



FDR/PRR of the L1Calo Readout Driver: Report

ATLAS Project Document. No. ATC – RD – ER – 0033	Institute Document No.	Created:	29-Sept-06	Page:	1 of 30
		Modified:	5-Jan-07	Version:	1.1

Report of the Final Design and Production Readiness Review

Level-1 Calorimeter Trigger Readout Driver

Abstract

The combined Final Design and Production Readiness Review of the Level-1 Calorimeter Trigger Readout Driver Module (ROD) was held at CERN on 15 August 2006. There were presentations and discussions of the hardware design, firmware design, performance in a wide variety of tests, and plans for production and testing of final modules.

The ROD has been available in a full-specification prototype version for some time. It was concluded that the design has been largely verified by extensive testing, and is basically sound. However, the ROD is a complex module that must be able to handle variable amounts of data from a number of very different trigger modules by the use of many firmware variants. It had not yet been possible to complete a few important tests, which were planned for the near future.

A large number of points, many quite detailed, were raised by the reviewers. Emphasis was placed on resolving those impacting on production of the hardware, but many useful comments concerning testing, firmware, data formats, and documentation were also submitted.

The procedure agreed was to list actions and recommendations, especially those needed before production can begin, and to identify individual reviewers to evaluate and discuss progress for groups of items. When all critical actions are completed production can begin.

A full compilation of recommendations and comments is included in this report, as is a list of editorial corrections to the ROD specification and other documents.

A summary of the results of the actions is now included.

<i>Prepared by :</i>		<i>Checked by :</i>		<i>Approved by :</i>
Eric Eisenhandler, Queen Mary, Univ. of London Appendices based on compilation of comments by Stephen Hillier, University of Birmingham		Philippe Farthouat, CERN Stefan Haas, CERN Stephen Hillier, Birmingham Murrough Landon, Queen Mary Kambiz Mahboubi, Heidelberg Gilles Mahout, Birmingham Uli Schäfer, Mainz		
<i>for information, you can contact :</i>	Eric Eisenhandler	Tel. +44 (0)1235 44 5562 +44 (0)20 7882 5056	Fax. +44 (0)1235 44 6733	E-Mail e.eisenhandler@qmul.ac.uk

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

TABLE OF CONTENTS

Table of Contents2

Purpose of the review.....3

Participants in the review3

Agenda and documentation.....3

Review outcome.....4

1 Introduction.....4

2 Summary of presentations and discussion5

 2.1 Introduction to the review5

 2.2 Design philosophy.....5

 2.3 Hardware design and version history6

 2.4 Firmware design and status6

 2.5 Test results6

 2.6 Production plans and tests.....7

 2.7 Comments made during the review session7

3 Conclusions, actions and recommendations8

 3.1 Main actions8

 3.2 Other recommended actions.....9

 3.3 Document updates9

 3.4 Procedure10

4 Follow-up on actions (added in version 1.1).....11

 4.1 Main actions11

 4.2 Other recommended actions.....13

 4.3 Remaining problems14

Appendix A: Detailed comments on ROD Specification, v.1.08a.....15

Appendix B: Detailed comments on ROD data formats18

Appendix C: Detailed comments on tests done, and ROD Test Plan20

 C.1: Hardware tests and performance20

 C.2: Production Test Plan draft 2.122

Appendix D: Editorial corrections to documentation23

 D.1: ROD Specification version 1.08a.....23

 D.2: Test Plan Document29

 D.3: Level-1/Level-2 Interface Document version 3.130

PURPOSE OF THE REVIEW

The overall goal of the review was to approve the final design and production of the Readout Driver (ROD) Module. To do that it was necessary to check that:

- The proposed design meets the Calorimeter Trigger and Data Acquisition requirements.
- The implementation is technically sound.
- There are adequate resources to run the firmware, with sufficient reserves for future developments.
- Tests have demonstrated correct functionality and compatible interfaces.
- Production planning and quality assurance are adequate, including checks at critical stages.
- A comprehensive test plan for production modules has been prepared.
- The module is adequately documented.

PARTICIPANTS IN THE REVIEWReview Panel

Eric Eisenhandler, Queen Mary, London (*chair*)
 Philippe Farthouat, CERN
 Stefan Haas, CERN
 Stephen Hillier, Birmingham
 Kambiz Mahboubi, Heidelberg
 Murrough Landon, Queen Mary, London
 Gilles Mahout, Birmingham
 Uli Schäfer, Mainz

ROD team

Norman Gee, RAL
 Bruce Barnett, RAL
 Weiming Qian, RAL
 Viraj Perera, RAL
 Ian Brawn, RAL (by telephone)
 Adam Davis, RAL
 Dave Sankey, RAL (not at review)
 Damien Prieur, RAL (not at review)

Ex officioFor information**AGENDA AND DOCUMENTATION**

The review agenda, with links to all the submitted documentation and all talks given at the review, is at:
<http://indico.cern.ch/conferenceDisplay.py?confId=4903>

The documents provided for the review were:

- **Readout Driver Specification version 1.08a**, by Norman Gee, Bruce Barnett, and Viraj Perera
- **Operational Test Report**, by Bruce Barnett: summary of test results using prototype RODS
- **PreProcessor Compressed S-Link Data Formats**, by Dave Sankey; also **Comments** on status of implementation
- **Data compression simulation study**, slides by Damien Prieur
- **Test plan for the ROD draft 2.1**, by Viraj Perera, Ian Brawn and Adam Davis: plans for testing production modules
- **Specification of the LVL1/LVL2 Trigger Interface version 3.1, ATL-D-ES-0003** (also includes material on the Muon and Central triggers); relevant for Region-of-Interest formats
- Schematic diagrams and detailed technical information:
<http://www.te.rl.ac.uk/esdg/atlas-flt/specs/9UROD/index.html>

REVIEW OUTCOME

1 INTRODUCTION

A combined Final Design Review and Production Readiness Review for the Readout Driver Module (ROD) of the ATLAS Level-1 Calorimeter Trigger was held at CERN on 15 August 2006. The reviewers were Eric Eisenhandler (Queen Mary, London; chair), Philippe Farthouat (CERN), Stefan Haas (CERN), Stephen Hillier (Birmingham), Murrough Landon (Queen Mary, London), Kambiz Mahboubi (Heidelberg), Gilles Mahout (Birmingham), and Uli Schäfer (Mainz). Talks were given by:

- Norman Gee (RAL): the overall design philosophy, how it fits into the Level-1 Calorimeter Trigger system, and the need to compress readout data from the PreProcessor
- Viraj Perera (RAL): the hardware design and version history of the ROD
- Ian Brawn (RAL) and Weiming Qian (RAL): the firmware design, current status, and plans, including work on the PreProcessor data compression by Dave Sankey (RAL) and Damien Prieur (RAL)
- Bruce Barnett (RAL): tests that have been done with the ROD, emphasising tests done in a system context
- Ian Brawn (RAL) and Adam Davis (RAL): plans for production of the RODs, testing of production modules, and quality assurance

The main aim of the review was to assess the readiness of the ROD hardware for production. The reviewers had raised a very large number of issues, questions and corrections. However, although very useful, many of these concerned the firmware, readout data formats, and documentation, with little or no impact on the hardware design or production. Therefore, in order to be able to get through the review session in a reasonable length of time, it was decided that the discussion would concentrate on points that were relevant to the hardware. The large number of other items would be dealt with within the context of the review and included in the report(s), but after the review session and mainly by e-mail and telephone.

Within this restriction of concentrating on the hardware, the agenda allowed ample time for discussion. In addition to the many comments and questions submitted and circulated in advance of the review, other points were raised on the spot. The talks were clear, and answered many of the items directly. The discussion was open and constructive.

Firm conclusions were reached, leading to a list of actions to be completed before production begins. A summary of these was circulated very shortly after the review session, in order to allow work on the critical items to go ahead without waiting for this full report on the review. These items are listed and described fully in this report.

Of the large number of points not impacting directly on the hardware (and often going into fine detail):

- Items regarding the **specification** of the ROD are listed in **Appendix A**
- Items concerning the **data formats** are listed in **Appendix B**
- Items on the **testing** done so far, and on the production **test plan**, are listed in **Appendix C**
- Corrections and comments on the **documentation** are listed in **Appendix D**: D.1 contains a large number of comments on the ROD Specification version 1.08A, D.2 a few comments on the Test-Plan Document draft 2.1, and D.3 has comments on the Level-1/Level-1 Interface Document.

