20th International Conference on Computing in High Energy and Nuclear Physics (CHEP2013) IOP Publishing
Journal of Physics: Conference Series **513** (2014) 012024 doi:10.1088/1742-6596/513/1/012024

Performance and development plans for the Inner Detector trigger algorithms at ATLAS

Stewart Martin-Haugh on behalf of the ATLAS Collaboration

Science and Technologies Facilities Council,
Rutherford Appleton Laboratory,
Harwell Oxford,
Didcot,
OX11 0QX

E-mail: stewart.martin-haugh@cern.ch

Abstract. A description of the algorithms and the performance of the ATLAS Inner Detector trigger for LHC Run 1 are presented, as well as prospects for a redesign of the tracking algorithms in Run 2. The Inner Detector trigger algorithms are vital for many trigger signatures at ATLAS. The performance of the algorithms for electrons is presented. The ATLAS trigger software will be restructured from two software levels into a single stage which poses a big challenge for the trigger algorithms in terms of execution time and maintaining the physics performance. Expected future improvements in the timing and efficiencies of the Inner Detector triggers are discussed, utilising the planned merging of the current two stages of the ATLAS trigger.

1. The ATLAS Inner Detector and Trigger during LHC Run 1 (2010-2012)

The ATLAS detector is one of two general-purpose detectors at the LHC [1] and is described in more detail in Reference [2] and references therein. It principally consists of an inner tracking detector, the Inner Detector (ID), electromagnetic and hadronic calorimeters and a muon spectrometer, as well as solenoidal and toroidal magnets. The ID plays a key role in the identification and measurement of objects, including electrons, muons, tau leptons and heavy flavour jets. The ID consists of three subdetectors: two silicon detectors (the Pixel and Semiconductor Tracker (SCT) detectors) and the Transition Radiation Tracker (TRT). The Pixel detector consists of three concentric layers of silicon pixel sensors, arranged radially (axially) in the barrel (endcap). The SCT barrel consists of four concentric layers of silicon microstrips, while the endcap consists of nine layers of silicon microstrips. The Pixel and SCT provide tracking over the range $|\eta| < 2.5$. The Transition Radiation Tracker (TRT) is a cylindrical detector extending to $|\eta| < 2.0$ consisting of 320,000 straw tubes filled with a Xe CO_2 O_2 gas mixture, around a central tungsten wire. A typical TRT track will have approximately 36 hits, which allows for improved estimation of track parameters when combined with the Pixel and SCT hits.

During the 2010-2012 running period (henceforth referred to as Run 1), the ATLAS detector recorded more than 5 fb^{-1} of proton-proton collision data at $\sqrt{s} = 7$ TeV 2010-2011 and more than 21 fb^{-1} of proton-proton collision data at $\sqrt{s} = 8$ TeV in 2012. These data were recorded with many additional interactions per bunch crossing. The mean number of interactions per bunch crossing, ($\langle\mu\rangle$) was 9.1 in 2011 and 20.7 in 2012. The presence of additional interactions presents a challenging environment – particularly for the ID trigger system, which must quickly decide which events are interesting for future analysis and which should be discarded.

1.1. The ATLAS Trigger System in Run 1

The ATLAS Trigger system as used in Run 1 is comprised of three levels, each progressively reducing the number of events considered. The Run 1 LHC bunch crossing frequency is 20 MHz. The first trigger stage, Level 1 (L1), uses specialised hardware systems in the calorimeter and muon detectors to accept events at a peak rate of 20–70 kHz between 2010 and 2012, with a decision time per event of < 2.5 µs. The L1 trigger identifies "Region of Interest" (RoI), regions in the detector potentially containing a physics object (calorimeter signals from electron or jet candidates, or hits in the muon detector from muon candidates). Detector data from the RoI(s) in the accepted event are then sent to Level 2 (L2), the second trigger level, which uses commodity PCs and software algorithms. The L2 output rate was between 3.5 and 6.5 kHz between 2010 and 2012. Events passing L2 are then sent to the Event Filter (EF), which has access to all sub-detectors and runs modified versions of the offline reconstruction algorithms. The EF output rate was 350 to 1000 Hz with an average event processing time of 0.3–1 s between 2010 and 2012. Collectively, the L2 and EF are referred to as the High Level Trigger (HLT).

