

First data with the ATLAS Level-1 Calorimeter Trigger

R. Achenbach², P. Adragna⁴, M. Aharrouche³, V. Andrei², B. Åsman⁶, B. M. Barnett⁵, B. Bauss³, M. Bendel³, C. Böhm⁶, J. R. A. Booth¹, J. Bracinik¹, I. P. Brawn⁵, D. G. Charlton¹, J. T. Childers², N. J. Collins¹, C. J. Curtis¹, A. O. Davis⁵, S. Eckweiler³, E. Eisenhandler⁴, P. J. W. Faulkner¹, J. Fleckner³, F. Föhlisch², C. N. P. Gee⁵, A. R. Gillman⁵, C. Göringer³, M. Groll³, D. R. Hadley¹, P. Hanke², S. Hellman⁶, A. Hidvégi⁶, S. J. Hillier¹, M. Johansen⁶, E.-E. Kluge², T. Kuhl³, M. Landon⁴, V. Lendermann², J. N. Lilley¹, K. Mahboubi², G. Mahout¹, K. Meier², R. P. Middleton⁵, T. Moa⁶, J. D. Morris⁴, F. Müller², A. Neusiedl³, C. Ohm⁶, B. Oltmann³, V. J. O. Perera⁵, D. P. F. Prieur⁵, W. Qian⁵, S. Rieke³, F. Rühr², D. P. C. Sankey⁵, U. Schäfer³, K. Schmitt², H.-C. Schultz-Coulon², P. Seidler², S. Silverstein⁶, J. Sjölin⁶, R. J. Staley¹, R. Stamen², M. C. Stockton¹, C. L. A. Tan¹, S. Tapprogge³, J. P. Thomas¹, P. D. Thompson¹, P. M. Watkins¹, A. Watson¹, P. Weber², M. Wessels², M. Wildt³

Abstract – The ATLAS Level-1 Calorimeter Trigger is one of the main elements of the first stage of event selection for the ATLAS experiment at the LHC. The input stage consists of a mixed analogue/digital component taking trigger sums from the ATLAS calorimeters. The trigger logic is performed in a digital, pipelined system with several stages of processing, largely based on FPGAs, which perform programmable algorithms in parallel with a fixed latency to process about 300 Gbyte/s of input data. The real-time output consists of counts of different types of physics objects, and energy sums. The final system consists of over 300 custom-built VME modules, of several different types. The installation at ATLAS of these modules, and the necessary infrastructure, was completed at the end of 2007. The system has since undergone intensive testing, both in standalone mode, and in conjunction with the whole of the ATLAS detector in combined running.

The final steps of commissioning, and experience with running the full-scale system are presented. Results of integration tests performed with the upstream calorimeters, and downstream trigger and data-flow systems, are shown, along with an analysis of the performance of the calorimeter trigger in full ATLAS data-taking. This includes trigger operation during the cosmic muon runs from before LHC start-up, and a first look at LHC proton beam data.

I. INTRODUCTION

The CERN Large Hadron Collider presents huge challenges for the trigger systems of its experiments. Proton bunches collide every 25 ns, and at full luminosity there will be an average of 23 proton–proton collisions per bunch-crossing. ATLAS will deal with this using a three-level selection scheme in order to select about 200 bunch-crossings per second for permanent recording.

The first level of trigger selection, Level-1, uses reduced-granularity information from the calorimeters and muon information from dedicated fast triggering detectors to reduce the 40 MHz bunch-crossing rate down to 75 kHz (and potentially 100 kHz) within a fixed latency maximum of 2.5 μ s. The second and third levels of the trigger system, known as Level-2 and the Event Filter respectively, are based on large-scale farms of commodity computers. Level-2 can access the full-granularity detector readout data from all components of ATLAS, and further reduces the rate to about 3.5 kHz within about 40 ms. To do this, it uses so-called “region-of-interest” (RoI) information from Level-1 to access potentially interesting regions in the readout data. Finally, the Event Filter works on fully built events, and can use software selections similar to the offline analysis. This cuts the rate to the desired 200 Hz within about 4 s.

The Level-1 Trigger cannot use computers because it must handle too much data, and must produce its results within a fixed and very short latency while the full readout data is being held on the detectors in front-end buffers. In fact, due to inevitable cable transmission delays from the detector and back to it, the Level-1 Trigger logic must do its work in around 1 μ s. The solution adopted is to use custom–built electronics, with as much programmable flexibility as possible provided by using ASICs, FPGAs and dedicated firmware to carry out the trigger algorithms.

Manuscript received 13 November 2008.

¹ School of Physics and Astronomy, University of Birmingham, Birmingham B15 2TT, UK.

² Kirchhoff-Institut für Physik, University of Heidelberg, D-69120 Heidelberg, Germany.

³ Institut für Physik, University of Mainz, D-55099 Mainz, Germany.

⁴ Physics Department, Queen Mary, University of London, London E1 4NS, UK.

⁵ STFC Rutherford Appleton Laboratory, Harwell Science and Innovation Campus, Didcot, Oxon OX11 0QX, UK.

⁶ Fysikum, Stockholm University, SE-106 91 Stockholm, Sweden.

Corresponding author: Eric Eisenhandler, Physics Department, Queen Mary University of London, London E1 4NS, UK. Email: e.eisenhandler@qmul.ac.uk.

The Level-1 Calorimeter Trigger has been supported by BMBF, Germany; The Swedish Research Council and the Knut and Alice Wallenberg Foundation, Sweden; and the Particle Physics and Astronomy Research Council (up to 2007) and the Science and Technology Facilities Council (from 2007), United Kingdom.

