Design and expected performance of the ATLAS trigger and event selection

S. George¹ on behalf of the ATLAS Collaboration

¹Department of Physics, Royal Holloway, University of London, Surrey TW20 0EX, UK

e-mail: S.George@rhul.ac.uk

Supported by the UK Particle Physics and Astronomy Research Council.

Abstract. The LHC and ATLAS present the T/DAQ system with a highly challenging environment: the unprecedented LHC rate of 10^9 interactions per second with large and complex detectors with $\mathcal{O}(10^8)$ channels to read out; the bunch crossing rate of 40 MHz requires a decision every 25 ns, while the event storage rate is limited to $\mathcal{O}(100)$ MB/s. Within these constraints, the ATLAS T/DAQ system must separate rare physics signatures from the overwhelming rate of background events. This paper gives an overview of the T/DAQ system and describes how the high level triggers are being designed to meet these challenges. Data bandwidth and processing times in the higher level triggers are reduced by region of interest guidance from the first level trigger and sequential steps in the reconstruction process. Flexibility is paramount in order to adapt to the changing luminosity, backgrounds and physics goals. This is achieved by simple, inclusive trigger menus and modular software design. Algorithms have been developed which provide the flexibility to control the trigger rates.

PACS: 07.05.Hd 07.05.Kf 29.90.+r

1 Introduction

1.1 LHC machine

The LHC will provide proton-proton collisions at a centre-of-mass energy of 14 TeV and design luminosity of $10^{34} \text{ cm}^{-2}\text{s}^{-1}$. At this luminosity there will be on average 23 collisions per bunch crossing. The bunch crossing interval will be 25 ns.

1.2 LHC physics programme

The primary goals of the LHC are to understand the mechanism for electroweak symmetry breaking and to search for new physics at the TeV energy scale, which will be accessible for the first time at the LHC. Both the Standard Model Higgs and those predicted by MSSM or extended versions of SUSY will be sought. Searches will also be made for other physics beyond the standard model, such as new particles predicted by super-symmetric models, compositeness and heavy gauge bosons. LHC will provide precision measurements such as the W and top masses, and parton density functions in new kinematic regimes. The large $b\bar{b}$

production cross section makes the LHC a highly productive B-factory, such that precision CP violation measurements can be made, and flavour oscillations and rare decays can all be studied with much higher precision than has previously been possible.

With the exception of B-physics, these processes can be characterised by high transverse momentum (p_T) leptons, jets or missing energy and this is exploited by the trigger.

1.3 ATLAS experiment

ATLAS is a general purpose experiment at the LHC. The detector, its performance and physics potential are described in [1].

Working out from the beam pipe, the inner detector consists of Silicon pixels, strips and a continuous transition radiation tracker, arranged to give seven precision, three-dimensional measurements and at least 36 further measurements of tracks within ± 2.5 units of pseudorapidity. The inner detector is inside a solenoid which gives a 2 T field.

Outside this is the calorimetry: a finely segmented Liquid Argon (LAr) electromagnetic calorimeter and a hadronic calorimeter using scintillating tiles in the barrel and LAr in the end caps. There is also a LAr forward calorimeter.

The muon spectrometer is built from four technologies: Thin Gap Chambers (TGC) and Resistive Plate Chambers (RPC) are fast enough for bunch crossing identification and used in the LVL1 trigger; Monitored drift tubes are positioned adjacent to these to provide precision measurements for the high level triggers and offline reconstruction. Cathode Strip Chambers (CSC) are used in the innermost layer of the endcaps. These detectors are arranged in three stations in the barrel and each endcap.

2 ATLAS Trigger and Data Acquisition System

2.1 Overview

The LHC and ATLAS present the T/DAQ system with a highly challenging environment. The primary challenge is to select rare physics signatures with high efficiency whilst rejecting the overwhelming majority of events. This can be illustrated by considering selection of the rare $H \rightarrow \gamma\gamma$ decay channel, which for a Higgs mass of ~ 100 GeV in the Standard Model is predicted to occur at about 10^{-13} of the LHC interaction rate. The approach to solve this problem is a three level trigger system, shown in figure 1 and more fully described in [2, 3].

The first level trigger (LVL1) decision is based on relatively coarse granularity calorimeter data and dedicated muon trigger stations. The second level trigger (LVL2) can get data at full granularity and can combine information from all detectors. Region of Interest (RoI) from LVL1 are used to reduce data requested to a few percent of the whole event in most cases. Specialised algorithms are used, optimised for fast rejection. The Event Filter (EF) refines the selection according to the LVL2 classification, performing a fuller reconstruction with more time

and more detailed alignment and calibration data. The EF reconstruction will be based on the use of offline algorithms.

