

DAQ and Level-1 Track Finding for the CMS HL-LHC Upgrade

F. Ravera* on behalf of the CMS Collaboration

Fermi National Accelerator Laboratory

E-mail: fravera@fnal.gov

The LHC will be upgraded to the High Luminosity LHC (HL-LHC) in the late 2020s in order to reach an instantaneous luminosity as high as $7 \times 10^{34} \text{ cm}^{-2}\text{s}^{-1}$, hence increasing the discovery potential of the machine. In order to preserve physics object performance in spite of large pile-up, the CMS detector will be significantly upgraded. A key component of the upgrade is the Outer Tracker detector that will be able to identify tracks with transverse momentum above $\sim 2 \text{ GeV}/c$ and provide them to the Track Finder boards, thus maintaining manageable trigger rates and good performance. One of the main challenges of the Level-1 track finding is being able to reconstruct charged particles trajectories from a 40 MHz collision rate with a few microsecond latency budget. Dedicated FPGA hardware systems have been developed for track finding to address this challenge. Another stringent requirement on the Tracker DAQ system is set by the unprecedented number of channels, reaching two billions for the Inner Tracker only. To handle this, the Tracker DAQ back-end boards will be equipped with commercial CPUs that will guarantee the system scalability and ensure an effective monitoring of the detector conditions. The DAQ proposal to handle this distributed computational power as well as the design choices of the Level-1 track finding are presented.

*The 28th International Workshop on Vertex Detectors - Vertex2019
13-18 October, 2019
Lopud, Croatia*

*Speaker.

1. Introduction

In the late 2020s the LHC machine will undergo a major upgrade to increase its peak instantaneous luminosity up to $7 \times 10^{34} \text{ cm}^{-2}\text{s}^{-1}$, the so-called High Luminosity LHC (HL-LHC) [1]. This upgrade will allow to deliver an integrated luminosity of up to 3000 fb^{-1} to the experiments within the ~ 10 years of foreseen data taking, opening up the possibility to perform Standard Model measurements with high precision and to extend the sensitivity of searches for beyond standard model theories.

In order to guarantee high performance, the entire CMS silicon tracking system will be replaced to meet the requirements of radiation tolerance, granularity increase, material budget reduction and data transfer to the L1 trigger. A detailed description of the upgrade can be found in the Tracker Technical Design Report [2].

The CMS Phase II tracker detector is divided into two main partitions: the Inner Tracker (IT) and the Outer Tracker (OT) (fig. 1). The first one is composed of 4 m^2 of silicon pixel modules with a coverage up to $|\eta| < 4$. The OT makes use of the so-called p_T -modules that allow identification of hits from high- p_T particles at module level and provide them at a rate of 40 MHz to the L1-trigger system.

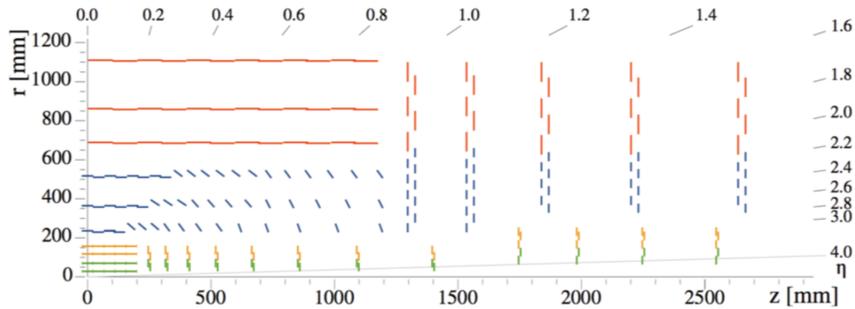


Figure 1: One quarter of the proposed CMS Phase-2 tracker layout in the $r - z$ plane. IT modules are shown in green and orange, blue and red lines show the location of Pixel-Strip (PS) and Strip-Strip (2S) OT modules, respectively.

The CMS experiment trigger system is structured in two layers: the real-time hardware-based Level-1 (L1) and the software based High Level Trigger (HLT). The first one currently relies on the information from the calorimeters and the muon system. If the event is accepted by the L1, the data from all the detectors is downloaded and passed to the HLT, which reconstructs the event and takes the decision to save the event for the offline analyses.