A block diagram of the Level-1 Trigger is shown in Fig. 1. The calorimeter and muon detector information is fed into trigger processors known as the Level-1 Calorimeter Trigger (“L1Calo”) and Level-1 Muon Trigger (“L1Muon”). Results from these, mainly in the form of multiplicities of objects exceeding various trigger thresholds, are sent to the Central Trigger Processor (CTP). The CTP implements the trigger “menu”, consisting of a set of logical conditions on the objects found (e.g. two jets having transverse energies above 100 GeV together with one muon having transverse momentum above 10 GeV). The result is a Level-1 Accept (L1A) signal if any of the items in the trigger menu are satisfied. At the same time, L1Calo, L1Muon, and the CTP all send region-of-interest (RoI) information to the Level-2 Trigger.

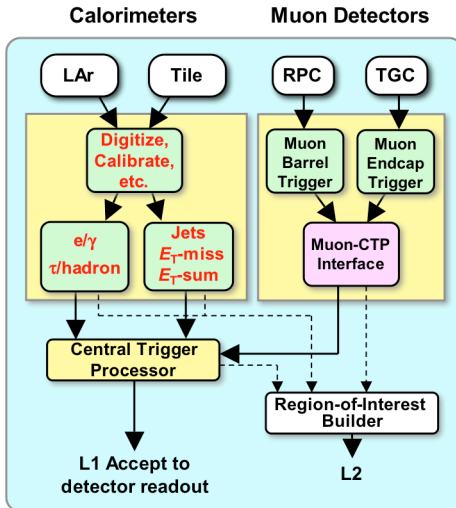


Fig. 1. Block diagram of the ATLAS Level-1 Trigger, showing the Calorimeter and Muon Triggers, Central Trigger Processor and Region-of-Interest Builder for the Level-2 Trigger.

II. THE LEVEL-1 CALORIMETER TRIGGER

A. General design features

The Level-1 Calorimeter Trigger (L1Calo) consists of three main subsystems: the PreProcessor (PPr) is followed by the Cluster Processor (CP) and the Jet/Energy-sum Processor (JEP), operating in parallel. In addition, and off the real-time path, Readout Driver modules collect and format data both for storage via the data acquisition system, and for region-of-interest data used by the Level-2 Trigger.

Fig. 2 is a simplified block diagram of L1Calo. It uses five main types of 9U VME modules: PreProcessor Modules, Cluster Processor Modules, Jet/Energy Modules, Common Merger Modules, and Readout Driver modules. A feature of the system is that the Common Merger and Readout Driver modules are used for several different applications by changing the firmware that is loaded.

In terms of digital signals, the data rate processed by the trigger is roughly 300 Gbyte/s. This is achieved by the use of massive parallelism in the logic. The $\sim 1 \mu\text{s}$ latency of L1Calo logic is fixed, and the operations are done using pipelined

processing. Most steps are carried out using a 40 MHz clock, but in a few places 80 and 160 MHz multiples are used.

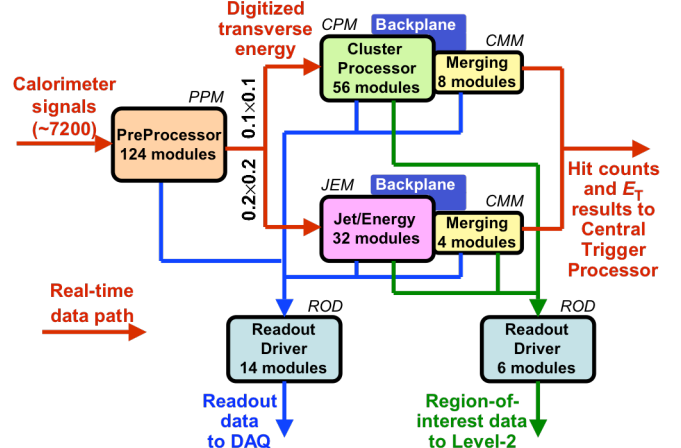


Fig. 2. Block diagram of the Level-1 Calorimeter Trigger, showing the main module types, the quantities of each used, and their connectivity.

The algorithms used demand extensive use of overlapping, sliding windows. This requires a great deal of signal fan-out and very complex connectivity, which shapes much of the design. In order to retain flexibility and the ability to change the way data are processed, the trigger only utilizes one L1Calo-specific ASIC, with most of the logic done by FPGAs.

B. Input signals

The input data from the calorimeters consists of about 7200 analogue trigger-tower signals, which arrive in the electronics cavern on 16-channel twisted-pair cables. Over most of the detector the trigger-towers are 0.1×0.1 in $\Delta\eta \times \Delta\phi$, but with some larger towers at the ends of the barrel and in the endcaps. Trigger towers are summed over the full depth of either the electromagnetic or hadronic calorimeters.

The calorimeter pulses are quite wide compared to the bunch-crossing interval. The rise-time is nearly 50 ns and the fall-time even longer. We will describe below the measures needed to assign pulses to the correct LHC bunch-crossing.

Just upstream of the Calorimeter Trigger, the signals pass through Receiver modules which adjust the gains so they are proportional to transverse energy (E_T). The Receivers are where the main energy calibration is applied.

C. The PreProcessor

The PPr consists of 124 PreProcessor Modules (PPMs), which provide the digital input data used by both the Cluster Processor and the Jet/Energy-sum Processor. Physically, the PPr occupies eight crates. Each PPM handles 64 trigger towers. A diagram indicating the processing steps is shown in Fig. 3.

