ATLAS NOTE

June 15, 2009

ATLAS Trigger and Data Acquisition: Capabilities and Commissioning

D.A. Scannicchio (CERN)

on behalf of the ATLAS TDAQ Collaboration [1)]

Abstract

The ATLAS trigger system is based on three levels of event selection that selects the physics of interest from an initial bunch crossing rate of 40 MHz to an output rate of $\sim$200 Hz compatible with the offline computing power and storage capacity. During nominal LHC operations at a luminosity of 10^{34} cm^{-2}s^{-1}, decisions must be taken every 25 ns.

The LHC is expected to begin operations with a peak luminosity of 10^{31} cm^{-2}s^{-1} with far fewer number of bunches, but quickly ramp up to higher luminosities. Hence, the ATLAS Trigger and Data Acquisition system needs to adapt to the changing beam conditions preserving the interesting physics and detector requirements that may vary with these conditions.

[1)] for the full author list see [5].

1 Introduction

The ATLAS experiment is one of the four major experiments aimed at studying high-energy proton-proton collisions at the Large Hadron Collider (LHC) at CERN.

The Trigger and Data Acquisition system (TDAQ) is organized as a three level selection scheme: the Level 1 (LVL1) which is based on custom electronics, the Level 2 (LVL2) and the Event Filter (EF), jointly referred to as High Level Triggers (HLT), which are software-based. A distinguishing feature of the LVL2 is that its selection is based on data from specific regions of the detector defined per event by the LVL1 trigger, the so called Regions of Interest (RoI). This minimizes the amount of data needed to calculate the trigger decisions thus reducing considerably the overall network data traffic. The initial LHC bunch crossing rate of 40 MHz must be reduced by the TDAQ to $\sim$200 Hz ($\sim$300 MB/s) compatible with the offline computing power and storage capacity. The LHC is expected to begin its operation with a peak luminosity of 10^{31} cm^{-2}s^{-1} with a relatively small number of bunches, but quickly ramp up to higher luminosities. The deployed trigger selection has to adapt to the changing beam conditions while preserving the interesting physics and satisfying varying detector commissioning requirements.

During 2008, a few months of ATLAS cosmic data-taking and the first experience with the LHC circulating beams provided an unprecedented testbed for the evaluation of the performance of the ATLAS DataFlow, in terms of functionality, robustness and stability, as well as its integration with the offline data processing and management.

This paper presents an overview of the TDAQ system and the status of the preparation of the trigger menu for the early data-taking and reports on the usage of the DataFlow infrastructure during the ATLAS cosmic and single LHC beam data-taking.

2 DataFlow

The DataFlow system [1] [2] is responsible for the collection and the conveyance of event data from the detector electronics to the mass storage, while serving the HLT processing farms, and it is based on a push-pull architecture.

The principal components of the TDAQ systems are shown in Figure 1. The movement of event data from detector to mass storage commences with the selection of events by the LVL1 trigger. For each accepted event the LVL1 trigger, via a dedicated data path, sends to the Region of Interest Builder (RoIB) the RoI information. The information contains the p_T and position (in eta and phi) of candidated objects and energy sums (E_T^{miss} and E_T).

The selected event data are also received into the corresponding ReadOut Buffer (ROB) contained in the ReadOut System (ROS) units where they are temporarily stored and provided, on request, to the subsequent stages of the DAQ/HLT system, see Figure 1.

The RoIB then assembles the RoI information from several sources of the LVL1 trigger into a single data structure and forwards it to one of the LVL2 supervisors (L2SV), which supervise the handling of events within the LVL2 trigger system. Each event is allocated to a LVL2 trigger processor (L2PU) where it is analyzed. The analysis is seeded by the objects located at the RoI identified by the LVL1. The L2PU returns to the L2SV the result of its analysis, accept or reject. This result is subsequently relayed by the L2SV to the DataFlow Manager, an application which manages the events during the event building process. In addition an L2PU also sends a summary of the performed analysis to a LVL2 trigger specific ROS (pROS).

