The ATLAS Trigger System

R. Hauser¹
on behalf of the ATLAS collaboration

Michigan State University
East Lansing
Michigan
USA
e-mail: rhauser@fnal.gov

Received:

Abstract. ATLAS is one of two general-purpose detectors at the next-generation proton–proton collider, the LHC. The high rate of interactions and the large number of read-out channels make the trigger system for ATLAS a challenging task. The initial bunch-crossing rate of 40 MHz has to be reduced to about 200 Hz while preserving the physics signals against a large background. ATLAS uses a three level trigger system, with the first level implemented in custom hardware, while the high level trigger systems are implemented in software on commodity hardware. This note describes the physics motivation, the various selection strategies for different channel and the physical implementation of the trigger system.

PACS: 25.70.Ef; 21.60.Gx; 27.30.+t

1 Introduction

The Large Hadron Collider (LHC) is the next-generation proton–proton collider currently under construction at CERN where it is planned to start taking physics data in 2007.

With a design luminosity of $10^{34}\text{cm}^{-2}\text{s}^{-1}$, a center of mass energy of $\sqrt{s} = 14\text{ TeV}$ and a bunch-crossing rate of 40 MHz the LHC allows for a wide physics program from searches for new physics to precision measurements of the Standard Model (SM).

The high bunch-crossing rate and the large number of read-out channels (10^8) provide a formidable challenge for the ATLAS Trigger/DAQ system. The about 23 interactions per bunch crossing and the corresponding high interaction rate of about 10^9 Hz require a highly efficient selection process. For example, the signal rate of a SM $H \rightarrow \gamma\gamma$ decay for a Higgs mass of 120 GeV is about 10^{-13} of the interaction rate.

In this note we present the ATLAS trigger system, starting from the physics motivation to a list of event-selection signatures and the details of the implementation including hardware and software aspects.

2 The ATLAS Experiment

ATLAS [1] is one of two general-purpose experiments at the LHC. Its inner detector is located inside a 2 T solenoid magnet and consist of a silicon pixel

detector, silicon microstrips (SCT) and a transition-radiation tracker (TRT). Outside of this is a liquid-argon electromagnetic calorimeter, and a hadronic calorimeter using scintillating tiles in the barrel and liquid-argon in the endcaps.

The muon spectrometer uses thin gap chambers (TGC) in the endcap and resistive plate chambers (RPC) in the barrel which are fast enough to be used in the first-level trigger (LVL1). For precision measurements monitored drift tubes (MDT) are used. Cathode strip chambers (CSC) make up the innermost layer of the endcaps.

The total number of read-out channels is of the order of 10^8 .

3 Physics Motivation

The ATLAS detector allows for a rich and far-reaching physics program [2]. Both precision measurements of SM parameters and searches for a wide variety of new physics beyond the SM will be possible. In addition there is a B-physics program complementing the physics done at specialized B factories.

Examples of physics that will be addressed by ATLAS include:

- The source of electroweak symmetry breaking.
- Precision measurements of SM parameters including top mass to 1 to 2 GeV, W mass to about 15 MeV, etc. The increased precision in these measurements will provide constraints on new physics.
- New physics searches including supersymmetry (SUSY), extra dimensions (ED), compositeness and new heavy gauge bosons (W'/Z') etc.

4 The ATLAS Trigger System

The ATLAS trigger system is composed of three levels (see Fig. 1). The first level is implemented in custom hardware. It reduces the 40 MHz input rate to about 75 kHz. The second and third (Event Filter) levels are referred to collectively as the High-Level Trigger system. They share an overall trigger selection framework, and differ mostly in the amount of event data they access and how they access it as well as in the complexity and speed of the algorithms.

While the LVL1 uses only coarse-grained calorimeter and muon information, the second-level trigger (LVL2) can use full-resolution, full-granularity data from all detectors and combine the information from different sub-detectors for the first time.