Simulations indicate that applying the current L1 trigger strategy to the conditions of the HL-LHC would cause a substantial degradation of the physics performance. Applying the present thresholds to events with 200 pile-up (PU) interactions per bunch crossing would result in a L1 trigger rate of 4 MHz, a factor of ~ 40 higher than the current rate [3]. Adding the L1-track information will provide an extra handle to keep the trigger rate within the planned rate of 750 kHz without a substantial increase of the thresholds. An example is shown in fig. 2 for isolated muons. In the high PU environment the track momentum measured by the muon system exhibits poor resolution resulting in an insufficient rate reduction even by increasing the p_T threshold. Adding

the information from the tracker (black points) a much steeper turn-on curve is obtained resulting in a more effective rate reduction as a function of the p_T threshold.

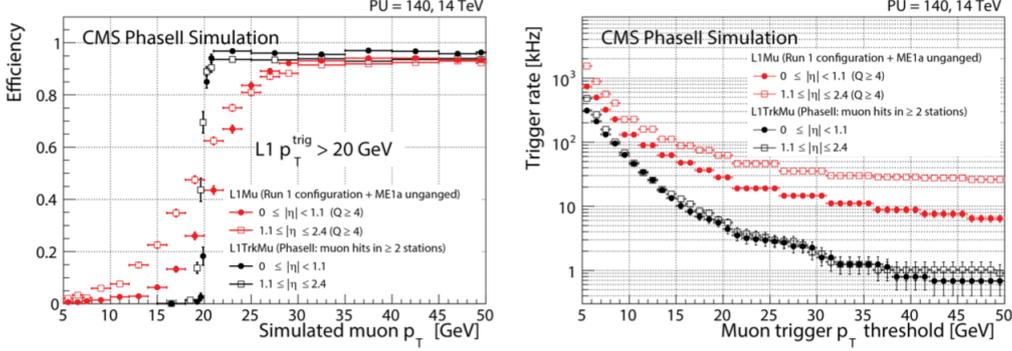


Figure 2: Single muon trigger efficiency with $p_T > 20$ GeV as a function of muon p_T (left) and the trigger rate as a function of the muon online p_T threshold (right). Red points refer to the stand-alone muon trigger, black ones to the trigger when the L1-track information is added. Regions of $|\eta| < 1.1$ and $1.1 \leq |\eta| \leq 2.4$ are indicated with solid and hollow markers, respectively [4].

This contribution provides an overview of the requirements and solutions for the data acquisition (DAQ) tracking system, and it describes the current progress on the L1-track finding.

2. The Phase II Tracker DAQ system

The tracker DAQ faces strong challenges at the HL-LHC, mainly driven by the number of read-out channels, the unprecedented particle rate due to the high PU and the trigger information needed at L1 from the Outer Tracker. In the following sections, the DAQ of the two tracker partitions are described.

2.1 Inner Tracker DAQ

The main challenge from the point of view of the IT DAQ arises from the particle rate. For the innermost layer, located at a radius of ~ 3 cm, this rate will be as high as 2 GHz/cm² (current CMS Pixel Detector maximum rate ~ 600 MHz/cm²).

In order to address both the radiation resistance and hit rate requirements, a CMS and ATLAS common development for the pixel readout chip (ROC) was carried out within the RD53 Collaboration [5]. The ROC will feature a 336×432 matrix of $50 \times 50 \mu\text{m}^2$ pixels size, which ensure an occupancy below 0.1%. The IT will feature modules made of 2×1 and 2×2 RD53 chips, indicated in green and yellow modules in fig. 1, respectively.

For each pixel hit, the chip provides the address and 4 bits of Time-Over-Threshold. In order to reduce the amount of data to be transmitted, pixels are readout in groups of 16 allowing to reduce the number of bits needed to address each of them, which results in a factor of ~ 2 compression. In order to handle the large amount of data to be transmitted, each ROC is equipped with four Aurora lines (protocol developed by Xilinx) of 1.28 Gb/s. This allows to address the requirements for the innermost layers. The final chip will implement the capability of merging data from multiple ROCs to reduce the number of optical fibers, hence reducing the material budget and the number of links

on the back-end electronics. In particular 2x1 and 2x2 modules will have the possibility to have one of the chips configured as the master, which will receive data at 320 MHz (640 MHz solution under investigation) from the other chips of the module, merge the data and ship them through one single Aurora line.

