PROCEEDINGS OF SCIENCE

The ATLAS High Level Trigger Configuration and Steering Software: Experience with 7 TeV Collisions

Simon George*
Royal Holloway, University of London, Egham, TW20 0EX, UK
E-mail: S.George@rhul.ac.uk

On behalf of the ATLAS Collaboration.

In 2010 ATLAS has seen the first proton-proton collisions at 7 TeV. Later this year a collision rate of nearly 10 MHz is expected. Events of potential interest for physics analysis are selected by a three-level trigger system, with a final recording rate of about 200 Hz. The first level (L1) is implemented in customised hardware, the two levels of the high level trigger (HLT) are software triggers. The selection is described by the Trigger Configuration in the form of menus, each of which contains more than 500 signatures. Each signature corresponds to a chain of algorithms which reconstruct and refine specific event features. The HLT Steering receives information from the Configuration system, dynamically creates chains and controls the execution of algorithms and flow of information during event processing. The Steering tests each signature on L1-accepted events, and those satisfying one or more test are recorded for later analysis. To save execution time, the Steering has a facility to cache results, avoiding later recalculation. To control rate, prescale factors can be applied to L1 or HLT signatures. Where needed for later analysis, the Steering has a test-after-accept functionality to provide the results of the tests for prescaled signatures. To maintain high data taking efficiency it is essential that the trigger can be dynamically re-configured in response to changes in the detector or machine conditions, such as the status of detector readout elements, instantaneous LHC luminosity and beam-spot position. This relies on techniques that allow configuration changes, such as L1 and HLT prescale updates, to be made during a run without disrupting data taking, while ensuring a consistent and reproducible configuration across the entire HLT farm. We present the performance of the steering and configuration system during collisions and the expectations for the first phase of LHC exploitation.

35th International Conference of High Energy Physics - ICHEP2010,
July 22-28, 2010
Paris France

*Speaker.

© Copyright owned by the author(s) under the terms of the Creative Commons Attribution-NonCommercial-ShareAlike Licence. http://pos.sissa.it/

ATL-DAQ-PROC-2010-053
30 November 2010

PoS(ICHEP 2010)487

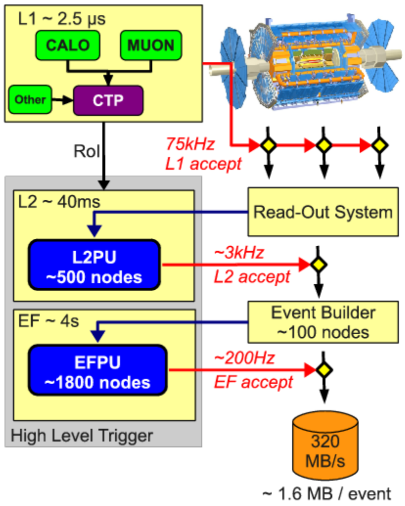

Figure 1: The ATLAS trigger with final design specifications for the rates, latencies and system sizes.

1. Introduction

During 2010, the Large Hadron Collider (LHC) at CERN has produced 7 TeV proton-proton collisions with increasing intensity. At the time of this conference in July, a peak instantaneous luminosity of 1.6 x 10^{30} $cm^{-2}s^{-1}$ had been achieved, and an integrated luminosity of 357 nb^{-1} delivered to ATLAS. These collision events were then detected in the ATLAS detector and recorded through the Trigger and Data Acquisition (T/DAQ) system.

ATLAS [1] is a general purpose detector with a diverse physics programme. With the LHC configuration of July 2010, a trigger is needed to reduced the collision rate from tens of kHz to 200 Hz. The ATLAS Trigger is illustrated in figure 1.

Level 1 is built from custom electronics which analyses data from calorimeter, muon, and minimum bias detectors and combines the results in the Central Trigger Processor. If an event is accepted, all the detectors are read out and buffered in the read-out subsystem. The positions in the calorimeter and muon detectors are sent to Level 2 as 'Regions of Interest' (RoI). Guided by these regions, the Level-2 trigger performs partial reconstruction of events using specialised software algorithms. Full reconstruction beyond the RoIs is performed in a small fraction of events. Events accepted by Level-2 are then fully assembled by the Event Builder and passed to a third level 'Event Filter' where reconstruction is performed using software tools drawn from the offline reconstruction software. Level-2 and the Event Filter are known collectively as the High Level Trigger (HLT). HLT software runs on commodity computers, rack mounted in a large 'farm' housed in one of the surface buildings above the ATLAS detector. Usually one instance of the software is run per CPU core. Figure 1 shows the final design specifications for the rates, latencies and system sizes. In July 2010 about 35% of the HLT computers were installed and the Level-1 output rate was kept at around 20 kHz. Further details of the ATLAS trigger performance and algorithms are given in [2].

This paper will describe some key elements of the ATLAS trigger and then focus on some of the new features introduced in the last year to maximise data taking efficiency.