After signal-conditioning to match the input range of the ADCs, and to add a DAC-controlled level for setting the pedestal, the bulk of the PPM logic is on small multi-chip modules (MCMs), each of which handles a group of four trigger towers. A photograph of an MCM is shown in Fig. 4.

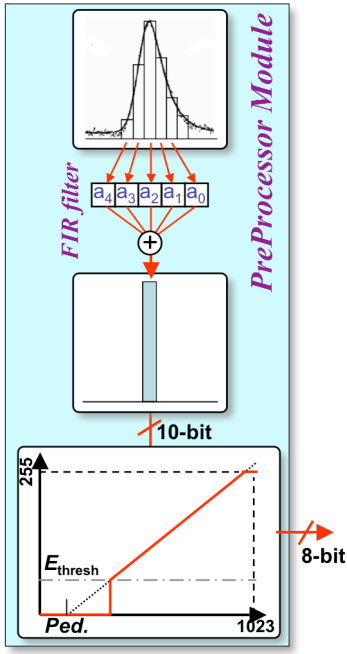


Fig. 3. Processing steps within the PreProcessor. After digitization, a digital filter and peak-finder associate the pulse with the correct bunch-crossing. A look-up table then does pedestal subtraction, noise suppression and final calibration, producing an 8-bit value for algorithmic processing.

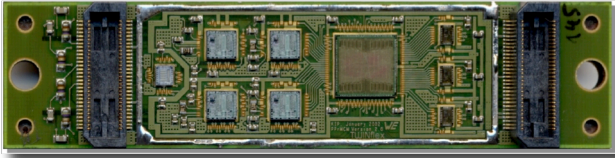


Fig. 4. Photograph of a PreProcessor Multi-Chip Module. The four ADC chips are left-of-centre, and the ASIC to its right. The three chips to the right of the ASIC are LVDS output drivers.

The signals are digitized every 25 ns to 10-bit precision. The scale is 0.25 GeV of E_T per count, with a saturation level of about 250 GeV. The ADC strobe timing can be adjusted in 1 ns steps in order to sample the pulse peak accurately.

The digital data must be converted from a series of ADC values covering each pulse into a single E_T value associated with the correct bunch-crossing, a process called bunch-crossing identification. This is done using a digital finite-impulse-response filter followed by a peak finder. On each bunch-crossing, five successive digitizations are multiplied by programmable coefficients (normally set to an approximation of the pulse shape), then the sum of these five products are compared to the sums from the two neighbouring bunch-crossings to see if there was a peak. Low-order bits are then discarded to produce a 10-bit transverse-energy value.

For saturated pulses, which might represent exciting new physics, the bunch-crossing identification is done by using comparators on the leading edge.

The 10-bit values are used as input addresses to a look-up table, which converts them to an 8-bit calibrated E_T value for use in the algorithmic trigger logic. The nominal scale for this is 1 GeV per count.

The look-up table does several useful operations simultaneously. First, it subtracts off the pedestal. Second, input values near zero can be suppressed in order to reduce noise. Third, it sets pulses deemed to be saturated to full scale, i.e. 255 GeV. Fourth, it can do a final calibration correction (and even apply a non-linear scale if necessary). Finally, it can be used to zero any problematic or faulty trigger towers.

The PPr ASIC includes automatic facilities for rate monitoring of either ADC or look-up table results. The results, which are useful e.g. for finding noisy trigger towers, are read out via VMEbus.

A final stage of the PPr sums 2×2 trigger towers to form coarser 0.2×0.2 regions for use in the Jet/Energy-sum Processor, which does not need the finest granularity.

Trigger towers are sent via LVDS serial cable links at 400 Mbit/s to the Cluster Processor, and the 0.2×0.2 regions are sent over similar links to the Jet/Energy-sum Processor. Signals near boundaries in azimuth are fanned-out to two crates in order to allow the algorithmic sliding windows to span these boundaries. In addition, since trigger towers cannot have peaks on two successive beam-crossings, a coding scheme allows each link to the Cluster Processor to carry two trigger towers, which halves the number of links needed.

D. The Cluster Processor

The CP consists of four crates, each housing 14 Cluster Processor Modules (CPMs) and two Common Merger Modules (CMMs). It implements the algorithm for electrons and photons (i.e. electromagnetic showers), and for τ leptons. The region covered is $|\eta| < 2.5$, which is the limit of the high-precision tracking.

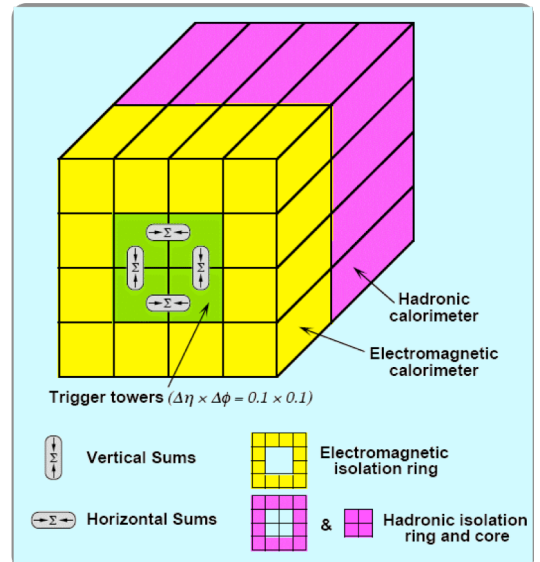


Fig. 5. The window and regions used for the electron/photon and τ algorithms.

