PROCEEDINGS OF SCIENCE

Architecture and Performance of the Upgraded CMS Calorimeter Trigger

Santeri Laurila[*†]
Helsinki Institute of Physics
E-mail: santeri.laurila@cern.ch

The CMS Level-1 calorimeter trigger was completely upgraded before the start of the 2016 data-taking. The upgraded system ran successfully during Run 2 and will continue running throughout Run 3. Its novel time-multiplexed architecture is made possible by large FPGAs connected by fast optical links, operating in MicroTCA chassis. We describe the architecture of the system and the hardware used to implement it. In particular, examples of advanced trigger algorithms enabled by the time-multiplexed design are discussed. Precise FPGA floorplanning allows placing all calorimeter trigger algorithms in a single board. The performance of the system is presented based on the 2016–2018 data taking of proton collisions at the LHC.

Topical Workshop on Electronics for Particle Physics TWEPP2019
2-6 September 2019
Santiago de Compostela - Spain

[*]Speaker.
[†]On behalf of the CMS Collaboration.

PoS(TWEPP2019)144

1. Introduction

The CMS experiment uses a two-level trigger system for online selection of collision events. The Level-1 trigger is a hardware system that selects events to be read out and sent to High-Level Trigger (HLT) at a maximum event rate of 100 kHz (out of the 40 MHz bunch crossing rate), with a latency of only 3.8 μs. The HLT is a software farm that selects $\approx$1 kHz of events for storage and offline analysis. The CMS experiment and its coordinate system are described in Ref. [1].

In order to achieve the best possible performance under the harsh luminosity conditions of Run 2, the CMS Level-1 calorimeter trigger was completed before the start of the 2016 data-taking [2]. This *Phase-1 upgrade* covered the hardware, firmware, the timing control system, and the trigger control software. The upgraded system collected data until the end of Run 2 in 2018, and will continue running throughout Run 3. Its novel time-multiplexed architecture is made possible by large Xilinx Virtex-7 690 FPGAs and fast optical links. Precise FPGA floorplanning allows placing of all calorimeter trigger algorithms in a single board. This enables precise evaluation of global event quantities and usage of efficient pileup mitigation techniques. For each event, the trigger algorithms reconstruct and identify *trigger objects* that correspond to electrons and photons (e/γ), hadronic tau leptons (τ_h), and hadronic jets. *Global energy sums* such as the missing transverse energy are also calculated. In the following, we describe the architecture of the system and the hardware used in it, and report its performance during the 2016–2018 proton-proton data taking at the LHC.

2. Architecture of the upgraded CMS calorimeter trigger

Both the Electromagnetic calorimeter (ECAL) and the Hadron calorimeter (HCAL) provide input to the calorimeter trigger. They have dedicated front-end electronics to shape the signal pulses, digitize them and sum them into *trigger towers* (TTs). In the barrel, the trigger towers have $\eta \times \phi$ granularity of approximately 0.087×0.087, corresponding to 5×5 ECAL crystals and one HCAL tower behind them. In the endcaps, the trigger tower size increases up to 0.17×0.17.

The two-layer architecture of the upgraded calorimeter trigger is illustrated in Figure 1. The first layer consists of 18 Calorimeter Trigger Processor 7 (CTP7) AMC cards [3], based on Virtex-7 FPGAs. The Layer-1 boards receive TTs from a specific geometrical region via 4.8 Gb/s or 6.4 Gb/s optical links, and take care of TT-level pre-processing of the data. This includes energy calibration and calculation of the TT transverse energy (E_T) as well as additional *feature bits*, such as the H/E ratio (the ratio of HCAL and ECAL energy deposits). The TT information is compressed into 16-bit words for efficient transmission to Layer 2.

The data corresponding to one event are collected from all Layer-1 boards and sent to one of the Layer-2 boards via 10 Gb/s optical links. Layer 2 consists of nine Xilinx Virtex-7-based Master Processor 7 (MP7) boards [3], placed in a single μTCA crate. The second layer works in a *time-multiplexed* manner: each Layer-2 board has access to a full event, and uses information from the all TTs to evaluate the event. The time-multiplexed approach is used in CMS also in DAQ and HLT systems, but the upgraded calorimeter trigger system is its first application in L1T.

