ATL-DAQ-2004-004
04 February 2004

The ATLAS *B*-Physics Trigger

Simon George *(Representing the ATLAS Collaboration)*

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

Abstract.
This paper gives an overview of the ATLAS Trigger and Data Acquisition (T/DAQ) system with an emphasis on *B*-physics capabilities. It describes recent work on how to maintain the *B*-physics programme within some constraints that have arisen: a higher target start-up luminosity, an incomplete configuration of the detector at start up, and cost constraints for the T/DAQ system. It also shows how the High Level Trigger (HLT) software has advanced and gives some results of new performance measurements.

INTRODUCTION

ATLAS is a general-purpose experiment for the Large Hadron Collider (LHC), currently under construction at the European Laboratory for Particle Physics (CERN). The LHC will collide protons with 14 TeV centre-of-mass energy within bunches crossing at 40 MHz. At start up in 2007, the target peak luminosity is $2 \times 10^{33}\,\mathrm{cm}^{-2}\mathrm{s}^{-1}$, rising to the full design luminosity of $1 \times 10^{34}\,\mathrm{cm}^{-2}\mathrm{s}^{-1}$ after a few years' running. There will be an average of 4.6 and 23 interactions per bunch crossing for initial and design luminosity respectively.

The ATLAS first-level trigger has to make a decision every 25 ns. The final event rate for mass storage must be reduced to ~ 200 Hz in view of the large event size ($O(1)$ MByte) and the available offline storage and processing resources. About 1% of collisions produce a $b\bar{b}$; the challenge for the trigger is to select those events of most interest for the ATLAS *B*-physics programme, within the limited trigger resources (processing power and network bandwidth) available. The ATLAS *B*-physics trigger is based on a lepton signature at the first level, which can optionally be accompanied by additional, lower transverse-energy signatures of leptons and jets. These signatures are refined in the higher trigger levels where specific decays are reconstructed. The ATLAS physics programme is described in Ref. [1]. More details on the design of the ATLAS T/DAQ system can be found in Refs. [2, 3].

The whole ATLAS detector is described in Ref. [4] and the inner detector in Ref. [5]. The largest volume part of the detector is the muon spectrometer incorporating toroid magnets with monitored drift tubes (MDT), resistive plate chambers (RPC) and thin gap chambers (TGC). Inside this is the calorimeter system, comprising a liquid-argon electromagnetic and hadronic endcap and electromagnetic barrel, and a scintillating-tile hadronic barrel. Within this is a solenoid providing a 2 T field and housing the inner detector (ID), which consists of silicon pixels nearest the beam pipe, then silicon microstrip detectors (SCT) at intermediate radii, and a straw-tube transition-radiation tracker (TRT) as the outermost part.

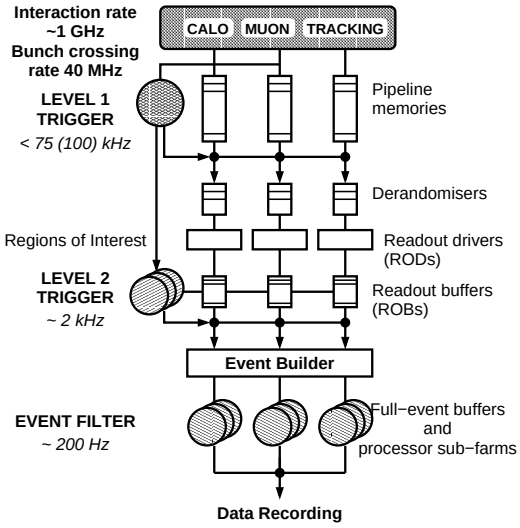

FIGURE 1. Overview of the ATLAS T/DAQ System.

THE ATLAS T/DAQ SYSTEM

Overview

ATLAS has a three-level trigger, shown in Fig. 1. The first-level (LVL1) has to operate at the bunch-crossing rate of 40 MHz, and takes a decision based on coarse-

granularity calorimeter towers and muon trigger stations in less than 2.5 μs, while detector data are held in pipeline memories. It reduces the event rate to less than 75 kHz, although this will be limited to 25 kHz LVL1 output for initial data taking due to resource limitations in the HLT/DAQ system. Accepted events are transferred from the pipeline memories of the detector to readout buffers (ROB) where they are held pending a second-level (LVL2) decision.

