Experience with the Custom-Developed ATLAS Offline Trigger Monitoring Framework and Reprocessing Infrastructure

Martin zur Nedden, *for the ATLAS Collaboration*

Abstract—The offline trigger monitoring of the ATLAS experiment is assessing the data quality and analyses those events where no trigger decision could be made. Within the offline monitoring, which is started shortly after the data acquisition has finished, the online data quality assessment is reflected. Additionally, a reprocessing system tests changes to the trigger software and configuration to ensure their stability and reliability before they can become operating. This note explains the activities performed to provide a flawless monitoring of the operation of the ATLAS trigger system and how to assess the quality of the recorded data.

I. INTRODUCTION

Powerful and sensitive monitoring is vital for such a complex system as the ATLAS trigger. Any occurring processing failures, misbehaviour of selection algorithms and data defects must be discovered immediately and made known to the relevant experts. Events for which the software-based trigger system cannot make a decision are automatically sent into a dedicated debug stream. The trigger decision is then taken offline, after these events have been recovered and analysed.

Any data not usable for physics analysis must be flagged accordingly. A complex system of data quality assessment has been developed which was used very successfully in 2011 and 2012. This is based on a first processing of a subset of data a few hours after recording, complemented by further monitoring performed during the bulk data processing approximately 48 hours later.

There is a steady flow of bug fixes, developments and improvements to trigger algorithms. In addition the trigger configuration changes constantly to adapt to the luminosity conditions as experienced in the first years of the Large Hadron Collider (LHC) operation. These changes need to be validated in the custom-developed reprocessing framework before they can become active in the online operation.

II. ATLAS TRIGGER SYSTEM

The three level trigger system[1] reduces the event rate from up to 40 MHz delivered by the detector to a recording rate of about 400 Hz. While maintaining high efficiency for physics processes is the central interest, a balance has to be kept between the individual trigger rates to share the

available bandwidth accordingly. In 2011 the level 1 (L1) trigger, realised as a hardware trigger, has reduced the event rate down to ≤ 60 kHz based solely on muon and calorimeter information with a latency of $2.5 \mu\text{s}$. The following two software-based higher levels, collectively known as the High Level Trigger (HLT), consists of the level 2 (L2) and the event filter (EF). A positive L2 trigger decision based on only 2% of the event information contained in a projective cone around regions of interest defined at L1 initiates the readout of the full event record and its transfer to nodes of the EF farm. The L2 output rate in 2011 was about 4-5 kHz with an average latency of about 40 ms. The EF applies optimised offline analysis algorithms to data. Events accepted by the EF are routed to permanent storage. The latency of the EF is about 4 s at an average output rate of about 400 Hz.

To ease physics analysis, the selected events are stored into one or more streams depending on which chain(s) caused the trigger. The streams are defined in terms of the trigger objects. Examples of streams include “muon”, “egamma” (for electrons and photons), “jets” (for tau-leptons, jets and missing transverse energy) and, in addition to the physics-motivated streams, those for calibration and debugging as well as an express stream. The latter fully overlaps with all other streams and contains a representative selection of events.

III. TRIGGER DATA QUALITY ASSESSMENT

The online and offline data quality (DQ) assessment is implemented in a dedicated Data Quality Monitoring Framework (DQMF) which is described elsewhere [2]. The DQ assessment is based on an automated analyses of appropriate histograms created and filled during data taking in the HLT trigger algorithms (online) or within the standard raw data reconstruction (offline). The online DQMF displays histograms of trigger objects in real time within a dedicated system to the ATLAS shift crew. The offline DQMF framework shows those on a web page for visual inspection after running over the express stream. It is also possible to implement and automatically apply algorithmic tests on the online and offline monitoring histograms in order to derive more automated DQ assessment results. The online DQ tests are reflected in the offline ones by digesting the same distributions with similar appropriate test algorithms to ensure that all possible issues can be caught. The DQ distributions can be overlaid and compared with a reference. Possible automatic algorithms to evaluate the DQ are checks on the RMS, Mean, as well as simple Fermi fits

Manuscript received June, 2012. This work was supported in part by the Federal Ministry for Education and Research of Germany.

Martin zur Nedden, Humboldt-Universität zu Berlin (e-mail: nedden@mail.cern.ch).

