PROCEEDINGS OF SCIENCE

# The ATLAS Trigger in Run 2: Design, Menu, and Performance

**Tamara Vazquez Schroeder, on behalf of the ATLAS Collaboration**
*McGill University*
*E-mail:* tamara.vazquez.schroeder@cern.ch

The ATLAS trigger system is composed of a hardware Level-1 trigger and a software-based high-level trigger. It was successfully operated during the first part of Run 2 (2015/2016) at the LHC at a centre-of-mass energy of 13 TeV. A comprehensive review of the ATLAS trigger design, menu, and performance in Run 2 is presented in this proceedings contribution, as well as an overview of the intensive preparation towards the second part of Run 2 (2017/2018).

*EPS-HEP 2017, European Physical Society conference on High Energy Physics*
*5-12 July 2017*
*Venice, Italy*

PoS(EPS-HEP2017)525

## 1. Introduction

The trigger system in the ATLAS detector [1] decides online whether or not to keep and record an event. Its successful operation has a crucial impact on the quality of the dataset used in physics analyses. During Run 1 of Large Hadron Collider (LHC) operations, the ATLAS trigger system operated efficiently at instantaneous luminosities up to $0.8 \times 10^{34}$ cm$^{-2}$s$^{-1}$ at centre-of-mass energies up to 8 TeV and collected more than three billion events. The trigger system was substantially upgraded in preparation for the increased collision energy, higher luminosity, and increased number of proton-proton interactions per bunch crossing (pile-up) expected in Run 2. Thanks to these improvements, the ATLAS trigger system has operated successfully during the first part of Run 2 collision data-taking (2015 and 2016), and is preparing intensively for the challenges of the second part of Run 2 (2017 and 2018).

An introduction to the ATLAS trigger and data acquisition system (TDAQ) is provided in Section 2. The ATLAS trigger menu strategy is explained in Section 3. The trigger rate predictions and HLT farm performance studies are discussed in Section 4. Section 5 summarises the validation cycle of the ATLAS trigger software. The online monitoring performance of the trigger is given in Section 6. Finally, an overview of the latest trigger signature performance results is provided in Section 7.

## 2. The ATLAS trigger and DAQ system

In Run 2, the TDAQ system consists of a hardware-based first level trigger (L1) and a software-based high-level trigger (HLT), reducing the 40 MHz collision input rate provided by the LHC to a rate of 1 kHz of events to be recorded [2]. The L1 trigger systems are implemented in hardware and use a subset of the detector information to reduce the rate of accepted events to 100 kHz with a fixed latency of 2.5 $\mu s$. Fast custom-made electronics find regions of interest (RoIs) using calorimeter and muon data with coarse information. The L1Calo subsystem uses a sliding-window algorithm [3] to find local transverse energy maxima up to $|\eta| < 4.9$ within two grids of trigger towers, each tower in the barrel covering $0.1 \times 0.1$ in $\eta \times \phi$. One grid comes from the electromagnetic calorimeters and one from the hadronic calorimeters. For Run 2, a new Multi-Chip Module was included which allows more flexible signal processing. L1Muon operates with fast Resistive Plate Chambers (RPC) in the barrel ($|\eta| < 1.05$) and Thin Gap Chambers (TGC) in the end-caps ($1.05 < |\eta| < 2.4$), and locates the coincidence between hits in different layers of the muon spectrometer. Since Run 2, coincidences with the inner detector have also been incorporated into the trigger logic. The performance is also improved with additional chambers in the feet of the barrel region and from the Tile calorimeter extended barrel region.

Both L1Calo and L1Muon output L1 trigger objects which encode the type, location, energy, and isolation status of identified objects. These are provided as inputs to the L1Topo subsystem - a new subsystem implemented in Run 2 - which performs geometric and kinematic selections on them in order to keep the L1 thresholds and dedicated trigger rates low. The final L1 decision is formed by the L1 Central Trigger Processor (CTP) using the L1 trigger objects and L1Topo output. In Run 2, the CTP has been operating with upgraded hardware to increase the number of triggers which can be processed in parallel.

