PROCEEDINGS OF SCIENCE

Design and performance of the upgrade of the CMS L1 trigger

Olivier Davignon* **on behalf of the CMS Collaboration**
University of Bristol, Bristol, United Kingdom
E-mail: davignon@cern.ch

During its second run of operation, the LHC delivered proton-proton collisions at a center-of-mass energy of 13 TeV with a peak instantaneous luminosity larger than 2×10^{34} cm$^{-2} \cdot$ s^{-1}, more than double the peak luminosity reached during Run-1 and far larger than the design value. The upgraded CMS Level-1 trigger is designed to improve the performance at high luminosity and large number of simultaneous inelastic collisions per crossing (pileup). During the technical stop at the beginning of 2016, all the electronic boards of the CMS Level-1 trigger have been replaced and the upgraded electronics tested, and commissioned with data. Smarter, more sophisticated, and innovative algorithms are now the core of the first decision layer of CMS: the upgraded trigger system implements pattern recognition and MVA (Boosted Decision Tree) regression techniques in the trigger boards for p_T assignment, pileup subtraction, and isolation requirements for electrons and taus. In addition, the new global trigger is capable of evaluating complex selection algorithms such as those involving the invariant mass of trigger objects. The upgrade reduces the trigger rate and improves the trigger efficiency for a wide variety of physics signals. In this paper, the upgraded CMS Level-1 trigger design and its performance are described.

The 39th International Conference on High Energy Physics (ICHEP2018)
4-11 July, 2018
Seoul, Korea

*Speaker.

PoS(ICHEP2018)198

At the CMS experiment [1], the trigger system is used to perform an online selection of physics processes of interest whose cross sections are typically orders of magnitude lower than the total proton-proton cross section. The trigger is implemented as a two-level system. First, the Level-1 (L1) [2] reduces the event rate from 40 MHz to 100 kHz with a total latency of 3.8 μs, and consists of custom electronics hardware. Then, the High-Level Trigger (HLT), based on software running on CPUs, subsequently reduces the event rate from about 100 kHz to 1 kHz. To maintain or improve its performance despite increasing LHC luminosity, the CMS L1 trigger subsystem was fully upgraded in 2016 (Phase-1 upgrade) [3].

1. Calorimeter trigger upgrade

The upgraded calorimeter trigger [4] is organized as follows. In ECAL and HCAL, the signals from multiple adjacent cells are combined into a "Trigger Tower (TT)". The ECAL and HCAL Trigger Primitive Generators in the Front-End boards combine and format the TT information into "Trigger Primitives". Those are then sent to one of the 18 Layer-1 CTP7 processing boards [5], which perform energy calibration. Every calibrated TT from a single bunch crossing is then sent by the Layer-1 cards for further processing by a *single* Layer-2 MP7 board [6]; each event is processed by one of nine boards in a round robin fashion. The Layer-2 builds the e/γ, τ-leptons, jets, and energy sum objects in the event. The Layer-2 benefits of a full view of the CMS calorimeters.

Level-1 jets are built as a collection of 9×9 TTs in η/ϕ located around a local maximum called the jet "seed". To estimate the contribution of pileup particles to the transverse jet energy, the energies from 4 strips (3×9 TTs) around the jet are computed; the E_{T} of the 3 lowest-energy strips are summed and subtracted from the jet E_{T}. Finally, the E_{T} is calibrated, using a jet-by-jet correction which depends on E_{T} and the jet position in η.

The L1 e/γ reconstruction (clustering) proceeds around a TT local maximum, and builds shapes dynamically from contiguous TTs to a maximal shape size of about 2×5 TTs in η/ϕ. The shapes least compatible with a genuine e/γ candidate are rejected. Additional identification criteria are defined: H/E, which quantifies the hadronic-to-electronic ratio of energy deposits; and isolation transverse energy $E_{\mathrm{T}}^{\mathrm{iso}}$, which quantifies the E_{T} deposit in the 6×9 TTs in η/ϕ around the e/γ from which the e/γ E_{T} is subtracted. To identify an isolated e/γ object, a cut is applied on $E_{\mathrm{T}}^{\mathrm{iso}}$ as function of $E_{\mathrm{T}}^{e/\gamma}$, the η position, and a pileup estimator called n_{TT}. The later is obtained by counting the number of TTs with $E_{\mathrm{T}}^{\mathrm{TT}} \geq 0.5$ GeV in the region with $\eta \in [-0.348, +0.348]$. The E_{T} of L1 e/γ is calibrated using a per-object correction function of E_{T}, the position in η and the reconstructed shape. A BDT regression technique is used offline to derive the set of corrections.

