PROCEEDINGS OF SCIENCE

Data Quality Monitoring in the ATLAS Inner Detector

Martin White[*]
On behalf of the ATLAS collaboration
Cavendish Laboratory, University of Cambridge E-mail: white@hep.phy.cam.ac.uk

This article describes the data quality monitoring systems of the ATLAS inner detector. A brief description of the ATLAS detector is provided along with a summary of the software frameworks available for developers. There follows a tour of the packages developed by each subsystem of the inner detector, and a discussion of recent tests.

Vertex 2008, 17th International Workshop on Vertex detectors
Utö Island, Sweden
July 28-August 1, 2008

[*] Speaker.

1. Introduction

The Large Hadron Collider (LHC) at CERN, Geneva, is the largest particle accelerator and collider in the world. Protons are accelerated to an energy of 7 TeV, and collide head on at four points around the 27 km circumference tunnel previously used by the LEP electron-positron collider. At each of the four collision points, activity is monitored by one of four primary LHC experiments: ATLAS, CMS, LHCb and ALICE. The largest of these- ATLAS and CMS- are specifically designed to search for new physics beyond the Standard Model, and each is a complex combination of subsystems designed to measure the momentum and energy of the particles produced in collisions. The LHC final design luminosity of 1×10^{34} cm^{-2}s^{-1} leads to immense challenges in event selection and offline data storage, in addition to creating an extremely harsh radiation environment for the detectors themselves. In such complex conditions, it is essential to develop accurate monitoring systems in order to calibrate and commission each detector and to track their performance over time. Furthermore, since detector problems can easily fake new physics signatures, it is vital to have fool-proof mechanisms for classifying data sets as good or bad in order that physicists can declare new discoveries with confidence.

This talk reviews the data quality monitoring systems of the ATLAS inner detector (ID). Following an 'ID-centric' review of the ATLAS detector in section 2 , the relevant ATLAS software frameworks are discussed in section 3. This is followed by a description of the monitoring packages developed by each subsystem community in section 4 and a final summary.

2. The ATLAS Detector

2.1 Detector description

The basic design criteria of the ATLAS experiment [1] include very good electromagnetic calorimetry for measurements of photons and electrons, and full-coverage hadronic calorimetry for jet and missing transverse energy measurements. Efficient tracking is expected at both high and low luminosity, thus fulfilling an important requirement for high-p_T lepton-momentum measurements, electron and photon identification and heavy flavour identification. The detector has a large acceptance in η and almost full coverage in ϕ, whilst also being able to trigger at low p_T thresholds.

Accurate measurements of particles in accelerator experiments rely on precise measurements of particle momentum and charge, whilst also requiring information such as whether a particle was produced at the primary interaction vertex or at a secondary vertex (the latter is important for b-tagging and tau reconstruction). The ATLAS inner detector is designed to perform all of these measurements by observing the bending of particle tracks in the 2 T magnetic field of the central solenoid. The field points along the beam axis and hence particle tracks are bent according to the transverse momentum of the particle.

The ATLAS inner detector (shown in figure 1) is divided into three separate parts: a semiconductor pixel detector that provides high granularity near the vertex region, a semiconductor tracker (SCT) that utilises silicon micro-strip technology, and a straw tube tracker (TRT) that provides continuous track-following with much less material per point. The detector as a whole consists of