After being accepted by the L1 trigger, events are buffered in the Read-Out System (ROS) and processed by the HLT. The HLT receives RoI information from L1, which can be used for regional reconstruction in the trigger algorithms. The HLT algorithms are executed on approximately 40000 CPU cores and reduce the rate of recorded full events to 1 kHz. Additionally, partial event building is used for trigger level analysis, detector monitoring, and calibrations of the ATLAS detector subsystems. In Run 2, the HLT readout and data storage systems have been fully upgraded. Furthermore, a new Fast TracKer (FTK) system [4], currently under commissioning, will provide global ID track reconstruction at the L1 trigger rate using lookup tables stored in custom associative memory chips for the pattern recognition. This hardware accelerated tracking will allow the use of tracks at much higher event rates in the HLT than is currently affordable using CPU systems.

After the events are accepted by the HLT, they are written into data streams, transferred to local storage at the experimental site, and exported to the Tier-0 facility at the CERN computing centre for offline reconstruction.

## 3. The ATLAS trigger menu

The trigger menu comprises the list of full L1 to HLT trigger selections (trigger chains) with prescale factors [5]. It reflects the physics goals of the collaboration, with high acceptance for beyond-the-Standard Model searches, as well as for Higgs boson and Standard Model precision measurements. The available data taking resources (L1, HLT and Tier-0) are also taken into account in the design of the trigger menu. In general, the trigger menu strategy is based on the following building blocks:

- **primary triggers**: used for physics measurements and typically run unprescaled;
- **support triggers**: used for efficiency and performance measurements, background estimates or monitoring, and typically running with a small rate;
- **alternative triggers**: running alternative online reconstruction algorithms;
- **backup triggers**: using tighter selections and therefore running with a lower expected rate, in case the rate of the main (primary) trigger becomes higher than allowed.

The trigger menu is designed for a specific peak luminosity. In 2016, the LHC exceeded its design luminosity of $1.0 \times 10^{34}$ cm$^{-2}$s$^{-1}$, reaching an instantaneous luminosity of $1.4 \times 10^{34}$ cm$^{-2}$s$^{-1}$. In 2017, the baseline menu was designed for an instantaneous luminosity of $2.0 \times 10^{34}$ cm$^{-2}$s$^{-1}$. Primary triggers are generally kept stable within a menu during data-taking. Furthermore, the trigger menu should be flexible enough to adjust to changing conditions during LHC ramp-up. Currently, over 3000 trigger chains are run to select events of interest and covering a large spectrum of physics objects and processes.

The trigger menu is deployed online with different prescale sets depending on the luminosity: as luminosity decreases throughout the fill, the bandwidth usage is optimised by increasing the rate of supporting triggers. As a result of the trigger menu used in 2016, the average uncompressed event size was 1.6 MB for a mean number of simultaneous interactions per proton-proton bunch crossing averaged over all bunches circulating in the LHC ($\langle \mu \rangle$) of 24.9.

## 4. Trigger rates and CPU usage

Understanding trigger rate predictions and HLT farm performance is essential for all menu developments and validation of HLT algorithms [6]. A special dataset, the so called EnhancedBias (EB) data stream, is collected every time data-taking conditions change and is used to provide rate predictions. For the EB dataset, events are selected by the L1 trigger system with higher energies and object multiplicities, and the selection bias is corrected for with event weights.

There have been several efforts to optimise the time and CPU consumption of HLT algorithms. In particular, there have been significant improvements in the timing for the ID track-based triggers in the HLT.

## 5. Trigger software validation

The full trigger menu and HLT software run offline over the EB dataset for algorithm validation. This software validation is performed on a weekly basis if there are significant changes in the software and menu, and involves expertise in trigger menus, HLT releases, software validation, and trigger signatures. For the validation, high memory consumption jobs are run on the Grid. The turnaround for a full validation is between 24 and 48 hours and requires up to 4 GB of memory usage per job. Once the jobs are finished, useful outputs are produced with reprocessing performance metrics (reconstructed observable distributions compared to reference, expected algorithm rates, etc.), which are then analysed by trigger signature and menu experts. Based on their feedback, the changes in the software release and trigger menu are accepted or rejected [7].

As of 2017, the CPU usage of several trigger chains has improved, and the release building and distribution have been automated and are done every night without manual intervention. These improvements reduce the length and memory consumption of the validation jobs, as well as the turnaround of the validation cycle.

## 6. Trigger monitoring performance

Once the trigger software and trigger menu are deployed online, distributions of HLT-level quantities are monitored. Automatic data quality (DQ) checks are applied based on standardised histogram analyses and comparisons to reference histograms. The trigger shifters in the ATLAS control room are able to track the performance of the HLT via red (alarm), yellow (warning) and green (OK) DQ evaluation. A similar procedure is followed offline to declare data good for physics. In Run 2, a menu-aware monitoring scheme makes it possible to update the monitoring configuration out-of-sync with software releases with very small latency of the order of 1 hour.