The overlapping, sliding windows that are the basis of both types of algorithms are illustrated in Fig. 5. The electron/photon algorithm demands that the sum of E_T in two adjacent e.m. trigger towers must be more than a programmable threshold. At the same time, there are three

isolation conditions: the sum of the 12-tower e.m. ring around the central 2×2 “core” of the window must be less than an e.m. isolation threshold, the sum of the 2×2 core of the hadronic region behind must be less than a hadronic isolation threshold, and the sum of the 12-tower hadronic ring must be less than a hadronic-ring isolation threshold.

The τ algorithm is similar, but takes account of τ decays depositing energy in the hadronic calorimeter. The E_T in pairs of adjacent e.m. trigger towers are added to the sum of the 2×2 hadronic core behind, and compared to a programmable threshold. Two of the three isolation conditions are also the same – the hadronic core sum is of course not used. Since much of the arithmetic and logic is the same as for the electron/photon algorithm, this is done using the same logic, within the same FPGA.

For both algorithms, it is possible for several overlapping windows to satisfy the criteria simultaneously. In order to resolve the ambiguity, and at the same time to define the location of the region-of-interest (RoI) associated with the shower, the transverse-energy sum of the central 2×2 core must be a local maximum compared to all eight of its overlapping neighbours.

There are eight dedicated e.m. “threshold sets” (each with its own shower and isolation E_T thresholds), and eight more threshold sets each of which can be programmed to carry out either the e.m. or τ algorithms. This means that the results from each CPM consist of a multiplicity of “hits” for each of the 16 threshold sets. Each multiplicity is a three-bit integer; if there are seven or more hits the full-scale value is transmitted.

The sliding, overlapping windows mean that a great deal of information must be shared between modules. By arranging each crate to cover 90° in ϕ and a full range in η , and for each module to cover the full range in ϕ and a narrow strip in η , the data only needs to be shared between adjacent modules. In order to limit the number of connections needed for this, the shared data are partially serialized and transmitted at 160 MHz on the custom backplane described below.

The results from each module are sent via the backplane to two CMMs, located at either end of the crate. One CMM adds up the total multiplicities in the crate of e.m. hits for the eight dedicated e.m. threshold sets, while the other one adds the multiplicities for the eight threshold sets that can be either e.m. or τ clusters.

Finally, the overall system-wide sums of multiplicities are carried out by the CMMs in one of the crates, which receive data from the other crates on cable links.

E. The Jet/Energy-sum Processor

The JEP consists of two crates, each housing 16 Jet/Energy Modules (JEMs) and two Common Merger Modules. It implements the algorithm for jets, including separate possibilities for forward jets. It also adds up both total and missing E_T and compares them to thresholds, and finally does the same for an approximation to total jet E_T .

The JEP uses 0.2×0.2 “jet elements” covering the full calorimeter depth, formed by summing corresponding e.m. and hadronic 2×2 regions just after they arrive from the PreProcessor. The region covered is $|\eta| < 4.9$.

The jet algorithm is illustrated in Fig. 6. The algorithm allows selection between three window sizes: 2×2 , 3×3 or 4×4 jet elements. Larger windows are more efficient for inclusive jets, while smaller ones are better for resolving multiple jets.

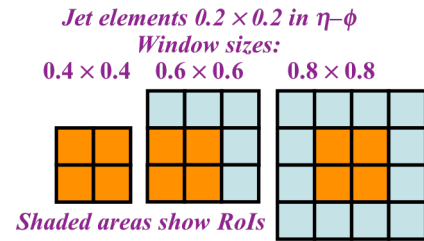


Fig. 6. The jet algorithm, showing the three window sizes available and their RoI positions. (For the 0.6×0.6 option there are four possible positions.)

As with the CP, in order to resolve overlapping windows that pass the same threshold, and to define the region-of-interest location, a 2×2 jet-element region within the window must be a local maximum. Data needing to be shared between adjacent JEMs are transmitted over the backplane at 80 MHz. By careful design the JEP and the CP use the same high-density backplane design.

There are eight independent jet thresholds, each with a selectable window size. There are also four thresholds on forward jets.

As with the Cluster Processor, each JEM counts the number of jet hits (to a maximum of seven) for each jet threshold, and transmits the result via the backplane to one of the two CMMs in the crate. This sums the crate multiplicities, and the CMM in one crate then forms the overall system multiplicities.

The JEMs also carry out a second important function, forming the total E_T on the module as well as the E_x and E_y components that are used to calculate missing E_T . These sums are transmitted in coded form to the second CMM in the crate, which does crate sums, and finally to one of the CMMs to form overall sums. The total scalar E_T is compared to four programmable thresholds. For missing- E_T , the E_x and E_y sums are used together as inputs to a look-up table which, in one step, effectively compares the resulting missing- E_T value to eight thresholds.

Finally, the JEP produces an approximate value for the total E_T in jets by using the numbers of jets passing each of the eight jet thresholds. This value is compared to four programmable thresholds.

F. Readout of trigger data and regions-of-interest

For all events accepted at Level-1 by the Central Trigger Processor, L1Calo provides a flexible selection of data to read out for monitoring and, if the event is accepted by the higher-level triggers, for long-term storage. In addition, region-of-interest data are read out for use by the Level-2 Trigger.

The selection of data to read out is programmable, and includes input data from the PPr and final results of the algorithms, as well as sufficient intermediate data to understand how the trigger was formed. Copies of the RoI data are also read out. Online and offline monitoring histograms allow rapid identification of any problems. In addition, a software simulation of the trigger allows all stages of the logic to be reconstructed from various points along the way, and to verify that L1Calo is functioning correctly. For testing and detailed fault tracing, all modules can provide readout data of their inputs and outputs. In addition, all modules have so-called playback memories near their inputs to allow offline testing of the operation of all electronics and digital links with simulated and test data.

