

Track Trigger at the High Luminosity LHC

Marco Trovato^{*†}

Northwestern University

E-mail: marco.trovato@cern.ch

During the High Luminosity LHC, both ATLAS and CMS detectors will need extremely fast charged-particle tracking at the trigger level to maintain manageable trigger rates and achieve their physics goals. An overview of the upgrades planned for the tracker sub-detectors and trigger systems is given, along with a description of the algorithms under study. Preliminary results from smaller scale demonstrator tests are also shown.

Sixth Annual Conference on Large Hadron Collider Physics (LHCP2018)

4-9 June 2018

Bologna, Italy

^{*}Speaker.

[†]Thanks a lot to Francesco Crescioli for the useful discussion.

1. Introduction

An important component of the High Luminosity (HL)-LHC programme is to take a closer look at the recently discovered Higgs boson. Precision measurements of its properties is a way to test the Standard Model pattern of couplings to elementary particles. Another goal of the HL-LHC programme is to search for new physics, which may provide answers to the still open questions in particle physics (e.g: what is the nature of the dark matter we observe in the Universe).

CERN's LHC accelerator complex will undergo major upgrades, planned for 2025, to increase the instantaneous luminosity up to around $7.5 \times 10^{34} \text{ cm}^{-2}\text{s}^{-1}$, which is needed to collect about 3000 fb^{-1} of data in about 10 years of operation. Such a data will amount to about 10 times the data which will be collected before the upgrade. HL-LHC will yield an average number of overlapping proton-proton collisions per bunch crossing (pileup or PU) between 140 and 200. At a 40 MHz collision rate, that would mean roughly 20 to 40 Tbps would need to be stored. However, only a small fraction of that data is of interest for further study. Therefore, both ATLAS and CMS physics programs for the HL-LHC require intelligent triggering scheme to be able to isolate collisions containing interesting physics scenarios. The trigger system would need to reduce the 40 MHz collision rate to a manageable data storage rate of about 7.5 kHz.

A major role in reducing the rate is played by the hardware trigger system. At the HL-LHC, the trigger rates for objects such as single muons, electrons, or jets will exceed the current front-end capabilities. Increasing the transverse momentum (p_T) trigger thresholds for these objects would negatively impact the physics potential of the HL-LHC. Integrating charged particle tracking from the silicon tracker into the hardware trigger will improve lepton identification and momentum measurements, and provide track isolation and vertex identification. These additional handles have the potential to reduce the rates, while maintaining reasonable trigger thresholds.

ATLAS and CMS have envisioned different strategies in how to embed a hardware tracking in the trigger data acquisition scheme. In the next sections ATLAS (Sec. 2) and CMS (Sec. 3) approaches to tackle the problem will be described.

2. ATLAS Track Trigger

ATLAS trigger will consist of a hardware trigger, followed by a software trigger ("Event Filter") running on a computer farm. The hardware trigger will operate in a single-level ("L0") mode.

The hardware track trigger, named "EFTrack", will act as co-processor for the Event Filter, as shown in Fig. 1, and will operate at 1 MHz Level-0 Accept accept rate. EFTrack is composed of two steps. The first step consists of tracking in Regions of Interest (RoIs), which are defined by the L0 muon and electromagnetic calorimeter triggers. RoIs cover approximately 10% of the detector volume. Regional tracking operates on track candidates of $p_T > 2 \text{ GeV}$. Reconstructed tracks and L1 information are used by the Event Filter to reduce the rate down to 400 MHz. Upon request, full tracker data can be used by EFTrack in the second step. Track reconstructed via global tracking in combination with offline-like quantities are exploited by the Event Filter to further reduce the rate to 10 KHz.

The current ATLAS tracker will be completely replaced during the upgrade. The new ITk will be all-silicon and consist of pixel detectors close to the beam pipe and strip detectors at larger radii.

Although not decided yet, the final layout will likely have a barrel with 5 pixel layers and 4 strip stereo-layers, and an endcap with coverage up to a pseudorapidity (η) of 4.

The EFTrack trigger will provide fast regional and global hardware track triggering. The system conceptually consists of three parts: pattern matching using Associative Memory (AM) chips [1], track fitting of the full-resolution hits using Field Programmable Gate Arrays (FPGAs), and selection of the best track candidates based on track parameters and the χ^2 of the fit. Upon a readout request, the inner tracking detector clusters within the RoI from the outermost 8 layers¹ would be passed to the EFTrack via the detector electronics. Clusters are initially processed using only coarse-granularity superstrip information from the inner detector: “superstrip”. Superstrip information is passed through the AM chips to check for matches against pre-loaded patterns, which consist of one superstrip per detector layer. Patterns are pre-computed using simulated single muons with realistic conditions. To improve coverage, and account for geometrical inefficiencies, some patterns use wildcards which allow certain layers to not contain a hit.

