Commissioning of the ATLAS High Level Trigger with Single Beam and Cosmic Rays

A. Di Mattia on behalf of the ATLAS Collaboration
Michigan State University - Department of Physics and Astronomy
3218 Biomedical Physical Science - East Lansing, MI 48824-2320

E-mail: dimattia@mail.cern.ch

Abstract. ATLAS is one of the two general-purpose detectors at the Large Hadron Collider (LHC). The trigger system is responsible for making the online selection of interesting collision events. At the LHC design luminosity of 10^{34} $cm^{-2}s^{-1}$ it will need to achieve a rejection factor of the order of 10^{-7} against random proton-proton interactions, while selecting with high efficiency events that are needed for physics analyses. After a first processing level using custom electronics based on FPGAs and ASICs, the trigger selection is made by software running on two processor farms, containing a total of around two thousand multi-core machines. This system is known as the High Level Trigger (HLT). To reduce the network data traffic and the processing time to manageable levels, the HLT uses seeded, step-wise reconstruction, aiming at the earliest possible rejection of background events. The recent LHC startup and short single-beam run provided a "stress test" of the system and some initial calibration data. Following this period, ATLAS continued to collect cosmic-ray events for detector alignment and calibration purposes. After giving an overview of the trigger design and its innovative features, this paper focuses on the experience gained from operating the ATLAS trigger with single LHC beams and cosmic-rays.

1. Introduction

ATLAS [1] is a general purpose detector built for collecting data at the Large Hadron Collider (LHC), the new accelerator machine of CERN that will produce p-p collisions with a centre of mass energy of 14 TeV at a luminosity of 10^{34} cm^{-2} s^{-1}. The design parameters of LHC will allow ATLAS to probe the phenomenology of the fundamental interactions up to the energy scale of 1 TeV. But, at the same time, will expose the detector to a harsh experimental environment. At the peak luminosity the background processes will produce an interaction rate of 10^9 Hertz, which must be reduced online by seven orders of magnitude in order to obtain a final event rate manageable in terms of storage space and offline processing.

The ATLAS trigger system achieves this rate reduction by using a multi-level architecture where the event selection decision is split into different stages (Trigger Levels) of increasing complexity and processing time. The Level-1 trigger [1], uses algorithms implemented on custom hardware boards which process only a subset of the event data (coming from the Muon Spectrometer [1] and the Calorimeter [1]) with coarse granularity. The output rate of the Level-1 is designed to be approximately 75 KHz, but upgradeable to 100 KHz, and the latency time is fixed at 2.5 μs. The High Level Trigger (HLT) [2] provides a software based event selection to further reduce the Level-1 rate to

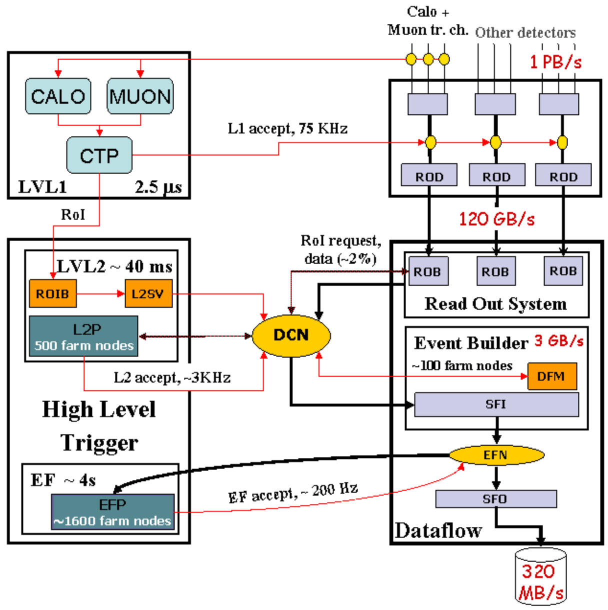

Figure 1. Schematic diagram of the ATLAS Trigger and DAQ system

approximately 200 Hz. At high luminosity, this event rate correspond to the rate of $pp \to W^+ \to l^+ \nu$, which shows how critical the role of HLT is in identifying interesting physics.