The DFM marshals the events during the Event Building (EB): it assigns the selected events to an Event Building node, SFI, which collects the event data from the ROSs and builds a single event data structure, the event. On completing the building of an event the DFM, notified by an SFI, informs all the ROSs to expunge the associated event data. The full event is then sent to the EF for further analysis.

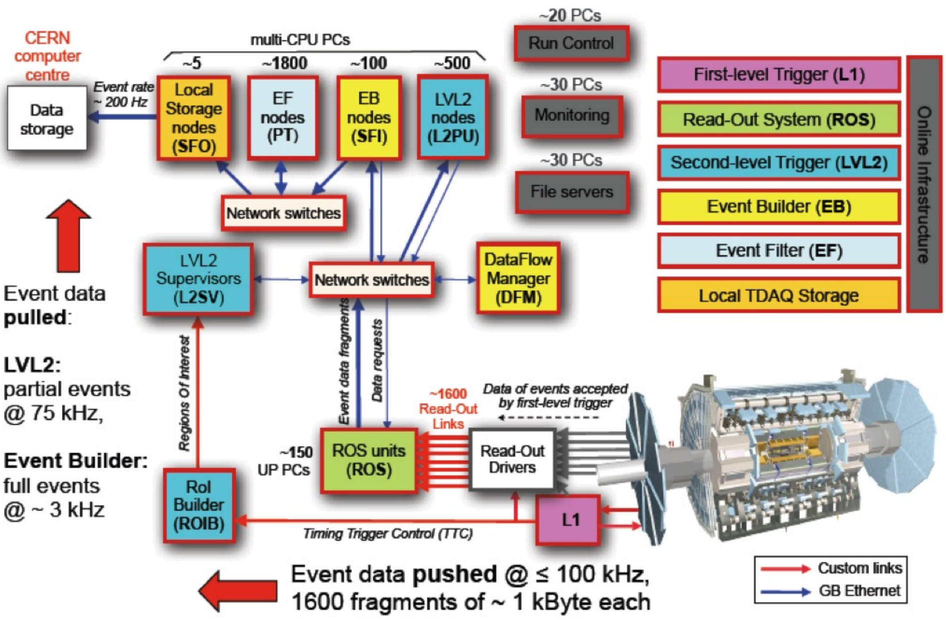

Figure 1: Diagram of the principal components of the ATLAS Trigger and Data Acquisition system.

Events selected by the EF are subsequently sent to the output nodes (Sub-Farm Output, SFO) of the DAQ/HLT system where they are stored in a local file system, according to a classification performed by the EF, and eventually transferred to CERN central data-recording facility. Conversely, those events not fulfilling any of the EF selection criteria are expunged from the EF system.

Regular TDAQ tests are performed in order to assess the system performance, beyond the operational conditions defined by cosmic data taking, and to evaluate new software releases. The present system size does not allow complete system performance studies at the final working point conditions, however it does allow for the evaluation of the overall functionality, scalability and performance studies of a subset of the system, for example event building.

The full system has been able to reach a LVL1 trigger rate of $\sim$70 kHz and a LVL2 trigger rate of 3.5 kHz. Many configuration parameters have to be taken into account when evaluating the performance of the DataFlow system: the event size, the number of nodes allocated for LVL2 or EF event selection, as well as the number of ROS, EB and SFO nodes. For example a test performed using a trigger selection optimized for a preliminary 10^{31} cm^{-2}s^{-1} menu which included a full scan of the Tile calorimeter (i.e. all information from calorimeter are used by the LVL2 trigger) resulted in a LVL1 trigger rate of $\sim$60 kHz limited by the ROS processing power: in this case the ROS PCs were subject to request rates up to 30 kHz, where in the final running conditions at most 12 kHz is foreseen.

Independent tests of the EB-EF scaling properties have also been performed and the results are shown in Figure 2. The throughput scales linearly with the number of event building nodes up to 55 and then it is constant where it is limited by the bandwidth to the EF, which in turn is a result of the current size of the EF, i.e. only one third of the foreseen EF farm has been installed. However the current installed processor power should be sufficient during the first running period of LHC.

