Dataflow Integration of the Calorimeter Trigger CPROD with the RoI Builder and the ROS.

10th September 2001

B. M. Barnett, J. Dawson, Y. Ermoline,
C. N. P. Gee, M.P.J. Landon, J. Schlereth

ATL-DAQ-2001-008
03/12/2001

Abstract

A first dataflow integration test was performed in April 2001 between the Level-1 Calorimeter Trigger Cluster Processor ROD (CPROD), the RoI Builder (RoIB), and the Readout Sub-System (ROS). The aim of the test was to send RoI information from the CPROD to the RoIB and ROS systems, and to understand the hardware interface between the systems concentrating on data performance and data integrity. The systems were controlled and monitored by separate diagnostic programs, and software integration within the ATLAS online framework was not included in the goals of these tests.

1 Introduction

The principal objectives of this integration test were:

- **verification of CPROD functions with the DSS module;**
- **transfer of individual events to RoIB and check content;**
- **running up to the maximum Level-1 Accept frequency limit (100 kHz);**
- **performance a long run to check for infrequent errors;**
- **measurement of system latency from L1A to RoIB;**
- **transfer of individual events into the PC ROS and check content; and**
- **completion of a combined ROS/RoIB Run**

The primary purpose of the tests was to check that data were transferred completely and correctly from the CPROD to the RoIB, with all data words checked. A range of different event fragments was generated, both to explore the phase space, and to ensure that complete fragments were not lost, double-counted, or reordered out of sequence.

2 Test Setup

The main components of the test are illustrated in Figure 1. They included the CPROD, RoIB, and ROS subsystems, TTC modules for timing and trigger signal generation, NIM logic, and instrumentation including a fast pulse generator. Other instrumentation, including a fast digital oscilloscope, is not shown in the diagram.

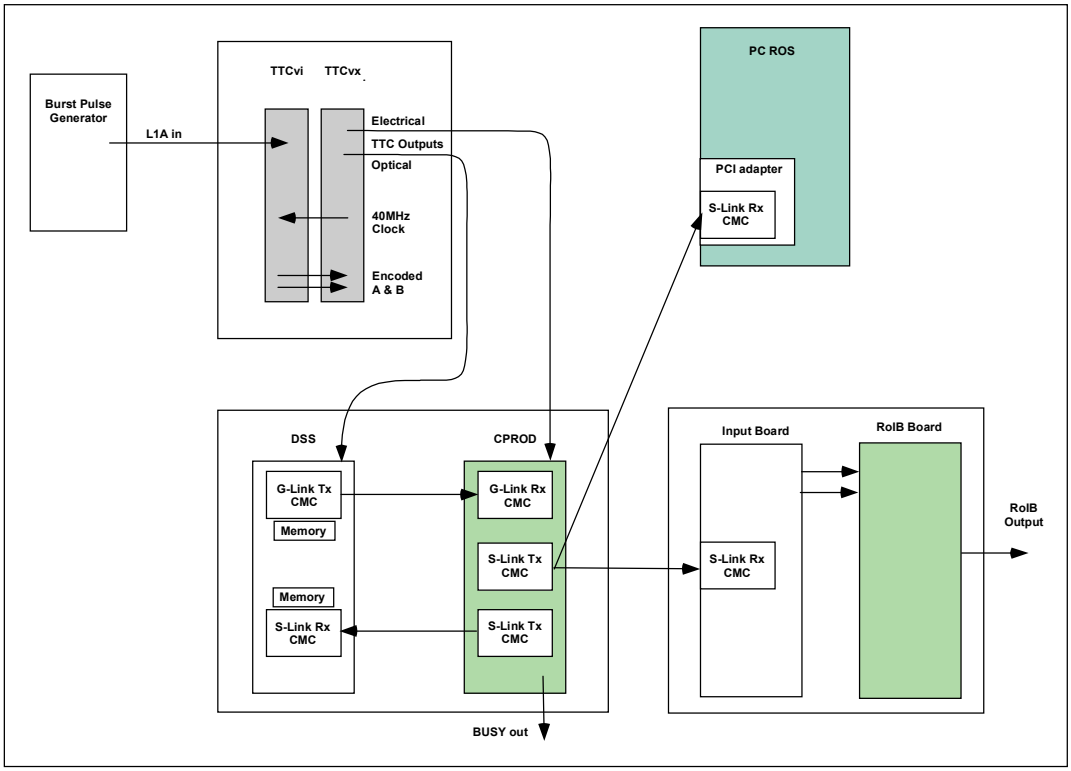

Figure 1: Schematic view of the test installation. All combinations of two of the three possible S-Link connections between RoIB, DSS and PC ROS were used at some point in the tests.