LVL2 has access to full-granularity data from all detectors (including the ID) by requesting them from the ROBs over the network. Information on the Regions of Interest (RoI) identified by LVL1 is transferred to LVL2, where it is used to reduce the amount of data requested to a few percent of the full event. Specialised trigger algorithms are used with an emphasis on fast rejection, to achieve the target average decision time of about 10 ms. Accepted events are classified by the signature(s) they match.

After a LVL2 accept, event fragments from the ROBs are transferred to the Event Builder at about 2 kHz. Full events are then distributed to the third-level, or event filter (EF), processing farms. Here the event selection is further refined according to the LVL2 classification by performing more sophisticated reconstruction and using more detailed calibration and alignment data. The processing can take on average about one second. The final output rate is of the order of 200 Hz.

Together, LVL2 and EF are known as the High Level Trigger (HLT) and have a large part of their infrastructure (hardware and software) in common.

Region of Interest mechanism

The Region of Interest (RoI) mechanism is a fundamental architectural principle of the ATLAS trigger.

The LVL1 event selection is based mainly on local signatures, identified at coarse granularity in the muon detectors and calorimeter. Significant further rejection can be achieved by examining the full granularity data from muon, calorimeter and inner detectors in the same localities.

The RoI is the geographical location of the LVL1 signature. It is passed to LVL2 where it is quickly translated into a list of corresponding ROBs. The data are requested sequentially, one detector at a time, so that only as much data are transferred as are needed. In most cases events are rejected and the data from all detectors in the RoI are not needed.

The RoI mechanism is a powerful way to gain additional rejection before event building. The benefit is an order of magnitude reduction in the dataflow bandwidth, at the small cost of more control traffic.

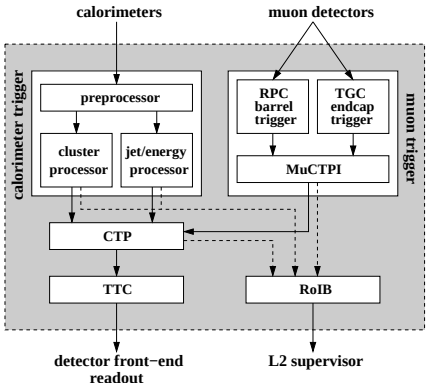

FIGURE 2. The LVL1 Trigger system.

Simulations of the LVL1 trigger show that there will be about 1.6 RoIs from triggers per event. These can be accompanied by RoIs with lower thresholds to give additional guidance to the HLT.

LVL1 Trigger

The LVL1 trigger works by identifying basic signatures that are consistent with interesting physics: muon tracks in the muon spectrometer, electromagnetic/hadronic/jet-like calorimeter clusters, and missing/sum transverse energy (E_{T}). The trigger decision is based on multiplicities and thresholds of these signatures.

LVL1 is a hardware trigger. The logic is implemented in programmable and custom electronics (FPGAs and ASICs). These will have programmable thresholds which can be configured to suit the expected conditions of a data-taking run. As shown in Fig. 2, there are separate muon and calorimeter triggers, whose results are combined in the Central Trigger Processor (CTP). If the event is accepted, signals from the CTP, muon and calorimeter triggers are sent to the Region of Interest Builder (RoIB) where they are prepared for LVL2. The CTP also signals the detector front-end readout via the Timing, Trigger and Control (TTC).

High Level Trigger

The HLT is built mainly from commodity equipment (PCs, Ethernet switches), with a few custom components such as the RoIB and ROBs in the DAQ system. The majority of the implementation is in software; indeed the HLT is a large and complex software-engineering project. Fig. 3 shows the main components and high-level interactions of the system that has been designed.

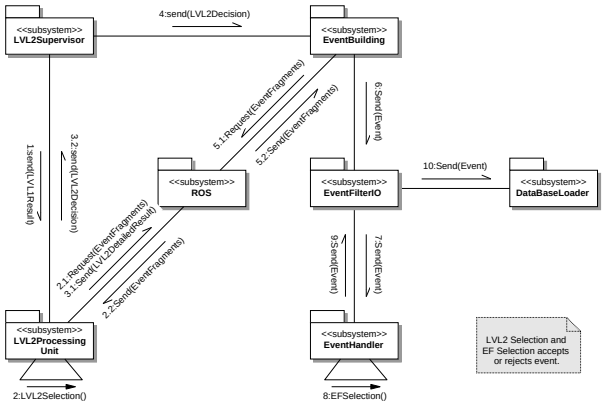

FIGURE 3. Overview of HLT components and the messages they exchange.

