Journal of Physics: Conference Series **331** (2011) 032009 doi:10.1088/1742-6596/331/3/032009

Commissioning of ATLAS Data Quality Infrastructure with First Collision Data

James Frost, on behalf of the ATLAS Collaboration

Cavendish Laboratory, University of Cambridge, J. J. Thomson Avenue, Cambridge, CB3 0HE, UK

E-mail: frost@hep.phy.cam.ac.uk

Abstract. During particle physics data taking, the data being recorded must be monitored and checked continuously and promptly. The data quality system of the ATLAS experiment at the Large Hadron Collider provides the means to investigate and monitor the data recorded. Over the course of 2010, this system has been commissioned successfully with first collision data. Data quality monitoring enables prompt investigation of the first pass of full event processing at the Tier-0 CERN computing centre, the validation of calibration and other reconstruction parameters, and any detector issues to be diagnosed and, wherever possible, fixed. Automatically filled histograms are checked by algorithms and against references, the results of which are stored as status flags in a database, and published onto a web server, where they can be inspected by shifters and experts. The data quality for approximately 100 data quality regions is determined, and propagated in a convenient format for data analysis teams.

1. Introduction

Knowledge of the quality of data recorded underpins all particle physics results. Careful monitoring of the data quality is necessary to understand data conditions and to enable the diagnosis and elimination of detector problems. During 2010, ATLAS has recorded $45.0\,\mathrm{pb}^{-1}$ of pp collisions, with a peak luminosity of $2.07 \times 10^{32}\,\mathrm{cm}^{-2}\mathrm{s}^{-1}$ (Fig. 1). Data-taking has been highly successful, with each of the ATLAS subdetectors have maintained high efficiency for physics data during stable beam collisions (Fig. 2). This paper describes the commissioning and status of the offline data quality monitoring system for the ATLAS experiment. ATLAS data-taking is divided into runs (nominally an LHC fill) and luminosity blocks (LB), discrete periods of time ($\sim$2 minutes) over which detector conditions are assumed to be constant. Data quality decisions are made and recorded using a 'traffic light' flag. This flag can be green, yellow, and red corresponding to good, caution and bad data quality, supplemented by grey (unknown - usually due to insufficient statistics) and black (subsystem off). Final data quality flags are derived after full reconstruction, with experts converging on the 'green' and 'red' states. This system allows shifters and experts to investigate the data promptly after they are recorded, in accessible formats. Information from fully reconstructed events processed at the Tier-0 farm, in addition to raw data quantities, can be investigated through histograms, available shortly after a run starts. Histograms produced undergo automatic checks, with the algorithms and thresholds being set per histogram, and with both result and histogram stored for future use. Top-level summary flags for the $\sim$ 100 data quality regions (often large section of subdetector) are stored in a database for subsequent data quality and analysis use. This allows the prompt

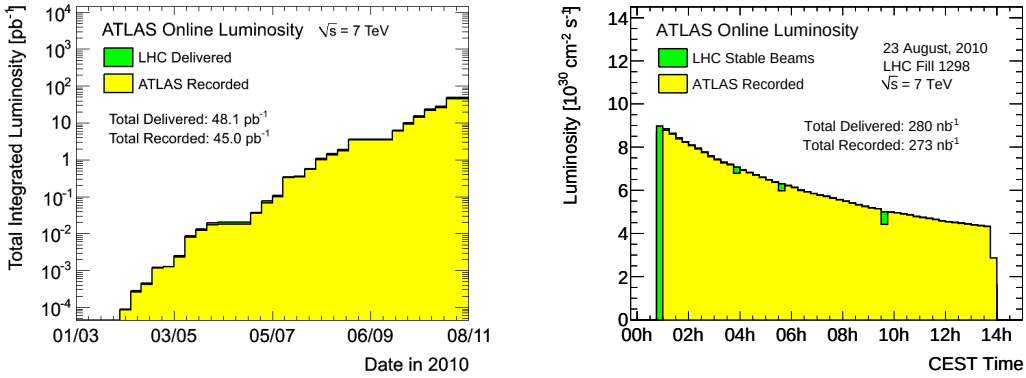

Figure 1. Cumulative luminosity versus week delivered to (green), and recorded by (yellow) during stable beams for 7 TeV centre-of-mass energy (left), and a sample instantaneous luminosity profile. The delivered luminosity accounts for the luminosity delivered from the start of stable beams until the LHC requests ATLAS to turn off the sensitive detectors to allow a beam dump or beam studies.