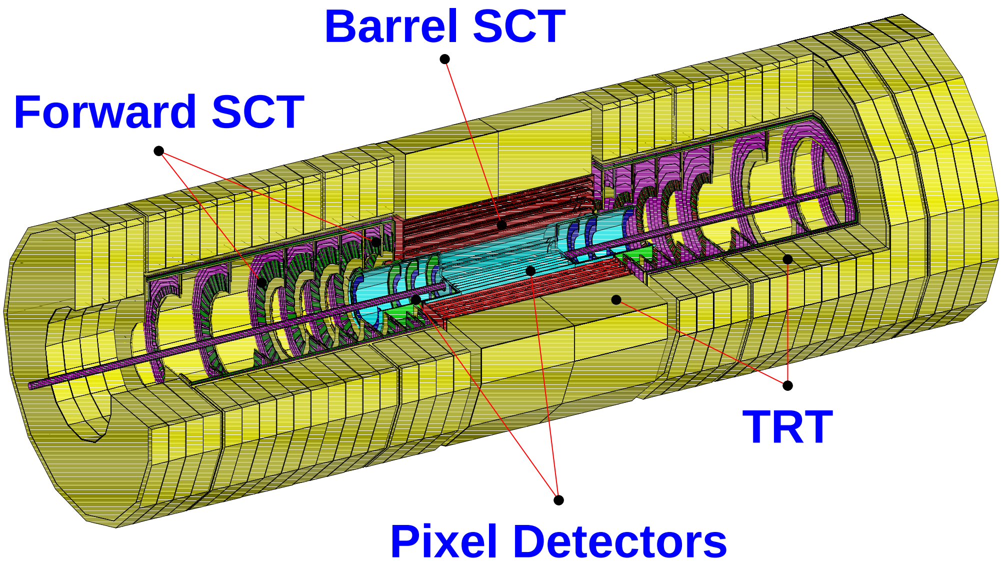

Figure 1: The ATLAS inner detector, taken from [2].

a barrel region and two endcap regions, and the design is such that no single component dominates the momentum measurement.

2.2 A brief guide to ATLAS dataflow

The path of data through the ATLAS experiment is illustrated in figure 2 and, for the uninitiated, is an enigmatic mess of acronyms. Essentially, data from the detector pass to a Read Out Driver (ROD) when the first level trigger (LVL1) has accepted an event. The ROD formats the data and sends a basic event fragment to a Read Out Buffers, which stores it until the level 2 trigger (LVL2) has reached a decision. If the event is good, a Read Out Subsystem moves the data block to the Event Builder (through a Sub Farm Input) which holds it until the Event Filter can apply the full ATLAS reconstruction software and produce hits, tracks and other derived quantities from the raw data. This represents the final result of the processing chain, and the events are subsequently written to disk through a Sub Farm Output.

The main aim of figure 2 is to illustrate the various points where one might choose to monitor data in ATLAS. Several monitoring frameworks have been developed, and each is capable of monitoring the data at any point in the data flow. Given that the rate of data flow drops at each stage of the trigger process, it is clear that there is often an advantage to monitoring data early in the chain if possible. For more sophisticated analysis, one must use data from the Event Filter, and hence must accept a lower rate of data collection.

3. ATLAS monitoring software frameworks

As is usual in large collaborations, several software frameworks have been developed in order to perform monitoring tasks in ATLAS. Some have merged over the years, but others have remained intact because they either serve a unique function or provide a vital cross-check when combined

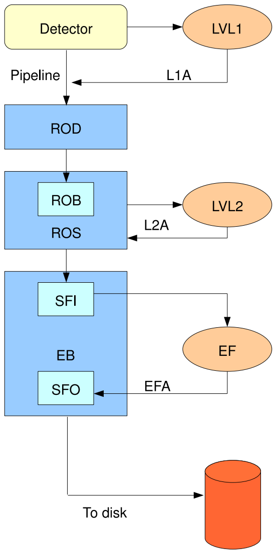

Figure 2: The path of data through the ATLAS experiment [3]. The abbreviations are defined in the text.

with another system. Before discussing these systems in depth, it is worth pointing out that there are three distinct ‘use cases’ for a general monitoring system:

1. **Online Running:** Whilst the detector is in operation, it is vital that shift workers and experts are able to observe the detector in real time in order to debug problems as they arise. Monitoring systems running online are forced to sample the data stream, and hence there will always be a limit to the statistics accumulated in online running.