In practice, however, the LVL2 trigger restricts itself to so-called Regions of Interest (RoI), small regions in pseudo-rapidity – azimuth ($\eta - \phi$) space centered on objects identified by LVL1. Data from RoIs make up a small subset of the full event. Pointers to these RoIs are provided by the LVL1 trigger. Data are accessed on demand from the buffers which store the event until the LVL2 decision is made. The LVL2 trigger reduces the 75 kHz rate from LVL1 to about 2 kHz.

After an event passes LVL2 the full event is built and sent to the Event Filter (EF). Algorithms in the EF can access the full event and will be derived from

Fig. 1. The ATLAS trigger system. See text for explanation of component names.

offline code. The algorithms may be *seeded* by the results of LVL2. A further reduction to about 200 Hz is achieved by the EF before events are put into mass storage.

ATLAS has tried to make its trigger as inclusive as possible, keeping the thresholds for fundamental objects sufficiently low to be sensitive to decay products of new particles and to leptons from W and Z decays. Low- p_T thresholds are important to understand the background as well as the shape of the spectrum down to low values of $p_T(E_T)$ for certain signatures.

The reason for using inclusive selections is to cover all possible topologies expected from new physics and not to bias the trigger by exclusive selections or topological criteria. Given the nature of the LHC as a discovery machine, the trigger should also be sensitive to presently-unforseen new physics.

Keeping the p_T thresholds reasonably low is driven by the desire to ensure a safe overlap with the discovery potential of Run II at the Tevatron. One also wants to keep a safety margin to refine and/or optimize selections offline with a more powerful analysis and relax cuts later for checks of systematics.

Finally there are large uncertainties in our understanding e.g. of the QCD cross sections, so one wants to keep a safety margin of at least a factor two to three in the rates.

Table 1 shows examples of physics signatures in terms of trigger objects and the related physics coverage.

Objects	Physics Coverage	Nomenclature
Electron	Higgs, W'/Z', ED, SUSY, W, top	e25i, 2e15i
Photon	Higgs, ED, SUSY	$\gamma 60i$, $2\gamma 20i$
Muon	Higgs, W'/Z', ED, SUSY, W, top	$\mu 20i$, $2\mu 10$
Jet	SUSY, compositeness, resonances	$j 400, 3j 165, 4j 110$
Jet + missing E_T	SUSY, leptoquarks	$j 70 + xE 70$
Tau + missing E_T	Ext. Higgs models (e.g. MSSM), SUSY	$\tau 35 + xE 45$

Table 1. Example of physics coverage for various signatures for $L = 2 \times 10^{33} \text{cm}^{-2} \text{s}^{-1}$. The last column refers to the multiplicity, trigger object type (electrons, photons, muons, jets and missing energy), E_T threshold and isolation used in the selection.

5 First-Level Trigger

$2 \times 10^{33} \text{cm}^{-2} \text{s}^{-1}$	Rate (Hz)	$10^{34} \text{cm}^{-2} \text{s}^{-1}$	Rate (Hz)
MU20	0.8	MU20	4.0
2MU6	0.2	2MU6	1.0
EM25I	12.0	EM30I	22.0
2EM15I	4.0	2EM20I	5.0
J200	0.2	J290	0.2
3J90	0.2	3J130	0.2
4J65	0.2	4J90	0.2
J60 + xE60	0.4	J100 + xE100	0.5
TAU25I + xE30	2.0	TAU60 + xE60	1.0
MU10 + EM15I	0.1	MU10 + EM15I	0.4
Others (pre-scales, calibration)	5.0	Others	5.0
Total	≈ 25		≈ 40

Table 2. LVL1 trigger menu for low and high luminosity. The quoted rates do not include any safety factor.