Data from each module will be transferred to the Low power GigaBit Transceiver (LpGBT) (up to 6 links per module), converted into optical by the Versatile Link+ (VL+) and sent to the back-end electronics at 10.24 Gb/s. Clock, fast-commands and programming are sent from the back-end electronics to the detector via a 2.56 Gb/s link. Both LpGBT and VL+ are CERN developed ASICs for the HL-LHC upgrades.

The back-end electronics, the so-called Data Trigger and Control (DTC), will be implemented on the ATCA standard. In the present design, each board features two Ultrascale Kintex FPGAs able to handle a total of 72 LpGBT links. After events are decoded and packed by the FPGA, data are sent to the central CMS DAQ at a maximum rate of 400 Gb/s per board. Besides the physics data taking, the IT will serve also as luminosity monitor, hence dedicated triggers will be added to the physics data stream. Data for luminosity monitoring will not exceed 200 Gb/s per board. The back-end boards will also host a System on Chip (SoC) which will be used for monitoring and to run calibrations and optimization procedures in order to guarantee the system scalability. Under these assumptions the IT will require 28 DTCs.

2.2 Outer Tracker DAQ

From the point of view of the DAQ, the main challenge of the Phase II Outer Tracker is the need of discriminating hits from high p_T tracks on detector and provide them to the L1-trigger system. This is made possible by the concept of the p_T module, which is composed by two sensors closely spaced (1.6 - 4.0 mm) and read out by the same electronics as sketched in fig. 3. For particles with higher p_T , the trajectory is less curved within the magnetic field, which results in small separation of hits in the two adjacent sensors of a module. The acceptance window translates into a p_T threshold (2 GeV/c is the baseline threshold). If a track produces hits within the accepted window, a local track segment, called stub, is formed. Stubs are transferred at 40 MHz to the back-end electronics that redirects them to the L1-tracking algorithm. All the other hits are stored in module buffers waiting to receive the L1-trigger. In case this happens, all the hits of the corresponding bunch crossing are sent to the back-end electronics and from there to the central CMS DAQ.

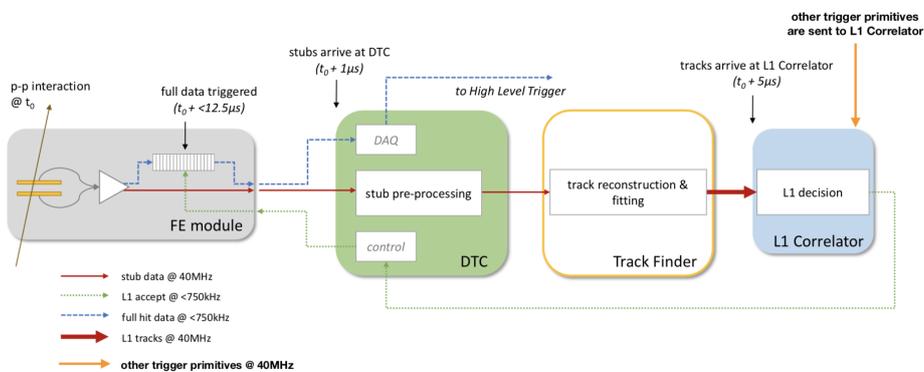


Figure 3: Sketch of the data flow (stubs and triggered data) for the p_T -module.

Two types of p_T modules will compose the OT: PS and 2S modules. In the latter, both the two strip sensors are readout by the same CMS Binary Chip (CBC) that, by matching the information from the top and bottom strips, reconstructs the stubs. In the PS module, the strip sensor is read out by the Short Strip ASIC (SSA), which sends the clusters to the Macro Pixel ASIC (MPA). The latter, bump-bonded to the pixel sensor, combines the information from the two layers and builds the stubs.

