

Commissioning Experience with the ATLAS Level-1 Calorimeter Trigger System

R. Achenbach¹, P. Adragna², V. Andrei¹, B.M. Barnett³, B. Bauss⁴, M. Bendel⁴, C. Boehm⁵, J.R.A Booth⁶, I.P. Brawn³, D.G. Charlton⁶, C.J. Curtis⁶, A.O. Davis³, E. Eisenhandler², P.J.W. Faulkner⁶, F. Föhlich¹, C.N.P. Gee³, C. Geweniger¹, A.R. Gillman³, P. Hanke¹, S. Hellman⁵, A. Hidvégi⁵, S.J. Hillier⁶, M. Johansen⁵, E.-E. Kluge¹, M. Landon², V. Lendermann¹, K. Mahboubi¹, G. Mahout⁶, K. Meier¹, V.J.O. Perera³, D.P.F. Prieur³, W. Qian³, S. Rieke⁴, F. Rühr¹, D.P.C. Sankey³, U. Schäfer⁴, K. Schmitt¹, H.-C. Schultz-Coulon¹, S. Silverstein⁵, R.J. Staley⁶, R. Stamen¹, S. Tapprogge⁴, J.P. Thomas⁶, T. Trefzger⁴, P.M. Watkins⁶, A. Watson⁶, P. Weber¹, E.-E. Woehrling⁶

¹Kirchhoff-Institut für Physik, University of Heidelberg, Heidelberg, Germany

²Physics Department, Queen Mary, University of London, London, UK

³STFC Rutherford Appleton Laboratory, Chilton, Oxon, UK

⁴Institut für Physik, University of Mainz, Mainz, Germany

⁵Fysikum, Stockholm University, Stockholm, Sweden

⁶School of Physics and Astronomy, University of Birmingham, Birmingham, UK

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 production of final modules started in 2006, and installation of these modules and the necessary infrastructure at ATLAS has been underway for some time, with the intention of having a full system in situ during 2007, before first collisions at the LHC.

The first experiences of commissioning and running the full scale system will be presented, along with results from integration tests performed with the upstream calorimeters, and the downstream trigger and data flow systems.

Index Terms — triggering, pipeline processing, real time systems, parallel architectures

I. INTRODUCTION

The ATLAS Level-1 trigger is designed to provide a trigger decision within a fixed time of 2 μ s in order to reduce the LHC bunch-crossing rate of 40 MHz down to a rate of less than 75 kHz of events to be retained for the second level of event selection. The Level-1 decision for physics events of interest is based only on reduced granularity calorimeter and muon detector data. The ATLAS Calorimeter Trigger is the part that processes the calorimeter information, which comprises an input of over 7000 analogue signals. The algorithms used have to be simple enough to be performed over a large number of input signals in this limited time, but

sophisticated and flexible enough to distinguish potentially interesting particle signatures from a large and, to some extent, unpredictable background.

These requirements have necessitated a design which incorporates several layers of parallel processing being performed in fixed latency pipelines. All of the physics algorithms are implemented in FPGAs to allow flexibility to allow for changing LHC conditions and trigger requirements. The nature of the current algorithms, which make extensive use of overlapping, sliding windows, mean that the ability to transfer large amounts of digital data between crates and modules is a crucial aspect of the system. Providing the necessary data throughput at each stage of the processing has been one of the major challenges of the system design. With the installation of the infrastructure and electronics of the final, full scale system at the ATLAS experiment now well underway, the feasibility of the design chosen, and the consequences of the choices taken can now be assessed.

The Level-1 Calorimeter trigger system consists of several designs of module, and each of these was extensively tested on a small scale at the prototyping phase. Both the generic module design and the testing programme, in the laboratory and at the ATLAS test-beam, have been described previously [1]. Production started for the majority of the modules in 2006. At the same time the crates and cabling were being installed at ATLAS. This has provided the opportunity to make tests on a far larger scale than previously possible. It also means that integration of the trigger system with the installed calorimeter detectors, high-level trigger and data flow system can be performed. Results from investigations made on complete crates of modules, as well as the outcome of tests of the several input and output interfaces, will be reported.