Inner Tracking Detectors			Calorimeters				Muon Detectors			
Pixel	SCT	TRT	LAr EM	LAr HAD	LAr FWD	Tile	MDT	RPC	CSC	TGC
99.0	99.9	100	90.5	96.6	97.8	94.3	99.9	99.8	96.2	99.8

Figure 2. Luminosity weighted relative detector uptime and good quality data delivery during 2010 stable beams in pp collision at $\sqrt{s} = 7\,\text{TeV}$ between March 30th and October 31st (in percentages). The inefficiencies in the calorimeters will largely be recovered in a future data reprocessing.

verification of derived calibration, alignment and reconstruction (e.g., beamspot) parameters, and the masking or fixing of any observed detector issues. These actions are derived from a processed 'express stream' containing a mixed subsample of the events and triggers, before full processing of the main bulk physics streams begins after the 36 hour 'calibration loop' period.

2. Data Processing, Reconstruction and Calibration Loop

The primary role of offline data quality monitoring is to review and verify promptly reconstructed ATLAS data from the Tier-0 computer farm at CERN. According to the ATLAS computing model [1], the first full reconstruction of data takes place at the Tier-0, though subsequent reprocessings use the resources of the Tier-1 centres. This reconstruction occurs shortly after the data are recorded, before they are staged onto tape and hence within a few days. A small subset of the data, the 'express stream', is reconstructed promptly at the Tier-0, as soon as the raw data and necessary conditions information are available, usually within an hour of the start of a run. Simultaneously, calibrations and alignment processes are run, using either the express stream, or dedicated calibration streams. Further reconstruction iterations on the express stream may be needed until the calibrations are signed off for full, 'bulk' reconstruction - consisting of main physics data streams, in combination with the express stream.

ATLAS data quality checks are primarily based on histograms, whose diskspace and memory requirements scale well with run length (or equivalently, the number of events). Histograms are produced by tools running within the ATLAS software framework Athena, and are able

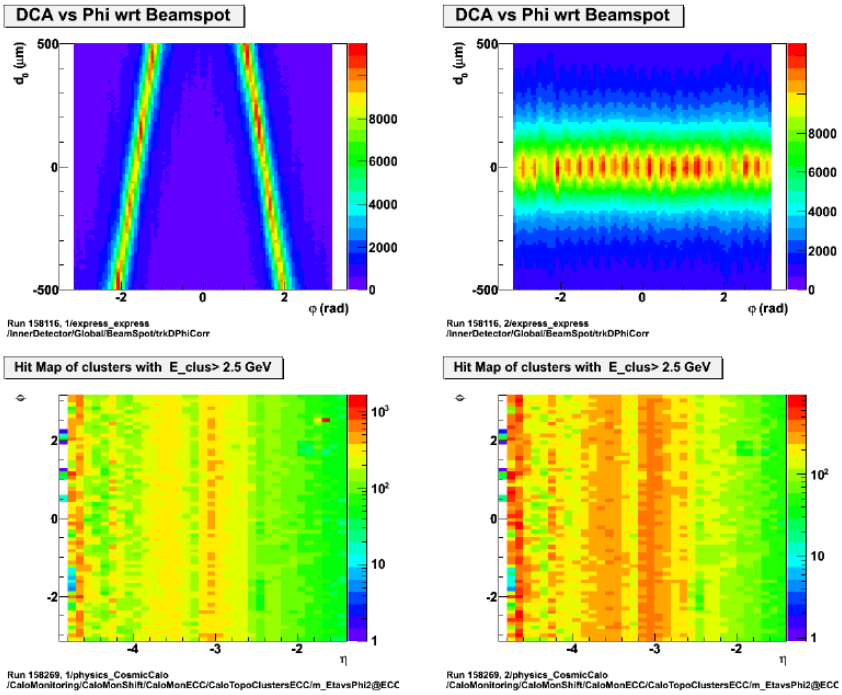

Figure 3. Histograms produced from data quality monitoring of the first (left) or second (right) express stream processing. The upper panes show the effect of the beamspot determination run during first pass processing on the distance of closest approach of tracks to the beamspot - a nominal beamspot is used for the initial pass, with the reconstructed beamspot input for subsequent passes.. The lower panes show the calorimeter cluster occupancy, and the ability to mask problematic channels during the calibration loop - the hot cell visible in the upper right corner of the initial processing being absent after the calibration loop.

to monitor both transient quantities from raw data and those stored in reconstructed files. Histograms can be generated for an entire run, or a shorter sub-run interval, and be made trigger-aware, taking entries only from events in which a specific trigger has been passed. The output format is that of a ROOT file.

