13 September 2024

All About the Rates: ATLAS Trigger Menu Analysis

Clarisse Prat

Summary

Report on the 13 weeks spent working as a Summer Student with the ATLAS Trigger Menu group under the supervision of Sofia Cella and Dr Stefanie Morgenstern.

The following poster was presented as part of the Summer Project: ATLAS Menu Poster

1 Introduction

Data-taking at any of the detectors at CERN is an enormous undertaking, and relies on complex and state-of-the-art Trigger and Data Acquisition (TDAQ) systems to select, save and process data. This is no different at the ATLAS experiment [1], where the data from the 40 MHz collision rate of the Large Hadron Collider (LHC) needs to be decreased to the manageable data output rate of 3 kHz (or 8 GB/s). This is no small feat, and relies on innovative engineering, computing and physics to ensure that the correct data is selected and saved for analyses to be carried out on.

One cornerstone of the process is the Trigger Menu, which dictates which triggers need to be passed for an event to be saved and sent for processing. This operational project was focused around analysing the Menu and developing tools for Trigger Experts and Developers to use to be able to analyse the rates of triggers and ensure optimal rate and resource usage.

2 The ATLAS Trigger System

The trigger system is a critical component of the ATLAS experiment and the data-taking process [2]. For Run 3, there are two levels to the system: Level-1 (L1) and the High-Level-Trigger (HLT), which work together to select which events are saved to be processed.

2.1 Level-1 Trigger

Data from the experiment needs to go through two important components in the L1 Trigger system as a first step in data acquisition: the Level-1 Calo (L1Calo) and the Level-1 Muon (L1Muon). **L1Calo** receives inputs from the calorimeter detectors: the Liquid Argon Calorimeter and the Tile Calorimeter. **L1Muon** receives inputs from the muon spectrometer. **L1Topo** is the next step in the L1 system and allows for more complex algorithms at L1 based on topological selections. It applies real-time kinematic and angular selections to Trigger Objects. Information from L1Topo is sent to the Central Trigger Processor (CTP). The **CTP** receives inputs from all of the L1 Trigger systems, and produces 512 L1 trigger items based on logical combinations of these inputs. The CTP combines this information in a format that is readable to the HLT, by identifying *Regions of Interest (RoIs)* with coarse information on the geometrical region in the muon spectrometer and the calorimeters, the type of object (corresponding to EM, tau, jet or muon) and the threshold passed (p_T, E_T).

2.2 High-Level-Trigger

The HLT reconstruction typically uses dedicated fast trigger algorithms to provide early background rejection, which is then followed by more precise and resource (CPU) intensive algorithms to make the final selection. The reconstruction algorithms reconstruct objects, and after which physics selections are applied. Athena, the ATLAS software framework [3], is used in the HLT reconstruction process.

2.2.1 HLT Chains

The HLT chain name contains all of the information of the chain configuration, in the format of `HLT_<HLT configuration>_L1<L1 seed threshold>`, and are categorised by their use cases. The **primary chains** are used for physics signals and in general all events passing these chains are recorded. The **back-up chains** are used in case of an unexpected luminosity increase and have higher thresholds compared to the primary chains. **Supporting chains** are chains that are used to support a physics analysis, (using data-driven background methods, efficiency measurements, etc.). Finally, **monitoring chains** are used to monitor the detector and to ensure good data quality, usually in the form of detector monitoring and calibration.

2.2.2 HLT Streams

The HLT chains are grouped together into **streams**, more formally defined as a collection of events or event fragments in the same dataset depending on the HLT decision. The design of the streaming model is inclusive (barring the debug streams), which means that an event can be written to multiple streams, although streams are designed to have minimal overlap.

Physics streams include full event building with information from all sub-detectors, with an event size of ∼2 MB and an output rate of ∼3 kHz. The **Main stream** is used for most physics analyses. **Delayed streams** contain events that are not automatically reconstructed at Tier-0 to rather use the resources during downtime. The **Enhanced Bias** stream is used to record events for trigger rate predictions and is biased towards high-p_T events likely to be selected by the Trigger.