II. THE ATLAS LEVEL-1 CALORIMETER TRIGGER ARCHITECTURE

The fast real-time output of the trigger system consists of counts of electron/photon-like, tau-like, or jet-like clusters above programmable transverse energy thresholds, as well as results of threshold comparisons on missing and total transverse energy to be sent to the Central Trigger Processor (CTP) [2]. However, all of the modules also have read-out capability, in order to verify the correct performance during normal operation. This read-out only occurs on events which pass the CTP Level-1 decision. On these events, additional location information on trigger objects (known as Region-of-Interest data) is also sent to the Level-2 trigger system.

The basic architecture of the whole of the Level-1 system was documented in an ATLAS TDR in 1998 [3]. Some evolution of the calorimeter trigger has taken place since then and a detailed description of the current design has been presented in [4]. A simplified schematic of the modules and dataflow is shown in fig. 1.

The real-time path consists of three subsystems: the Preprocessor (PPr), Cluster Processor (CP) and Jet/Energy-sum Processor (JEP). The Preprocessor system consists of 124 Preprocessor Modules (PPM), which provide the input data used by both the CP and JEP systems. Physically, they are 9U VME modules which fit into eight crates. The module input consists of analogue pulses, mostly corresponding to 0.1×0.1 sums in eta/phi space, from the ATLAS calorimeters. These input signals are often referred to as trigger towers. These towers are digitized and energy is assigned to the correct bunch-crossing from which each pulse originated. Finally, lookup tables perform the E_T calibration for these trigger towers and these form the basis of the digital trigger decision. Data are sent downstream to the CP and JEP systems using LVDS 400 Mbit/s serial link chipsets in order to reduce the I/O requirements on cables and pins.

The Cluster Processor (CP) system consists of 56 Cluster Processor Modules (CPM) which identify and count electron/photon and tau candidates. The final sums are performed in 8 Common Merger Modules (CMM), and sent to the CTP. Together, these occupy four 9U VME crates. The Jet/Energy-sum processor (JEP) consists of 32 Jet/Energy Modules (JEM) which count jet candidates and make missing and total transverse energy sums, with the final results again being summed in 4 CMMs. The JEP system fits into two 9U crates, which are identical to those of the CP system. Both systems require the exchange of a large volume of data between neighbouring modules, for which a common custom backplane has been designed. This backplane contains over 22,000 pins, through which digital signals with speeds of up to 400 Mbit/s differential and 160 Mbit/s single-ended are propagated.

The read-out and Region-of-Interest data is handled by 20 Read-out Driver modules (ROD). These receive signals from all of the other modules via optical links running at a maximum of 800 Mbit/s using the Agilent G-link protocol. The data is reformatted into standard ATLAS event fragments,

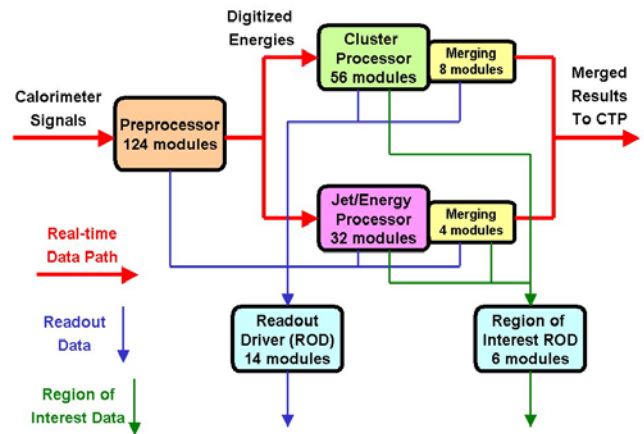


Fig. 1. Module types, numbers and connectivity in the complete Level-1 Calorimeter Trigger System

and transmitted on optical links using the ATLAS S-Link protocol.