Full data quality monitoring is run during each of these reconstruction passes. Histograms produced from the first pass can be compared easily with those of later passes to validate calibrations. Figure 3 shows sample histograms from first (left) and second (right) pass reconstruction. The upper row shows the effect of the beamspot determination on tracking variables, whilst the lower row shows the ability to mask hot calorimeter cells before the second, bulk processing. While reconstruction is ongoing, the available files are merged to sum the available statistics every 10 LBs, and again at the end of reconstruction; histograms for individual blocks of 10 LBs (a period of approximately 20 minutes) are also produced. During the merging step, post-processing algorithms may be run on the histograms produced, with the final, merged histogram files at the end of the run registered and accessible as a Grid dataset.

3. Data Quality Monitoring Framework

ATLAS online and offline data quality monitoring use a common infrastructure called the Data Quality Monitoring Framework (DQMF) [2, 3]. The primary concept is a 'tree' of data quality decisions, where the result flags of checks on individual histograms (the 'leaves'), propagate up through the directories of the tree. Each check consists of an input data histogram, decision

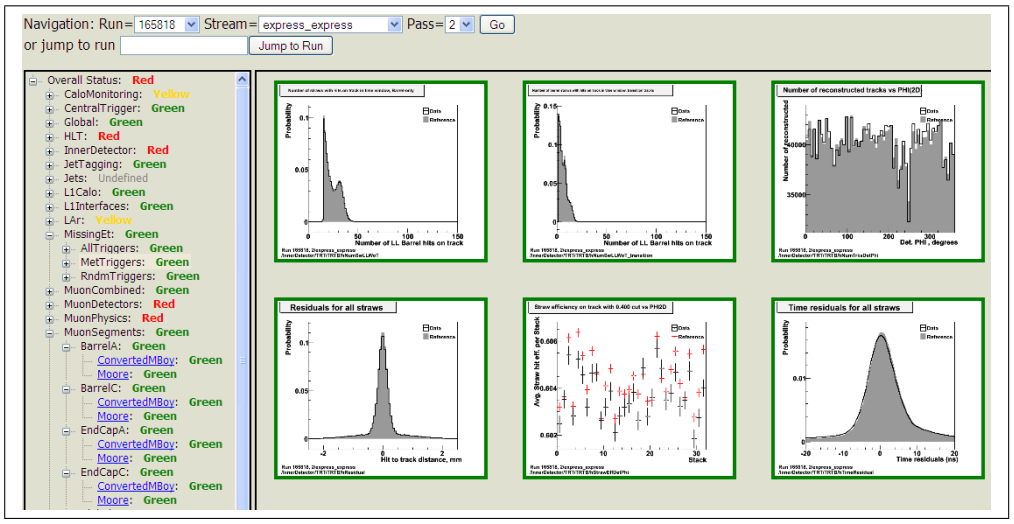

Figure 4. The ATLAS data quality web display. A specific directory of histograms and results is shown in the right pane; each histogram is compared to references, and has a border reflecting the status flag colour for its data quality checks. The left shows the data quality tree, its navigable subdirectories and their status derived from their daughter histograms.

algorithm, thresholds corresponding to a specific (e.g. red) status flag, and usually a reference histogram, useful to aid shifters and as an algorithm input. Histogram checks run the gamut of complexity, from simple ones to ensure a histogram is empty (for use with e.g. error histograms), through checks of means and variances, to shape comparisons via χ^2 or Kolmogorov-Smirnov tests. Relevant information about the histograms can be published in addition, whilst 'summary' algorithms determine how histogram results propagate up the data quality tree. The offline infrastructure runs as a standalone program run at the Tier-0, called **'han'**, which takes as input the merged ROOT histogram files and a binary configuration file which codifies all the checks and references. The output 'han' ROOT file is archived and contains everything necessary to reproduce the results. For further technical details, the reader is referred to [3].