2.1 CPROD Sub-system

The CPROD subsystem consisted of a Data Source Sink (DSS) module and a CPROD module, housed in a 6U VME crate controlled by a Concurrent PSE/P34 single-board computer running the Linux (RedHat 6.1) operating system.

2.1.1 DSS Module

The DSS module [1] is an FPGA-based test module providing two CMC sites, each associated with VME-accessible memories. In these tests, the first CMC site was fitted with a plug-in transmitter board containing four H-P G-Links [2] . The DSS was also equipped with an interface to the Trigger Timing and Control (TTC) [4] subsystem.

The DSS module firmware was configured to generate packets of serial data on each of the G-Link links following each Level-1 Accept (L1A) signal received via the TTC system. The serial packets were formatted to mimic the RoI output expected from four Cluster Processing Modules in the production trigger system. The packet content was obtained from an internal DSS memory, which was preloaded with data corresponding to a range of different RoI contents at different stages during the tests.

The second CMC site on the DSS module held a plug-in S-Link [3] receiver module. Data arriving over the S-Link were captured in a 32bit x 32k deep memory and could later be read from VME.

2.1.2 CPROD Module

The CPROD module [5] is a prototype, performing all the functionality of the planned production module, but with four G-Link input channels rather than the 18 needed in the production trigger. In these tests, the CPROD received RoI packets from DSS over the G-Link links after each L1A. This data was converted to 32-bit parallel words, one representing each RoI datum, and words with none of the 16 threshold bits set were discarded. Surviving words were written to FIFOs, from where they were collected to build an ATLAS-standard S-Link event fragment [6] . The CPROD offers two CMC sites for S-Link output modules, and can send the RoI fragments over either of the S-Links. It can also send duplicate copies of the same fragments in parallel over both S-Links. Logic is provided to send data at the rate of the slower link (i.e. the link full flags are ORed), but can be disabled independently for either link if required for diagnostics.

The CPROD also has an interface to the TTC system, enabling it to obtain the common 40 MHz clock, L1A, L1A number and bunch-crossing number needed for the event header.

The CPROD used several FIFO memories as temporary data stores to decouple the data rates in the TTC and G-Link interfaces from that on the S-Links. The module could process incoming and outgoing events concurrently, and could store several events if these arrived at a higher instantaneous rate than the

output S-Link could accept. Due to the components used, the module could contain no more than 255 events at any instant (many more than needed during normal ATLAS running, when no more than about 8 events need be stored). A BUSY signal with a programmable threshold was available to monitor the occupied depth in the main data FIFOs.

2.2 RoIB Sub-system

The RoIB system [7] consisted of an input card and the main RoIB card. The input card received data over S-Link from the CPROD, and then prepared and sent two identical copies to the main RoIB card, which required at least two input fragments to assemble composite RoI S-Link output packets. The 12U RoIB VME crate also contained two single-board VME computers, one controlling the RoIB system and the other receiving and checking the composite RoI packets. These computers ran the LynxOs operating system.

The RoIB system included several useful testing features. These included the option to generate test data locally within the input card, and also the possibility to connect the incoming S-Links directly to the RoIB data sink computer in place of the RoIB composite output S-Link.

2.3 ROS Sub-system

The ROS system consisted of a conventional 186 MHz Pentium-based PC, with an S-Link to PCI interface. The tests focussed on used of the ODIN hardware implementation and the ROS S-Link interface library. These present the hardware and software S-Link interface to the remainder of the ROS, and their use to verify flow-control and data integrity aspects thus constitutes a valid test of link inter-operability with the ROS. Other components of the ROS were not included in these tests.

2.4 S-Links

Three different physical implementations of the S-Link were available for use in the tests. All tests with the RoIB used an electrical cable link developed at Argonne National Laboratory. The links from CPROD to DSS were the CERN electrical S-link, and the link to the ROS used the ODIN optical S-Link. It was not possible to use the ODIN links with the RoIB, as a required reset signal was not implemented. This will be included in a planned RoIB revision.