Note that the Read Out Subsystem (ROS) contains the ROBs.

To optimise processor usage, the LVL2 processing unit can run several worker threads at once on a dual-processor computer. Each worker thread will run a copy of the event-selection software which processes events independently. This allows idle cycles to be used to process another event while waiting for the ROS to return event data.

Event Selection Software

LVL2, EF and the offline software ('Athena') share a common event-selection software suite. It is used both to perform the real event selection online, and to emulate the HLT offline. The latter is also used as a development environment for trigger algorithms.

The necessary differences between these environments are hidden behind interfaces, so for example trigger algorithms have no knowledge of whether event data are being accessed from a file, database, memory, or ROBs over a network. Thus the principle of re-use of offline software as far as possible has been established. In current prototypes, offline software provides the basic framework and services, the data unpacking mechanism, part of the event data model and detector description. LVL2 has its own specialised algorithms and data preparation. All software used in LVL2 must meet the additional constraints of being multi-thread safe and efficient (i.e. while there are available processor cycles, the event input rate that can be sustained by one processor should scale more or less linearly with the number of threads, showing that performance is not limited due to locks). The EF re-uses offline reconstruction algorithms; to facilitate this they are currently being redesigned in the form of a 'toolkit' as described in Ref. [6].

A tool to aid integration of offline software and trigger algorithms with the LVL2 test beds has been developed: 'AthenaMT' provides a single-PC test environment and automatically sets up emulation of the full LVL2 system including the multiple worker threads. Simulated data for tests have a fully simulated detector response and the data are formatted as they are expected to come from the readout electronics. This ensures that the additional processing time to deal with data in this format is taken into account.

Both LVL2 and EF use a common 'steering'. This is a RoI-driven framework which calls the algorithms necessary to find features in subdetectors and combine these to establish hypotheses. Early rejection of events is achieved by comparing the event with a Trigger Menu (lists of signatures) after each step in the processing: the event is rejected if it is no longer consistent with any signature in the menu.

B-PHYSICS TRIGGER STRATEGY

At the peak target luminosity of $2 \times 10^{33}\mathrm{cm}^{-2}\mathrm{s}^{-1}$ in the initial running (and the eventual target of $1 \times 10^{34}\mathrm{cm}^{-2}\mathrm{s}^{-1}$), the B-physics trigger will be limited to a di-muon trigger. This is based on muons found in LVL1 with a p_{T} threshold of around 6 GeV. At LVL2 and EF, the regions of interest will be confirmed in the precision muon chambers, the tile calorimeter, then extrapolated to the ID. The trigger selection is based on identifying candidates for specific exclusive or semi-inclusive decays, for example $J/\psi \to \mu^+\mu^-$ and $B_{d,s} \to \mu^+\mu^-(X)$.

As the luminosity falls to around $1 \times 10^{33}\mathrm{cm}^{-2}\mathrm{s}^{-1}$, further B-triggers can be added based on a single muon trigger plus at least one additional trigger from the semi-inclusive reconstruction of specific decay candidates, for example $J/\psi \to e^+e^-$, $B \to h^+h^-$ and $D_s^+ \to \phi\pi^\pm$. Two strategies have been investigated for these additional triggers: 'RoI-guided' and 'full-scan'. In both cases the LVL1 muon is confirmed at LVL2 using the tile calorimeter and ID.

The baseline, RoI-guided strategy extends the fundamental architectural principle of the ATLAS trigger to B-physics. The LVL1 trigger provides information on low-E_{T} EM and jet RoI in addition to the muon trigger. At LVL2, tracks are reconstructed in the ID within these RoI and the reconstructed track information is used to search for candidates with specific decays. Track reconstruction and B-decay candidate selection is repeated in the EF, using the LVL1 RoIs again or the LVL2 tracks as seeds. The EF selection is refined by fitting tracks to include secondary vertex information.

This approach significantly reduces the resources

needed for the B-physics trigger compared to the second approach. However, there could prove to be too many low-threshold jet and EM RoIs, and/or the inefficiency for identifying RoIs corresponding to the B-decays of interest could be unacceptably low, although initial studies based on fast simulation are encouraging. Studies based on full simulation and the full-scan strategy show that, for example, tracks from $D_s^+ \to \phi\pi^\pm$ are contained within a cone of about 1.2 radians by 1 unit of pseudorapidity.