The three-level architecture keeps the data collection and event building bandwidth under control: as a result of the RoI guidance, the LVL2 network only has to deal with requests for (on average) a few percent of the event fragments held in readout buffers, resulting in a maximum bandwidth of a few GB/s. Only events accepted by LVL2 are passed to the event builder, requiring a similar bandwidth to LVL2.

The total rejection of the trigger is by 5 to 6 orders of magnitude. At LVL1 the calorimeter triggers are dominated by jets, and muon triggers by the subset of jets which contain either a semileptonic b decay or a π/K decay in flight to a muon. After the event filter, leptonic W and Z decays form a significant fraction of the remaining events.

2.2 LVL1 Trigger

The LVL1 trigger works by looking for the basic signatures of interesting physics, characterised by high- p_T . These are: muons; electromagnetic, tau/hadronic and jet clusters; missing and scalar sum transverse energy. The trigger decision is based on requiring these objects with various transverse energy thresholds and multiplicities or combinations. The same objects are made available to the LVL2 trigger to provide RoI guidance.

Due to the extreme rate and latency requirements, LVL1 is a hardware trigger. The logic is implemented in a mixture of custom and programmable electronics (ASIC and FPGA) with programmable thresholds to allow flexibility.

2.2.1 LVL1 Muon Trigger The muon trigger looks for track segments in the outer station (or middle station for low p_T triggers in the barrel) then projects back a window to the origin to search for compatible segments in inner stations. The width of the window is determined by the trigger threshold. The low threshold trigger with p_T above 6 GeV requires coincidence in two layers while three layers are required for the higher threshold (20 GeV). There are six programmable thresholds in total.

The LVL1 accept rate for $p_T > 6$ GeV at luminosity $10^{33} \text{ cm}^{-2}\text{s}^{-1}$ is estimated by simulation to be about 23 kHz. The majority of this (17 kHz) is background muons coming from decays of pions and kaons in flight. The rest mainly come from b and c quark decays, including 2 kHz which have a p_T below the threshold due to the resolution of the trigger measurement.

2.2.2 LVL1 Calorimeter Trigger The LVL1 calorimeter trigger uses ~ 7000 dedicated, relatively coarse calorimeter readouts with towers of 0.1 units in pseudorapidity by 0.1 radians in azimuthal angle. The towers have two layers: the electromagnetic and hadronic calorimeters. A sliding window algorithm is used to find towers which satisfy criteria for electromagnetic, tau/hadron and jet clusters.

As an example, the electromagnetic trigger, which is the highest rate LVL1 trigger from the calorimeter, will accept events at 20 kHz with a 30 GeV transverse energy threshold at a luminosity of $10^{34} \text{ cm}^{-2}\text{s}^{-1}$. As well as the E_T threshold, electromagnetic and hadronic isolation provide powerful jet rejection.

2.3 High Level Triggers

The LVL2 Trigger and EF are collectively known as the High Level Trigger. As at LVL1, the event selection is based on simple, inclusive, high-transverse-momentum signatures; the aim is to avoid biases and remain as open as possible to new physics. The selection criteria are derived from the physics analysis requirements. As an example, the trigger menu with only the un-prescaled items for selection at low luminosity comprises the following high- p_T signatures¹: $e25i$, $2e15i$, $\gamma60i$, $2\gamma20i$, $\mu20i$, $j360$, $3j150$, $4j100$, $j60 + xE60$. Overlapping triggers, forced accepts and prescaled triggers will be used to check efficiency. Special calibration and alignment triggers from detectors are also foreseen.

Some B-physics signatures are also included in the low luminosity trigger menu. B-physics cannot be selected by the high- p_T signatures above because the rate is too high; the inclusive rate for b decays to muons with $p_T > 6 \text{ GeV}$ is around 2 kHz – far in excess of the final HLT output rate. A selection based on partial reconstruction of the required decay modes must therefore be done in the HLT. This is especially challenging as it requires tracking, vertexing and particle identification down to low p_T .

3 High Level Trigger Event Selection

The HLT event selection strategy will be implemented in software, running on (mainly) commodity hardware. Current test beds are based on PC farms interconnected by switched fast/Gigabit Ethernet.

The decision has been taken to design the event selection strategy and software for LVL2 and EF in a coherent way with a common selection strategy and common core software to implement this in order to keep maximum flexibility. The main difference between them will come down to performance of algorithms and data access.

3.1 Selection Strategy

The general requirements of the HLT selection are that it must be flexible, fast, robust, efficient and have sufficient rejection power. Clearly some of these are in opposition with one another, so there is plenty of scope for optimisation. The approach devised for the HLT is based on the key concepts of seeding and sequential steps.