From this description, it can be seen that the difficult task of extracting trigger objects is achieved by making use of a highly parallel, multi-stage processing design. One of the consequences of this architecture is that there is a requirement for large bandwidth data transmission, both between modules in different crates, and between modules sharing a common backplane. The rationale behind the choices made on how best to split the processing between modules, and the practical consequences for installation, will be described below.

III. BREAKDOWN OF THE TRIGGER SYSTEM INTO A PARALLEL ARCHITECTURE

In many ways, the ideal solution to designing a suitable trigger processor would be to have a very compact and powerful processor which could receive all the inputs and produce the final outputs. Compared to a large parallel system with several stages, this has many advantages, such as:

- Low latency
- No need for fast data transmission between stages
- Better capacity for data sharing

However, the volume of input data, and the complexity of the processes that have to be applied to the signals makes a very compact system impossible with current technology. Within the trigger system, the analogue signals require some analogue processing prior to digitization before any digital algorithms can be performed. One consideration is that combining such a mixed analogue and digital system requires care and it would be difficult to miniaturize without compromising the analogue performance.

After digitization and energy calibration, the data consists of over 7000 8-bit words, equating to about 2300 Gbits/s to feed into a digital processor. Even if a processor could be found with the power to perform the algorithms, there is a limit to the i/o capability. Clearly it is necessary to parallelize the processing, both at the analogue and digital stages.

Physical limitations were also important in determining the final architecture. Due to demanding requirements of signal quality and low cross-talk on the input signals, the cables used

for the analogue signals are large and inflexible, and the connectors also are large. Each cable carries 16 signals, and it is only possible to fit four of these into a single 9U module. In the chosen design, 496 of these input cables are needed, requiring 124 modules to receive them.

A. From system to subsystems

The first stages of the trigger processing are entirely parallel. These consist of the analogue processing, digitization, bunch-cross identification and energy calibration of the trigger towers. It was therefore natural to design a module (the Preprocessor module) which performed all these in parallel, with the smallest size possible. The PPM was successfully designed to handle 64 input channels (16 towers from four cables). To miniaturize to this level required the design of an MCM containing the FADCs and an ASIC. Using these modules, a system of 8 crates of PPMs is needed for the final full scale system.

However, the physics object identification required algorithms to be applied in overlapping windows of 4x4 towers. For example, the electron/gamma/tau algorithm is shown in fig. 2. This type of algorithm necessitates information sharing between nearby towers, which would be impossible for the already complex preprocessor system. Therefore the application of physics algorithms to the digital data is separated out from the analogue to digital preprocessing. Another natural split occurs between the CP and JEP systems, as the JEP system needs lower granularity towers as input (0.2x0.2 rather than 0.1x0.1 in eta/phi). Thus the trigger splits naturally into three subsystems.

The disadvantage of the processor systems being separate from the preprocessor is that another data transmission stage is required. This is provided by LVDS links running at

400 Mbits/s. Along with the processing power required, the physical limit of the density of cables into the processor systems is one of the determining factors of the number of modules and crates needed. Another major determining factor is the density of signal transmission across the custom backplane in the CP and JEP systems.

B. Consequence of overlapping windows

The algorithms used in both the CP and JEP systems use a 4x4 array of inputs throughout the complete eta/phi array of towers. This means that for every tower processed, an 'environment' of towers around it is needed for the calculation. Each of these environment towers will also be processed, possibly in the same module or crate, but possibly not. Thus, as long as the processing has to be split into separate pieces of hardware, then some of the input information has to be duplicated into one or more other modules. Again the scale of the processing needed, and the physical volume of input cables means that a parallel system is necessary. However, care has to be taken not to parallelize too far, as the consequences of duplication of data may outweigh the advantages of the division.