The next stage of the process is to apply a track fitting procedure to the full granularity hits for patterns which have been matched. This stage selects tracks for a global decision and is performed within FPGAs. The track fitting is performed by applying a set of fit constants to relate the cluster positions to the 5 track parameters. A linear approximation, whose constants are produced by a principal component analysis on a large sample of simulated muons, is exploited. Candidate tracks are accepted or rejected based upon the value of the computed χ^2 .

In the case of global tracking the information available in the innermost layers is also used: tracks and clusters are extrapolated and re-fitted in the FPGA using full hit information.

Efficiencies, as predicted in simulation for muons with $p_T > 10$ GeV, are close to 100%, and the resolution after global tracking is comparable to the offline one.

EFTrack system builds upon the FTK experience [2]. Many independent units handle different $\eta - \phi$ regions of the detector. The building blocks of each unit or “tracking processor (TP)” are ATCA boards located in ATCA shelves with full mesh backplane. Two types of building block units are envisioned:

- Associative Memory TP (“AMTP”): a master Virtex-7 FPGA plus two pattern recognition mezzanines (“PRMs”) on it. Each PRM features 12 Associative memories (“AM06”, [1]) and a KU060 FPGA. AM06 is a 65 nm dedicated ASIC with 120K pattern and runs at 100 MHz. The final targeted Associative memory is the AM09, which is a 28 nm ASIC with 384K patterns running at 250 MHz.
- Second-Stage TP (“SSTP”): two track fitting mezzanines on it with two Kintex Ultrascale FPGA on it. The design is still under discussion.

Each TP is composed of 6 AMTP and 1 SSTP. Each TP communicates with the Event Filter through a dedicated interface.

Further details on the ATLAS track trigger are available at [3]

¹two-sided strips count as two layers

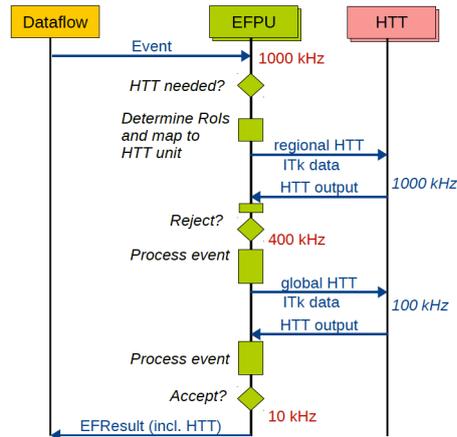


Figure 1: Diagram showing the interaction between the Event Filter processing unit (EFPU), Dataflow system and the hardware tracking (HTT). After receiving the event data from the Dataflow system, the EFPU decides, based on L0, if regional hardware tracking is needed for this event. If so, requests are sent to the relevant HTT units along with the necessary ITk data. After receiving the HTT result the EFPU continues with the event processing including a possible second request for global hardware tracking based on the L0 decision or additional EF-based selections.

3. CMS Track Trigger

Hardware tracking in CMS is part of the first-level trigger (“L1”) of the upgraded system. L1 trigger decision must occur within $12 \mu\text{s}$. Track trigger is expected to receive data from the tracker front-ends at 40 MHz, and must reconstruct track in approximately $4 \mu\text{s}$, in order to allow for correlating tracks with other physics objects and achieving a trigger decision.

The track trigger system is part of the upgraded CMS tracker, which is entirely made of silicon. The new tracker consists of a pixel detector (not part of the trigger decision) and an outer tracker that has a central barrel and two endcaps. The outer tracker will have the unique feature of “ p_T modules” that provide p_T discrimination at the level of the front-end readout electronics. Requiring $p_T > 2 \text{ GeV}$ tracks provides a factor of $\sim 10\%$ data reduction.

There are two possible implementations of L1 tracking in CMS, all relying on commercial FPGAs:

- the “tracklet” approach, a road-search algorithm (Sec 3.1);
- a Hough-transform based approach.

In this document, we will focus on the tracklet approach. Details on the other Hough-transform approach can be found elsewhere [4].

3.1 Tracklet Approach

The tracklet approach starts by forming track seeds (“tracklets”) from pairs of stubs in adjacent barrel layers or endcap disks. The tracklet is an initial estimate of the tracklet parameters, which are calculated from these two stubs and the interaction point.

The track parameters of the tracklets are then projected to other layers and disks to search for consistent stubs. A track candidate is found if at least four stubs have been matched. The tracking efficiency, as expected from an integer-based C++ emulation of the algorithm, is very close to 100% (Fig. 2).

A linearized χ^2 fit is performed using all stubs of the track candidate. The track fit uses pre-calculated derivatives and the projection-stub differences. The linearized χ^2 fit improves the initial tracklet parameters giving the final track parameters: p_T , η , the longitudinal distance to the detector origin (z_0) and the azimuthal angle (ϕ_0) at the closest approach, and optionally the impact parameter of the transverse plane (d_0). Because seeding is performed for different layers/disks configurations, a single track may be found several times: duplicated tracks are removed by comparing the number of independent and shared stubs.

