Available on CMS information server

CMS CR -2023/305

The Compact Muon Solenoid Experiment

Conference Report

Mailing address: CMS CERN, CH-1211 GENEVA 23, Switzerland

06 December 2023 (v4, 08 December 2023)

Level-1 Tracking at CMS for the HL-LHC

Sara Fiorendi for the CMS Collaboration

Abstract

The success of the CMS physics program at the High-Lumi LHC requires maintaining sufficiently low trigger thresholds to select processes at the electroweak scale. With an average expected 200 pileup interactions, the inclusion of tracking information into the Level-1 trigger will be critical to achieve this goal while maintaining manageable trigger rates. The main challenges of reconstructing tracks in the Level-1 trigger are the large data throughput at 40 MHz and the need for a trigger decision within 12.5 μs, out of which 4 μs is for track finding. The CMS outer tracker for HL-LHC uses modules with closely-spaced silicon sensors to read out only hits compatible with charged particles with transverse momentum above 2 GeV ("stubs"). These are used in the back-end L1 track-finding system, based on commercially available FPGA technology, to form tracks. The algorithm uses stubs in adjacent layers to form track seeds and performs a final track fit based on a Kalman filter. In this contribution the status of the CMS Level-1 track finding and its implementation are discussed. Simulation studies of the estimated performance are shown and the developments of hardware demonstrators described.

Presented at *CTD2023 8th International Connecting the Dots Workshop*

Proceedings of the CTD 2023
PROC-CTD2023-XX
December 8, 2023

Level-1 Tracking at CMS for the HL-LHC

Sara Fiorendi

On behalf of the CMS Collaboration,
University of Tennessee, Knoxville

ABSTRACT

The success of the CMS physics program at the High-Lumi LHC requires maintaining sufficiently low trigger thresholds to select processes at the electroweak scale. With an average expected 200 pileup interactions, the inclusion of tracking information into the Level-1 trigger will be critical to achieve this goal while maintaining manageable trigger rates. The main challenges of reconstructing tracks in the Level-1 trigger are the large data throughput at 40 MHz and the need for a trigger decision within 12.5 μs, out of which 4 μs is for track finding. The CMS outer tracker for HL-LHC uses modules with closely-spaced silicon sensors to read out only hits compatible with charged particles with transverse momentum above 2 GeV ("stubs"). These are used in the back-end L1 track-finding system, based on commercially available FPGA technology, to form tracks. The algorithm uses stubs in adjacent layers to form track seeds and performs a final track fit based on a Kalman filter. In this contribution the status of the CMS Level-1 track finding and its implementation are discussed. Simulation studies of the estimated performance are shown and the developments of hardware demonstrators described.

PRESENTED AT

Connecting the Dots Workshop (CTD 2023)
October 10-13, 2023

1 Introduction

The Large Hadron Collider (LHC) will be updated to the High-Luminosity LHC (HL-LHC) [1] expected to start operation in 2029. It will provide the experiments with unprecedented amounts of data. These data will allow to extend the discovery capacity in the searches for new physics and for rare standard model (SM) processes, as well as to improve Higgs boson and SM precision measurements. On the other hand, the running conditions will be extremely challenging: with an instantaneous luminosity that will increase up to $7.5 \times 10^{34}\,\mathrm{cm}^{-2}\mathrm{s}^{-1}$, the expected average number of overlapping proton collisions (pileup) per event will reach 200, with a corresponding increase in particle density. The Phase-II upgrades of the CMS detector [2] have been designed in order to maintain its excellent detection abilities, and even to improve its performance with respect to the current one.

The inclusion of tracker information in the Level-1 (L1) trigger plays a crucial role to pursue this goal. It will improve the transverse momentum resolution and particle identification performance. Tracking information at L1 will also make possible to reconstruct the primary interaction vertex to help with pileup mitigation and will allow to carry out object-vertex association. Furthermore, the presence of reconstructed tracks at L1 will extend the physics reach of the CMS experiment by opening the possibility to trigger on physics processes that previously could not be selected in the first trigger stage, like, for instance, B-physics processes with full hadronic final states, or exotic long-lived particles with decay topologies including displaced trajectories.

