The Compact Muon Solenoid Experiment
Conference Report

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

15 December 2017

The CMS Level-1 Trigger system for LHC Run II

Alex Tapper for the CMS Collaboration

Abstract

During LHC Run II the centre-of-mass energy of pp collisions has increased up to 13 TeV and the instantaneous luminosity has progressed towards $2 \times 10^{34} \text{cm}^{-2}\text{s}^{-1}$. In order to guarantee a successful and ambitious physics programme under these conditions, the CMS trigger system has been upgraded. The upgraded CMS Level-1 trigger is designed to improve performance at high luminosity and large number of simultaneous inelastic collisions per crossing. The trigger design, implementation and commissioning are summarised and performance results are described.

Presented at *CHEF2017 Calorimetry for the High Energy Frontier 2017*

PREPARED FOR SUBMISSION TO JINST

CALORIMETRY FOR THE HIGH ENERGY FRONTIER

2-6 OCTOBER 2017

LYON, FRANCE.

The CMS Level-1 Trigger for LHC Run II

A. Tapper on behalf of the CMS collaboration

Imperial College London & The European Organization for Nuclear Research (CERN)

E-mail: a.tapper@imperial.ac.uk

ABSTRACT: During LHC Run II the centre-of-mass energy of pp collisions has increased up to 13 TeV and the instantaneous luminosity has progressed towards $2 \times 10^{34} \text{cm}^{-2}\text{s}^{-1}$. In order to guarantee a successful and ambitious physics programme under these conditions, the CMS trigger system has been upgraded. The upgraded CMS Level-1 trigger is designed to improve performance at high luminosity and large number of simultaneous inelastic collisions per crossing. The trigger design, implementation and commissioning are summarised and performance results are described.

KEYWORDS: Trigger algorithms, Trigger concepts and systems (hardware and software)

Contents

1	Introduction	1
2	Design, implementation and commissioning	1
3	Trigger algorithms and performance	2
3.1	e/γ finder	2
3.2	τ finder	3
3.3	Jet finder and energy sums	4
3.4	Invariant mass	5
4	Summary	5

1 Introduction

At the LHC, production cross sections for the most interesting processes are many orders of magnitude lower than that of the total pp cross section. Since the experiments cannot store all data generated in pp collisions, the experiments must select interesting events online. This is the function of the trigger systems of the experiments. In the CMS experiment this is implemented as a two-level system, with the first level (Level-1) based on custom hardware and with a software-based second level running on a large farm of computers.

With the start of LHC Run II, trigger rates have increased due to the increase in luminosity, the increase in centre-of-mass energy, and by the higher pile-up (the number of interactions per bunch crossing). The CMS detector electronics limit the Level-1 (L1) trigger rate to 100 kHz. The CMS physics programme requires the sensitivity for electroweak scale physics and for TeV scale searches be maintained at the level it was in Run I. This requirement motivated the upgrade to the CMS Level-1 trigger system, described in this paper, which focuses on the calorimeter trigger, given the topic of the conference.

2 Design, implementation and commissioning

The legacy CMS Level-1 trigger is described in detail elsewhere [1, 2]. The overall aim of the upgrade was to implement a flexible, maintainable system, capable of being adapted to the evolving CMS physics programme. Figure 1 shows schematics of the overall Level-1 trigger system, and the calorimeter trigger system in more detail. The key conceptual change planned to the design of the Level-1 calorimeter trigger was the removal of boundaries by streaming data from single events into one FPGA. The global trigger was also expanded to allow many more possible conditions and more sophisticated quantities, to give a richer menu of possible trigger conditions [3].

Figure 1. Schematics of the upgraded L1 trigger, showing the complete trigger system (left) and the calorimeter trigger in detail (right).

The key technology changes may be summarised as follows: legacy VME-based electronics were replaced with microTCA, a modern telecoms standard, making the electronics more maintainable and compact in many cases; the upgraded system uses the latest generation of FPGAs, Xilinx Virtex 7, in a reduced number of common electronics cards; parallel copper links were replaced in almost all cases with serial optical links, which improves maintainability, and allows link speeds to increase from 1 Gb/s to 10 Gb/s; large optical patch panels with single fibres were replaced with custom-made commercial solutions (Molex Flexplane technology) which are more compact and maintainable; finally, the online software was rewritten to have more common code, modern libraries, and to be more easily maintained.