In order to ensure that various points are not neglected, each of these areas will be followed and monitored by one of the reviewers.

2 SUMMARY OF PRESENTATIONS AND DISCUSSION

2.1 Introduction to the review

Eric Eisenhandler summarised the aims of the review. As the main goal was to approve production of the ROD, he proposed concentrating on issues affecting the hardware and suggested a procedure to deal with the many detailed comments (most of which would not affect the hardware) by telephone and e-mail, in such a way that delays to the schedule would be minimised. He mentioned some of the main issues raised by the ROD people (who had a pre-review action list) and by the reviewers before the review session:

- Hardware ‘flaky’? – some FPGA balls have become disconnected
- G-link inputs – run 18 simultaneously, check for errors and stability
- S-link outputs – run simultaneously; try different switch modes
- Flow-control buffers, and monitoring buffers in Switch FPGA – check they are big enough
- Compression firmware – make sure it fits into FPGAs; possibility to use bigger devices (at higher cost)
- Monitor FPGA – check transfers ok, external memory and PCI interface also ok
- RoIB – test interface, especially flow control
- S-Link rear-transition module – more details of its specification, status and proposed production tests
- Production-test plan – only covers initial standalone, low-level tests, not system and interface tests
- Level-1/Level-2 Interface document – updated version needed urgently

2.2 Design philosophy

Norman Gee introduced the ROD design and requirements. He explained the design choices and how they fit into the Level-1 Calorimeter Trigger. Each trigger module stores readout data in FIFOs, with output sent in serial format by G-link chips over optical fibres to the RODs. ROD outputs use S-Links to the readout buffers (ROBIN); depending on the amount of data to be read out each ROD uses between one and four of the S-Link outputs housed on a rear-transition module. The RODs are organised to each handle a full crate of trigger modules; this requires up to 18 G-link inputs. There are 14 trigger crates in the system, but as separate RODs are used to handle RoI information from the Cluster and Jet/Energy-sum Processors a total of 20 RODS is needed, housed in two VME-64x 9U crates.

The main components on the ROD are the optical G-link inputs, five Input FPGAs each handling four inputs (except the last, which handles two), the Switch FPGA which routes output data to between one and four of the output S-Links (depending on data volume), and a Monitoring FPGA.

Readout data from trigger modules is aimed to verify digital processing and inter-module links, so both input and results are read out. In the case of the PreProcessor, digitised FADC outputs and data from the lookup tables after bunch-crossing identification and pedestal subtraction are read out. In all cases, multiple bunch-crossings around the one that produced the level-1 accept can be read out to verify timing and, in the case of FADCs, understand the pulse shapes. Data from the PPMs, where the need to read out multiple bunch-crossings is most important, must be compressed without loss in order to keep readout data volumes acceptable. Data from subsequent processing is less voluminous and simple zero-suppression is sufficient. The situation for regions-of-interest (RoIs), which are sent to the Level-2 RoI Builder, is simpler since only once bunch-crossing is read out.

ROD output data can be monitored, either directly via VME or indirectly via a feed to the Monitor FPGA.

2.3 Hardware design and version history

Viraj Perera delved into the details of the ROD design; see his talk for these. He also described the history, starting with a number of four-channel 6U prototypes that were extensively tested and used, followed by four full-specification 9U prototype modules that have been tested for quite some time. After the first 9U prototype, the Switch FPGA was upgraded from XC2VP20 to the footprint-compatible XC2VP30. The 9U design had a minor update for the third and fourth modules, including wider JTAG coverage, but no major changes to the original design have been necessary.

One ROD has had FPGA connection problems, diagnosed as some BGA solder-ball pads coming loose from tracks on the board. This seems to be due to the board somehow suffering mechanical damage, as the affected balls are near corners of the chips. At the time of the review discussions with the ROD manufacturer were underway, in order to ascertain whether the problem was an isolated incident or inherent in the manufacturing process.

2.4 Firmware design and status

Ian Brawn and Weiming Qian discussed the firmware for the ROD. This is stored in compact-flash memory on the ROD, using Xilinx System ACE. The Input FPGA has to handle and format data from the different types of trigger modules, both for DAQ and ROIs, and requires 13 different firmware versions. A lot of the effort needed to write and test the firmware is due to this very complex requirement. However, by doing this the trigger only needs one ROD hardware design, rather than several. Downstream of the ROD the firmware is the same for all the various types of data.

Ian gave an overview of the firmware in all the FPGAs and its status, including some estimates of how much of the FPGA resources are being used. He also described how the people working on the firmware are organised and what tools they use.

There were two areas not yet in a satisfactory state. The Monitor FPGA had received very little attention – the hardware interfaces were not tested, and no firmware had been written. This was due to lack of available effort. In order to sign off the hardware design, at least the Monitor FPGA interface connectivity should be checked.

The other open question concerned the PPM compression firmware. At the time of the review it was not yet clear whether it would fit into the existing Input FPGAs (XC2VP20) with sufficient resources remaining for future developments. The footprint-compatible XC2VP30 would almost certainly be adequate, but with five Input FPGAs per module the extra cost would be considerable. In addition, there was a possibility of very long lead times to obtain either type of FPGA. It was therefore considered to be a very high priority to clarify the situation by making rapid progress on finishing this firmware, and to test it in the hardware. Once the FPGA types were agreed they could be ordered, in advance of production.

2.5 Test results

Bruce Barnett summarised tests of the 9U ROD modules. RODs have been available for some time, and an extensive range of tests has been done, with readout data from all the different types of trigger modules. Data formats have been largely verified. Flow control of data output from the S-Links has been tested at high event rates. G-link stability has been checked with up to 16 inputs active. S-Link stability has been checked with all four S-Links operating.

Bruce also listed some tests that had not yet been done but which were planned for the near future. More work was needed with the CMMs. As already mentioned, PPM compression firmware in a reasonably final form had not yet been tested in the hardware. Tests of the ROD interfaces to the RoI Builder and the

ROBINs were still required. Playback mode, for testing internally without a trigger module feeding input data, had not yet been tried.

Finally, some concerns were mentioned and work still to be done was listed. The flow control has many variable parameters and is complex to configure in a stable mode. Some of the firmware has been unstable under certain conditions. Error conditions and flags need to be sorted out. The whole area of event sampling and monitoring must be addressed. Work is needed on the System ACE firmware storage to keep track of which firmware configuration is loaded. Run checkpoints are not yet dealt with.

2.6 Production plans and tests

Ian Brawn talked about production plans and series tests of the final modules, then Adam Davis described in more detail the standalone tests that will be done. Manufacture is done under a framework agreement with three firms. All three had already done extensive work for RAL with consistently good results. Both PCB manufacture and component assembly are done under a 'one-stop shop' arrangement, with the firms taking full responsibility for problems – including the cost of any expensive FPGAs if they are irreparably damaged.

The standalone tests include JTAG, to check most of the ROD module's connectivity. Then a variety of configuration and register tests is done. Optical links and various data paths are tested. The TTC and CANbus facilities are checked. After that system tests will be done, but these were not yet spelled out.

2.7 Comments made during the review session

Discussion during the review session was extensive, with too many items raised to list here. Many of the points made were adequately answered. Below are listed a selection of points that lead to actions or that must be taken into account.

- Concern that G-link chips might become unavailable; however, a large number of spares had already been purchased.
- What would happen if a crate of trigger modules loses power? The G-links would lose lock or there would be no data, so the ROD would see that. It was suggested to think about what circumstances might leave a G-link up but with no data coming in, and how to deal with that.
- What would happen if the trigger-type arrives too late? Perhaps there should be a bit to indicate that the default trigger-type was used.
- Playback mode is implemented but not yet tested. Weiming thought it would be more robust if it could start in the Switch FPGA rather than the Input FPGAs. ROD tests have used trigger modules as data sources, so the need for playback has not been obvious. However, it's an ATLAS requirement to have it.
- An agreed labelling system is needed for compact flash cards used as System ACE firmware storage. It must indicate type of module, and firmware configuration and version. It should also be possible to read electronically what is on the card. It is also desirable to be able to write data to these cards via VME.
- The schematics for the ROD should be reviewed when all modifications have been put in.
- The possibility of running the ROD at 40.00 MHz rather than 40.08 MHz was raised. Basing the clock on a crystal would avoid jitter problems in the TTC. A check might be made that the chips would all run in this way, and perhaps the unused DES2 clock lines could be modified to carry 40.00 MHz instead.
- A connector should be made so that the PCI connectivity for monitoring can be checked. The Monitor interface to external RAM should also be checked.