2 The CMS trigger and tracker upgrades

The CMS experiment uses a two-stage trigger system [3] to identify and select the events of interest: a hardware-based Level-1 trigger stage is followed by the High Level Trigger based on commercial computers. While maintaining this two-level structure, the entire CMS trigger system will be replaced [4, 5] for the HL-LHC in view of the larger amount of data to be processed. The L1 trigger will have a latency of $12.5\,\mu$s and a maximum output rate of 750 kHz, to be compared with the current $4\,\mu$s and 100 kHz, respectively. Figure 1 shows a schematic overview of the upgraded L1 trigger architecture. The global track trigger will reconstruct primary vertices and track-only objects. L1 track information will also be integrated into muon and calorimetric objects by a two-stage correlator trigger, that will perform a simplified Particle Flow event reconstruction. The final decision to store or discard an event will be taken by the global trigger, based on the objects provided by previous stages.

In view of the HL-LHC operation also the CMS tracking detector will be replaced [6] to cope with the increased luminosity. The new detector will feature enhanced granularity and pseudo-rapidity coverage, and higher radiation tolerance. The inner region will be instrumented with thin silicon pixel sensors, arranged in four cylindrical layers in the barrel region, and eight small plus four large disk-like structures in the forward direction on both sides of the interaction point, extending the acceptance up to $|\eta| \approx 4$. The Outer Tracker (OT) will consists of silicon "$\mathrm{p_T}$ modules" arranged in six cylindrical barrel layers and five endcap double disks. These modules are capable of performing a local measurement of the track transverse momentum ($\mathrm{p_T}$) by correlating charged particle hit positions in their top and bottom sensors, and reject signals from tracks below a predefined threshold. Two kinds of $\mathrm{p_T}$ modules will be employed: the pixel-strip (PS) module has one sensor with macro-pixels (1.446 mm long, $100\,\mu$m pitch), and the other with strips (2.4 cm long, $100\,\mu$m pitch); the strip-strip (2S) module, in which both sensors feature strips (5 cm long, $90\,\mu$m pitch). In both cases the data from the two silicon sensors in the module is finally processed by a single ASIC to perform the correlation of the hit positions in the two sensors and to form short track segments (*stubs*) for clusters consistent with a $\mathrm{p_T} > 2\,\mathrm{GeV}$, which are used as input for the L1 track finding algorithm.

The geometry of the detector is shown in Fig. 2, where the two types of modules are highlighted in blue (PS) and red (2S). The modules in the forward region of the three inner layers are tilted such that charged particles from the interaction point will traverse the modules in a direction approximately perpendicular to the sensor plane, thus increasing the reconstruction efficiency and reducing the number of modules needed to provide hermetic coverage.

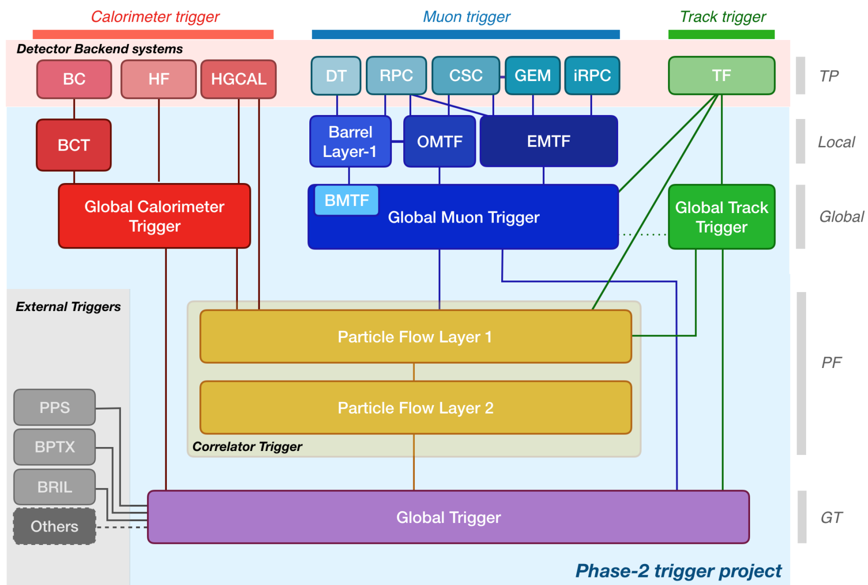

Figure 1: Diagram of the upgraded L1 trigger design [4]: the trigger receives inputs from the calorimeters, the muon spectrometers and the track finder (TF). In particular, the TF provides tracks to various parts of the trigger system, including the global track trigger and the correlator trigger, which is dedicated to Particle Flow (PF) reconstruction. All objects are sent to the global trigger issuing the final L1 trigger decision. The various levels of processing are indicated on the right: trigger primitives (TP), local and global trigger reconstruction, PF trigger reconstruction and global decision.