The Layer-2 boards host the *trigger algorithms*, described in Ref. [4], reconstructing the trigger objects and energy sums. Precise FPGA floorplanning allows all calorimeter trigger algorithms for

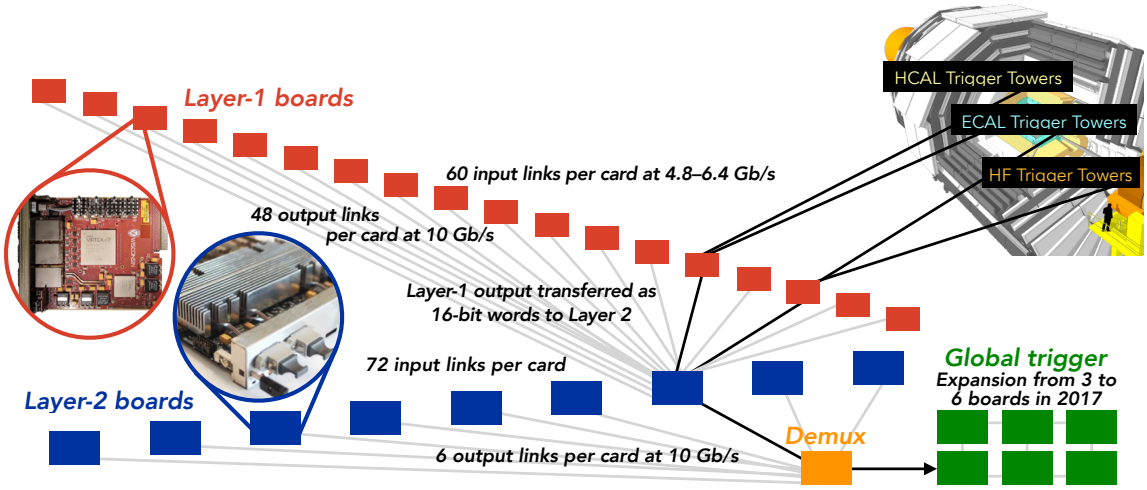

Figure 1: The architecture of the upgraded calorimeter trigger system. In Layer 1, 18 CTP7 boards pre-process the trigger tower information regionally and send it via fast optical links to Layer 2, which contains 9 MP7 boards. TheLayer 2 functions in a time-multiplexed way, processing each full event in one of the nine boards. Finally, a "demux" card de-multiplexes the output for the global trigger. Examplary data flows from two regions are shown as black lines, and the other links as grey lines.

different objects to run in a single MP7 board. Finally, a *de-multiplexing* MP7 board collects the data from a Layer-2 node, serializes it, and sends the 12 highest-energy objects of each type and the energy sums to the global trigger.

Upon arrival to the global trigger, these input data are synchronized to the LHC clock. The global trigger, also based on MP7 boards, combines the calorimeter information received from the de-multiplexing node with the muon trigger outputs, and performs the final *L1 Accept* decision based on the active trigger menu.

In 2017, the global trigger was expanded from 3 to 6 boards. The larger amount of logic resources in the global trigger allows implementation of complex multi-object trigger paths targeting specific event topologies, previously possible only at HLT level or offline. For example, a specific algorithm that targets Higgs boson production via vector boson fusion has been developed [5].

3. Trigger algorithms and their performance

The jet reconstruction is seeded by a TT that exceeds a pre-programmable E_T threshold of 4 GeV. The jet energy is determined by summing the TT energies in a 9×9 TT window centered around the seed. The effect of pileup is estimated with a "chunky donut" algorithm [4]. The jets are calibrated as a function of the transverse momentum (p_{T}) and η using a dedicated look-up table. Efficiency curves for a set of single-jet algorithms used in 2018 are shown in Fig. 2 (left).

