Upgrading the ATLAS Level-1 Calorimeter Trigger using topological information

This article has been downloaded from IOPscience. Please scroll down to see the full text article.

2010 JINST 5 C12046

(<http://iopscience.iop.org/1748-0221/5/12/C12046>)

View [the table of contents for this issue](#), or go to the [journal homepage](#) for more

Download details:

IP Address: 137.138.124.142

The article was downloaded on 11/07/2011 at 15:41

Please note that [terms and conditions apply](#).

RECEIVED: October 29, 2010

REVISED: November 25, 2010

ACCEPTED: December 6, 2010

PUBLISHED: December 21, 2010

TOPICAL WORKSHOP ON ELECTRONICS FOR PARTICLE PHYSICS 2010,
20–24 SEPTEMBER 2010,
AACHEN, GERMANY

Upgrading the ATLAS Level-1 Calorimeter Trigger using topological information

Y. Ermoline¹

*Department of Physics and Astronomy, Michigan State University,
BPS Building, East Lansing, MI 48824-2320, U.S.A.*

E-mail: Yuri.Ermoline@cern.ch

ABSTRACT: The ATLAS Level-1 Calorimeter Trigger (L1Calo) is a fixed latency, hardware-based pipelined system designed for operation at the LHC design luminosity of $10^{34}\text{cm}^{-2}\text{s}^{-1}$. Plans for a several-fold luminosity upgrade will necessitate a complete replacement for L1Calo (Phase II). But backgrounds at or near design luminosity may also require incremental upgrades to the current L1Calo system (Phase I). This paper describes a proposed upgrade to the existing L1Calo to add topological algorithm capabilities, using Region of Interest (RoI) information currently produced by the Jet and EM/Hadron algorithm processors but not used in the Level-1 real-time data path.

KEYWORDS: Trigger concepts and systems (hardware and software); Digital electronic circuits

¹On behalf of the ATLAS TDAQ collaboration [1]

Contents

1	Introduction	1
2	Ideas on limited short-term topological upgrade	2
2.1	Limitations of current system and possible improvement	2
2.2	Modifications required for limited upgrade	2
2.3	Common Merger Module modifications	3
3	Recent R&D towards CMM++ specification and design	4
3.1	Backplane data transfer	4
3.2	Latency survey	4
3.3	Optical links study	4
3.4	Technology study	5
3.5	Firmware study	5
4	Conclusion	5

1 Introduction

The L1Calo trigger (figure 1) is a fixed-latency, 40 MHz pipelined digital system [2]. Its input data comes from about 7200 analogue trigger towers of reduced granularity, mostly 0.1×0.1 in $\Delta\eta \times \Delta\phi$, from all the ATLAS electromagnetic and hadronic calorimeters.¹ The L1Calo electronics has a latency of less than a microsecond, resulting in a total latency of about $2.1 \mu\text{s}$ for the L1Calo chain including cable transmission delays and the Central Trigger Processor (CTP) processing time.

The Cluster Processor (CP) identifies candidate electrons, photons and τ 's with high ET above programmable thresholds and, if desired, passing isolation requirements. The Jet/Energy-sum Processor (JEP) operates on 'jet elements' at the somewhat coarser granularity of 0.2×0.2 in $\Delta\eta \times \Delta\phi$ to identify jets as well as produce global sums of total, missing, and jet-sum ET . Both the CP and the JEP count 'hit' multiplicities of the different types of trigger objects, and send them, together with tower energy information (total and x,y components), to the two Common Merger Modules (CMMs) in each crate.

The 'crate' CMMs (one at each end of the main block of modules in each crate) processes the results from the CPMs or JEMs to produce results over the entire crate, and send them to a 'system' CMMs in order to produce system-wide results. These final results are sent on cables to the CTP. Upon receiving the L1Accept signal from the CTP, the Region of Interest (RoI) information is sent to the Data Acquisition (DAQ) system through ReadOut Drivers (RODs).

¹ η is pseudo-rapidity and ϕ is azimuthal angle around the beam axis.

Figure 1. L1Calo trigger system and modifications required for limited upgrade.

2 Ideas on limited short-term topological upgrade

2.1 Limitations of current system and possible improvement

The present trigger capabilities allow to make selections on counts of objects of various types (for example 2 jets > 40 GeV), and even separate counts of objects (eg MET > 50 GeV & 2 jets > 40 GeV). However, there is presently no provision for spatial correlation of different objects, or for differentiating jets and em/tau clusters identified in the different subsystems but originating from the same energy deposits.

