Overview of the HL-LHC Upgrade for the CMS Level-1 Trigger

Claire Savard[1,*] on behalf of the CMS Collaboration

[1]University of Colorado, Boulder

Abstract. The High-Luminosity LHC will open an unprecedented window on the weak-scale nature of the universe, providing high-precision measurements of the standard model as well as searches for new physics beyond the standard model. Such precision measurements and searches require information-rich datasets with a statistical power that matches the high-luminosity provided by the Phase-2 upgrade of the LHC. Efficiently collecting those datasets will be a challenging task, given the harsh environment of 200 proton-proton interactions per LHC bunch crossing. For this purpose, CMS is designing an efficient data-processing hardware trigger (Level-1) that will include tracking information and high-granularity calorimeter information. Trigger data analysis will be performed through sophisticated algorithms such as particle flow reconstruction, including widespread use of Machine Learning. The current conceptual system design is expected to take full advantage of advances in FPGA and link technologies over the coming years, providing a high-performance, low-latency computing platform for large throughput and sophisticated data correlation across diverse sources.

1 Introduction

In 2029, the Large Hadron Collider (LHC) – the world's largest and most powerful particle accelerator – will have completed an upgrade which increases the instantaneous luminosity of proton-proton collisions to 7.5×10^{34} cm^{-2}s^{-1}. This luminosity increase will allow for more statistics ($\sim$ 4000 fb^{-1}) which will in turn help in searches for new and rare physics. For example, the High-Luminosity LHC (HL-LHC) is expected to produce > 15 million Higgs bosons per year which will allow for more thorough analyses on Higgs properties and mechanics.

A consequence of increasing the LHC luminosity is an increase in additional proton-proton interactions per LHC bunch crossing (called pileup), making each collision event more complicated than ever before. The algorithms and hardware used for careful selection of interesting physics events therefore need to be updated to accommodate the harsher conditions present in each event. This work will give an overview of the updates being made to the CMS Level-1 Trigger (L1T) System – part of the event selection system for the CMS detector [1] – including a few detailed examples on novel L1T subsystems and algorithms and how they help extend our physics reach.

*e-mail: claire.savard@colorado.edu

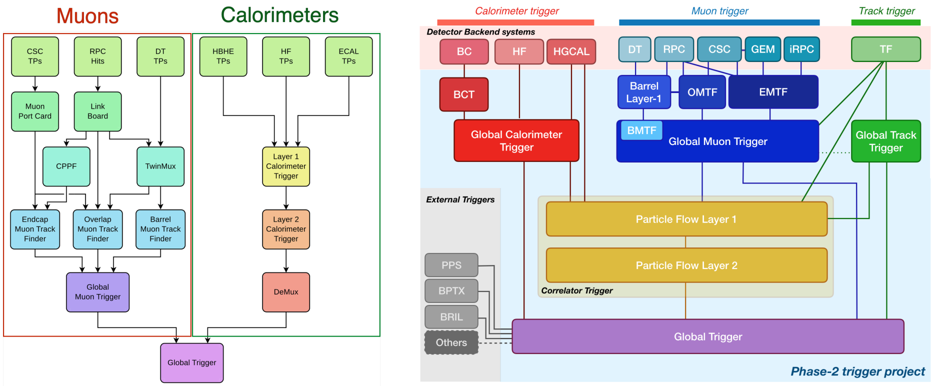

(a) (b)

Figure 1: CMS Level-1 Trigger schematic for (a) the current system and (b) the proposed upgraded system [2].

2 CMS Level-1 Trigger

The CMS detector consists of many individual subdetectors with unique technologies for probing specific elementary particles. In brief, the silicon tracker is the innermost component which reconstructs charged particle paths passing through, the calorimeters measure the energies of electrons, photons, and hadrons, and the muon system provides more precise muon measurements. Each of these subdetectors feed information into the L1T system for informative event selection. The current L1T structure can be seen in Figure 1a, and the proposed upgrade is shown in Figure 1b. In the L1T upgrade, the different components of the calorimeter, muon and global track trigger systems will feed into an intermediate correlator trigger before sending all information collected and processed to the global trigger for final selection.