PoS(ICHEP 2010)487

2. Trigger Configuration and High Level Trigger Steering

The software framework of the HLT, known as the 'Steering', and the 'Trigger Configuration', have previously been described in [3]. They were designed to provide a dynamic, flexible system that can cope with a wide range of operational scenarios. A brief overview follows.

The Trigger Configuration provides a coherent description of the way the entire trigger system is set up, both Level-1 and the HLT. It compromises a database, software services to access it, and tools to view and modify the configuration. This paper concentrates on the HLT configuration. At the most abstract level the trigger is defined in terms of a 'menu', which is a list of physics-like objects with multiplicities and thresholds. For example, 'mu6' is a muon with at least 6 GeV transverse momentum , '2e5' means two electrons above 5 GeV. These are known as 'chains', because they are built up from a series of intermediate 'steps', each characterised by one or more 'algorithms'. The same signatures may be used in many chains. At each level, chains are linked to the previous level, so Event Filter chains are linked to Level-2 chains, and Level-2 chains to items in the Level-1 menu. Each chain has a prescale factor associated with it at each trigger level. Prescaling reduces the frequency at which the chain will be processed, so 1 means not prescaled at all, while 10 should reduce the rate by a factor of 10 in addition to any intrinsic rejection power. The choice to prescale an HLT chain or not in a given event is taken at the start of the Level-2 and Event Filter processing of the event using a random number draw. By convention a prescale factor less than or equal to zero disables the chain. Prescales can be decided upon and changed during a run to manage the rate, in response to changes in LHC and detector conditions.

A typical menu consists of around 200 chains and 500 algorithm instances each for Level-2 and Event Filter. The set of prescales is stored separately and contains a number for every chain. A complete configuration is uniquely identified in the database by three numeric keys: the super master key for the full configuration and prescale keys for Level-1 and the HLT.

The Trigger Configuration provides a recipe for how to process events. The HLT Steering is data driven: it applies this recipe to the specific data encountered in a given event. The two HLT levels work in pretty much the same way. The first step in either HLT level is to activate chains based on the results of the previous trigger level. The aim of the HLT steering is then to process all active chains until they can be rejected for some reason. If all chains have been rejected, the event itself can be rejected without further processing. If on the other hand the event is fully processed and at least one chain ended with a positive decision, it is accepted and at Level-2 passed to the Event Filter, or from the Event Filter recorded for offline analysis. The first processing step is to apply prescales to the chains. Some chains will already be rejected at this point. After that the steering will work through the remaining chains, processing the first step of each one, then the next step, and so on. Within a step, algorithms will be run to perform reconstruction, typically in a RoI which may be refined at each step, and to test a hypothesis. Reconstruction examples include basic things like tracks and calorimeter clusters, and combining these into electron, tau, and muon candidates. Hypotheses include applying a transverse momentum threshold or testing matching criteria between two objects. If a hypothesis succeeds the step of the chain is confirmed and the processing of the chain continues. If it fails that's the end of the chain.

The steering employs an intelligent caching mechanism so that if it sees an algorithm is to be run on the same data more than once, it uses the cached result from the first run. Data requests

PoS(ICHEP 2010)487

at level-2 are also cached. The steering is instrumented to monitor algorithm times, decisions and trigger rates.

3. Recently developed procedures and tools

The ability to update HLT prescales during a run was introduced at the start of 2010 data taking after it proved very useful in Level-1 in 2008/9. It allows the trigger operator to respond to changing LHC and detector conditions. Examples include running the HLT with most chains disabled before the LHC declares stable beams, temporarily disabling triggers whose rates spike due to noise, and optimising the output bandwidth as luminosity falls during the lifetimes of a LHC fill. Restarting a run to apply any of these changes would result in several minutes of dead time.

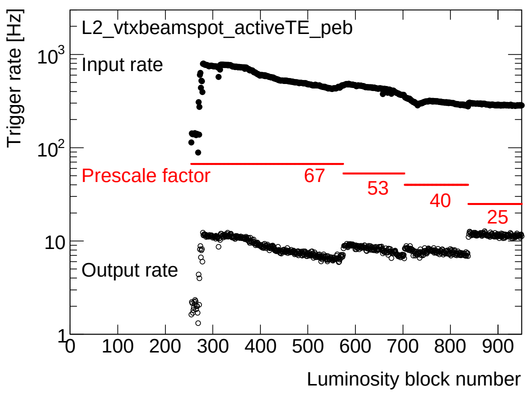

Figure 2: Here the L2_vtxbeamspot_activeTE_peb chain (an arbitrary choice to illustrate the point) is initially disabled, then enabled at luminosity block (LB) 255 when LHC stable beams are declared. As the beam intensity falls, the input rate drops by almost a factor of 4. The prescale factor is reduced, without restarting the run, in a few steps to keep the output rate between about 6 to 12 Hz.

