, if needed?

(WQ) Page 17, bullet 1 - What is the maximum programmable current of the controller? How does that compare to the peak current on PPM? Any chance that the PPM peak current could exceed that maximum limit?

(WQ) Page 17, bullet 3 - Does this mean MCM can work at 120 Celsius degree continuously?

(NG) Page 13, RGTM - Is there actually a register where the status of the RGTM can be read?

(DC) Page 18 - Is the slot 11/12 poor-cooling problem fixable? Could it be worse in the USA15 set-up?

(EE) Pages 18, 19, Figure 17 - The MCM temperature plots are very interesting. Although the situation seems well under control, I wonder whether the "valley" in the centre of the crate could be smoothed out by a small, simple air deflector or two in the gap between the fans and the bottom of crate.

(SHa) Page 19 - I didn't understand the comment about the "service plug-on", a picture of how this fits in could be helpful.

(DC) MCM temperatures - you mention that they get into the 60 degrees C range, and that previously MCMs have recovered from being heated to 120C. But what is the maximum operating temperature that you are comfortable with? Can we quantify what is the danger point?

(SHa) Page 21 - The initialization problems observed with the PHOS4 are worrying. We also use the PHOS4 on some of the modules of the central trigger (only few devices per system though) and the initialization procedure is somewhat different to the one you describe: the PHOS4 is supplied via a power switch is only enabled once the reference clock is stable and the I2C is high.

You can probably also adjust the duty cycle of the 40MHz clock by putting a Thevenin termination on the output of the crystal oscillator.

Special clock buffer chips (instead of the general-purpose tri-state buffers used) would also give a lower duty-cycle distortion, however this would require a design change. You could however try using pin-compatible buffers from a different logic family, e.g. LVT instead of LVC.

A concern using the clock from the FPGA DLL instead of the crystal oscillator clock is that this clock signal would only be available significantly later (after the FPGA is configured). Has this approach been tested?

(NG) Page 22, 23 - Will you please confirm that the G-Link will run from a low-jitter 40.00 MHz crystal, not from the TTC?

(WQ) Page 23, para 1 - The last statement is not correct. The production ROD will be able to operate its G-link receiver at frequencies other than 40MHz. The G-link simplex mode II has been successfully tested on current ROD and will be incorporated into production ROD design. Should you want to change the PPM readout frequency, is the current RTGM capable of doing this without any hardware change?

(WQ) Page 24, Figure 21 - Is it the plan to use the home-brew VME-CPU in the final PPM system in USA15? Why put VME-CPU in slot 3 rather than slot 1? I think some VME controller functions need to reside in slot 1, namely interrupt handler, although such functions may not be used in PPM crate.

(SHa) Page 24 - It would be nice to have more details on what tests have been done with the full crate, for instance has a significant fraction of the LVDS cables been connected in the tests?

(NG) Page 24 - The picture shows Homebrew in crate. Is it clear that the production test must be done with Concurrent (as the system will run at CERN)?

(EE) Pages 24-25 “Full Production” (Planning, Testing, ATLAS Check-out) – They seem to me to be the most inadequately described items in the document; hopefully we’ll hear more at the review.

(NG) Page 25, "Atlas Check-out", point 3 - We agreed that crates should be sent when complete. However, the crates are needed for LVDS cabling, so a complete crate should not mean "Complete and full of tested modules".

(PF) The TTCdec includes a PLL which is I think necessary. Is the clock delivered by this board used to drive the G-Link? If yes the PLL must be used.

(SHi) There is no mention of changes to the output pinning of the V8 and W8 LVDS fan-out signals as was requested by me at the IDR. Was this actually done? Actually I've checked with the schematics (Page 41) and from what I can see, this has been done, which is good news.

E THE PRE-PROCESSOR MODULE (PPM V2.1) SCHEMATICS

http://www.kip.uni-heidelberg.de/atlas/L1/PP_Module/PPM21_schematics.pdf

(SHa) Page 9 - Analog ground is connected via an inductor (I assume that these can be replaced by wire links if needed) to the chassis ground, the link to the to the system (digital) ground can be made via jumpers at the MCM (page 24). What is the configuration that will be used in the system?