After the initial Level-1 selection, the event data from the various sub-detectors is held in separate memory buffers and waits for a decision to determine if it is assembled into a complete event for the final selection or discarded. This intermediate trigger step of the HLT, called Level-2, analyzes the data using fast algorithms, which perform an approximated reconstruction of the physics candidates. The trigger algorithms make use of the full detector data but access only the detector regions that have been flagged by the Level-1 as those containing the physics candidates. These regions are called Regions of Interest (RoI). This RoI-based access requires only 2% of the event data to make the Level-2 decision, thus limiting the required bandwidth of the dataflow; in contrast, to move the full event data into the processing farm would have required a data throughput that reaches the performance limit of the current commercial equipment. The trigger design foresees that the Level-2 will reduce the Level-1 rate by approximately a factor of 30 with a latency of approximately 40 ms.

The data, after being accepted by the Level-2, are assembled into the full event and delivered to the Event Filter to make the final decision on the event. The algorithms employed at this stage are the same as those used in the offline reconstruction software and make use of the fully detailed detector condition data. The increase of the reconstruction accuracy provides a decrease in the Level-2 accept rate by about a factor 10. But the use of a more sophisticated procedure requires more time to process the events, with the approximated latency of the Event Filter estimated to be about 4s.

A diagram of the ATLAS Trigger/DAQ is shown in figure 1. The detector data flows through the approximately 1600 Read Out Divers (RODs) into the Read Out System (ROS). This is undertaken by 150 PCs hosting several ROBIN boards which buffer the data within the Level-2 decision. The event RoI regions are extracted from the Level-1 data by the RoI Builder (ROIB) and then sent to a Level-2 Supervisor (L2SV), which assigns the event to a Level-2 Processing Unit (L2PU). The L2PU is the trigger processor to be executed in the Level-2 processing farm, which contains about 500 processing nodes. In the case of a “Level-2 accept” the result is stored into the pseudo ROS (pROS) as part of the event data and then is sent through the Event Builder and the Sub Farm Input (SFI) to the “Event Filter farm”, together with the ROS’s data under the supervision of the Data Flow Manager (DFM). A pool of approximately1600 processing nodes implements the Event Filter farm, where the Processing Task (PT) is executed to make the final decision on the event.

2. The High Level Trigger system

The HLT provides software based selection through parallel online processing on thousands of nodes. The system stability depends on the interplay among the performance of the trigger algorithms, the selection menu, the computing resources, and the data transfer to the Level-2 and Event Filter farms. The optimal working parameters depend upon the evolution of the system size and the experimental conditions. Estimates made on large scale prototypes take into account assumptions on performance scalability that may not be reliable.

The HLT selection software is similar to an offline analysis in that reconstruction algorithms with good physics performance are crucial to achieve the required rate reduction. This reflects the complexity of the trigger configuration and requires the online implementation of functionalities, like the access to updates of detector condition data, typically needed by the offline code.

Before data taking, the overall performance of the HLT was constantly checked with dedicated technical runs [3,4] where Monte Carlo events were injected into the dataflow and processed by the HLT configured for physics selections. It was possible to assess the compliance of the system with respect to the design parameters (latencies, timeouts, etc.) and to test the required operational functionalities such as event streaming and monitoring. The system used in the data taking period of 2008 is summarized in the next subsections.

2.1. The hardware

Apart from the Level-1 hardware, the RODs and the RoIB, all the components of the ATLAS Trigger and DAQ system consists of standard PCs running a Scientific Linux distribution with several add-ons by CERN. Currently, machines with two quad-core 2.5 GHz processors are employed. HLT is designed to exploit the full potential of these multi-core processors by running a trigger application (L2PU or PT) on each core. Although the trigger algorithms use the same condition and configuration data within a run, memory sharing among the trigger applications is not possible because the software does not support multi-threaded operations. Therefore the machines are equipped with 2 GBytes memory per core, which is enough to contain the full memory usage of a trigger application. During the 2008 data taking 850 processing nodes, hosted on 27 racks, were operated to implement the HLT farms. This corresponded to 35% of the final HLT system. Monitoring of machine power status and remote machine rebooting is done via IPMI, an additional board connected to each machine and always powered on. The machine resources and the operational status are monitored remotely via a custom system based on NAGIOS [5]. The partitioning of the farm is fully configurable, as within a day it is possible to change the assignment of the computing resources to Level-2, Event Filter and DataFlow components in order to cope with unexpected increase of the workload in one of these systems. This has been successfully tested in runs for the calorimeter calibration, where the event size was increased by a factor of seven with respect to the design size.