The Level-2 Trigger uses the RoI data for essential guidance. This means that it only needs to access full-precision detector data for regions that contain objects of potential interest.

The readout and RoI data are handled by 20 Readout Driver modules (ROD). These receive data from all of the other modules, reformat into standard ATLAS event fragments, and transmit the data into the ATLAS data-flow infrastructure. All RODs use the same hardware design, but can handle data from different types of modules – both for DAQ and RoIs – by using different firmware.

III. COMMISSIONING THE LEVEL-1 CALORIMETER TRIGGER

A. Analogue and digital links

The full trigger system has been installed since the end of 2007. An important part of the installation and commissioning, which took a great deal of effort, was the large number of input cables and inter-module links. Despite their bulk, these had to be installed in a tidy way that allows for fault repair and, most crucially, minimizes their length and hence the trigger latency.

Incoming analogue signals from the calorimeters arrive on 16-way twisted-pair cables. Before reaching the PreProcessor, the signals must pass through analogue Receivers, and in many cases through various types of patch panels. That required 776 short cables of the same type. The analogue cables needed careful routing and custom-built supports because they are very stiff, heavy and awkward to handle.

Digital signals from the PreProcessor to the Cluster and Jet/Energy-sum Processors are sent serially, on parallel-pair cables at 400 Mbit/s. There are about 1900 cables, each carrying four signals. These again needed careful routing, handling and support, because the cables and connectors are somewhat fragile, the cable volume is considerable and tightly packed, and they are plugged directly into the high-density backplanes of the digital processors.

There have been very few problems with either the analogue or digital cable links. This is at least partially due to careful preparation, support structures, and installation.

Within the Cluster and Jet/Energy-sum Processors themselves, inter-module connectivity for the complex

arrangement of overlapping and sliding windows, as well as results from CPMs and JEMs to CMMs, is handled by the densely packed custom backplane already mentioned. This has about 1150 pins per 9U slot, and a total of about 22,000 pins on each backplane. There were a few initial problems with the backplane due to faulty pin installation and bent pins when plugging modules in, but these were repaired. Further damage is avoided by taking more care with module removal and installation, which is also done as little as possible.

Finally, readout of DAQ and RoI data to the RODs is done by optical fibres running at 800 Mbit/s. One fibre from each trigger module handles DAQ data, and another handles RoI data (except for PPMs, which do not generate RoIs). Once more, careful handling and installation has kept problems to a very low level.

All in all, the number of problems at any time on all the various types of links is typically a small fraction of 1%.

B. Pedestals and noise

ADC pedestals can be set by means of DACs at the inputs of the PreProcessor. This allows separate adjustment of each trigger tower. Great care has been taken with the grounding and shielding of trigger-tower signals in order to minimize noise. Doing pedestal measurements while the rest of ATLAS is operating allows noise levels to be checked. In case of noisy trigger towers, the source of the noise is tracked down and, hopefully, repaired. Towers with unacceptable noise can be masked out of the trigger.

The expected level of r.m.s. noise is about 0.4–0.5 GeV in energy. What we see in the trigger depends somewhat on the conversion to E_T and where in the electronics chain this is done. At this time, more than 99% of all trigger towers behave as expected, and efforts continue to identify and fix the remainder.

In recent cosmic-ray running (see below), the PreProcessor look-up tables have been operated with almost all towers having a noise cut-off of about 1 GeV. This yields very low trigger rates, with most triggers fired by genuine cosmic rays rather than noise.

C. Coarse and fine timing

The timing of signals into L1Calo, the internal timing within L1Calo, and the timing of results to the Central Trigger Processor must all be set correctly in order to achieve correct operation. There are two main aspects of the timing to consider.

The first, which affects all parts of the system, is “coarse” timing, mainly in 25 ns steps, with respect to the LHC 40 MHz bunch-crossing clock that is the basic heartbeat of L1Calo. This is typically done using FIFOs. Input analogue signals have to be synchronized to the same bunch-crossing, allowing for time-of-flight within the detector, different analogue cable lengths, etc. Timing is done in the PPMs so that all signals from a given bunch-crossing reach the look-up tables together. Further timing-in must be done for all inputs to the CPMs, JEMs, and CMMs, due to differing cable lengths and processing times.

In the Cluster Processor, the data sharing between neighbouring modules is done at 160 MHz, and in the Jet/Energy-sum Processor it is done at 80 MHz. The timing of these is therefore more critical, but can still be set up by standard procedures.

Everywhere in the system, the DAQ and RoI readout data are stored in FIFOs until a Level-1 Accept signal arrives, and so readout pointer offsets must be derived and set.

Typically, all these “coarse” timing settings are set up initially using calorimeter pulser systems, and are quite stable once set up. However, some of the timing settings are somewhat different when running with cosmics or physics.

The analogue input signals require more difficult “fine” timing settings, in order to achieve maximum energy resolution and to associate small pulses as accurately as possible with the correct LHC bunch-crossing. This is done by setting the ADC strobes to a precision of a few nanoseconds. The strobe timing is set by CERN Phos4 ASICs, allowing each trigger-tower strobe to be separately adjusted in 1 ns steps. Initial set-up is done using calorimeter pulser systems, by stepping the strobe timing in order to determine the location of the pulse peak. A typical plot from such a run is shown in Fig. 7.

