

Implementation of the ATLAS trigger within the multi-threaded AthenaMT framework

Alison Elliot* on behalf of the ATLAS Collaboration

Queen Mary University of London

E-mail: alison.elliott@qmul.ac.uk

Athena is the software framework used in the ATLAS experiment throughout the data processing path, from the software trigger system through offline event reconstruction to physics analysis. The shift from high-power single-core CPUs to multi-core systems in the computing market means that the throughput capabilities of the framework have become limited by the available memory per process. For Run 2 of the Large Hadron Collider (LHC), ATLAS has exploited a multi-process forking approach with the copy-on-write mechanism to reduce the memory use. To better match the increasing CPU core count and the, therefore, decreasing available memory per core, a multi-threaded framework, AthenaMT, has been designed and is now being implemented. The ATLAS High Level Trigger (HLT) system has been remodelled to fit the new framework and to rely on common solutions between online and offline software to a greater extent than in Run 2. Presented is the implementation of the new HLT system within the AthenaMT framework, which will be used in ATLAS data-taking during Run 3 (2021-2023) of the LHC.

*XXIX International Symposium on Lepton Photon Interactions at High Energies - LeptonPhoton2019
August 5-10, 2019
Toronto, Canada*

*Speaker.

1. Motivation for ATLAS software upgrades

In order to process the data recorded at the ATLAS [1] detector in the upcoming Run 3 of the LHC, the reconstruction software will need to operate at the highest capabilities possible. The trend in commercial computing power is away from single core CPUs and towards multi-core systems. As seen in Figure 1, the levelling in CPU frequency and single-core performance, combined with the increase in number of logical cores, means that the maximum processing power is from a multi-threaded software system.

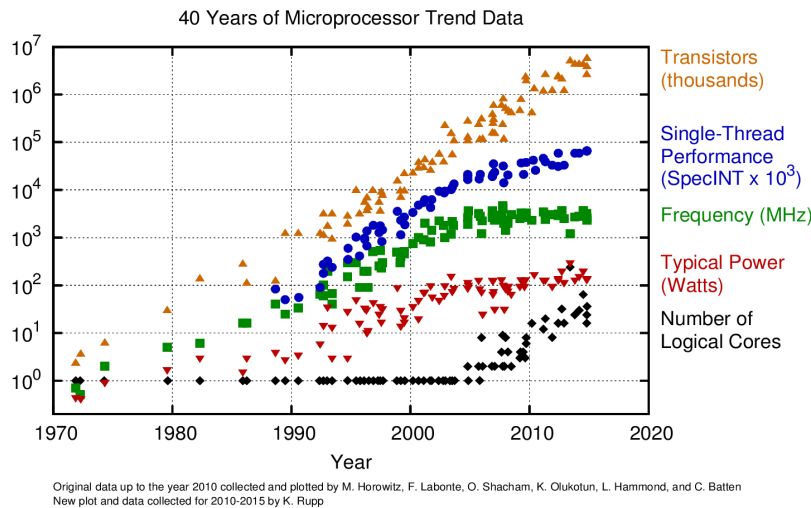


Figure 1: A plot of 40 Years of Microprocessor Trend Data [2]. The plot shows that the performance and frequency of single-threaded CPUs have been levelling off for more than a decade. However, the number of cores along with the number of transistors are increasing exponentially in accordance with Moore’s law.

The LHC, in normal operations for Run 2 and Run 3, collides protons at 40 MHz inside the ATLAS detector. This produces too much data to record, so ATLAS employs a two level trigger system. First, the level one (L1) hardware trigger reduces the data to 100 kHz with custom built electronics. The second level trigger is known as the High Level Trigger (HLT). The HLT is made up of 40k physical cores. It is a software trigger that further reduces the rate to about 1 kHz.

2. ATLAS software

The ATLAS software framework, Athena [3], uses dynamically assembled sequences of algorithms in order to perform various event processing workflows needed for simulation, reconstruction, and online triggering. In Run 2 the processing was organised with sequential execution of algorithms, and trigger steering code was written exploiting the knowledge of the input and output of each algorithm. In Run 3, in order to take advantage of the latest trends in computing, algorithms will be evaluated in parallel.

The multi-threaded framework, AthenaMT, will provide infrastructure for parallel execution of algorithms and parallel event processing, both inter- and intra-event, as in Figure 2. AthenaMT is based on the Gaudi Hive [5], a modern high energy physics data framework, that contains a

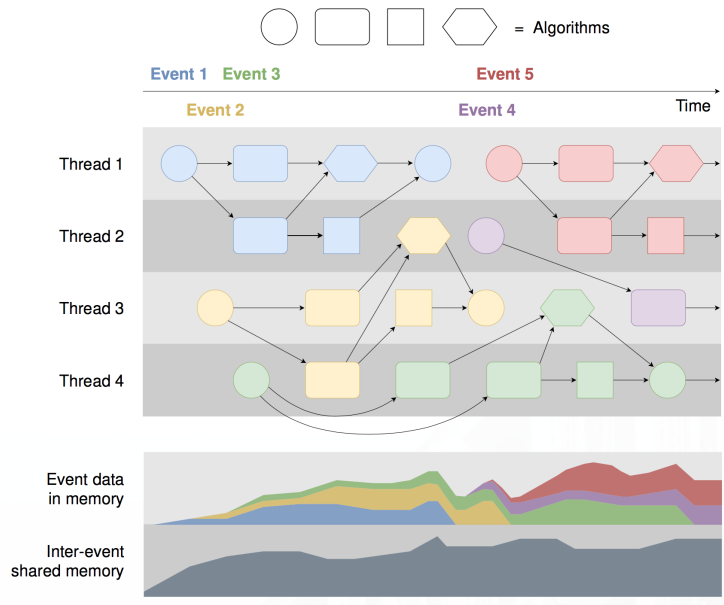


Figure 2: Illustration of inter- and intra-event multi-threaded processing in Athena, where the parallel processing is illustrated with five events and their algorithms in four threads. The colours represent different events, and the shapes represent different algorithms. [4]

