


The Compact Muon Solenoid Experiment  
**Conference Report**

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



11 November 2016 (v2, 15 November 2016)

# The CMS Level-1 Calorimeter Trigger for LHC Run II

A. Zabi, F. Beaudette, L. Cadamuro, O. Davignon, T. Romanteau, T. Strebler, M. Cepeda, J. B. Sauvan, N. Wardle, R. Aggleton, F. Ball, J. Brooke, D. Newbold, S. Paramesvaran, D. Smith, J. Taylor, C. Foudas, M. Baber, A. Bundock, S. Breeze, M. Citron, A. Elwood, G. Hall, G. Iles, C. Laner, B. Penning, A. Rose, A. Shtipliyski, A. Tapper, T. Durkin, K. Harder, S. Harper, C. Shepherd-Themistocleous, A. Thea, T. Williams, S. Dasu, L. Dodd, R. Forbes, T. Gorski, P. Klabbers, A. Levine, I. Ojalvo, T. Ruggles, N. Smith, W. Smith, A. Svetek, J. Tikalsky and M. Vicente for the CMS Collaboration

## Abstract

Results from the completed Phase 1 Upgrade of the Compact Muon Solenoid (CMS) Level-1 Calorimeter Trigger are presented. The upgrade was completed in two stages, with the first running in 2015 for proton and Heavy Ion collisions and the final stage for 2016 data taking. The Level-1 trigger has been fully commissioned and has been used by CMS to collect over 43 fb<sup>-1</sup> of data since the start of the Large Hadron Collider (LHC) Run II. The new trigger has been designed to improve the performance at high luminosity and large number of simultaneous inelastic collisions per crossing (pile-up). For this purpose it uses a novel design, the Time Multiplexed Trigger (TMT), which enables the data from an event to be processed by a single trigger processor at full granularity over several bunch crossings. The TMT design is a modular design based on the uTCA standard. The trigger processors are instrumented with Xilinx Virtex-7 690 FPGAs and 10 Gbps optical links. The TMT architecture is flexible and the number of trigger processors can be expanded according to the physics needs of CMS. Sophisticated and innovative algorithms are now the core of the first decision layer of the experiment. The system has been able to adapt to the outstanding performance of the LHC, which ran with an instantaneous luminosity well above design. The performance of the system for single physics objects are presented along with the optimizations foreseen to maintain the thresholds for the harsher conditions expected during the LHC Run II and Run III periods.

Presented at *TWEPP-16 Topical Workshop on Electronics for Particle Physics*

PREPARED FOR SUBMISSION TO JINST

16<sup>TH</sup> TOPICAL WORKSHOP ON ELECTRONICS IN PARTICLE PHYSICS  
26-30 SEPTEMBER 2016  
KARLSRUHE, GERMANY

## The CMS Level-1 Calorimeter Trigger for the LHC Run II

---

A. Zabi,<sup>a,1</sup> F. Beaudette,<sup>a</sup> L. Cadamuro,<sup>a</sup> O. Davignon,<sup>a</sup> T. Romanteau,<sup>a</sup> T. Strebler,<sup>a</sup> M. Cepeda,<sup>b</sup> J. B. Sauvan,<sup>b</sup> N. Wardle,<sup>b</sup> R. Aggleton,<sup>c</sup> F. Ball,<sup>c</sup> J. Brooke,<sup>c</sup> D. Newbold,<sup>c</sup> S. Paramesvaran,<sup>c</sup> D. Smith,<sup>c</sup> J. Taylor,<sup>e</sup> C. Foudas,<sup>d</sup> M. Baber,<sup>e</sup> A. Bundock,<sup>e</sup> S. Breeze,<sup>e</sup> M. Citron,<sup>e</sup> A. Elwood,<sup>e</sup> G. Hall,<sup>e</sup> G. Iles,<sup>e</sup> C. Laner,<sup>e</sup> B. Penning,<sup>e</sup> A. Rose,<sup>e</sup> A. Shtipliyski,<sup>e</sup> A. Tapper,<sup>e</sup> T. Durkin,<sup>f</sup> K. Harder,<sup>f</sup> S. Harper,<sup>f</sup> C. Shepherd-Themistocleous,<sup>f</sup> A. Thea,<sup>f</sup> T. Williams,<sup>f</sup> S. Dasu,<sup>g</sup> L. Dodd,<sup>g</sup> R. Forbes,<sup>g</sup> T. Gorski,<sup>g</sup> P. Klabbers,<sup>g</sup> A. Levine,<sup>g</sup> I. Ojalvo,<sup>g</sup> T. Ruggles,<sup>g</sup> N. Smith,<sup>g</sup> W. Smith,<sup>g</sup> A. Svetek,<sup>g</sup> J. Tikalsky and <sup>g</sup> M. Vicente<sup>g</sup> on behalf of the CMS collaboration