Valeria Bartsch, University of Sussex (e-mail: V.Bartsch@sussex.ac.uk).

and Kolmogorov distribution shape tests of the histograms. The optimisation of the mean and the threshold values for the automatic checks is an iterative process so that the assessment becomes more adequate with time as experience with data is accumulated.

The DQ distributions of the trigger signature objects can be roughly divided into

- transverse energy, E_T , or transverse momentum, p_T , distributions. E_T or p_T values are often used as threshold values in the trigger decisions, thus the distributions of these parameters are vital for the DQ assessment. The thresholds of the trigger requirements change with increasing luminosities. The reference plots for the E_T or p_T distributions need to be updated whenever the thresholds change.
- pseudorapidity, η , and azimuthal angle, ϕ , distributions of the number of trigger objects at different trigger levels can quickly identify any geometry dependence of the trigger efficiency and correlate trigger problems with detector issues. The distributions can be used to determine the exact location of the deviating parameter. A further diagnosis might relate the location of the deviation to a detector malfunction, e.g. a dysfunctional part of the calorimeter, or an algorithm problem. In case of a significant change in the observed parameter two dimensional η - ϕ maps can be used to diagnose the issue.
- invariant mass distributions. Di-muon invariant mass distributions show a clear e.g., J/ψ peak which can be picked up on trigger level. This distribution is clearly the most sophisticated distribution to be used on trigger level and is quite independent on the trigger configuration, luminosities or the running conditions.
- multiplicity distributions. The number of tracks or jet multiplicity can be checked against a reference. Generally, multiplicities depend on the luminosity and therefore their reference needs to be updated whenever there is a major luminosity change.

In case of any issue with the trigger operation or trigger algorithms appropriate DQ defects need to be set. The defects are either tolerable or intolerable. Defects can be vetoed when selecting the data used in a physics analysis. Intolerable defects are usually recorded when a defect is clearly an issue for most physics analyses and the events cannot be used. Tolerable defects are set if the perceived impact on the physics results is much smaller and the data can be taken into account with the given care [3].

In 2011 most of the defects set for the HLT were tolerable. Only 3% of data collected had an intolerable defect, see Fig. 1. The percentage of the tolerable defects is much higher. The high percentage reflects the fact that tolerable defects were used to document information about the trigger performance in the DQ defects. They are available for use in lists of good data taking periods for physics analyses if required and can safely be ignored in most cases.

IV. TRIGGER REPROCESSING

The trigger selection algorithms can be run offline to validate changes in the HLT, namely: software changes in the

L1			HLT						
Muon	Calo	CTP	electron	photon	muon	tau	jet	b-jet	missing E_T
99.0	100	99.8	99.3	99.3	100	99.9	98.6	99.9	99.3
Luminosity weighted relative fraction of good trigger data quality delivery during 2011 stable beams in pp collisions at $\sqrt{s}=7$ TeV between 13 March and 30 October (in %)									

Fig. 1. Luminosity weighted relative fraction of good quality data delivery by the various components of the ATLAS trigger system during LHC fills with stable beams in proton-proton collisions, and after turning the tracking detectors on. Data taking period was between March 13th and October 30th, 2011. The slight inefficiency in the HLT jet trigger is due to a misconfiguration associated with specific calorimeter conditions that occurred during two short data taking periods. However, all high p_T jets were still successfully triggered and recorded.

release (e.g. bug fixes), major changes of the trigger menu (e.g. new trigger algorithms or thresholds) and condition changes (e.g. improved alignment or calibration). Small changes are assessed by an automatic standardised nightly software test on a single computing node, with memory usage test and a count of the number of events accepted. However, if the proposed change influences data taking (e.g. if it is planned to deploy a new trigger software release into operation) the changes are tested by a trigger reprocessing that reruns the full HLT decision. Since a malfunction of the trigger software can cause unrecoverable losses of events, it is vital to ensure the reliability of the new software release. In addition the reprocessed data are reconstructed to obtain the offline objects and to allow a more thorough validation of the reprocessing and assessment of the trigger decision taken by the new trigger software under test. Sometimes the reconstruction step is skipped in case of trivial changes. The data from the trigger reprocessings and reconstructions are validated with similar tools as those used for the usual DQ assessment described in Sect. III. The reprocessings are run over about one million events of a data set collected with a very loose trigger selection, thus delivering a large enough data sample to validate the changes. Furthermore, the loose trigger requirements ensure an unbiased sample. The event selection by the HLT has to be deterministic otherwise changes between the trigger reprocessing to test and the reference one are expected.