3 Results

3.1 Stand-alone Tests of the CPROD

For much of the testing, including stand-alone tests, a sequence of closely spaced L1A signals was generated by a burst-mode pulse generator and distributed by the TTC system. On each L1A, previously generated and serialised RoI data were selected and transmitted by the DSS, received, deserialised and reformatted in the CPROD, transmitted over S-Link and received into the second DSS port,

and captured in DSS memory. The number of L1As per burst was chosen so that all S-Link RoI packets generated in one burst by the CPROD could be recorded in the DSS memory without overwriting.

For each burst, the recorded S-Link data were read out by the online computer and compared with the expected fragment content. All events in the burst were different, with a cyclical pattern used to check correct CPROD processing of a variety of RoI contents. The test program could not check the bunch crossing number copied into each event from the TTCrx by the CPROD, but it did check that the event number increased monotonically from 1 through the run.

It was found that this test could routinely be run overnight without error at burst L1A frequencies beyond 800 kHz, exceeding the required 100KHz by a large factor. It should be noted that the DSS could sustain incoming S-Link data at the full speed of the S-Link, so did not normally use the S-Link flow control features.

3.2 Low-Rate Tests CPROD → RoIB

For testing with the RoIB, the second CPROD S-Link port was connected to the RoIB input card using an ANL S-Link. Reference copies of the expected S-Link RoI packets were written to files and transferred to the RoIB computer before each test, enabling the RoIB software to check all received events in full except for event number and bunch crossing number.

Single event fragments were first tested individually, and correct reception in the RoIB confirmed by manual checking.

Events were then transferred in bursts, and the L1A frequency gradually raised. Errors were detected, which on investigation proved to be related to incomplete implementation of the S-Link protocol. This is discussed further in section 3.4.

Testing was continued at lower frequencies where the protocol error did not appear. An overnight run with 124 byte events containing 16 RoIs was successfully completed without error. The events were generated in bursts of 1024 2mS apart in a 3S cycle, which averaged 418 Hz. In total, 2.1×10^{7} events, corresponding to 2.1×10^{10} bits, were sent and received without error. The 1024 events in each burst were all different in their RoI content. Events were not lost, duplicated nor corrupted in any way.

3.3 Latency Measurements

The latency of the system was measured using the above configuration while monitoring various test points using a digital oscilloscope. The test points included:

1. **the L1A signal from the output of the TTCvi module;**
2. **the L1A signal from the TTCrx chip on the DSS module;**
3. **the G-Link data available signal (DAV*, pin 70) on the DSS module. This signal was asserted during G-Link transfers;**

4. **the S-Link control bit (#LCTRL, pin 19) on the S-Link input to the RoIB input card. This signal was asserted to identify link control words and the beginning and end of each event fragment;**
5. **The BUSY signal at the CPROD front panel.**

3.3.1 Operation of CPROD Busy

Data entering the CPROD from the four serial G-Links is processed and placed into four FIFOs to await readout. A Busy threshold inside the CPROD is constantly compared (in the firmware) with the occupied depth of each FIFO, leading to assertion of the front-panel BUSY signal whenever one or more of the FIFOs is filled up to or beyond the threshold level. In normal running, BUSY would be used to inhibit further triggers if the FIFOs were at risk of overflow. However, in these tests the signal also provided a useful tool to monitor the CPROD performance.

Figure 2 shows the behaviour of the BUSY signal when the threshold was set to the artificially low value of 3. This threshold was reached as soon as 3 RoIs had entered the FIFOs, and remained asserted until the depth of the last FIFO to be emptied fell below 3 again. For events with 16 or 64 RoIs, the theoretical minimum busy times (to transfer 61-3=13 or 64-3=61 RoIs over the S-Link) are 325 ns or 1525ns. These are in good agreement with the measured times of approximately 400 and 1650 ns.