<sup>a</sup>*Laboratoire Leprince-Ringuet, Ecole Polytechnique, IN2P3-CNRS, Palaiseau, France,*

<sup>b</sup>*CERN, European Organization for Nuclear Research, Geneva, Switzerland*

<sup>c</sup>*University of Bristol, Bristol, United Kingdom*

<sup>d</sup>*University of Ioánnina, Ioánnina, Greece*

<sup>e</sup>*Imperial College, University of London, London, United Kingdom*

<sup>f</sup>*Rutherford Appleton Laboratory, Didcot, United Kingdom*

<sup>g</sup>*University of Wisconsin, Madison, Wisconsin, U.S.A.*

*E-mail:* [Alexandre.Zabi@cern.ch](mailto:Alexandre.Zabi@cern.ch)

---

<sup>1</sup>Corresponding author.

ABSTRACT: Results from the completed Phase 1 Upgrade of the Compact Muon Solenoid (CMS) Level-1 Calorimeter Trigger are presented. The upgrade was performed in two stages, with the first running in 2015 for proton and heavy ion collisions and the final stage for 2016 data taking. The Level-1 trigger has been fully commissioned and has been used by CMS to collect over  $43 \text{ fb}^{-1}$  of data since the start of the Run II of the Large Hadron Collider (LHC). The new trigger has been designed to improve the performance at high luminosity and large number of simultaneous inelastic collisions per crossing (pile-up). For this purpose it uses a novel design, the Time Multiplexed Trigger (TMT), which enables the data from an event to be processed by a single trigger processor at full granularity over several bunch crossings. The TMT design is a modular design based on the  $\mu$ TCA standard. The trigger processors are instrumented with Xilinx Virtex-7 690 FPGAs and 10 Gbps optical links. The TMT architecture is flexible and the number of trigger processors can be expanded according to the physics needs of CMS. Sophisticated and innovative algorithms are now the core of the first decision layer of the experiment. The system has been able to adapt to the outstanding performance of the LHC, which ran with an instantaneous luminosity well above design. The performance of the system for single physics objects are presented along with the optimizations foreseen to maintain the thresholds for the harsher conditions expected during the LHC Run II and Run III periods.

KEYWORDS:  $\mu$ TCA, Trigger, Time-Multiplexed-Trigger, Algorithms, Firmware

---

## Contents

<b>1</b>	<b>Introduction</b>	<b>1</b>
<b>2</b>	<b>Upgrade to the Level-1 trigger system for the Run II of the LHC</b>	<b>2</b>
2.1	Conceptual choices for the upgraded trigger	2
2.2	Advanced Level-1 data processors	4
2.3	The overall upgraded trigger infrastructure	4
<b>3</b>	<b>The improved Level-1 Calorimeter trigger algorithms</b>	<b>5</b>
3.1	The electron, photon and tau lepton finders	6
3.2	The jets and energy sums finders	6
3.3	The algorithm firmware implementation	7
<b>4</b>	<b>Commissioning, operations and performance of the calorimeter trigger</b>	<b>8</b>
<b>5</b>	<b>Conclusions</b>	<b>9</b>