A possible solution could be to include jet/cluster position information (RoI) in the real time data path (in the current system it is available only for the DAQ system) and to use this information to add topology-based algorithms at Level 1. For example, identification of spatial overlap between e/tau clusters and jets, usage of the local jet Et sum to estimate energy of overlapping e/tau objects and calculation of invariant transverse mass would be possible. Some of them require only local information, others need global information.

2.2 Modifications required for limited upgrade

For several years the calorimeter electronics and the trigger hardware up to and including the Cluster Processor Modules (CPMs) and Jet/Energy Modules (JEMs) will remain essentially unchanged. The hardware components on which the system is built are about 5 to 7 years old and don't allow much freedom for modifications; many parts are already obsolete. A limited modification of the algorithm processor firmware can be done in order to extract the extra information, but in order to run topological algorithms a new module must be designed. This module will replace the so-called "common merger" module (CMM) in the current system.

In the proposed new architecture (figure 1) the CP and JEP systems (and L1 Muon trigger processor) transmit additional RoI information (which is not currently used in the real-time data path) taking advantage of the higher bandwidth potential inherent, but not used, in the backplane to the re-designed CMM module (CMM++). The data transfer rate over the crate backplane can be increased from 40 Mbit/s to 160 Mbit/s. A topological processor (TP) performing more sophisticated algorithms on the combined feature set and sending results to CTP can be added at later stage.

Figure 2. Current CMM and proposed CMM++.

2.3 Common Merger Module modifications

The current CMM module [3] processes the results from the CPM or JEM modules to produce results over the entire crate, and send them to a “system” CMM in order to produce system-wide results. These final results are sent on cables to the CTP. A “crate” FPGA on each CMM receives backplane data and produces crate/wide sums of identified features. A “system” FPGA collects crate results over LVDS cables and sends the trigger output to the CTP. On an L1A, data and ROI are sent via G-Links to the DAQ system. All CMMs (figure 2) have identical hardware and several different firmware variants allowing to perform different functions.

The current L1Calo trigger system must remain unchanged for the next few years and it would be desirable that the system modifications can be done and tested in parallel with the running system. The CMM++ development scenario assumes that the module can be a drop-in replacement for the CMMs, with the ability to perform additional logic on top of providing all the necessary backwards compatible CMM functionality (figure 2). The two basic requirements are that this module should:

- provide all the necessary functionality to replace a current CMM (electrical interfaces, programming model, data formats),
- be able to transmit all the backplane data, received from upgraded CPM/JEM modules, onwards to the TP over optical links (with or without internal processing) and to receive the data via optical links in order to implement internal topology processing without TP.

A desirable extra feature would be that the module could perform extra processing (and possibly output extra trigger bits) which would act as a test bed for future trigger algorithms. Development of such a module can be staged in the following way:

- hardware design with all present interfaces plus the optical links for the new topological processor, use one large FPGA, adaptation of the current CMM firmware to the new hardware for initial use,
- upgrade CPM/JEM and CMM++ modules firmware using new data format and 160 Mb/s data transfer, incrementally add new functionality, data supplying to the topological processor connected to the CTP.

3 Recent R&D towards CMM++ specification and design

In order to prepare the CMM++ specification, several feasibility studies are currently under way in different areas, namely — backplane data transfer rates tests, latency survey, optical link and FPGA technologies studies, study of transferring current firmware to new hardware.

3.1 Backplane data transfer

The CMM receives over the backplane from the CPM/JEM modules up to 400 signal lines on each 40 MHz clock cycle. The CMM++ interface to CPM/JEM shall be able to run in the backward compatible (CMM) version first and then shall be upgraded to 160 Mb/s data transfer and a new data format. The deployment of the CMM++ module requires firmware modifications in the CPM and JEM in order to collect the RoI information, generated in the modules, and to send it to the CMM++ over the crate backplane.

The Backplane and Link Tester module (BLT) [4] was built to qualify the backplane transmission lines for increased data rates between CPM/JEM and CMM++ modules. As a result of the backplane transmission test, stable data transmission at 160Mb/s was achieved with source termination of the data lines and the forwarded clock line sink terminated. Therefore, for each 40 MHz clock cycle 96 bits of data (24×4) can be transferred from each CPM/JEM module to CMM++ module. The 25th signal line on the backplane will be used to forward the encoded clock/parity onto the CMM++.

