

The ATLAS Run-2 Trigger Menu for Higher Luminosities: Design, Performance and Operational Aspects

Arantxa Ruiz Martínez^{1,a} on behalf of the ATLAS Collaboration

¹*Department of Physics, Carleton University, Ottawa, ON, Canada*

Abstract. The ATLAS experiment aims at recording about 1 kHz of physics collisions, starting with an LHC design bunch crossing rate of 40 MHz. To reduce the massive background rate while maintaining a high selection efficiency for rare physics events (such as beyond the Standard Model physics), a two-level trigger system is used. Events are selected based on physics signatures such as the presence of energetic leptons, photons, jets or large missing energy. The trigger system exploits geometrical information on candidate objects, as well as multi-variate methods to carry out the necessary physics filtering. In total, the ATLAS online selection consists of thousands of different individual triggers. A trigger menu is a compilation of these triggers which specifies the physics algorithms to be used during data taking and the bandwidth a given trigger is allocated to. Trigger menus reflect not only the physics goals of the collaboration for a given run, but also take into consideration the instantaneous luminosity of the LHC and limitations from the ATLAS detector readout and offline processing farm. For the 2017 run, the ATLAS trigger has been enhanced to be able to handle higher instantaneous luminosities (up to $2.0 \times 10^{34} \text{ cm}^{-2}\text{s}^{-1}$) and to ensure the selection robustness against higher average multiple interactions per bunch crossing.

In these proceedings, we describe the design criteria for the trigger menus used for Run 2 at the LHC. We discuss several aspects of the process, from the validation of the algorithms, the fine-tuning of the prescales, and the monitoring tools that ensure the smooth operation of the trigger during data taking. We also report on the physics performance of a few trigger algorithms.

1 Introduction

The trigger system is a fundamental component of the ATLAS [1] experiment at the LHC [2], since it is responsible for the online event selection for data storage. In the LHC Run-1 (2009-2013), the ATLAS trigger system [3] operated efficiently at instantaneous luminosities of up to $8 \times 10^{33} \text{ cm}^{-2}\text{s}^{-1}$ and at centre-of-mass energies of up to 8 TeV, collecting more than three billion events. During the LHC shutdown (2013-2015), the trigger system underwent major upgrades to cope with the extremely challenging conditions expected for the upcoming years. In the LHC Run-2 (2015-2018), the increased centre-of-mass energy to 13 TeV and higher instantaneous luminosities, surpassing the initial design luminosity of $10^{34} \text{ cm}^{-2}\text{s}^{-1}$, lead to a significant increase of rates, exceeding the capabilities of the original (Run-1) trigger system.

^ae-mail: aranzazu.ruiz.martinez@cern.ch

2 The ATLAS Run-2 Trigger and Data Acquisition

The ATLAS experiment is a multi-purpose particle detector with a forward-backward symmetric, cylindrical geometry and nearly 4π coverage in solid angle. The interaction point is surrounded by an inner detector, a calorimeter system, and a muon spectrometer.

The ATLAS detector in Run-2 uses a two-level trigger system: a hardware-based first level trigger (Level-1 or L1) [4] and a software-based high level trigger (HLT) [5]. A schematic view of the different components of the upgraded ATLAS trigger and data acquisition system is shown in Figure 1.

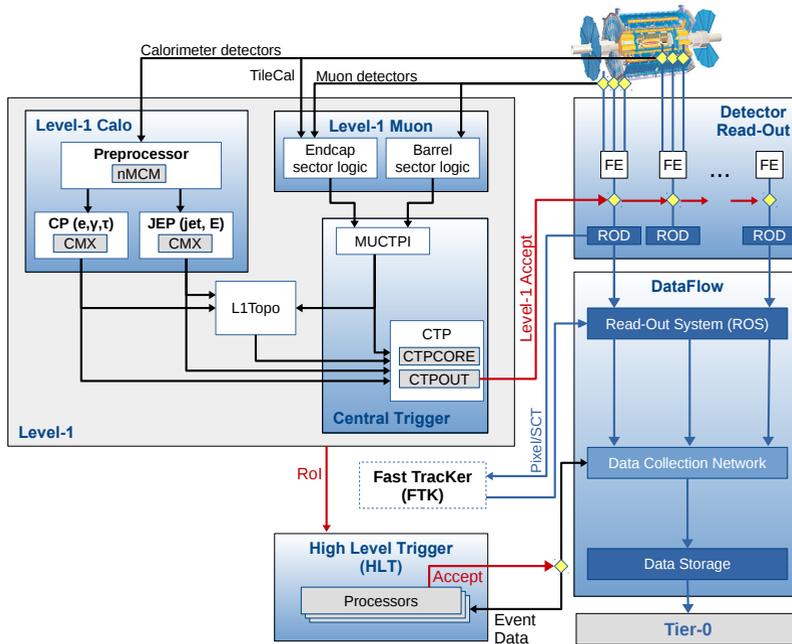


Figure 1. Schematic illustration of the ATLAS Run-2 Trigger and Data Acquisition system in 2017.