During the HL-LHC era, the L1T will receive data at a rate of 40 MHz and will need to make selections within a 12.5 μs latency to output at a rate of 750 kHz. Within these tight constraints, the L1T will run sophisticated algorithms and reconstruct useful physics objects on FPGA-based hardware. All upgrades are made with the goal of maintaining current physics reach given an increase in pileup. Furthermore, we will increase the coverage and probe physics phase spaces that were previously inaccessible through the capabilities of the new system; from increased granularity and acceptance of the detectors, improved electronics, incorporating new trigger systems and utilising innovative techniques including Machine Learning.

2.1 Architecture

All firmware for the L1T system will be held on FPGAs due to the hardware's computing speeds and re-programmable nature. All FPGAs used will be the same chip, currently expected to be Xilinx VU13P. The FPGAs will be placed onto PCB boards with ATCA form factor. The functions performed by the different board families are different and therefore have unique firmware. Information is transferred between boards through 25 Gb/s optical fibers.

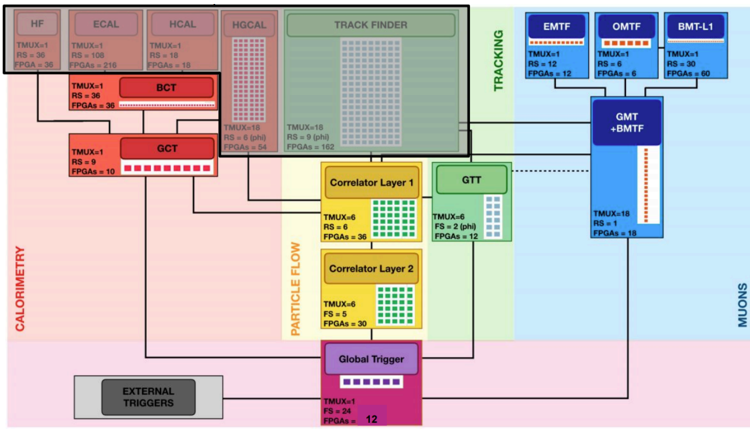

Figure 2: The proposed L1T architecture mapping indicating the full layout of the PCBs, FPGAs that they contain, and optical links between the boards. The gray shaded section outlined in black at the top left indicates the trigger boards associated with the detector backend system which creates trigger primitive objects used further down the trigger pipeline [2].

Figure 2 shows a map of the proposed architecture for the upgraded CMS L1T system. The different highlighted sections indicate the separate trigger subsystems, such as the section labeled "muons". With each highlighted section, the dark boxes indicate PCBs, such as the box labeled "EMTF". Within each PCB box, smaller boxes laid out like a grid indicate different FPGAs assigned to that PCB. The optical links are represented by a black line which shows the flow of information downward.

2.2 Subsystem upgrades

There are 4 independent data processing paths in the upgraded L1T, as seen in Figure 1b: calorimeter trigger, muon trigger, global track trigger, and correlator trigger. When compared with the current L1T structure shown in Figure 1a, we can see that two of the subsystems, namely the correlator and track triggers, are completely new. The calorimeter and muon triggers have also been given major upgrades to provide more accurate information of the physics within the corresponding sub-detector.

The most important subsystem components are summarized below:

- **Calorimeter trigger** The calorimeter trigger receives information from the electromagnatic and hadronic calorimeters. The main calorimeter trigger objects built are: electrons, photons, jets, hadronically decaying taus, and various energy sums. In comparison to the current calorimeter trigger, the updated calorimeters, including the new High Granularity Calorimeter, will provide higher granularity, which will allow for high-resolution clusters and identification variables that increase the accuracy in the trigger objects.
- **Muon trigger** The muon trigger aims to identify and reconstruct muon tracks passing through the CMS detector. This trigger system will take in information from both the muon system and tracking system to create the muon trigger objects. In comparison to the current muon trigger, the upgrade will extend its reconstruction coverage from $|\eta| < 2.4$ to $|\eta| < 2.8$ and add in new primitive information from the CMS silicon tracker to those from the separate muon track finders in different regions of the muon system.