2.2. The Event Selection Software

The execution of the Event Selection software is undertaken by the L2PU and by the PT, which are final state machines that respond to the commands from the run control, the application which manages the DAQ operations. The software framework [4] of these applications was built by reusing part of the offline software components in order to allow for operating both the Level-2 and the Event Filter algorithms in a common environment. In particular all the interfaces to the data and the code providing the detector condition data (e.g. alignment of geometry and calibration) are implemented by the same offline tools. This eases the development and the study of the selection algorithms, optimizing the manpower and increases the long term maintainability of the code. However the use of pure offline components in the online environment sometimes conflicts with the latency limit, thus requiring their replacement with highly optimized code, especially at Level-2. Specific software tools are also needed to allow for the update of condition data only before the start of run transition (i.e. no database access during the event processing). Moreover the use of databases for configuring the selection software in a large parallel system such as the HLT poses a serious problem to the overall boot-up time when thousands of applications access the database at the same time. In order to avoid a direct connection from every application, a proxy mechanism has been put in place on every farm rack, which caches the meta data request coming from the processing nodes in the rack.

The trigger selection proceeds in steps, called selection sequences, where a Feature Extraction algorithm reconstructs the physics candidate and a Hypothesis algorithm validates it. The execution of the selection sequences is done by the HLT steering, which processes them in a trigger chain with the aim to define the full properties of a candidate entering the trigger menu. The menu specifies not only the physics selection, hence the sequences and the configuration of the related algorithms, but also the processing order of the sequences into the trigger chain. The latter is crucial for optimizing the overall trigger latency and is chosen to prioritize the fastest algorithms because the steering aborts the processing of a chain as soon as a selection sequence step fails (early rejection). Further optimization of the trigger latency is performed by the HLT Steering code, which avoids running the same algorithms on a given RoI when there is an overlap among trigger chains. The HLT steering also allows for applying pre-scales and pass-through factors to the trigger chains in order to keep the event rate under control.

The complete ATLAS trigger menu contains up to a hundred selection chains, thus requiring the handling of a thousand parameters for configuring the event selection software. Because these parameters, especially the pre-scale and pass-through factors, are updated constantly to follow the evolution of both the run condition and the detector status, a bookkeeping of the full configuration is needed to keep track of all the changes made. This is mandatory for reproducing the online running conditions of the HLT when it runs offline. The configuration data are stored in the Trigger Database and the bookkeeping is implemented through four database keys, each one referring to particular sets of configuration parameters. The database also performs consistency checks of the trigger menu, assessing that there are no HLT chains orphaned from LVL1 menu item. A user interface, the Trigger Tool, has been also developed and provides access functionalities according to account rules, which are different for shifters, experts and offline users.

After two months of cosmic data taking, 200 database keys and 10000 configuration parameters, mainly concerning algorithm settings and pre-scales, were required to follow the evolution of the detector condition and of the LVL1 trigger rate. The database was vital to keep track of all the changes.

2.3. The trigger monitoring

Online monitoring of the trigger selection is done with histograms in order to assess the quality of the trigger reconstruction and to show the event flow in the trigger chains. The histograms for the trigger quality are implemented independently in each trigger algorithms, being strictly related to the reconstruction steps, while those for the trigger rate are provided centrally by the HLT steering. The statistics related to the trigger applications are published via the Information System into the data control network and are collected by the Gatherer [6] applications that concentrate and serve it to the online monitoring tools. The number of histograms to be published has been reduced to the minimum needed in order not to overload the task of the Gatherers. This also eases the tasks of non expert shifters.

The Trigger Presenter shows the instantaneous trigger rate for each Level-1 trigger item and for each HLT selection chain. This tool was used by the shifter to trace the origin of an increase in the event rate and to react by adjusting a pre-scale factor. Left side of figure 2 shows an example of history plot of the trigger rate. The algorithm monitoring is done via the Online Histogram Presenter [7] that also performs automatic checks on the histograms for data quality assessment. Since the HLT is the first place in the online system where data from different detectors are combined and reconstructed, trigger monitoring may help to understand the detector behaviour before the offline data quality monitoring. An example of online monitoring histograms made by the Level-2 muon trigger is shown on the right side of figure 2: it plots the distribution of the residuals between the muon trigger chamber hits and the muon precision hits. This histogram is sensitive to the timing synchronization of the two detectors and

Raw 16% 16.37
Prescale 16% 16.37
Forced Acc 16% 16.37

Figure 2. Left: history of the Event Filter trigger rate shown by the Trigger Presenter. Right: display of a monitor histogram from the Level-2 muon trigger.

provided a lot of feedback to those monitoring the detector during the early phase of the test combined runs.