The Level-1 trigger uses custom electronics to determine Regions-of-Interest (RoIs) in the detector, taking as input coarse granularity calorimeter and muon detector information. The Level-1 Calorimeter trigger makes use of reduced granularity information from the electromagnetic and hadronic calorimeters to search for electrons, photons, hadronic taus and jets, as well as missing transverse energy (E_T^{miss}). The Level-1 Muon trigger system, which consists of a barrel section and two endcap sections, provides fast trigger signals from the muon detectors for the Level-1 trigger decision. The Level-1 Topological trigger (L1Topo) system allows to apply (so-called) topological event selections at the Level-1 trigger combining kinematic information from multiple calorimeter and muon trigger objects, such as angular separation, invariant mass requirements, or global event quantities such as the sum of the transverse momenta of all Level-1 jet objects. The Central Trigger Processor (CTP) forms the Level-1 trigger decision receiving the inputs from the Level-1 Calorimeter and Level-1 Muon triggers as well as several other subsystems such as the Minimum Bias Trigger Scintillators, the LUCID Cherenkov counter and the Zero-Degree Calorimeter. The Level-1 trigger reduces the event rate from the LHC bunch crossing rate of approximately 30 MHz to 100 kHz. The decision time for a Level-1 accept is $2.5 \mu\text{s}$.

At HLT, sophisticated offline-like reconstruction algorithms are executed in either the RoI defined at Level-1 or the whole event using full granularity detector information, being the data retrieved on demand from the readout system. Most HLT triggers use a two-stage approach with a fast reconstruction step to reject the majority of events and a slower precision reconstruction for the remaining events. However, due to the unification of the Level-2 and Event Filter computer clusters in Run-2, there is no fixed bandwidth or rate limitation between the two steps. The HLT reduces the rate from the Level-1 output rate of 100 kHz to approximately 1 kHz on average for physics events with full event-data within a processing time of about 200 ms, making use of a computer farm of up to approximately 40,000 CPU cores.

Events accepted by the HLT are transferred to the Tier-0 computing centre for offline event reconstruction to be used later on in physics analyses.

2.1 Trigger menu, rates and bandwidth

The event selections implemented at the HLT are referred to as chains, as they chain together a sequence of different algorithms. The complete set of all trigger chains is called the trigger menu. Each chain specifies a set of Level-1 trigger seeds which if present will activate the chain. The rates are adjusted via prescale sets, where the prescale denotes the probability of a chain being activating when seeded. Chains are comprised of selection steps, where a step is a sequence of algorithms. The Level-1 trigger menu consists of 512 single items and combinations, while the HLT menu is composed of $O(1000)$ chains.

The trigger menu for 2017 data taking has been designed for a peak luminosity of $2 \times 10^{34} \text{ cm}^{-2} \text{ s}^{-1}$. For early data taking and lower peak luminosities, looser selections were deployed in some cases. The triggers menu strategy is based on the following trigger types:

- primary triggers, used for physics analyses, typically running unprescaled
- support triggers, used for efficiency and performance measurements or monitoring
- alternative triggers, used to study alternative online reconstruction algorithms
- backup triggers, used for emergency cases with tighter selections and lower rates
- calibration triggers, used for detector calibrations

The trigger menu has to respect all trigger constraints at Level-1 and HLT and make the optimal distribution of the available bandwidth and CPU resources. The distribution is driven by the physics priorities of the experiment, and the number of clients satisfied by a particular trigger, where very specialized triggers can aim at $O(1 \text{ Hz})$, while triggers serving several analyses can aim at $O(10 \text{ Hz})$ or even more as in the case of the single electron and muon primary triggers. The HLT bandwidth assigned to the support triggers is about 20%.

