The ATLAS Trigger Core Configuration and Execution System in Light of the ATLAS Upgrade for LHC Run 2

Abstract

During the 2013/14 shutdown of the Large Hadron Collider (LHC) the ATLAS first level trigger (L1) and the data acquisition system (DAQ) were substantially upgraded to cope with the increase in luminosity and collision multiplicity, expected to be delivered by the LHC in 2015. Upgrades were performed at both the L1 stage and the single combined subsequent high level trigger (HLT) stage that has been introduced to replace the two-tiered HLT stage used in Run 1. Because of these changes, the HLT execution framework and the trigger configuration system had to be upgraded. Also, tools and data content were adapted to the new ATLAS analysis model.

The ATLAS Trigger

Due to bandwidth and storage limitations not all proton collisions can be stored permanently for analysis. The ATLAS trigger is a two-tiered system to select the most interesting collisions with a output rate of 1kHz.

The first tier (Level 1 trigger, L1), implemented in custom hardware, selects events and identifies "Regions of Interest" (RoIs) from low granularity calorimeter and muon system data.

The second tier (high level trigger, HLT) is run on a commodity PC farm. Starting from the RoIs, events are selected by processing the full-granularity data of all subdetectors.

ATLAS Trigger and Data Acquisition System for Run 2

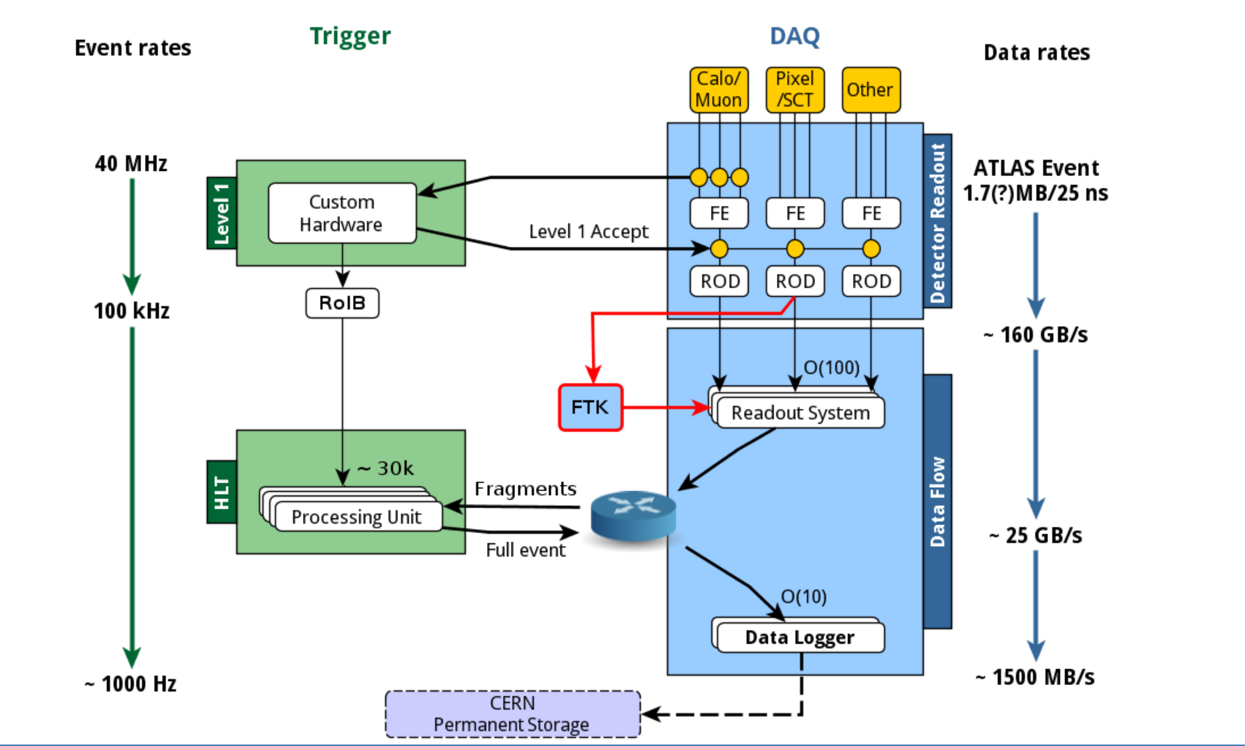

Trigger Configuration

The trigger configuration is stored in a database schema. It stores lists of trigger definitions ("menus") as well as hardware (L1) and software (HLT) configuration. A typical menu consists of several hundred trigger sequences.

The configuration functionality has been extended to support the upgrades to the L1 and the HLT shown below.

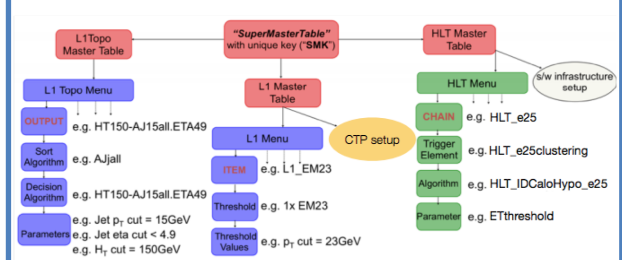

part of run 2 trigger config database schema

LHC Run 2

After the first Long Shutdown (LS1), data-taking at the LHC resumes in 2015. Trigger rates will increase due to higher energy and luminosity, requiring a more stringent and sophisticated selection. The greater number of simultaneous collisions (pile up) results in larger event data sizes and consequently higher bandwidth in the DAQ system and longer processing times in the HLT.

