

Triggering in ATLAS in Run 2 and Run 3

Yasuyuki Okumura*

on behalf of the ATLAS collaboration,

International Center for Elementary Particle Physics, The University of Tokyo

E-mail: yasuyuki.okumura@cern.ch

The ATLAS experiment at the LHC can record about 1 kHz of physics collisions, out of the LHC design bunch crossing rate of 40 MHz. To achieve a high selection efficiency for rare physics events while reducing the significant background rate, a two-level trigger system is used. The event selection is based on physics signatures, such as the presence of energetic leptons, photons, jets or missing energy. In addition, the trigger system can exploit algorithms using topological information and multivariate methods to carry out the filtering for the many physics analyses pursued by the ATLAS collaboration. In Run 2, around 1,500 individual selection paths, the trigger chains, were used for data taking with specified rate and bandwidth assignments. We will give an overview of the Run 2 trigger menu and its performance, which supports a broad physics program. We present the tools that allow us to predict and optimise the trigger rates and CPU consumption for the anticipated LHC luminosity. They are essential components to react to the changing LHC conditions and data taking scenarios. As an outlook to the upcoming ATLAS data-taking period in Run 3 from 2022 onward, we present the design principles and ongoing implementation of the new trigger software in a multi-threaded framework AthenaMT together with some outlook to the expected performance improvements.

*** *The European Physical Society Conference on High Energy Physics (EPS-HEP2021), ****

*** *26-30 July 2021 ****

*** *Online conference, jointly organized by Universität Hamburg and the research center DESY ****

*Speaker



1. Introduction

Online event selection provided by the trigger system is exploited to selectively store events of interest within a bandwidth limitation of the data acquisition (DAQ) system and to realise efficient use of computational resources for offline event reconstruction. The trigger system of the ATLAS detector [1] consists of a Level-1 (L1) and a High Level Trigger (HLT) [2]. The L1 is a hardware-based processor system consisting of dedicated custom-made electronics. The designed rate of the events that the L1 accepts is 100 kHz at maximum, and the L1 algorithms run with a fixed latency of about $2.5 \mu\text{s}$. The HLT exploits software-based object reconstruction running on commodity CPU farms with networking built with commercial technologies. The HLT algorithms perform online reconstruction with respect to Region of Interests (RoIs) given by the L1 algorithm, to achieve fast reconstruction, and select events to be recorded to permanent storage. The recording rate of HLT is typically 1-1.5 kHz on average.

The LHC instantaneous luminosity has evolved through Run 2 with the typical peak luminosity of $0.5 \times 10^{34} \text{ cm}^{-2}\text{s}^{-1}$ in 2015, $1.3 \times 10^{34} \text{ cm}^{-2}\text{s}^{-1}$ in 2016, $1.6 \times 10^{34} \text{ cm}^{-2}\text{s}^{-1}$ in 2017, and $2.0 \times 10^{34} \text{ cm}^{-2}\text{s}^{-1}$ in 2018. The expected peak instantaneous luminosity during the Run 3 is $2.0 \times 10^{34} \text{ cm}^{-2}\text{s}^{-1}$ with the luminosity levelled at the peak for 6-10 hours, depending on the conditions of the LHC machines. Both the L1 hardware systems and HLT algorithms have evolved throughout Run 2 and will be further improved for Run 3. A robust and flexible operation scheme has been a key technique along with precise monitoring and prediction methodology for operation costs. In this article, we discuss the improvements of triggers through Run 2 in Section 2, the developed operation scheme in Section 3, and new features for Run 3 in Section 4, and Section 5 concludes the discussion.

2. Evolutions of triggering in ATLAS through Run 2

Improvement of the L1 muon trigger is an important example among those made for the L1 trigger system during Run 2, integrating new hardware and firmware in the system. The L1 muon trigger improvement exploits additional coincidence conditions with the inner stations of muons and tile hadronic calorimeter [3]. The requirement rejects fake contributions owing to charged particles not originating from the interaction point. Figure 1 (a) shows the event rejection thanks to the additionally required coincidence with the signal of the hadronic calorimeter in the outermost layer, where the region in $|\eta|$ is between 1.05 and 1.3. Another example of the L1 system improvements is an introduction of a new FPGA-based L1Topo processor, which has been in operation since 2017 [4]. The L1Topo allows us to require topological conditions in the L1 trigger stage such as $\Delta\eta$, $\Delta\phi$, ΔR , and invariant mass using the L1 trigger objects provided by the L1 calorimeter and muon algorithms. Figure 1 (b) shows rate reduction owing to the additional topological requirements for a L1 trigger designed to select events coming from the $J/\psi \rightarrow \mu\mu$ process.