The energy sums calculated at Level-1 include the trigger tower scalar sum E_{T}, the jet energy scalar sum H_{T}, and the missing E_{T} (MET). The Level-1 MET is reconstructed as a vector sum of all TT energies that exceed predefined thresholds. For optimized pileup mitigation, the thresholds depend on η and the TT size, and on the amount of pileup, estimated from the number of active TTs. Efficiency curves for the MET trigger algorithms used in 2018 are shown in Figure 2 (right).

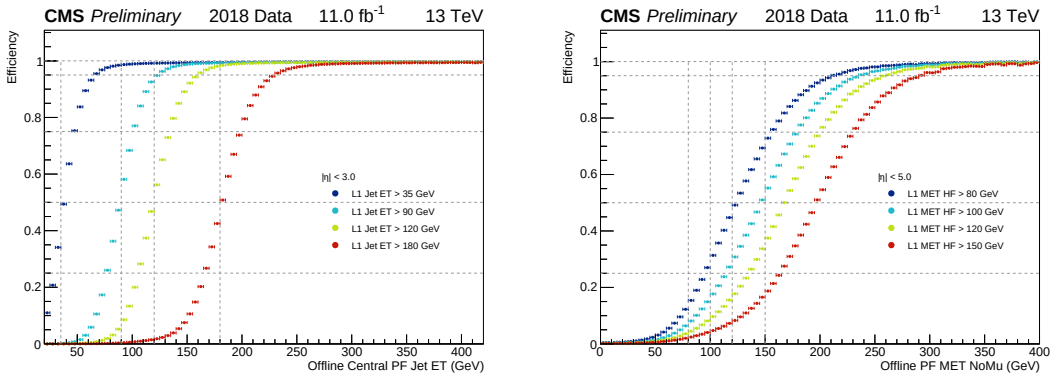

Figure 2: Left: Efficiency curves for the Level-1 single jet trigger algorithms with different E_{T} thresholds, as a function of offline-reconstructed anti-k_{T} transverse energy. Right: Efficiency curves for the upgraded Level-1 MET trigger algorithms with different MET thresholds, as a function of offline-reconstructed type-1 corrected MET with muons excluded. In both cases, the efficiency is measured using an unbiased sample selected with a single-muon trigger. [6]

The e/γ candidates are identified using *dynamic clustering*, seeded by a TT with $E_{\mathrm{T}} > 2\,\mathrm{GeV}$. The clusters are built iteratively, adding neighboring TTs with $E_{\mathrm{T}} > 1\,\mathrm{GeV}$ up to a maximum cluster size of 8 TTs. The clusters are categorized based on the shape of their energy distribution, used to distinguish between e/γ candidates and hadronic jets. Efficiency curves for a set of single-e/γ algorithms used in 2016–2018 are shown in Fig. 3 (left). The improvement of efficiency in 2018 compared to the previous years reflects the accumulated experience in optimization of the isolation criteria.