Visualisation
The histogram analysis framework produces 20,000 histograms (100 MB) per trigger stream (usually between 6 and 10) for each run, which are checked and flagged by algorithms. Additionally 700 histograms (6 MB) are produced for each ten LB interval. The final data quality histogram results, in the merged ROOT file, are uploaded to a web display, which is implemented as a *CherryPy* web application, and available to shifters and experts to find and diagnose any detector issues present. Fig. 4 shows a sample web display from the second express stream processing of one of the 2010 collision runs. Initially these were stored as images and static HTML webpages, but the disk and CPU requirements per run became prohibitive. Given that few histograms are accessed frequently, and many not at all, a solution was implemented whereby the image files are generated on request and cached for some time. This is achieved through *PyROOT* functions (e.g. make a histogram image) using a pool of daughter processes.

4. Propagation to Analysis: Virtual Flags and Good Run Lists
In addition to the automated decisions made by both offline and online DQMF outlined in Section 3, data quality decisions are input by people. Currently, decisions are made by automated systems from DQMF, the online shifter in the control room and the detector control system

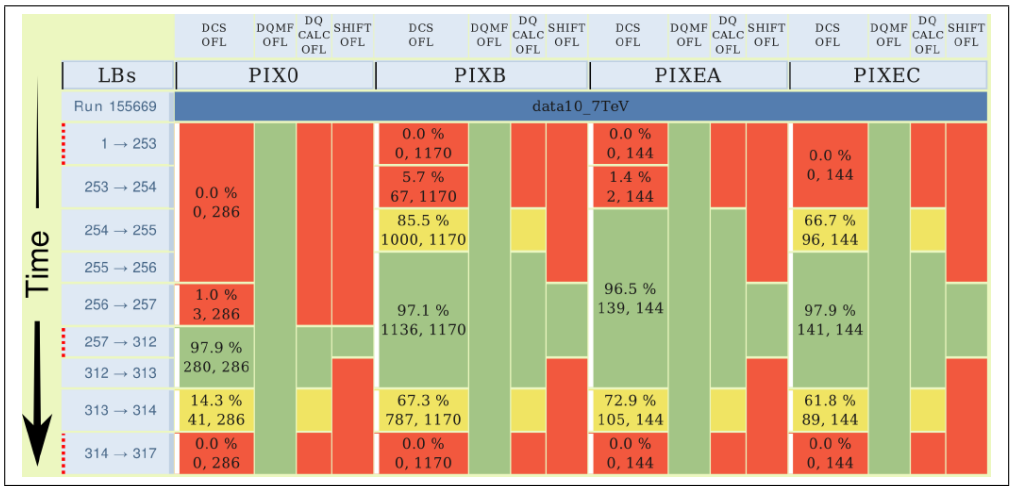

Figure 5. Data quality flags for the pixel regions (Layer 0, Barrel, Endcaps A & C) as determined by the automated data quality systems (*DCSOFL, DCMFOFL*), the data quality calculator (*DQCALCOFL*) and the final decision of the expert offline shifter (*SHIFTOFL*). The 'warm start/stop' procedure is followed, whereby detector modules are turned fully on after the publication of the LHC stable beam flag, and return to standby at the end of stable beams. The percentage of modules ready, their number and the total in configuration are shown. Note the vertical time axis is compressed.

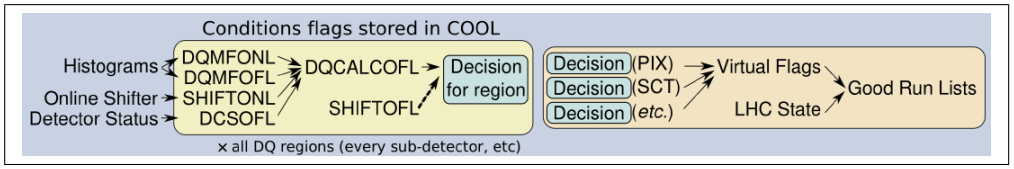

Figure 6. The propagation of flags from the various data quality systems through to analysis.

(DCS) [4], which monitors the detector hardware status. Figure 5 shows an example of the flagged decisions for the Pixel subsystem throughout one run. For each of these data quality systems, the top-level flag for each data quality region is stored in the ATLAS conditions database [1], which uses LCG COOL technology [5]. These flags are combined (according to subsystem-specific logic) into a proposed decision flag for each data quality region by the 'DQ Calculator', which can be overridden by an expert offline shifter (Fig. 6). Currently all of these 'primary' flags are checked by shifters, but in the future, as the automated systems become more trusted, a signoff record will be necessary for the flags.