---

## 1 Introduction

The Large Hadron Collider (LHC) has resumed its operation during the spring of 2015 after a two-year long shutdown. The Run II of the LHC is progressing with an increased center-of-mass energy of 13 TeV and an instantaneous luminosity that rapidly reached the record of  $1.5 \times 10^{34} \text{cm}^{-2} \text{s}^{-1}$  on the 17<sup>th</sup> of October 2016. The average number of pile-up reached about 50 for this particular run. These running conditions are well beyond the design specifications of the machine and more than ever, the physics program relies on the trigger system to select all potentially interesting collision events amongst the millions occurring per second [1]. The CMS trigger system is organised in two consecutive steps: the hardware-based Level-1 (L1) trigger utilizes coarse energy deposits in the calorimeters and signals in the muon systems to reduce the rate from 40 MHz to about 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. The output rate of the HLT is about 1500 Hz and the overall reduction factor achieved is  $O(10^5)$ . The CMS electromagnetic (ECAL), the hadronic (HCAL) and the forward hadronic calorimeters (HF) perform measurements of particles' transverse energy ( $E_T$ ) that are transmitted to the Level-1 Calorimeter Trigger in the form of primitives. These primitives are combined to reconstruct calorimeter trigger objects, which corresponds to electrons, photons,  $\tau$  leptons, jets and energy sums. The increase of the luminosity experienced in 2016 is expected to continue in order to reach the unprecedented target of  $2 \times 10^{34} \text{cm}^{-2} \text{s}^{-1}$  during the LHC Phase-1 period. The Phase-II will then start around 2024 with an expected luminosity of

$5 \times 10^{34} \text{cm}^{-2} \text{s}^{-1}$  along with 140 pile-up events. Until then, no fundamental modification of the current primitive generation scheme is required. One of the main constraints comes from the total Level-1 bandwidth of 100 kHz that will remain unchanged. In order to avoid a significant increase in trigger energy thresholds, which would be detrimental for physics, an upgrade of the L1 trigger system has been required [2].

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

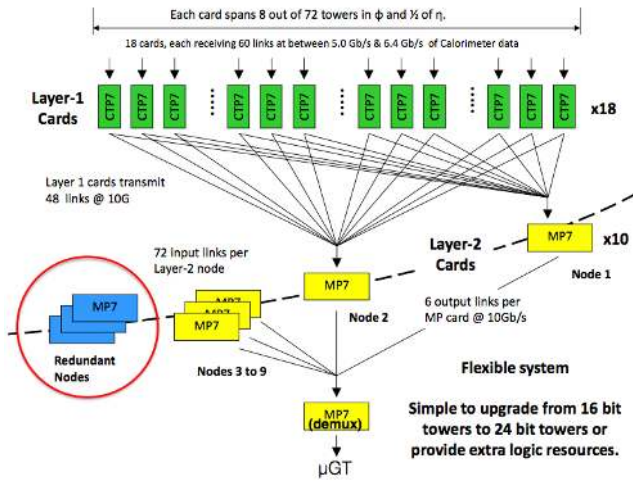
The upgrade to the Level-1 calorimeter trigger system is certainly motivated by the need to preserve the trigger acceptance for physics. The system installed during Run I would not have been capable of maintaining the thresholds low enough for any decent physics program to be pursued. For example, a single electron trigger of 20 GeV threshold would give a rate equivalent to almost half the total Level-1 bandwidth while jets and energy sums trigger rate would scale rapidly with pile-up and could not be sustained at all. In these intense conditions, the implementation of pile-up mitigation techniques was required already at Level-1 to reach acceptable performance. The Calorimeter trigger upgrade for Phase-I has been performed in two stages to insure that CMS would rely on an efficient trigger system right from the start of Run II. The first stage was a partial upgrade that went online during the spring of 2015 [3]. It has allowed collecting successfully the first  $3 \text{ fb}^{-1}$ . The second stage was a full upgrade of the trigger system that started operations in March 2016. It is referred as the upgraded system in the rest of this document.

### 2.1 Conceptual choices for the upgraded trigger

The performance required to select efficiently collision events in a much harsher environment than the LHC Run I lead to the conceptual choices discussed below. In order to sustain single physics object triggers, higher level of background rejection needed to be achieved through the implementation of sophisticated reconstruction and identification algorithms using the full tower granularity. Large FPGAs such as Xilinx Virtex-7 have been chosen to provide enough computational power. The evaluation of global quantities such as missing transverse energy or pile-up energy can be performed precisely by removing all boundaries. A total of 1152 10 Gbps high-speed optical links were installed to collect rapidly all calorimeter primitives hence providing a full field view of the detector. An increased selectivity of the system is required to cope with the constant raise of instantaneous luminosity and pile-up. The micro Global Trigger ( $\mu\text{GT}$ ) was designed to be expandable to many more possible conditions and more sophisticated quantities, to give a richer menu as the HLT. A flexible and modular architecture based on the  $\mu\text{TCA}$  telecom standard has been selected to instrument the upgraded system. All hardware was replaced including the timing control system and all software and databases.

