Triggering on electrons, jets and tau leptons with the CMS upgraded calorimeter trigger for the LHC RUN II

TOPICAL WORKSHOP ON ELECTRONICS FOR PARTICLE PHYSICS 2015,
SEPTEMBER 28TH – OCTOBER 2ND, 2015
LISBON, PORTUGAL

Triggering on electrons, jets and tau leptons with the CMS upgraded calorimeter trigger for the LHC RUN II

A. Zabi,[a,1] F. Beaudette,[a] L. Cadamuro,[a] L. Mastrolorenzo,[a] T. Romanteau,[a] J.B. Sauvan,[a] T. Strebler,[a] J. Marrouche,[b] N. Wardle,[b] R. Aggleton,[c] F. Ball,[c] J. Brooke,[c] D. Newbold,[c] S. Paramesvaran,[c] D. Smith,[c] M. Baber,[d] A. Bundock,[d] M. Citron,[d] A. Elwood,[d] G. Hall,[d] G. Iles,[d] C. Laner,[d] B. Penning,[d] A. Rose,[d] A. Tapper,[d] T. Durkin,[e] K. Harder,[e] S. Harper,[e] C. Shepherd-Themistocleous,[e] A. Thea[e] and T. Williams[e]

[a]*Laboratoire Leprince-Ringuet, Ecole Polytechnique, IN2P3-CNRS,*
 Palaiseau, France
[b]*CERN, European Organization for Nuclear Research,*
 Geneva, Switzerland
[c]*University of Bristol,*
 Bristol, United Kingdom
[d]*Imperial College, University of London,*
 London, United Kingdom
[e]*Rutherford Appleton Laboratory,*
 Didcot, United Kingdom

 E-mail: Alexandre.Zabi@cern.ch

ABSTRACT: The Compact Muon Solenoid (CMS) experiment has implemented a sophisticated two-level online selection system that achieves a rejection factor of nearly 10^5. During Run II, the LHC will increase its centre-of-mass energy up to 13 TeV and progressively reach an instantaneous luminosity of 2×10^{34} cm^{-2} s^{-1}. In order to guarantee a successful and ambitious physics programme under this intense environment, the CMS Trigger and Data acquisition (DAQ) system has been upgraded. A novel concept for the L1 calorimeter trigger is introduced: the Time Multiplexed Trigger (TMT). In this design, nine main processors receive each all of the calorimeter data from an entire event provided by 18 preprocessors. This design is not different from that of the CMS DAQ and HLT systems. The advantage of the TMT architecture is that a global view and full granularity of the calorimeters can be exploited by sophisticated algorithms. The goal is to maintain the current thresholds for calorimeter objects and improve the performance for their selection. The performance of these algorithms will be demonstrated, both in terms of efficiency and rate reduction. The callenging aspects of the pile-up mitigation and firmware design will be presented.

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

Contents

1 Introduction

The search for new physics crucially relies on the performance of the trigger system used to select the most interesting collisions amongst the millions occurring per second [1]. The CMS trigger system is organised in two consecutive steps: the hardware-based Level-1 (L1) trigger utilises coarse energy deposits in the calorimeters and signals in the muon systems to reduce the rate from about 40 MHz to 100 kHz; this is followed by the software-based High Level Trigger (HLT), implementing selection algorithms based on finer granularity and higher resolution information from all sub-detectors in regions of interest identified at L1. The output rate of the HLT is about 1000 Hz. The CMS electromagnetic calorimeter (ECAL) provides a precise measurement of the energies and positions of incident electrons and photons for both triggering and offline analysis purposes. ECAL and HCAL energies are combined to reconstruct hadronically decaying τ leptons, particle jets and energy sums. Run I of the LHC (2010–2012) already reached an instantaneous luminosity of nearly $8 \times 10^{33}\,\mathrm{cm^{-2}\,s^{-1}}$ with p-p collisions at $\sqrt{s} = 8\,\mathrm{TeV}$, 50 ns bunch-spacing and up to about 40 pile-up events per bunch-crossing, almost double the design pile-up. Nevertheless, the trigger system performed extremely well. Run II started in Spring 2015 with p-p collisions at $\sqrt{s} = 13\,\mathrm{TeV}$, with 25 ns bunch-spacing. In 2016, the instantaneous luminosity is expected to reach up to $2 \times 10^{34}\,\mathrm{cm^{-2}\,s^{-1}}$ and the number of pile-up events may be up to 70 per bunch-crossing. To avoid a significant increase in triggering energy thresholds, which would be detrimental for physics, an upgrade of the L1 trigger system is required [2].