3.2 Latency survey

The maximum ATLAS L1 trigger latency was defined as 2500 ns (100 bunch crossings, BCs) — from colliding beams to L1Accept arrival to detector front-end electronics. The new topology algorithms require additional latency to process the data and it was decided to measure the actual latencies of different parts and paths of the L1Calo system.

The total L1Calo latency of 36 BCs was measured in the counting room in the installed complete L1Calo system. The detailed breakdown was performed in the test setup, which allows access to individual modules/cables. These measurements provided an insight into the used latencies in different parts of the L1Calo system [5]. A maximum possible reserve for the upgrade in complete L1 system is about 18 BCs by the expense of increased dead time.

3.3 Optical links study

In order to investigate a possibility to exploit inexpensive 30 Gbit/s link using commercial components (Xilinx Spartan-3 FPGA, 10Gb Ethernet transceiver TLK3114SC, SNAP12 Tx/Rx pair) a prototype was built.

It is driven by the LHC TTC clock, jitter reduced with LMK03033CISQ clock conditioner. To run links synchronously with LHC clock, alignment characters in some LHC bunch gaps are sent for link maintenance.

3.4 Technology study

The CMM++ module will be based on new components (modern FPGAs, high-speed optical links). In order to acquire experience with the new technologies, the GOLD (Generic Optical Link Demonstrator) is under design.

Its main features include new, state-of-art FPGAs (Xilinx XC6VLXxxxT and Xilinx XC6VHX380T), optical links (6.4–10 Gbit/s) with SNAP12 and non-SNAP12 compliant transceivers, manufactured by Avago, optical backplane connectors, new ATCA form factor.

3.5 Firmware study

The CMM++ development scenario assumes that the module can be initially a drop-in replacement for the CMMs. This implies the adaptation of the current CMM firmware to the new hardware in order to provide full backward compatibility and testing with the current system.

A first attempt was made to port Jet CMM firmware to a Virtex 6 device (XC6VHX565T-2FF1924). The aim was to use existing VHDL with minimal changes, to update architecture-specific features, to estimate I/O requirements and to produce a realistic user constraint file for timing simulation.

Results so far demonstrate that existing VHDL is easily ported, uses $\sim 2\%$ of available resources and by emulating G-Link in FPGA, can keep I/O count below 600/640 pins.

4 Conclusion

Performed simulations show that rates for electrons and jets (especially of larger size) will be hardly kept within the current L1 trigger rate budget even near the design luminosity. Trigger algorithms may need to be augmented to reduce rates and improve selectivity even before the current L1Calo trigger system is replaced.

A promising solution consists in adding topological algorithms based on relationships among triggered objects and thereby reduce the L1 rate. This modest modification to the current system will have low impact on other ATLAS components.

The R&D projects and studies are well under way: backplane data transfer rates, technology demonstrator, firmware studies and others. Preliminary results of these studies look promising and proposed modifications will improve the performance of the current ATLAS L1Calo trigger system.

References

- [1] ATLAS TDAQ collaboration, *The ATLAS Trigger/DAQ Authorlist, version 4.0*, ATL-DAQ-PUB-2010-002, CERN, Geneva (2010), <http://cdsweb.cern.ch/record/1265604>.
- [2] R. Achenbach et al., *The ATLAS Level-1 Calorimeter Trigger*, 2008 *JINST* **3** P03001.
- [3] I.P. Brawn and C.N.P. Gee, *Specification of the Common Merger Module*, v1.8, ATL-DA-ES-0021, 18 July 2008, http://hepwww.rl.ac.uk/Atlas-L1/Modules/CMM/CMM_V1.8.pdf.
- [4] B. Bauss, U. Schäfer and C. Schröder, *Backplane Tester Module*, v0.3, 11 December 2008, http://www.staff.uni-mainz.de/baushh/Backplanetester/Sheets_Vers1.0.pdf.
- [5] Y. Ermoline, *Latency Measurement*, ATL-DA-OR-0004, 24 November 2009, https://edms.cern.ch/file/1100580/0.1/Latency_measurements.doc.