Hardware Trigger Processor for the ATLAS MDT System

Thiago Costa de Paiva[1] on behalf of the ATLAS TDAQ collaboration

Many of the physics goals of ATLAS in the High Luminosity (HL) LHC era, including precision studies of the Higgs boson, require an unprescaled single muon trigger with a 20 GeV threshold. This can only be achieved by incorporating Monitored Drift Tube (MDT) information into the L0 trigger. We are developing a low-latency hardware trigger processor that will fit candidate muon tracks in the drift tubes in real time, improving significantly the transverse momentum (p_T) resolution provided by the dedicated trigger chambers. We present a novel pure-FPGA implementation of a Legendre transform segment finder, an associative-memory alternative implementation, an ARM (Zynq) processor-based track fitter, and compact ATCA blade architecture. The ATCA architecture is designed to allow a modular, staged approach in the development of the system and exploration of alternative technologies.

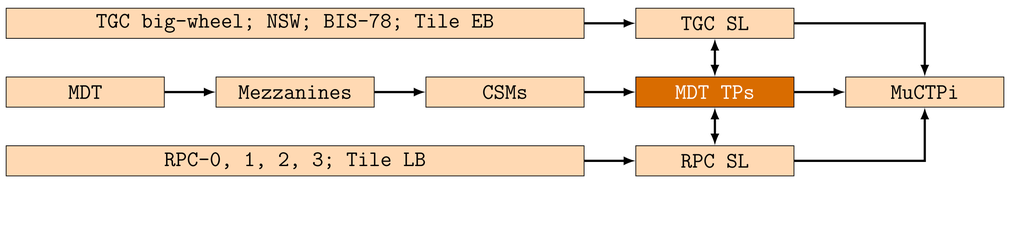

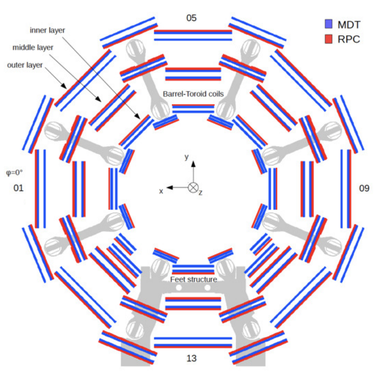

Hardware Description

Raw hits are received on three groups of GBT links from inner, middle and outer MDT stations. Region of Interest (ROI) data are received on a single fiber per sector from the muon Sector Logic (SL). Track segments found by the station processors are passed to the track fitter, which transmits the fitted track parameters to the global muon trigger. All hits are buffered independently in the DAQ buffer and are matched to time windows surrounding L0/L1 trigger accept signals with matching hits sent to the FELIX system. There are a total of 192 copies of this logic (16 sectors, A/C/Barrel/Endcap, Inner/Middle/Outer) implemented in a set of 64 ATCA blades.