2 Upgrade of the Level-1 trigger system for Run II

As the LHC restarts and operates at higher luminosity, the current CMS trigger system will not be capable of maintaining the thresholds required for the CMS physics programme. For example, a

double-electron trigger, with thresholds of 13 GeV and 7 GeV in E_T for the two electrons respectively, had a L1 rate of 5 kHz in 2012; this would increase to about 50 kHz for the expected Run II conditions. A single electron trigger of 18 GeV threshold would give about 40 kHz, compared to 6 kHz during Run I. In these intense conditions, the implementation of pile-up mitigation techniques is required already at L1 to reach acceptable performance.

Modern technologies offer an effective solution to achieve these goals. The trigger primitives generated by the detector will be transmitted by newly installed high-speed optical links (4.8 to 6.4 Gb/s) replacing the existing copper cables (1.2 Gb/s), to a new system based on the μTCA electronics standard. The system is based on custom designed AMC (Advanced Mezzanine Card) with Xilinx Virtex-7 FPGAs. In the TMT approach [3], these FPGAs use 10 Gb/s transceivers to gather information from the entire calorimeter for each event in a single FPGA, where sophisticated algorithms may be implemented. The complete view of the calorimeter will allow the trigger to compute global quantities such as the average energy density that can be used to estimate the pile-up level. The full calorimeter granularity provided to the algorithms is used to enhance the energy and position resolution of Level-1 candidates. The upgraded calorimeter trigger architecture is described in [4].

3 Improved selection algorithms at Level-1

The algorithms described here have been designed to thoroughly exploit the global view of the calorimeters and the full trigger tower (TT) granularity provided by the upgraded trigger system. The main goal is to get as close to the offline selection performance as possible by introducing innovative reconstruction techniques at firmware level. The Level-1 trigger system is a synchronous electronics system and therefore all algorithms implemented in the electronics must be of fixed latency. This is in contrast with offline reconstruction algorithms which are typically iterative. Lepton signatures are reconstructed with a dynamic clustering approach instead of a sliding window while particle jets are reconstructed with optimum size. The improved response results in sharper efficiency curves and cutting on dedicated isolation variables helps controlling the rate. Along with better reconstruction, identification and isolation techniques, the pile-up energy must be evaluated as well. What is required for the calorimeter trigger is an estimate of the energy that should be subtracted from the measured energy for each calorimeter object and additionally a means to remove objects which originate from pile-up particles. For the Phase 1 upgrade no information from tracking detectors is available to assist in these tasks so they must be accomplished using only the calorimeter data. Several approaches are possible and have been explored. The selection techniques must be able to evolve with the CMS physics programme and be flexible enough to remain robust in any changing conditions. The e/γ, τ lepton and jet algorithms will be presented along with their performance.