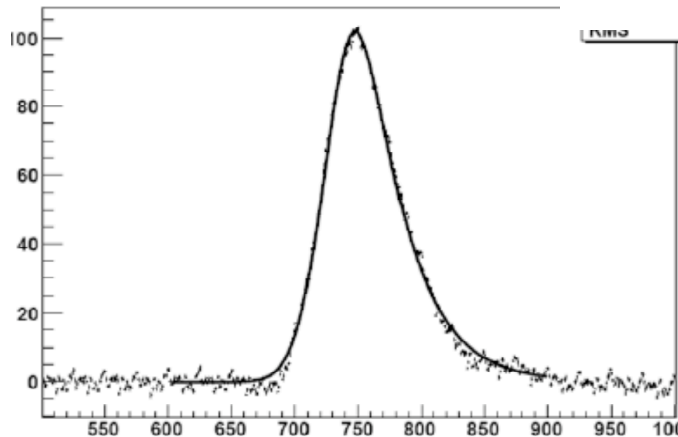


Fig. 7. Fine timing scan, in 1 ns steps, of an analogue input pulse.

However, for physics running the timing will be slightly different, and even the pulse shape will change somewhat. In physics running it will not be possible to do these so-called Phos4 scans, or to achieve the same level of statistical accuracy. Therefore, a procedure for fitting many individual pulses of varying amplitude with a mixture of Gaussian and Landau shapes has been developed, as illustrated in Fig. 8. The validity of this approach is indicated in Fig. 9, which shows the correlation between timing settings derived using the two methods.

D. Energy calibration

A critical aspect of the operation of L1Calo is the energy calibration of the input signals. This is initially done in special calibration runs, using the calorimeter charge-injection systems set up to produce pulses at a variety of energies.

To check the initial values, Fig. 10 compares e.m. and hadronic energies during a cosmic run, as measured by L1Calo and by the full-precision readout data from the calorimeters. The agreement is already fairly good. Some of the spread is due to the fact that cosmics do not arrive synchronous with the LHC bunch-crossing clock, and that means that the ADC strobes are not ideally timed.

Further work to improve L1Calo’s energy calibration will be done, by improving the ADC strobe timing, and by doing more calibration runs. When there is collision data, the calibration will be tuned further by using physics data

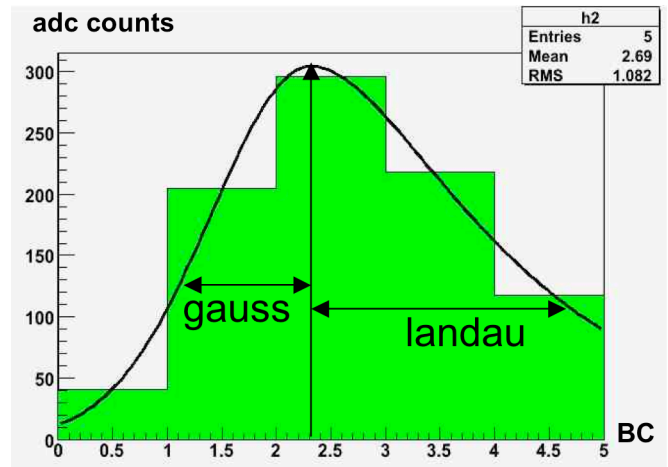


Fig. 8. Fit of Gaussian and Landau distributions to a single analogue input pulse, digitized for five successive bunch-crossings.

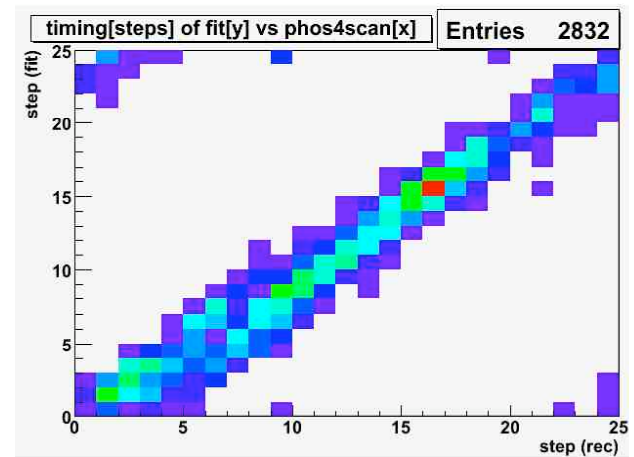


Fig. 9. Comparison of timing settings derived from pulse fits (vertical axis) and fine-step timing scans (horizontal axis).

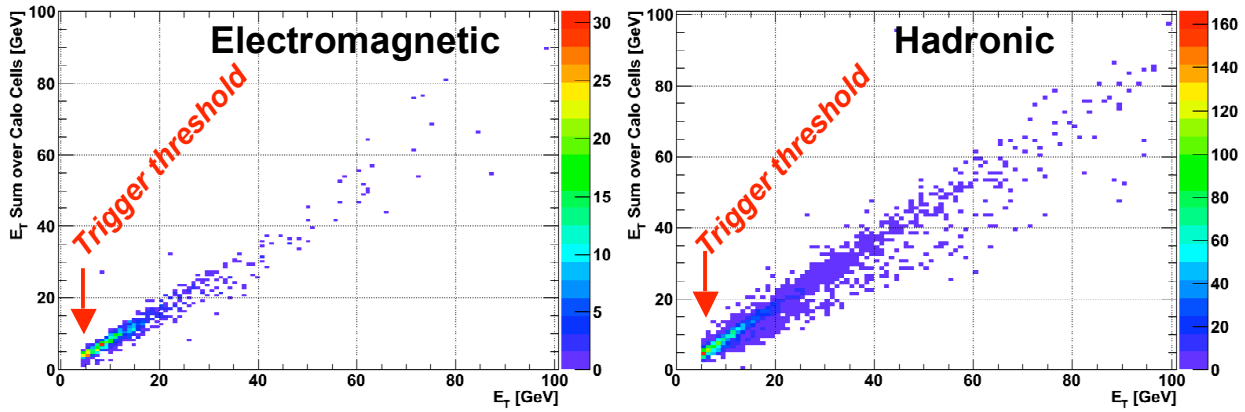


Fig. 10. Comparison of transverse energy measured by L1Calo (horizontal axis) with precision measurement by calorimeters (vertical axis).