- Ball-grid connectivity problems have been seen on CTP modules at CERN. They appeared after temperature cycling of the module. A suspected cause is incorrect temperature for soldering, and possibly incorrect temperature profiling of the module.
- There was a comment that the module does not have enough test points for doing diagnostics easily.
- The pros and cons of upgrading the Input FPGAs in order to be sure there are adequate resources when compression firmware is running were discussed. The only objection is the higher cost, which is considerable. If it can be shown convincingly that the smaller FPGAs are adequate we should use them.
- Although a full-crate test is not possible before production, more demanding tests to try should include all 18 G-links running at once, and two RODs running in adjacent crate slots.
- Can ECR be set to an arbitrary value, in addition to being reset to zero? This should be checked.
- If there is +5 V supplied to the S-Link card it should be removed, it is dangerous.
- What range of G-link data rates can be used – might want to change in the future?
- It seems as if a different compression scheme has to be invented for each different number of bunch-crossings being read out, or for different noise and occupancy situations. This is too inflexible.

3 CONCLUSIONS, ACTIONS AND RECOMMENDATIONS

The conclusions of the review are generally favourable, and the ROD design is mature. However, a number of items that should be tested, investigated, or implemented before production begins are listed as actions, and these are detailed below. Note that some of these actions were already specified by the ROD people before the review. The results of these actions will be circulated, discussed and approved by e-mail.

3.1 Main actions

1. The timescale for the ROD is very tight, and delivery of the FPGAs may be a holding item – **FPGA availability should be checked**. In case of long lead times, **FPGAs should be ordered as soon as possible**, consistent with the next item.
2. In order to decide which FPGA types to order, the critical determinant is to know whether bigger Input FPGAs are needed to accommodate the PPM compression firmware. **A demonstration of compression firmware using the present algorithm and formats, running reasonably well in the hardware, is urgent and should be done within a few weeks of the review**. This will allow a **decision on FPGA type** to be made and an order to be placed.
3. Unless there are unforeseen problems, the aim should be to **freeze the hardware design by mid-October** in order to start the production process. The **number of pre-production modules** should be discussed – it might be useful to build more than two.
4. Problems with some **connections to FPGAs becoming bad must be understood and resolved with the manufacturer** before final production. The cause may be incorrect temperature distribution over the module during soldering of the chips. It is essential to: (a) **determine whether this is the cause of the problems seen on some RODs**, (b) if not **whether the problem is likely to occur in other RODs**, and (c) **make sure this will not occur in production RODs**. Temperature-cycling a working module is a recommended test.
5. A **test of flow control between the ROD and the RoI Builder** should be carried out, up to a rate of 75 kHz and running for several hours.
6. A **test of flow control between the ROD and the final ROS** (i.e. ROBIN) should be carried out, up to a rate of 75 kHz and running for several hours.

7. A **test with all 18 ROD inputs running simultaneously and data being read out to a ROS** should be done, at high rate checking the link stability and at reduced rate with detailed analysis of the data. **Power consumption** and **on-board temperatures** should be measured in this mode.
8. A **test with two RODs running in adjacent crate slots** should be done.
9. A **review of the final design schematics by an external engineer** should be done, after all changes have been made.
10. Some items related to the Monitor FPGA should be **checked** for basic functionality: **monitor and instruction RAM, fast links, PCI connectivity, and if possible Rocket I/O**.
11. **Termination of outputs to the S-Links should be investigated**. Examine the signals, and if they need termination try to do it using Xilinx DCI. External resistors are a last resort. In addition, consider adding an **output buffer**.

3.2 Other recommended actions

12. Fix the situation in which the **VME** cannot be accessed via the CPLD without configuring the VME FPGA.
13. Extend the **test plan documentation** to include higher-level **system tests**.
14. Make sure that the values used for **timeout limits**, notably on **trigger type**, are long enough.
15. An analysis of the consequences for the hardware if **one voltage supply fails** but the others stay up should be done.
16. **Test points** giving access to critical signals (clocks, high-speed signals) should be added to the board.
17. Check that the VME is configured correctly, in VME-standard and configuration spaces, according to the **geographical address** in all crate slots and with different crate numbers.
18. Check that the **ECR value** can be set to an arbitrary value, and check that the reset behaviour is correctly specified and implemented.
19. It is desirable to have a **playback facility**, and ideally it should be in the Input FPGAs. If possible it should be tested before production.
20. If not difficult to implement, consideration should be given to changing clock lines to the FPGAs that carry the (unused) DES2 clock to distribute the **crystal clock** instead or, if easy, to add new crystal clock lines and keep DES2.
21. If **+5 V** is supplied to the **S-Link card** it must be **removed**.
22. Check the **Aux output**.

3.3 Document updates

23. It is important to produce an updated version of the **Level-1/Level-2 Interface document**, as soon as possible. See Appendix D.3 for comments.
24. The **ROD specification** document should be updated, following the comments in Appendices A, B, and D.1.
25. The **test-plan document** should be updated, following the comments in Appendices C and D.2.

3.4 Procedure

The ROD people should inform the review chairman of progress on the actions listed above. The information will then be circulated to the reviewers in order to agree on approval for production.

In order to check that the many detailed items in the appendices are dealt with, the chairman will ask individual reviewers to follow progress on groups of items, and to discuss and resolve anything controversial with the ROD expert(s) concerned. A summary will be circulated when the review follow-up is considered to be completed and signed off.

4 FOLLOW-UP ON ACTIONS (ADDED IN VERSION 1.1)

This section contains a summary of the results of the actions required before signing off the module for pre-production. The actions are almost all completed as far as feasible, and pre-production RODs were therefore ordered on 22 December 2006. Work is well underway on the documentation, and will be followed up as described in sections 3.3 and 3.4. Additional problems that appeared during the testing are discussed briefly in section 4.3.

4.1 Main actions

1. *The timescale for the ROD is very tight, and delivery of the FPGAs may be a holding item – **FPGA availability should be checked.** In case of long lead times, **FPGAs should be ordered as soon as possible, consistent with the next item.***

Done.

2. *In order to decide which FPGA types to order, the critical determinant is to know whether bigger Input FPGAs are needed to accommodate the PPM compression firmware. **A demonstration of compression firmware using the present algorithm and formats, running reasonably well in the hardware, is urgent and should be done within a few weeks of the review.** This will allow a **decision on FPGA type to be made and an order to be placed.***

The compression firmware has been improved and its resource utilisation is now about 50%, plus about 30% for common logic, so there are now adequate resources for future developments. There is an ample speed margin. The compression firmware has been tested successfully in hardware. Therefore, it has been agreed that the smaller FPGAs currently in use are sufficient.

3. *Unless there are unforeseen problems, the aim should be to **freeze the hardware design by mid-October** in order to start the production process. **The number of pre-production modules should be discussed – it might be useful to build more than two.***

The design was frozen before all required tests were completed, so it did not cause any delay. It was decided to build four pre-production modules instead of two.

4. *Problems with some **connections to FPGAs becoming bad must be understood and resolved with the manufacturer** before final production. The cause may be incorrect temperature distribution over the module during soldering of the chips. It is essential to: (a) **determine whether this is the cause of the problems seen on some RODs**, (b) **if not whether the problem is likely to occur in other RODs**, and (c) **make sure this will not occur in production RODs.** Temperature-cycling a working module is a recommended test.*

Despite dismounting chips and examining detached solder pads, it is still not clear what caused the problem; a mechanical shock to the affected module seems to be a more likely cause than a general manufacturing problem. Additional bracing will be used temporarily during manufacture of the modules to prevent warping when the boards are heated. The thermal cycling test has not yet been done. Note that three other modules have not had problems.

5. *A **test of flow control between the ROD and the RoI Builder** should be carried out, up to a rate of 75 kHz and running for several hours.*

The test was done with the cooperation of Jinlong Zhang. The ROD provides dual S-Link outputs, one to the DAQ (ROS) and a copy to the RoIB. One can construct a number of test topologies. Our test-trigger generation allows us to simulate our hardware, so that event fragments arriving at the ROS can be checked for integrity.

With a single input the RoIB acts as a large buffer, so the L1Calo simulation need not be aware of its presence. A single CPM was configured in playback mode to generate from 0 to 7 RoIs per trigger. Before starting the RoIB tests, we ran both S-Link streams into the L1Calo CP-RoI ROS.

