

New development in the CMS ECAL Level-1 trigger system to meet the challenges of LHC Run 2

F.Thiant¹, Y.Geerebaert, F.Magniette, T.Romanteau, A.Zabi, A.Zghiche,
on behalf of the CMS Collaboration

Laboratoire Leprince-Ringuet - CNRS

École Polytechnique, 91128 Palaiseau, FRANCE

E-mail: thiant@llr.in2p3.fr, geerebaert@llr.in2p3.fr,
magniette@llr.in2p3.fr, romanteau@llr.in2p3.fr, zabi@llr.in2p3.fr,
zghiche@llr.in2p3.fr

Abstract - The CMS Electromagnetic Calorimeter (ECAL) provides energy sums to the Level-1 Calorimeter Trigger at a rate of 40 MHz. The processing of these trigger primitives (TPs) is performed by dedicated trigger concentrator cards (TCCs) located in the CMS service cavern. Updates to the functionality of the TCCs were required to respond to the challenging experimental conditions of LHC Run 2, where the center-of-mass of proton-proton collision energy was 13 TeV and the peak instantaneous luminosity of the proton beams reached $2 \times 10^{34} \text{ cm}^{-2} \text{ s}^{-1}$. A new algorithm, termed the Cumulative Overflow Killing Engine (COKE), has been developed and implemented via software and firmware updates to the TCCs in order to automatically detect and mask noisy or problematic TPs via configurable thresholds. The auto-recovery of the TCCs has also been improved, to manage the Single Event Upsets (SEUs) from the front-end electronics. This allows the detector to trigger efficiently without direct expert intervention, and the thresholds can evolve with evolving LHC conditions.

Topical Workshop on Electronics for Particle Physics (TWEPP2018)
17-21 September 2018
Antwerp, Belgium

¹Speaker

© Copyright owned by the author(s) under the terms of the Creative Commons Attribution-NonCommercial-NoDerivatives 4.0 International License (CC BY-NC-ND 4.0).

1. Introduction

The CMS Electromagnetic Calorimeter (ECAL) provides energy sums to the Level-1 Calorimeter Trigger at 40 MHz. The processing of these trigger primitives (TPs) is performed by dedicated trigger concentrator cards (TCCs) located in the CMS service cavern. The ECAL TCCs transmit approximately 4000 TPs every LHC bunch crossing to the Level-1 trigger to form electron/photon, tau and jet candidates, as well as calorimeter energy sums. The functionality of the TCCs must be updated to respond to the challenging experimental conditions of LHC Run 2, in which the center-of-mass energy of proton-proton collisions was 13 TeV and the peak instantaneous luminosity of the proton beams reached $2 \times 10^{34} \text{ cm}^{-2} \text{ s}^{-1}$. New features have been added during Run 2, firstly to mitigate the effect of noisy signals from the ECAL on-detector electronics, which would otherwise affect the triggering of CMS at high luminosities, and secondly to improve the auto-recovery from Single Event Upsets (SEUs).

A key element of this is a new algorithm to automatically detect and mask noisy or problematic TPs via configurable thresholds. This allows the detector to trigger efficiently without expert intervention, with thresholds that can be tuned to cope with evolving LHC conditions. This algorithm, termed the Cumulative Overflow Killing Engine (COKE), has been developed and implemented via software and firmware updates to the TCCs.

In this report, descriptions of the CMS Level-1 Trigger and readout systems are given. The algorithms used to form the ECAL Trigger Primitives are described, along with the new methods developed to monitor and automatically mask noisy or problematic trigger towers, and to detect and automatically recover from SEUs. A quantification of the gain in operating efficiency during LHC collisions during Run 2 is also provided.

2. The CMS Level-1 Trigger System

The CMS trigger system [1-2] is designed to perform filtering of events from LHC collisions in two steps. First, the hardware-based Level-1 (L1) trigger exploits coarse transverse energy (E_T) deposits in the calorimeters and signals in the muon systems to reduce the rate from 40 MHz to about 100 kHz. Secondly, the software-based High-Level Trigger (HLT), that runs selection algorithms based on finer granularity and higher resolution information from all sub-detectors, further reduces the rate to 1 kHz, providing an overall reduction factor of $O(10^5)$.

The Level-1 Calorimeter Trigger consists of two stages. In the first layer (L1-CaloLayer1), trigger primitives from the CMS ECAL and Hadronic Calorimeter (HCAL) are transmitted to 18 Calorimeter Trigger Processor boards (CTP7). A Time-Multiplexed-Trigger architecture (TMT) is used to send the pre-processed and formatted data to 9 Layer-2 (L1-CaloLayer2) Master Processor cards (MP7) which run algorithms to form calorimeter trigger objects: electrons, photons, τ leptons, jets, and energy sums. The L1 Calorimeter Trigger and L1 Muon Trigger objects are sent to the Global Trigger (μ GT), where the final L1 trigger decisions are taken. The L1 Calorimeter Trigger system uses high speed optical links (10Gb/s), and the CTP7 and MP7 boards use the μ TCA telecom standard and are equipped with the latest generation of Xilinx Virtex-7 FPGAs.