In order to reduce the potential risk to CMS data taking the upgrade was commissioned in parallel with the legacy trigger data taking in 2015. Trigger inputs from the calorimeters were duplicated (in FPGAs and optically). The new system was run parasitically with CMS data taking, recording results but not issuing triggers towards the end of 2015, and became the primary trigger for CMS in 2016. This paper reports the performance of the trigger during the early part of the 2017 run, where the level of pile-up was significantly higher than in 2016.

3 Trigger algorithms and performance

3.1 e/γ finder

The upgraded e/γ finding algorithm [4, 5] brings a number of improvements over the legacy algorithm. In particular the new algorithm is based on dynamic clustering, allowing the dimensions of the energy clustered to form the e/γ candidate to adapt to the pattern of energy deposited in the electromagnetic calorimeter. This gives improved energy containment for showering electrons and photon conversions without being affected significantly by pile-up. Clusters produced using this algorithm have better energy resolution than those formed by the legacy algorithm.

In addition the shape of the clusters may be used, along with the fraction of the total energy measured as electromagnetic, to discriminate between real e/γ candidates and background from

Figure 2. Trigger efficiency for the single isolated e/γ trigger as a function of offline E_T (left) and number of reconstructed vertices (right).

hadronic jets. The clusters also provide an energy-weighted position measurement which may be used to correlate objects, for example in invariant mass calculations.

Isolation annuli (removing the footprint of the e/γ candidate) are constructed around each candidate cluster and are used to implement an isolation energy requirement, as a function of cluster E_T , pseudorapidity and the pile-up estimated event-by-event using central detector occupancy.

Figure 2 (left) shows the measured efficiency for a single e/γ trigger with $E_T > 38$ GeV as a function of the offline E_T , for 2016 and 2017 tunings and (right) the efficiency as a function of the number of vertices reconstructed in the event, which is a good proxy for the pile-up level [6]. The efficiency was measured using the tag-and-probe technique on a sample of $Z \rightarrow ee$ events and shows good efficiency with a sharp response, due to the excellent energy resolution and is insensitive to pile-up. The tuning for 2017 improved the trigger efficiency despite the higher pile-up compared with 2016.

3.2 τ finder

The upgraded τ -lepton finding algorithm [4, 5] is based on the same dynamic clustering as the e/γ finder described above, optimised for τ leptons. In addition neighbouring clusters may be merged to recover multi-prong τ decays. An isolation requirement may also be placed on the τ candidates, which includes whether two clusters were merged as input.

Figure 3 (left) shows the efficiency for single τ triggers with $E_T > 26, 30$ and 34 GeV, which are representative thresholds used in di- τ triggers for 2016 and 2017, as a function of the offline E_T [7]. Figure 3 (right) shows the efficiency for a single isolated τ trigger with $E_T > 30$ GeV, as a function of the offline E_T , for the two isolation tunings used in 2016, showing slightly lower efficiency, due to a tighter isolation in the latter part of the year, when the luminosity and pile-up were higher. The efficiencies were measured using the tag-and-probe technique on a sample of $Z \rightarrow \tau\tau$ events, where one τ decays hadronically and one to a muon, and shows good efficiency with a sharp response, due to the excellent energy resolution.

Figure 3. Trigger efficiency for single τ (left) and single isolated τ (right) triggers as a function of offline E_T .

Figure 4. Efficiencies for central (left) and forward (right) single jet triggers, as a function of offline E_T .

3.3 Jet finder and energy sums

The jet finding algorithm [4, 5] benefits from the increased spatial granularity of inputs compared to the legacy trigger. The algorithm is a sliding window based sum, with dimensions matched to the $R=0.4$ jet finding used offline by CMS. A jet is formed by searching for a seed energy deposit above a chosen threshold and summing the energy deposits around the seed. A veto mask is applied to remove duplicate jets and areas outside the size of typical jets are used to estimate and subtract the energy within the jet from pile-up. Jet energies are calibrated as a function of jet E_T and η .