A typical arrangement is illustrated in fig. 3. A single JEM module only fully processes a 4x8 matrix of core towers. In order to do this however, it requires a full environment of 7x11 towers. That is, almost 1.5 times the core number of towers have to be duplicated from data that is in the core region of another JEM.

The amount of fanout required rises quickly as the core region is reduced. Some examples are shown in table 1. The choice of size therefore becomes a compromise between the advantages of spreading the processing over more parallel processors, against the extra processing and connectivity needed to provide and route the increasing quantity of fanout




Fig. 2. Cluster Processing algorithm requiring a 4x4 tower array in two layers (electromagnetic and hadronic)

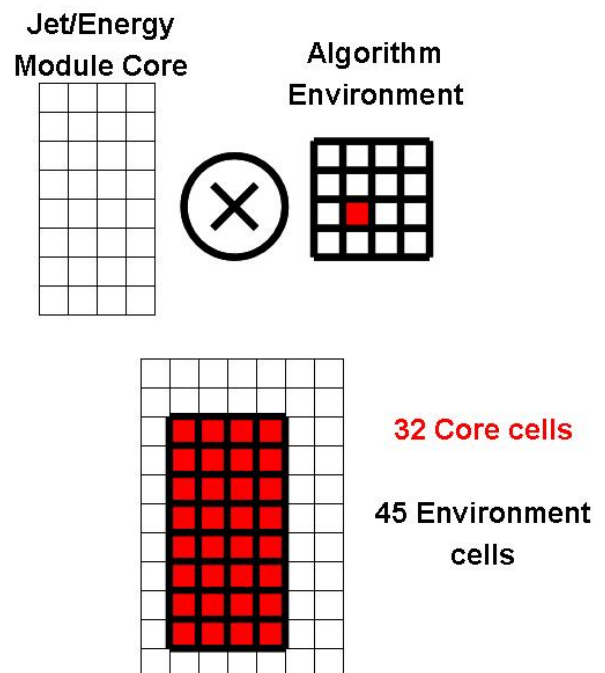


Fig. 3. Core and environment towers for an individual Jet/Energy Module

data. The final decisions made were a delicate balance of what was thought to be possible with the FPGAs and supporting circuitry at the time of finalizing the system design. As shown in fig. 2, the JEP system splits its 32x32 input array into 4x8 regions per JEM, requiring 32 modules which fit into two 9U VME crates. The CP system has higher granularity input over a 50x64 array, and each CPM fully processes a 4x16 core, so 56 modules are needed, fitting into four similar 9U VME crates. However, the 4x2 entry in the table is also relevant to the CPM, since internally, the processing is performed on 8 FPGAs each with a core area of 4x2. Clearly even more fanout of information to each FPGA is required to support this modularity, and routing of these tracks on the CPM means that this PCB has a very complex and dense internal connectivity [5].

TABLE I
DUPLICATION AS A FUNCTION OF CORE SIZE

Core Size	Core Towers	Environment Towers	Duplication, as a percentage of core size
4x16	64	69	108
4x8	32	45	141
4x4	16	33	206
4x2	8	27	338

C. Direct and Fanout Connectivity

As mentioned above, the tower data is sent to the processor crates on LVDS cables. Some of the necessary fanout is performed through this LVDS connectivity. The preprocessor system produces copies of data near the phi boundaries of the processor regions, and sends these copies to different processor modules. Thus fanout in one plane is achieved, and this accounts for about 25% of all the duplication required.

The remainder of the fanout and extra connectivity is performed in the processor systems themselves. Copies of 75% of the input towers are sent through a custom backplane to the neighbouring modules. The LVDS signals themselves also come through the backplane, so tower connectivity takes the majority of the real estate of the backplane. In all, these signals, with their associated grounds, account for the majority of the pins on the backplane and contribute to the total of about 1150 pins per CPM or JEM slot, and about 22,000 pins per backplane. The delicate, high density nature of the custom backplane is a direct consequence of the need to try to process as large a region as possible in each module, and it is clearly a critical item in the system design.