2. **Offline running:** Once data has been collected on disk, one can perform much more thorough and detailed studies using the ATLAS full reconstruction software. This can in principle use a large number of machines (using GRID computing for example), and the output can be stored in offline databases in order to record detector parameters and classify the quality of datasets.

3. **Semi-offline running:** In ATLAS, there is a 24 hour delay between live data taking and the commencement of the full offline processing, and this can be used to store parameters that are needed *before* the full reconstruction is run. As an example, one might want to store, and subsequently veto, noisy detector components that might otherwise disturb the track reconstruction algorithms.

There are two dominant online monitoring systems in ATLAS. The GNAM framework [3] is capable of monitoring at all levels shown in figure 2 but is mostly used for obtaining fast feedback using relatively raw data such as hits, errors, bunch crossing information and very simple derived objects. This is computationally quick and can be performed at lower trigger levels to maximise the event rate. A second monitoring system uses the ATLAS Athena software framework. This is typically run at the Event Filter stage of the trigger process, leading to a much lower rate of approximately 1 Hz. In addition, some detectors have custom systems built into their RODs.

The online monitoring tools publish ROOT histograms in an online histogram server which can subsequently be displayed by a series of collaboration wide tools such as the Online Histogram Presenter [4]. A subset of histograms is also sent to an automatic checking system that is designed to highlight basic problems in the form of a traffic light reporting system.

Most offline monitoring in ATLAS uses the Athena software framework, often with the same code as the online system. The greater processing power available offline enables more detailed studies to be performed. There are also custom tools available in some subsystems that were developed before the use of Athena became more widespread.

The offline tools store histograms in ROOT files, as well as sending histograms to an offline version of the automatic histogram checking system described above. For permanent reference, output from the monitoring systems can be stored in a COOL database.

All systems have had ample opportunity for recent tests both in combined commissioning runs with cosmic rays and in staged challenges that pipe simulated data through the complete data flow chain (the so-called ‘Full Dress Rehearsal’ challenges).

4. Monitoring systems by subsystem

4.1 Pixel detector

The pixel detector community use three different monitoring systems, each for a different purpose. Online, their GNAM based system samples data from their ROSs at a rate of approximately 1 kHz, and monitors the total number of hits in each ROS (including the total number per event), errors flagged by each ROS and the occupancy for pixel layers. An Athena based system samples events from the SFI and reconstructs pixel clusters with a maximum rate of about 20 Hz. This plots occupancies at the module level as well as time over threshold information, the number of pixels per cluster and time dependent quantities. There is also a built in histogramming function at the hardware level, but this has yet to be used in any organised way.

Offline, their Athena code runs in an expanded version to provide more thorough data quality checks. The online and offline functionality has been tested extensively in the recent tests, and all packages were found to be in good shape. The focus has now shifted from code development to training a shift team for full time data taking.

4.2 SCT

Physicists working on the SCT detector also have a variety of monitoring systems at their disposal. A GNAM based system samples data fragments from the ROSs and plots errors, hits and hit coincidences (i.e. primitive spacepoints) in the form of maps and trends. This provides very fast basic feedback online and has been widely used during this year’s commissioning work to time in the detector and spot noisy detector components.

An Athena online monitoring system (whose development is dealt with in detail in [5] samples data from the SFIs and uses full event data (hits, clusters, tracks, etc) to provide more detailed checks at a much lower rate of approximately 1 Hz. The SCT Athena monitoring system is unique in having amendments that allow dynamic rebooking of histograms. This allows an expert to examine each strip of the SCT in real time during live data taking without exceeding the limited bandwidth allocation for sending histograms to and from the online server. The system automatically updates a display for shift workers to show the most problematic modules at the strip level. The wider system monitors noise occupancies, errors, efficiencies and track parameters such as residuals, pulls and fit χ^2 values. The system is running smoothly, and is being used to commission the detector and train shift workers for next year.