scheduling component for algorithms allowing for the above mentioned parallelism to optimise the throughput. In order to achieve that, the central scheduler will use specific data handles to process inputs and outputs of each algorithm.

It is preferable that online HLT efficiency be high with respect to offline selection. Using the same Athena code ensures the online and offline requirements are very similar, so that the trigger efficiency can be very high in the phase space defined by the offline requirements. Due to differences in the HLT framework this required adapting interfaces between the HLT and the offline code. In the Run 3 framework this has been accounted for during the design phase, and therefore offline algorithms in the HLT are streamlined for processing. This redesign has entailed a nearly complete rewrite of the HLT software.

3. The High Level Trigger in AthenaMT

The HLT will run a combination of the multi-process and multi-threaded approaches re-using offline algorithms and configurations. In Run 2 the multi-process approach, based on forking after initialisation (AthenaMP), was used to reduce the overall memory requirements thanks to the copy-on-write mechanism. This means the move to AthenaMT will likely be less of a reduction in processing power for the HLT compared to offline. The performance of the offline code is the main motivation for the move to multi-threading. However, the move to a common multi-threading environment will allow offline and HLT code to share algorithms and common configurations.

One of the primary principles of resource reduction is to allow for early event rejection. This means that if the event does not satisfy intermediate hypotheses that are interesting for physics, the algorithms further down the line do not have to be executed. If none of the steps are passed, then the

whole event can be rejected and resources freed up to process the next event. The algorithms that allow for early event rejection are illustrated in the data flow diagram in Figure 3. Filter algorithms can either reject event completely or pass certain Regions of Interest (ROIs). Hypothesis algorithms implement selections and are specific for each type of signature.

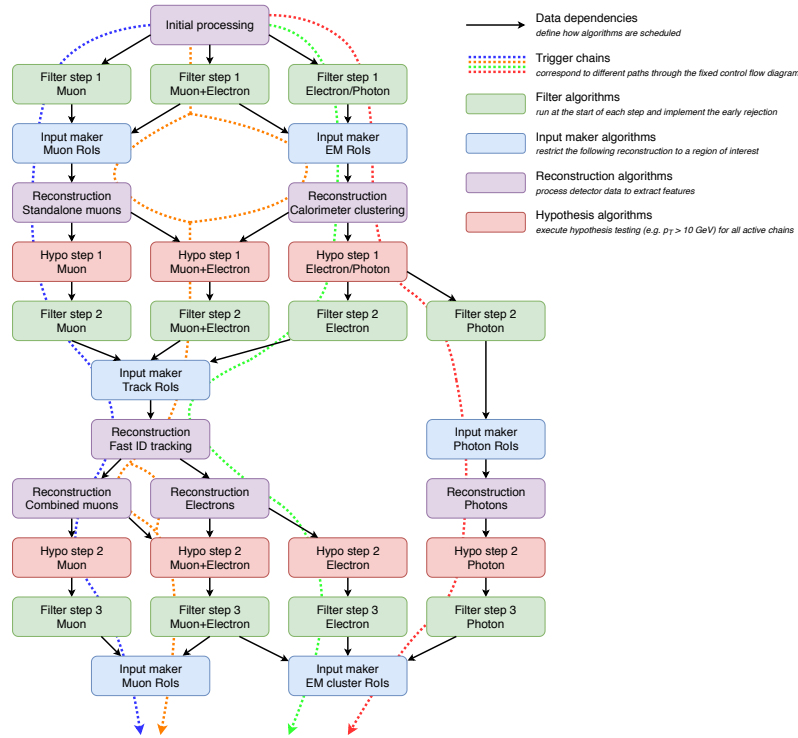


Figure 3: Data flow diagram showing electron and muon selection steps sent through HLT algorithms. [4]

4. Outlook

Work is ongoing in the migration to AthenaMT both online and offline. At the time of writing these proceedings, large scale tests are validating the new approach from the perspective of selection quality as well as achieved throughput. If successfully certified, this setup will allow ATLAS online and offline software to take advantage of the trends in computing to obtain the best processing performance possible for Run 3 of the LHC.

References

- [1] ATLAS Collaboration, JINST **3** S08003, 2008
- [2] K. Rupp, <https://www.karlrupp.net/2015/06/40-years-of-microprocessor-trend-data/>, 2015
- [3] ATLAS Collaboration, Athena Release 22.0.1, <https://doi.org/10.5281/zenodo.2641997>, 2019
- [4] R. Bielski, ATLAS Collaboration, ATL-DAQ-PROC-2019-004, 2019
- [5] M. Clemencic, D. Funke, B. Hegner, P. Mato, D. Piparo and I. Shapoval, J. Phys. Conf. Ser. 608 012021, 2015