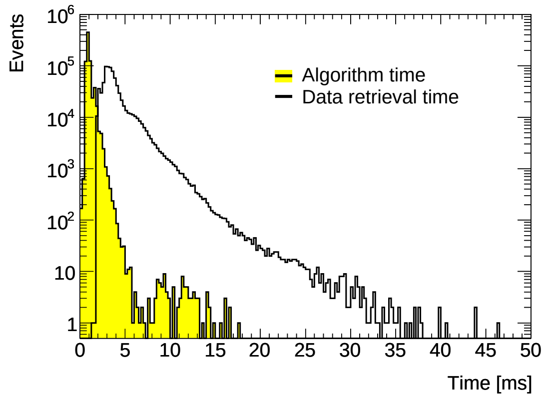

Figure 3: Times for an arbitrary algorithm from an arbitrary run and conditions, as recorded by the HLT resource monitoring software tool. The time is broken down into CPU time running the algorithm and the time to retrieve Region-of-Interest data over the Level-2 network. The aim of this plot is purely to illustrate the detailed monitoring of times within the HLT that is possible.

ATLAS divides data in luminosity blocks (LB). They are typically 2 minutes long and during this time conditions changes should not be made; LBs are the atomic units for data quality. Prescale changes requested by the trigger operator are therefore applied on the next LB boundary. The Level-1 CTP is used to flag when a change is required, since this is a single point through which all events seen in the HLT must flow. The mapping of LBs to prescale sets is recorded in the data-taking conditions database for later offline access. The HLT processes keep a cache of prescale sets and will extend it to include a new one when it is first seen. An example of prescale changes to a particular chain during a run is shown in figure 2.

Two tools have been written to view the Trigger Configuration. The TriggerTool is a Java application that interacts directly with the Trigger Configuration database. It provides visualisation of the menu and full configuration details, comparison of different configurations and prescale sets, and searches across all configurations. Experts also use it to edit the Trigger Configuration and

PoS(ICHEP 2010)487

prepare prescale sets. An interactive web-based view of the configuration is aimed at people who just want to see the status and configuration of a chain they are interested in for a given menu.

It is very useful to be able to monitor online CPU usage and rates, including rejected events: it helps to identify bottlenecks and provides data that can be used to extrapolate trigger performance to new scenarios. The challenge was to implement this with minimal overhead or impact on the T/DAQ system and to handle the large number of events, chains, algorithms and HLT software instances. The solution was to store minimal information for every event (L1 accept time stamp, L1 and HLT decisions) then sample detailed timing and other data from the instrumented HLT steering at some lower rate (about 1 in 10). This is buffered and written out every 100 events as calibration data. Prompt offline processing then produces ntuples and a summary on a web page for each run. Figure 3 shows an example of the detailed data available from this monitoring.

New trigger configurations and software versions are tested using 'enhanced bias' data – events triggered only by a prescaled Level-1 minimum bias trigger. The HLT software is re-run offline on these data and the resulting monitoring data is used to check the expected online resource usage, rates, etc. This has proved to be a powerful technique for validation of the many changes needed to keep the trigger optimal as increasing luminosity was delivered by the LHC.

The same monitoring framework and analysis code is used for offline studies of the trigger performance. One application is the extrapolation of trigger rates to higher luminosities, which requires knowledge of event by event trigger decisions. It can predict trigger rates for the three levels of the trigger system within a few minutes, a capability which was of critical importance for the development of trigger menus during the rapid luminosity increase of the LHC during 2010.

4. Conclusions

ATLAS has a reliable and flexible trigger that has proved successful in taking 7 TeV collisions data in 2010. Trigger Configuration tools allow easy checking and comparison of configurations used online. HLT prescales can be changed during a run to cope with changing LHC and detector conditions. HLT resource monitoring gives a detailed picture of online CPU usage and rates, a powerful diagnostic tool and a way to test new software and configurations. These features have all contributed to the impressive data recording efficiency of ATLAS, which stood at 94% [4] of the LHC-delivered 'stable beams' luminosity at the 20th of July 2010.

References

[1] The ATLAS Collaboration, *The ATLAS Experiment at the CERN Large Hadron Collider*, JINST, 3 (2008) S08003 (2008) [doi: 10.1088/1748-0221/3/08/S08003]

[2] J. Baines for the ATLAS Collaboration, *Performance of the ATLAS Trigger with Proton Collisions at the LHC, in these proceedings.*

[3] N. Berger et al, *The ATLAS High Level Trigger Steering*, in proceedings of *International Conference on Computing in High Energy and Nuclear Physics, Victoria, Canada, 2 - 7 Sep 2007*, J. Phys. Conf. Ser. 119 (2008) 022013

[4] P. Onyisi, on behalf of the ATLAS Collaboration, *Operation of the ATLAS detector with first collisions at 7 TeV at the LHC, in these proceedings.*

PoS(ICHEP 2010)487