The HLT selection is also monitored offline by matching the physics objects reconstructed by the trigger algorithms with the offline result. This happens at the CERN Tier0 cluster, where the events recorded are reconstructed by the offline code within one day of data taking. More detailed monitoring and debugging is done at the CERN Analysis Facility, which is designed to reprocess 10% of the recorded event, including re-running the Event Selection code. Here, the quality of the HLT decision can be assessed by reprocessing the events in the minimum bias stream or events taken in pass-through mode. The coherence among the offline and the online trigger processing is also checked.

3. The experimental environment

In 2008, the data taking for single beam and cosmic was used to commission the HLT. These two data taking methods had different purposes and required the trigger to be operated in different ways. Also the physics was different for each case and this was reflected in the HLT achievements. Full details about the Level-1 trigger in operation are given in [8].

The LHC machine was operated to provide several kinds of beam, yielding different activity in the detector. The highest activity was provided by the so called “beam-splash” events. These originate from the collision of the proton beam with the tertiary collimators, placed 140 meters away from the ATLAS interaction point. When the collimators were in the closed position a large number of charged particles from the beam crossing flooded the ATLAS experimental hall. Among the Level-1 subsystems used for triggering on these events the main contribution to the trigger rate came from the Minimum Bias trigger, implemented using scintillators put on the external surface of the calorimeter cryostat. The LHC machine also provided stable circulation of beam for a maximum period of half an hour. In order to improve the trigger efficiency and to get rid of the cosmic trigger rate the Beam Pick Up [9] trigger, an electrostatic detector placed 170 meters away from the ATLAS experimental hall, was employed. Although this mode didn’t yield a significant activity in the detector, it allowed for the stress test of the Level-1, as a Level-1 rate up to 11 KHz could have been sustained.

The priority for this data taking mode was to synchronize the Level-1 subsystems and to have the detector readout well timed in. In order to improve the robustness of the first data taking with beam, the HLT selection algorithms were not operated online, and the HLT was used online only for event streaming based on Level-1 information. The plan was to perform HLT commissioning offline by reprocessing the subset of beam-related events containing muon or calorimeter RoIs. The trigger code proved to be well prepared for the event data reprocessing as the first feedback from the HLT was available within hours of the first beam related activity in ATLAS. Unfortunately, due to the short operation period and to the low efficiency for firing RoIs in the single-beam events (the tracks were non pointing to the nominal interaction region) only about a thousand events were available to be

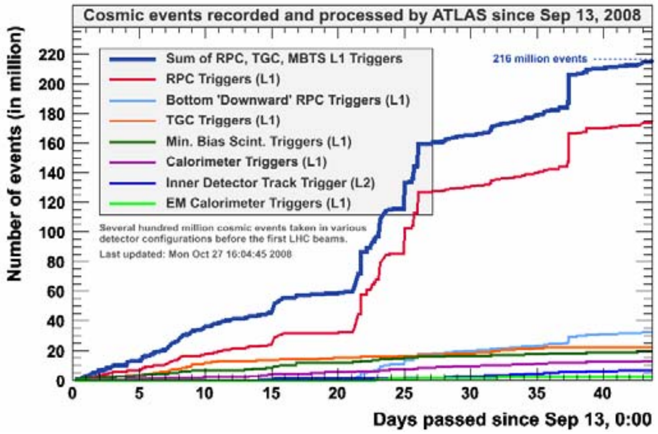

Figure3. History of the event data taking since the LHC accidents.

reprocessed by HLT. Thus only limited commissioning work could be accomplished with the data from the single beam running in September 2008, although much useful experience was obtained.