The CPLD supply seems to be generated by two diodes in series from the 5V supply. This is not very accurate; is there a specific reason not to use an LDO regulator instead?

(SHa) Page 14 - Can the dedicated clock inputs on the CPLD be left open? A series resistance of 1K for the JTAG pins seems very high.

(SHa) Page 17 - It seems that it's only possible to use the JTAG interface either on the MCMs or the FPGA/CPLDs/TTC (selection via solder jumpers?). Is it not necessary to access both? Also there are many devices on the MCM chain and no additional buffers. Is this working properly? We have had problems with long JTAG chains, especially when they include devices from different vendors.

F CHANNEL MAPPING

http://www.kip.uni-heidelberg.de/atlas/L1/PP_Module/ppm_channel_mapping.pdf

(SHi) I very much like this document, but as the DRAFT status suggests, it is not quite complete or correct. Here are a few comments (some of them repeated from the last time I saw a version of the document):

I think I have only (so far) found a couple of genuine errors, the rest of my comments being suggestions for improvements and additional information. The errors are associated with the good old V8 and W8 outputs (surprise, surprise). I think what is shown here corresponds to the (incorrect) specs for these being on rows 6 and 20, and so not to the actual (correct) implementation derived from the schematics, where they appear on rows 7 and 21. The odd thing is that I think I saw a version of this document where this was correct, so I don't know how the error got re-introduced.

On the eta/phi mapping issue, although it may be correct that this document shouldn't cover all the horrible endcap mapping issues for eta/phi, it might be nice to have the eta/phi mapping for a 'standard' PPM. What I mean by this is that for 112 out of the 124 PPMs, the eta/phi mapping is constant and relatively un-messy. So it might be nice for general understanding of the easy central area to include these eta/phi mappings in the form 1-4 (or 0-3) and 1-16 (or 0-15 starting with 1,1 (or 0,0) as the lowest eta/phi cell. Thinking again about the endcap nastiness, Victor produced some very nice diagrams for his talk at the RAL joint meeting. This document might be a good place to include those diagrams and so cover all the horrible mapping issues.

A minor point on numbering. You have numbers such as A1 and A22. I often find it nicer, if there is a need for two digits, to make even the single digit numbers the same length – i.e. A01, A02 etc. It just makes tables/labelling look more consistent. Also on digital numbers, I believe Florian and others tend to number from 0 not 1, so maybe D00-D63 makes more sense to most people than D01-D64.

Another minor point on the tables. On the LVDS labels you have repeated the names in capitals and non-capitals for PPM/ CPM or JEM labels – e.g. (WX/wx). While this is no doubt following a convention from the Cabling Bible, I'm not sure that this convention is either intentional, or likely to be followed through in software or hardware labels. I don't think there's any need to distinguish the capitalisation, just the letters themselves matter. So you can probably save a little space in the tables by just having one version, and I don't care if you chose this to be capitalised or not.

A more substantial issue is the way columns 4 and 5 are done in the table (MCM channel map and digital num). I actually find this a little confusing, as I was expecting to follow an individual line through from the left hand side of the table through to the right. However, the meaning of each line changes for columns 4 and 5. While I understand (now) that this is implied by the different ordering of the analogue channels in columns 2 and 4, I did find this rather confusing at first. I wonder if it would make more sense to just get rid of column 4, and reorder column 5 so that the correspondence between these numbers and the analogue channels in columns 1 and 2 is more explicit. You could add a small column for MCM channel number instead, which would be a repeating 1/3/4/2 (or should it be A/C/D/B?).

G PRODUCTION AND TEST ISSUES

(EE) What is the production schedule; it's not really shown

(EE) What tests will be done at Heidelberg? Who will do them? How long might they take?

(EE) What tests are needed at CERN to integrate the modules? Who will do them (especially in the new situation regarding effort)? How long might they take?

(PF) The problem of soldering of the MCM sockets should be solved. Reading your documentation it looks like you got a number of problems with them with a 20-board production. I don't think you can afford so many errors with the full production. Have you got a discussion with the assembly company in order to understand what the problem(s) are? Some changes my be needed.

