

Operation of the CMS Level-1 Calorimeter Trigger in High Pileup Conditions and Motivations for Phase-2

A. Bundock^{*†}

Imperial College London, UK

E-mail: aaron.bundock@cern.ch

To maintain high trigger efficiencies and stable rates during significant changes to beam conditions throughout 2017, the CMS Level-1 calorimeter trigger required dynamic and flexible operation. Successfully running since 2015, utilising Xilinx Virtex 7 690 FPGAs and 10 Gbps optical links, the versatile design has enabled quick adaption to improve algorithms to mitigate large rates from high pileup and changes in detector response, and as the LHC responded to a number of unexpected challenges. Operational experience and lessons learned will be discussed, and how they will inform important decisions in the design and implementation of the Phase-2 trigger upgrade.

Topical Workshop on Electronics for Particle Physics (TWEPP2018)

17-21 September 2018

Antwerp, Belgium

^{*}Speaker.

[†]On behalf of the CMS Collaboration

1. Introduction

In 2017, the LHC delivered an instantaneous luminosity of $2 \times 10^{34} \text{ cm}^{-2} \text{ s}^{-1}$, double the nominal design performance, which generated a number of challenges, both for the LHC to maintain stable beam conditions, and for CMS to run successfully with highly efficient triggers sensitive to a wide range of interesting physics signatures, and keep within the 100 kHz Level-1 (L1) trigger bandwidth. The ability of CMS to adapt to changing beam conditions was significantly enhanced by using 0.94 Tbps input + 0.94 Tbps output bandwidth optical processor boards with large Xilinx Virtex-7 FPGAs. These provide substantial logic resources with a flexible architecture that can be quickly utilised to add new features and develop existing algorithms as conditions change.

A number of challenges arose during 2017 data taking. An increase in the number of protons per bunch crossing and reduced beam emittance raised considerably the number of pileup interactions, which revealed a strong pileup dependence for the rate of a number of important L1 trigger algorithms. A problem developed in one of the LHC sectors that made it difficult to maintain stable beam conditions with nominal long bunch trains, and to allow operation at full luminosity required switching to smaller bunch trains, which induced large variations in the level of out-of-time pileup, resulting in higher trigger rates. The high luminosity also resulted in increased levels of radiation damage to detectors, particularly the forward regions of the calorimeters. Each of these challenges were met in various ways by adapting the L1 trigger algorithms, and provide useful experience to inform the development of the Phase-2 CMS L1 trigger upgrade for the HL-LHC.

2. The CMS L1 Calorimeter Trigger

The Phase-1 upgrade of the L1 calorimeter trigger has adopted a novel time-multiplexed architecture consisting of two layers, with higher granularity calorimeter data for the full detector being processed in a single main processor (MP), providing extra latency for more sophisticated algorithms. Layer 1 consists of 18 CTP7 pre-processor boards that receive trigger primitives optically from regions of ECAL and HCAL, sum the transverse energies and further format the data ready for processing at Layer 2, which consists of 10 MP7 boards that perform the object identification and compute global energy sums (Figure 1). Both boards are built around high performance Xilinx Virtex 7 FPGAs, which provide the large IO and logic resources that are required [1]. The MPs receive data on 72 links, each link containing 2 out of 72 ϕ segments with alternate links carrying data from the positive and negative η -halves of the calorimeters. A Molex Flexplane patch panel provides the interface between Layer 1 and Layer 2, fanning out each pre-processor for a given detector region to each MP.

The data protocol used between Layer 1 and Layer 2 is a 16-bit word comprising the sum of ECAL and HCAL transverse energies, their ratio, and various feature flags, allowing the separate calorimeter energies to be computed. Each link carries two of these words, with the full 32-bit word computed at a clock frequency of 240 MHz. The MP identifies up to 12 jet, e/γ , and τ candidates, and calculates the total transverse energy and jet energy sums, along with their x and y components. The transverse momenta and $\eta - \phi$ co-ordinates of the particle candidates and the energies of the sums are sent from each MP over 6 links to the demux board (also an MP7), which calculates the vector sums and forwards data to the global trigger.