The full-scan strategy does not use RoIs to seed searches for B-decay products. Instead, LVL2 track reconstruction is performed in the full acceptance of the SCT and pixels and the resulting tracks are used to form $B \to h^+h^-$ and $D_s^+ \to \phi\pi^\pm$ candidates. Given additional resources, the TRT can be scanned for tracks in order to identify $J/\psi \to e^+e^-$ candidates. The EF can perform another full scan or use the LVL2 tracks as seeds. This approach is expected to be more efficient than the baseline above, but at a significantly higher resource cost.

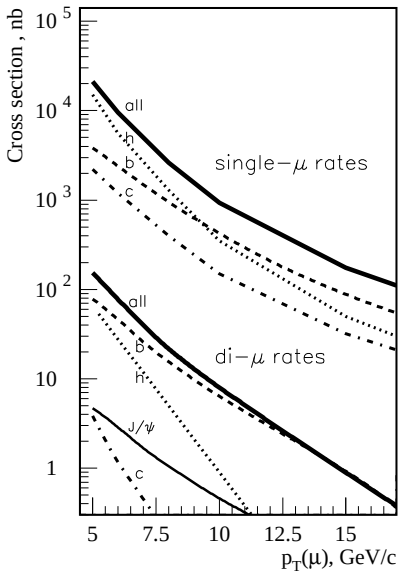

FIGURE 4. Single and di-muon cross sections.

Start-up conditions

When the LHC starts up, the initial beam conditions and the state of the ATLAS detector will not be optimal. Anticipated problems include: luminosity will vary somewhat from fill to fill, variable beam-related background, the ATLAS detector will be incomplete (see Ref. [5]), the detector will need understanding and tuning, and the T/DAQ processing capacity and bandwidth may be limited.

The target peak luminosity for the LHC initial running is $2 \times 10^{33}\,\mathrm{cm}^{-2}\mathrm{s}^{-1}$. However, luminosity will fall by a factor of about two during the coast. The T/DAQ system will clearly be built with sufficient capacity to meet the requirements of triggering on 'discovery' physics at the target peak luminosity, which will in particular require more processing power. When the luminosity drops, demand on trigger resources will also fall and some spare capacity can be used for B-physics triggers.

To cope with the initial running conditions, trigger algorithms must be robust with respect to noise, misalignment, and be able to operate effectively with the reduced detector that will be available at start-up time. The configuration must be flexible to adapt thresholds, pre-scales and other parameters to cope with varying noise, luminosity, etc.

B-physics Trigger Rates

The estimated B-physics trigger rates for the signatures discussed in this paper are shown in Table 1. The left column applies to higher luminosities where the strategy is based solely on di-muon triggers, while the right column shows the increased coverage planned for lower luminosities. Further details can be found in Refs. [2, 7]. Note that these are only a small fraction of the total LVL2 and EF accept rates.

Muon triggers

The LVL1 muon trigger has logic designed for both low and high-p_T thresholds; low-p_T is the one used for B-physics. It uses data from the two inner RPC stations (barrel) and two outer TGC stations (endcap); each station contains two layers and each layer has two projections. The algorithm starts with hits in the outer station and looks for coincident hits within a window (determined by the p_T threshold) in the inner station. Three out of four layers must contain hits in each projection.

The main single-muon background comes from π/K decays in flight. These can mostly be rejected at LVL2, by matching muon tracks to ID tracks and applying sharp p_T cuts. Fig. 4 shows how the prompt-muon cross section falls steeply with increasing muon p_T (and the background drops off even more quickly), which means that the trigger rate can be controlled by fine-tuning of the threshold. Simulations predict that a p_T threshold of 6 GeV has a LVL1 rate of about 20 kHz at $1 \times 10^{33}\,\mathrm{cm}^{-2}\mathrm{s}^{-1}$.

For di-muon triggers, lower thresholds are possible if the rates are low enough: $p_T > 5$ GeV in the barrel, 3 GeV in the endcaps. At $1 \times 10^{34}\,\mathrm{cm}^{-2}\mathrm{s}^{-1}$ the di-muon

TABLE 1. Estimated B-physics Trigger Rates.