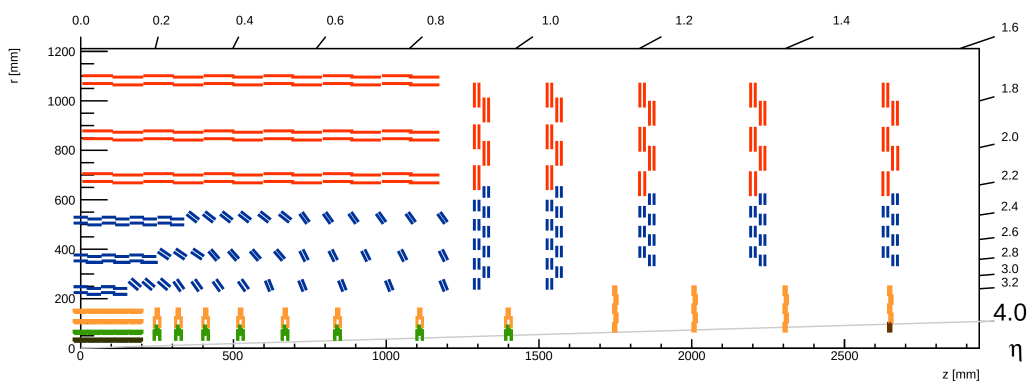

Figure 2: Sketch in the $r - z$ plane of one quarter of the upgraded tracker for HL-LHC, showing the Outer Tracker (in blue and red), as well as the Inner Tracker detector (in black, green, and yellow).

3 Level-1 track finding algorithm

The L1 track finding is implemented on commercially available field-programmable gate arrays (FPGAs) and extensively exploits parallel processing in order to deal with the high data rate and large combinatorics. The detector is split in nine sectors taking advantage of its natural segmentation in ϕ, and further parallel processing is envisaged also within a given ϕ sector. To increase the time available for data processing, the system is also time multiplexed by organizing the data such that eighteen identical copies of the track finding processor are run in parallel.

The track finding sequence consists in a road search algorithm which performs the pattern recognition, followed by a Kalman filter for the final track fit. Pairs of stubs from adjacent layers or disks of the tracker form seeds (called *tracklets*); eight possible combinations of layers and disks are currently considered in the seed generation step. The seeding is performed in parallel and in multiple combinations of layers to ensure coverage and redundancy. Once a seed is created, an initial track trajectory is estimated based on the tracklet information and imposing the track provenance from the beam spot (in the transverse plane). Projections to other layers are then computed, and compatible stubs are associated to the tracklet if they are found within predefined narrow windows. The removal of duplicate tracks happens next, where tracks sharing a certain number of stubs are merged. Most duplicates are generated by the same particle trajectory being reconstructed starting from different seed combinations. The remainder are caused by multiple stub combinations being associated with a single particle, seeded by the same combination of layers. Finally, the candidate track is fit with a Kalman Filter algorithm [7]: starting from an initial estimate of the track parameters and their uncertainties, it iteratively updates their values adding the information from the stub in the next layer (disk). The process continues until a minimum number of stubs has been added to the trajectory.

The expected tracking performance has been estimated using simulated events with $t\bar{t}$ decays in an environment with 200 pileup interactions. Examples of the expected L1 tracking performance are illustrated in Fig. 3: the left plot shows the expected L1 track reconstruction efficiency as a function of pseudorapidity for particles with $\mathrm{p_T}$ larger than 2 and 8 GeV. The efficiency is larger than 90% for both cases across the full η range. The right plot displays the resolution along the direction of the beam line ($\mathrm{z_0}$) as a function of pseudorapidity, for intervals that contain 68% and 90% of all tracks with $\mathrm{p_T} > 2$ GeV. The $\mathrm{z_0}$ resolution is about 1 mm in the barrel ($|\eta| < 0.7$) and degrades towards the endcap region.