3.1 Electron and photon trigger algorithm

The improved e/γ algorithm [5] is described in figure 1. Clusters are seeded by local maxima of energy above a fixed threshold. The maximum size of the clusters is limited (at most 8 trigger towers can be clustered) in order to minimize the impact of pile-up energy deposits while including most of the electron or photon energy. An extended region in the ϕ-direction is used to obtain a better
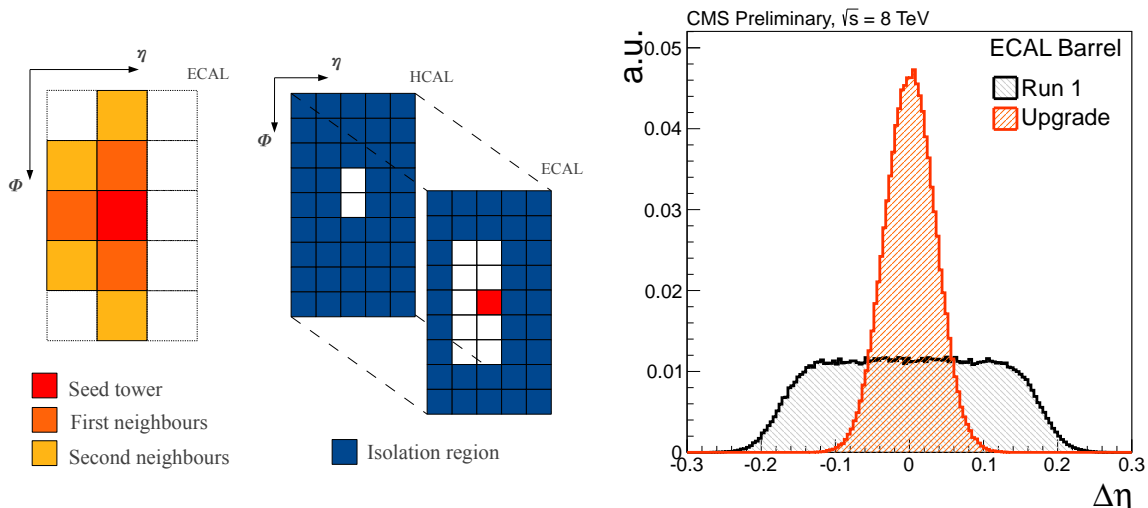
Figure 1. Left: the L1 e/γ clustering and isolation. A candidate is formed by clustering neighbour towers (orange and yellow) if they are linked to the seed tower (red). A candidate is considered as isolated if the E_T in the isolation region (blue) is smaller than a given value. The position resolution of the e/γ L1 trigger candidates in the ECAL barrel (right).

containment of the shower since electron and photon showers spread mostly along the ϕ-direction due to the magnetic field. Benefitting from the enhanced granularity, the e/γ candidate position can be computed as an energy-weighted average centered on the seed tower. As seen on figure 1, this results in a factor 4 improvement with respect to the Run I algorithm (centre of a 4×4 TT region). Additional background rejection is achieved by introducing a shape veto on the large variety of clusters produced by the dynamic clustering. The sum of the ECAL E_T of the seed and clustered towers is taken as the raw E_T of the cluster. A calibration derived from $Z \to ee$, is applied to this raw energy with factors depending on the η-position of the seed tower, the shape and the cluster E_T.

The E_T deposited in a 5×9 TTs isolation region displayed on figure 1, is computed excluding the footprint of the e/γ candidate. The threshold is a function of η and depends on an estimator for the number of pile-up interactions. This estimator corresponds to the number of TTs above a certain E_T threshold produced in the 8 central η rings of the calorimeters. The isolation threshold is currently tuned to reach 90% of trigger efficiency, constant as a function of pile-up and η.

The performance of this algorithm is compared with the Run I trigger using a $Z \to ee$ 2012 data sample (events selected with a tag-and-probe method) for the efficiency and a zero bias sample from a special high pile-up fill for the rates. The efficiency curves as well as the expected rate are shown in figure 2. The turn-on curve obtained is sharper than that of the Run 1 algorithm due to the recovery of energy lost through bremsstrahlung using a dynamic clustering at the TT level. The energy deposited by electrons is better clustered and leads to a better energy resolution of about 30%. In the endcaps, the upgrade improvement comes from the ability of the clustering to adapt to the peculiar geometry along with a more precise energy calibration. The rate of the single electron trigger can be reduced further by using the isolation critera with a 10% efficiency loss. The shape veto can discriminate between e/γ and jet clusters reducing the fake rate, whilst keeping the efficiency loss neglible.
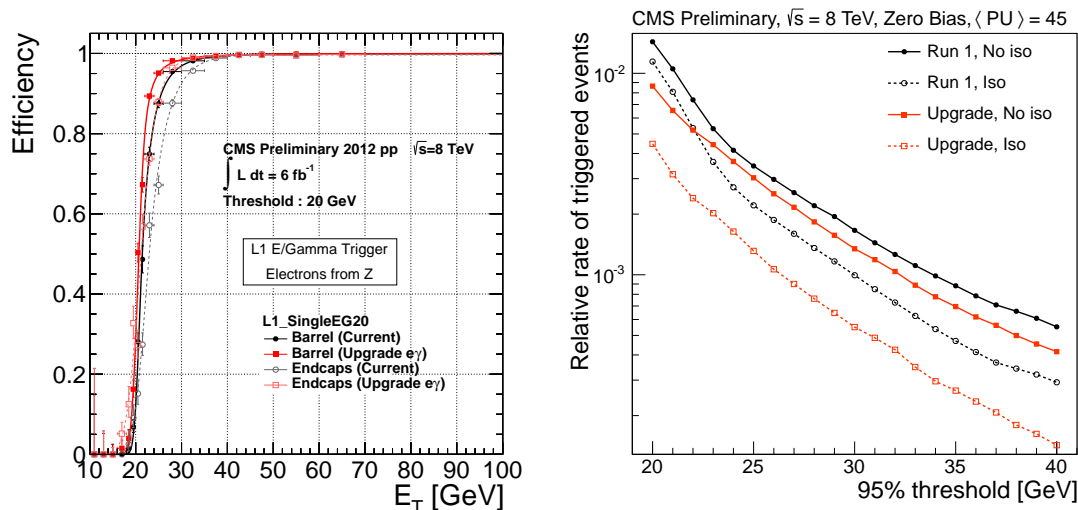
Figure 2. Electron trigger efficiency for a 20 GeV threshold at L1 as a function of the offline reconstructed E_T in the EB and EE (left). The efficiencies obtained with the current and the upgraded algorithms are shown with and without isolation criteria. Relative rate (right) of triggered events from 8 TeV zero bias data for an average pile-up of 45, obtained with the current (Run I) and the upgraded algorithms, both with and without their respective isolation requirements.