	$2 \times 10^{33}\,\mathrm{cm}^{-2}\mathrm{s}^{-1}$		$1 \times 10^{33}\,\mathrm{cm}^{-2}\mathrm{s}^{-1}$	
Trigger	**LVL2**	**EF**	**LVL2**	**EF**
$B_{d,s} \to \mu^+\mu^-(X)$ $J/\psi \to \mu^+\mu^-$	200 Hz	small 10 Hz	100 Hz	small 5 Hz
$D_s^+ \to \phi\pi^\pm$	–	–	60 Hz	9 Hz
$B \to \pi^+\pi^-$	–	–	20 Hz	3 Hz
$J/\psi \to e^+e^-$	–	–	10 Hz	2 Hz
Total	**200** Hz	**10** Hz	**190** Hz	**20** Hz

trigger rate should be below 1 kHz for a 6 GeV threshold, only a small fraction of the total LVL1 rate. This is dominated by heavy-flavour decays but is subject to large uncertainties due to its sensitivity to very low p_{T} muons in the endcap and a small fraction of single muon events that suffer from double counting. Di-muon triggers are used to select decays such as $B^0 \to J/\psi(\mu^+\mu^-)K_S^0$ and $B_{d,s} \to \mu^+\mu^-(X)$.

At LVL2, muon triggers are first confirmed using the precision chambers (MDT). Better track measurement gives a tighter threshold which, due to the steeply falling muon p_{T} spectrum, significantly reduces the trigger rate. Further rejection is obtained by extrapolating tracks to the ID, requiring a track with closely matching z, ϕ, and p_{T}. Track matching helps to reject muons from π/K decays and using the ID further improves p_{T} resolution for prompt muons. In the barrel this has been fully simulated and studied in detail. At luminosity of $1 \times 10^{33}\,\mathrm{cm}^{-2}\mathrm{s}^{-1}$, the LVL2 rate for a single-muon trigger with a threshold of 6 GeV and $|\eta| < 1$ is about 2 kHz, about half of which is remaining π/K decays and the rest heavy-flavour decays. At $2 \times 10^{33}\,\mathrm{cm}^{-2}\mathrm{s}^{-1}$, a higher threshold of about 8 GeV gives similar rates. Extrapolated to the full detector, the resulting rate is estimated to be about 5 kHz.

The EF will perform near offline-quality track reconstruction, vertex fit and mass cuts, for example to select $J/\psi \to \mu^+\mu^-$.

LVL1 jet and EM cluster trigger

The LVL1 calorimeter trigger uses $\sim$ 7000 dedicated, relatively coarse calorimeter '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 clusters which satisfy criteria for EM, tau/hadron and jet selections. A E_{T} threshold is applied, with isolation in both layers to provide powerful jet rejection for EM/tau/hadron triggers.

The average multiplicity of LVL1 EM and jet RoI in events with a muon trigger is important as it determines the fraction of the detector that must be read out in order to make the B-trigger decision. This has been studied using a sample of $b\bar{b}$ events containing a muon with $p_{\mathrm{T}} > 6$ GeV, produced by a fast simulation which contains a detailed simulation of the calorimeter and LVL1 electronics. With a jet E_{T} threshold of 5 GeV there are on average about 2 RoIs per event. An EM cluster threshold of 2 GeV yields on average about 1 RoI per event.

Hadronic final states

Tracks, found in the SCT and pixels by either by the RoI-guided or full-scan strategies, are used to reconstruct the B-decays $B \to h^+h^-$ and $D_s^+ \to \phi\pi^\pm$. Studies show that the RoI-guided method would be efficient for a B hadron with $p_{\mathrm{T}} > 15$ GeV. At LVL2, kinematic and topological cuts are made to reduce combinatorial background. At EF there is more processing time so the resolutions of reconstructed track parameters are better and a vertex fit is available. The rate can therefore be further reduced by tighter mass cuts, and using decay-length and vertex-quality cuts. Studies show that the tracking algorithms are very robust with respect to a missing middle pixel layer (as in the initial detector configuration) and anticipated levels of misalignment at LVL2.