IV. INSTALLATION OF THE FULL SCALE SYSTEM

Some of the physical limitations imposed by cables have already been mentioned. Although the architecture was designed to accommodate these, it was not until installation began at ATLAS that real confirmation of the feasibility could be obtained. The cabling process itself was difficult and time-consuming, but is now almost complete.

The incoming analogue cables had to be cut precisely to length in order to achieve the joint goals of minimizing



Fig. 4. One quarter of the analogue cable inputs into the preprocessor system. Supporting stocks are located at about 2/3 of cable height

latency and providing a well structured cable management system, where it would be possible to remove individual modules or whole crates without the need for excessive re-cabling. This was done using a system of stocks to hold the heavy cables in place, and provide strain relief for the input connectors. A fraction of the analogue cabling infrastructure can be seen in fig. 4, showing the careful individual routing of each cable. Finding a suitable route for all of the input cables, including accommodating several patch panels for rearranging the data input mapping, stretched the rack and cable tray mechanics to their limits.

A similar challenge was encountered with the LVDS cable



Fig. 5. Full LVDS input cabling at the back of a JEP crate

routing. Both the provision of cable trays, and the access to the back of the processor crates around the crate infrastructure (power supplies, cooling, etc) proved to be just adequate for the mass of cables that are required to be routed to the back of the preprocessor crates, particular in the case of the JEP crates, which require more input data. An example is shown in fig. 5, where the top crate is a fully cabled JEP crate. Note that like the analogue cables, a system of strain relief is used to protect the cables and input connectors from the weight of the cables. Better solutions to the routing of these cables, using a different configuration of power supplies, could probably be found if long term reliability problems are found with this arrangement.

V. FULL SCALE CRATE TESTS

Apart from the mere physical challenge of the full scale cabling exercise, the electronics also needed testing a full scale. Before 2006, only small scale partial crate setups could be built. With the production runs of most of the major custom modules beginning in 2006, it became possible to construct full scale crate tests of each of the subsystems, and try them out in the ATLAS installation environment.

A. Preprocessor system

In some ways, the preprocessor system was the least interesting of the full crate tests, since the real time path of these modules works entirely in parallel. However it was important to verify that it was possible to run a full crate within the limits of the power supplies and cooling systems supplied, since the preprocessor MCMs are quite power hungry. This was demonstrated, and temperate profiles of the components could be derived from the on-board CAN monitoring information. This showed that all MCMs were running well within their tested operational temperature range.

With both the preprocessor and processor system, there were some initial problems with slowly oscillating power supplies as the crates were loaded with more modules. This was found to be due to a feedback mechanism caused by the power supply's voltage compensation mechanism using remote sensing, combined with the inductance of the power leads and capacitance of the modules. This has been overcome by better routing of thicker power leads, and the problem can always be overcome by changing to local sensing of the power supply.

B. Processor system

Running a full processor crate is also demanding in terms of energy, but more interestingly, it probes a new regime of system architecture that could not be fully tested before. In order for data transfer between the modules in a processor crate to take place without errors, all of the modules must be carefully timed in at appropriate phases, so that the data strobes occur at the correct time across the full crate. Beforehand, it was only possible to test the concurrent connectivity of a few modules, now the ability to perform the correct timing setup for a complete crate could be proved.

Not only do the CPM and JEM modules need to be able to transfer data between neighbouring modules, but they also all have to send results data to two Merger modules (CMM) in each crate. Thus the CMM must be able to cope with differing input phases, since the propagation time of signals along the backplane is significant at the data rates used. This could also now be tested with a crate full of modules.

Complete crates of both the CP and JEP systems were built in the custom crates at ATLAS, and the timing behaviour in both was found to scale well to the full system. However, problems were observed with a small minority of the fanin/fanout signals via the backplane. These problems were soon identified as being faults in the backplane, as described below.