CBC and MPA send stubs (triggered data) via five (one) 320 MHz lines per chip to the Concentrator Integrated Circuit (CIC) in 2S and PS modules, respectively. The CIC aggregates data from eight CBCs or MPAs, constructs data packets and transmits them via the LpGBT via five or six lines for stubs and one for triggered data. The transmission frequency of the CIC can be configured at frequencies of 320 MHz or 640 MHz depending on the expected module data rate. Each module is equipped with a LpGBT and a VL+ to combine data from two CICs and transfer it via optical links at 5.12 or 10.24 Gb/s to the OT DTCs.

As for the IT, the OT DTCs will be implemented on the ATCA standard with similar hardware characteristics (72 LpGBT links, two Ultrascale FPGAs and SoC). The available data bandwidth will be 100 Gb/s per board to the central DAQ and 800 Gb/s per board to the L1 track finding. 216 boards will be required for the Outer Tracker.

3. Level-1 track finding at CMS for the HL-LHC

At a PU of 200 interactions, about 15000 stubs are generated per event resulting in a total data rate $O(20)$ Tb/s. The total latency foreseen for HL-LHC is $12.5 \mu\text{s}$ of which $4 \mu\text{s}$ are available for the track reconstruction. The remaining time is used to extract the stubs from the detector and provide back the L1 accepted ($1 + 1 \mu\text{s}$), correlate information from the various CMS detectors participating to the L1 trigger ($3.5 \mu\text{s}$) [3] and $3 \mu\text{s}$ are kept as contingency.

The current solution investigated by CMS has been derived from two independent approaches which are both fully based on FPGA: the Time Multiplexed Track Trigger (TMTT) and the Tracklet algorithm. Both the algorithms were investigated on a hardware demonstrator fulfilling the requirements of performance and resource utilization. In the following the two approaches are described as well as the solution currently under development.

3.1 TMTT algorithm

The Time Multiplexed Track Trigger [6] takes its name from the need of duplicating the HW that performs the tracking so that each board will run on $1/n$ events where n is the time multiplexing factor. In the proposed approach the multiplexing factor is 18. Furthermore, the detector is divided into eight sectors (octants) in azimuthal angle (ϕ). Each processing board receives data from two adjacent octants. In this way stubs are replicated, but an individual L1 tracking board does not need to share data with the others and the system is highly parallelised.

Stubs are associated to track candidates by means of a Hough Transform (HT). The algorithm, as illustrated in fig. 4, assumes, with good approximation, that the track trajectory in the $r - \phi$ plane is a circle. For a given stub with coordinate (r, ϕ) belonging to a track originated from the beam spot, track parameters p_T and ϕ_0 (the azimuthal angle of the track at the origin of the $r - \phi$ plane) to first order must satisfy $\phi_0 = \phi + r \times q / p_T$, where q is the charge of the particle. Therefore each stub

(r, ϕ) is mapped into a straight line into the Hough space $(q/p_T, \phi_0)$. Stubs belonging to the same track will intersect in the same point of the Hough space, thus constituting the track candidate. In practice, the Hough plane is implemented in the FPGA as a 32×64 matrix and a track candidate is found if at least 4 stubs are associated to the same cell. The method is intrinsically redundant against detector inefficiencies and the selected matrix cell provides the seed of the track.

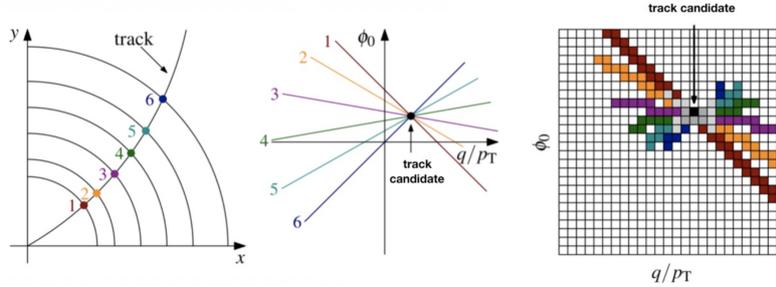


Figure 4: Illustration of the track finding algorithm via Hough transform. Stubs in the $r - \phi$ plane (left) map into straight lines intersecting in the Hough space (center) that is implemented as a matrix in the FPGA (right) .