The L1 hadronically-decaying τ reconstruction shares the same clustering and isolation computations as e/γ. A potential merging of two neighbouring clusters (named *primary*/*secondary*) has been implemented to trigger efficiently on multi-prong τ decays. An isolation cut, as a function of E_{T}, η, and n_{TT}, is applied to discriminate τ from QCD jets. The n_{TT} dependence of the isolation cut is used to ensure the stability of the τ trigger efficiency. The E_{T} is also calibrated using a per-object correction function of E_{T}, the position in η, a simple estimate of the H/E, and a flag indicating whether the τ is merged or not. A BDT regression technique is used offline to derive the set of corrections. The performance of the L1 τ algorithm was measured in Run-2 data for τs from $Z \to \tau_\mu \tau_h$ decays selected using a tag & probe technique. The excellent stability of the efficiency with respect to pileup, measured in 2017 data, is demonstrated in Fig. 1 (left).

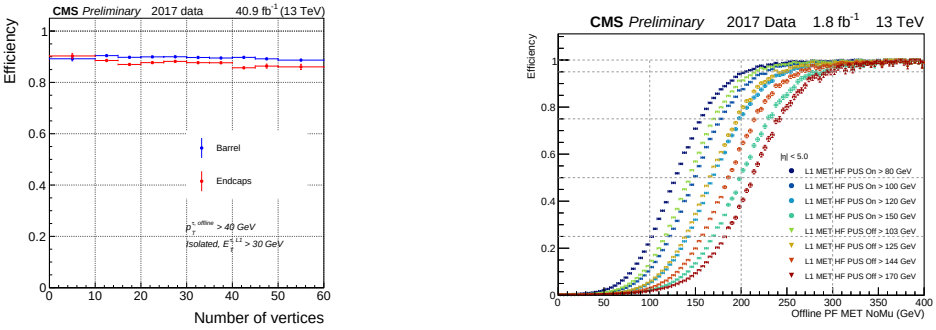

Figure 1: (Left) L1 τ trigger efficiency vs. the number of offline reconstructed vertices, for an L1 selection requiring at least one L1 isolated $E_{\mathrm{T}} > 30$ GeV τ candidate matched to a well-identified offline τ with $p_{\mathrm{T}} > 40$ GeV. Taken from CMS-DP-2018/006 [7]. (Right) Efficiencies for the upgraded L1 MET trigger as a function of offline MET. Efficiency curves with and without pileup subtraction (PUS) applied are compared for thresholds that give the same rate: L1 MET PUS on > 80/100/120/150 GeV correspond to PUS off > 103/125/144/170 GeV, respectively. Taken from CMS-DP-2018/004 [7].

The L1 missing transverse energy (MET) is defined as the negative vector sum of $\vec{E}_{\mathrm{T}}^{\mathrm{TT}}$ for all TTs up to $\eta = 5$. To reduce the rate, and ensure the stability of the efficiency with pileup, energy deposits below a dynamic η-dependent E_{T} threshold, calculated as a function of n_{TT}, are excluded from the sum. The L1 MET efficiency, as measured in 2017 data, using an unbiased sample triggered by muons, is shown in Fig. 1 (right).