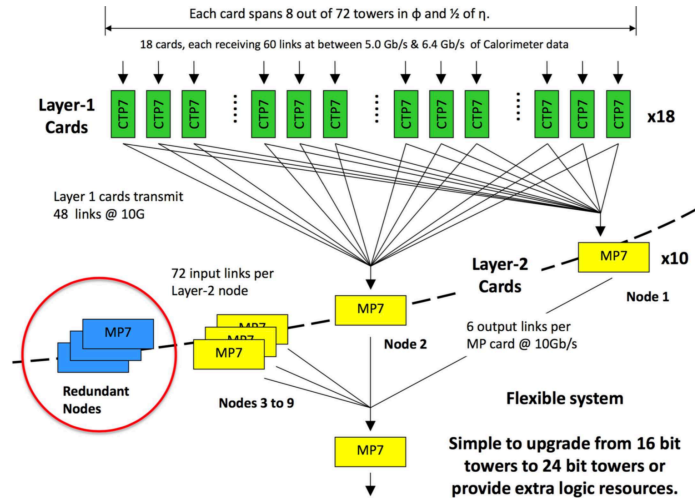


Figure 1: The CMS L1 calorimeter trigger architecture.

3. Mitigating Pileup Dependence

One of the challenges of taking high quality data at a hadron collider is reducing as much as possible the effect of many simultaneous low energy pileup collisions superimposed on top of the hard scatter process that needs to be measured. In the high level trigger and offline, this is facilitated using information from the tracker, however at L1 this is not yet available. Peak average pileup increased from ~ 45 in 2016 to ~ 55 in 2017, revealing the strong pileup dependence of rates for missing transverse energy, jet sum and low threshold multi-jet algorithms, with some rates doubling with an increase of ~ 5 in pileup. Maintaining L1 trigger performance required the implementation of additional pileup mitigation. The use of large Virtex 7 FPGAs in the calorimeter trigger processor boards provides capacity for many run-time configurable LUTs, that are used mainly for calibration and algorithm optimisation. A bus consisting of clock and data lines receives an instruction via the online trigger software and programs the LUT contents from the configuration key into the FPGA. This provides a very flexible interface that allows algorithms to be tuned and adapted quickly as LHC conditions change.

Pileup interactions typically produce low energy deposits in the calorimeters, and so to first order one of the simplest pileup correction methods is to reduce contributions to the L1 algorithm from low energy trigger towers. For global sums such as missing transverse energy however, this can significantly affect efficiency, thus low energy towers should only be omitted from the sum calculation for events with high pileup, when performance will be improved and not reduced. The flexibility of the firmware allowed for a LUT to be implemented that returns an energy threshold only above which trigger towers are included in the algorithm. The LUT takes as input the η of the tower and an estimation of the pileup in the event calculated from the trigger tower occupancy in the central region of the calorimeter. This LUT was derived and tuned using Zero Bias and Single Muon events, for determining the trigger rate and efficiency respectively. The improvement in performance for 2017 data can be seen in Figure 2, which shows a gain of ~ 10 GeV at 95% efficiency for the same rate, increasing sensitivity to a range of interesting physics signals.




Figure 2: L1 missing transverse energy trigger efficiencies in 2017 Single Muon data, for L1 thresholds that provide the same trigger rate with (circles) and without (triangles) pileup mitigation [2].

In 2017, one of the LHC sectors caused frequent beam dumps due to electron clouds produced by the beam interacting with gas contamination around a pumping port. To prevent this, the filling scheme was changed from 48 to 8 bunch trains, which further increased the pileup dependence of the rates due to the calorimeter front end dynamic pedestal subtraction not being as effective for shorter bunches. The ability to easily deploy a re-tuned pileup mitigation LUT provided a mechanism for the trigger to quickly adapt, allowing L1 thresholds to be maintained despite difficult beam conditions.

