Available on CMS information server

CMS CR -2024/010

The Compact Muon Solenoid Experiment

Conference Report

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

07 January 2024 (v2, 15 January 2024)

CMS Level-1 Trigger Upgrade

Jose Enrique Palencia Cortezon for the CMS Collaboration

Abstract

To achieve the physics goals of the High Luminosity LHC program, the CMS Phase-2 upgrade of the Level-1 trigger system utilizes modern technologies to enhance the physics selectivity already at the hardware level of the data acquisition system. The intense hadronic environment corresponding to 200 simultaneous collisions per bunch crossing, imposes serious challenges to the system requirements. To profit from the extended coverage and increased granularity of the upgraded CMS detector, the latency of the system is extended to 12.5 μs. The use of tracking and high-granularity calorimeter information becomes possible for the first time at Level-1. The maximum output bandwidth is 750 kHz. Modern processors are used to implement sophisticated algorithms including machine learning-based approaches to target the selection of specific final states.

Presented at *LHCP2023 The 11th annual conference on Large Hadron Collider Physics*

PROCEEDINGS OF SCIENCE

CMS Level-1 Trigger Upgrade

Jose Enrique Palencia Cortezon (for the CMS Collaboration)[a,*]

[a] *Universidad de Oviedo, Instituto Universitario de Ciencias y Tecnologías Espaciales de Asturias (ICTEA), Spain*

E-mail: enrique.palencia@uniovi.es

To achieve the physics goals of the High Luminosity LHC program, the CMS Phase-2 upgrade of the Level-1 trigger system utilizes modern technologies to enhance the physics selectivity already at the hardware level of the data acquisition system. The intense hadronic environment corresponding to 200 simultaneous collisions per bunch crossing, imposes serious challenges to the system requirements. To profit from the extended coverage and increased granularity of the upgraded CMS detector, the latency of the system is extended to 12.5 μs. The use of tracking and high-granularity calorimeter information becomes possible for the first time at Level-1. The maximum output bandwidth is 750 kHz. Modern processors are used to implement sophisticated algorithms including machine learning-based approaches to target the selection of specific final states.

The Eleventh Annual Conference on Large Hadron Collider Physics (LHCP2023)
22-26 May 2023
Belgrade, Serbia

[*]Speaker

1. Introduction

These proceedings describe the ongoing developments and plans towards the upgrade of the CMS [1, 2] Level-1 trigger (L1T) for the High-Luminosity Large Hadron Collider (HL-LHC) [3]. In its ultimate configuration, the HL-LHC will reach unprecedented performance in terms of instantaneous luminosity (7.5×10^{34} cm^{-2}s^{-1}) leading to 200 simultaneous collisions per bunch crossing and potentially delivering a total integrated luminosity of up to 4000 fb^{-1} after ten years of operations, scheduled to start in 2029. This previously unmatched amount of data opens the door to a rich and ambitious physics program including both high-precision measurements and searches for physics beyond the Standard Model (BSM).

In order to fully exploit the HL-LHC running period, major consolidations and upgrades of the CMS detector are planned. Along with the sub-detector upgrades, a complete replacement of the trigger (L1 and HLT) and data acquisition (DAQ) system, with increased throughput, is planned. The selection algorithms will perform a rate reduction leading to an output bandwidth of 7.5 kHz.

2. The L1 Phase-2 upgrade

The Phase-2 upgrade of the L1T system is designed not only to maintain the efficiency of the signal selection to the level of the Phase-1 performance, but also to significantly enhance, or enable, the selection of any possible new physics manifestations that could lead to unconventional signatures. The Phase-2 upgrade of the trigger and DAQ system will keep a two-level strategy while increasing the L1 maximum rate to 750 kHz. The total latency will be increased to 12.5 μs to allow, for the first time, the inclusion of the tracker and high-granularity calorimeter information. Moreover, a longer latency will enable higher-level object reconstruction and identification, as well as the evaluation of complex global event quantities and correlation variables to optimize physics selectivity. The implementation of sophisticated algorithms using particle-flow (PF) reconstruction techniques or machine learning (ML) based approaches can now be contemplated. In addition to these features, a 40 MHz scouting system harvesting the trigger primitives produced by sub-detectors and the trigger objects produced at various levels of the trigger system is proposed.