The ECAL TPs are sent from the TCCs to the L1-CaloLayer1 via a customized mezzanine, the Optical Synchronization and Link Boards (oSLBs), and 4.8 Gbps optical links. The

distribution of fast signals (L1-Accept, data acquisition (DAQ) start, DAQ stop, etc.) to all the sub-systems is managed by the Timing and Control Distribution System (TCDS), and the Clock and Control System (CCS) relays the broadcasting of these signals to the ECAL on-detector and off-detector boards [4].

3. ECAL TPs, Level-1 Trigger and readout system

The ECAL trigger primitives are generated by the on-detector front-end boards and sent to the off-detector TCC boards at a rate of 40MHz, via 800 Mbps serial optical links. Each trigger primitive in the ECAL barrel (EB) represents the transverse energy of 5x5 crystals, referred to as a Trigger Tower (TT). The ECAL trigger system comprises 3072 front-end and 108 TCC boards. Each ECAL barrel (EB) TCC receives 68 16-bit words, each representing the TT E_T encoded using 10 bits, with an additional 6-bits for internal TCC processing. The EB TCCs use a look-up table to transmit a compressed 8-bit TP E_T in addition to one bit (named the fine grain veto bit) for electromagnetic shower characterization. Similarly, ECAL Endcap (EE) TCCs receive up to 48 12-bit pseudo-strip E_T sums (of up to 5 crystals) and compute and transmit 8-bit TT E_T sums (of up to 5 pseudo-strips) [3]. The TCC outputs are transmitted to the L1-CaloLayer1, following TP time alignment via the oSLBs.

In addition, the TCCs provide the Selective Readout Processor (SRP) with Trigger Tower Flags (TTF), computed using two (high and low interest) E_T thresholds. Following the reception of a “L1 trigger accept” signal, the Data Concentrator Cards (DCC) read out the pipelined data from the front-end cards and the TTF values from the SRP determine the volume of data to be sent to the CMS DAQ. For a given tower TT, if E_T is below the low SR threshold (TTF = 0), only channels with an energy above the zero-suppression threshold are read out. If E_T is between the two SR thresholds (TTF = 1), all the channels in the corresponding readout unit (RU) are read out, while if E_T is higher than the high SR threshold (TTF = 3), all the channels of a 3x3 RU matrix, centered on the high interest tower, are read out. Using the TTF flags provided by the TCCs, the data volume is reduced by a factor of 20.

4. New developments for Run 2

4.1 Deployment of auto-masking algorithms

In LHC Run 1, the ECAL trigger system relied on human intervention whenever a problematic link or a noisy tower caused high trigger rates. This often required the stopping of a run, expert diagnosis and “on the fly masking” of the problematic regions. The whole procedure took 10 minutes at best. During LHC physics runs, this produced a loss of data taking efficiency, referred to as downtime. With increasing LHC luminosity, incidents such as this had to be mitigated with an automatic masking procedure. Two methods were developed, addressing different problems. First the “Link Swing Detector” (LSD) was developed in 2015 and integrated into the firmware of the TCCs. Its purpose was to detect defects that affect the 800 Mbps optical serial links used between the on-detector and off-detector electronics. A Cumulative Overflow Killing Engine (COKE) was also implemented in the EB TCCs to detect and mask noisy TPs. A noisy ECAL trigger channel creates high L1 trigger rates as well as dead time in the data acquisition, sometimes blocking data taking.

The LSD and COKE methods monitor and mask problematic links or noisy TTs in real time. COKE was completely validated and deployed in 2017. The masking is performed either by firmware or software and is periodically reset to restore the operation of transient problematic TTs. In 2017, a second version of COKE was developed to monitor and control the channels located in EE, in particular those close to the LHC beam axis, where increased noise is observed due to the accumulated irradiation of the EE crystals. This required the implementation of multiple thresholds that depend on pseudo-rapidity, η , in order to provide efficient suppression of problematic channels with no impact on physics signals.