Recorded events are gathered in data streams, depending on their primary use case and their specific offline reconstruction needs. The HLT streams are classified in different categories: the main physics stream which contains all triggers for physics analyses, the *B*-physics and light states (LS) stream which contains triggers specific to *B*-physics analyses, the express stream which is used for fast data quality monitoring and run sign-off, and other minor streams with physics applications, such as zero-bias events, the trigger level analysis stream, the detector calibration streams, and the detector monitoring streams. The main physics stream constitutes around 75% of the total HLT output bandwidth.

Chains require either full event building (EB) or partial EB with only subdetector information for recording into data streams, such as the chains under the trigger level analysis stream, which is a novel idea to circumvent the bandwidth limitation using partial event building. HLT low transverse

momentum (p_T) jet triggers are highly prescaled in the main physics stream, while higher rates can be afforded and therefore a large gain in statistics can be achieved by recording a small fraction of the data ($< 5\%$ of the standard event size) for jets with $p_T < 400$ GeV. This is important to increase the sensitivity reach of the low-mass dijet searches.

Figure 2 shows the output rates of some representative unprescaled Level-1 triggers as a function of time for a run taken in 2017. The rates have an exponential decay with decreasing luminosity during an LHC fill. The rates periodically increase due to LHC luminosity re-optimisations, while dips are due to deadtime and spikes are caused by detector noise. The Level-1 single object trigger requiring an electromagnetic cluster with transverse energy $E_T > 22$ GeV and hadronic and electromagnetic isolation requirements (EM22VHI) is the item which gives the highest rate in the Level-1 menu.

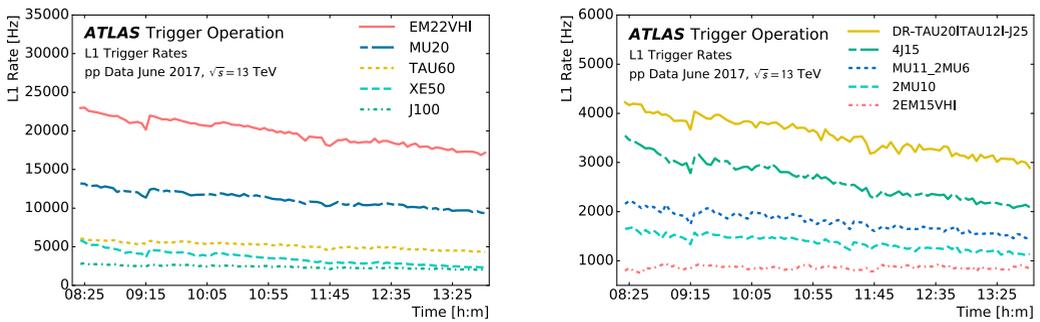


Figure 2. Level-1 trigger rates as a function of time in a fill taken in June 2017 [6]. The rates of some representative unprescaled Level-1 single (left) and multi (right) object triggers are shown. These trigger items are based on electromagnetic clusters (EM), muon candidates (MU), tau candidates (TAU), missing energy (XE) and jet candidates (J). The trigger name indicates the trigger threshold in GeV and the object multiplicity as well as other details of the selection such as: varying thresholds as a function of pseudorapidity (V), hadronic isolation (H), isolation (I), and separation in η and ϕ (DR).

The L1Topo was activated and commissioned in 2016 and is used by default in several primary triggers in 2017. Figure 3 shows the rates of a B -physics trigger requiring two muons with and without applying a topological selection at Level-1 as a function of the number of luminosity blocks for a run taken in 2017. This L1Topo item requires dimuon invariant mass of $2 < m_{\mu\mu} < 9$ GeV and angular separation of $0.2 < \Delta R_{\mu\mu} < 1.5$, and leads to an overall rate reduction of approximately a factor of four.

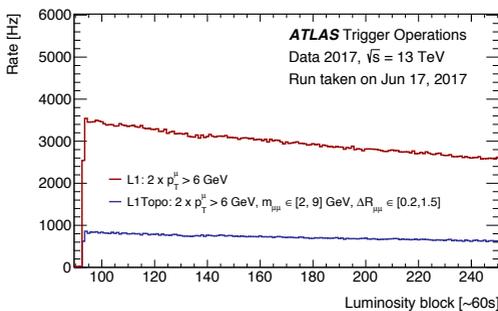


Figure 3. Rates of the Level-1 trigger item requiring two muons with $p_T > 6$ GeV with and without L1Topo requirements as a function of the number of luminosity blocks, which on average correspond to 60 s, in a run taken in June 2017 [6]. The topological selection applied in the L1Topo item requires dimuon invariant mass of $2 < m_{\mu\mu} < 9$ GeV and angular separation of $0.2 < \Delta R_{\mu\mu} < 1.5$.