3.2   Selecting tau leptons

The algorithm developed is aiming at reconstructing efficiently hadronically decaying τ leptons at hardware trigger level [6]. Depending on the decay mode, several decay products may be producing more than one cluster spatially separated along the ϕ direction due to the magnetic field. Although the footprint of the τ lepton energy deposit is larger than that of an electron, the dynamic clustering developed for e/γ is perfectly adapted to reconstruct these individual clusters which can subsequently be merged. The calibration scheme is using separately the ECAL and HCAL energies and combines them linearly to optimize the response. The parameters are computed for various bins in p_T and η and stored in a Look-Up-Table (LUT). The isolation energy is derived in a similar way than for e/γ candidates. The τ footprint is subtracted from the isolation energy which is then compared to a threshold depending on η. An additional shape veto LUT is also produced in order to discard background-like clusters from the list of possible τ candidates.

The performance of the τ lepton finder algorithm has been assessed on Monte Carlo simulation samples produced with $\sqrt{s}$ = 13 TeV, a bunch spacing of 25 ns and 40 average pile-up interactions. The performance of the 2016 upgraded L1 algorithm is evaluated with respect to hadronic τ decays reconstructed offline using a particle flow based technique. The trigger efficiency as a function of the offline τ reconstructed p_T is displayed in figure 3. The 2016 upgrade algorithm shows superior triggering efficiency performance than Run 1 which does not reach a 100% plateau due to the implementation of a strict veto requirement and isolation on its candidates. Figure 3 shows the trigger efficiency on signal against the background rejection. For the same background rejection a significantly higher efficiency is achieved. Considering a target rate of 3 kHz for the double hadronic τ trigger, the L1 thresholds for the upgrade algorithm range between 30 (29) GeV and 42 (40) GeV depending on the working point of the isolation without (with) shape veto.
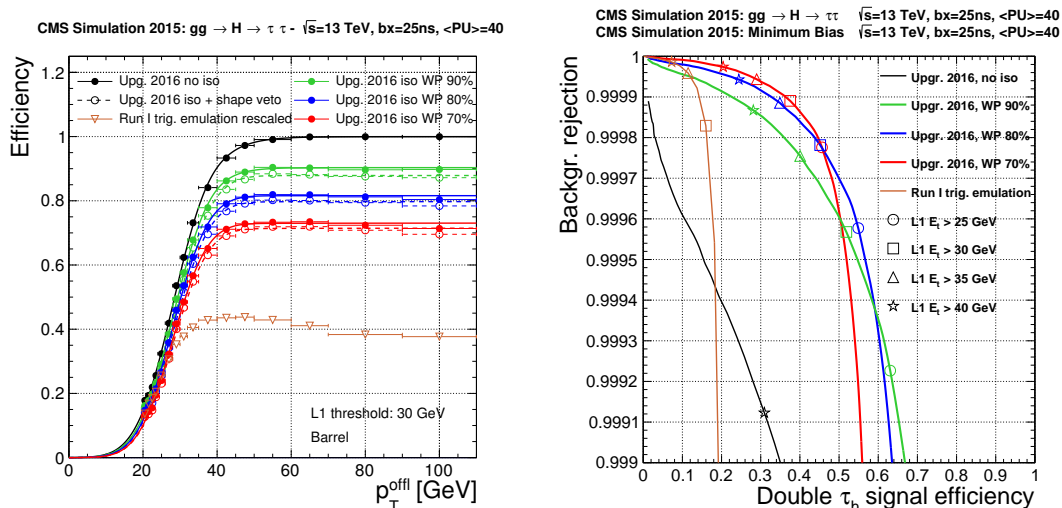
Figure 3. The Level-1 2016 upgraded trigger efficiency (left). The black line is computed for the upgrade trigger in absence of isolation, and the green, blue and red lines correspond to the 90%, 80% and 70% isolation efficiency working points (WP). The dashed lines correspond to additional requirements on the shape veto. The orange line denotes the Run I algorithm. Rate against efficiency for the double τ hadronic trigger at Level-1 (right). Various WP are shown as well as the Run 1 curve in orange.