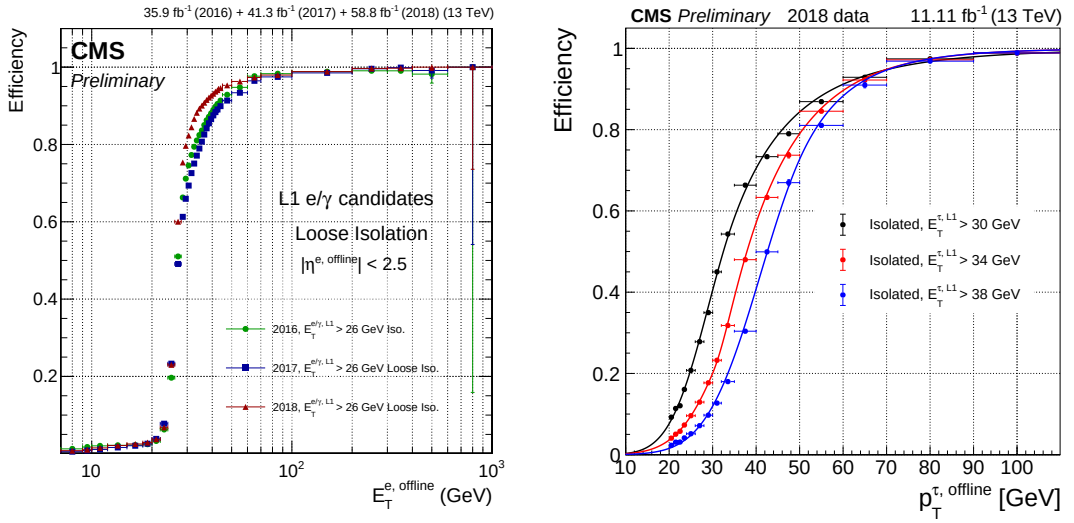

Figure 3: Left: Efficiency curves for the Level-1 single-e/γ trigger algorithms with different E_{T} thresholds in 2016–2018, requiring a loose isolation of the e/γ candidate, as a function of offline-reconstructed electron supercluster transverse energy. Right: Efficiency curves for the Level-1 single-τ_{h} trigger algorithms with different E_{T} thresholds in 2018, requiring isolation of the τ_{h} candidate, as a function of offline-reconstructed τ_{h} p_{T}. [6]

The τ_h candidates are identified using a dedicated algorithm, similar to the e/γ algorithm. Since the fingerprint of a τ_h decay in the calorimeters can correspond to several e/γ clusters, the τ_h identification algorithm is allowed to merge together several energy clusters if certain proximity conditions are fulfilled [7]. An efficiency curve for a set of single-e/γ algorithms used in 2018 are shown in Fig. 3 (right).

For e/γ and τ_h candidates, the position is determined as an energy-weighted average of the TT positions, yielding a factor of four improvement in spatial resolution compared to Run-1 algorithms that used the position of the seed TT.

The isolation of e/γ and τ_h candidates is defined in a 6×9 TT window centered around the candidate, and the reconstructed candidate energy is subtracted from the sum. Look-up tables are utilized for defining pileup-dependent isolation criteria, and for the energy calibration of the candidates.

4. Conclusions

The upgraded CMS Level-1 calorimeter trigger maintained excellent performance throughout Run 2, as illustrated in the efficiency curves shown above. The trigger worked in a robust way despite changes in the LHC filling schemes and ageing of the calorimeter systems. A more comprehensive collection of performance measurement results can be found in Ref. [6]. The experience gained in the implementation and operation of complex electronics systems for triggering, based on large FPGAs and fast optical links, will be utilized in the upcoming *Phase-2 upgrade* of the trigger system for the HL-LHC [8].

References

[1] CMS Collaboration, *The CMS experiment at the CERN LHC*, *JINST* **3** (2008) S08004.

[2] CMS Collaboration, *CMS technical design report for the Level-1 Trigger upgrade*, CMS Technical Design Report CERN-LHCC-2013-011, CMS-TDR-012, 2013.

[3] K. Compton et al., *The MP7 and CTP-6: Multi-hundred Gbps processing boards for calorimeter trigger upgrades at CMS*, *JINST* **7** (2012) C12024.

[4] CMS Collaboration, *Triggering on electrons, jets and tau leptons with the cms upgraded calorimeter trigger for the LHC Run II*, *JINST* **11** (2016) C02008.

[5] CMS Collaboration, *The CMS Level-1 tau lepton and vector boson fusion triggers for the LHC Run II*, *PoS* **EPS-HEP2017** (2017) 773.

[6] CMS Collaboration, *Level-1 Calorimeter Trigger Performance*, CMS Detector Performance Summary CMS-DP-2018-040, 2018.

[7] CMS Collaboration, *Tau lepton trigger and identification at cms in Run-2*, *Nucl. Part. Phys. Proc.* **287-288** (2017) 107.

[8] CMS Collaboration, *The Phase-2 Upgrade of the CMS L1 Trigger Interim Technical Design Report*, Tech. Rep. CERN-LHCC-2017-013. CMS-TDR-017, 2017.