IV. FIRST COSMIC AND PROTON-BEAM RESULTS

L1Calo was not designed to be able to trigger on cosmic muons. Minimum-ionizing particles are only one or two counts above pedestal in the ATLAS Tile Calorimeter, and give even lower pulses in the Liquid Argon Calorimeters. It is therefore gratifying that we have been able to set up an efficient high-energy cosmic trigger which is dominated by real events rather than noise. This cosmic trigger has also been very useful in setting up other parts of ATLAS.

Using an electron/photon E_T threshold of 3 GeV E_T produces a trigger rate of about 1 Hz. At the same time, τ and

jet thresholds of 5 GeV produce rates of about 4 Hz. Note that these thresholds are well below what would be used for physics running.

Fig. 11 shows an event display for a clean cosmic event triggered by the L1Calo τ and jet triggers. Fig. 12 shows a distribution of measured cosmic-event hadronic energies, which indicates that the majority of cosmic triggers are genuine and not due to noise.

Finally, in Fig. 13 we show an event display for the first event induced by an LHC proton beam that was recorded in ATLAS. This so-called “beam splash” event produced a trigger in L1Calo, as indicated on the display.

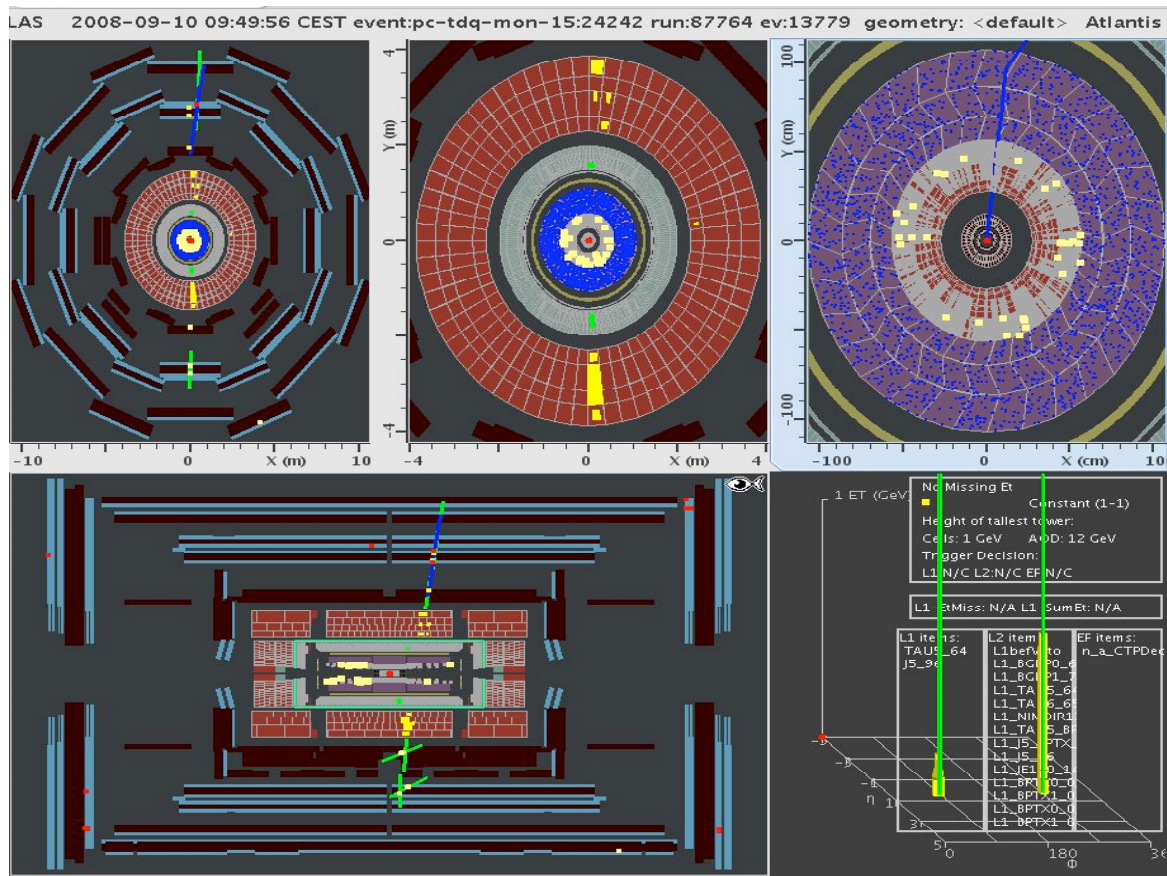


Fig. 11. A cosmic event triggered by L1Calo τ and jet triggers..