We have also been upgrading the HLT software algorithms in terms of both the performance and CPU resource usage throughout Run 2. Among many important new features introduced during Run 2, improving b-jet reconstruction algorithms, which require computationally expensive tracking, is a typical effort to achieve stable HLT operation [5, 6]. An introduction of a super-RoI approach for online tracking (see Figure 2 (a)) allows us to avoid duplicating track reconstruction in




Figure 1: (a) Pseudorapidity distribution of the L1 RoIs (η^{RoI}) that satisfy the 20 GeV requirement (L1_MU20) after the additional coincidence with tile hadronic calorimeter coincidence in $1.05 < |\eta| < 1.3$ (solid black line) with a reference (blue triangles) that shows the η^{RoI} without the tile coincidence. The difference represents the gain in the fake reduction by the additional coincidence conditions [3]. (b) The rate of L1 triggers selecting two muons, each with transverse momentum above 6 GeV with (blue) and without (red) the additional L1Topo requirements on the invariant mass and ΔR of the two L1 muon objects [3, 4].

overlapping regions of RoIs. These efforts for the optimisation of software have reduced the CPU resource usage by $\sim 30\%$. The Run 2 b-tagging algorithm has fully exploited the new pixel layer, which provides a significant gain in the performance. Furthermore, an optimal use of multivariate analysis improves the rejection of light-flavoured jet backgrounds by $\sim 50\%$ (see Figure 2 (b)).

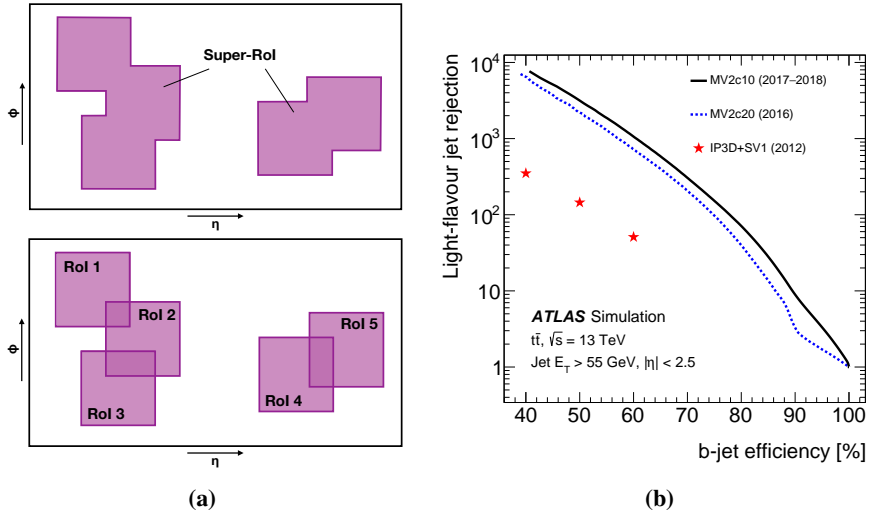


Figure 2: (a) The concept of a super-RoI approach to avoid duplication of online tracking in overlapping regions of multi-RoIs [5, 6]. (b) The expected light-flavour jet rejection factors as a function of b-jet efficiency for the MV2c10 tagger (black line, used in data taking in 2017 and 2018) and MV2c20 tagger (blue dotted line, used in 2016). As a reference, performance of the Run 1 algorithm (red stars, “IP3D+SV1”) is also shown [5].

3. Operation models of the ATLAS trigger

The series of the L1 and HLT algorithms are called trigger chains. We have a mechanism to control the rates of accepted events for individual chains, denoted as prescales. Evolving the collection of trigger chains and controlling the prescales allow us to handle the trigger rates for the increasing instantaneous luminosity throughout Run 2. Further, dynamic control of prescales during data taking is used to fully exploit the available bandwidth of L1 readout and the HLT output. Figure 3 (a) shows the total L1 trigger output rate as a function of time in an LHC fill, with the instantaneous luminosity overlaid. In the fill, the peak luminosity was $2.0 \times 10^{34} \text{ cm}^{-2}\text{s}^{-1}$ and the peak pile-up, $\langle\mu\rangle$, which is defined as the average of number of collisions per bunch crossing, was 56. The allocation of the readout bandwidth for individual L1 algorithms is dynamically changed with configurable L1 prescales. Particularly, at the end of fill, the allocation of the L1 jet algorithms for Trigger-Level Analysis [7] is significantly extended to realise the most efficient use of the available L1 bandwidth with the low-luminosity conditions. Various data streams are defined in ATLAS DAQ for dedicated purposes and different optimal schemes of offline reconstruction, such as Physics Streams, Express, Debug, Calibration, Trigger-Level Analysis, and Monitoring streams, and events that pass the HLT selection are classified to streams according to the satisfied chains. The output composition of various data streams dynamically changes to optimally use the available bandwidth (Figure 3 (b)).




Figure 3: (a) The total L1 trigger output rate and (b) the recording bandwidth at the HLT for individual data streams as a function of time in an LHC fill with a peak luminosity of $2.0 \times 10^{34} \text{ cm}^{-2}\text{s}^{-1}$. The LHC instantaneous luminosity is also shown for reference [7, 8].