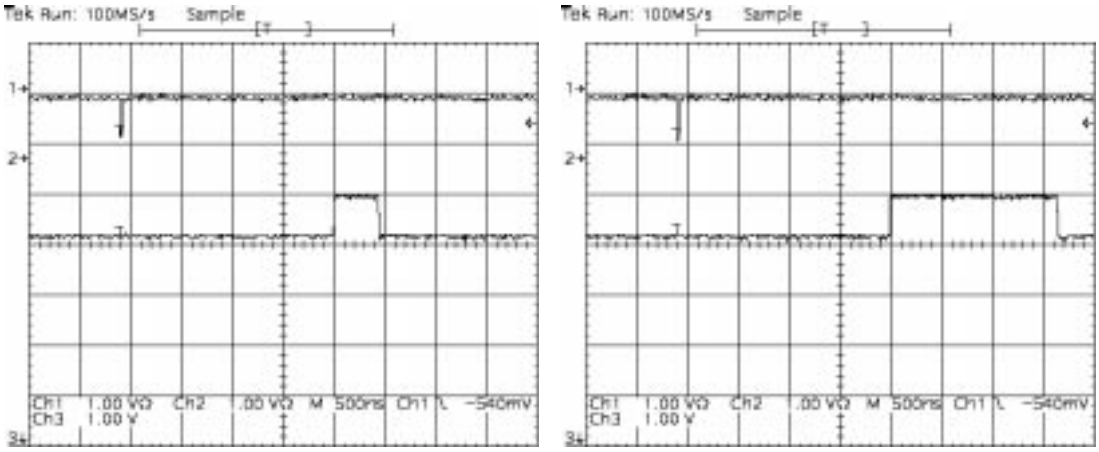

Figure 2: Timing relationship between L1A (upper trace) and BUSY signal (lower trace), for 16 RoIs (left) and 64 RoIs (Right). The busy threshold is set at 3.

3.3.2 Readout Latency

The time from assertion of L1A to the end of the S-Link event transmission includes five components:

1. **TTC signal encoding, transmission, and decoding time from TTCvi to the DSS module;**
2. **DSS time to extract data from internal memory and prepare for transmission;**
3. **G-Link transmission time from the DSS to the CPROD;**
4. **internal processing time in the CPROD;**

5. transmission time on the S-Link.

The start of this sequence is shown in Figure 3, where the measured TTC transmission time between traces 1 and 2 was 150ns. Transmission on the G-Link to the CPROD started as the DAV signal went active 440 ns later, trace 3. The G-Link remained active for 550 ns, matching the expected time for the 22-bit RoI frame to be transmitted

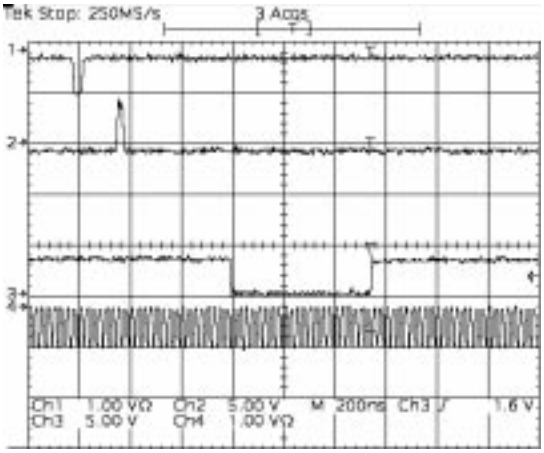

Figure 3: Timing of L1A at TTCvi (top trace) ,L1A received from TTCrx(2nd trace), and G-Link DAV signal (3rd trace). The 4th trace shows the 40MHz TTC clock. The timebase is 200 ns per division, as confirmed by the eight 40MHz clocks per major division.

The remainder of the readout sequence is illustrated in Figure 4, for events with 16 and 8 RoIs. Each event fragment includes a 9-longword header and a 6-longword trailer, so the total event lengths are 31 and 23 longwords respectively. Transmission on the S-Link started 2150 ns after the original L1A. The expected times to transfer the events on a 160 Mbyte/sec link are 775ns and 525 ns, matching the measured times between the control signals within measurement errors. As shown in the right hand plot, LFF was asserted for approximately 1700ns, slightly exceeding the 1533ns RoIB time to empty the input card memory at 15MHz.