In this approach, a Kalman filter (KF) is used to fit the tracks. Stubs associated to the same matrix cell are iteratively added to the fit improving in each step the precision of the track parameters. Stubs that are inconsistent with the extrapolation are skipped. Due to the granularity of the HT matrix, the same track may lay on adjacent cells, hence duplicates may be created. To remove those, tracks whose fitted parameters do not correspond to the ones of the HT cell in which they were seeded are discarded.

The algorithm was tested at the end of 2016 on a hardware demonstrator used to address the requirements for one time multiplexed slice and one ϕ octant. The latency, measured as the time between the arrival of the first stub in the track finder and the transmission of the first track, is $\sim 3.5 \mu\text{s}$, well within the $4 \mu\text{s}$ of latency budget. Based on the demonstrator test and considering the possible optimization of the algorithm, it was concluded that two Kintex Ultrascale FPGAs (as available on the DTC) would be able to host the firmware for one octant and one time multiplexed slice.

Physics performance was evaluated by injecting simulated $t\bar{t}$ events with an average PU of 200 into the hardware demonstrator and results were compared with a software emulator. Good tracking efficiency (greater than 95% for most of the p_T and η range) was achieved. p_T and vertex z resolutions of 1% and 2 mm, respectively, were obtained for central tracks ($|\eta| < 1.1$). Very good agreement in performance was found between the hardware demonstrator and the software emulator.

3.2 Tracklet algorithm

As for the TMTT approach, the Tracklet algorithm multiplexes events both in space and time [7]. For the hardware demonstrator [8], the detector is divided into 28 ϕ regions. No stub replication is performed among ϕ sectors, meaning that the boards have to share data; however, the sectors are built such that all tracks with $p_T > 2 \text{ GeV}$ lay only in two of them, hence boards need to share data only among contiguous neighbors. The time multiplexing factor n is driven by a balance of cost,

efficiency, and needed processing power. For 28 sectors, $n = 6$ was chosen, leading to each sector receiving data every 150 ns.

The track finding is based on a road search, as sketched in fig.5. The algorithm starts by selecting a pair of adjacent layers that will be used to create the track seed (tracklet). Seeding is done in multiple layer/disk pairs in parallel which ensures redundancy for detector inefficiencies. In order to better parallelise the processes and reduce the number of possible combinations, stubs are organized in ϕ and z regions within the ϕ sector called Virtual Modules (VM). Only stub combinations belonging to VM pairs compatible with high p_T tracks create tracklets. Assuming that the tracklet was originated from the beam spot, the initial track parameters are calculated and the track is projected on the remaining layers/disks. In each of them, the stub closest to the projection is associated to the track and its residuals are calculated.

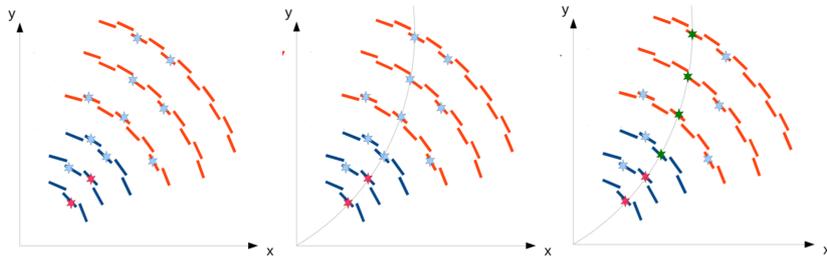


Figure 5: Illustration of the tracklet algorithm. Stubs from two adjacent layers, in this case layer 1 and 2, are chosen to form the track seed or tracklet (left). By requiring that the tracklet belongs to a particle originating from the nominal interaction point, the trajectory is projected on the other layers (center) and their stubs are associated to the track candidate (right).

A linearised χ^2 is used to extract the track parameters from the stub coordinates and residuals. In order to implement the fit algorithm into the FPGA, all complex calculations involving derivatives are pre-computed and stored in look-up tables. Duplicated tracks, mostly due to the multiple seeding, are removed by comparing tracks in pairs and counting the number of independent and shared stubs. Tracks that do not have at least three unique stubs are considered duplicates and the ones with higher χ^2 over the number of degrees of freedom are dropped.