C. Processor backplane

The custom backplane for the CP and JEP systems consists of a multi-layer PCB routing signals of several types between modules and also through to the back of the backplane into which connectors of various types are plugged. The slots reserved for CPM or JEM modules contain over 1150 pins, and in all the backplane comprises 22,000 pins. The vast majority of these pins are dedicated to LVDS input or fanout of tower signals between modules.

In backplane production, connectors with pins are pushed through the drilled holes in the PCB, either to make contact with the PCB tracks or to go through to the other side for the external connections. In a very small percentage of cases, one of these pins had missed the hole and instead been compressed under the connector itself. An example is shown in fig. 6. In these cases, clearly the pin would not make contact as required. In some cases the damaged pin also managed to short a nearby pin, thus compromising its function too. The frequency of these problems was such that almost all backplanes had at least one fault.

However, once the cause of the problems had been identified, it was possible to scan for problem pins in the laboratory, and incorrect connectors could be replaced. This process of repairing the backplanes is still ongoing, but so far the indications are that the replacement process works well, and fixed backplanes should function as new.




Fig. 6. X-Ray of a damaged backplane pin. The damaged pin can be seen as a curled up shadow behind the backplane

VI. INTEGRATION WITH EXTERNAL ATLAS SYSTEMS

In order to prepare for real data taking, a series of increasingly challenging integration steps with the rest of ATLAS has begun. All external interfaces have now been tested at some level, though typically only with a small subset of the final modules required.

A. Calorimeter input

The long cables carrying trigger tower signals from the ATLAS cavern to the electronics barrack have been laid during the course of the last year. This gives us access to genuine calorimeter trigger signals for the first time since the ATLAS test-beam in 2004. Without beam, there is only calibration or cosmic ray data to observe. The detector calibration systems are very useful for connectivity tests, initial timing setup and developing calibration techniques.

Signals have been seen from both calorimeters (the Liquid Argon Calorimeter[6] and the Tile Calorimeter[7]), testing both the PPMs, and the analogue chain that feeds them. More detailed work has been performed using the Tile calorimeter and its calibration system. Firstly the connectivity and mapping has been verified using tuned patterns of signals. Timing and energy calibrations have also been performed using a different, controllable set of pulses with differing energy. Typical results of such a run can be seen in fig. 7, where the pulse shapes can clearly be seen as a result of adjusting the FADC strobe in steps of 1ns. The pulse shapes seen are as expected, and the size and linearity of the response can be derived from these plots.

B. Trigger result output to CTP

The other fast real-time interface in the system is the output to the CTP. From each subsystem, this only consists of a handful of result bits at a signal speed of 40 MHz, thus this interface should not be a difficult challenge. Tests were performed with the output of one CMM connected to a CTP input module. Some minor problems were found, some due to faulty cables, and others due to small firmware bugs. All known problems have now been fixed, and the tests will be repeated soon.

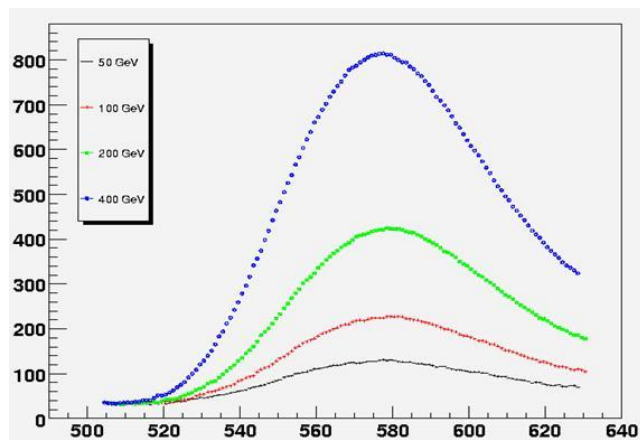


Fig. 7. Superposition of results of fine timing scan using calorimeter input pulses of different energies

C. Read-out data to ATLAS read-out system (ROS)

The standard route for read-out data in ATLAS is through a detector specific Read-out Driver (ROD) into a ROBIN on a ROS system via a Slink fibre[8]. The final Level-1 calorimeter trigger ROD has not yet entered full production, but comprehensive tests have been performed on the fully functional pre-production modules. One of these was to run known data through the ROD to a final ATLAS ROS (through the final installed fibres) at a high Level-1 accept rate, thus fully exercising the flow control feedback mechanism. The data gathered by the ROS was verified, and no corruption occurred over a long period.