4. Responding to Detector Ageing

The high luminosity delivered by the LHC has resulted in significant levels of radiation damage to the detectors. For ECAL, by the end of 2017 the forward-most PbWO4 crystals required amplification factors of up to 50 due to loss of transparency (Figure 3), bringing noise levels up to an RMS of ~ 1 -2 GeV, above the 0.5 GeV minimum energy for trigger primitives to be included in the calorimeter trigger algorithms. To mitigate this effect in 2018 required adaption of the calibration scheme to progressively zero suppress ECAL trigger primitives as a function of η using programmable LUTs.

Particle-level response calibrations are applied to the ECAL and HCAL trigger primitives in the Layer 1 pre-processor boards using scale factors derived from simulation and applied differentially in E_T and η . The granularity of these calibrations was increased at low E_T so that ECAL trigger primitives below a chosen threshold at high η could be calibrated to zero. These thresholds were determined by optimising the efficiencies of calorimeter trigger algorithms at fixed L1 rate using 2017 data. The performance for forward jets and missing E_T was improved significantly. These zero suppression thresholds can be easily updated as the detector response evolves to keep trigger rates low and maintain performance.

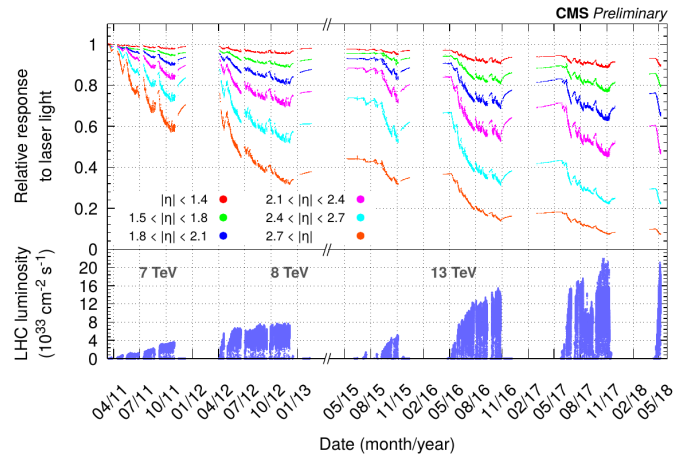


Figure 3: Relative response to laser light in ECAL, averaged over crystals in bins of η , from 2011-2018 [3].

5. Implications for Phase-2

The ability of the calorimeter trigger to adapt to the various changes in operating conditions mentioned in this report has allowed the CMS experiment to continue taking high quality data throughout Run 2. The experience gained will help achieve the same quality of data taking throughout Run 3 and provide useful feedback for the design and implementation of the L1 trigger upgrade for Phase-2 [4]. The level of pileup expected at the HL-LHC is ~ 140 -200, creating an even more challenging environment in which to efficiently select interesting events for offline analysis.

With tracker data available to the L1 trigger [5], resilience to pileup will be vastly improved, by allowing vertex reconstruction and particle flow algorithms at L1. The upgrade of the CMS endcap calorimeter [6] will increase dramatically the lateral granularity and provide fine sampling of longitudinal shower development to give excellent separation of individual particle showers. Precision time measurement will help to reject energy from pileup, and the fast shaping time of the front-end electronics will minimise out-of-time pileup contributions. These detector upgrades combined with a flexible, highly configurable multi-layered L1 trigger implementing a wide variety of both simple and robust, and complex high-performance algorithms, comprised of fast optical, large logic capacity processors will enable CMS to trigger efficiently at the HL-LHC.

References

- [1] K. Compton et al., "The MP7 and CTP-6: multi-hundred Gbps processing boards for calorimeter trigger upgrades at CMS" 2012 JINST 7 C12024
- [2] CMS Collaboration, "L1 jets and sums performance with 2017 dataset" CERN-CMS-DP-2018-004
- [3] CMS Collaboration, "CMS ECAL with 2017 data" CERN-CMS-DP-2018-015
- [4] CMS Collaboration, "Phase-2 Upgrade of the CMS Level-1 Trigger" CERN-LHCC-2017-013
- [5] CMS Collaboration, "Phase-2 Upgrade of the CMS Tracker" CERN-LHCC-2017-009
- [6] CMS Collaboration, "Phase-2 Upgrade of the CMS Endcap Calorimeter" CERN-LHCC-2017-023