Figure 4 shows the HLT output rates of the main streams as a function of time for a run collected in 2017. The rates periodically increase due to change of prescales to optimise the bandwidth usage or LHC luminosity re-optimisations, while dips are due to deadtime and spikes are caused by detector noise. Higher HLT output rates are achieved by partial EB of some dedicated streams.

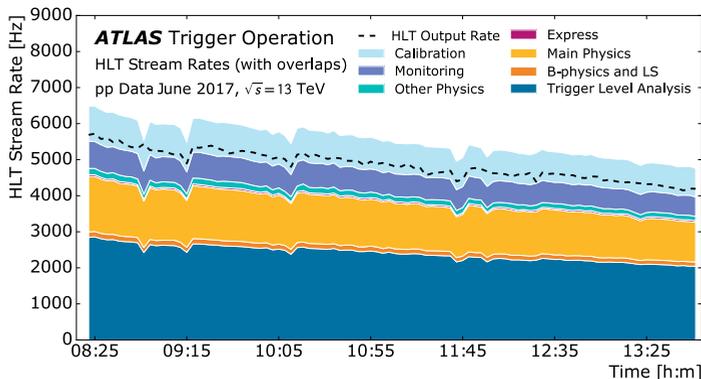


Figure 4. HLT physics stream rates as a function of time in a fill taken in June 2017 [6]. Higher HLT output rates are achieved by partial EB of some dedicated streams.

Figure 5 shows the HLT trigger rate of the main physics stream as a function of time for a run acquired in 2017, with the breakdown into the different signatures used to group the chains: electrons, muons, jets, *b*-jets, missing transverse energy, taus, photons, *B*-physics and their combinations. Overlaps are only accounted for in the total main physics stream rate. Due to decreasing luminosity during the run, the average rate is about two thirds of the peak rate. Events for physics analyses using full event building are recorded at an average rate of approximately 1 kHz.

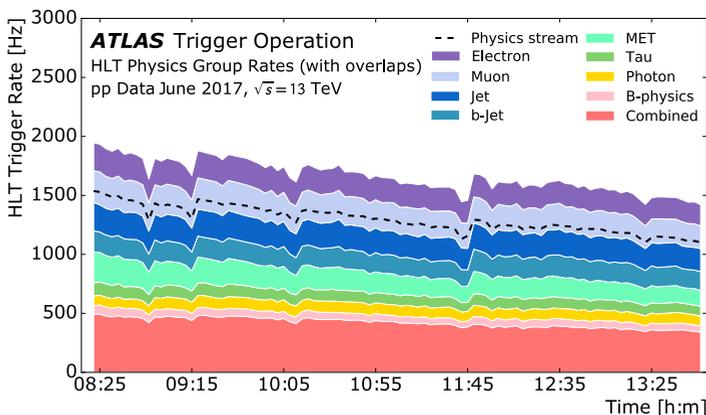


Figure 5. HLT physics trigger group rates as a function of time in a fill taken in June 2017 [6]. The rates of the individual trigger groups specific to trigger physics objects are shown. Each of the groups contain single and multi triggers of the same object. Overlaps are only accounted for in the total main physics stream rate.

2.2 Trigger software validation and operational aspects

Frequent updates are needed on the trigger software to adapt to the changing running conditions. Rock-solid software validation procedures are in place to guarantee a stable and reliable trigger operation. The trigger software validation cycles are carried out on a regular basis to include all the developments needed in new releases for online data taking. These validations can happen for special cases, such as to prepare for new pileup conditions, to move from proton-proton to heavy ion runs, or specifically to test improved algorithms, modification of parameters or new signature calibrations. Many groups and related experts are involved in these validation cycles.

This process typically begins compiling all the requests for software updates, bug fixes for problems found either in previous validation cycles or through feedback from online usage. New software tags to address these issues are provided by the different signature groups. After that, all the new tags are collected by the HLT software release expert who builds the release. Then the trigger menu expert uses this new release to build an updated menu. Subsequently the release expert deploys the new release to the grid in preparation for a large scale reprocessing of enhanced bias data (sample of events selected using the main Level-1 physics triggers without further HLT selection). The next step is therefore to launch the reprocessing of enhanced bias data to test the HLT software. First the raw enhanced bias data are processed emulating the HLT, followed by the reconstruction of the accepted data and by the calculation of trigger efficiencies and resolutions for validation. Once the data have been processed, the software validation expert makes available the output results to the signature experts and requests their feedback. If the software release is signed off then it can be deployed for data taking in ATLAS. Otherwise, new bug tickets are opened and the process re-starts. Overall this cycle typically takes approximately one week to be completed. Enhanced bias data trigger reprocessings are also used to obtain trigger rate predictions at a given instantaneous luminosity in preparation of the trigger menu to distribute the bandwidth allocation among the different triggers.