(DC) FINAL PPM Page 15 "Assembly of PPMs": you say that the 0.5mm pitch MCM socket soldering is tricky to repair, and has been found faulty in the pre-production. What action is being taken to avoid this in the rest of the main production? Will the assembly company pay closer attention?

(SHi) In the tests of the LVDS outputs, is the check on the data just at the level of parity correctness, or is it checked that the correct data from the right channel is appearing in the right place?

(SHi) I would like to see all the LVDS outputs checked with a CPM or JEM or both before more production is done. It would also be good to see a full load on the LCD in these tests, i.e. with all of the LVDS outputs for the three typical PPM configurations connected (central, PPM8 and PPM8a).

(SHi) Have there been any tests of the CAN functionality? If not, should some be done before full production is started?

(SHi) How do the results of the crate test power consumption compare with the predicted figures. (Maybe I missed this in the write-up.)

(SHi) What's the status of testing of readout with L1As that are very close together? Is there any suggestion of FPGA limitations in this area?

(WQ) Since now we have a full crate of PPMs and a full crate of CPMs and 9U RODs, a full crate test of PPM in realistic environment should be pursued.

(PF) A detailed production and testing schedule would be welcome.

(DC) The timescale for production is not very clear to me from the document. We will need to talk about what is reasonable, and what the expected test rate would be. On this issue, for a PRR it would be nice to see a more detailed test plan for the modules, with a procedure defining what tests will be done in sequence, with time estimates (if successful, at least, plus perhaps any information so far available on time needed for any rework, and likely fractions needing rework based on the first 20).

H PRODUCTION/TESTING PLANNING

http://www.kip.uni-heidelberg.de/atlas/L1/DISCUSS/FDR_PRR_Planning.pdf

The following series of comments and questions was compiled shortly after the review, once the reviewers had been able to make a careful assessment of the above planning document. The list of issues was submitted to the Heidelberg group, and summarised versions of Paul Hanke's responses are shown after each point.

i) Suggestions for accelerating the schedule

a) The time to completion of the full production and test programme to some extent obviously depends on the number of modules, which it is proposed to increase from 140 to 160. Although erring on the side of safety, this could easily add some time to the schedule, depending on the organisation of batches. Consideration should be given to reducing this number if it would save time (it would clearly save valuable effort for testing).

*The aim is to obtain 128 functioning PPMs (8 crates*16 PPMs) as soon as possible directly from the production line. The experience of the 20 pre-series PPMs indicates that the initial yield of such modules is 75-80%, so at least 160 PPMs would be needed to yield 128 modules that require no re-work.*

b) Rather than using a 9U ROD for testing the readout, requiring effort from ROD experts, etc, it would be simpler, and therefore faster, to use a DSS with associated G-link daughter-card.

This will certainly be considered, if there is DSS-type hardware available for loan. Software (e.g. C-code), which can be tied into a test-package, must also be considered.

c) Commission a second company to manufacture and assemble the CAN daughter-cards in parallel with the other board production, rather than wait for slots with a single company.

The assembly company has already started production of 160 LCDs and 160 CAN Controllers (November-/early December 2006), and this work should be complete by 08 December 2006 when the 160 PPM PCBs are scheduled to arrive. No hand-assembly work is required for these daughter-cards. Engaging, checking and setting-up with a second company would create an additional administrative overhead (and new risk), for which technical people would be required, but none is spare.

d) Explore the possibility of accelerating the manufacturing schedule by paying a premium for faster turnaround. Although the cost may be high, this could possibly save a few weeks.

Realistically, it is not worth it. For the PCB production, to save two weeks would increase the cost by a factor of 2-4. For board assembly, the SMD work is automated so the time taken is negligible. The electro-mechanical and QA work is labour-intensive and therefore inherently time-consuming. The experience and skill for the careful, thorough work which is needed here cannot necessarily be purchased easily.

e) When placing the contracts for production, ensure that the companies already have all components available, or can show credible lead-times. Some devices often develop very long lead-times, which can clearly damage planned schedules.

Agreed – this is self-evident. One relevant issue here is to ensure that the Parts Lists are really complete, and that nothing has been overlooked.

f) Rather than expending valuable effort in Heidelberg in repairing faulty boards after delivery, explore the possibility of getting the assembly companies to do this work and only deliver fully-working boards.