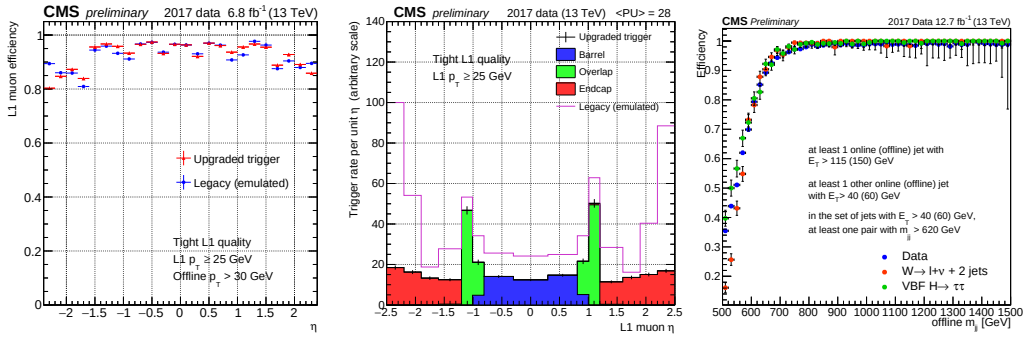

Figure 2: (Left) L1 μ tigger efficiency vs. geometrically matched and well-identified offline $p_{\mathrm{T}} > 30$ GeV muon η for a L1 muon p_{T} threshold of 25 GeV, compared between the upgraded and legacy systems. (Center) Distribution of number of background μ per unit η passing a p_{T} threshold of 25 GeV, built by the three track finders in the upgraded L1 μ trigger, compared with the emulated legacy trigger. Both taken from CMS-DP-2017/041 [7]. (Right) Efficiency of the L1 VBF-trigger as a function of the offline m_{jj}, estimated as the fraction of analysis-like offline events passing the L1 VBF-trigger selection. The efficiency is evaluated on 2017 data, on $W \to l\nu$ +2 jets and VBF H$\to \tau\tau$ simulations. Taken from CMS-DP-2018/005 [7].

2. Muon trigger upgrade

Through its upgrade, the Level-1 muon triggers moved from a muon detector-based scheme (DT, RPC, and CSC) to a geometry-based system of Muon Track Finders: Barrel (BMTF) [8], Overlap (OMTF) [9], and Endcap (EMTF) [10], covering the pseudorapidity range $0 < |\eta| < 0.83$, $0.83 < |\eta| < 1.23$ and $1.24 < |\eta| < 2.4$, respectively. The goal of the new architecture is to bring in together data from the 3 complementary detector technologies early in the track finding procedure to improve p_{T} resolution. The system also takes advantage of the redundancy to achieve higher

efficiency and better rate reduction. Both the OMTF and EMTF use pattern-based track finding and assign the p_T of a given track using look-up-tables (LUT). In EMTF, the LUT is derived offline through a multivariate BDT technique. The OMTF and EMTF are based on MTF7 processing boards. In BMTF, a road search extrapolation track finder is used to reconstruct muon candidates. After reconstruction, muons from all three systems are processed in the Global Muon Trigger (GMT), which provides transverse momentum, a quality assessment of the muon candidate, and in order to remove duplicate candidates, a dedicated cancellation between different track finders.

The efficiency of the upgraded L1 muon trigger system was measured on unbiased muons from $Z \to \mu\mu$ events of 2017 data, selected using a tag & probe technique. It is compared to the Run-1 system ("Legacy"), whose performance is assessed by re-emulating the Legacy decision on the same events. Unbiased events from 2017 data are used to estimate the rates. Fig. 2 demonstrates the superiority of the upgraded system, achieving similar efficiencies at a much lower rate.

3. Global Trigger upgrade

The Global Trigger (GT) [11] receives all objects from the calorimeter trigger and muons from the GMT, and makes the final decision whether to pass an event to HLT or not. The GT performs a combination of objects, building algorithms tailored for physics. The GT is implemented on 6 MP7 boards and accommodates large *trigger menus* of up to 512 algorithms. A flexible trigger menu loading allows CMS to adapt to the evolving luminosity and pileup conditions from the LHC. For example, twice as many cross-triggers (e.g. electron+τ) were implemented in 2017 as in 2016.

The large processing power enables complex correlation algorithms, such as invariant mass computations. It was used to design the first L1 trigger algorithm explicitly looking at the Vector Boson Fusion (VBF) Higgs production topology called the "VBF-trigger". It relies on the tagging of VBF's forward/backward jets, and applies an invariant mass cut on them. The VBF-trigger yields large acceptance gains for VBF analyses. The performance of the VBF-trigger is shown in Fig. 2 (right). This new algorithmic capability opens interesting prospects for the design of algorithms whose philosophy and performance approach those in the HLT.

References

[1] S. Chatrchyan *et al.* [CMS Collaboration], JINST **3**, S08004 (2008).

[2] V. Khachatryan *et al.* [CMS Collaboration], JINST **12**, no. 01, P01020 (2017).

[3] A. Tapper *et al.* [CMS Collaboration], CERN-LHCC-2013-011, CMS-TDR-12.

[4] A. Zabi *et al.*, JINST **12**, no. 01, C01065 (2017).

[5] A. Svetek *et al.*, JINST **11**, no. 02, C02011 (2016).

[6] G. Iles, J. Jones and A. Rose, JINST **8**, C12037 (2013).

[7] CMS Collaboration, https://twiki.cern.ch/twiki/bin/view/CMSPublic/L1TriggerDPGResults.

[8] J. Ero *et al.*, JINST **11**, no. 03, C03038 (2016).

[9] W. Zabolotny, A. Byszuk, JINST **11**, C03004 (2016)

[10] A. Madorsky, JINST **12**, no. 07, C07010 (2017).

[11] M. Jeitler *et al.*, IEEE Trans. Nucl. Sci. **62**, no. 3, 1104 (2015).