After the LHC accident, data taking was reverted to cosmic running and detector commissioning. In cosmic data taking mode, events were mainly selected by the Level-1 muon trigger and accepted by the HLT with a pass-through trigger. Problems in the event recording and handling at Tier0 were avoided by controlling the trigger rate with pre-scale factors. Since the rate provided by the muon barrel trigger was approximately a factor of 100 bigger than that provided by the endcap trigger, these two triggers were handled separately by the Central Trigger Processor [8]. This allowed for different pre-scale factors on the two triggers. The HLT Cosmic Menu was run online in parasitic mode and some of the algorithms played a role selecting events suitable for the calibration and alignment of the detectors. The physics selection menu was also run in parallel. Figure 3 shows the number of the events recorded since the LHC accident on September 19. The biggest event sample was selected by the Level-1 Muon barrel trigger, but there has been a small contribution from the calorimeter trigger and the Level-2 data stream that used the Level-2 tracking algorithms. The large change in the slope signifies the end of the data taking for the calorimeter calibration, which caused the reduction of the event size of a factor of 7, thus allowing lowering the pre-scale factors applied to the Level-1 muon trigger.

3.1.1. Issues on cosmic data

The ATLAS Muon Spectrometer is built of three measurement stations (inner, middle and outer), equipped with trigger and precision chambers, that allows for both the trigger and the p_T measurement of the muon track. In normal operation mode, the Level-1 muon trigger fires on trigger hit patterns compatible with a track heading towards the ATLAS Interaction Point. The trigger hits are provided by the Resistive Plate Chambers (RPC) [1] in the barrel and the Thin Gap Chambers (TGC) [1] in the endcap, which are typically adjacent to the precision chambers, the Monitored Drif Tube (MDT) [1]. This is exploited by the HLT muon tracking algorithms, which search for MDT hit clusters associated with trigger hits. On the contrary, the selection of the cosmic muons was performed unbiased in the longitudinal detector plane (r-z view), because unbiased tracks were required for computing the alignment of both Inner Detector and Muon Spectrometer. Figure 4 shows a diagram of the RPC trigger towers and how the unbiased selection logic works. Typically a selected track misses the high-p_T portion of the trigger tower and will be likely rejected by the Level-2 muon tracking algorithm because it fails to reconstruct the precision hits along the full Muon Spectrometer. The efficiency on this type of tracks can be partially recovered by loosing the pointing requirement, but hits left outside the RoI are lost for reconstruction. This effect is even more dramatic when considering the Level-2 tracking for the Inner Detector, since the track may completely miss the Silicon Detectors.

With respect to physics collision tracks, a cosmic event crosses the upper part of the detector with a reverse path and has a limited time resolution because the timing information comes from the muon trigger chambers and not from the beam clock. Therefore, a looser timing of the detector readout latency had to be used in order to maximize the fraction of the cosmic muon hits in the event data. Bad event timing, together with the poor knowledge of detector calibration, caused unphysical conversion of the time measurement of the drift detector in some event.

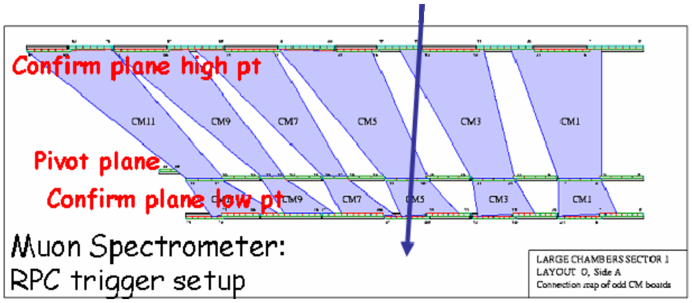

Figure 4. Schematic diagram of the trigger coincidences implemented on the Barrel RPC station.

4. The commissioning work

The HLT system operates on detector data so it is affected by any change in the experimental environment. Since it handles the DAQ of the experiment the system must be very stable and well assessed before operating it in the data taking environment. At the same time, the system is required to be flexible and responsive to requests coming from the detector subsystems that may need the help of the HLT to accomplish their duties (detector calibration, alignment).

During the 2008 data taking some of the detectors components were still in the commissioning phase. This constituted a formidable stress test for the full software infrastructure of the HLT, which had to cope with atypical operation of the Level-1 trigger and with incomplete detector readout. In such an environment, increasing the sensitivity on the detector condition of the trigger monitoring was mandatory in order to understand the source of the problems. An incorrect software description of the detector readout and cabling could cause errors in data access, especially in Level-2 which is RoI driven, mimicking missing detector data. At the same time, the trigger reconstruction may fail because of the presence of a bad or noisy detector module in the RoI region. The robustness of the code with respect to data corruption and missing data was also tested. At the beginning of data taking the HLT failure rate was about 1% of the total events. Continual improvements in the code allowed for the reduction of this rate by more than factor of 10.