Engaging the assembly company for this type of work will definitely be considered, assuming that the balance between logistics and repair-effort justifies the action.

g) "Batching" of the production boards should be explored further, in order to produce at least some modules earlier than currently proposed. This may involve paying some price premium.

Only the manual assembly-work is batched, and the module testing phase (mainly identifying those PPMs which are initially good) can just be synchronised to that, as discussed in "The FINAL Pre-Processor Module (PPM) for the ATLAS Level-1 Calorimeter Trigger". *The more time-consuming debugging work on the remaining PPMs would be carried out at a later time.*

h) The highest priority should be given to delivering working modules to CERN. Valuable effort (and time) should not be spent on fixing faults on a few modules (assuming that the yield is high enough to produce a sufficient number of initially good modules).

This issue is partially addressed under the response to item **1g** *above. The primary objective is to install 128 fully-working PPMs in USA15, and only after this to check out the other modules in order to build up a stock of fully-working spares.*

ii) Further questions and clarifications

a) There is not very much detail about the proposed test programme for daughter-cards, of which there will be a very large number (almost 1,000?). More details should be provided.

This effort is ongoing, and the details will be provided at the same time as the software. The requirements are known, and the infrastructure will come together.

b) Exactly what tests will be carried out on the production PPMs? Will there be any repetition of any parts of the MCM tests, but this time on populated modules?

It is unnecessary to re-test the MCMs in fine detail, but their functionality must be verified. The KIP group will define a relevant sub-set of tests (or even run them all, if time is not the issue). Each PPM will spend ample time under power in the test-rig.

c) There is some concern about the state of readiness of the test software. How much is still to be written. and how much debugging is still required?

There is a software nucleus still existing from the MCM test programme, and many pieces of code are around, but putting all this together and debugging it needs eager people to work now.

d) There is some apparent disagreement between the 2-3 PPMs per day which can be electrically tested (page 2) and the 4 PPMs per day shown in the project plan (page 5). This should be clarified.

It is not a contradiction if the 75-80% yield expectation is met (see response to **1a** *above), as there should be 15-16 initially-good PPMs to test electrically out of every 20 modules delivered.*

e) The current yield of 75% of good PPMs from the manufacturer is already rather low. What plans are in place to deal with a further reduction in yield, and how could that be prevented?.

f) Is it planned to maintain an historical database record of which daughter-cards have been used to populate each motherboard?

The PPr SQL Data-Base was prepared from the start for this purpose, although minor extensions may be needed once some test-experience has been gained. MCMs will be classed as either working or non-working, whereas other items could be repaired. The relevant database fields could perhaps be: "Repair date", "Test date", "OK?", "Installed on", Installation date".

iii) General comments (including installation/commissioning issues)

a) All of the production crates should be prepared and delivered to CERN as soon as they become available, in order that the LVDS cabling can proceed (although special precautions should be taken to maintain their cleanliness). They should not wait for the delivery of tested PPMs, which can be shipped separately to CERN.

Agreed. The importance of the LVDS cabling installation is accepted, and fully-prepared crates will be installed in USA15 as soon as possible. This process has already begun (although the two crates currently installed are not yet completely prepared).

b) Only well-tested PPMs should be delivered to CERN, as fixing faults could be very time-consuming in the less benign environment at CERN.

Agreed. The aim is that no repair work will be carried out in USA15 or in Bat. 3150. In the case of a PPM failure, a meaningful error-report will be generated and the complete module will be exchanged for a spare one.

c) Very careful planning of the installation and commissioning of tested PPMs in USA15 will be essential to maximise the efficiency of achieving a functioning trigger as soon as possible.

Agreed.

I TEXTUAL COMMENTS - THE PRE-PROCESSOR MODULE (PPM) (SPECIFICATIONS)

(EE) This document is essentially the PPM specification, but it needs material in some areas. "Status report" type material needs to be removed or changed, especially since it's often obsolete. Also, it sometimes reads a bit like it was "pasted" together - sometimes repeats things, and occasionally is out-of-date.

(EE) To be added:

- S/W model (as mentioned in the other document; but isn't that done in Section 4 already?)