For analysis users, data quality decisions are most relevant when they relate to analysis-level physics objects. Each physics object is likely to require good quality data from a number of data quality regions. For example, electrons in the barrel of the detector will require the data quality to be good for magnets, tracking subsystems, tracker alignment and the barrel parts of both electromagnetic and hadronic calorimeter systems. This is achieved by defining 'Virtual Flags', combinations of primary flags, each tailored to physics objects for analysis. Consequently, each physics analysis team can require a set of Virtual Flags, corresponding to the physics objects it uses, to be of good (i.e. 'green') data quality. These green virtual flags, supplemented by those for luminosity, trigger systems and LHC information can be combined to make a 'good runs list' of runs and luminosity block ranges suitable to that analysis. These are generated by central tools, in an ATLAS standard XML format for input and export.

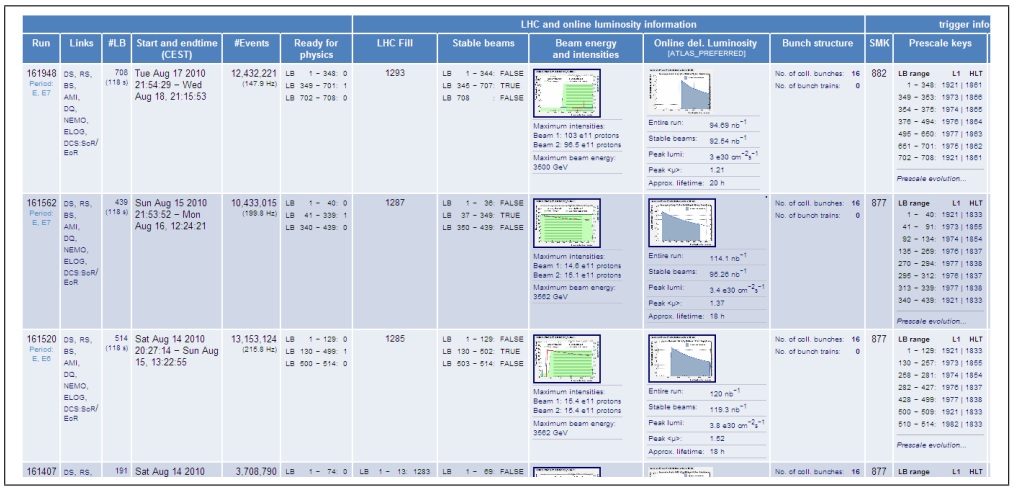

Figure 7. Run query webpage output from a query of 7 TeV collision runs taken between the 14th and 18th of August 2010.

5. Querying Runs

It is important that both analysis users and detector experts are able to view, search for and filter runs for a variety of LHC, data quality and run information. This is achieved through the ATLAS run query webpage, a highly flexible tool for investigating runs and their properties. The results of a sample search are shown in Figure 7. The run query webpage displays, and allows users to filter runs from, information including the number of events, LHC status (including stable beam flag, fill number, beam energy, bunch configuration and intensity), luminosity (displayed at run and LB level, with profile), trigger menu and keys, individual trigger rates, output streams and rates, and data quality of both primary and virtual flags.

6. Summary

Over the last year, the ATLAS data quality monitoring infrastructure has been commissioned rapidly and performed well in 7 TeV collision runs. The ATLAS data quality system has provided monitoring of Tier-0 reconstruction, producing status flags and algorithm results for a large number of histograms, made speedily accessible to users, primarily via an easily navigable webdisplay. Quick monitoring feedback and calibration loop procedures have enabled timely data quality decisions and accurate reconstruction parameters, producing high quality physics data for analysis teams. Data quality decisions are propagated to analysis users in a convenient way through 'virtual flags' tailored to physics analysis objects, providing each analysis with a 'good run list' of suitable runs and luminosity blocks. Tools exist to search, filter and view runs and their properties.

References

[1] Duckeck, G. *et al.* [ATLAS Collaboration], 2005. ATLAS Computing: Technical Design Report. *Preprint* CERN-LHCC-2005-022.
[2] Corso-Radu A., Hadavand H., Hauschild M., Kehoe R. and Kolos S. 2008, *IEEE Trans. Nucl. Sci.* **55** 417
[3] ATLAS Collaboration. AtLAS offline data quality monitoring, *Journal of Physics: Conference Series*, **219**(4):042018,2010.
[4] ATLAS Collaboration. Data quality from the Detector Control System at the ATLAS experiment. *Journal of Physics: Conference Series*, **219**(2):022037, 2010.
[5] Valassi A *et al.*, http://lcgapp.cern.ch/project/CondDB/