The Athena monitoring package also forms the basis for the semi-offline and offline monitoring strategies. For semi-offline processing, the SCT use a stripped down version of the Athena package (with tracking turned off) to log noisy and dead strips in a COOL database. For pure offline running, the package runs the full Athena reconstruction and plots similar quantities to the online package, though with greatly increased statistics.

4.3 TRT

The TRT use a combination of an Athena monitoring package and an earlier custom tool known as TRTViewer which has remained in use due to its excellent visual representation of the TRT. In its most recent incarnation, TRTViewer acts simply as display program for the output of the Athena monitoring package, or for histograms produced using a standalone ROS level monitoring package. Quantities of interest include the high level and low level occupancies of the detector (which has two thresholds), the average drift time, the number of tracks and the track residuals. Variables are plotted in a variety of ways, including vs time and vs ϕ angle.

The TRT monitoring systems are already interfaced with the ATLAS automatic data quality checker, and have performed well in recent commissioning tests.

4.4 Global Inner Detector Monitoring

In addition to monitoring systems developed for each subsystem, there is an Athena based monitoring package designed to investigate issues that involve combinations of sub-detectors such as detector synchronisation (in the form of bunch crossing and level 1 trigger information), correlations in detector occupancy, track parameters for global tracks, and angles between track segments for different detectors. This can run either offline or online, in which case it samples the event stream from the SFI.

The tool is split into three packages:

1. An alignment monitoring tool focusing on hit efficiencies, residuals and track properties.

2. A beam spot monitoring tool that compares the measured beam spot with reconstructed vertex positions.

3. A performance monitoring tool that looks at some classic particle resonances, and assesses the detector performance by comparing the mass, width and yield to the expected values.

As would be expected during early commissioning with cosmic rays, the first of these has proved most useful so far, though the other two tools continue to be developed using simulated data.

5. Summary

The detectors of the LHC are amongst the most complicated apparatus ever built, and accurate monitoring systems are essential to evaluate their performance, provide feedback to shift workers and track changes in detector parameters over time.

The monitoring systems of ATLAS use a variety of custom software frameworks, and comprise detailed monitoring packages for each sub-detector and a global tool designed to monitor synchronisation, alignment and global performance. All systems have been tested in commissioning runs involving cosmic rays and in simulation exercises, and all are at the final stages of debugging and commissioning.

As a final note, it is worth stating that the arrival of beam at the LHC will no doubt provide surprises for the monitoring systems, most immediately involving a jump in the average occupancy of the detectors above which problems must be detected. This must be used to reevaluate the performance of the algorithms that are already in place, but is not expected to cause major problems for any of the systems described in this article.

6. Acknowledgements

I am hugely grateful to Dr Richard Brenner for his invitation to present the talk on which this article is based. I thank the organisers of the Vertex 2008 conference for their spectacular hospitality, and the UK based Science and Technology Facilities Council for their partial funding of the visit. Finally, I would like to thank my colleagues in the ATLAS inner detector collaboration for helpful comments and discussions during the preparation of this final draft.

References

[1] G. Aad et al. *The ATLAS Experiment at the CERN Large Hadron Collider*. *JINST*, **3**:S08003, 2008.

[2] *ATLAS Inner Detector Technical Design Report*. CERN/LHCC/97-16 and CERN/LHCC/97-17, 1997.

[3] P. Adragna et al. *GNAM: A low-level monitoring program for the ATLAS experiment*. *IEEE Trans. Nucl. Sci.*, 53:1317–1322, 2006.

[4] P. Adragna, G. Crosetti, M.D. Pietra, A. Dotti, R. Ferrari, G. Gaudio, C. Roda, D. Salvatore, F. Sarri, W. Vandelli, and P.F. Zema. *The GNAM monitoring system and the OHP histogram presenter for ATLAS*. *Real Time Conference, 2005. 14th IEEE-NPSS*, pages 4 pp.–, June 2005.

[5] Martin John White. *Searching for supersymmetry with the ATLAS detector*. CERN-THESIS-2008-052.