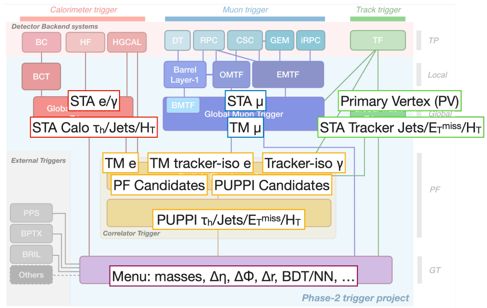

Figure 3: Schematic showing where on the trigger system certain trigger objects are reconstructed. STA is stand-alone, meaning it only take information from one sub-system. TM is track-matched, meaning it takes in information from the tracker and another system. PF and PUPPI combine information from all sub-systems.

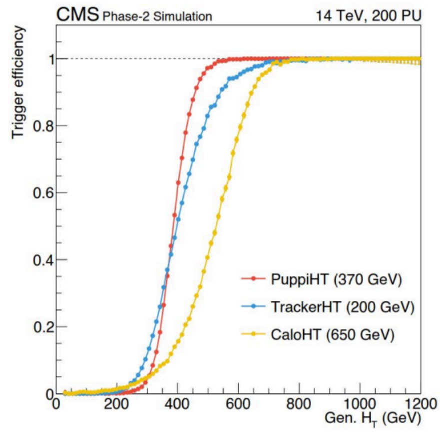

Figure 4: The trigger efficiency as a function of H_T for $t\bar{t}$+pileup events with a fixed rate of 10.5 kHz for H_T reconstructed in the global track trigger (blue), calorimeter trigger (yellow), and correlator trigger (red) [2].

- **Global track trigger** The Global Track Trigger (GTT) is a new trigger system being introduced in the upgrade that takes in information from the track finder module which reconstructs charged particle tracker tracks. This trigger system will then use these tracks to build high-level track objects, such as jets, vertices, and jet H_T. These trigger objects can help significantly in tasks like pileup mitigation. Further details and examples are given in Section 3.1.
- **Correlator trigger** The correlator trigger is another new trigger system. It's function is to aggregate all information processed from the previous three upstream trigger systems (calorimeter, muon, and global track triggers) to achieve the best possible trigger performance on challenging physics topologies. There are two essential algorithms on the correlator trigger: Particle Flow identifies and reconstructs all particles given all sub-detector information, and Pileup Per Particle Identification (PUPPI) mitigates pileup effects. The correlator trigger ultimately creates more accurate trigger objects like hadronic taus, jets, missing transverse energy, and H_T.

A summary of where the trigger objects are created can be found in Figure 3. Figure 4 shows how the trigger efficiency changes when you look at the different H_T objects reconstructed in different parts of the level-1 trigger system. The H_T object reconstructed in the correlator trigger, "PuppiHT", which can reconstruct physics objects with more overall knowledge from all upstream trigger subsystems, achieves the highest trigger efficiency.

2.3 40 MHz scouting system

The upgraded CMS trigger system will also have a scouting system that runs separately but in parallel to the trigger system connected through spare optical links. The scouting system takes in a subset of trigger primitives and objects at 40 MHz to monitor trigger system performance, which is very useful as the global trigger is reprogrammed several times a year. This continual monitoring allows for real-time trigger component diagnostics. Furthermore, new trigger algorithms can be prototyped and tested rapidly using the scouting system.

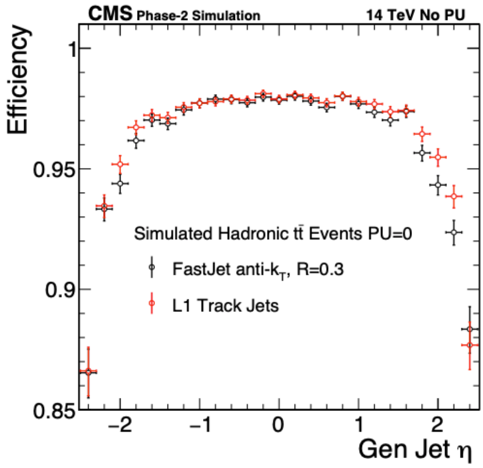

Figure 5: The jet finding efficiency of reconstructed jets within ΔR < 0.4 of generator jets as a function of η for a proposed reconstruction algorithm "L1 Track Jets" and an alternative method "FastJet" used as a benchmark for performance [2].