---

¹Notation for the physics signatures: the first letter indicates the type of object, e.g. μ for muon, j for jet and xE for missing energy, the number following that gives the p_T threshold, the letter i at the end implies isolation criteria are applied, and an initial number gives the multiplicity if greater than one.

LVL1 RoIs are used as seeds for the LVL2 reconstruction to reduce the amount of data accessed. LVL2 only attempts reconstruction in parts of the detector flagged as RoIs, except in the case of B-physics which is described later. This means that typically only a few percent of the event fragments are needed by the LVL2 processor. Seeding of the EF reconstruction by LVL2 results is also being considered as it can reduce processing time compared to completely unguided reconstruction.

The objects in the trigger menu are built up in stages: a series of algorithms are run in a sequence to reconstruct, say, an electron candidate. This sequence is broken into several steps. After each step the reconstruction of an object can be discontinued if it no longer matches the original hypothesis. If this in turn means that the event is no longer viable for acceptance, then it can be rejected at this stage. The order of the sequences can be optimised so the steps providing the highest rejection come first, to reduce average computation time and the amount of data accessed. The emphasis is on fast rejection; it doesn't matter if the few events which pass the trigger take a relatively long time to be processed as long as the majority can be rejected quickly.

Management of the seeding and algorithm sequences is known as steering; this has been identified as a major component of the software. The data produced in the reconstruction sequences are stored in a structured way to facilitate the steering; this is known as the event data model. A data manager component abstracts data access to hide the differences between LVL2, EF and offline, so that the software may be run in all these environments. Figure 2 shows the context of these components.

3.2 Software Development

The software to make the HLT event decision will be large, complex and long-lived. It is on the critical path for data taking, so reliability, maintainability and flexibility are crucial. In response to this it was decided to take a thorough engineering approach to the software, as is standard practice with trigger hardware and detectors. The software therefore follows a well-defined development process[4] in which quality assurance (QA) plays an important part. The object-oriented analysis and design paradigm was chosen.

The software development process is common throughout T/DAQ. The main phases are: prototyping, requirements capture, high-level analysis and design, detailed design, implementation, testing and integration. This is used iteratively with reverse feedback normal. The key QA strategy is to produce documents at each step which are formally inspected by peers of the authors.

The high-level design phase has recently been completed and the detailed design and implementation is underway. Figure 2 shows the domain decomposition and package relations identified by the high level design.

3.3 Performance Optimisation

The performance of the HLT has a strong influence on the cost of both the T/DAQ system and, indirectly, the offline computing of ATLAS. The physics

performance (efficiency for signal physics and rejection of background events) must be optimised to achieve the best physics reach within the available costs and system performance (data size and computing time).

The number of HLT processors needed depends on the input rate multiplied by the computing time. The data collection network bandwidth depends on the input rate multiplied by the amount of data being transferred.

The number of processors for offline reconstruction depends on the HLT output rate. The product of the output rate and event size relate to the volume of mass storage required.

4 High Level Trigger Examples

Two examples are given to demonstrate the concepts described above.

4.1 Electron Trigger

The HLT electron trigger begins with guidance from LVL1 EM RoIs. Candidate electron and photon clusters are identified by transverse energy and shower shape in the calorimeter. Electron candidates prompt a track search in the inner detector close to the cluster. Energy/momentum matching and position matching of track and cluster are required; this gives powerful rejection against photons from π^0 decays (figure 3). Electron identification can be improved by looking for transition radiation detected in the TRT. Bremsstrahlung recovery for electrons and conversion recovery for photons are possible ways that the selection can be further refined if necessary.

This selection sequence can be done first in LVL2 and then in the EF, but some selection paths may be more optimal than others as shown in figure 4. It can be seen that while all paths give about the same rate in the end, limiting the scope of LVL2 and relying more on the EF for rejection is more efficient. However, this comes at the cost of increased event building bandwidth and computing time, since EF algorithms are slower. The boundary between LVL2 and EF is to be flexible and included in the overall optimisation of the HLT.

At high luminosity, this strategy reduces the electron trigger rate from 21.7 kHz output by LVL1 to 114 Hz final output rate from the HLT. The composition of these accepted event is 40% from $W \rightarrow e\nu$, 13% from b and c decays to $e\nu$ and the remaining 47% from fakes and conversions. Figure 5 shows that the rate can be controlled by fine tuning the threshold around 30 GeV without significant loss in efficiency of $Z \rightarrow e^+e^-$: the rate drops more quickly than the efficiency. This is partly due to the presence of the double electron trigger with an unchanged threshold.