In total, 23 full time equivalent people, providing expertise in many areas, were required to operate the HLT during the data taking of 2008. The synergy of so many was crucial for the success of the infrastructure commissioning. The trigger algorithms also took advantage of this data taking period, as some of their functionalities could have been tested with cosmic events.

4.1.1. Running the e/γ trigger slice

In the cosmic running, the Level-1 calorimeter trigger was operated with a very low energy threshold in order trigger on cosmic muons that lose energy in the Landau tail region of the energy absorption spectrum. Both Level-2 and Event filter algorithms were run on a sample of few thousand events, where the muons faked an e/γ signature. The trigger reconstruction discriminated against these energy clusters with high efficiency, and fakes occurred mainly because of the very low energy threshold applied by the Level-1 and because of poor pedestal calibration. Perfect agreement was found between the Level-2 and the Event Filter results. Despite the low event statistics, the e/γ trigger was also seen to be sensitive to the detector status. Some failure of the Level-2 cluster reconstruction appears in correspondence with hot spots in the detector, as confirmed by the detector data quality studies. Full details about the operation of the e/γ trigger in the 2008 data taking can be found in [10].

4.1.2. Running the Muon trigger slice

The muon trigger slice [11] reconstructs muons in the full ATLAS detector. However, in the cosmic menu, the HLT muon trigger was setup to have selection chains perform muon track identification only in specific subdetectors. This allowed the use of most of the muon trigger algorithms, even when a bad condition in a subdetector caused one of the algorithms to fail. MuFast performs the Level-2 standalone reconstruction, is present in all the selection chains, and was always executed online in every run. MuFast also provides the parts of the data necessary for the calibration of the full muon spectrometer in a custom data stream. This procedure, embedded within the track reconstruction, was successfully used in all the cosmic runs [12].

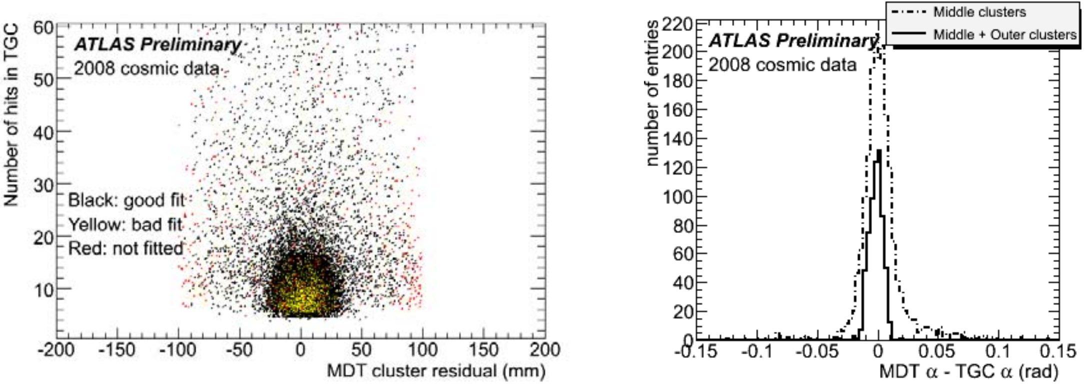

Figure 5. MuFast performance for endcap cosmic events. Left: distribution of the residual, with respect to the trigger seed direction, of a selected MDT cluster as a function of the number of trigger hits. Right: difference between the polar angle (α) of the trigger data fit (TGC) and the polar angle of the precision data hit (MDT).

The physics performance of the muon trigger slice was not able to be assessed using the cosmic events. The lack of pointing tracks posed a problem not only to the pattern recognition but also to the estimate of the muon p_T, which is performed in Level-2 using pre-calculated look-up tables derived from Monte Carlo. This also inhibited the combination of Muon Spectrometer data with Inner Detector data to form a complete muon track because the track extrapolation procedure also makes use of look-up tables at Level-2.