Although currently only computing resources within the PANDA framework [4] local to CERN are used, the framework will allow the utilisation of distributed resources in future. This move has become necessary because of the larger data sets expected due to the increasing LHC luminosities. The developments toward the distributed resources are still ongoing.

V. DEBUG STREAM TREATMENT

Events where the HLT could not make a decision within the allocated time end up in the dedicated debug stream. The presence of an event in the debug stream usually indicates a problem in some aspect of the online system. However, presence of events in the debug stream does not necessarily imply a misbehaviour at collision conditions; timeouts may be caused at the presence of a very busy event and in that case the timeout ensures the system robustness. At the current phase of the experiment the debug stream is a way to spot all

Fig. 2. ATLAS data streams as recorded by the HLT. The data streams can be roughly divided into physics streams, calibration streams, debug streams and the express stream and several streams dedicated to special purposes. Whereas the physics, calibration and the express stream can be immediately used for physics analyses, calibration of the detector and a quick assessment of the recorded data, the debug stream is caused by a missing HLT decision. The debug stream data set is automatically recovered and if recovered streamed in the same manner as the physics streams. For most runs the percentage of events in the debug stream is lower than one per mille. Only 0.1% of the debug stream events remain unrecovered.

possible problems and weaknesses of the online system. Most of the debug stream events are caused by misconfiguration of the trigger software which do not happen often, however can send a lot of events into the debug stream at a given time. Other sources of debug stream events are timeouts or errors at L2 or EF. For most runs (where a run is a short data taking period often corresponding to a proton fill in the LHC) the percentage of events in the debug stream is lower than one per mille. Having dedicated tools to deal with the debug stream aims at identifying problems swiftly and reducing the turn-around time for fixing these problems. To achieve that, there is a documented procedure followed by an offline trigger expert in all events of the debug stream of each run. The goal is to achieve quasi real-time handling of the debug stream.

The raw data of the debug stream events is recorded and the trigger decision is rerun offline. If the recovery is successful the data streams obtained online (e/γ , μ , jets) are recorded as pictured in Fig. 2. It is up to the physics groups to decide how to deal with the recovered data streams.

A thorough analysis of the debug stream events before and after recovery aims to find trigger decisions which have failed, correlation between detector and data acquisition issues, and non-recoverable events. Errors in the trigger decision algorithms are followed up and lead to improved trigger algorithms. Issues with the non-recoverable events are fixed quickly and improved for the next data run. Only 0.1% of the debug stream events remain unrecovered.

VI. SUMMARY

The motivation for developing a custom online and offline trigger monitoring and reprocessing system is to provide a fast and reliable feedback on any issues with the ATLAS trigger system which do not become apparent during data taking and to evaluate the influence of any issues of the ATLAS trigger on the data quality of the collected data.

In 2010 and 2011 the thorough offline data quality checks were challenging because of the frequent changes of the instantaneous luminosity provided by the LHC which causes constant changes of the ATLAS trigger configuration. Although the algorithms and thresholds of the trigger needed to be constantly adapted only about three percent of all data had an intolerable defect resulting from the ATLAS trigger system. This is achieved by the strict validation of any trigger changes before they are used online with the help of trigger reprocessing. In addition after each run has finished events with a failed HLT decision were analysed in detail and recovered for physics analyses.

In 2012, after about two years of data taking the offline monitoring framework and reprocessing infrastructure have evolved into an essential and reliable component of the experiment operation.

REFERENCES

- [1] S. George [ATLAS Collaboration], PoS ICHEP **2010** (2010) 487.
- [2] S. Kolos, A. Corso-Radu, H. Hadavand, M. Hauschild and R. Kehoe, "A software framework for Data Quality Monitoring in ATLAS," J. Phys. Conf. Ser. **119** (2008) 022033.
- [3] T. Golling, H.S. Hayward, P.U.E. Onyisi, H.J. Stelzer, P. Waller, "The ATLAS Data Quality Defect Database System," Eur. Phys. J. **C72** (2012) 1960
- [4] T. Maeno [ATLAS Collaboration], "PanDA: Distributed production and distributed analysis system for ATLAS," J. Phys. Conf. Ser. **119** (2008) 062036.