4.2 B-physics trigger

The B-physics trigger is to be used mainly at low luminosity, $10^{33} \text{ cm}^{-2}\text{s}^{-1}$. The selection steps in the HLT start with the LVL1 low- p_T muon trigger which has a rate of 23 kHz. At LVL2, the muon reconstruction can be refined by using the muon precision chambers to improve the transverse momentum resolution. The

Selection	Rate (Hz)
B-physics	~60
Electron	~20
Photon	~20
Muon	~15
Jet	~25
Jet + E_T^{miss}	~20
Total	~160

Table 1. Example trigger rates after the HLT, for physics triggers at low luminosity.

inner tracker can be used to search for tracks close to these muons and hence reject many of the pion and kaon decays which do not match the muon tracks so well (figure 6). These two steps reduce the rate to around 5 kHz.

After this, semi-exclusive decay reconstruction must be done to make further rejections, by performing an unguided track reconstruction in the inner detector. This is necessary because the momentum of tracks required is too low to be seeded by the calorimeter or muon detectors. It is however possible to use the z-intercept of the triggering muon to constrain the z-vertex of the track search, with a good probability of eliminating pile-up tracks.

Channels studied include $B_d \rightarrow J/\psi(e^+e^-)K_s^0$, $B_d \rightarrow \pi^+\pi^-$, and $B_s \rightarrow D_s(\phi(KK)\pi)\pi$. At LVL2 the reconstruction is simpler, for example only the J/ψ is reconstructed, not the K_s^0 or secondary vertices. Figure 7 shows the reconstructed mass peak for $J/\psi \rightarrow e^+e^-$ from simulated B_d decays. Reconstruction can be refined in the EF by using more time consuming algorithms like secondary vertex reconstruction, to cut on decay length, and better track reconstruction, which permits tighter mass cuts and more exclusive selections.

5 Trigger Rates

An example of the output rate after the HLT at low luminosity is shown in table 1. Note that these rates are subject to uncertainties of the order of a factor of two to three in the cross sections, however there are effective ways to control the rate as shown above. Pre-scaled physics triggers, calibration and monitoring triggers have yet to be added.

6 Conclusions and outlook

The ATLAS collaboration is rising to the challenges of designing an effective trigger and event selection strategy. The LVL1 trigger is well advanced and now starting full slice tests of final prototypes with full functionality. The HLT designs are being matured in preparation for the HLT, DAQ and DCS Technical Design Report which will be submitted to the LHCC around the end of 2002.

A viable approach for the HLT has already been demonstrated for most key physics signatures. This design is from the outset robust and flexible to respond to changes in cross sections and conditions, which cannot be known accurately before the LHC starts. The computing time and data access are minimised by seeding and sequential steps to give fast rejection. Now a major software project is underway to develop a coherent and maintainable framework and algorithms for the HLT.

In the next year, performance studies will be extended, rate evaluations will be made in greater detail and the selection strategies will be further optimised.

References

1. ATLAS Collaboration: ATLAS Detector and Physics Performance Technical Design Report, CERN/LHCC 99-14 and 99-15, 1999.
2. ATLAS Collaboration: ATLAS First Level Trigger Technical Design Report, CERN/LHCC 98-14, 1998.
3. ATLAS Collaboration: ATLAS High-Level Triggers, DAQ and DCS Technical Proposal, CERN/LHCC 2000-17, 2000.
4. Atlas TDAQ and DCS Connect Forum:
<http://atlas-connect-forum.web.cern.ch/Atlas-connect-forum/>

Fig. 1. Schematic of the ATLAS T/DAQ system.

Fig. 2. High level design of the software, identifying domains. Arrows show use relationships between domains. The event data model is an abstract concept which underpins several domains. Components outside the shaded circle are external.

Fig. 3. Source of electron trigger candidates at LVL2. The open bars show the cause of triggers after the calorimeter trigger, when photons dominate the rate. Matching the calorimeter to a track, shown by the hatched bars, is very effective in discriminating against these photons.

Fig. 4. Several paths through LVL2 and EF for electron triggers, demonstrating different possibilities for the LVL2/EF boundary.

Fig. 5. Trigger rate and $Z \rightarrow e^+e^-$ efficiency as a function of E_T threshold for LVL2 and EF.

Fig. 6. Improvement in rejection of non-prompt muons at LVL2: efficiencies with muon spectrometer alone (left) and when information from inner detector is also used (right).

Fig. 7. Reconstructed mass of $J/\psi \rightarrow e^+e^-$ candidates at LVL2 from simulated $B_d \rightarrow J/\psi(e^+e^-)K_s^0$ signal events. The open histograms show the mass distribution for all combinations of oppositely charged tracks considered, whilst the filled part of the histograms show the distributions for the best combination in the event. The relative contributions from correct and false combinations are indicated by the hatched and cross-hatched areas, respectively. Dashed vertical lines show the mass window that is used to select these events.