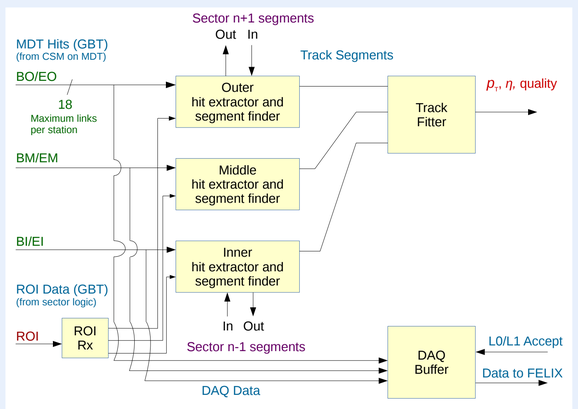

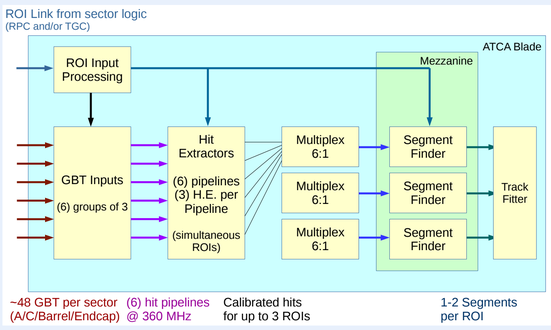

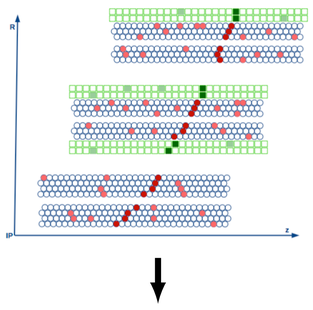

Hit Extraction

An ROI is determined in the barrel (endcap) SL from a coincidence of hits in the RPC (TGC) trigger chambers likely to be originating from a single track. These hits are used to reconstruct an ROI segment per MDT station. Matching of MDT hits to ROIs will be performed in an FPGA on the blade. Tube coordinates are transformed to convenient station-local coordinates. All MDT hits falling within a window centered on the ROI are matched by identifying the unique pair of MDT tube IDs from the innermost and outermost MDT layer with respect to the IP. These windows are different for each station. The matched MDT hits are then calibrated (drift time converted into distance) and sent to the sector processors on mezzanine boards.

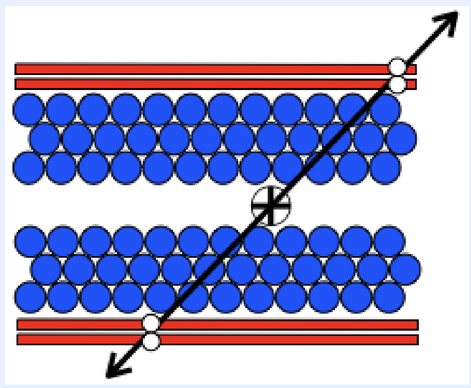

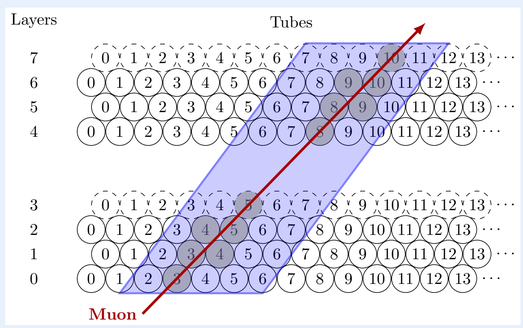

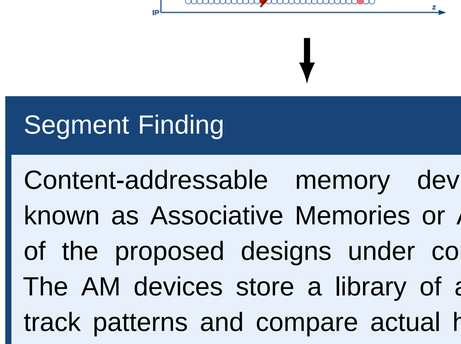

Segment Finding

Content-addressable memory devices (also known as Associative Memories or AM) is one of the proposed designs under consideration. The AM devices store a library of all possible track patterns and compare actual hits against the track patterns, producing a low-resolution segment candidate. A subsequent FPGA-based fitting stage follows.

The second approach uses FPGA logic to implement a Legendre transform based segment finder. This logic evaluates in parallel a total of 128 possible track segment angles for each MDT hit, calculating in a fast FPGA pipeline the offset of each track candidate from an arbitrary origin for each angle. The (angle, offset) pairs are used to fill a 2D histogram, with the maximum peak in the histogram representing a likely track where a number of tubes ”agree” on the position and angle. As part of the filling process, the 128 highest-occupancy bin locations are maintained, so finding the overall histogram maximum requires only a few clock cycles.