The CMS L1 Phase-2 trigger system is designed to benefit from the new features provided by the upgraded sub-detectors to sustain a high efficiency of physics event selection in the very high luminosity regime. The functional diagram of the architecture and data flow of the Phase-2 trigger system is presented in Fig. 1. With the 12.5 μs latency, not only information from the calorimeters and muon detectors is used (as in the Phase-1 system), but the information from the new tracker and high-granularity endcap calorimeter can also be included. The key feature of the proposed system is the introduction of a correlator layer, which implements sophisticated algorithms producing higher-level trigger objects resulting from the combination of the information of multiple sub-detectors to achieve enhanced selectivity, approaching that of the offline reconstruction. To achieve optimum flexibility of the design with the required robustness, four independent data processing paths are implemented: tracking, calorimetry, muon systems, and PF techniques. This division reflects the need to generate complementary types of trigger objects to achieve the best physics selectivity.

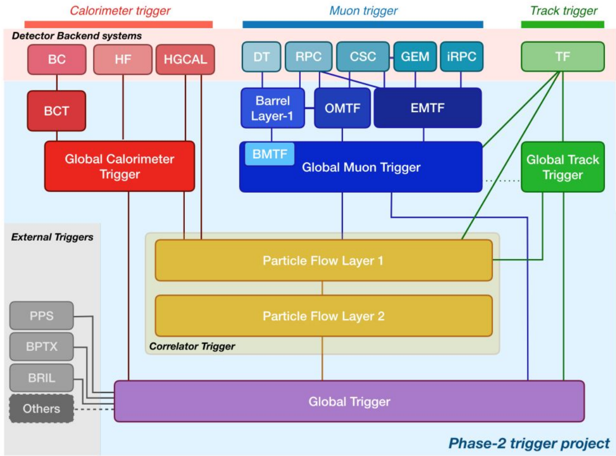

Figure 1: Functional diagram of the CMS L1 Phase-2 upgraded trigger design [3].

3. Recent developments

The inclusion of tracking into the L1T offers the possibility of developing an algorithm to identify jets originating from bottom quarks (b-tagging) for use in the L1T for the first time at CMS. With the use of the outer tracker information only, sent to L1T as fully reconstructed tracks, b-tagging can be achieved in hardware with promising performance. In order to make effective use of the new information CMS will implement specialized versions of the PF and Pileup Per Particle Identification (PUPPI) algorithms. The particles reconstructed using the PUPPI algorithm will then be used to construct higher level objects, such as jets. An algorithm using a neural network (NN) has been implemented on the Correlator Layer-2 system where jets are built has been developed [4]. This implementation is also capable of operating within the budgeted latency required for the CMS detector readout. The algorithm takes as input the PUPPI particles in each jet, and is capable of discriminating between jets originating from bottom quarks and jets originating from light quarks or gluons. As shown in Fig. 2 (left), the b-tag NN trigger increases the efficiency for events with low di-Higgs mass by up to a factor of 1.5 over the QuadJet+H_T or Jets+Muon triggers.

Part of the upgrade will see the L1T use charged particle tracks within the full outer silicon tracker volume as an input for the first time and new algorithms are being designed to make use of these tracks. One such algorithm is the primary vertex finding which is used to identify the hard scatter in an event and separate the primary interaction from additional simultaneous interactions. A novel approach to regress the primary vertex position and to reject tracks from additional soft interactions, which uses an end-to-end NN, is shown [5]. This NN possesses simultaneous knowledge of all stages in the reconstruction chain, which allows for end-to-end optimisation. This network improves the performance of the baseline approach in the primary vertex regression and track-to-vertex classification. A quantised and pruned version of the NN is deployed on an FPGA to match the stringent timing and computing requirements of the L1T.