## 7. Trigger signature performance

The improvements in the L1 and HLT systems are reflected in the performance of the trigger objects produced. Some examples of these improvements in the 2017 dataset are presented in this section.

Figure 1 (left) shows the efficiencies for HLT large-radius ($R$) single-jet triggers as a function of the leading offline trimmed jet $p_{\mathrm{T}}$ for jets with $|\eta| < 2.0$ and jet mass above 50 GeV [8]. The trimming procedure removes soft contamination from pile-up in large-*R* jets. Blue circles represent a trimmed large-*R* jet trigger with a $p_T$ threshold of 420 GeV. Adding an additional 30 GeV cut on the jet mass significantly suppresses the QCD dijet background, allowing a lower $p_T$ threshold of 390 GeV, while retaining nearly all signal-like jets with a mass of above 50 GeV. This is shown in green triangles.

Figure 1 (right) shows the efficiencies for an unprescaled (small-*R*) single-jet trigger with three different calibrations applied to jets in the HLT [8]. Offline jets are selected with $|\eta| < 2.8$. The calibration steps applied in 2016 data are represented in green (open squares); the updated calibration applied in 2017, utilising only calorimeter information, can be seen in red (closed circles); and in blue (open circles) the updated calibration with track information is shown. The extra calibration steps present in 2017 include global sequential corrections and the application of in situ corrections. The Global Sequential Calibration (GSC) corrects jets according to their longitudinal shower shape and associated track characteristics without changing the overall energy scale. Since tracking is not guaranteed to be available for all jet thresholds, options are provided with and without the track-based corrections. The data-driven $\eta$-intercalibration correction is the most important in situ correction added, and fixes differences in jet response as a function of $\eta$. Together, these additional corrections allow for improved agreement between the scale of trigger and offline jets as a function of both $\eta$ and $p_T$, and thus the trigger efficiency rises much more rapidly.

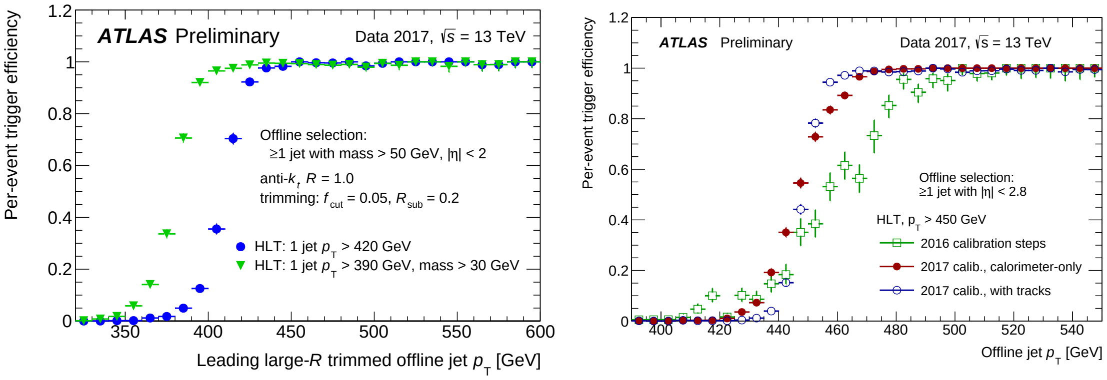

**Figure 1:** Left, efficiencies for HLT large-*R* single-jet triggers are shown as a function of the leading offline trimmed jet $p_T$ for jets with $|\eta| < 2.0$ and jet mass above 50 GeV. Two large-*R* jet triggers from the 2017 menu are shown [8]. Right, efficiencies for an unprescaled small-*R* single-jet trigger with three different calibrations applied to jets in the HLT [8].

Pile-up mitigation is the main challenge for missing transverse energy ($E_T^{miss}$) triggers. The "mht" algorithm, based on the sum of the $p_T$ of HLT jets, was the default algorithm in 2016. In 2017, the so called "pufit" algorithm is the new baseline, where pile-up is estimated event-by-event and subtracted [2]. Figure 2 (left) shows the trigger cross section as a function of $\langle\mu\rangle$, for the "mht" and "pufit" algorithms. The "pufit" algorithm reduces the trigger cross section significantly compared to "mht" for high pile-up [9].