Although this was a demanding test of the ROD functionality, it still requires testing in a larger integrated ATLAS environment to check all the features – for example the correct insertion of the bunch crossing number, and trigger type into the data.

D. Region of Interest data to the ATLAS RoIB

The Region of Interest Builder (RoIB) has been available at ATLAS for some time, and tests of the ROD to RoIB interface were also performed in a similar way to the ROS. Known data could again be run from the level-1 ROD, but this time it was routed through the RoIB before going into a ROS. Again, a check was made for data corruption and none was found. Using this technique, the flow control feedback fed from the ROS through the RoIB to the ROD, and so demonstrated that the throttling mechanism was also working via this route.

VII. CONCLUSION

Many of the challenges of installing the ATLAS Level-1 Calorimeter Trigger in situ have already been met. By mid-2007, much of the custom hardware should also be installed, with the remainder arriving later in the year. This should easily be in time for any beam collisions at LHC.

Many of the modules have been tested in their final configuration at ATLAS, and tests of larger scale setups than previously possible have been performed. No problems that require major rework have yet been found. Integration with all directly interfacing systems in ATLAS has been performed, again with no major problems found.

However, there is still much work to be done both in expanding the test setups internally to include large numbers of modules in each sub-system, and in starting to interface to larger parts, and different parts, of the calorimeters. Only a small part of the totality of the calorimeters has yet been explored, and more work will be required to fully understand and calibrate these signals.

It is expected that the Level-1 Calorimeter Trigger will soon join with full ATLAS integration runs and possibly cosmic data taking exercises. Though the cosmic ray signals may be of little direct interest, being small and unsynchronized, it will still be a useful learning experience of integrating with the rest of the ATLAS system.

ACKNOWLEDGMENT

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.

REFERENCES

- [1] R. Achenbach et al., *Pre-production Validation of the ATLAS Level-1 Calorimeter Trigger System*, IEEE Trans. Nucl. Sci., Vol. 53, p859-863, 2006
- [2] P.B. Amaral et al., *The ATLAS Level-1 Central Trigger System*, IEEE Trans. Nucl. Sci., Vol. 52, p1217-1222, 2005
- [3] ATLAS Level-1 Trigger Group, *ATLAS First Level Trigger Technical Design Report*, ATLAS TDR-12, CERN/LHCC/98-14, CERN, Geneva, 1998
- [4] J. Garvey et al., *The ATLAS Level-1 Calorimeter Trigger Architecture*, IEEE Trans. Nucl. Sci., Vol. 51, p356-360, 2004
- [5] J. Garvey et al., *Use of an FPGA to Identify Electromagnetic Clusters and Isolated Hadrons in the ATLAS Level-1 Calorimeter Trigger*, Nucl. Instrum. Meth., A512, p506-516, 2003
- [6] P. Schacht et al., *The ATLAS Liquid Argon Calorimeter: Status and expected performance*, Nucl. Instrum. Meth., A535, p466-471, 2004
- [7] P. Adragna et al., *The ATLAS Hadronic Tile Calorimeter: From construction towards physics*, IEEE Trans. Nucl. Sci., Vol. 53, p1275-1281, 2006
- [8] H.P. Beck et al., *The base-line DataFlow system of the ATLAS Trigger and DAQ*, IEEE Trans. Nucl. Sci., Vol. 51, p470-475, 2004