From the list of particles reconstructed in the Correlator, where the PUPPI algorithm is also applied to suppress particles that originate from pileup interactions, some final state objects, such as jets, are reconstructed. The Seeded Cone (SC) jet algorithm [6] is an approach to reconstructing jets from these particles using the full available granularity. It has been implemented in the Xilinx Ultrascale+ FPGA used in the Correlator system. The algorithm latency is within the 1 μs constraint of the subsystem. The algorithm begins by finding the highest p_T particle in the event, then finding all particles within a cone of this seed in (η, ϕ) to form the first jet. The constituents of the jet are removed from the processing before the process is repeated in order to find up to 16 jets in the event. The jet axis is computed as the sum over the constituents of the p_T and the p_T-weighted sum of their η and ϕ. Figure 2 (right) shows similar performance of the SC and anti-k_t jet algorithms, both with a radius parameter of 0.4 and using the same L1T PUPPI candidates as inputs.

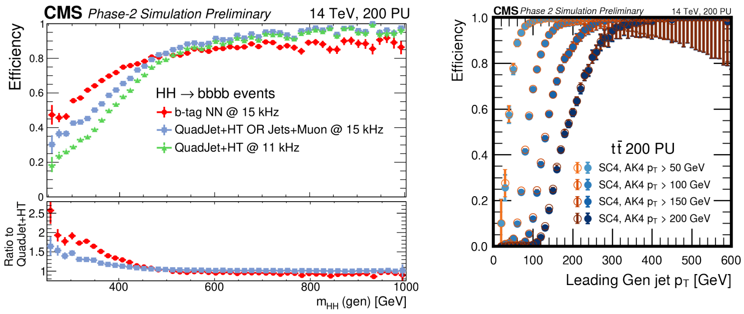

Figure 2: Left: Trigger efficiency in HH → bbbb events as a function of the di-Higgs mass at generator level [4]. Rigth: turn-on curve with different online jet thresholds [6].

A Continual Learning (CL) approach [7] is investigated to solve the problem of a changing environment where a ML is constantly updated using a stream of labelled data. In the case of the L1T hardware, this could be from the unbiased 40 MHz scouting data that is used as truth level training data for a model. This would be a solution to the fill-level changes in detector conditions where generating a large Monte Carlo (MC) dataset is not feasible. The CL solution has the advantages of not needing a large MC dataset, so can react quickly to different conditions compared to the timescale of a large MC production campaign. It also can perform small updates to a stable model footprint as compared to training a fresh model which could have variations in the quantisation or pruning, which is especially important in the low-resource, low-latency environment of the L1T.

4. Physics reach

The CMS Phase-2 physics program plans to fully exploit the HL-LHC to perform searches for BSM physics, achieve unprecedented high-precision SM measurements, including a significantly improved characterization of the Higgs Boson sector. A broad spectrum of physics analyses will become possible with the exceptionally large HL-LHC data samples and the new capabilities offered by the detector upgrade, such as its extended coverage. More advanced selection algorithms are already at trigger level to maintain the effective selection of electroweak-scale processes with 200 pileup events, including Vector Boson Scattering or Vector Boson Fusion production, rare B-meson decays (based on usage of tracks in the L1 trigger for the first time), long-lived particles, etc. Low mass resonances could also be identified with a dedicated scouting system.

References

[1] CMS Collaboration, *The CMS Experiment at the CERN LHC*, JINST **3** (2008) S8004, doi:10.1088/1748-0221/3/08/S08004.

[2] CMS Collaboration, *Development of the CMS Detector for the CERN LHC Run 3*, accepted by JINST (2023), arXiv:2309.05466.

[3] CMS Collaboration, *The Phase-2 Upgrade of the CMS Level-1 Trigger*, CMS-TDR-021.

[4] CMS Collaboration, *Neural network-based algorithm for the identification of bottom quarks in the CMS Phase-2 Level-1 trigger*, CMS-DP-2022-021.

[5] C. Brown et al, *Neural Network-Based Primary Vertex Reconstruction with FPGAs*, 2023 J. Phys.: Conf. Ser. 2438 012106, doi:10.1088/1742-6596/2438/1/012106.

[6] CMS Collaboration, *Jet Reconstruction with the Seeded Cone algorithm in the CMS Phase-2 Level-1 Trigger*, CMS-DP-2023-023.

[7] CMS Collaboration, *Continual Learning in the CMS Phase-2 Level-1 Trigger*, CMS-DP-2023-022.