The LVL1 trigger [3] is implemented in custom hardware. It reduces the event rate from the 40 MHz bunch-crossing rate to about 75 kHz. Since a decision cannot be reached in the 25 ns between two bunch crossings the detectors store the event data in pipelined buffers until the LVL1 decision is made. The pipelines allow a fixed latency of up to $2.5 \mu\text{s}$ for a trigger decision, after which accepted events are forwarded to detector-specific Read Out Drivers (ROD). These drivers in turn send their data to detector independent Read Out Buffers (ROB) where the data are stored until a LVL2 decision is reached.

The LVL1 trigger only uses information from the calorimeters and muon detectors to reach its decision. Results from both the calorimeter and the muon triggers are combined in the Central Trigger Processor (CTP) which makes the final decision based on multiplicities of identified trigger objects for various p_T thresholds and information on global-energy variables (e.g. missing E_T). The decisions are forwarded via the Timing, Trigger and Control (TTC) system to the front-end electronics. For accepted events the LVL1 trigger sends information

to the Region Of Interest Builder (RoIB) which assembles a list of RoIs for the event, to be used by LVL2. The information covers both RoIs that were used in making the LVL1 decision (*primary RoIs*) and, possibly, additional RoIs that were identified (*secondary RoIs*).

The overall structure and data flow inside the first level trigger is shown in Fig. 2.

Fig. 2. LVL1 Trigger Overview.

5.1 Calorimeter Trigger

The LVL1 calorimeter trigger uses a pre-processor followed by a Jet/Energy sum processor and a Cluster processor. The latter identifies electron and photons as well as taus and hadrons.

The analogue sums are performed on the detectors. There are a total of about 7200 relatively-coarse granularity trigger towers; there are separate towers for the electromagnetic (EM) and hadronic (HA) calorimeters of size $\Delta\eta \times \Delta\phi = 0.1 \times 0.1$ in the region $|\eta| < 2.5$. Forward calorimeters are treated with even coarser granularity.

For the electron/photon identification a sliding-window algorithm sums all neighbouring trigger towers and finds the maximum tower in a 0.2×0.2 region.

The 12 EM towers surrounding the found cluster in a 0.4×0.4 region and all 16 hadronic towers behind it are used to apply isolation criteria.

The tau/hadron trigger uses the same inputs but a slight variation on the algorithm. It considers both electromagnetic and hadronic towers in the central 2×2 core. The isolation criteria uses only the outer 12 towers around the core in both EM and HA.

The jet trigger algorithm uses a granularity of 0.2×0.2 in η and ϕ for its window algorithm and sums in depth over both electromagnetic and hadronic calorimeters. A local E_T maximum in a 2×2 cluster (0.4×0.4 in $\Delta\eta \times \Delta\phi$) is used to identify the RoI position of the jet, while a separate trigger cluster is used to measure the total E_T of the jet. The latter has a programmable size for each threshold setting and can be 0.4×0.4 , 0.6×0.6 or 0.8×0.8 in $\Delta\eta \times \Delta\phi$.

Finally, the LVL1 trigger computes both missing and total E_T as well as x and y components of the missing E_T , which can be used in the trigger decision.

Fig. 3. LVL1 electron/photon Algorithm.

5.2 Muon Trigger

The LVL1 muon trigger uses only RPC (in the barrel) and TGC (in the endcap) information for making its decision (Fig. 4). It makes use of three stations with two layers each in the barrel. The algorithm requires a coincidence of hits in different layers within a road where the width of the road is related to the p_T threshold applied. For the detection of low- p_T muons a successful trigger condition consists of three hits in the four inner layers while the high- p_T trigger requires an additional hit in the outer station. A similar algorithm is used in

the endcap. The programmable coincidence logic allows multiple thresholds to be used at the same time (three each for the high and low- p_T algorithms). The Muon-to-CTP Interface (MuCTPI) forwards the results from the barrel and endcap triggers to the CTP and the RoIB.

Fig. 4. The LVL1 muon trigger.