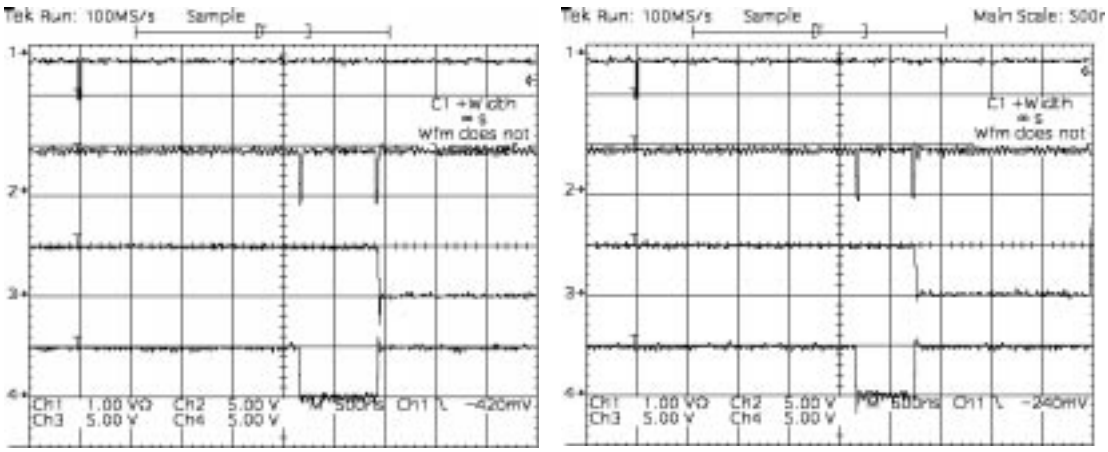

Figure 4: Readout latency in the DSS/CPROD system for 16 RoI (left) and 8 RoI (Right) events: L1A (top trace), S-Link control words (2nd trace), S-Link write enable (4th trace), and LFF (3rd trace) asserted by the RoIB when the event has been received. The timebase is 500ns/division.

The complete sequence of timing is shown in Figure 5. All the times relating to the level-1 trigger are expected to remain constant independent of the number of RoIs in the output fragment, except for the S-Link timing. which should rise by 25ns per extra RoI word .This is because the G-Link readout always carries all the information needed for the maximum number of RoIs.

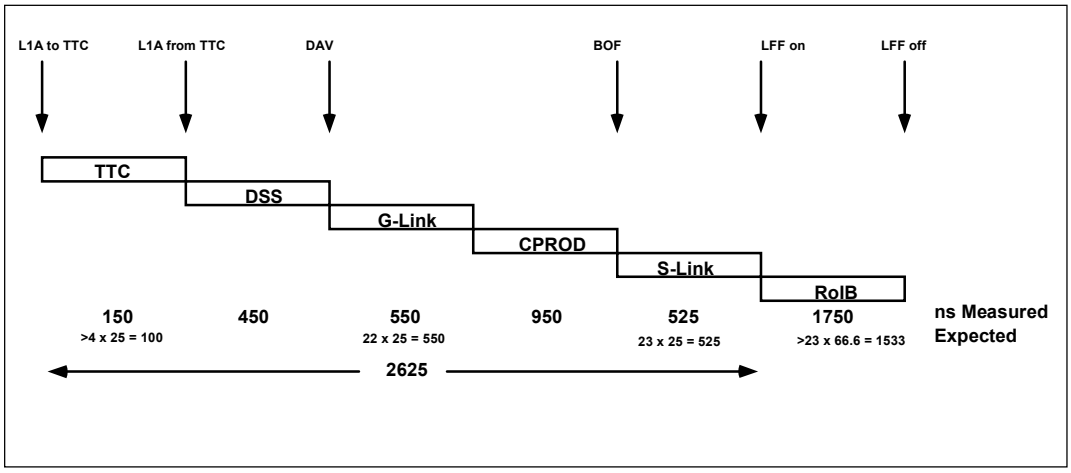

Figure 5: Overall timing from L1A to end of RoIB input card readout.

3.4 S-Link Flow Control Issues

The S-Link protocol specification requires that, when the data sink has no more room for input data, it asserts S-Link signal UXOFF# low. The Link Destination Card (LDC) then transmits an XOFF to the Link Source Card (LSC), which stops data transmission. The last data packet from LSC travels via the LDC to the data sink, where it may arrive a considerable time after UXOFF# low is asserted. The data sink *must* have room to receive this data.

The RoIB prototype input card used in the tests asserted UXOFF# immediately upon detection of the end of event marker, and ignored any further S-Link data until the event had been moved to the main RoIB board. Meanwhile, the CPROD could carry on transmitting data for several words after receiving LFF#. The combination of these two faults meant that errors occurred if events were not far enough apart. Improved CPROD and RoIB firmware will resolve this problem in future.