3.3 Jet and energy sums algorithms and their performance

The jet reconstruction algorithm described here is based on a similar square-jet approach as used in Run I but considering a 9×9 TT sliding window centered on a local maxima [7]. The window size chosen here is coherent with a cone radius of 0.4 for the offline jet reconstruction algorithm. In order to avoid double counting of jets the central TT energy is required to satisfy the inequalities illustrated in figure 4. A jet candidate is discarded if any of the other TTs in the square have an energy deposit of either greater than or greater than or equal to that in the central TT. The veto condition applied is antisymmetric along the diagonal of the square to prevent TTs with the same energy from vetoing one another. These conditions avoid double counting of overlapping jets without significant inefficiency for jet finding. Any TT that satisfies these conditions is considered as a potential jet center. The jet candidate energy is the sum of all 9×9 TTs energies. As seen on figure 4, the algorithm shows excellent agreement when compared with the anti-k_t algorithm used for offline reconstruction. Global quantities are also computed on full calorimeter granularity such as HT, the jet-based equivalent of the total E_T.

The pile-up subtraction is designed to operate on an event-by-event basis, allowing the performance to change dynamically based on differing pile-up conditions. A local pile-up correction technique called "chunky donut" was selected and proven to be efficient. Figure 4 illustrates the chunky donut area that is used to estimate the local pile-up energy density to be subtracted from the jet energy. Fluctuations that originate from the presence of other jets in the vicinity or the absence of pile-up is mitigated by the extension of the donut area to four 9×3 strips around the jet and dropping the highest and lowest energy sides. Figure 4 clearly illustrates the nice correlation of the donut energy with the number of interactions and consequently to the level of pile-up.
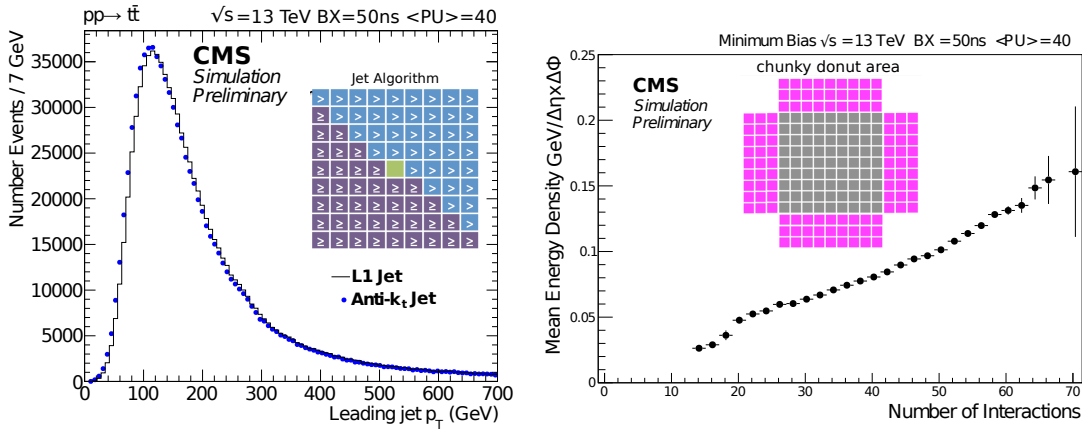
Figure 4. The 9 × 9 sliding window (left) used in the Level-1 jet algorithms. Any TT that satisfies these conditions is considered as a potential jet center. The histogram is a comparison between anti-k_t and the Level-1 jet algorithms. Anti-k_t is using a distance parameter of 0.4 and for the sake of performing a fair comparison; it is run on the same Level-1 TT inputs. The calorimeter area (shown in magenta) around the Level-1 jet considered for the chunky donut algorithm that is used to estimate the local energy density from pile-up (left). Correlations of the donut energy with the number of interactions is displayed.