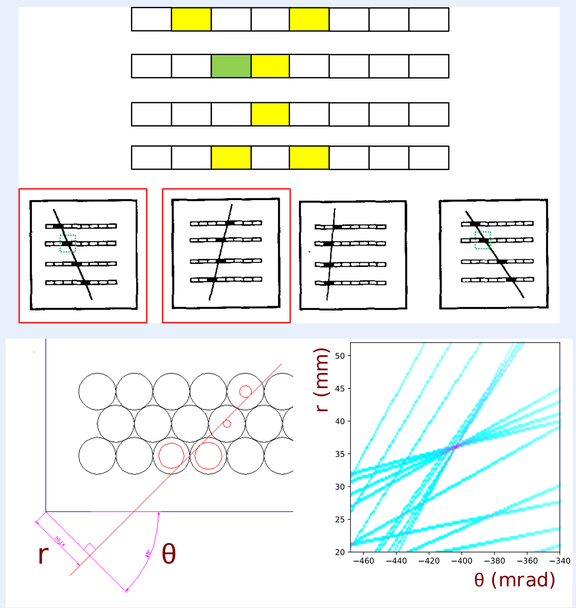

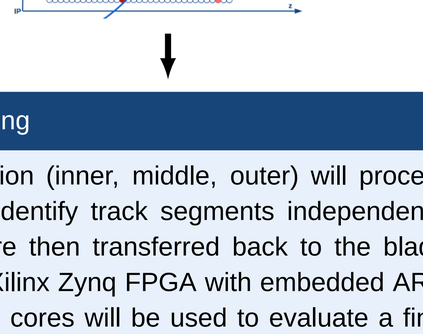

Track Fitting

Each station (inner, middle, outer) will process hits and identify track segments independently. All hits are then transferred back to the blade, where a Xilinx Zynq FPGA with embedded ARM processor cores will be used to evaluate a final parameterized track fit.

Depending on how many segments can be reconstructed per muon candidate in the different MDT stations, muon p_T can be determined using two different methods.

- If three segments are found, each in a different MDT station, the positions of these can be combined to measure the track curvature by calculating the sagitta out of the three points (3-station method).
- Otherwise, two segments in different MDT stations still can be combined to extract the p_T by measuring their deflection angle (2-station method).

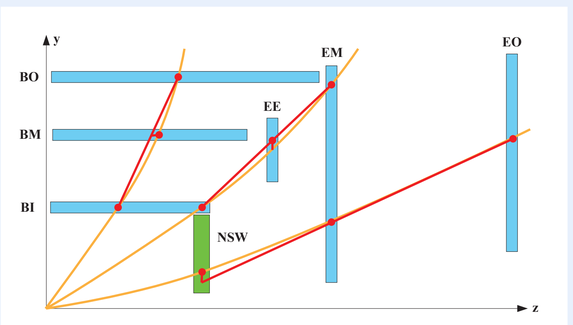

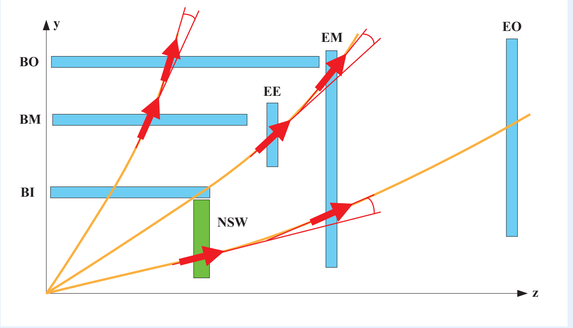

Hardware Implementation

The ATCA blade provides basic services including module management, power conditioning, base Ethernet and firmware management. The blade also contains 72 optical receivers and 72 optical transmitters, capable of operation up to 10 Gbps. The blade contains one Xilinx Ultrascale-class FPGA which will handle the reception of the MDT data via GBT links, ROI information, hit extraction and calibration, and transmission of hits to the mezzanine board(s). Segment data is transferred back from the mezzanine board to the blade for track fitting. In addition, the blade FPGA will transfer the MDT hits to the ATLAS DAQ via FELIX. An external DDR memory device may be required for buffering of DAQ data.

A Xilinx Zynq ^{TM} FPGA/CPU combination chip provides an Ethernet interface and might also be used to implement certain track fitting algorithms. The Zynq device requires RAM for its operating system as well as an interface to a μSD card or other flash filesystem storage.