4.2 Functionalities and implementation

Optical links transfer the transverse energies of a TT in EB, or a pseudo-strip in EE, from the front-end boards to the TCCs. Each link is associated with an LSD module that monitors the status bits of the deserializers inside the TCC, and to a COKE module that assesses the output data with individually customized programmable parameters. LSD errors are of two types: “StuckAtOne” and “SwingDetect”. The permanent error “StuckAtOne” is declared if the status bits reflect an unsynchronized link or link error state, and/or invalid data observed during a long programmable delay. If a programmable transition number on the “LinkError” status signal is reached during a sliding and programmable time window (between $1.6\mu\text{s}$ and $1636.8\mu\text{s}$) then LSD will state a “SwingDetect” error. The masking request is the logical OR of “SwingDetect” and “StuckAtOne”. The binary information is propagated to registers that concatenate the same information for all channels. In addition, this system allows the selective masking of a “StuckAtOne” event and includes the possibility of fast automatic masking by firmware.

In the first version of COKE, the deserialized data are passed through two algorithms, “SuccDetect” and “CumulDetect”, which use the same programmable E_T threshold value. The detectors operate at the LHC clock frequency (40MHz) during a sliding time window of 16 LHC clocks. The first algorithm checks whether the E_T threshold is reached in X successive clock cycles. The value of X is programmable from 2 to 8 clocks. The second algorithm keeps a cumulative total, Y , of the occasions during 16 clock cycles in which the E_T threshold is exceeded. Note that this value means a transition number where the threshold is reached (event) and not a number of clock cycles. The value of Y is programmable from 1 to 16. The second version of COKE incorporates four algorithms; two of which (Cumulative-1 and Successive-1 detection) are associated with the same programmable threshold (Thd1) value, the third (Occupancy detection) with another threshold (Thd2) and the last (Cumulative-3) with a threshold (Thd3). In this version, the main and important modification is the ability to mask a fault condition with a combination of the four algorithms using AND, OR, and NOT operators. For a given TT, one can differentiate between the presence of low energy noise, a real trigger, and a “high” energy noisy tower. Each of these cases is also characterized by its repeatability in the time window of 16 clock cycles. Finally, all the masking requests are inputs to a programmable lookup table (LUT), which delivers the final masking request, using the programmed settings of each trigger channel, such as the priority between competing requests. The masking inputs of the LUT are stored in a 4-bit vector and can be monitored online.

4.3 Auto-recovery procedure for TCC

The detection of a problematic optical link results in the zeroing of the transverse energy recorded by the tower or pseudo-strip, and the masking status is maintained until a control signal from the TCDS system (termed “BgoResync”) is received, or a specific instruction is executed in the TCC to remove the masks. The auto-masking decision is stored in the data stream for every TP using special TTF flags. In EB, masked towers are indicated by TTF flag 4. In EE, where the masking is performed per pseudo-strip, TTF values can be 4, 5 or 7, depending on whether the other TT channels have low, medium or high-interest data, respectively. A global monitoring of the TT masking occupancy is performed offline with CMS Data Quality Monitoring system using the TTF value.

Online TCC monitoring software provides more detailed reports of masking status, including the error type information (stuck, swing, noisy). The TCC monitoring is controlled by “ecalView”, a web application designed to monitor the status of the ECAL off-detector electronics boards [4]. Its main purpose is to report the occurrence of link errors, and check for pathological cases, where several neighboring towers are found to be in error. A failure pattern is formed by a minimum of three neighboring towers, including diagonals, and declared as a Single Event Upset (SEU). An auto-recovery procedure is initiated, the run is temporarily paused, the auto-masking suspended, links resynchronized, and link errors cleared. This procedure avoids running with significant parts of the detector being masked for long periods. Such on the fly recovery from an SEU, creates ~20-40 seconds of downtime during an LHC physics run, while a stopped state implies at least a downtime of ~3-5 minutes where the run is stopped, the system is reset and fully reconfigured before the data acquisition is started again.

5. Conclusion

The implementation of new automatic ECAL trigger tower masking and auto-recovery procedures was crucial for maintaining smooth ECAL operations during LHC Run 2. The operating efficiency was improved despite the much higher LHC beam intensities. The lost luminosity from ECAL trigger issues was reduced by a factor 10 between 2017 and 2018, making the ECAL trigger system one of the lowest sources of lost data in CMS during Run 2. The firmware and software improvements to automate fault detection, masking and recovery has proved its ability to improve the data-taking at CMS, and face the consequences of higher luminosity.

References

- [1] CMS Collaboration, “CMS Technical Design Report for the Level-1 Trigger Upgrade”, CERN-LHCC-2013-011 ; CMS-TDR-12
- [2] A.Zabi et al, “The CMS Level-1 Calorimeter Trigger for the LHC Run II”, JINST **12** C01065 (2017)
- [3] P.Paganini and the CMS collaboration 2009 J. Phys.: Conf. Ser. **160** 012062
- [4] P.Siddireddy for the CMS Collaboration, “The CMS ECAL Trigger and DAQ system: electronics auto-recovery and monitoring”, CMS CR-2018/084