This allowed us to verify that the ROD hardware was correct and agreed with the simulation. Almost 100 million events were processed by the ROS, with more than 14 million checked against the simulation. We observed four events with corrupt payloads, but can attribute these to the behaviour of the ROD trigger-type timeout handling. When a trigger-type timeout occurs in the ROD, the wrong BC number is found in the fragment. This new bug has only been observed here as we are now testing with very high trigger rates (115 kHz peak), but it is not in any likelihood due to the ROD–ROS interface itself. (see section 4.3). The acquisition rate was 37 kHz, so the system was running in a regime of heavy flow control.

Next, we connected the ROD to an RoIB input card, and looped back the RoIB output into the second of our ROS ROBIN inputs. As anticipated, no flow control was exerted by the RoIB, as its large buffers isolate the XOFF from the ROS from our ROD. Heavy flow control affected the direct connection from the first ROD S-Link to the ROS. Almost 30 million events were processed, with 3.5 million checked with no errors.

Finally, we deactivated the second ROD link and ran RoI data directly into the RoIB, with a direct connection between that and our CP-RoI ROS – only one link active. In this configuration, ROS flow-control passes back through the RoIB and eventually reaches the ROD through that path. The acquisition rate was 53 kHz. Fifty million fragments were processed and almost seven million checked without error.

The ROS was then tuned to process 60% of incoming fragments, with the result that the accepted level-1 rate rose to about 83 kHz. The peak L1A rate was 115 kHz so the system was 72% live, a regime with heavy flow-control (as confirmed by the frequent S-Link XOFF conditions at the ROD). Eighty million events were processed, and almost 6.5 million checked. Of these, eight events exhibited the trigger-type-timeout error described earlier.

Two further tests checked that longer fragments (22 words – corresponding to the neutral formatting of CPM RoI data) and high instantaneous rate (6 triggers each orbit separated by 5 BCs each) produced no errors. In the former test, 8.5 million events were processed, and of these 1.2 million were analysed without error. In the latter test, 32 million events were processed and 4 million analysed without error.

During the tests, no ROS fragment framing errors were observed – indicating that headers and trailers were well-formed. *(Summary by Bruce Barnett)*

6. *A test of flow control between the ROD and the final ROS (i.e. ROBIN) should be carried out, up to a rate of 75 kHz and running for several hours.*

We have demonstrated data transmission of fragments (payload sizes 100 and 348 bytes) from our RODs to two L1Calo ROSs in USA15. Fragments from these two sources were built into events by a third L1Calo ROS. When all fragments were processed and built, we ran at about 45 kHz. Processing and building only a quarter of the events allowed running at a trigger rate of about 140 kHz average. Processing of half of the events ran at about 85 kHz, with heavy flow control on the S-Links.

The tests were short, but one longer run handled 200 million events without ROS errors. (This means that the fragments were correctly formatted, but doesn't guarantee the payload integrity.) Finally, we did a short test checking events – about 200,000 were checked without error before losing access to the ROSs. *(Summary by Bruce Barnett)*

7. *A test with all 18 ROD inputs running simultaneously and data being read out to a ROS should be done, at high rate checking the link stability and at reduced rate with detailed analysis of the data. Power consumption and on-board temperatures should be measured in this mode.*

Done, in full-crate test-rig at Birmingham.

8. *A test with two RODs running in adjacent crate slots should be done.*

The test is mainly to see if data errors arise due to any crosstalk between the modules. RODs were

installed in slots 17 and 18, with an auxiliary J0 backplane. Only one S-Link rear transition module was available. It was discovered that the S-Link module had to be filed slightly in order to fit with the auxiliary backplane. Eight G-link inputs from PPMs were fed to one ROD. The second ROD, which had the S-Link module, received G-link signals from five CPM DAQ and five CPM RoI outputs. Two S-Links were used.

After solving various link problems, a problem with setting up the second ROD due to unreliability of VME writing appeared – see section 4.3 below.

After much pain, the “noise” ROD and the “tested” ROD were operational, at around 67 kHz instantaneous, 10 kHz average, with vector 19 (4 random bits set) on the PPMs and a CPMhits-ramp (vector 3) in the CPMs. After 90.8 million events, there were 70 errors, but these were all of the trigger-type-timeout type which arose during tests with the RoIB and the ROS at about one per million events (see section 4.3).
(Summary by Bruce Barnett)

9. *A review of the final design schematics by an external engineer should be done, after all changes have been made.*

Done, by Richard Staley.

10. *Some items related to the Monitor FPGA should be checked for basic functionality: monitor and instruction RAM, fast links, PCI connectivity, and if possible Rocket I/O.*

RAM connections and PCI interface both checked. It was felt that an adequate test of Rocket I/O on the available timescale was not feasible. However, this was not strictly required.

11. *Termination of outputs to the S-Links should be investigated. Examine the signals, and if they need termination try to do it using Xilinx DCI. External resistors are a last resort. In addition, consider adding an output buffer.*

The signals seem satisfactory, and Xilinx DCI will be used.

4.2 Other recommended actions

12. *Fix the situation in which the VME cannot be accessed via the CPLD without configuring the VME FPGA.*

Done and tested.

13. *Extend the test plan documentation to include higher-level system tests.*

To be done; not critical for pre-production.

14. *Make sure that the values used for timeout limits, notably on trigger type, are long enough.*

The default value and available range can and will be increased in the firmware.

15. *An analysis of the consequences for the hardware if one voltage supply fails but the others stay up should be done.*

Done; no serious consequences found.

16. *Test points giving access to critical signals (clocks, high-speed signals) should be added to the board.*

Done.

17. *Check that the VME is configured correctly, in VME-standard and configuration spaces, according to the geographical address in all crate slots and with different crate numbers.*

Done.

18. *Check that the ECR value can be set to an arbitrary value, and check that the reset behaviour is correctly specified and implemented.*

An arbitrary value can be set in the firmware. A VME register should be added to control it.

19. *It is desirable to have a **playback facility**, and ideally it should be in the Input FPGAs. If possible it should be tested before production.*

Playback is implemented but not yet tested; any problems are likely to be firmware issues.

20. *If not difficult to implement, consideration should be given to changing clock lines to the FPGAs that carry the (unused) DES2 clock to distribute the **crystal clock** instead or, if easy, to add new crystal clock lines and keep DES2.*

Adding the crystal clock to the existing networks: All the FPGAs have enough internal resources to handle another clock domain. However, most of their clock pins have been used for (non-clock) I/O. This option would therefore require adding a new clock driver, level converters for all of these clock lines, and reworking the Switch FPGA pin-out (and therefore PCB routing). Reworking the Switch FPGA routing is arduous and hazardous, which rules this option out.

Replacing CLK40DES2 with crystal clock: This requires one extra clock driver, because the Monitor FPGA doesn't currently receive CLK40DES2. We'd also need level conversion, or to change the Monitor pin-out (and thus PCB routing). This requires less work than the previous option, but still would need some changes to existing (working) routing on the PCB.

The probability that the crystal clock will be required for these FPGAs seems no higher than the probability that CLK40DES2 would be needed. TTC jitter is mentioned as a motivation for these changes. The rest of the trigger runs off the TTC clock, and if we can trust it for the CPM (some of which runs at 160 MHz) we should trust it for the ROD. Those parts of the ROD which are susceptible to jitter – the G-links and the Switch-Monitor Gigabit links – already have dedicated crystal clocks of 40 MHz and 100 MHz respectively.

Another possible motivation, not mentioned above, is to increase our bandwidth by increasing the frequency of the crystal clock. However, it is not necessary to re-route the crystal clock to do this. Currently the crystal clock is fed to the G-links, and via them to the front half of the Input FPGAs (the formatting logic). This front section of the ROD has an asynchronous interface to the rest of the logic and can already be sped up independently of it, by upgrading the crystal clock. The rest of the ROD (the rear of the Input FPGAs plus the Switch and Monitor FPGAs) currently uses CLK40DES1. The clocks within all these FPGAs can be stepped up using the digital clock managers to the desired frequency (not just multiples of 40 MHz). **Therefore, the current hardware, is capable of running at speeds other than 40.08 MHz.** (Summary by Ian Brawn.)

21. *If +5 V is supplied to the S-Link card it must be **removed**.*

Done.

22. *Check the **AUX output**.*

Not yet done.

4.3 Remaining problems

In the course of completing the actions above, a few possible problems drew attention. It was necessary to make sure they do not have hardware implications.