When not sending data to the EF processors, the achieved EB throughput is almost double that of the design specifications.

The SFO is the only component which is regularly used at the design working point and even beyond:

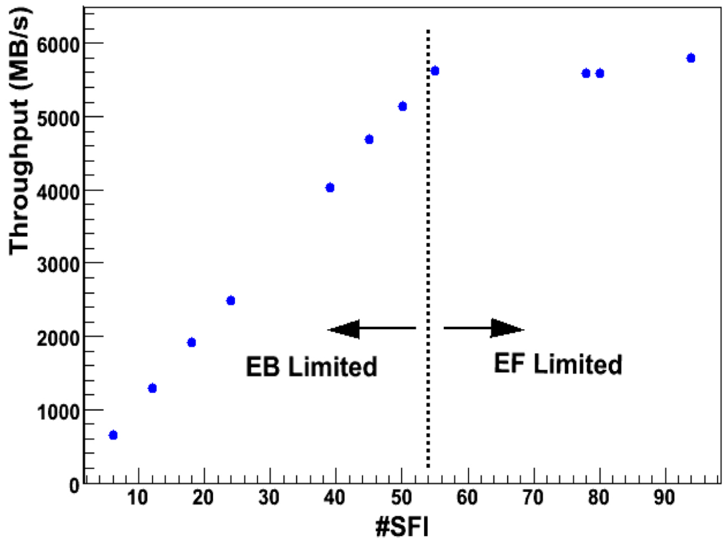

Figure 2: Event Building scalability.

the full farm is in fact able to sustain an aggregate I/O rate of 550 MB/s, while the design value is ∼300 MB/s. Roughly one petabyte of data, distributed over 650 thousand files, have been handled by the farm during the 2008 ATLAS cosmic data-taking (July-October 2008).

3 Trigger

The ATLAS trigger [1] [2] is composed of three levels of event selection that must reduce the output event storage rate to ∼200 Hz (about 300 MB/s) from an initial bunch crossing of 40 MHz. A large rejection against QCD processes is needed while maintaining high efficiency for low cross-section physics processes that include searches for new physics as shown in Figure 3. The rate estimates discussed in the following are based on simulations and are subject to several source of uncertainties which include lack of knowledge of the exact cross-sections, detector performance and beam related background conditions.

3.1 Level 1

The LVL1 trigger system [3] receives data at the full LHC bunch crossing rate of 40 MHz and must reduce the output rate to ∼40 kHz during ATLAS start-up (∼75 kHz at nominal luminosity) within 2.5 μs using its dedicated access to data from the calorimeter and muon detectors.

The total number of allowed LVL1 trigger items that can be deployed at any time is 256. Each of these LVL1 items, programmed in the Central Trigger Processor [3], is a logical combination of the configured LVL1 thresholds, whose number is also limited.

Furthermore, for each of the 256 LVL1 items, a prescale factor N can be specified (where only 1 in N events is selected and passed to the HLT for further consideration). As the peak luminosity drops during a fill, the LVL1 prescale value can be adjusted to keep the output bandwidth saturated without stopping and restarting a data-taking run. A given data-taking run is sub-divided into time intervals called luminosity blocks that provide the smallest granularity at which various data will be monitored and available for physics analysis. The length of the luminosity block represents stable conditions in the data taking. The