- Various figures and accompanying descriptive material from the "Review" document, for example:
  - Figures 4, 5, 6, 10, 12, 19
  - Appendix on CAN (and schematic diagram)
  - Material on firmware storage and flash-RAM loading
  - Material on PHOS4 and its problems
  - Material on hot-swapping

In other words, some of the explanatory material of the "Review" document should be merged in.

(WQ) Throughout the document, several "should", "would", "could" and "has to" are used as below, which makes it unclear whether a function is implemented or not:

- Page 18, Section 3.4.7, para 2 - "Placement of these components **should** ..."
- Page 18, Section 3.4.7, last para - "Hence, the dissipated power of the Pre-Processor Module **should not** ..."
- Page 20, note box - "It **should** be possible to ...", "The trigger **could** ...", "Using available resources on J1 **would** ..."
- Page 24, para 2 - "VME 'block read' **should** be able to..."
- Page 36 para 1 - "selective readout **could** be implemented..."
- Page 23 para 2 - "The readout merger **has to** ..."

(As for the Review document, I'd like to know what has been done and what is not, rather than being given some possibilities.)

(EE) Page 2, section 1.1, item 2 - Formally, the name of the JEP is "Jet/Energy-sum Processor" (although for the JEM we dropped the "sum"). I like to leave "sum" in for the JEP because it actually describes the type of trigger it's producing.

(EE) Page 2, section 1.1, next-to-last paragraph, line 5 - "**Programmable**" not "Programmbale"

(EE) Page 2, section 1.1, next-to-last paragraph, line 7 - "**auxiliary**" not "auxilliary"

(EE) Page 3, section 1.3, last line - "**Timing, Trigger and Control**" not "Trigger Timing and Control"

(EE) Page 4, section 2.1, Figure 1 - It would help if the outputs on the right were labelled "CPM" and "JEM".

(NG) Section 2.1.1, para 1, line 2 - By charge-balanced, do you mean a/c coupled? Does this refer to the incoming signal, or the coupling on the PPM? The spec should define both.

(NG) What is the expected signal amplitude, and from what impedance? What impedance does the PPM present?

(NG) Page 5, para C - I don't understand the statement about the working range.

(EE) Page 8, section 2.2, Fig. 4 - Obsolete layout; please update it.

(EE) Page 8, section 2.3, last sentence - A specification document shouldn't be giving the current status; delete.

(EE) Page 9, section 2.4 – It would be clearer if it mentioned the auxiliary backplane and how it connects all the J0 slots.

(EE) Page 10, section 2.6, para above Figure 5 - The sentence starting "In this example" is NO LONGER TRUE, and in any case the diagram doesn't indicate anything about allocation of the channels. Delete the sentence, and in the following paragraph start with: "**A particular**" rather than "This particular".

(EE) Page 10, section 2.6, TCM paras after "Important Note", last line – "**panel**" not "panle"

(EE) Page 12, section 3.3, para 5, last sentence - NO LONGER TRUE; all cables to a crate of PPMs come from either below or above, they are not split.

(WQ) Page 13, Figure 7 - Front panel led G-link (RGTM) locked does not agree with the description "DAQ readout activity". If this led is for G-Link lock status, it is better to use green colour.

(WQ) Page 13, Figure 7 - Front panel led PPM READY (clock from TTCdec): I wonder here why use 100 ms stretched clock from TTCdec as PPM READY indicator. Even without TTC input, the TTCdec still output a clock signal. So here, TTCReady from TTCdec should be used to indicate PPM READY.

(EE) Page 16, section 3.4.5, bullet after "Note" - "**synchronous**" not "synchroneous"

(NG) Page 18, para in box - What fraction of the various RemFPGA resources are used?

(WQ) Page 18, Table 1 - It is said in text above the table that some figures are measured and other are estimates. But in the table, it not indicated which is estimated and which is measured.

(EE) Page 19, section 3.4.7, italics just before "Local Power Control" - Now that you have a full crate, has this now been done?

(EE) The PPM specification should spell out the BC-MUX scheme details. People should not be expected to go back to the TDR (with its discussion of two possible options, and which we chose).