- VME writes in some crate slots (e.g. 18) did not always work. This occurred only at CERN, and after extensive effort was shown to be due to a faulty TCM.
- Problem writing to playback memory. Further investigation needed, but quite possibly related to the VME writing problem just mentioned. Seems unlikely to be due to ROD hardware.
- Occasional errors (roughly one per million events at high rates) with trigger-type timeout. A firmware fix from Weiming Qian is awaiting test.

APPENDIX A: DETAILED COMMENTS ON ROD SPECIFICATION, v.1.08AGeneral

I guess the same modifications as done (or planned) on CMM will be done on ROD's TTC lines. (KM)

Sect. 1.1

Use of Agilent chip: what's the DIV1/DIV0 configuration for the ROD? Do you set the DIV1/DIV0 via VME? Does it need to be different when a CPM is plugged in? This is true for the DSS with G-link Rx, but it could be history now. (GM)

Sect. 2.1 (f)

Does this mean neutral format or something else? (ML)

Sect. 2.1 (i)

I don't remember the purpose of this, why not send data to ROS for all trigger types? (ML)

Sect. 2.2

Is there any provision to monitor the incoming raw data, and hence to check the data compression algorithms? (PF)

Table 2

Aren't PPM compressed and uncompressed formats sufficiently different to be identified with different entries in this table? Also, there is no differentiation between the compression levels. Is this really what we want? I suspect it's left-over from the days when we thought a single piece of firmware could implement all of the different compression algorithms. (ML, IPB)

Sect. 2.2.1, final para

Is it really sensible for the ROD itself to calculate these values - in many other cases software has to set the registers (and sometimes it'll have to override values in test setups) so why not just let the software do it. In fact is this already what is being done? (SJH)

Similarly: I thought the source ID is obtained from the S/W. Can it be deduced from the type of F/W loaded? (If JEM DAQ type then 0x74) (GM)

Sect. 2.2.3

What is the size of the output FIFOs to the individual S-Link channels? Although the channels work independently they appear to share the BCID, event number and trigger type buffers, so unless the per-channel buffers are large enough, flow-control on one output channel could block the others. (SH)

Sect. 2.2.4, para 1

Is ECR needed to reset the G-link timeout flag even if link comes back before that? (ML)

Sect. 2.2.4, para 4

If formatted data arrives late for one event is it flushed before the next event? Or what? (ML)

Sect. 2.2.5

How does the ROD recover from the arrival of a late fragment (i.e. after the timeout)? (SH)

Sect. 2.2.5

Is the DAV gap independent of the number of slices? (GM)

Sect. 2.2.9

Do I understand correctly that ECR is required to recover from errors which don't necessarily involve mismatch of BC/L1 numbers, e.g. G-link problems? Is this desirable? What if the expected high rate of ECR turns out not to be needed in reality? (ML)

Sect. 2.3.3

It's not clear whether the S-Link headers and trailers (including the ATLAS event format header and trailer) are included in the rate calculations or not. In any case the start and end control words don't seem to be taken into account. (SH)

Sect. 2.3.3

Why 5%? I thought ATLAS will run with a 10% occupancy rate. We can easily check the figures of the table, at least for the CPM DAQ, by using several CPMs and ramping up the occupancy rate. Test vectors exist with the occupancy rate passed as a parameter. We can also check the rate between compressed and uncompressed. (GM)

Table 6

When you say "Unused", what happens if a board is connected? The S-link packet could still be correctly assembled? (GM)

Sect. 3.4.4

Why is the event number buffer handled differently than the BCID and trigger type buffers (buffer is read when the ROD receives the DAV* signal, not when the first input data block is available)? (SH)

Sect. 3.6.2

What's the status of the alternative proposal from Robert McLaren for ROD module IDs? (ML)

Sect. 3.8, bottom para

Have you tried to read more than 2 RODs in the same crate? I was not able at B'ham to read a ROD positioned in a different slot than the one at the address 0x2c0000. Could be a VME mapping problem? (GM)

Sect. 3.12

I found that I can't download the JEM DAQ F/W using the VME register. Any conflict possible between downloading the FPGA via VME and using the crate/slot address? (GM)

With System ACE the only thing you have to be careful with is to write physically on the CF card what the F/W version on it is, and for which board type. (GM) We need some way of identifying them to avoid confusion between JEM, CMM and ROD CF cards. TC labels!?! (ML)

The JTAG connector used by the ROD is different for the connector used by the JEM, also different from the connector used by the CPM...maybe we should specify which cable to use. (GM)

Is there any long-term plan to use the planned facility to be able to update them in place? This is done for CMS FEDs and the CMS code has recently been taken and extended by the ATLAS SCT group. (ML)

There seem to be 8 different firmware versions for the ROD, is the difference only in the input FPGA, or are other parts affected as well. Is there a scheme to handle common interfaces and versioning? (SH)

Sect. 3.5.1

This is not have a very detailed specification for the RTM. It needs more design details. (EE)

Sect. 4.6

Can't find a register to switch the CAN from Burn mode to Run mode. Could be handy once you've got several boards in the crate and can't access the jumper. (GM)

Sect. 4.6.6

F/W bug report: the bits ACE done and ACE Error are swapped. (GM)

Sect. 4.7.10

Can the mode select be selected knowing the F/W type loaded rather than by using the S/W? (GM)

Sect. 4.7.12, bits 0-3 description

Is the maximum value still meant to be 5? How does this fit with the current PPM numbering scheme? (SJH)

Sect. 4.8.6, ECR count register

The 8-bit ECR counter is readable and can be set to 0 (page 72). It is necessary that it can be set through VME at any value for fixing synchronization problems. Is this possible? (PF)

Sect. 5.1, item 4

I don't think this is a true description of the current implementation, though it may be desirable. As I understand it, all firmware variants currently ignore the control register. (SJH)

Sect. 7.3, item 1

Quantity needed: my table has 3 spares, you say 4 (i.e. total of 27 not 26) - we should decide. (EE)

Appendix A

It seems that the ROD can supply 3.3V as well as 5V to the S-LINK transition module, this is somewhat dangerous, since the the 3.3V HOLA cards will be destroyed if they are supplied with 5V (by setting the jumper on the transition module wrong). Since the 5V S-LINK cards are now essentially obsolete it might be safer to remove this supply voltage. (SH)

APPENDIX B: DETAILED COMMENTS ON ROD DATA FORMATS

First three comments on PPM Data Compression notes, others on ROD specification.

Sect. 2, para 2

"Pending resolution of a scheme ...". Isn't this solved by the L1Calo-specific header in the S-link payload? This is the very reason why that header was introduced. If it's inadequate, then we should rethink. (SJH)

What about compression algorithms for other than 5 FADC + 1 LUT? Can use uncompressed for fewer FADC slices? Also, 7 + 1 has been mentioned at least for diagnostics. (ML)

I would not be so much afraid about data compression because we'll start without any I guess and then we limit the L1 rate to 75 kHz not 100 kHz. When it arrives we'll know more about the system and very likely we can relax on the number of slices to be read-out. (PF)

Sect. 2.2.9, final para

Though I'm not advocating sorting, I wonder if it would make sense to somehow prioritize CMM RoIs so that they don't get left off if the Max RoI count is reached. Somehow losing the missing ET RoI because there were lots of JEM RoIs seems wrong, it would be good if we could guarantee to deliver the ET RoI. On the other hand, this is probably academic, since any event above the max RoI count should probably just be automatically passed by level-2 anyway. (SJH)

Sect. 2.5, final para

As far as I know, the current implementation of the native formats either does zero suppression or doesn't, there's no switching. Can we confirm this? If so, then this requirement ("it must be possible") is not being met. Is this a problem, or are we happy with always zero suppressing (for JEMs, CPMs etc)? (SJH)

Sect. 3.6.1

Are you sure you can use the minor event format version to identify the firmware release? The ROS s/w may expect the version number to be 0x03000000 as described in the event format document. (SH)

We have at times discussed how to identify changes in RoI format versions. Is this the only place to do it? Have we finally decided not to include a user data header word in our RoI fragments? (ML)

Sect. 3.6.2

The module-ID for the DAQ and RoI output of an RoI ROD will be identical (and will not follow this scheme) if the two slink outputs are a pure copy of each other. If this is the case you may have complaints from the HLT as they use the module-ID as a unique identifier (we have got this problem in the CTP, which also uses the same ID to send data to the ROS and the RoIB). (PF)

Sect. 4.7.14, register description

As I understand it, all firmware variants currently ignore this register (to be checked). Maybe it will be used in the future (e.g. firmware variants that switch on/off zero suppression, PPM compression with threshold) but if not, then this register, or some bits of it, are redundant. (SJH)