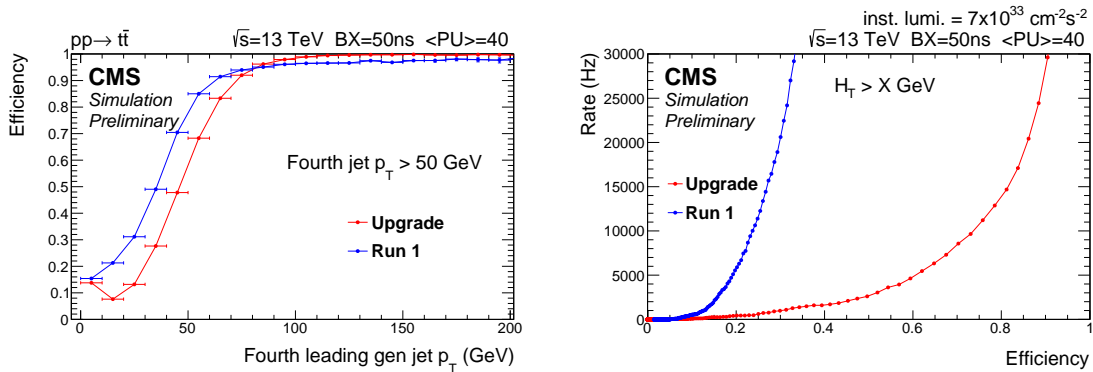

Figure 5. Level-1 Jet trigger efficiency as a function of the offline generated jet p_T for the fourth-leading jets in MC $t\bar{t}$ simulated events (left). Level-1 HT trigger rate against efficiency for an instantaneous luminosity 7×10^{33} cm^{-2} s^{-1}.

Jet energy calibration is expected to depend on the jet p_T and η. A dedicated LUT is derived from QCD di-jet MC by matching Level-1 jets to offline ones. The performance studies of the jet and energy sum algorithms were based on $t\bar{t}$ MC samples with an average of 40 pile-up interactions. A comparison is made to the same quantities produced by the Run 1 system as it was in 2012. Figure 5 shows the Level-1 jet trigger efficiency as a function of the fourth-leading offline jet p_T for a threshold of 50 GeV. This trigger efficiency is shown against the rate (for an instantaneous luminosity of 7×10^{33} cm^{-2} s^{-1} with 50 ns bunch crossing interval) on the same figure. The rate vs efficiency is shown for the global quantities HT on figure 5. The Level-1 upgraded jets algorithm with pile-up subtraction gives substantially better performance than the Run 1 trigger.

4 Firmware implementation and first commissioning results

The firmware implementation is particularly challenging as the electron finder along with the τ lepton and jet finders must fit within a single XC7V690T Xilinx FPGA. The firmware also includes core firmware, which comprises all the necessary logic to control the 72 input/output optical serial links. It also includes the configuration registers, the input pattern buffers and output spy buffers that should be accessible by software. The core firmware represents at this stage a total of 22% of the chip. The software interface is based on the IPBUS standard using libraries such as μHAL developed at CERN [8].

The calorimeters have a granularity of 72 TTs in the ϕ direction and 41 TTs in each of the positive and negative η direction (including here ECAL, HCAL and the forward calorimeter). The TMT architecture allows therefore the data to be rearranged in geometrical order (spatially-pipelined). The algorithms are thus fully pipelined and process the data at the incoming rate starting on the reception of the first data word. For the 32 bits received on each link, the internal computing frequency achieved is 240 MHz.