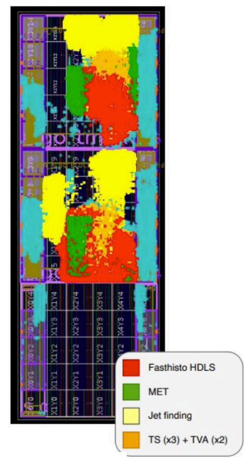

Figure 6: An example floor plan of the GTT showing the placement and space required for different algorithms. "Fasthisto" is a vertexing algorithm and "MET" stands for $p_{\mathrm{T}}^{\mathrm{miss}}$.

3 L1T Upgrades Examples

In this section we expand on some selected features and applications of the upgraded L1T.

3.1 Trigger subsystem: Global Track Trigger

The GTT is a new subsystem of the trigger system that builds track-level physics objects. The charged particle track inputs are reconstructed in the Track Finder (TF) with silicon outer tracker information at 40 MHz. These tracks are stored in a 96-bit track word that contains information about different track properties, such as the track ϕ, tanλ, z_0, n_{stub}, quality, etc.

Charged particle tracks are combined in various ways in the GTT to create a few different physics objects. Tracker jets are made up of tracks clustered within a certain vicinity of one another, and H_{T} is built from the tracker jet collection per collision. Primary vertices can also be created from clusters of tracks along the z axis (beamline) of the CMS detector. Missing transverse energy ($p_{\mathrm{T}}^{\mathrm{miss}}$) can be calculated by aggregating the track transverse momenta in each collision and determining the amount missing.

Figure 5 shows the performance of the track jet reconstruction algorithm in comparison to another, more computationally intensive, algorithm used as a benchmark for performance. Figure 6 shows an example floor plan of the GTT on a VU9P FPGA. Current algorithm implementations on the GTT FPGAs deplete about 15% of the Lookup Table (LUT), 20% of the Flip Flops (FF), 25% of the Block RAM (BRAM), and 3% of the Digital Signal Processing Elements (DSP) available.

3.2 Algorithm: End-to-End Neural Network Vertexing

A proposed method of creating track-level primary vertices on the GTT is to use a Neural Network. The input to the algorithm is a full collection of Level-1 tracker tracks, and the output

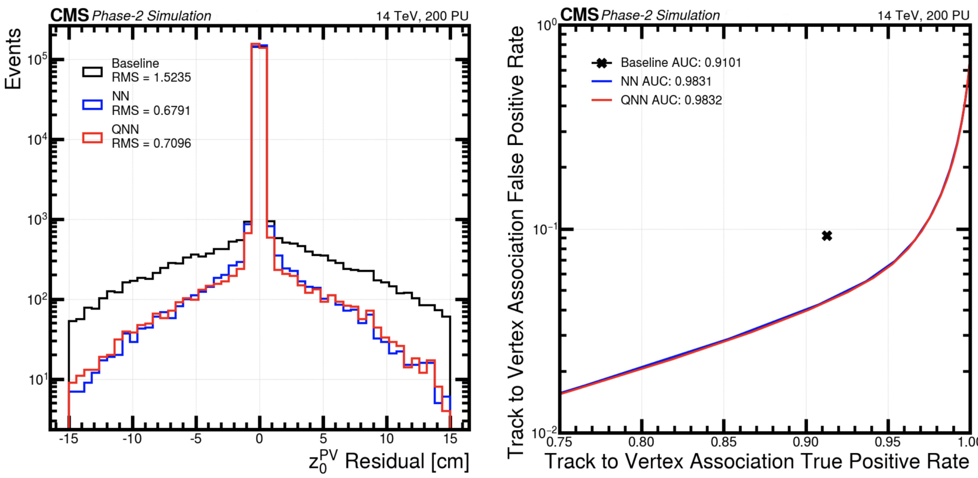

(a) (b)

Figure 7: Comparison of primary vertex reconstruction algorithms along the z_0 residual (a) and when looking at the ROC curves for track to vertex association (b) [3].

is a likelihood that each track belongs to the PV. This network is referred to as "end-to-end" because tracks are input without any initial track selection, and so no additional computing is required on the input as is common with vertexing algorithms. This Machine Learning approach to vertexing is a suggested upgrade to the baseline approach, which clusters tracks along the beamline (z) to determine which tracks originate from the same vertex.