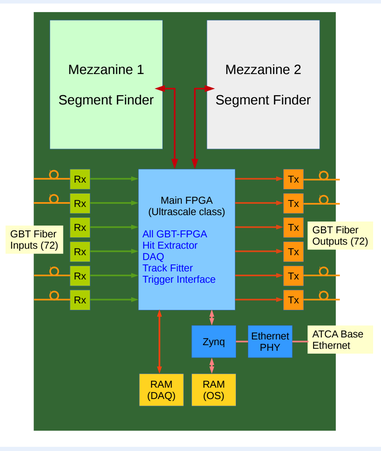

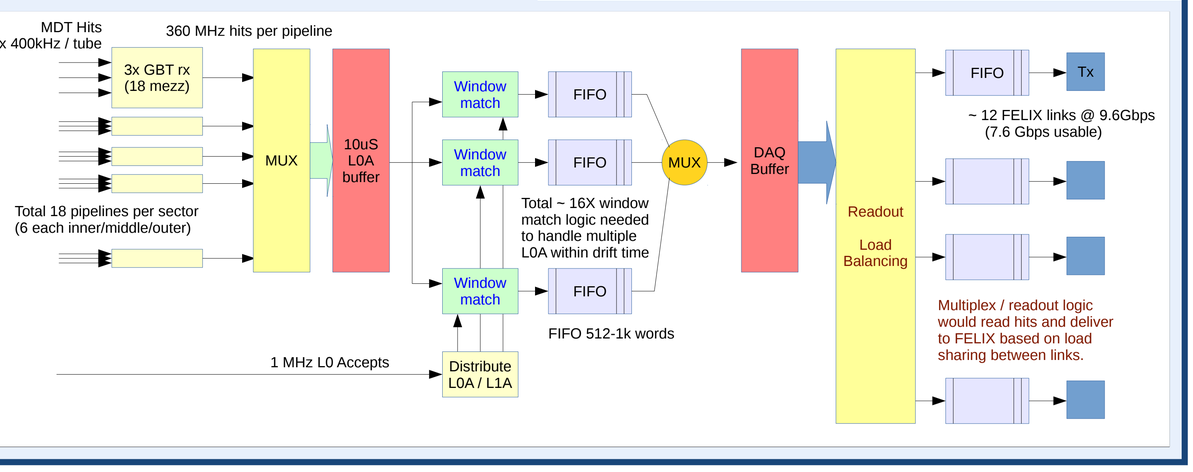

MC framework and Latency Studies

Monte Carlo (MC) simulated event samples are used to aid in the development of the algorithms and measure the expected trigger rates and efficiencies under the HL-LHC luminosity and pileup conditions.

In this project, MC events are also used as input to a cycle-accurate hardware simulation for the digital processes up to the segment finding mezzanines. The simulation is necessary to understand the behavior of the MDT readout/trigger chain since events are random and the processing in the front-end electronics adds uncertainties to the data delivery.

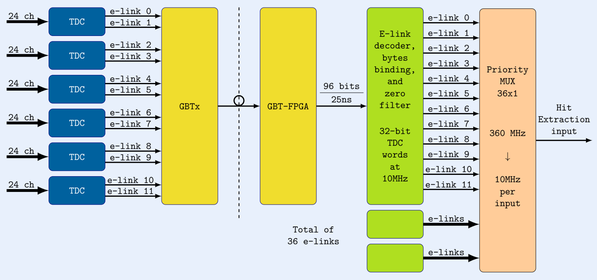

Conclusion and Future Work

Studies shown that the hardware trigger processor will be able to improve the purity of the triggers and sharpen the p_T turn-on to meet the rate limit for single muon triggers for the HL-LHC.

Our best estimate for the total latency for the L0-MDT trigger processor (once all the data has arrived to the board) is less than 1μs to process up to 100 MDT hits within an ROI and provide a high-resolution p_T measurement as output.

A detailed conceptual design with extensive simulation studies is being prepared now, to be published in the ATLAS TDAQ Technical Design Report. A first generation of hardware prototypes is planned for 2018-2019, with a full system ready for installation in approximately 2024.

Poster presented at TWEPP 2017 – Santa Cruz, CA, US.

[1] tcpaiva@cern.ch