1.2. The ATLAS Inner Detector Trigger System in Run 1

The EF trigger algorithms are identical to those used for offline reconstruction, but reconfigured for faster event processing. The L2 algorithms were developed specially for use in the trigger: more details may be found in Reference [3]. The final 2012 configuration used three different algorithms:

- Strategy A starts from identification of a primary vertex candidate, then finds tracks consistent with this primary vertex.
- Strategy B uses triplets of spacepoints corresponding to possible track roads as a seed for fast track finding.
- Strategy C is an implementation of offline or EF-style tracking.

Successful tracks passing one of these strategies can be extended into the TRT and then refitted to improve track parameter estimation.

2. Measurement of Electron Trigger Tracking Efficiency in 2012 Data

The performance of the electron tracking in the ID trigger was measured using $Z \to ee$ candidates in 2.5 fb^{-1} of 2012 data. A "tag and probe" method was used in order to avoid bias. The event selection trigger imposed tight tracking and calorimetry cuts on the tag electron and looser calorimetry cuts on the probe electron. No tracking selection cuts were applied to the probe electron in the trigger to avoid bias. The invariant mass between the tag and probe electrons was constrained to 70 GeV $< m_{e^+e^-} <$ 110 GeV in order to reduce contamination from fake electron candidates. Good quality offline electron tracks were then matched to the probe electron calorimeter cluster found in the trigger, and the efficiency measured by finding how often the trigger identified a track within $\Delta R = \sqrt{\Delta\eta^2 + \Delta\phi^2} < 0.03$ of this good electron candidate. The tag electron was also matched to an offline electron track in order to further increase the purity of the sample selected. The tracking efficiency was then defined as the fraction of offline probe electron tracks with a matching track identified in the trigger. Figure 1 shows the tracking efficiency for this measurement. The efficiency is shown separately for the L2 and EF trigger tracking (red and black respectively). Figure 1(a) shows the measured tracking efficiency as a function of p_T: the electron tracking is $> 99\%$ efficiency across the entire p_T range, without any decrease in performance at higher p_T. Figure 1(b) shows the efficiency as a function of $\langle\mu\rangle$: again, there is no degradation of performance at the highest values of $\langle\mu\rangle$ encountered.