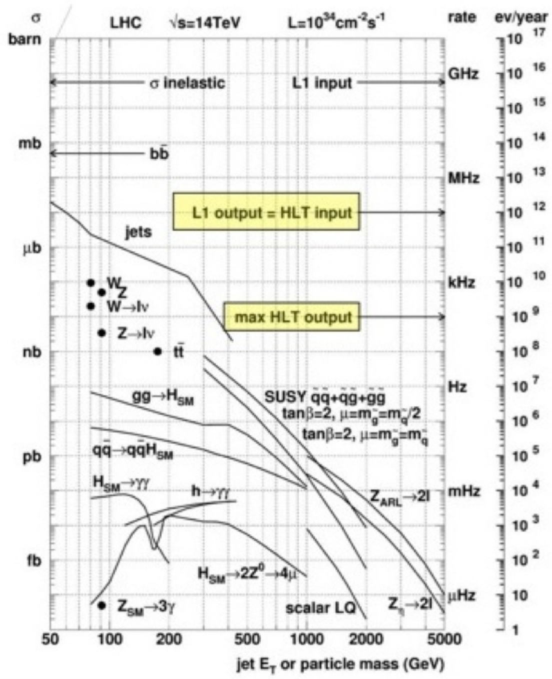

Figure 3: Expected event rates for several physics processes at the LHC design luminosity.

trigger configuration, including the LVL1 prescale settings, remains unchanged within the luminosity block and adjustments to LVL1 prescale factors will be made on luminosity block boundaries.

3.2 High Level Triggers: Level 2 and Event Filter

The High Level Triggers are selection algorithms that are executed on farm of commodity PCs: 500 and 1800 nodes are foreseen for the LVL2 and the EF respectively.

At LVL2 the selection is largely based on RoI. Selectively accessing, reformatting and analyzing the detector data only within the RoI greatly reduces the processing time, and subsequently the processing power, of the LVL2 system, and the data movement to the LVL2 system. The selection algorithms confirm and refine the LVL1 result using fine-grained detector data and improved calibrations constant. In addition to the muon and calorimeter data it also uses in the selection data available from the tracking chambers.

The LVL2 system must reduce the output rate from $\sim$40 kHz ($\sim$75 kHz at design luminosity) to $\sim$3 kHz. The foreseen average event processing time is 40 ms which includes the time to transfer data from the ROSs.

The final online selection is performed by the EF that must reduce the output rate to $\sim$200 Hz (corresponding to $\sim$300 MB/s) with an average processing time per event of 4 s.

Unlike LVL2, which uses specialized algorithms optimized for timing performance, the EF typically uses the same algorithms as the offline analysis. The use of the more complex pattern recognition algorithms and the full set of calibration and alignment constants helps in providing the additional rejection needed.

The EF, in addition to selection, classifies the events according to a predetermined set of event streams and the result of this classification is added to the event structure. Eventually the events are stored to different files on the basis of the EF classification.

3.3 Data streams

ATLAS has adopted an inclusive streaming model according to which the selected events are written locally by SFO to one or more files on the basis of the EF classification.

During the early running the proposed streaming configuration consists of four data streams: electron or photons; jet, tau and Etmiss; muons and minbias. Each stream consists of events that pass one or more trigger signatures. The stream names indicate the type of trigger objects they will contain and events passing certain topological triggers could be written to more than one stream.

The streams have been chosen so as to have approximately the same proportion of events and to keep the total overlap (event duplication across streams) to less than 10%. The stream definitions, based on the experiences gained during initial running, will be optimized to minimize the total overlaps and vary according to luminosity.

In addition to the four physics streams, an express and a calibration stream have also been defined. The express stream is defined to contain events that enable a rapid check on the quality of the data being recorded. Events in the express stream will also appear in one of the foreseen data streams. The calibration stream is defined to contain data needed for detector alignment and the determination of the detector energy scale.

4 Trigger Menu

A recipe for triggering on various physics processes is provided by trigger menus: tables of signatures that are fully specified by thresholds and selection criteria for various physics object at each of the three trigger levels.

A new signature, before being included into a trigger menu, is carefully evaluated: its physics goals (or commissioning or calibration), the efficiency and the background rejection it provides to meet these goals and the consumed bandwidth are taken into account.

The trigger menu also includes additional signatures for trigger validation, monitoring, calibration and measuring the performance of the physics triggers.

A set of trigger menus for the start-up phase have been developed taking into account all these considerations.

Of course the trigger menus will evolve as the understanding of the detector and trigger evolve and as physics requirements mature.

Once data taking operations begin, dedicated data samples for further menu optimization and rate estimations will be collected. Trigger rates have been estimated using a sample of simulated minimum-bias events that, in order to design the trigger menu for low luminosity phase, contains seven million non-diffractive events with cross-section of approximately 70 mb.