At Level-2, the reconstruction of MDT clusters and the measurement of the track bending performed by MuFast were tested with cosmic data. In figure 5, the x-axis of left plot shows the distance of the precision chamber clusters from the track reconstructed with the trigger data and the y-axis shows the number of trigger hits found in the TGC RoI. Despite the fact that the cluster reconstruction is less accurate when seeded by a large number of trigger hits, as happens for a non-negligible fraction of events, the efficiency provided is close to the design goal of 99%. In the study shown, the efficiency was 93%. It was less than the design efficiency because of bad detector conditions. A 2% inefficiency originated from the bad calibration of some MDT tubes, which yields an unphysical conversion of the MDT drift time into space, while the remaining 5% originated from chambers taken out from the readout or having noisy readout module. As with the e/γ trigger, the Level-2 muon is sensitive to the detector status.

Another measurement used in the Level-2 tests is the polar angle, α, defined as the angle between the slope of the linear fit performed with Muon Spectrometer data in the bending plane (r-z view) and the vector pointing from the ATLAS interaction point to the track. The relative alignment between the trigger and the precision detector can be checked by matching two measurements of this polar angle; one performed with the trigger data and the other performed with the precision data. The right plot of figure 5 shows the distribution of the difference between the two α measurements obtained selecting the data sample with different pattern recognition criteria. The measurement made using only the data from the MDT middle station does not provide a good match. The presence of a reconstructed cluster in the outer station allows for the refinement of α, and a performance similar to the Monte Carlo prediction to be achieved, but it is limited by the knowledge of the misalignment of the MDT detector. Since α is used to measure the track bending in the endcap, this also indicates that good performance can also be achieved for the muon p_T estimation.

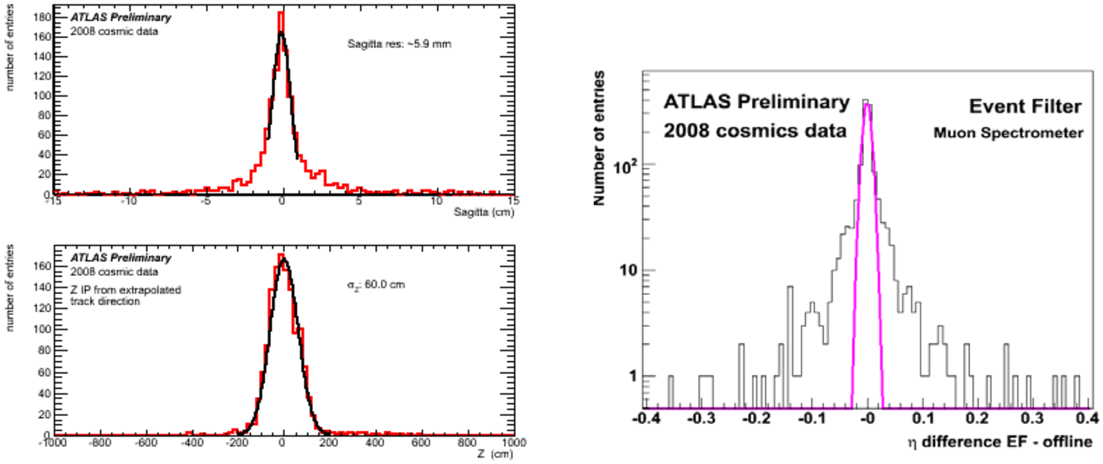

Figure 6. Left top: muon track sagitta reconstructed by MuFast in the barrel from data taken without magnetic field. Left bottom: intersection of the reconstructed muon track direction with the beam line. Right: offline match of the pseudorapidity η reconstructed by TrigMuonEF.

In the barrel, MuFast uses the sagitta measurement to estimate the muon p_T. The top left plot of figure 6 shows the sagitta reconstructed from three MDT clusters using data from a run taken without magnetic field. The performance obtained (5.6 mm) constitutes the intrinsic limit of the p_T estimation that can be achieved by the MuFast algorithm. This limit, as in the endcap, is dominated by the misalignment of the MDT detector, but is expected to be reduced approximately to 1 mm once the alignment correction is used at Level-2. The present limit is negligible with respect to the expected resolution for tracks below 11 GeV in p_T. The bottom left plot of figure 6 shows the intersection of the direction of the reconstructed track with the beam line. This quantity is distributed around the ATLAS interaction point, demonstrating the ability of pattern recognition to select pointing tracks.