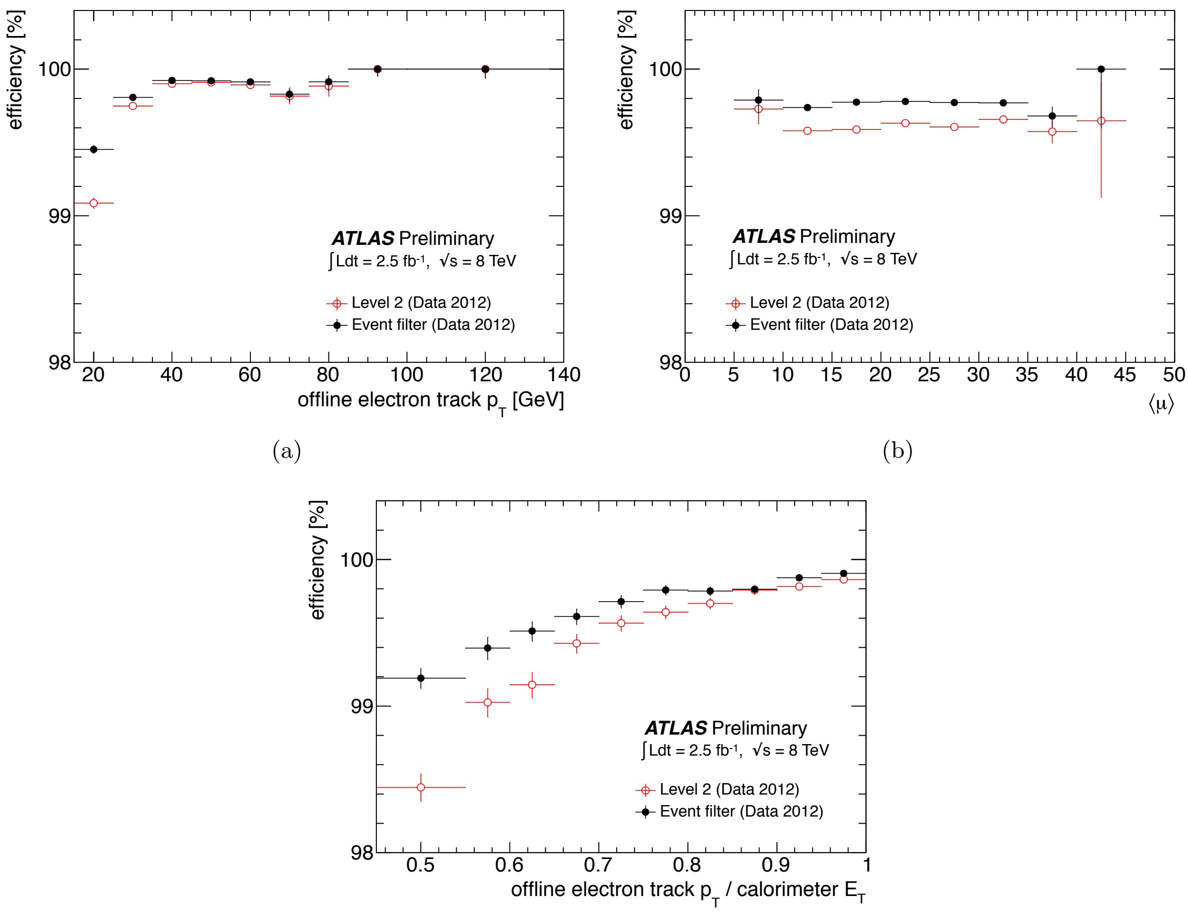

(a) (b) (c)

Figure 1. The electron trigger tracking efficiency as a function of p_T (a), $\langle \mu \rangle$ (b) and the ratio of track p_T to calorimeter E_T (c). From Reference [4].

Figure 1(c) shows the efficiency as a function of the ratio of track p_T to calorimeter E_T. Since the calorimeter is sensitive to radiation from electrons and photons, this ratio is a measure of the amount of bremsstrahlung an electron has undergone. Bremsstrahlung represents a challenge for tracking since it leads to changes in track curvature which must be incorporated into the track fitting model. Even for electrons losing 50% of their momentum in bremsstrahlung, the trigger tracking is over 98% efficient at L2, and over 99% efficient at the EF.

3. The ATLAS Inner Detector Trigger Upgrade

At the time of writing, the ATLAS Detector is being upgraded to meet the challenge presented by the LHC Run 2 data taking conditions. The programme of upgrade activities (referred to collectively as the Phase-I Upgrade) is documented fully in an upcoming Technical Design Report [6]. The two stage HLT used in Run 1 (comprised of L2 and the EF) will be merged to form a single trigger stage. This has the immediate advantage of reducing data preparation requirements: in Run 1, data were prepared separately for the L2 and EF systems. Additionally, there is significant scope to redesign the structure and inter-operation of the different trigger algorithms. The target event processing time in the merged HLT is likely to be approximately 200 ms, to compare to the 75 ms (L2) and 1 s (EF) timings in Run 1. The redesigned ID trigger will use fast track reconstruction, similar to that in the current L2, and a more detailed track reconstruction as in the current EF but seeded from the individual track candidates identified by