Trigger Level Analysis (TLA) streams are streams with reduced event content, but still part of the Physics streams. These streams only store specific physics objects reconstructed by the HLT instead of the full ATLAS event information. As a result, there is a much smaller event size and a much higher rate per HLT chain possible, since the HLT output is limited by the bandwidth and not by the rate itself.

3 The ATLAS Trigger Menu

The ATLAS collaboration has a dedicated physics programme which dictates what physics is being searched for, and the Trigger system is responsible for selecting events aligned with the programme. This is all put together in the **Trigger Menu**, a collection of trigger chains and the corresponding set of rules that defines the rate that each chain can write out. Events are only selected to be recorded if they satisfy the conditions of one or more chains on the Menu. The rules that define the output rate of each chain come in the form of a **prescale value** (or simply a prescale) that controls the rate of accepted events to stay within rate, bandwidth and CPU constraints. These prescales are computed using the Trigger Rulebook.

The Menu is constantly being updated and fine-tuned according to the running conditions and new developments. There are physics priorities aimed at taking a balanced dataset for analyses. Supporting triggers for calibration and efficiency measurements also need to be included. Online and offline resources available also contribute to the design of the Menu and it varies with luminosity and time. Additionally, there are different menus defined for specific purposes and conditions: the proton-proton and Heavy-Ion physics menus differ considerably, there is a Monte Carlo menu used in event generation and there are Cosmic and Standby menus for technical stops and machine development.

4 It's All About the Rates

The rates of chains and streams in the Trigger Menu can be analysed to ensure that the rate limitations of the HLT are not exceeded and that the Menu is operating optimally. This is done using the Rates Analysis package, a suite of tools for trigger rates calculations, which is found in Athena. Rate predictions are run on an Enhanced Bias dataset, the results of which are used by the Rulebook to calculate the prescales that are applied in data-taking. The actual rate output of different chains and streams can then be analysed to control or improve the resource allocation to each chain. The Trigger Menu is periodically being updated to include new triggers and optimise streaming. Each time this is done, the rates analysis should be performed to check the output rates of updated streams and chains. The process is repeated on a regular basis, and a script that can perform an analysis of the rate usage of different chains and streams is a useful tool to facilitate these analyses. This would be very useful for developers of new triggers to assess the rate contribution of new trigger chains and overlap with existing chains

As mentioned previously, the streaming model is inclusive and so chains can be written to multiple streams, and the same events can be recorded in different streams. The *overlap* between two streams is a measure of how many events are written into the streams before being processed, and the overlap should be minimised as to not waste resources. The contribution of a chain to a stream can also be determined, as well as the overlap between two chains. It is also possible to create a set of chains and find the unique rate of the set with respect to the entire Menu. Prescale files can be changed when fed into the script, depending on what is needed.

The project involved writing a script that adds four different functionalities to be used on top of the Rates Analysis package. The four functionalities are calculating the:

1. overlap between two streams
2. unique contribution of a chain to a stream
3. overlap between two chains
4. unique rate of a set of chains

4.1 Overlap between Two Streams

By performing the analysis on the chosen streams, Trigger Menu Experts can decide whether or not chains from a chosen stream should be moved into a different stream to avoid wasting resources due to duplication. This is done by finding the *unique stream rate and overlap between streams* of the selected stream. To do this, the script enables all chains in Stream A (which as a default is set to Main, but can changed to any other stream) and Stream B (which can be any other stream), determines the Global HLT Rate for the combination, and then subtracts the rate when only chains in Stream A are enabled.

The script takes two flags `--stream0` and `--stream` as arguments and creates two prescale files: one where only Stream A is enabled, and one where Stream A and Stream B are enabled. Stream A is the reference stream (passed by `--stream0`), and as a default is set to Main. A different reference chain can be chosen by passing it as an argument after this flag. A list of streams can be passed as an argument to `--stream`, and the script iterates over the list. It then calls the `run_RatesAnalysis.py` script (already included in the Rulebook and which calculates the rates of all the chains in the Menu, based on a specific prescale set) and runs the script for each combination of Stream A and Stream B, and writes a text file with the Global HLT Rate for Stream A only, the Global HLT Rate for Stream A and B, and calculates the difference between the rates to get the unique rate of Stream B. For bookkeeping and keeping track of which streams were compared, the script creates a directory for each pair that is compared, and inside that has two directories for the `stream0` and `stream` RatesAnalysis outputs. The information for the unique rate and the overlap is stored in a text file in the directory for each iteration.