The method was validated on a hardware demonstrator in the late 2016. Two independent implementations were developed for half of the tracker barrel and for a quarter of the barrel plus the endcap for three contiguous ϕ sectors. This allows testing the feasibility of sharing stubs among boards. Results of the implementations were compared with a software emulator for both isolated muon and $t\bar{t}$ events with a PU of 200, finding an excellent agreement. The time between the arrival of the first stub in the track finder and the transmission of the first track was measured to be $\sim 3.3 \mu\text{s}$ and with appropriate improvements the code can fit in a Kintex Ultrascale FPGA.

Having validated the C++ emulator with the HW demonstrator, performance of the algorithm was evaluated with the software simulator. As for the TMTT, good tracking efficiency ($> 95\%$) was measured and p_T and vertex z resolutions of 1% and 2 mm, respectively, were obtained for central tracks in $t\bar{t}$ events with 200 PU.

3.3 Hybrid solution

The current approach under investigation by CMS is based on a hybrid solution which merges

the two efforts: the tracklet method for track finding and the KF for track fitting based on the stubs. The detector is divided in nine ϕ nonants and the time multiplexing factor is set to 18, requiring a total of 162 DTC boards to run the L1-tracking. Stubs from two adjacent sectors are sent to the same L1-tracking hardware to avoid exchange of data among the boards.

The hybrid solution allows a list of important improvements: due to the capability of the KF to check for all possible combinations of matched stubs to build and fit the best track, wider search windows can be used in the matching stage and track candidates that share three or more stubs are merged into a single candidate removing duplicates prior to the fit, drastically reducing the number of track candidates fed into the KF; stubs forming the seed tracklet do not need to be fitted by the KF, saving latency; by removing the constraint of the track coming from the beam spot and creating tracklet from three adjacent layers, it is possible to develop triggers on displaced tracks to be used for long-lived particle searches. Current status of the development can be found in Ref. [10].

Unlike the previous approaches, this solution is being mainly developed in High Level Synthesis (HLS) language developed by Xilinx. This allows for a faster code implementation also by non-experts and also allows its application as a software emulator, since HLS code can run also outside of an FPGA.

The algorithm emulator was tested with events with PU ranging from 0 to 300. Figure 6 shows that the efficiency in reconstructing tracks with $p_T > 2$ GeV in simulated $t\bar{t}$ events is $> 95\%$. The efficiency in reconstructing muons with $p_T > 2$ GeV and electrons with $p_T > 10$ GeV is greater than 97% and 90% (fig. 7), respectively. With the same simulation it was also demonstrated that the L1-tracking algorithm works well up to a PU of 300, showing that there is a significant margin for possible increase of instantaneous luminosity in HL-LHC.

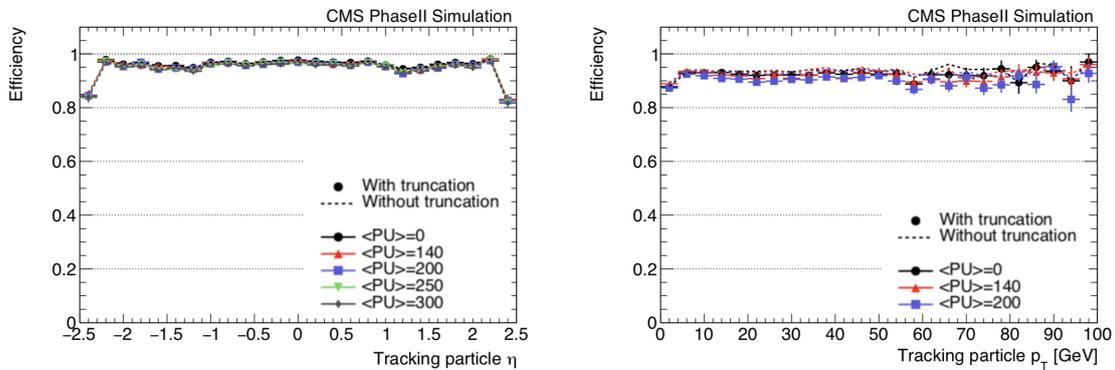


Figure 6: Track finding efficiency as a function of η (left) and p_T (right) for particle with $p_T > 2$ GeV, produced in $t\bar{t}$ events with PU of 0, 140, 200 (250 and 300 for efficiency vs η). Plots show results with and without the fixed latency cut-off (truncation) included in the simulation [10].