Figure 13, Ns2 and Ns1 fields

Are these still split 4/4 or rather 5/3 for PPM purposes? Also note order of LUT and FADC is inconsistent with description in table 15. (SJH)

Table 15, note on module number

What (slot) number is it? Physical position in crate or relative position of a group of boards of same type? I thought the latest, i.e. from 1 to 14 in case of the CPM. (GM)

Table 15, note on seqno

Number of first ASIC channel only for PPM formatted (redundant?), PPM compression uses it for variant version. (SJH)

Sect. 5.4.1, final table

The CD error field is described as a whole MCM disabled. So how can we tell if one channel is disabled, rather than a whole MCM. This is a far more common situation in the forward region. The internal PPM format knows about individual channels, so it seems useful to me for this to get through the G-link at least, even if the ROD can't cope with it. Maybe a future ROD PPM firmware variant could do something about it. (This is really a PPM format question, but the ROD spec is the only place where the current format is documented, and this question is potentially critical to ROD compression algorithms.) (SJH)

Figure 19, sub-block header

According to Dave's compression note, the ASIC number is no longer stored, instead it's the firmware variant number (SJH)

Figure 31, Energy RoI type words

The same comment applies to same diagram in L1/L2 doc. I think these are out of date with respect to the ET RoI (and others in consequence). I remember that we agreed (ECR 04) that the ET RoI type should be 0110 to avoid conflict with header/status words. The other two (Ex and Ey) should be extended by a bit to 0100 and 0110, and the parity bits shifted down by one bit. (SJH)

Sect. 5.15, final sentence

Internal Data Transfer Error is always set to zero, but from cross-reading with the L1/L2 interface document, it would appear to me that an appropriate use of this would be an OR of the three error bits in the private half of status word one. The note on IDTE there says "can be further explained in the private section - data may be incorrect". This seems to me a perfect general description of the various link/parity errors in the private section. (SJH)

Figure 33 (and accompanying table)

The description of the private section of BOTH status words does not agree with the L1/L2 doc. I would guess that this one is more up to date, but all bits defined in both documents should be considered for inclusion. The ROD spec has 2 that the L1/L2 doc doesn't, and vice-versa. Even the 2 that they agree on have different bit assignments. (SJH)

APPENDIX C: DETAILED COMMENTS ON TESTS DONE, AND ROD TEST PLAN

In this Appendix we first list in Appendix C.1 comments and questions related to the testing done so far, as described in the Operational Test Report. Appendix C.2 contains comments and suggestions for the production testing, as described in the Test Plan for the ROD production.

C.1: Hardware tests and performance

1. Please describe the situation regarding FPGA balls that have apparently become disconnected - is this a general problem we should be worrying about, is there a simple fix in manufacturing process, etc. (EE)
2. Are all known hardware-related faults understood and resolved? Which parts of the hardware are not (or not intensively enough) tested? (KM) It is not clear to me from the test report that everything that could be sensitive to the hardware has been checked though. (ML)
3. No test yet with all 18 input G-links active? This would presumably be the worst test for power consumption. E.g. could power supplies be just at their limit with the demonstrated 16-input test? (NB: Operation Test Report p10 refers to the "ROD's 17 input[s]", not 18!). (ML) Have you got any loss of synchronization problems with links running simultaneously? Do you know the admissible jitter on the clock? (PF)
4. Have all 18 inputs have been tested separately with full path through the module (i.e. all tracks) being shown to be functioning correctly (i.e. more than just JTAGED). (Am I correct in assuming that the 16= input Bham test was not actually collecting and checking data, but was just for stability?) (ML)
5. Are there any estimates (even rough) of bit-error rate limits on the links? What is the longest period that the ROD has been tested/running in a demanding setup, with many links active? (KM)
6. Have you measured the current consumption of the board, in various modes of operation? Have you used LabView and checked the current consumption of the FPGAs and the temperature of the G-links when all active? Anything to worry about? (GM, KM)
7. If, as stated, not all crate slots have been tested, may it still be possible that some geographical address lines are not working? (ML)
8. Has any test been made of the AUX output? (ML)
9. Have the logic analyser pins been tested? (GM)
10. Is the logic to drive Busy tested in all respects. There are many settings with correlations that need to be fine tuned for smooth operation. Again, all thinkable inconsistent settings should be tested. (KM)
11. Has the ROD CAN interface been tested? Has the CAN controller seen F/W and has it talked to the TCM in the crate? Any results on the signal quality across backplane? Needs checking before production. (KM)
12. Connection between the monitoring FPGA and the switch FPGA, and from monitoring to the outside world? Any test on this yet? Is it planned to be used at all? (KM)
13. Any block transfer or interrupt mechanisms which might be used and but not touched upon yet? (KM)
14. Has playback mode been (implemented and) tested? Do we need to test this before production? I think the ATLAS ROD requirements demand support for playback to test ROD-ROS connections. Are there any control lines on the board which are only used for setting up and/or running the as-yet untested playback mode? E.g. VME or TTC broadcast commands? (ML, KM)

15. Should the checkpoint handling be tested, and if so how (need to send the correct trigger type, see ROD specification sect. 3.4.10). (SJH) But note: Although it is still requested to be implemented, the check point is more or less abandoned. (PF)
16. Has the ECR behaviour been tested with long G-link readout? (GM)
17. Have all three switch modes been tested? With how many inputs, how many outputs and for how long? Checks for independence of the four S-link outputs (flow control, IDs, etc). Errors in one, e.g. long time XOFF really don't affect the others? (ML)
18. Has the timeout logic been tested in hardware by, as suggested in the ROD specification (sect. 2.2.4), unplugging a G-link? Does the ROD recover properly? If not, it should be done. (SJH, GM) Different timeout scenarios both on the input (G-Link) and on the output (S-Link) need be tested by simulating faults on links. It would be of importance to see what happens to the corrupted event and to events following it, i.e. to test recovery from faults. (KM)
19. Have any of the following tests of errors and failure modes been done? (ML):
 - Temporarily missing DAV/link on one G-link: error bits, recovery after link comes back up, no garbled data during resynchronisation process
 - XOFF if S-link cable fails (or is removed, or disabled at ROS)
 - Treatment if incorrect BCID from one, many or all inputs: error-bit set, recovery, diagnostic info in registers
 - Correct handling of ECR, BCR
 - Timeout/status bit if trigger-type not received
 - Checking of internal LIID (incremented by L1A) with LIID from TTCrx (I think Thorsten had some scheme to get RODs to make such checks in case of faulty TTCvi or fibres?)
20. General firmware robustness: no need for module reset or power cycle to recover errors (e.g. as was needed on 6U ROD) (ML)
21. At what rates have the different modules that are read out by the ROD been tested? Which parts are not tested? (On CMM for instance). (KM)
22. Firmware combinations: check that jet-CMM firmware in FPGA 5 can have two different algorithms for the two channels? (ML)
23. How much realistic testing of interfaces is needed before we're confident to produce all modules? Beyond the RoI Builder, and standalone input module and output to ROS tests, should more be done. I didn't feel entirely happy about the 6U ROD until after the test-beam. Can we do something similar with a few RODs in a semi-realistic L1Calo slice before going into major production. It needn't be quite as ambitious as the test-beam, but running with detectors and checking data synchronisation and read-out stability would be good. Ideally 5 RODs, one of each type (PPM, CPM DAQ and RoI, JEP DAQ and RoI) would be involved, each fed by at least one of the appropriate modules. Of course, the desirability of such a test should be weighed against the needs of the production schedule. (SJH)
24. It is necessary to perform integrated tests with for instance calorimeters with (a close to) final setup (final components) to make sure that event alignment/synchronization at ROD level is performed properly. (KM)

C.2: Production Test Plan draft 2.1

1. Seems trivial, but during the test you should also check that the LEDs on the front panel are working correctly. (I recall the TTCrx LED didn't work in previous RODs when I used them at B'ham.) (GM)
2. How do you check that TCM and ROD communicate ? Do you need special F/W in the TCM? (GM)
3. The test-plan document doesn't mention tests of the AUX or BUSY outputs as part of production tests. How do you test them? (ML, SH)
4. Are there any plans for a burn-in test on the VME communication? (KM)
5. In production tests, would it be worth trying to use the S-link diagnosis modules as well as a DSS to verify the correct performance of the S-link outputs? (SJH)
6. The production test plan does not address testing of the S-link rear-transition module (EE, GM)
7. Nothing is said about ROD testing beyond the basic standalone production tests (EE)

APPENDIX D: EDITORIAL CORRECTIONS TO DOCUMENTATION

In this Appendix we list detailed editorial comments and corrections for the “formal” documents, i.e. those which will continue to be updated and referred to beyond the review. Items that have an impact on the ROD specification and work are not listed here; these are purely corrections and typographical errors in the documentation. The documents covered are the **ROD Specification** (Appendix D.1), the **Test-Plan document** (Appendix D.2), and the **Level-1/Level-2 Interface Document** (Appendix D.3).