The performance comparing the baseline approach, floating point Neural Network approach, and quantized Neural Network approach can be found in Figure 7. The Neural Network approach is more accurate in determining the correct PV z_0 location than with the baseline algorithm. In order to place the algorithm on an FPGA, all internal numerical values must be quantized with fixed precision. Figure 7 shows that the quantized version of the network (red) does not lose any performance in comparison to the floating-point version (blue).

3.3 Physics Reach: Exotic Higgs

The Phase-2 L1T will maintain the current physics reach, as well as extend that reach into physics realms previously inaccessible with the current L1T. For example, many beyond standard model (BSM) signals are not capable of being explored due to bandwidth constraints of the current L1T. One such BSM signal is the exotic Higgs boson decay $h \to \phi\phi \to 4j$ where j stands for jets. The ϕ scalars are long lived and therefore decay to jets far from the collision point. In the Phase-2 upgrade of the L1T, the H_T trigger threshold can be reduced while maintaining the acceptable L1T rates, meaning that more high-jet multiplicity events such as these can be recorded.

The GTT will be integral in the triggering of this signal through the reconstruction of displaced jets made up of displaced tracks coming from the TF. Displaced tracking is an exciting addition to the TF as it will allow for more accurate investigation of many BSM events that may be produced. With the addition of displaced tracking and a reduced H_{T} trigger threshold, the CMS L1T will be able to trigger on an exotic Higgs boson decay.

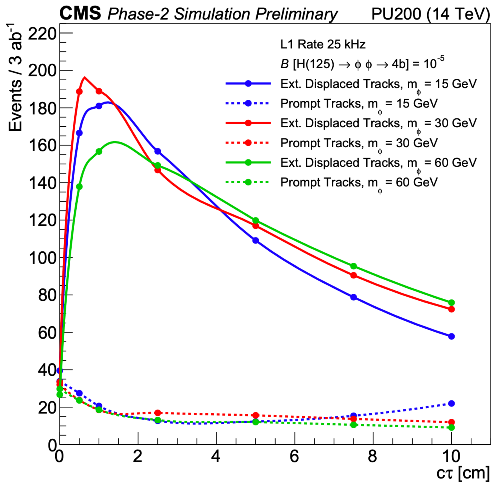

Figure 8: The expected number of signal events that will pass the H_T trigger for different m_ϕ and when displaced tracking is included in addition to prompt tracking or not [2].

Figure 8 shows the expected number of signal events that the upgraded L1T will be able to accept if displaced tracking is included in addition to standard prompt tracking and for a variety of different ϕ masses. We see that the addition of displaced tracking (upper solid lines) is key to ensuring that we can capture this signal.

4 Summary

The next-generation HL-LHC will increase the statistics for all physics searches, but also produce more complicated events due to increased pileup. To maintain the same level of physics reach as the current L1T, and also extend beyond the current capabilities, the L1T system will be upgraded to allow for the reconstruction of more sophisticated, offline-like, physics objects. The proposed trigger will be fully FPGA-based, and will host more sophisticated algorithms using techniques like Machine Learning while staying within the desired rate and latency constraints. A number of benchmark tests which prove the feasibility of the phase 2 L1T design have been studied, and there is ongoing work to continue to improve upon the current baseline model.

References

[1] *The CMS experiment at the CERN LHC*. CMS Collaboration. JINST 3 (2008) S08004, doi:10.1088/1748-0221/3/08/S08004. url: https://cds.cern.ch/record/1129810.

[2] *The Phase-2 Upgrade of the CMS Level-1 Trigger*. CERN-LHCC-2020-004, CMS-TDR-021. CMS, 2020. url: https://cds.cern.ch/record/2714892.

[3] *Neural Network-Based Primary Vertex Reconstruction with FPGAs for the Upgrade of the CMS Level-1 Trigger System*. C. Brown et al 2023 J. Phys.: Conf. Ser. 2438 012106