5.3 Performance

Table 2 shows the LVL1 trigger menu for low and high luminosity together with the expected event rates. As can be seen the total expected rate at $L = 10^{34} \text{cm}^{-2}\text{s}^{-1}$ is about 40 kHz giving a safety factor of about two compared to the design input rate for LVL2 of 75 kHz.

With the planned trigger thresholds for the LVL1 calorimeter trigger, the rate is dominated by the single EM trigger which has a rate of more than 20 kHz at a luminosity of $10^{34} \text{cm}^{-2}\text{s}^{-1}$.

The majority of muons passing the trigger come from π and K decaying in flight and from b and c quark decays (see Fig. 5).

6 High-Level Trigger

High-Level Trigger (HLT) [4] refers to both LVL2 and the EF system. Since both share the same trigger selection framework the boundary between them is very flexible and many algorithms can be moved rather easily from one to the other. This is especially true for algorithms which were written with an online environment in mind, though less so for algorithms which are derived from offline code and might require access to up-to-date calibration and alignment data.

Fig. 5. Differential cross-sections for inclusive muon production.

The following section describes the common selection-software framework and outlines the differences between LVL2 and EF. All software is written in C++ and runs under the Linux operating system.

6.1 HLT Selection Software

The HLT selection-software framework is based on the ATLAS offline framework Athena [5] which in turn makes use of GAUDI [6]. This allows easy integration into the offline environment for developing, debugging and physics-performance studies. On the other hand the strict separation of interfaces and implementation in the framework allows one to replace the various offline services with an on-line version where required. Examples are error logging, access to configuration databases, etc., and also access to the event data itself.

The HLT event-selection software defines additional requirements on trigger algorithms compared to offline algorithms. These are especially important in the case of LVL2, where data access must be formulated in terms of a Region of Interest, where the data will be requested by the algorithm via the network from the Read Out System.

The HLT selection software uses two principles to steer the triggering process: seeding and sequential processing.

Seeding refers to the fact that each trigger level will use information from the previous stage to guide the processing.

Sequential processing means that the algorithm controlling the overall execution of the trigger algorithms (called the *steering*) arrives at its final decision

by a sequence of steps where, after each step, the event can be potentially rejected. This early rejection saves processing time and reduces the latency for events where a decision can be made based on some rather simple and quickly-calculated properties. The more time-consuming algorithms are only run at a later stage. An event is only accepted if it passes through all selection steps.

Thanks to the RoI mechanism and the use of sequential selection, the LVL2 trigger accesses on average only about 2% of the full event data.

The selection process starts with an initial set of trigger elements which in the case of LVL2 are derived from Regions of Interest. At each step the steering uses a table-driven approach to select the right algorithm to run for each trigger element, then compares the resulting set of trigger elements against the trigger menu. When a decision can be made that the existing trigger elements can no longer satisfy any of the signatures in the trigger menu the processing can stop and the event can be rejected.

At LVL2 the algorithms access event data only through a well-defined interface, specifying a RoI and detector type. In the offline and EF implementation of this interface this is translated into accessing the correct part of the full event in memory. In the LVL2 implementation, this is translated into one or more accesses to the Read Out System. Consequently, for events rejected at LVL2 about 98% of the event data will stay in the ROBs of the Read Out System without ever being transferred across the network.

A representative HLT trigger menu is shown in Table 3 for both low and high luminosity.

Selection	$2 \times 10^{33} \text{cm}^{-2} \text{s}^{-1}$	$10^{34} \text{cm}^{-2} \text{s}^{-1}$
Electron	$e25i, 2e15i$	$e30i, 2e20i$
Photon	$\gamma60i, 2\gamma20i,$	$\gamma60i, 2\gamma20i$
Muon	$\mu20i, 2\mu10$	$\mu20i, 2\mu10$
Jets	$j400, 3j165, 4j110$	$j590, 3j260, 4j150$
Jet + E_T^{miss}	$j70 + xE70$	$j100 + xE100$
Tau + E_T^{miss}	$\tau35i + xE45$	$\tau60 + xE60$
Muons + Electrons	$\mu10 + e15i$	$\mu10 + e15i$
B physics	$2\mu6$ with $m_B, m_{J/\psi}$	$2\mu6$ with m_B

Table 3. HLT Trigger Menu.