(EE) Page 19, section 3.4.7, "Local Power Control" section - We hear about the possibility of detailed error codes using the ATMEGA, but how much is actually implemented and is there a list? (i.e. is this alphanumeric display actually doing anything useful?)

(EE) Page 20, section 3.4.7, Note at end of section - May be valid, but no longer possible as TCM design is finished. Delete.

(EE) Page 21, section 4, para 2, line 3 - "**auxiliary**" not "auxilliary"

(EE) Page 21, section 4.1, para 1, last line - "**loading**" not "laoding"

(EE) Page 21, section 4.1, para 3, line 4 - "**available**" not "avialable"

Note that the information in this paragraph is more or less repeated, with more detail, at the end of section 4.3.1.)

(NG) Page 21, Section 4.1, bottom - I think this is fine, so make black not grey. Are these address bits also used for TTC and CAN addresses? Please give the precise recipe used for TTC and CAN addresses.

(SHi) Page 21, final sentence - I don't think the crate IDs are any longer in doubt. PPM crates will be 0-7 and the assignment is given in the cabling document.

(WQ) Page 22, Figure 13 - Only 2 Mbyte address space allocated for 8 MByte Flash-RAM?

(NG) Page 23, Section 4.2.1, para 1 - Are the (36-32) = 4 bits not accessible from VME ever used. If not, please say so. I think what is meant is that all addressing on the module is in longwords.

(EE) Page 23, section 4.2.1, box - Please bring this up to date; current situation and plans? What about section 4.3.2?

(NG) Contents of the box - The VME model has to be documented for the production module, down to the level of meaning of individual bits (so that someone coming in a few years with no knowledge of the module can work with it).

(NG) Page 24, Section 4.3.1, bullet 1 - You've just said that the DACs and Phos4 are write only. So there are indeed write-only registers (EE) Page 25, section 4.3.2, para 5, last line - "**laid**" not "layed"

(NG) Page 26ff, list of registers - This is a long impressive list. In several places there is greyed-out text or question marks - clearly to be resolved. Where is the content of the AnIn registers defined?

(NG) Page 26ff, list of registers - Your question: Is the PPM.DAC format correct?

(NG) Page 26ff, list of registers - I don't understand the addresses. For example, the picture shows AnIn starting at VME byte address 0x17FFC0. The next one should be at 0x17FFC4. In the table, the first reg is 0x17FFC0, but the second says 0x17FFC1. What do you mean in the heading by "Offset/byte*4"? Much better just to give the byte offset, as we do in the other module specs.

(NG) Page 28, PprDAC_Channel_1 etc - I thought the DAC registers were Read-only?

(NG) Page 28, PprSIFr_ChannelAB - Text should explain what happens if you write or read these registers.

(NG) Page 29, Table 3 caption - Back to MCM numbering 1-16 !! Should be 0-15.

(NG) Page 29, Channel registers - Why not give them names which say what they do? e.g. PPrChannelSatLevel rather than PPrChannelControlReg[xx].

(NG) Page 30 - Does the PPM provide memory for TTC dump command? If so, where?

(NG) Page 31-35, Register descriptions - The big thing missing here is a description of what individual bits do. e.g. does VMEReset perform a full reset of everything on the board, or just reload the VME FPGA? There are many bits where the exact function isn't clear (at least to me).

(NG) Page 34 - At PPM.CR33, there's a reference to ‘override’ mode. What is overridden?

(NG) Page 34 - After PPM.CR33 there follows a list of TTC registers. Do these duplicate what we've already seen?

(NG) Page 35 - Sections need numbering.

(NG) Page 35, para 1 - Isn't this just a complicated way of saying that the RemFPGA receives commands and data from VME and writes the data to the destination device via the appropriate bus?

(NG) Page 35, para 2 - Surely you don't mean that registers can only be read if there are L1As?

(NG) Page 35, third para - what happens if a Read or Write is done while "blocked"? Reading and (sometimes) writing registers while the module is active is essential for debugging. Although the numbers of timeslices can be set randomly as described, the code in the RemFPGA can't cope unless they are all the same, and nor can the ROD.