In addition, it is important to monitor the trigger rates, trigger performance and CPU consumption during online data taking. Online monitoring provides real-time analytic data on high level quantities such as the number of CPU cores available to accept new events or the average event processing time. HLT resource usage is checked using cost monitoring tools during the trigger execution. These data are used by the trigger community to optimise, debug and develop the trigger menu. The offline monitoring provides much greater level of detail on resource usage within the HLT farm as well as monitoring histograms for physics and performance.

3 Trigger performance with 2017 data

Due to the excellent LHC performance, a peak luminosity of $1.74 \times 10^{34} \text{ cm}^{-2}\text{s}^{-1}$ and an average of 32.2 interactions per bunch crossing ($\langle\mu\rangle$), also known as pileup, were achieved in the first months of 2017 data taking. The trigger performance has been evaluated with the first proton-proton collision data acquired in 2017 and is shown below for different signatures.

Electrons are selected at HLT using likelihood-based identification criteria which takes as input electromagnetic shower shape and tracking information with different selections available (very loose, loose, medium and tight) [7]. Figure 6 shows the efficiency of the single electron trigger HLT_e28_lhtight_nod0_ivarloose, which applies a 28 GeV threshold, a likelihood-based tight identification and a track-based isolation at HLT, obtained with respect to the offline reconstructed electron candidates as a function of their E_T and the average number of interactions per bunch crossing. The offline reconstructed electrons are required to pass a likelihood-based tight identification and be isolated. The efficiency is measured using the $Z \rightarrow ee$ tag-and-probe method in data and Monte Carlo simulation, with no background subtraction applied.

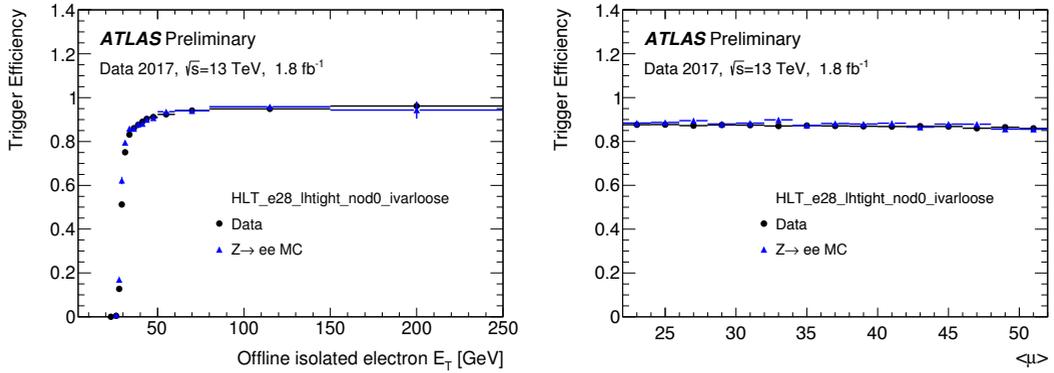


Figure 6. Efficiency of the HLT_e28_lhtight_nod0_ivarloose trigger with respect to the offline electron candidates as a function of their E_T (left) and the average number of interactions per bunch crossing (right) [8]. The efficiency is measured using the $Z \rightarrow ee$ tag-and-probe method in data and Monte Carlo simulation, with no background subtraction applied. Only statistical uncertainties are shown.

Photons are reconstructed at HLT using different working points (loose, medium and tight), applying a set of η - and E_T -dependent rectangular cuts on identification variables [9]. Figure 7 shows the efficiency of the leg of the diphoton primary trigger HLT_g35_medium_g25_medium_L12EM20VH with respect to the offline reconstructed photon candidates as a function of their E_T and the average number of interactions per bunch crossing. The trigger applies a 25 GeV threshold and medium identification criteria. The photon candidates reconstructed offline are required to be within $|\eta| < 2.37$, excluding the transition region between the barrel and endcap electromagnetic calorimeters at $1.37 < |\eta| < 1.52$, to pass the tight identification selection and be isolated. The efficiency is measured using the bootstrap method in data and $H \rightarrow \gamma\gamma$ Monte Carlo simulation, with no background subtraction applied.