4.1.1 Example: Overlap between the stream0 and three other streams

The example illustrates the use of the script by analysing the stream overlap between the Main stream and three other streams, running over 100 000 events from data processed with Athena release 24.0.52. The Main stream needs to be processed immediately during data-taking, and if the overlap between the streams are high, it can be considered to move chains to one of the other streams.

The results are shown below:

Stream	**Rate (Hz)**	**Error (Hz)**	**Unique Rate w.r.t. Main (Hz)**
Main	2160.1456	53.1707	
Stream 1	347.3337	22.6622	240.7915
Stream 2	475.1429	96.4235	449.3020
Stream 3	6858.5519	202.172	5907.5888

We see that the overlap between Main and Stream 1 is 30.67%, while Stream 2 has an overlap of 5.43% and Stream 3 has an overlap of 13.87%.

4.2 Unique Contribution of a Chain to a Stream

The unique contribution of a chain to a stream can be used as a survey to see if any chains should be moved to a different stream. If a primary chain in the Main stream has an unusually high rate, the chain could be moved to a delayed stream to optimise resource usage. The events would then be processed at a later time when resources are available. This is done as to not slow down the the processing of the Main stream data. The setup for this analysis was similar to that of the stream overlap, but the script enables chains in the the given stream (with the chain selected enabled), gets the prescales for the chains and then determines the Global HLT Rate of the stream. The script then gets another prescale file in which the given stream is enabled but the chain is *disabled* and subtracts this rate from the Global HLT Rate. The result is then the unique rate contribution of the chain.

An extra functionality of the script is that chains can be added using RegEx by passing the argument `regex` to the flags `--chain`, `--chain_comp` and `--chain_set`.

4.2.1 Example: Jet Chain

In this example, we found the overlap between a jet chain and the given stream. The analysis ran over 500k events in the processed dataset for Athena release 24.0.52. The chain requires a single jet ($p_T > 0$ GeV), with a central eta tag (eta value between 0 and 240) and requires that the chain reconstructs particle flow jets which needs full scan fast track finding. The chain is only fed to the Main stream.

	Rate (Hz)	**Rate Error (Hz)**
Chain	38.4361	0.9789
Stream and Chain	2518.1376	38.8825
Stream w/o Chain	2479.7015	38.811
Unique Contribution	38.4361	—

The unique rate of the chain is taken from the menu with the chain enabled, and is given as 38.4361 Hz. This is the same value as the chain rate, since the chain under consideration was only fed to the Main stream.

4.2.2 Example: DarkJet Chain

Considering a chain that is similar to the one outlined in the previous example except that it has a Dark Jet requirement, the unique contribution of the Dark Jet chain to the Main stream was analysed. Note that the chain is not fed into the Main stream, but rather the DarkJetPEBTLA stream.

	Rate (Hz)	**Rate Error (Hz)**
Chain	2.3314	0.052
Stream and Chain	2494.3401	44.9878
Stream w/o Chain	2494.3401	44.9878
Unique Contribution	0.0000	—

4.2.3 Example: Muon Chain

In this example, consider a muon chain that is fed into the BPhysDelayed stream and the Main stream.

	Rate (Hz)	**Rate Error (Hz)**
Chain	22.4568	5.7983
Stream and Chain	2348.8865	54.587
Stream w/o Chain	2344.4024	54.5254
Unique Contribution	4.4841	—

The muon chain selected has a unique contribution of 0.19% to the Main stream.

4.3 Overlap between Two Chains

Calculating the overlap between two chains is a useful tool for Trigger Experts and Developers of new triggers to see if events pass multiple chains and if the overlap between the chains is too large. There is one prescale file where only one chain is enabled and everything else in the menu is disabled, and another prescale file where both chains are enabled. The script has the option of comparing multiple chains to a single reference chain. The chains are entered as a list and the user does not need to run the script for each comparison chain.