4. Conclusions

In order to maintain high physics performance at the HL-LHC, the CMS tracking detector will be fully replaced and the trigger and DAQ system has to meet challenging requirements. In particular the Inner Tracker will need to handle a hit rate up to 3 GHz/cm² and the Outer Tracker will have to provide stubs at 40 MHz to the back-end electronics.

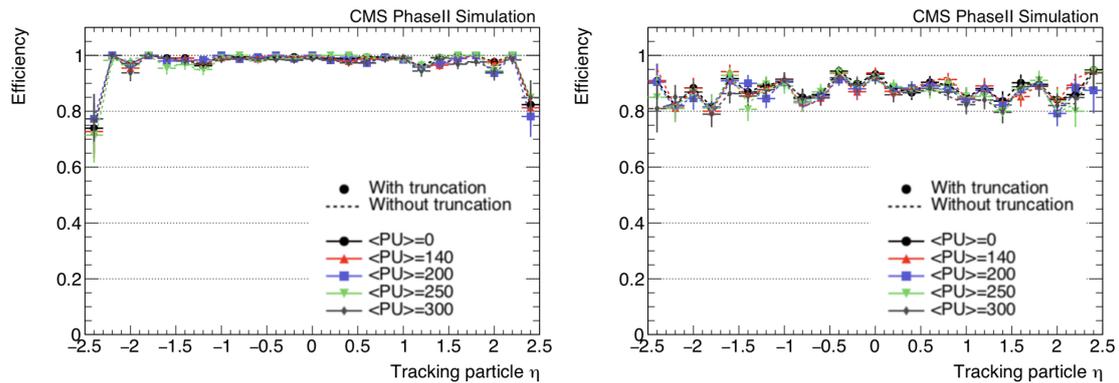


Figure 7: Track finding efficiency as a function of η for muon with $p_T > 2$ GeV (left) and electron with $p_T > 10$ GeV (right) with PU of 0, 140, 200, 250 and 300. Plots show results with and without the fixed latency cut-off (truncation) included in the simulation [10].

A key feature of the tracker upgrade is the capability to reconstruct tracks at L1. The currently considered approach is derived from two separate full-FPGA developments that were tested on hardware demonstrator showing excellent performance in terms of efficiency, track parameter resolution, latency and resource utilization. The solution investigated will open up the possibility of searches not accessible with the current detector extending the CMS physics program at HL-LHC.

References

- [1] F. Baggins et al., *High-Luminosity Large Hadron Collider (HL-LHC): Preliminary Design Report*, CERN-2015-005, (2015)
- [2] CMS Collaboration, *The Phase-2 Upgrade of the CMS Tracker*, CERN- LHCC-2017-009, CMS-TDR-014, (2017)
- [3] CMS Collaboration, *The Phase-2 Upgrade of the CMS Level-1 Trigger*, CERN-LHCC-2017-013 CMS-TDR-017, (2017)
- [4] CMS Collaboration, *Technical Proposal for the Phase-II Upgrade of the CMS Detector*, CERN-LHCC-2015-010, LHCC-P-008, CMS-TDR-15-02 (2017)
- [5] RD53 Collaboration, *RD Collaboration Proposal: Development of pixel readout integrated circuits for extreme rate and radiation*, CERN-LHCC-2013-008, LHCC-P-006 (2013)
- [6] R. Aggleton et al, *An FPGA based track finder for the L1 trigger of the CMS experiment at the High Luminosity LHC*, JINST **12** P12019 (2017)
- [7] E. Bartz et al., *FPGA-based tracking for the CMS Level-1 trigger using the tracklet algorithm*, CMS NOTE -2019/005, arXiv:1910.09970
- [8] E. Bartz et al., *FPGA-Based Tracklet Approach to Level-1 Track Finding at CMS for the HL-LHC*, EPJ Web Conf. **150** 00016 (2017)
- [9] T. James, *Level-1 track finding with an all-FPGA system at CMS for the HL-LHC*, in proceedings of CTD/WIT 2019, under publication, arXiv:1910.12668
- [10] T. James, *Level-1 track finding with an all-FPGA system at CMS for the HL-LHC*, [physics.ins-det] PROC-2019-027, arXiv:1910.12668