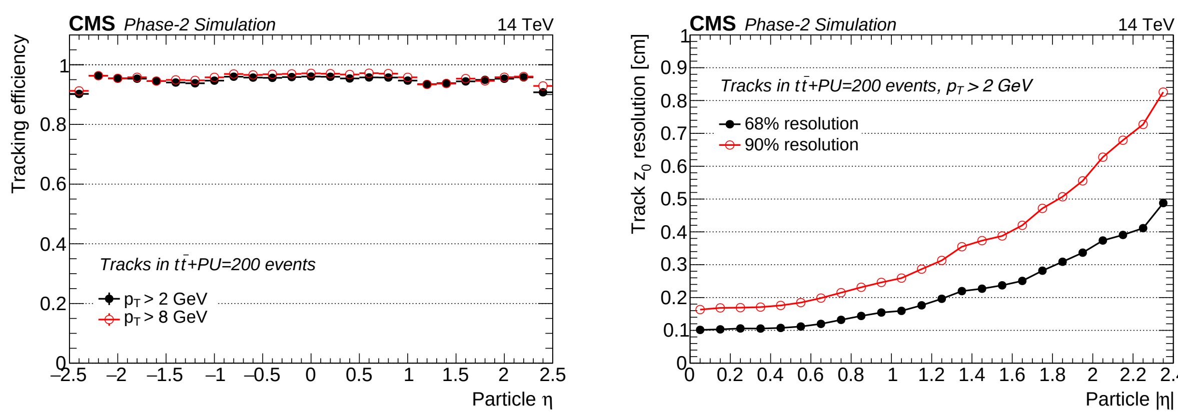

Figure 3: Left: L1 tracking efficiency as a function of η for tracks above 2 GeV (filled black markers) or 8 GeV (open red markers) for tracks in $t\bar{t}$ events overlaid with an average pileup of 200. The efficiency is defined with respect to truth-level particles that produced stubs in at least four layers/disks. Right: Track $\mathrm{z_0}$ resolution as a function of $|\eta|$. The $\mathrm{z_0}$ resolutions correspond to intervals that encompass 68% (filled black markers) or 90% (open red markers) of all tracks with $\mathrm{p_T} > 2$ GeV [4]

.

An additional track quality module will be run after the Kalman Filter step in order to reduce the number of spurious tracks, which could potentially have a negative impact on the L1 trigger decision, for instance invalidating the missing momentum calculation. Three alternatives have been explored for this purpose: a simple cut-based selection considering a minimum number of stubs, the track minimum $\mathrm{p_T}$ and normalised χ^2, and an indicator of the consistency between the track $\mathrm{p_T}$ and the bend measured in the individual stubs; a gradient boosted decision tree (GBDT) and a neural network (NN), which both use three track coordinates (ϕ, η and $\mathrm{z_0}$), two variables describing the detector hits associated with each track, and three goodness-of-fit variables as input features. It has been demonstrated [8] that the two machine learning approaches have comparable performance and both outperform the cut-based selection, as illustrated in Fig. 4; given the lower FPGA resource usage of the GBDT when compared to the NN, this approach is adopted for the track quality estimation.

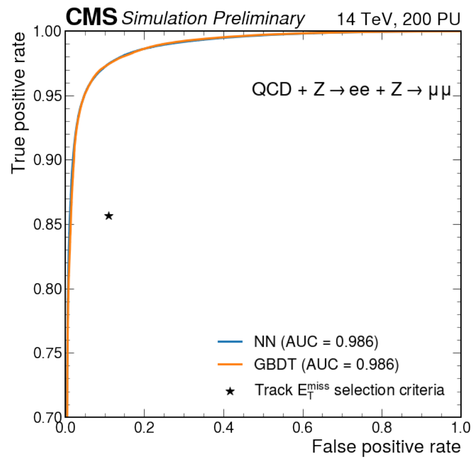

Figure 4: Comparison of the performance of the neural network (NN, blue line), gradient boosted decision tree (GBDT, red line), and the cut-based (indicated with a star) L1 track selection criteria [8]. The true positive rate measures the fraction of real tracks that were correctly identified, and the false positive rate measures the fraction of fake tracks that were incorrectly identified. AUC is a measure of the model accuracy.

4 Hardware platforms and firmware implementation

Custom carrier boards for specific FPGAs were designed for different parts of the trigger chain. The Apollo platform [10] will be used for the track finder processing boards.