In the following sections we present some examples for the HLT strategy for the case of $L = 1 \times 10^{34} \text{cm}^{-2} \text{s}^{-1}$.

6.2 Electron Trigger

The sequence of the HLT electron trigger begins with EM RoIs from the LVL1 trigger. In the first step they are used to find clusters in the calorimeter to which transverse-energy and shower-shape criteria are applied. Most events are already rejected at this stage. For electron candidates, a track search in the inner detector in the vicinity of the cluster is performed. A large rejection against photons from π^0 decays is achieved by requiring matching of energy/momentum and position

of cluster and track (see Fig. 6). The transition-radiation tracker can be used to further discriminate electrons from pions. Bremsstrahlung recovery for electrons and conversion recovery for photons are further steps which are likely to be performed in the EF. The EF can also redo some of the steps done in LVL2 using additional calibration and alignment information which are available here.

Fig. 6. Rejection gained by track match of EM cluster.

The exact order in which these steps are executed and if they are done in LVL2 or the EF leads to different paths for the selection process as seen in Fig. 7. Although the selection efficiency is generally higher when using the EF for rejection, this has to be weighted against the increased rate and bandwidth out of LVL2 and the computing resources required at the EF level. The common HLT selection software used in both systems provides the flexibility to optimize the overall HLT system taking both physics and system parameters into account.

Starting with an LVL1 EM-cluster trigger rate of about 22 kHz, the HLT reduces this to about 114 Hz. The remaining sample is composed of electrons from $W \rightarrow e\nu$ (40%), from b and c decays to $e\nu$ (13%) and fakes and conversion (47%).

As an example, the execution time of the LVL2 calorimeter algorithm is in the order of 2 ms on a 2 GHz PC, excluding data access and preparation time. The LVL2 tracking algorithm takes about 3 ms. This is well within the overall latency budget of about 10 ms foreseen for LVL2.

6.3 Muon Trigger

The muons accepted by the LVL1 trigger include a high rate of muons from π and K decaying in flight (see Fig. 5), many of which have true p_T below the nominal threshold. The first step of the HLT muon trigger is to try to confirm the

Fig. 7. Electron trigger: possible selection paths through the High-Level Trigger.

LVL1 muons and reject fakes. The algorithm uses MDTs in addition to RPCs and TGCs, and achieves a p_T resolution of about 5.5% at low p_T and about 4% in the high p_T case. The trigger efficiency is about 90% for muons with p_T above the trigger threshold (see Fig. 8 and 9). Overall the muon-spectrometer algorithm reduces the LVL1 rate by factors of 2 and 10 for the low and high- p_T cases respectively.

A further strong rejection is achieved by combining the LVL2 muon-spectrometer results with information from the precision tracker. The algorithm takes advantage of the fact that, for those muons originating from π and K decays, the p_T measured in the inner detector and the muon system will differ significantly. Applying p_T -matching criteria achieves another significant reduction compared to the LVL2 muon algorithm alone (Fig. 10). Following this, isolation criteria in the calorimeter can be applied to further reject muons from semi-leptonic b and c decays.

The trigger rates for muons with $p_T > 6$ GeV are shown in Table 4.

6.4 B-Physics trigger

The B-physics trigger must be able to identify decay channels of interest against the large background of $b\bar{b}$ events. Examples are $B_d \rightarrow \pi^+\pi^-$ and $B_d \rightarrow J/\psi K_s$ with J/ψ decaying to either e^+e^- or $\mu^+\mu^-$ for CP-violation studies, $B_s \rightarrow D_s\pi$ and $B_s \rightarrow D_s a_1$ for studies of B_s oscillations, and rare decays like $B_{d,s} \rightarrow \mu^+\mu^- X$. Other areas of study include b-hadron production measurements.