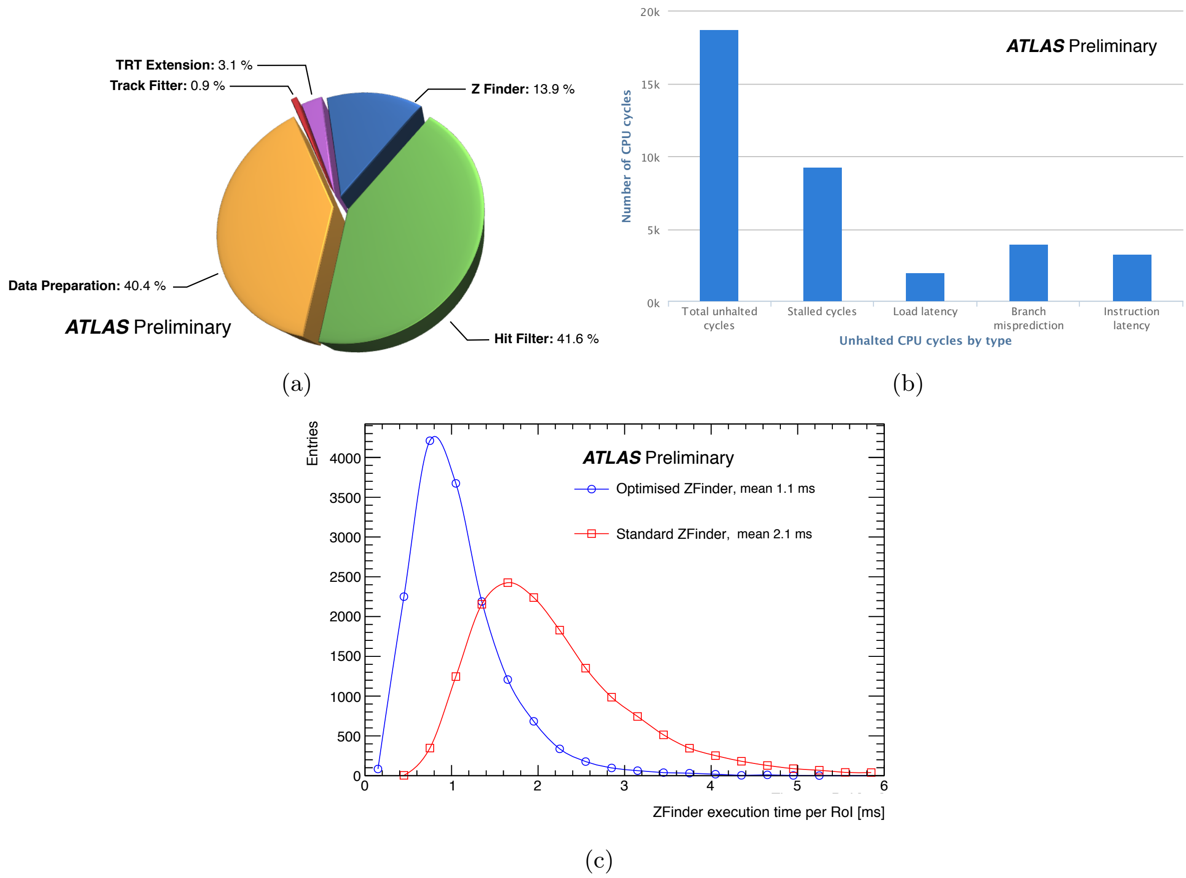

Figure 2. Profiling information for an L2 tracking algorithm. In (a), the proportion of total execution time taken by various modules in an L2 tracking algorithm. In (b), the total number of CPU cycles spent in the `Z Finder` module, either processing information or idle. In (c), the effect of optimisation on `Z Finder` timing. From Reference [5].

the fast track finding stage. In addition, a fast track trigger using custom hardware (known as the Fast TracKer or FTK) will be installed in 2016 [7] and will provide tracks at the L1 output rate. Several strategies for use of FTK tracks in the HLT are currently under investigation, including simply refining the FTK track parameter measurements and using FTK tracks to seed full offline-style combinatorial tracking. In addition to the upgrades to the trigger and data acquisition systems, a new layer of the Pixel detector, the Insertable B-Layer (IBL) [8] is being installed. This will be situated 25.7 mm from the beamline, while the current innermost layer is situated 50.5 mm from the beamline. The IBL is expected to improve the tracking and vertexing resolutions, b-tagging performance and increase the robustness of track reconstruction against fake tracks. In particular, an online impact parameter resolution of approximately 10 μm should be possible. An MC study of resolution in ID trigger tracks demonstrating improvement due to the IBL is presented in Reference [5].