The tracklet algorithm is implemented in firmware as alternated processing steps and memory modules. The seeding and propagation steps of the algorithm are being implemented with a High Level Synthesis (HLS) language developed by Xilinx, while the memory modules and the Kalman Filter steps are written in standard Hardware Description Language (HDL). The algorithm implementation is targeting a 250 MHz clock frequency.

An end-to-end demonstration of the track finding chain has been conducted on a narrow ϕ slice and its output has been compared to software emulation. This implementation, in which only one barrel seed combination was considered, successfully met the timing requirements. A full barrel project (which represents $\approx 2/3$ of the full project) has also been tested on a single Virtex Ultrascale 13 Plus (VU13P) FPGA, using the aggregation of the processing steps that will be used in the final design. The system showed high congestion and required a careful floorplanning in order to meet timing requirements. For the final project the algorithm will be implemented on two VU13P FPGAs.

5 Conclusions

The upgrades foreseen for the CMS experiment for the HL-LHC will offer the opportunity to extend the current physics reach despite the challenging pileup conditions. The future Outer Tracker will be instrumented with p_T modules able to perform on-detector data filtering, enabling the possibility of performing track reconstruction at the first trigger stage. The track finding algorithm that will be executed consists of a tracklet pattern recognition step, followed by a Kalman Filter to estimate the track parameters and a track-quality evaluation module to improve the purity of the reconstructed tracks. The full algorithm chain needs to run within a latency of $4\,\mu$s. A reduced version of the system has been demonstrated in hardware, and further tests are ongoing towards running the complete chain for the full detector.

ACKNOWLEDGEMENTS

This material is based upon work supported by the U.S. Department of Energy, Office of Science, Office of Basic Energy Sciences, Energy Frontier Research Centers program under Award Number DE-SC0020267.

References

[1] G. Apollinari, I. Bjar Alonso, O. Brning, P. Fessia, M. Lamont, L. Rossi and L. Tavian, "High-Luminosity Large Hadron Collider (HL-LHC): Technical Design Report V. 0.1," CERN Yellow Report CERN 2017-007-M [doi:10.23731/CYRM-2017-004].

[2] CMS Collaboration, "Technical Proposal for the Phase-II Upgrade of the CMS Detector," CERN-LHCC-2015-010 [https://cds.cern.ch/record/2020886].

[3] CMS Collaboration, "CMS: The TriDAS project. Technical design report, Vol. 2: Data acquisition and high-level trigger," CERN-LHCC-2002-026 [https://cds.cern.ch/record/578006]

[4] CMS Collaboration, "The Phase-2 Upgrade of the CMS Level-1 Trigger," CERN-LHCC-2020-004 [https://cds.cern.ch/record/2714892]

[5] CMS Collaboration, "The Phase-2 Upgrade of the CMS Data Acquisition and High Level Trigger," CERN-LHCC-2021-007 [https://cds.cern.ch/record/2759072]

[6] CMS Collaboration, "The Phase-2 Upgrade of the CMS Tracker," CERN-LHCC-2017-009 [doi:10.17181/CERN.QZ28.FLHW]

[7] R. Frühwirth, "Application of Kalman filtering to track and vertex fitting," Nuclear Instruments and Methods in Physics Research Section A: Accelerators, Spectrometers, Detectors and Associated Equipment Volume 262, Issues 2-3 [https://doi.org/10.1016/0168-9002(87)90887-4]

[8] C. Savard for the CMS Collaboration, "Level-1 Track Quality Evaluation at CMS for the HL-LHC," Proceedings of 41st International Conference on High Energy physics - PoS(ICHEP2022) [doi:10.22323/1.414.0962]

[9] Rose A. et al., "2018 Serenity: an ATCA prototyping platform for CMS Phase-2," Topical Workshop on Electronics for Particle Physics (TWEPP2018) PoS 343 11 [doi:10.22323/1.343.0115]

[10] A. Albert, J. Butler, Z. Demiragli, K. Finelli, D. Gastler, E. Hazen, J. Rohlf, S. Yuan, T. Costa De Paiva and V. Martinez Outschoorn, *et al.* "The Apollo ATCA Platform", PoS TWEPP2019 (2020), 120 doi:10.22323/1.370.0120