Fig. 8. LVL2 Muon trigger efficiency at low luminosity for a 6 GeV LVL1 threshold. μ FAST is the LVL2 muon algorithm described in the text.

Fig. 9. LVL2 Muon trigger efficiency at high luminosity for a 20 GeV LVL2 threshold.

Fig. 10. Improvement when combining LVL2 muon algorithm with inner detector track information.

Contribution	LVL2 Muon	Combined
K/π decays	3.0	1.0
b decays	0.9	0.7
c decays	0.5	0.4
Total	4.4	2.1

Table 4. Trigger rates in kHz for LVL2 muon-spectrometer algorithm alone and combined with precision tracker for $L = 1 \times 10^{33} \text{cm}^{-2} \text{s}^{-1}$, a 6 GeV p_T threshold and $|\eta| < 1.0$.

The original B-physics trigger strategy of ATLAS required a $p_T > 6$ GeV muon at LVL1, which, if confirmed at LVL2, was followed by a search for low- p_T tracks in either the full TRT or the full SCT. The track search, not guided by RoIs from LVL1, implied a significant increase in the size of the LVL2 farm compared to the high p_T trigger.

Given the doubling of the initial target luminosity and the fact that a large part of the HLT resources may be deferred, the B-physics strategy has been recently re-assessed.

For the single-muon trigger the p_T threshold has been raised to 8 GeV and will be combined with either a full scan in the pixel and SCT, or an additional LVL1 Jet or EM RoI. At LVL2, tracks can be reconstructed inside the RoI, vastly reducing the resources required compared to a full scan, although with some loss of efficiency.

The di-muon trigger uses a 6 GeV p_T threshold at LVL1, leading to a rate of about 200 Hz at $2 \times 10^{33} \text{cm}^{-2} \text{s}^{-1}$. Muons from π and K decays can be suppressed at LVL2 using information from the precision tracker as described above. In the EF, tracks can be refit and selections can be made based on mass and decay-length cuts. The final trigger rate is about 10 Hz.

For hadronic final states an alternative to the full-scan approach consists

in requiring an additional low- E_T LVL1 jet which defines an RoI to be used for track reconstruction in the inner detector. This reduces the data volume to consider by about a factor of 10 with similar savings in execution time. An E_T threshold of about 5 GeV seems to give reasonable jet multiplicities (about two RoIs per event), based on an initial study.

The same approach has been followed for muon-electron final states where the LVL1 calorimeter trigger condition is replaced with a low- E_T EM RoI. Fast simulation studies show that reasonable RoI multiplicities and efficiencies can be obtained with an E_T threshold of about 2 GeV. At LVL2 the electron candidates are confirmed, followed by a track search in the SCT, pixels and TRT guided by the RoI information.

6.5 Summary of Trigger Rates

Table 5 summarizes the expected rates out of the HLT system for a luminosity of $2 \times 10^{33} \text{ cm}^{-2}\text{s}^{-1}$.

HLT signature	Rate (Hz)
$e25i$	40
$2e15i$	< 1
$\gamma60i$	25
$2\gamma20i$	2
$\mu20i$	40
$2\mu10$	10
$j400$	10
$3j165$	10
$4j110$	10
$j70 + xE70$	20
$\tau35i + xE45$	5
$2\mu6$ with vertex, decay length and mass cuts (m_B , $m_{J/\psi}$)	10
Others (pre-scaled, exclusive, monitor, calibration)	20
Total	≈ 200

Table 5. HLT output rates at $2 \times 10^{33} \text{ cm}^{-2}\text{s}^{-1}$.

6.6 HLT/DAQ Architecture

The HLT system is implemented as a LVL2 processor farm and multiple EF subfarms after the event building stage. LVL2 deals with a much higher event rate, but only a fraction of the full event data. The EF works on full events but at a much lower rate. Overall the network bandwidth needed is in the same order of magnitude for both systems.