D.1: ROD Specification version 1.08aGeneral

Globally edit "Atlas" to be "ATLAS". (EE)

General comments about TABLES:

Many of the tables use more lines than needed in some of the rows. They could be made to look much more attractive, and take up less space, by slightly widening the affected columns.

If paragraph format in tables is left justified rather than fully justified, can avoid the ugly "type spreading out" syndrome seen in some entries.

Tables should be set up so they do not break across pages, unless of course they are more than a full page long. I've noted below two that break unnecessarily, but as pagination shifts others might do so too. (EE)

Sect. 1, para 2

It's unfortunate that references occur in the order 1, 9, 5, 6, 7, 8, and then (in sect. 1.1) go back to 2. Can they be rearranged without too much trouble? (Also see comment on duplicated refs. below.) (EE)

Sect. 1.3, para 2, lines 5 and 7

"modules ... sample up to five consecutive timeslices". But in fact ALL modules can read out MORE than five slices - there is no hard limit! Reword to something like: "modules ... can sample a programmable number of consecutive timeslices. Although there is no hard limit, we normally plan on a maximum of five timeslices for all modules except the PPM". (EE)

Fig. 4

Label says "S-Link Swtch"; should be "... Switch" (missing "i"). (EE)

Fig. 5

Switch FPGA is now XC2VP30, not XC2VP20. (EE)

Sect. 2.2, line 1

"There are six TYPES OF sources of ..." (EE)

Table 2

Types 7 and 8 would be better called "CPM RoI" and "JEM RoI", to show the types of modules they come from, and consistency with other entries in the table. (Also see comment on Sect. 5.11.) (EE)

The column header "source" is mislabelled - it's actually the field value. (IPB)

Table 3

Types "CP RoI" and "JET RoI" are better called "CPM RoI" and "JEM RoI", to show the types of modules they come from, and consistency with other entries in table. (Also see comment on Sect. 5.11.) (EE)

Dave's note suggests the PPM compressed value here should be 1 and not 3, check and correct if so. It also makes the table almost redundant if true, as only the neutral mode differs from 1. (SJH)

Sect. 2.2.1, first three paras

There's a lot of repetition here, at various levels of detail (looks like paragraphs written at different times with little reference to others). It therefore reads strangely. Could be organised more logically. (SJH)

Sect. 2.2.1, para 3, line 7

"... Bunch-Crossing number (BCN) present ...", i.e. define the abbreviation, since "BCN" is used a few lines later but has not been defined. Alternately, in line 12, replace "BCN" with "Bunch-Crossing number". (EE)

Sect. 2.2.1, para 3, line 10

"sets" not "set" (EE)

Sect. 2.2.1, para 4, line 2

Replace "BC number" by "BCN" or "Bunch-Crossing number", depending on what you did about the last comment. (EE)

Sect. 2.2.1, last para, line 3

"G-Link numbers" is a bit ambiguous, could refer to how many - it might be better to say "G-Link assignments" or similar. (EE)

Sect. 2.2.2, para 1, line 4 AND para 2, line 1

As above, "up to five consecutive timeslices" is not accurate, reword as "a variable number of consecutive timeslices" (EE)

Sect. 2.2.9, para 3, penultimate line

Insert a comma after "Event Type" (?) (EE)

Sect. 2.2.10, para 1, line 4

"indicates" not "indicate" (EE)

Sect. 2.3.1

The information about the PPM configuration is repeated, and in general reads badly. Perhaps the detailed comment justifying the chosen PPM slice configuration should come after explaining the slice numbers for other modules. (SJH)

Sect. 2.3.2

Packet length of 309. I think this figure is historical, it should be different (smaller) now. The argument still stands however of course. (SJH)

Sect. 2.3.2, final para

As shown from table 3, 5 ticks is not really an appropriate gap to talk about nowadays - recalculate for 1 tick, true for all native modes. (SJH)

Table 4

The CP CMM frame length was "corrected" to 27 from 35, but it should still be 35, it's just that many bits aren't used. The same is true for other variants (bits not being used that is), but you haven't given a modified frame value to suit them. (SJH)

Sect. 2.3.3, para 2

Unfortunate that you've chosen 1/5 for PPM readout here and 1/3 for the previous rates. While I know why you've chosen 1/5 here, wouldn't it be better to go back and do the previous table with 1/5 too? (SJH)

Sect. 2.3.3, final para

While this information is separated from the rest of the section by a table, there's no disguising the fact that its content is pretty similar to the previous sentence. Are both needed? (SJH)

Table 5

Please add some explanation how you arrive at the CP and JEP uncompressed rates. (EE)

S-Links/ROD for RoIs shown as 1, but should indicate in TABLE that identical data are sent down TWO S-Links, one to RoIB and one to DAQ. (EE)

Table 6

Should indicate in TABLE that CMM RoIs are from JEP system-summing crate only, not the other crate. (It's said in the text, but someone looking at the table for reference might well not see that.) (EE)

Fig. 10

Solid line from input FPGA 5 should only join that from FPGA 4, not 3? (ML)

Sect. 2.4, final sentence

Slink port 3 should be 2 now maybe? (numbered 0-3) (SJH)

Sect. 2.4.1, para 1, final sentence

Ditto above comment. (SJH)

Tables 6 and 7

G-link numbers should be 0-17 nowadays? (SJH)

Sect. 2.4.2, first para

Twice the numbers 1 and 3 are mentioned. These should be 0 and 2 nowadays? (SJH)

Sect. 2.4.2, final para

Slot numbers for PPMs are wrong. Should be 6-12, 14-20 (SJH)

Sect. 2.5.2, last two bullets

This description is rather historical, probably needs updating in the light of Dave's recent work. (SJH)

Sect. 3.1, line 2

"via optical links" (plural). (EE)

Sect. 3.1, line 4

"such as" is probably redundant now - either they are or they aren't. (SJH)

Sect. 3.2, para 1, line 3

The JEM does not use the HDMP1024, but instead the compatible 16-bit HDMP1032. So delete "JEM", and add a short statement such as "The JEM uses the compatible HDMP1032 (Tx)." (Is it worth pointing out that the HDMP1024 can do 16 or 20 bits, and we use both modes?) (EE)

Sect. 3.3, first sentence

Reference for 6U ROD should be [17] not [19]. (SJH)

Sect. 3.3, lines 5-6

I believe these FIFO buffers are actually 8k and not 11k deep. (EE)

Sect. 3.4, para 3, line 1

Again, aren't they really 8k not 11k deep? (EE)

Sect. 3.4, last para

Switch FPGA is now XC2VP30, not XC2VP20. (EE) Also note: What has driven the choice of that FPGA? The amount of memories more than the speed or the number of I/Os? (GM)

Sect. 3.4.1

Globally edit "event type" to be "trigger type" in this section; there is no such thing as "event type". (EE)

Sect. 3.5.1, fig. 11, and accompanying text

This picture is the old CERN transition module. Why don't we have a picture of our own? (SJH)

Sect. 3.5.1, para 2, last sentence

This implies all outputs have headers, whereas I think only one has this capability? (SJH)

Sect. 3.6, line 4

Bad reference - should be [21] not [6] (SJH)

Table 8

Minor format version number is non-zero for 9U ROD (text ok but table should not have "= 0x0") (ML)

Sect. 3.6.6, para 1, line 4

"The user header is format is shown ..." (EE)

Sect. 3.8, bullet 6

"TCM-VME" not "TCMA". (EE)

Table 10

It would be very helpful to indicate (e.g. in the caption) that the other ROD crate is the same, with only the crate number different. Otherwise it might not be at all obvious; the text doesn't say it. (EE)

Sect. 3.11

"reloading" should be "reload"? (SJH)

Sect. 3.14

There is very little information on the PCB (e.g. number of layers, etc.) (SH)

Sect. 3.14.3

n=1-18 should be n=0-17? (4 instances of this) (SJH)

And S-links 1-4 should be 0-3? (EE)

Sect. 3.14.7

Notes 6-8 should be notes 1-3 (SJH)