Muon-electron final states

To select channels such as $B^0 \to J/\psi K_S^0$ with $J/\psi \to \mu^+\mu^-$ or $J/\psi \to e^+e^-$ with opposite side electron or muon tag respectively, electron identification is required. ATLAS provides two options for this: use the RoI-guided strategy to find silicon tracks guided by EM RoIs, or perform full, unguided reconstruction of tracks in the TRT. The RoI-guided method is much faster, since typically only about 0.3% of the full ID is reconstructed. However, the lowest nominal threshold possible in the calorimeter with acceptable RoI multiplicity is 2 GeV and this is not efficient until a higher energy, significantly higher than the minimum threshold possible with a full scan of the

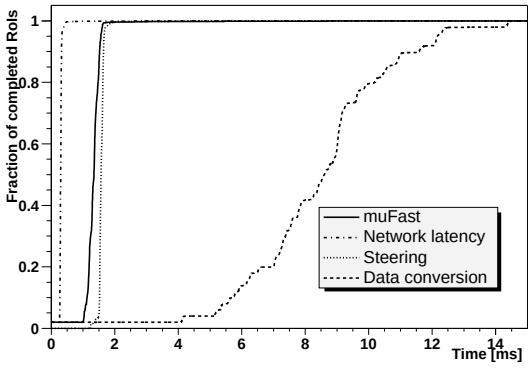

FIGURE 5. Testbed measurements of the LVL2-muon trigger processing time.

TRT. The efficiency to find a separate RoI for both e^- and e^+ with $p_T > 3\,\text{GeV}$ is about 80%.

PERFORMANCE

Testbed performance measurements

Software to perform the LVL2 muon reconstruction in the barrel of the muon spectrometer was set up on a LVL2 farm node, a PC with dual 2.2 GHz Xeon processors, with separate PCs running the LVL2 supervisor process and ROS emulation.

As shown in Fig. 5, the algorithm time itself is very small, and the time to access data over the network is also modest, since typically data are fetched from 1-2 ROBs. Data preparation time dominates the processing time despite the fact that only data within the RoIs need to be prepared – in the absence of the RoI mechanism the time would be much larger. This time is spent decoding binary data from the RPC and MDT detectors, converting logical addresses to geometrical positions taking into account the complex geometry of the muon detectors, and instantiating collections of objects, using the 'offline' framework and services. There is still considerable scope for code optimisation. The results are pessimistic because of a conservative level of radiation in the cavern and the fact that only muon events, not background events, have been simulated. Even so, these results already show adequate performance, since computers will be bought in 2006 when 8 GHz is expected to be the commodity processor speed.

Resource estimates

A lot of care has been taken to ensure that the studies provide a realistic basis for resource calculations, such as simulation of realistic raw data, timing of network requests and data preparation, and the reduced rate of later steps in the sequential-processing scheme. This has lead to the conclusion that the LVL2 budget of $\sim$10 ms will be achievable for high-p_T 'discovery' physics. The time for muon and calorimeter triggers, for which studies are most advanced, is already acceptable; this work will be extended to cover all detectors and triggers.

The target time of 10 ms at 25 kHz for LVL2 yields a requirement of 250 processors, scaling to 750 for the full 75 kHz LVL1 rate. Note that the number of EF processors required is much more.

CONCLUSIONS

The latest picture of ATLAS and LHC at start up is less favourable for B-physics than in the past, but ATLAS has responded with a variety of flexible trigger schemes to make the most of it. The RoI-based strategy allows a full programme at modest resource cost, with slightly reduced trigger efficiency. ATLAS will take advantage of beam-coast and lower-luminosity fills to increase the coverage of B-physics at lower luminosities. Algorithms are robust enough for the initial detector and conditions. For the muon spectrometer and calorimeter, there is now full simulation of realistic raw data and a full chain of algorithms to retrieve, unpack and process them.

Further studies of the RoI strategy will be made and more of the software will be tested and optimised through test bed and test beam deployment. Thus the T/DAQ system will be prepared for commissioning in 2006 and ready for first collisions in 2007 when it is hoped to record B-physics data from ATLAS.

The ATLAS HLT, DAQ and Controls TDR (Ref. [2]) was recently recommended for approval by the LHCC.

REFERENCES

1. Eerola, P., ATLAS B Physics Performance Update (2003), in these proceedings.
2. ATLAS Collaboration, *ATLAS High Level Trigger, Data Acquisition and Controls Technical Design Report*, CERN/LHCC/2003-22, 2003.
3. ATLAS Collaboration, *ATLAS Level-1 Trigger Technical Design Report*, CERN/LHCC/98-14, 1998.
4. ATLAS Collaboration, *ATLAS Detector and Physics Performance Technical Design Report*, CERN/LHCC/99-14, 1999.
5. Moser, H.-G., Status of the ATLAS Inner Detector (2003), in these proceedings.
6. Parodi, F., ATLAS software (2003), in these proceedings.
7. Baines, J., *Nucl. Phys. Proc. Suppl.*, **120**, 139–144 (2003).