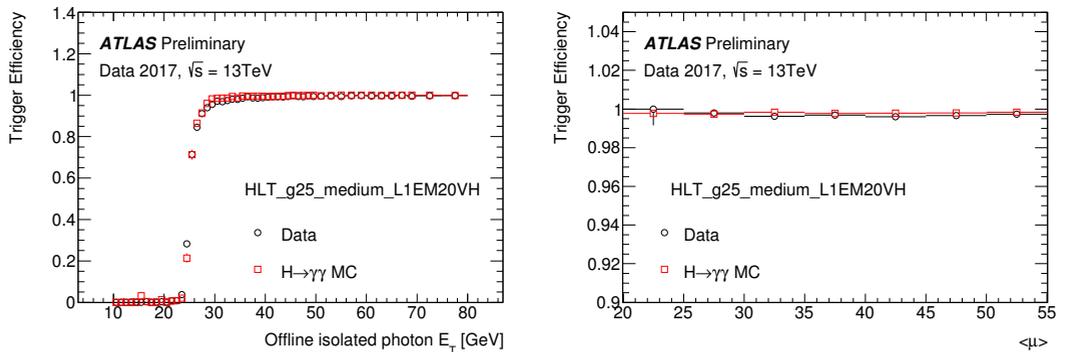


Figure 7. Efficiency of the HLT_g25_medium_L1EM20VH trigger with respect to the offline photon candidates as a function of their E_T (left) and the average number of interactions per bunch crossing (right) [8]. The efficiency is measured using the bootstrap method in data and $H \rightarrow \gamma\gamma$ Monte Carlo simulation, with no background subtraction applied. Only statistical uncertainties are shown.

Muons are reconstructed at HLT combining the inner detector and muon spectrometer information. Figure 8 shows the efficiencies of the lowest unprescaled single muon triggers at Level-1 and HLT with respect to the offline muon candidates requiring to pass a medium quality requirement [10], as a function of their p_T and ϕ , and the number of reconstructed vertices. The HLT_mu26_ivarmedium trigger applies a 26 GeV threshold and track-based isolation, while HLT_mu60 applies a 60 GeV threshold and no isolation requirements. The efficiency is measured using the $Z \rightarrow \mu\mu$ tag-and-probe method in data, with no background subtraction applied. The efficiencies in the barrel ($|\eta| < 1.05$) are lower than in the endcap ($1.05 < |\eta| < 2.4$) due to the reduced barrel geometrical acceptance.