	Run 1	Run 2
Luminosity ($cm^{-2}s^{-1}$)	$8x10^{33}$	$1.7x10^{34}$
Energy (TeV)	8	13
bunch spacing (ns)	50	25
collisions/bunch crossing	20.7	43

Level 1 Upgrades

L1Topo: The new FPGA-based topological trigger selects events based on topological features such as invariant masses and angles. A decision can be reached within 100ns within L1 time constraints.

L1Calo: The preprocessing module has been upgraded to cope with increased simultaneous collisions and improve ability to isolate electron, photon and jet candidates.

Central Trigger Processor: The number of input channels was increased and the number of programmable L1 decisions doubled from 256 to 512.

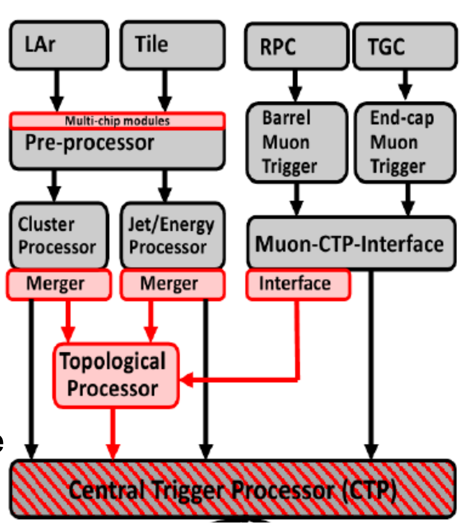

HLT Upgrades

Merged HLT: the Level-2 (L2) /Event Filter (EF) split in the HLT of Run 1 has been replaced by a more flexible merged design that allows for more efficient resource allocation and minimizes duplicate fetching of data from buffers. Also, Event Building can be scheduled statically or dynamically at any point in the HLT reprocessing.

Optimized DAQ: The network architecture was changed to a single network with 6 Tbps bandwidth to reflect L2/EF merger. Each Readout Module has 2x10Gbps links into the network to manage a 100kHz L1 rate and an increased number of HLT data requests.

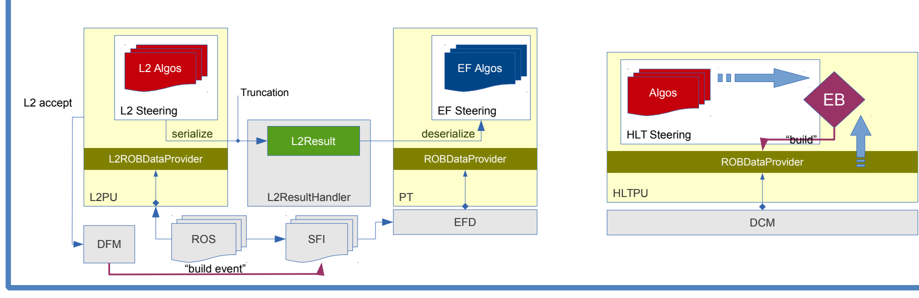

Trigger Analysis Tools

In light of the unprecedented dataset sizes of Run 2, the reconstruction data format has been upgraded to be directly ROOT-readable. It is thus a feasible format for fast end-user analysis and avoids the need for storage-intensive derivative formats such as the ones produced in Run 1.

The trigger event data model (EDM) and trigger analysis tools have been adapted to this new format. This will enable analyzers to access a wider range of trigger information than what has been available using derived data in Run1. For example, not only will the final decision be available but also the full decision-graph leading to it.

Multiprocessing in the HLT

The HLT is embarrassingly parallel but the number of independent processes that can be run on one computer is limited by their memory requirements. The large amount of conditions data in common between processes leads to the following solution that makes use of copy-on-write semantics. A single mother process prepares memory with event-independent data and forks HLT child processes. Memory pages are copied on demand so that many more processes can be run per computer.

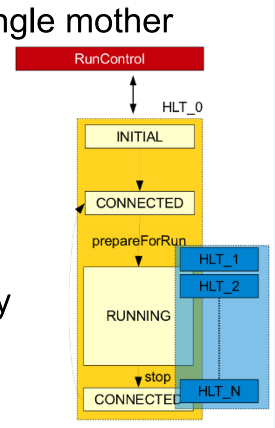

Data-Scouting Stream

The trigger data is a small fraction of the total event data but its resolution is almost comparable with full offline reconstruction.

In Run 2 an additional data-scouting stream is introduced for events that would have too high a rate to be included in the normal physics streams without heavy prescalings. For these events, only the trigger reconstruction results will be stored.

This approach provides an opportunity to allow for high-statistics samples for both calibration purposes (e.g. muons) and some searches for new physics.

Cost Monitoring

The Cost Monitoring framework tracks relevant usage statistics (rates, data requests, execution time per event etc.) to assess resource utilization by trigger algorithms. This is a crucial tool for the optimization of the trigger. It has been redesigned in Run 2 to work with the merged HLT. Data is accessible via a new web interface.

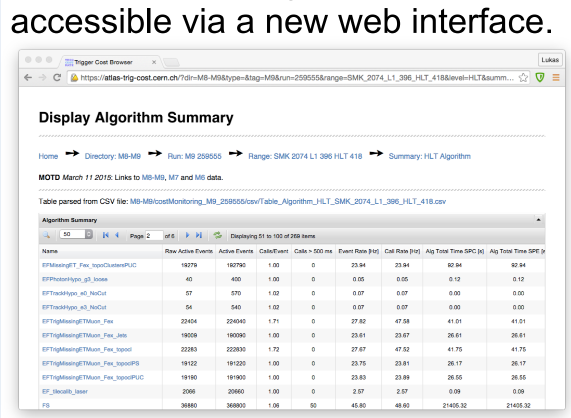

Lukas Heinrich, for the ATLAS Collaboration