The ATLAS *b*-jet trigger uses a boosted decision tree (BDT) algorithm to separate *b*-jets from light and *c*-jet backgrounds. The BDT algorithm was re-optimised in 2017 to improve the *b*-tagging performance [10]. Figure 2 (right) shows the performance of *b*-tagging algorithms, measured using

$t\bar{t}$ Monte Carlo events, in terms of $c$-jet rejection as a function of $b$-jet efficiency. The expected performance of the $b$-tagging algorithm for $b$-jet triggers in 2017 data-taking (green solid line) is compared to $b$-tagging algorithms used for $b$-jet triggers in 2016 (red solid line). The $c$-jet rejection of the $b$-tagging algorithm of $b$-jet triggers improved considerably in 2017 and is much closer to that of offline $b$-jets (purple dotted curve).

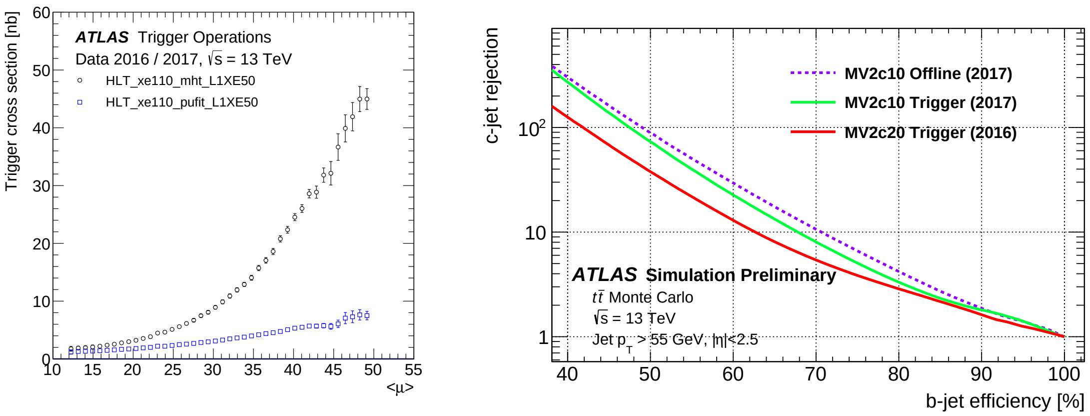

**Figure 2:** Left, trigger cross section for the main $E_{\mathrm{T}}^{\mathrm{miss}}$ trigger reconstruction algorithms used in 2016 ("mht") and 2017 ("pufit") as a function of $\langle\mu\rangle$ [9]. Right, performance of $b$-tagging algorithms in terms of $c$-jet rejection as a function of $b$-jet efficiency [10].

Electron, photon, and muon trigger efficiency performance has also been excellent so far in 2017, showing a sharp turn-on curve as a function of the energy or $p_{\mathrm{T}}$ of the triggered object.

## 8. Conclusion

The trigger hardware and software have been modified and improved to cope with the challenges expected during LHC Run 2. The trigger was successfully commissioned in 2015 and it has smoothly operated during 2016 despite the very challenging LHC conditions. Impressive improvements were made in preparation for the expected highest ever luminosities and pile-up in the 2017/18 LHC run, and are already reflected in the early 2017 trigger performance results. Further improvements, such as the full integration of the FTK in the ATLAS trigger system, are expected in 2018.

## References

[1] ATLAS Collaboration, JINST 3 (2008) S08003.

[2] ATLAS Collaboration, Eur. Phys. J. C 77 (2017) 317.

[3] R. Achenbach et al., JINST 3 (2008) P03001.

[4] ATLAS Collaboration, ATLAS-TDR-021 (2013).

[5] ATLAS Collaboration, ATL-DAQ-PUB-2017-001 (2017).

[6] ATLAS Collaboration, ATL-DAQ-PUB-2016-002 (2016).

[7] Robert Keyes on behalf of the ATLAS Collaboration, ATL-DAQ-PROC-2016-040 (2016).

[8] ATLAS Collaboration, ATL-COM-DAQ-2017-063 (2017).

[9] ATLAS Collaboration,
https://twiki.cern.ch/twiki/bin/view/AtlasPublic/MissingEtTriggerPublicResults.

[10] ATLAS Collaboration, ATL-COM-DAQ-2017-062 (2017).