The Level-2 calorimeter algorithm that searches for low-p_T muons was tested on a run taken with no magnetic field. In order to reduce the effect of the electronic noise, the energy cut in the calorimeter cells was put to 300 MeV. Despite the low reconstruction efficiency due to the lack of track pointing the performances were consistent with the Monte Carlo predictions. The reconstructed tracks of cosmic muons had an up-down distribution as expected for pointing tracks able to pass through the full calorimeter system. The measurement of the energy deposition peaked at approximately to 2.5 GeV, which is consistent with the energy loss of a minimum ionizing particle.

The performance of the muon Event Filter was checked by comparing the reconstruction of the track position with the output of the offline reconstruction. The right plot in figure 6 shows the result of this comparison as a function of the for the pseudorapidity η. Altough this Event Filter algorithm is the same as the one used in the offline reconstruction, some non-gaussian tails in the reconstruction can be observed. These are most likely due to different settings for handling the RoI-based strategy with respect to the offline reconstruction.

4.1.3. Running the Inner Detector trigger reconstruction

As opposed to for the Muon System, less than 3% of the track selected by the Level-1 muon trigger could be used for the Inner Detector alignment procedure. Most of the tracks missed the silicon detector and this forced the implementation of a dedicated selection stream that used the Level-2 Inner Detector tracking for selecting tracks that passed through the full Inner Detector volume. The standard tracking algorithms provide efficiency for reconstructing tracks with transverse impact parameter less than 5 mm. The track for the alignment is not required to have a small impact parameter because no bias on the direction is preferable. Therefore significant modifications were implemented on the pattern recognition of the tracking algorithms in order to get rid of the pointing requirement. The trigger reconstruction was setup to select top-bottom tracks, which leave at least 3 silicon hits on both the lower and the upper halves of the detector. The data from noisy detector modules had to be identified and removed from the reconstruction, as they yield a lot of fakes lowering the efficiency for the good track selection.

Three different algorithms SiTrack, IDscan and TRT Segment finder have been used for this selection, as they provided complementary efficiency for different values of the track impact parameter and p_T. The combined efficiency of all three algorithms was better than 99% for all cosmics passing through the Inner Detector barrel volume and approached 100% for high p_T. The fake selection was less than 1% after the removal of the data from noisy modules. The HLT selection stream mode of these algorithms has been used online in the cosmic running and collected more than 300 million of cosmics providing high quality tracks for use in the Inner Detector alignment. Further details about the use of the Inner Detector trigger in cosmic data taking are given in [13].

5. Conclusions

The HLT was fully tested, as well as some of the algorithms that contributed to the cosmic event selection. The operation mode provided a good balance between stability and responsiveness to detector requests. The successful running in 2008 has provided invaluable experience, and improvements made based on this experience will ensure ATLAS is better prepared for running with collisions at the end of 2009.

6. References

[1] The ATLAS Collaboration, G. Aad et al., “The ATLAS Experiment at the CERN Large Hadron Collider”, JINST 3 (2008) S08003.

[2] ATLAS High-Level Trigger Data Acquisition and Controls Technical Design Report, CERN/LHCC/2003-022 (2003).

[3] A. dos Anjos, “The DAQ/HLT system of the ATLAS experiments”, proceedings of the *12th International Workshop on Advanced Computing and Analysis techniques in Physics Research*, Erice, Italy 2008 (ACAT 08).

[4] W. Wiedenmann, “The ATLAS Online High Level Trigger Framework: Experience reusing Offline Software Components in the ATLAS Trigger” this proceedings.

[5] A. Zaytsev, “System Administration of ATLAS TDAQ Computing Environment”, these proceedings.

[6] P. Renkel, “The Gatherer - a mechanism for integration of monitoring data in ATLAS TDAQ”, these proceedings.

[7] A. Dotti, “The Online Histogram Presenter for the ATLAS Experiment: a modular system for histogram visualization”, these proceedings.

[8] T. Pauly, “The ATLAS Level-1 Central Trigger System in Operation”, these proceedings.

[9] C. Ohm, “The ATLAS beam pick-up based timing system”, these proceedings.

[10] D. Damazio, “Atlas High Level Calorimeter Trigger Software Performance for Cosmic Ray Events”, these proceedings.

[11] A. Ventura, “The Muon High Level Trigger of the ATLAS Experiment”, these proceedings.

[12] A. De Salvo “ATLAS MDT remote calibration centers”, these proceedings.

[13] M. Sutton, “Commissioning of the ATLAS Inner Detector Trigger”, these proceedings.