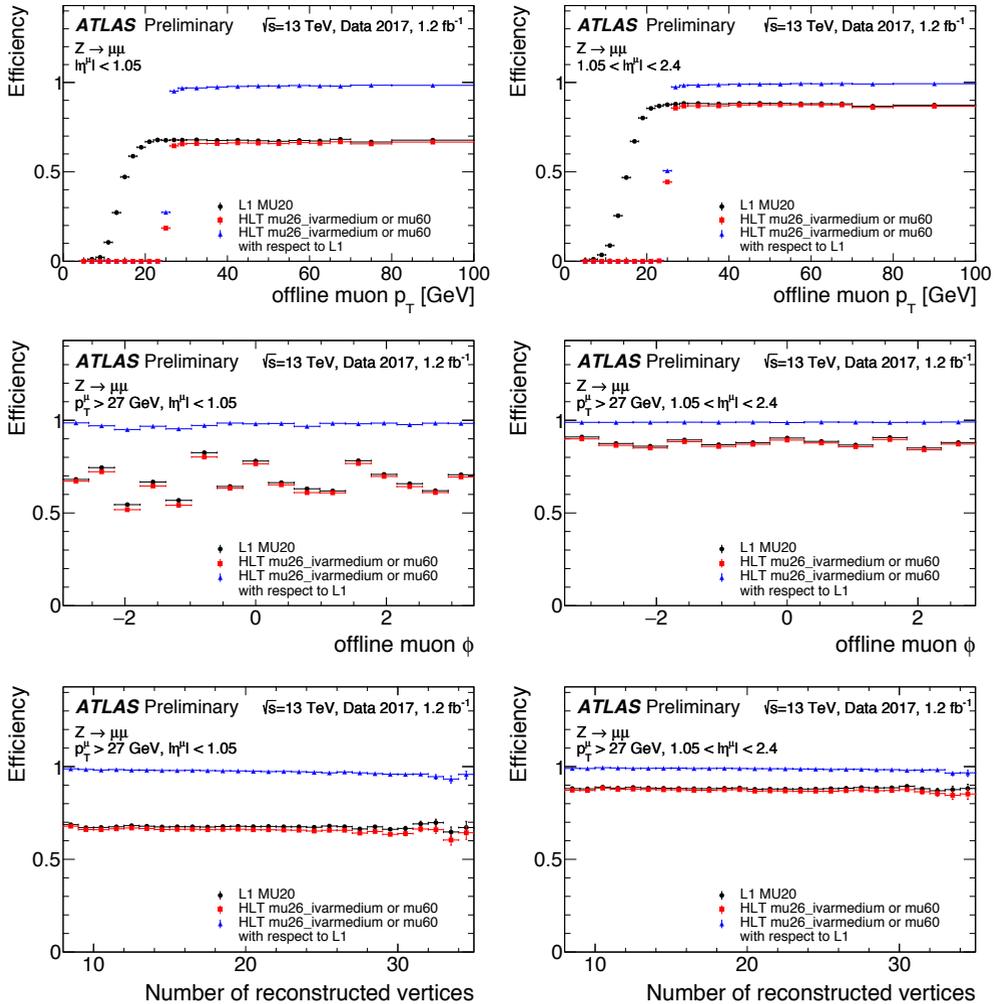


Figure 8. Efficiency of the logical OR of HLT_mu26_ivarmedium and HLT_mu60, and the corresponding Level-1 trigger (L1_MU20) with respect to the offline muon candidates as a function of their p_T (top) and ϕ (middle) and the number of reconstructed vertices (bottom) in the barrel (left) and endcap (right) detector regions [11]. The efficiency is measured using the $Z \rightarrow \mu\mu$ tag-and-probe method in data, with no background subtraction applied. Only statistical uncertainties are shown.

The jets are reconstructed at HLT using the anti- k_t jet finding algorithm [12] employing different distance parameters, $R = 0.4$ for small- R jets and $R = 1.0$ for large- R jets. In addition, jet substructure techniques, such as trimming [13], are applied to the large- R jets to mitigate contamination from soft radiation (initial state radiation, multiple parton interactions and pileup). Figure 9 shows the per-event trigger efficiency of a trimmed large- R dijet trigger as a function of the second leading offline trimmed jet mass. The efficiencies are measured using the bootstrap method in 2017 data.

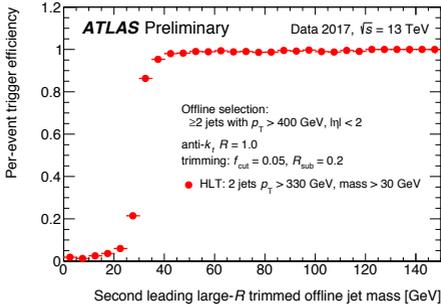


Figure 9. Efficiency of a trimmed large- R dijet trigger as a function of the second leading offline trimmed jet mass for jets with $p_T > 400$ GeV and $|\eta| < 2$ [14]. The trimmed large- R dijet trigger requires $p_T > 330$ GeV and an additional 30 GeV cut on the jet masses of the leading and second leading selected trimmed trigger jets. The efficiencies are obtained using the bootstrap method in data. Only statistical uncertainties are shown.

The main challenge of the E_T^{miss} trigger is the pileup mitigation. There was a significant effort to mitigate the impact on the trigger rates and deliver more efficient pileup suppression algorithms for 2017 data taking. Figure 10 compares the trigger cross sections for the baseline triggers used in 2016 (HLT_xe110_mht_L1XE50) and 2017 (HLT_xe110_pufit_L1XE50) as a function of the average number of interactions per bunch crossing.

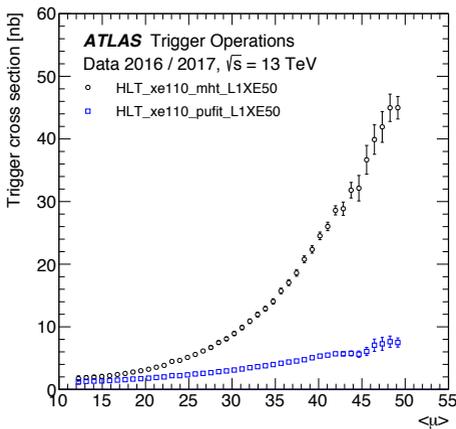


Figure 10. Trigger cross section measured using the online rate and luminosity, comparing the main E_T^{miss} trigger algorithms used in 2016 (“mht”) and 2017 (“pufit”), as a function of the average number of interactions per bunch crossing [15]. These triggers are used as representative benchmarks of the 2016 and 2017 data-taking campaigns, respectively.