4.1 Trigger menu at $L = 10^{31}$ cm^{-2}s^{-1}

The start-up phase of LHC will be an extended period of time, more than six months, during which the luminosity will increase up to 10^{31} cm^{-2}s^{-1} with a low number of bunches in the machine. These conditions will be ideal for commissioning the trigger and detector system, as well as for the initial data taking which will be dedicated to high cross-section Standard Model signatures.

Different menus are being prepared with increasing complexity to match the different running conditions [4].

During the ATLAS start-up and early running phase the focus of the trigger selection strategy will be to commission the trigger and the detector and to ensure that established Standard Model processes are observed. Low p_T thresholds and loose selection criteria at each selection stage will be deployed and many triggers will operate in pass-through mode, which entails executing the trigger algorithms but accepting the event independent of the algorithmic decision. This allows the trigger selections and algorithms to be validated to ensure that they are robust against the varying beam and detector conditions that are hard to predict before data-taking.

The total output rates, that is the cumulative rate accounting for overlaps between the trigger groups as well, for each trigger level are shown in Table 1. The estimated rates are well within the available TDAQ bandwidth, although there are large uncertainties due to the use of simulations that extrapolate from 2 TeV center-of-mass energy of the Tevatron to the 14 TeV expected for the LHC.

Trigger Level	**Rate at L = 10^{31} cm^{-2} s^{-1}**
LVL1	6470 Hz
LVL2	427 Hz
EF	94.2 Hz

Table 1: Total trigger rates estimated at LVL1, LVL2 and EF. The total rate accounts for overlaps.

Experience with early data will allow validation of these estimated rates and further optimization of the trigger algorithms and menus and will improve the ability to estimate rates in the high luminosity regime.

Combining object types at LVL1 provides a mechanism to control the rates while maintaining low enough thresholds to meet the physics goals of the trigger. Event though single object trigger may suffice for low luminosity running, it is necessary to deploy the multi-object triggers at low luminosity to validate them and ensure their reliability as the LHC moves to higher luminosities.

As the luminosity increases, also the use of higher thresholds, isolation criteria and tighter selections at the HLT become necessary to reduce the background rates while achieving selection of interesting physics with high efficiency.

Complex trigger signatures with multiple observables, higher p_T thresholds and tighter selections will be deployed at the design luminosity to maintain the EF output rate at $\sim$200 Hz. With increasing luminosity, higher LVL1 thresholds have to be introduced at the expense of some of the lower thresholds. Tighter HLT selections will also be implemented to achieve the required rejection as running in pass-through mode or with loose selections will not be more possible.

5 Conclusions

ATLAS TDAQ system has been successfully used in cosmic and single LHC beam data taking and the results obtained, backed-up by complementary performance tests, have allowed to validate the architecture of the ATLAS DataFlow and to demonstrate that the system is robust, flexible and scalable enough to cope with the requirements of the ATLAS experiment.

The trigger menu for the LHC start-up phase is designed to enable the rapid commissioning and preparation for the high regime using low p_T thresholds and loose selections. Observations with early data will therefore allow validation of the estimated rates and extrapolations to higher luminosity. Trigger items and their performance have been studied in details at both low and high luminosity.

References

[1] ATLAS Collaboration, Detector and Physics Performance Technical Design Report, CERN/LHCC/99-14/15 (1999)

[2] ATLAS Collaboration, The ATLAS Experiment at the CERN Large Hadron Collider, JINST 3 (2008) S08003

[3] ATLAS Collaboration, Level 1 Trigger Technical Design Report, CERN/LHCC/98-14 (1998)

[4] ATLAS Collaboration, Expected Performance of the ATLAS Experiment: Detector, Trigger and Physics, (2008) CERN-OPEN-2008-020

[5] ATLAS TDAQ Collaboration, TheATLASTrigger/DAQ Authorlist, version2.1, ATL-DAQ-PUB-2009-004, CERN,Geneva,2009,http://cdsweb.cern.ch/record/1195753