Figures 11 and 12 show a schematic outline of the HLT/DAQ architecture. One (or more) central switch is used for each of LVL2 and event building. The use of commodity networking interfaces and switches allows for flexible arrangements during commissioning or when adding new processing nodes. For example, a

Fig. 11. A schematic view of the LVL2 architecture.

single switch can be used for both LVL2 and event building in the initial phase and can be replaced with multiple central switches later when the bandwidth requirements increase.

All processing nodes are commodity PCs running the Linux operating system and all software is written in C++. Gigabit Ethernet is foreseen for all network links in the system. The only custom hardware deployed in the HLT system is the ROI Builder and the input part of the Read Out System (ROBin), both of which have to run at the full LVL1 rate of 75 kHz. The ubiquitous use of standard hardware allows us to take advantage of the progress in this sector in industry. Delaying the purchase of components until the latest possible date will allow us to have the most up-to-date hardware at the time when the experiment requires it. A clear upgrade path exists for all components in the long run if necessary, e.g. in the form of 10 Gigabit/s Ethernet for the networking.

6.6.1 Read Out System The Read Out System (ROS) stores the event data until the LVL2 trigger has made its final decision.

The Read Out Buffers (ROB) receive their input via uni-directional read-out links (ROL, the actual implementation uses S-LINK [7]) from the detector-specific Read Out Drivers (ROD) after an event has been accepted by LVL1. The current design foresees that each ROB has four input links and is implemented as a standard PCI card with 64 MBytes of buffer memory per link. A prototype with two input links currently exists.

Fig. 12. A schematic view of the event building and EF architecture.

The ROB implements a flexible buffer manager which answers request for event data from LVL2 or the event builder. While the high-speed input part is implemented in FPGAs, an additional PowerPC CPU is available for handling more complex operations.

The interface to the LVL2 trigger and the event builder consists of multiple Gigabit network links. Two variants are currently under study, with a decision for the final design to be taken by the end of 2003 (see Fig. 13):

In the first variant a number of ROB PCI cards are put into a standard rack-mounted PC with multiple PCI buses. Requests are received by the main CPU and event data are read out via the PCI bus and then sent via one of two Gigabit links to either the LVL2 system or an event builder node.

In the second variant each ROB uses an on-board Gigabit NIC directly to receive and answer the requests for event data. This variant requires an additional layer of smaller concentrating switches to avoid the need for a central switch with a large number of (expensive) ports. However, since the PCI bus is not used for data transfers, a larger number of ROBs can be housed in a single PC. The PCI bus is still used for control and monitoring operations.

6.6.2 LVL2 The LVL2 trigger receives RoI information from LVL1 through the RoI Builder. The RoI Builder has eight S-LINK inputs receiving RoI information from the various LVL1 trigger components. It combines these data on an event per event basis and forwards them on a round-robin basis to one of

Fig. 13. The bus and switched based variants of the Read Out System.

multiple LVL2 supervisor nodes.

A LVL2 supervisor node is responsible for scheduling events to a subset of the full LVL2 farm. It assigns a LVL2 processing node for the event, forwards the RoI information and keeps track of the decision or timeouts in case of errors.

Finally, the LVL2 supervisor sends lists of accepted and rejected events to the Data Flow Manager (DFM) of the event building system. It groups the decisions to reduce the message rate with a slight increase in latency for the Read Out System.

If an event is accepted at LVL2, the LVL2 processing node will send the full details of the selection process to a so-called Pseudo-ROS. This is a node which behaves in all respects like a normal ROS to the event building system but has no detector inputs. Instead it keeps track of the LVL2 results so when the event is built, the LVL2 result appears like another subdetector.

LVL2 processing nodes are commodity PCs. They are connected to the central LVL2 switch through a set of concentrating switches to reduce the number of ports on the central switch. The actual data throughput per processing unit is small enough that about five of them can use a common concentrating switch and still share a single Gigabit uplink.