The “pufit” algorithm computes the E_T^{miss} at HLT as the negative of the p_T vector sum of all calorimeter topological clusters corrected for pileup. The pileup correction is done by grouping the clusters into coarser towers which are then marked as pileup if their E_T falls below a pileup-dependent threshold. A simultaneous fit to both classes of towers is performed, taking into account resolutions, making the assumption that the contribution of the pileup to E_T^{miss} is zero. The fitted pileup E_T density is used to correct the above-threshold towers.

The “mht” algorithm computes the E_T^{miss} at HLT as the negative of the p_T vector sum of all jets reconstructed by the anti- k_t jet finding algorithm taking as input calorimeter topological clusters. These jets have pileup subtraction and jet energy scale calibration applied.

Figure 11 shows the efficiency of the E_T^{miss} triggers (HLT_xe110_mht_L1XE50 and HLT_xe110_pufit_L1XE50) as well as the corresponding Level-1 trigger (L1_XE50) as a function of the offline reconstructed E_T^{miss} (modified to count muons as invisible). The events shown are taken from data with a $W \rightarrow \mu\nu$ selection to provide a sample enriched in real E_T^{miss} .

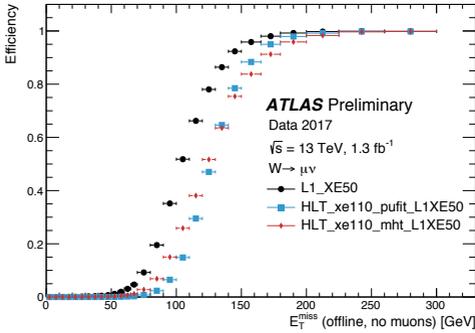


Figure 11. Efficiency of the HLT_xe110_mht_L1XE50 and HLT_xe110_pufit_L1XE50 triggers as well as the corresponding Level-1 trigger (L1_XE50) as a function of the offline reconstructed E_T^{miss} modified to count muons as invisible [15]. The events are taken from data with a $W \rightarrow \mu\nu$ selection to provide a sample enriched in real E_T^{miss} .

4 Conclusions

The trigger system has been successfully operated in 2017. The initial design luminosity of $10^{34} \text{ cm}^{-2}\text{s}^{-1}$ has been surpassed and the trigger menu has been prepared for twice this value and the high pileup conditions in 2017 data taking. Rock-solid well-established validation procedures are set up to ensure a smooth trigger operation. The first performance studies of various trigger signatures have been presented using $\sqrt{s} = 13 \text{ TeV}$ data acquired in 2017.

References

- [1] ATLAS Collaboration, JINST **3**, S08003 (2008).
- [2] L. Evans and P. Bryant (editors), JINST **3**, S08001 (2008).
- [3] ATLAS Collaboration, Eur. Phys. J. C **77**, 317 (2017).
- [4] ATLAS Collaboration, CERN-LHCC-98-014, <http://cdsweb.cern.ch/record/381429>.
- [5] ATLAS Collaboration, CERN-LHCC-2003-022, <http://cdsweb.cern.ch/record/616089>.
- [6] <https://twiki.cern.ch/twiki/bin/view/AtlasPublic/TriggerOperationPublicResults>.
- [7] ATLAS Collaboration, ATLAS-CONF-2016-024, <http://cdsweb.cern.ch/record/2157687>.
- [8] <https://twiki.cern.ch/twiki/bin/view/AtlasPublic/EgammaTriggerPublicResults>.
- [9] ATLAS Collaboration, Eur. Phys. J. C **76**, 666 (2016).
- [10] ATLAS Collaboration, Eur. Phys. J. C **76**, 292 (2016).
- [11] <https://twiki.cern.ch/twiki/bin/view/AtlasPublic/MuonTriggerPublicResults>.
- [12] M. Cacciari, G. P. Salam, G. Soyez, JHEP **04**, 063 (2008).
- [13] D. Krohn, J. Thaler and L.-T. Wang, JHEP **02**, 084 (2010).
- [14] <https://twiki.cern.ch/twiki/bin/view/AtlasPublic/JetTriggerPublicResults>.
- [15] <https://twiki.cern.ch/twiki/bin/view/AtlasPublic/MissingEtTriggerPublicResults>.