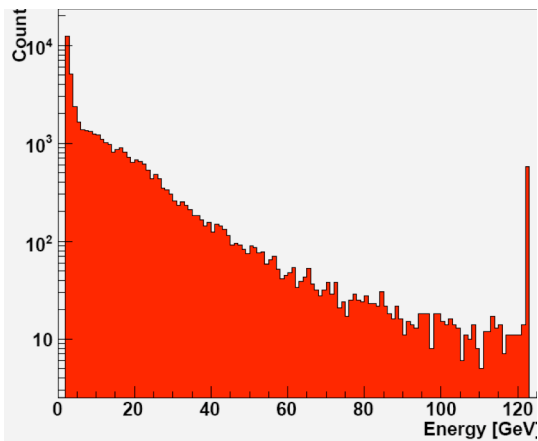


Fig. 12. Hadronic calorimeter energy spectrum of cosmic events. (The highest bin includes overflows.)

V. SUMMARY AND OUTLOOK

The ATLAS Level-1 Calorimeter Trigger has been installed and is working well. A number of important steps have been achieved: integration with the calorimeters, correct working of the algorithmic logic, integration with the Central Trigger Processor and Data Acquisition, trigger formation, and provision of regions-of-interest to Level-2.

The system has been tested in integrated runs with the rest of ATLAS over long periods, both for calibration and for

triggering on cosmic – and even injected proton beams. Events with significant energy have been used to cross-check the trigger decision against the data read out from the calorimeters themselves.

There are very few known problems, and they are all minor.

Although much fine tuning and work to make operation smoother and more reliable is still needed, the system is already running reliably and well. It is therefore ready for physics data.

VI. ACKNOWLEDGEMENTS

We wish to acknowledge the work of the ATLAS TDAQ community in providing the underlying online software and infrastructure for triggering, read-out and dataflow. We would also like to thank the ATLAS calorimeter communities, in particular those working on the trigger-tower builders and receivers, for their efforts to provide genuine input signals to the trigger. We are grateful to Björn Gosdzik, who produced Fig. 7. Finally, the successful installation of the infrastructure and cabling would have been impossible without the careful work of many technicians connected to the institutes involved.

REFERENCES

- [1] R. Achenbach et al., *The ATLAS Level-1 Calorimeter Trigger*, 2008_JINST_3_P03001.

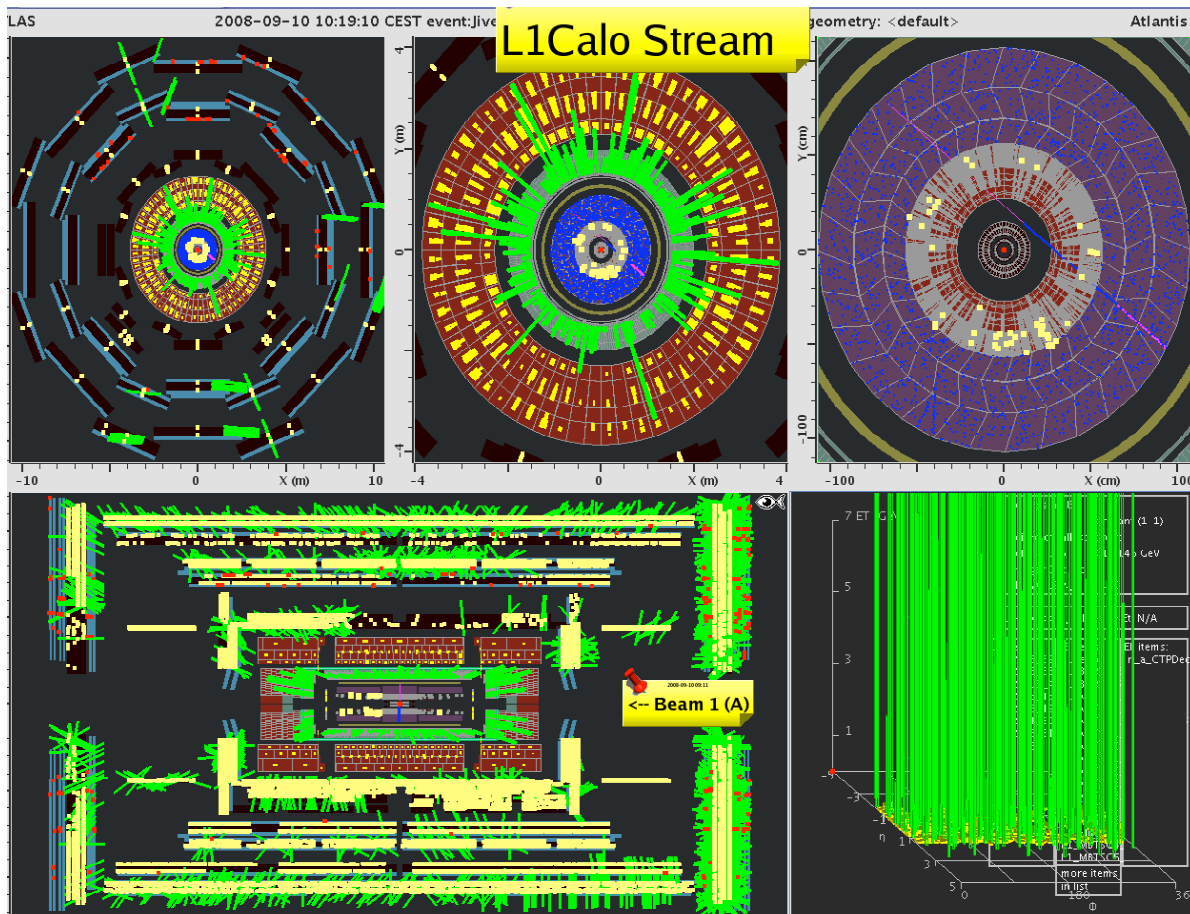


Fig. 13. First LHC proton beam event recorded in ATLAS.