3.1. Profiling and optimisation of ID Trigger algorithms

The ID tracking algorithms used in Run 1 will form the building blocks for the Run 2 trigger tracking algorithms, so significant effort is being invested in identifying where improvements in speed and memory usage in the Run 1 algorithms would be most beneficial. Such efforts are focused on areas of code which contribute significantly to the overall running time of the trigger, as identified by profiling tools. The tools used for this work are the Generic Optimization Data Analyzer (`GOoDA`) [9] `perf` [10] and `Callgrind` [11]. Figure 2(a) shows the proportion of total execution time by various modules from Strategy A - the first of the L2 tracking algorithms described in Section 1.1. It is clearly most beneficial to focus optimisation effort on three modules (`Z Finder`, `Hit Filter` and `Data Preparation`) which together account for 95% of the total execution time. The `perf` tool provides very detailed profiling information: Figure 2(b) shows the total number of CPU cycles spent either processing information (total unhalted cycles) or idle for a variety of other reasons (stalled cycles). Modern CPUs preload upcoming instructions while processing current instructions. If there is a branch in program structure (i.e. a point where there is a choice of two or more instructions), the branch predictor will try to guess which statement is likely to be executed next so that it may be preloaded. Branch mis-prediction can incur a significant performance penalty: in this case it accounts for 21% of all stalled cycles. Figure 2(c) shows a factor of two (2.1 ms to 1.1 ms) improvement in `Z Finder` timing after re-ordering loops to avoid branch misprediction. The potential for similar gains elsewhere is currently under investigation.

4. Conclusions

Results from performance studies with Run 1 data have been presented, showing robustness of trigger tracking performance in high pile-up environments, and high efficiency in the presence of significant bremsstrahlung for high-p_T electron tracks. Progress in ID trigger software development has also been discussed, building on Run 1 successes and preparing for Run 2. The ongoing development for the tracking trigger will play a significant role in enabling the ATLAS trigger to meet the challenges of Run 2 and future data taking.

Bibliography

[1] Evans L and Bryant P (editors) 2008 *JINST* **3** S08001
[2] ATLAS Collaboration 2008 *JINST* **3** S08003
[3] Bernat P 2012 Architecture and Performance of the Inner Detector Trigger of the ATLAS detector Tech. Rep. ATL-DAQ-PROC-2012-032 URL `https://cds.cern.ch/record/1458060`
[4] ATLAS Collaboration 2013 High Level Trigger Tracking Public Results URL `https://twiki.cern.ch/twiki/bin/view/AtlasPublic/HLTTrackingPublicResults`
[5] ATLAS Collaboration 2013 Studies for the development of the Inner Detector trigger algorithms at ATLAS Tech. Rep. ATL-DAQ-PUB-2013-002 URL `http://cds.cern.ch/record/1602918`
[6] ATLAS Collaboration 2013 Technical Design Report for the Phase-I Upgrade of the ATLAS TDAQ System Tech. Rep. CERN-LHCC-2013-018. ATLAS-TDR-023 URL `http://cds.cern.ch/record/1602235`
[7] ATLAS Collaboration 2013 Fast TracKer (FTK) Technical Design Report Tech. Rep. CERN-LHCC-2013-007. ATLAS-TDR-021 URL `https://cds.cern.ch/record/1552953`
[8] ATLAS Collaboration 2010 ATLAS Insertable B-Layer Technical Design Report Tech. Rep. CERN-LHCC-2010-013. ATLAS-TDR-19 URL `https://cds.cern.ch/record/1291633`
[9] Calafiura P, Eranian S, Levinthal D, Kama S and Vitillo R 2012 *Journal of Physics: Conference Series* **396** 052072 URL `http://stacks.iop.org/1742-6596/396/i=5/a=052072`
[10] Vitillo R A et al 2013 Linux profiling with performance counters URL `https://perf.wiki.kernel.org/index.php/Main_Page`
[11] Seward J et al 2004 URL `http://www.valgrind.org`