4.3.1 Example: Jet Chains

In this case, the prescales are set to 1, however the user could change the values in the input prescale file. The example considers jet chains with the same L1 seed that requires a minimum p_T of 160 GeV, with the reference chain having a p_T cut of 360 GeV at the HLT level, and the three comparison chains having increasing p_T cuts. We see that the

	Weighted PS Rate (Hz)	**Unique Rate (Hz)**
Reference Chain	90.0203	45.8392
Comparison Chain 1	44.1811	0.0000
Reference Chain	90.0203	54.7757
Comparison Chain 2	35.2447	0.0000
Reference Chain	90.0203	62.7814
Comparison Chain 3	27.2390	0.0000

weighted PS rate (the rate of the chain after the prescales have been applied) for the reference chain is the same in each comparison. The weighted PS rate of the comparison chain, plus the unique rate of the reference chain, adds up to the weighted PS rate of the reference chain. The comparison chains have no unique rate, because if an event passes their selection criteria, it also passes the selection criteria of the reference chain.

It follows that the weighted PS rate of the comparison chains decrease as the selection criteria of the chains become stronger, and subsequently we see that the unique rate of the reference chain increases as the jet p_T cut increases. The percentage overlap of the reference chain with the three comparison chains is 49.07%, 39.15% and 30.25%.

4.3.2 Example: Similar Electron Chains with Different Selection Criteria

In this example, the reference electron chain has a higher L1 selection criteria (a p_T greater than 28 GeV) than the comparison chain (a p_T greater than 26 GeV), and a higher HLT p_T threshold (30 GeV as opposed to 28 GeV).

	Weighted PS Rate (Hz)	**Unique Rate (Hz)**
Reference Chain	864.4192	678.3463
Comparison Chain	227.7408	41.6680

The rates show that there is indeed an overlap of 26% between the two chains.

4.4 Unique Rate of a Set of Chains

The unique rate of chains is displayed in the standard Rates Analysis results, but the way that it is determined means that if an event triggers multiple chains, it does not contribute to the unique rate of the chain. It then appears that a set of chains has no unique rate although this set of chains selects events which are not selected by any other chain in the menu. In analogy to the streams analysis, we define our set of chains as a 'stream' and perform the rates analysis on the full menu, and then the rates analysis when all of the chains in the set are disabled. The difference in the Global HLT Rate is then the unique rate of the set. This is in particular useful for developers planning to add multiple chains to the menu, e.g. a mix of primary and support chains targeting a specific analysis signature.

4.4.1 Example: hitdv

The set of chains contain the tag `"hitdv"`, meaning we look for hits with displaced vertices. For this set of 20 chains, the script was on 500k events from processed AODs. The unique rate of the hitdv set is 0.1330 Hz for the changed and unchanged prescale files.

	Global HLT Rate [Hz]	**Global HLT Rate Error [Hz]**
Full Menu Enabled	9762.7216	116.9259
Chains Disabled	9762.5886	116.9224
Unique Set Rate	0.1330	

5 Conclusion

Analysing the chains and streams in the Trigger Menu is a critical task to ensure that the correct data is selected and saved by the ATLAS experiment. This script adds useful and important functionalities to the Rates Analysis package that developers and Menu Experts can use to analyse the Trigger Menu. Overlap between two streams and two chains, unique contributions of chains to streams and the unique rate of a set are important add-ons that would help in Menu monitoring and the development of new triggers.

References

[1] ATLAS Collaboration. "The ATLAS Experiment at the CERN Large Hadron Collider: a Description of the Detector Configuration for Run 3". In: *Journal of Instrumentation* 19.05 (May 2024). ISSN: 1748-0221. DOI: 10.1088/1748-0221/19/05/p05063. URL: http://dx.doi.org/10.1088/1748-0221/19/05/P05063.

[2] ATLAS Collaboration. "The ATLAS Trigger System for LHC Run 3 and Trigger Performance in 2022". In: *Journal of Instrumentation* 19.06 (June 2024). ISSN: 1748-0221. DOI: 10.1088/1748-0221/19/06/p06029. URL: http://dx.doi.org/10.1088/1748-0221/19/06/P06029.

[3] ATLAS Collaboration. "Athena, Version 22.0.1". In: (Apr. 2019). DOI: 10.5281/zenodo.2641997. URL: https://doi.org/10.5281/zenodo.2641997.