3.5 Tests with the ROS

For testing with the ROS, the S-Link interfaces were re-cabled such that the CPROD was interfaced to the ROS via an optical "ODIN" S-Link pair. Software in the ROS used the ROS S-Link access library to obtain events, which were then checked as in the RoIB. It was found that:

- **Event frequencies of up to 20kHz could be sustained into the ROS with full event checking. To reach this rate, the ROS-hosted readout and verification software was compiled with full optimisation. The PC was observed to be**

CPU bound;

- **Instantaneous L1A frequencies of up to 660kHz could be sustained, in bursts of 127 events.**

Very careful control of the CPROD was necessary. The DSS module needed to be reset after each burst of events, and the checking software required an exact repeating sequence of 127 events, except for the incrementing L1A number and changing bunch crossing number. It was therefore not possible to use the BUSY signal to suppress L1As if the CPROD memories became full, so it was essential to wait for the CPROD FIFOs to empty completely before triggering the next event burst. This constraint had not been fully appreciated before the test, and has implications for future development of the DSS firmware and other supporting test hardware.

The second output S-Link from the CPROD to the DSS was active during these tests, providing a powerful diagnostic capability in understanding those errors that did occur.

3.6 Combined tests with the RoIB and ROS

A series of short tests were made with data transmitted over S-Link both to the RoIB and to the ROS. This represented exactly the connectivity to be used in the production trigger, where the RoI fragments will be sent both to the RoIB and to the main DAQ readout system.

The data rates again had to be restricted so that S-Link flow control did not operate on the CPROD-RoIB link. Running with bursts of 127 different events spaced 1ms apart, several runs of about 2M events were performed, after which the first data error typically appeared. The same errors were detected by software in the RoIB and ROS. Investigation revealed a firmware problem with one of the CPROD FIFOs. This was understood, but it was not possible to obtain a new firmware version before the end of the tests.

This test nevertheless established that the CPROD could transfer S-Link data concurrently to two of its downstream neighbours.

4 Benefits of the Tests

Valuable experience of all sorts was gained during these tests. Much of it relates to specific features of the firmware versions in the modules, and will lead to improved logic. There were also new insights into the techniques needed to set up and control the event flow at high rates, and into ways of tackling combined testing in general. A few of the generally useful points are summarised here:

- **It can be hard to pin down problems unless documented test points and LEDs are provided as a diagnostic aid for key signals like L1A, LFF, and DAV.**
- **Modules using S-Link should be tested with the S-Link diagnostic cards to exercise the protocol in full.**
- **Module specifications should include statements about action to be taken if unexpected inputs are received.**

- **With suitable tools, it is possible to transfer FPGA designs between institutes and burn new EPROMs.**

5 Conclusion

The goals of the integration of CPROD, RoIB, and ROS were not achieved in full. Nevertheless, the participants agree that the test was essential to a proper understanding of the interface between the Level-1 and the Level-2 systems, and revealed problems which had not been detected in previous tests.

A low rate of RoI fragments sent from the CPROD to the RoIB at up to 400Hz could be sustained. No errors were detected in a long run cycling over 1024 different events checked in detail. The flow control on the S-Link interface did not work correctly. The latency from an L1A to the RoI fragment being sent to the RoIB input card is about 2.6 μs.

Important firmware design problems were found in the CPROD and RoIB modules, but they are relatively easy to correct with firmware changes. A joint test with the Muon trigger awaits completion of the MICTP module, and joint tests with the MUCTPI, using the TTC system, will be needed in future.

6 Acknowledgements

The authors would like to thank Georges Schuler and Ralf Spiwoks of the CERN level-1 trigger group for their friendly support in these tests and for the use of laboratory facilities.

7 References

[1] **DSS Specification, http://hepwww.rl.ac.uk/Atlas-L1/Modules/Modules.html**

[2] **G-Links, see physical layer components via http://www.semiconductor.agilent.com:80/cgi-bin/morpheus/home/home.jsp**

[3] **CERN S-Link Specification, http://his.web.cern.ch/HIS/s-link.html**

[4] **CERN TTC-VMEbus Interface (TTCvi), http://ttc.web.cern.ch/TTC/intro.html**

[5] **CPROD Specification, http://hepwww.rl.ac.uk/Atlas-L1/Modules/Modules.html**

[6] **C. Bee et al., The Event Format in the ATLAS DAQ/EF Prototype -1, ATLAS-DAQ-98-129**

[7] **R.E. Blair et al., A Prototype RoI Builder for the Second Level Trigger of ATLAS Implemented in FPGAs, ATLAS note ATL-DAQ-99-016, December 1999, and R. E. Blair et al., A Prototype RoI Builder for the Second Level Trigger of ATLAS Implemented in FPGAs, Proceedings of the Fifth Workshop on**

Electronics for LHC Experiments, Snowmass, Colorado, September 20-24, 1999, page 323