Upon reception, the TTs are combined to form basic blocks of 3×1 TTs which are then used to form larger objects such as 3×3, 3×9 and 9×9. These objects are the base components of the jet finder algorithm, the donut pile-up estimator and the lepton isolation region. A TT energy threshold is implemented to identify potential cluster seeds while a specific veto procedure is used to perform that particular step for jets as seen on figure 4. Global quantities are calculated as TTs are being received. At this early stage the 8 central η rings are summed to estimate the pile-up level for lepton isolation. In order to reduce the amount of resources required by the implementation of several adders in the lepton cluster logic, only "quality flags" are set for each of the incoming TT. These flags are determined by predefined criteria such as seeding, sharing or trimming. The cluster energy is thus computed as the sum of TTs with non-zero quality flag. The identification, H/E, shape identification criteria and energy calibration are implemented as LUTs. τ lepton candidates are built from e/γ-type clusters. Another LUT is consulted to perform cluster merging if required. A precise FPGA floor planning scheme has been developed to efficiently perform a place and route process and to guarantee that timing constraints will be satisfied after modification of VHDL sources during the development process. The fully pipelined firmware approach provides an efficient way to localize the processing, reduce the size and number of fan-outs, minimize routing delays and eliminates register duplication. The algorithm firmware is now loaded into the Layer-2 processors and comparisons of the outputs produced with the emulator are provided on figure 6. These show an excellent agreement of p_T and η distributions with the expected outputs from simulated MC data patterns.

5 Conclusions

The TMT approach provides an efficient way to build sophisticated algorithms as described in this document. The Level-1 e/γ, tau leptons, jets and energy sums algorithms have been presented and the performance reached are substantially higher than that of the Run 1 system. The upgraded trigger system is now fully deployed in the CMS service cavern and all the links amongst subcomponents have been installed and validated. The latency of the upgraded system has also been determined [4].

Figure 6. Level-1 e/γ upgraded algorithm firmware output distributions for p_T (top left) and η (top right) produced on $Z \to ee$ MC events by the Layer-2 firmware. Level-1 jet upgraded algorithm firmware output distributions for p_T (bottom left) and η (bottom right) produced on $t\bar{t}$ MC events. All results are compared to the expected outputs from the simulation.

The algorithm firmware is now deployed on the production system. The outputs obtained from injected MC patterns were compared with the expected simulation and shown to be in excellent agreement. The system is currently running in parallel alongside the existing calorimeter trigger for full commissioning with real collision data.

Acknowledgments

References

[1] CMS collaboration, *CMS TriDAS project: technical design report. Volume 1: The trigger systems*, CERN-LHCC-2000-038, CMS-TDR-6-1 (2004).

[2] CMS collaboration, *CMS technical design report for the level-1 trigger upgrade*, CERN-LHCC-2013-011, CMS-TDR-12 (2013).

[3] M. Baber et al., *Development and testing of an upgrade to the CMS level-1 calorimeter trigger*, 2014 *JINST* **9** C01006.

[4] B. Kreis et al., *Run 2 upgrades to the CMS level-1 calorimeter trigger*, in *Topical Workshop on Electronics for Particle Physics 2015*, Lisbon Portugal, 28 Sep–2 Oct 2015.

[5] CMS collaboration, J.B. Sauvan, *Performance and upgrade of the CMS electron and photon trigger for Run 2*, *J. Phys. Conf. Ser.* **587** (2015) 012021;
CMS collaboration, F. Beaudette et al., *The CMS level-1 calorimeter trigger electronics system for the LHC Run II*, in *Topical Workshop on Electronics for Particle Physics 2014*, Aix en Provence France, 22–26 Sep 2014.

[6] CMS collaboration, L. Mastrolorenzo, *The CMS level-1 tau algorithm for the LHC Run II*, CMS-CR-2014-309 (2014);
CMS collaboration, L. Cadamuro, *The CMS level-1 tau algorithm for the LHC Run II*, PoS(EPS-HEP2015)226.

[7] CMS collaboration, *L1 calorimeter trigger upgrade: jet and energy sum performance and commissioning status*, CMS-DP-2015-051, CERN-CMS-DP-2015-051 (2015).

[8] C.G. Larrea et al., *IPbus: a flexible Ethernet-based control system for xTCA hardware*, 2015 *JINST* **10** C02019.