Inside a LVL2 processing unit multiple threads are used to process events in concurrently. Each thread has its own steering controller and handles exactly one event at a time. When event data are accessed from the ROS, a thread sends out the requests and then puts itself to sleep until all data have arrived. In the meantime another event can be processed. This effectively hides the typical latency of about 150 to 180 μs until the reply for a data request is received.

6.6.3 Event Building Each LVL2 Supervisor sends lists of accepted and rejected events to the Data Flow Manager node. The task of the Data Flow Manager consists of sending a *Clear* message to the Read Out System to delete rejected events from the buffers and of assigning an event building node to each accepted event.

Event building nodes are called SFIs (Sub Farm Input) since they pass on the full event to subfarms of the EF. Each SFI is connected directly to the central event building switch via a Gigabit Ethernet link as well as to an EF subfarm by a second network interface.

At any given time an SFI will typically handle several events. It sends out data requests to each ROS and combines the replies into a single event. It applies a traffic-shaping algorithm using a credit system to avoid overloading its input link.

Using today's technology about 60 event-building nodes will be needed for the full ATLAS system assuming that each link in the system runs at most at 60-70 % of its full capacity.

6.6.4 Event Filter The EF subfarms receive full events from the event building nodes. The communication protocol supports multiple SFIs per subfarm, so the size of the subfarms is flexible and can be adjusted as needed.

The event is received into a shared memory buffer by one process while the HLT software runs in separate processes. The software framework is flexible enough to accomodate a variety of *Processing Tasks*. These include not only trigger algorithms but also calibration or monitoring tasks.

Accepted events are sent to a Sub Farm Output (SFO) node. It is responsible for forwarding the event to mass storage. The current implementation used in the testbeam writes a series of disk files for each run and also interfaces to the CASTOR (CERN Advanced Storage Manager) system [8].

Due to the modular design of the software the pieces can be arranged in various ways for installation or commissioning tasks. For example, the SFI and SFO functionality can be combined if no processing at the Event Filter level is required or the SFO functionality can be integrated into an EF task.

7 Summary

The ATLAS trigger system has been designed to meet the challenging task of reducing the large event rates in the LHC environment to an output rate of about 200 Hz, while being highly efficient for the interesting physics channels. The LVL1 trigger hardware has been successfully prototyped and is nearing production. The HLT consists of a flexible software framework which will allow one to optimize the selection process. All hardware components of the HLT/DAQ system have been studied in testbeds with a size of up to 10 % of the final system and it has been shown that a solution based mostly on commodity components will provide the performance necessary for ATLAS. The results of these studies form the basis of the recently published HLT/DAQ Technical Design Report [9].

References

1. ATLAS Collaboration: ATLAS Technical Proposal, CERN/LHCC 94-43, 1994.
2. ATLAS Collaboration: ATLAS Detector and Physics Performance Technical Design Report, CERN/LHCC 99-14 and 99-15, 1999.
3. ATLAS Collaboration: ATLAS First Level Trigger Technical Design Report, CERN/LHCC 98-14, 1998.
4. ATLAS Collaboration: ATLAS High-Level Triggers, DAQ and DCS Technical Proposal, CERN/LHCC 2000-17, 2000.
5. Athena: <http://atlas.web.cern.ch/Atlas/GROUPS/SOFTWARE/OO/architecture>.
6. G. Barrand et al.: GAUDI: A Software Architecture and Framework for building HEP Data Processing Applications, CHEP 2000 proceedings, Padova, 2000.
7. O. Boyle et al.: The S-LINK Interface Specification, <https://edms.cern.ch/document/110828>.
8. CASTOR - CERN Advanced Storage Manager, <http://castor.web.cern.ch/castor>.
9. ATLAS Collaboration: ATLAS High-Level Triggers, Data Acquisition and Controls Technical Design Report, CERN/LHCC 2003-22, 2003.