Figure 4 shows the efficiencies of the Level-1 jet trigger for central (left) and forward (right) pseudorapidities, as a function of offline jet E_T , using a data sample triggered with an independent, single muon, trigger [8]. Level-1 and offline jets are matched geometrically. The thresholds shown are representative of those used in various conditions in 2016 and 2017. The results show a sharp efficiency turn-on with a well calibrated E_T scale.

Figure 5 (left) shows the efficiency for the missing transverse momentum (MET) trigger, calculated as the vector sum of all energy deposits and evaluated using the same sample as the jet trigger [8]. The higher thresholds are representative of those used in the 2017 run. Figure 5 (right)

Figure 5. Efficiencies for missing transverse momentum triggers, for different L1 thresholds (left), and illustrating the gain from pile-up mitigation (right).

illustrates the pile-up mitigation scheme used in 2017, in which an event-by-event zero-suppression was applied to energy deposits entering the sum, which was a function of pseudorapidity and the pile-up estimated in each event. As an example, the application of this scheme reduces the rate of the missing transverse energy trigger with a threshold of 125 GeV by a factor of two, allowing the threshold to be decreased to 110 GeV and also decreases the pile-up dependence of this trigger quantity.

3.4 Invariant mass

The extended capabilities of the upgraded L1 global trigger allow both a higher number of trigger conditions and more sophisticated quantities to be calculated. One example of such an algorithm is the calculation of invariant mass, which is used in a dedicated trigger for the vector-boson fusion (VBF) production of Higgs bosons, with subsequent decay to τ -lepton pairs [7]. This process results in a final state with two low- E_T jets, separated by a large rapidity gap, and a central high- E_T pair of τ leptons from the Higgs decay. Using a threshold on the invariant mass of the jet pair allows the jet- E_T thresholds to be kept low, since the rapidity gap results in a large invariant mass, even for events in which the jets are low E_T .

Figure 6 illustrates the complementarity of the VBF trigger with respect to the di- τ trigger, yielding a significant increase in acceptance for this key Higgs process.

4 Summary

Run II at the LHC is a very challenging environment to search for new physics and measure the properties of the Higgs boson. The newly installed CMS Level-1 trigger upgrade tackles these challenges and has run reliably, providing CMS with high-quality physics data since the start of the 2016 run. The system is based on state-of-the-art, FPGA based, very high bandwidth processors with sophisticated, programmable algorithms. Performance results from the 2016 and early 2017 runs have been documented and demonstrate good performance. The design of the trigger allows flexibility to evolve with the requirements of the CMS physics programme and will also provide lessons for the design of the trigger upgrade planned for high luminosity LHC running.

Figure 6. Distributions of numbers of signal events in simulation, for events satisfying VBF and di- τ triggers, as a function of the offline sub-leading τ E_T (left) and VBF jet invariant mass (right). The increase in efficiency can be seen in both offline quantities.

References

- [1] CMS Collaboration, “The TriDAS Project Technical Design Report, Volume 1: The Trigger Systems”, CMS-TDR-006, CERN-LHCC-2000-038 (2000). <https://cds.cern.ch/record/706847>
- [2] CMS Collaboration, “The CMS trigger system”, JINST 12 (2017) P01020.
- [3] CMS Collaboration, “CMS Technical Design Report for the Level-1 Trigger Upgrade”, CMS-TDR-012, CERN-LHCC-2013-011 (2013). <https://cds.cern.ch/record/1556311>
- [4] A. Zabi et al., “Triggering on electrons, jets and tau leptons with the CMS upgraded calorimeter trigger for the LHC RUN II”, JINST 11 (2016) 02, C02008.
- [5] B. Kreis et al., “Run 2 Upgrades to the CMS Level-1 Calorimeter Trigger”, JINST 11 (2016) C01051.
- [6] CMS Collaboration, “Level-1 E/Gamma and Muon performance on 2017 data”, CERN-CMS-DP-2017-024 (2017). <https://cds.cern.ch/record/2273270>
- [7] CMS Collaboration, “Level 1 Tau trigger performance in 2016 data and VBF seeds at Level 1 trigger”, CERN-CMS-DP-2017-022 (2017). <https://cds.cern.ch/record/2273268>
- [8] CMS Collaboration, “Level-1 jets and energy sums trigger performance with part of the 2017 dataset”, CERN-CMS-DP-2017-040 (2017). <https://cds.cern.ch/record/2286149>