Handling the CPU resource usage and trigger rates has been a key role of the trigger operation for stable operation throughout Run 2 [7]. We have established a framework to monitor and predict the CPU usage and data flow over the network. Figure 4 (a) shows distributions of processing times per event for the topological clustering of calorimeter data and the inner-detector electron track identification as examples of the cost monitoring of the HLT algorithms. A special dataset that consists of events selected by representative Level-1 algorithms without bias of HLT selection is collected and used for the precise trigger rate estimation, called “enhanced bias data sample”. The dataset allows us to gain statistical power over orders of magnitude in the high p_T regime of the L1 spectrum for the rate estimation with various threshold conditions as shown in Figure 4 (b). Appropriate event weights are applied to achieve a proper trigger rate estimation, removing

the L1 selection bias. The enhanced bias datasets are collected with representative LHC conditions for each data taking period, especially when the instantaneous luminosity and $\langle\mu\rangle$ has changed significantly. The accuracy of the prediction has been confirmed to be compatible with its statistical uncertainties by a comparison between the observed and predicted trigger rates for the major chains.




Figure 4: (a) Monitoring of process time per event of the topological clustering and inner-detector electron track identification [7]. (b) L1 rate prediction as a function of the L1 threshold conditions. The enhanced bias dataset is used predict the L1 rate precisely with statistical power over a wide-ranged spectrum [7].

4. New features of triggering in ATLAS for Run 3

The L1 and HLT triggers are being upgraded for Run 3 starting in 2022. For the L1 calorimeter trigger, new FPGA-based feature extractors are introduced by upgrading the backend electronics, and upgraded frontend electronics extends the online readout capability of the electromagnetic calorimeter. The individual energy deposits in four electromagnetic calorimeter layers are available in the online algorithms with an improved granularity in η - ϕ . The additional information related to the shower evolution improves the selectivity for the electrons and photons in the L1 calorimeter trigger [9, 10]. The hadronic object reconstruction in the L1 stage is also significantly improved with the new processor. Besides the L1 calorimeter upgrade, the L1 muon system is upgraded to improve the inner coincidence functionality with new detectors such as the New Small Wheel [10, 11].

The software framework with which we process the ATLAS data from the HLT to the offline reconstruction, Athena [12], will be reorganised in a multi-threaded framework for Run 3, AthenaMT. It is implemented with a concurrent task scheduler to support both the intra- and inter-event parallelism. It supports efficient memory use by sharing data between parallel processes [13]. Further, common use of the AthenaMT framework between HLT and offline reconstruction code maximises the sharing of algorithms. The HLT algorithm developments for the Run 3 are in progress for various trigger signatures. Among them, one remarkable new feature is running the HLT tracking in the full detector region for certain signatures in the trigger selection. It provides improvement in online energy calibrations and coherency with offline reconstruction. We have achieved a significant speedup of the online track reconstruction to make this possible. Machine learning extends filtering

on pixel detector doublet space points in seed processing, which results in a significant speedup without major efficiency loss at the expected peak instantaneous luminosity [14].

5. Conclusion

Throughout Run 2, to cope with the increase of the LHC instantaneous luminosity, we have been improving the L1 hardware system and HLT algorithms for triggering in ATLAS. Additionally, we have established a model for the robust trigger operation as well as a precise prediction methodology for the resource usage and trigger rates for both L1 and HLT. For triggering in Run 3, we are upgrading the L1 trigger hardware system. The new multi-threaded Athena framework (AthenaMT) will maximise the efficiency of memory usage with massive parallelism, and the common use of the AthenaMT will maximise the sharing of algorithms between offline and online reconstruction. The HLT algorithms are being developed for Run 3 to improve the resource usage and performance and to realise new features such as an online tracking in the full detector region.

References

- [1] ATLAS Collaboration, JINST 3 (2008) S08003.
- [2] ATLAS Collaboration, Eur.Phys.J.C 77 (2017) 5 317, CERN-EP-2016-241.
- [3] ATLAS Collaboration, JINST 15 (2020) P09015, CERN-EP-2020-031.
- [4] ATLAS Collaboration, CERN-EP-2021-040, arXiv:2105.01416 [hep-ex].
- [5] ATLAS Collaboration, CERN-EP-2021-032, arXiv:2106.03584 [hep-ex].
- [6] ATLAS Collaboration, CERN-EP-2021-076, arXiv:2107.02485 [hep-ex].
- [7] ATLAS Collaboration, JINST 15 (2020) 10 P10004, CERN-EP-2020-109.
- [8] ATLAS Collaboration, <https://twiki.cern.ch/twiki/bin/view/AtlasPublic/TriggerOperationPublicResults>.
- [9] ATLAS Collaboration, <https://twiki.cern.ch/twiki/bin/view/AtlasPublic/L1CaloTriggerPublicResults>.
- [10] ATLAS Collaboration, CERN-LHCC-2013-018, ATLAS-TDR-023, <https://cds.cern.ch/record/1602235>.
- [11] ATLAS Collaboration, CERN-LHCC-2013-006, ATLAS-TDR-020, <https://cds.cern.ch/record/1552862>.
- [12] ATLAS Collaboration, 2019, Athena (22.0.1), Zenodo <https://doi.org/10.5281/zenodo.2641997>
- [13] Rafal Bielski on behalf of the ATLAS Collaboration, J. Phys.: Conf. Ser. 1525 012031 (2020).
- [14] ATLAS Collaboration, <https://twiki.cern.ch/twiki/bin/view/AtlasPublic/HLTTrackingPublicResults>.