(NG) Page 36 - I would like more details of the event and histogram access. The text implies that the copy of readout data may or may not be produced, depending on how busy the RemFPGA is. How does the user know?

(NG) Page 36 - Outside DAQ, how does the PPM send data to the CPU?

(NG) Page 36 - In the table, what are the roles of the LUT and FADC reference and readback sections? What is the format of all this information?

(NG) Page 37, section 4.4.3 - At least the terms Mem, Local, Global, Chx and Pipeline mask should be explained here.

(NG) Page 39 - Diagram (and calculations) out of date compared with ROD specifications (and what is implemented in PPM).

(WQ) Page 39, Figure 16 - This diagram is obsolete.

(SHi) Page 39 and start of page 40 - Most of this information is seriously out of date. Needs bringing into line with current G-link format.

(EE) Pages 40-41, sections 5 and 6 - Out of date. The paragraphs on testing at KIP are OK and useful (might need updating?), but the other stuff on planning needs to be brought into the present and rewritten like a specification, not a status report.

(NG) Page 40, section 4.5.2 - Does the PPM have control of the G-Link Tx modes. Which registers?

(NG) Page 40, section 4.5.2 - No mention here that we are running with a crystal at 40.00 MHz.

(EE) Page 40, section 4.5.2, para 2, line 3 - "**synchronous**" not "synchroneous"

General comment - this section hopes for a 16-bit not 20-bit solution. But what we have works! Please rewrite this section to conform to reality.

(NG) Am I right that the current implementation uses the SRAM as a cache of register values in the ASICs. This needs to be clearly explained, and was given to check that register values have actually arrived in ASICs. If I write to a register in a missing MCM, then read back the same register, what do I get back - I think still the same value?

(NG) Page 42, Figure 17, and Page 48, Figure 19 – Figures are hard to read.

(SHi) Page 46 - Table 13 still holds the (hopefully) old version of V8 and W8 as presented at the IDR. It was pointed out then that the JEM signals on pins 6 and 20 should really be on 7 and 21. I hope that this was fixed, and the schematics seem to confirm this.

(NG) Page 48, table at bottom - Backplane is **VME64xP**, not VME64x.

(NG) Page 48, section E - ...user-defined pins **IN ROWS C AND D** to...

(NG) Page 49, section F - Now we use the aux backplane, not wiring.

(NG) Page 49, section F - ...standard **VME64xP** backplane (not VME64x)

(NG) Page 49, section F - Please describe the diagnostic board and list the signals and what they mean.

(NG) Page 49, section F, Table 15 caption - **VME64xP** backplane (not VME)

(EE) Pages 50-51, References:

- 2: "**Timing, Trigger and Control**" not “Trigger, Timing and Control”
- 4, 5, 10, 12, 14, 15: Give EDMS document number and URL, not "our" web page; in EDMS you get the latest posted version.
- 13: ?

(NG) Page 51, References - Something is missing just before refs [12] and [14].

(NG) Page 51, Ref [15] - Current document as reviewed is 1.08 and dated June 2006.

(NG) Page 54, Glossary:

- *Corrections:*
  - G-Link: ...chip from AGILENT (add reference)
  - S-Link: **160 Mbyte/sec link**....
  - **CANBus**, not CanBus

(EE) Page 54, Glossary:

- Add entries for CPM and JEM
- *Corrections:*
  - G-Link: "**Gigabit rate serial transmit/receive chipset**" (not just transmit; add "**serial**")
  - JEP: "**Jet/Energy-sum Processor**"
  - LHC: "**Collider**" not "Coillider
  - RGTM: maybe add "**RGTM-O - Optical version**"
  - TTC: **Timing, Trigger and Control** (not Trigger Timing ...)

(NG) The following information is missing:

- I think there should be links to manufacturers’ specifications, at least for the FADC device.

- There is no description of DCS implementation except in the "PPM Modifications" note, where the DCS daughterboard appears for the first time.
- The requirements are listed in sec 2.5. There should be a statement about device used and a complete list of the monitoring voltages, currents, temps available to it.
- Details of RemFPGA algorithms are not described: data processing, error detection, handling missing channels...
- There's no statement about VME64xP configuration space, apart from the assertion that the AM code is supported.