Sect. 4.6.3, bits 6 and 7 descriptions

Identical descriptions - presumably there must be some difference. (SJH)

Sect. 4.6.5 bit 0

"Writing a 1 to this bit sets ...". Add the missing information. (EE)

Sect. 4.6.5, bits 4-11 description

Incorrectly sub-titled bit 0. (SJH)

Sect. 4.6.11, register description

Doesn't seem to be any description of CAN register A, does it do anything? (SJH)

Sect. 4.7.2, register description

Wrong description of function, repeated from Sect. 4.7.1. Also bits 2, 3 sub-titled bit 0 and 0 again. (SJH)

Sect. 4.7.6, bit 0-12 description

Maximum value should be 0x1FFF, or maybe it's just bits 0-11? (SJH)

Sect. 4.7.7, bit 0-3 description

Maximum value should be 0xF, or maybe it's bits 0-8? (SJH)

Sect. 4.7.8, bit 0-8 description

Maximum value should be 0x1FF, or maybe it's bits 0-3? Also description says event FIFO, should be BCN FIFO? (SJH)

Sect. 4.7.18, register description

"bit fields are input to card" - should be output? (SJH)

Sect. 4.8, register map

Is register at 0x0018 still used for anything? Source IDs are now done on Slink basis (0x1004 for example), so this should be removed to stop confusion. Also there are 7(!) registers named "s-link n LFF Time", 0x40-0x48 and 0x01000, 0x11000 etc. Are these different, or are some redundant? (SJH)

Sect. 4.8, line 1

"during the design phase" is now out of date, so reword. (EE)

Sect. 4.8.7, bits 16-23 description

Format version is now 3.0 not 2.4. (SJH)

Sect. 4.8.14

Table should not break across pages. (EE)

Sect. 4.8.18, bits 0-5 description

Should be bits 0-15. Also bits 24-31 description mislabelled as 0-1. (SJH)

Sect. 4.8.25, register description, final part

"but are actively used...". Should this read "but are NOT actively used..."? Also bits 28-31 are now defined to be 0xF, not zero I think. (SJH)

Sect. 4.9, line 1

"during the design phase" is now out of date, so reword. (EE)

Table 13, subdetector IDs

CP RoI should be 0x73, JEP DAQ should be 0x74, JEP RoI should be 0x75. (SJH)

Sect. 5.2, table 15

Table should not break across pages. (EE)

Figure 15 and nearby text

This figure is not easily comprehensible:

(a) "BCID result" should be labelled "LUT output" - BCID has not been mentioned at all (and in a ROD context might be understood to be the DAQ definition rather than our digital filter), and the text (as well as other nearby figures and tables) very consistently calls it "LUT output", so use that name.

(b) It is not at all obvious why the 8-bit LUT output and the 10-bit FADC data are all shown as 11 bits. Please say in the text that the extra bits are flags related to the bunch-crossing identification. In the figure it would be best to show what the bits are, just like in fig. 16; if not the reader should be referred to the penultimate line of fig. 16. (EE)

Sect. 5.4.1 line above unnumbered table at bottom of p.87

"are AS follows" (EE)

Sect. 5.4.1, first table info

PB is repeated twice, one of them should be SB. (SJH)

Sect. 5.4.1, second table info

Is missing a description of AFF. (SJH)

Sect. 5.4.2, para 2, last 2 lines

"to accommodate 3 FADC and 1 LUT timeslices giving reduced data volume", i.e. omit "a", and "reduced" not "reduce". (EE)

Sect. 5.4.3, para 2, line 3

"five" not "Five". (EE)

Sect. 5.4.4, last line

the correct reference is [28] not [27] (EE)

Sect. 5.5.1, para 3, last line

Replace "... at present in the interests of system stability." by "... since there is no need for 5-slice CPM readout at full rate." (EE)

Sect. 5.11 section title

Should be "CPM ..." not "CP ...", both for consistency with other titles (e.g. Sect. 5.12) and with Tables 2 and 3 (see comments above). (EE)

Sect. 6.2.1

Probably out of date, why not just refer to test plan(s) (SJH)

Sect. 7.3

Item numbers 2 to 4 are bold, but should not be. (EE)

Sect. 7.3.1

Very unclear! Why only one crate (which is a standard item)? And what kind of software is meant? This section should either be expanded/clarified or just removed. (EE)

Sect. 7.5

first para is not accurate. "Customer" does attend FDR (which is combined with the PRR in this case). (EE)

Sect. 7.6, line 1

Insert a comma after "If required" (EE)

Sect. 7.7

Insert a comma after "design tools" (EE)

Sect. 7.9, line 2

"Deliverables" ref. is sect. 7.3, not 4.2. (EE)

Sect. 8, ref. 5

PreProcessor specification is currently being rewritten, so the comment about it being "(under revision)" should remain. But please give the web address anyway so the reader can find what information there is:
<<http://hepwww.rl.ac.uk/Atlas-L1/Modules/Modules.html#PPM>>

Note that using this URL allows both old and new (when ready) versions to be found. (EE)

Sect. 8

There appear to be a couple of repeated references: [10] and [26]; [18] and [27]

Also the way EDMS references are presented is a bit inconsistent. I prefer the way it's done in references 26-28 (just the document number, no versions, specific names). (SJH)

D.2: Test Plan DocumentSect 3.1

"make sure links..." (ML)

Sect 6.3

"ROD LEDs" (not "CMM LEDs"!) (ML)

Sect 6.4-5

Are the LED names correct for the ROD? I don't see them listed on page 41 of the ROD spec. (ML)

D.3: Level-1/Level-2 Interface Document version 3.1Sect. 3.3, final sentence

"minor version numbers ... still to be defined". This is not quite true for L1Calo - in the ROD spec it already says that 6U ROD outputs should have this as 0, and 9U ROD are non-0, starting at 1 right now. (SJH)

Also: We have at times discussed how to identify changes in RoI format versions. Should the L1/L2 document include something about minor format versions in our ROD fragment headers? Have we finally decided not to include a user data header word in our RoI fragments? (ML)

Sect. 6.2.1, after description of 2 crate bits

"Each crate occupies a region of ± 2.6 ". A misleading number, it's either 2.8 (for outer CPM module boundary) or 2.5 (for actual region processed). (SJH)

Sect. 6.2.2.1, near end

"eta = -3.2 to -2.6 ". Both figures wrong: 3.2 understandable but misleading - physical eta range is 4.9 due to the irregularities at end; -3.2 would be true if range were normal. 2.6 is incorrect, should be 2.4. (SJH)

Sect. 6.2.2.1, para after last correction

This description of the usage of these 5 bits, really split, as shown in the figure, into frame and local coordinate, is almost certainly too simplistic. The frame locations etc are described in the JEM spec, and almost certainly don't conform to this description. (SJH)

Sect. 6.2.2.1, final sentence

"with the outer loop over crate number". Crate loops are not mentioned in the CPM case, and probably sensibly so, since the order will depend on the order of the slinks into the RoIB and possibly dependant on re-ordering within the RoIB. However this would be a desirable order, if we can arrange it correctly. (SJH)

Figure 8 AND Sect. 6.2.2.2

I think these are out of date with respect to the ET RoI (and others in consequence). Note the figure is also incorrect (the same figure) in the ROD spec. I think that we agreed that the ET RoI type should be 0110 to avoid conflict with header/status words. The other two (Ex and Ey) should be extended by a bit to 0100 and 0110, and the parity bits shifted down by one bit. (SJH)

Sect. 6.2.2.3, final bullet

Misplaced, but more importantly incorrect. 0010 should be 101. (SJH)

Sect. 6.2.5, para 2

This description of the private section of the first status word does not match the description in the ROD spec. Both the order and number are different. This description is missing the CMM parity error bit. (SJH)

Sect. 6.2.5, para 3

This description of the private section of the second status word does not match the description in the ROD spec. Both the content and number are different. This one has saturation and overflow bits (not sure how you separate saturation from overflow?) and the ROD spec just has Trigger Type timeout. (SJH)

Figure 9, both private fields

Needs updating to match ROD spec (or vice-versa) as noted above. (SJH)

Sect. 6.2.6.5

The new trigger configuration DB adds a global energy scale parameter to be loaded into the PPMs to set the conversion of FADC counts to GeV. Expected range from 1 to 4. (ML)

Sect. 6.2.7, first bullet point

"... are zeroed for two clock cycles". This is too simplistic. Either don't go into details, or say something like "any data that could have been corrupted by the error are zeroed". (SJH)