To address the challenging amount of data and limited processing time available, the tracklet hardware configuration massively parallelizes the data processing. The main parallelization is the division of the detector into sectors in the $r - \phi$ plane. The current project uses 28 ϕ sectors. To allow for more time for data processing, the whole system is time multiplexed by a factor of 6. This choice is driven by a balance of cost, efficiency and needed processing power.

The tracklet algorithm is implemented in the firmware as several processing steps. Currently all processing steps are synchronized to a single master 240 MHz clock. By construction, the system is fully pipelined and operates at a fixed latency: when a new event arrives the previous event moves to the next processing step. Because of that, a given step can only perform a fixed number of operations. If the time limit is reached, processing on the remainder of the data is truncated. The effect of truncation on the system is minimal (Fig. 2).

The full tracklet algorithm, including all processing and transmission steps, has been implemented in firmware. Two complete implementations, one for half the barrel ($+z$) and one for a quarter of the barrel plus the forward endcaps, were used to demonstrate the feasibility of this approach for the full η range of the detector. In the final system, each sector processor will be an ATCA blade with a Virtex Ultrascale+ FPGA. Conversely, the current demonstrator system is made of four μ TCA blades, named “CTP7” boards [5]. Each CTP7 has a Xilinx Virtex-7 FPGA and a Xilinx Zynq-7000 SoC processor for configuration and outside communication. The CTP7 boards were developed for the current CMS L1 trigger. An AMC13 [6] card provides synchronization between the boards with a central clock distribution. The inter-board communication uses 8b/10b encoding with 10 Gbps links. The demonstrator system is shown in Fig. 2.

The demonstration at the end of 2016 was successful: the total measured latency is about 3.3 μ s, well within the latency budget of 4.0 μ s.

Further details on the CMS track trigger are available elsewhere [7].

4. Summary and Outlook

For the HL-LHC both CMS and ATLAS tracker sub-detectors will be upgraded to all-silicon systems. In CMS the tracker will feature new trigger capabilities. For both experiments the trigger chain will be redesigned to cope with the hard conditions during the HL-LHC era and to be able to achieve the challenging physics goals. Tracking in hardware will be the key of the new triggering

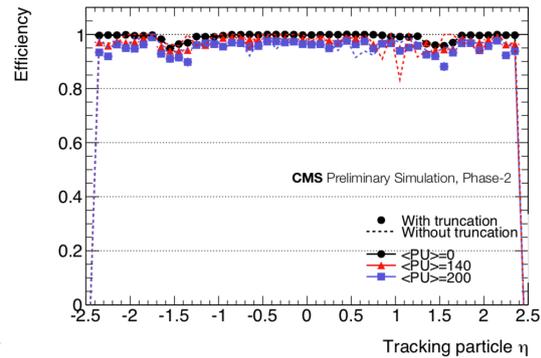


Figure 2: Image (left) of the tracklet demonstrator system: CPT7 boards are placed into μ -TCA crates. Efficiency as a function of η for muons (right). Results are shown for three different average pileup scenarios: 0 (black), 140 (red), and 200 (blue). Efficiencies are shown with (filled points) and without (dashed lines) truncation included.

systems and will provide greatly improved momentum measurements, thus enabling rate reducing triggers without increasing p_T thresholds. In ATLAS, the hardware track trigger will assist the software-based trigger, after the first-level trigger. The hardware tracking, mostly performed in regions of interest, is composed of associative memories for the pattern recognition steps, and commercial FPGAs for the track fitting and duplicate removal steps. In CMS, the hardware track trigger will operate at the first-level trigger. Two approaches based on commercial FPGAs are being investigated. The first results from preliminary smaller scale systems demonstrate that tracking particles at L1 is feasible even with today's technology.

Both experiments keep on studying the feasibility of embedding hardware tracking in the trigger chain. Full-blown demonstrators are expected to happen between 2019 and 2021.

References

- [1] A. Annovi, M. M. Beretta, G. Calderini, F. Crescioli, L. Frontini, V. Liberali, S.R. Shojaii and A. Stabile, JINST 12, C04013 (2017)
- [2] ATLAS Collaboration, CERN-LHCC-2013-007, ATLAS-TDR-021, Geneva (2013)
<https://cds.cern.ch/record/1552953>
- [3] ATLAS Collaboration, CERN-LHCC-2015-020, LHCC-G-166, Geneva (2015)
<http://cds.cern.ch/record/2055248>
- [4] R. Aggleton *et al.*, JINST 12, P12019 (2017)
- [5] A. Svetek *et al.*, JINST 11, C02011 (2016)
- [6] E. Hazen, A. Heister, C. Hill, J. Rohlf, S.X. Wu and D. Zou, JINST 8, C12036 (2013)
- [7] J. Chaves, JINST 9, C10038 (2014)