One of the main key technological changes is the implementation of the original TMT architecture [4] visible on Figure 1. The upgraded trigger is organised in two consecutive processing layers. The layer-1 is composed of 18 Calorimeter Trigger Processor boards (CTP7) [5] that are used to perform the pre-processing and data formatting. Tower level operations are executed such as the sum of ECAL and HCAL energies and energy calibration

as well as the computation of the H/E ratio <sup>1</sup>. The 9 Layer-2 master processor cards (MP7) [6] host all calorimeter algorithms that find particle candidates and compute global energy sums. Each MP7 has access to a whole event at trigger tower (TT) granularity. Because both the volume of incoming data and the algorithm latency are fixed, the position of all data within the system is fully deterministic and no complex scheduling mechanism is required. The benefit of time multiplexing is the removal of boundaries. Redundant processing nodes allow to accommodate sufficient time for complex trigger algorithms to run. These algorithms are fully pipelined and start processing as soon as the minimum amount of data is received. The trigger candidates are sent to a demultiplexer board, also an MP7, that formats the data for the upgraded  $\mu$ GT [7]. The  $\mu$ GT is responsible for allowing the event to be further scrutinized by the HLT.



**Figure 1.** The upgraded Time-Multiplexed-Trigger architecture (left). The outputs of the Layer-1 CTP7 corresponding to one event are transmitted to single processing nodes in the second layer, Layer-2. The Calorimeter Trigger Processor card (CTP7, bottom left) and the Master Processor card (MP7, bottom right).



<sup>1</sup>The ratio of HCAL and ECAL energy for a given tower is used as a veto to discriminate hadrons from electrons energy deposits

## 2.2 Advanced Level-1 data processors

Both the CTP7 and MP7 boards are Advanced Mezzanine Cards (AMC) from the  $\mu$ TCA telecom standard. They have been designed as generic-processing engines equipped with multiple Gbps transceivers. The CTP7 board is visible on Figure 1(bottom left). It is equipped in its core with a Xilinx Virtex-7 XC7VX690T FPGA and a total of 67 Input and 48 output 10 Gbps optical links instrumented by a combination of Avago miniPods (MPO connector) and microPods (Pluggable CXP). The particularity of that AMC board is the addition of a Xilinx Zynq SoC XC7ZQ45 FPGA (Dual ARM Cortex-A9 CPU) that is running embedded linux to handle all communication and support functions. This feature greatly simplifies the addressing and interfacing with the chip. The MP7 is visible on Figure 1(bottom right) and is featuring a Xilinx Virtex-7 XC7VX690T FPGA placed underneath the heatsink. A total of 72 Rx and 72 Tx 10.3 Gbps optical links provided by two sets of Avago miniPODs connected through 4 MPO connectors are placed on the front panel. An Atmel 32-bit MMC supporting microSDHC interface is used handle firmware upload ( $2 \times 144$  Mb 550 MHz QDRII and SRAM). The data throughput and computational power of these cards are sufficient to surpass the performance of the Run I system, which contains over hundreds of older FPGAs distributed across multiple custom boards. These AMC boards are housed in Vadatech VT892  $\mu$ TCA crates, and additional serial and LVDS electrical I/O occurs via the backplane. Figure 2 displays the overall infrastructure with the 3 Layer-1 and the Layer-2 crates. Each of the crates hosts an AMC13 card that provides clock and timing Signals, DAQ readout including monitoring as well as the L1 trigger decision via LVDS. Data to configure the lookup tables (LUTs) and registers are sent via Ethernet to a NAT-MCH  $\mu$ TCA Carrier Hub. In the case of the MP7 a serial link following the IPbus [8] protocol standard and using libraries such as  $\mu$ HAL developed at CERN is used.


## 2.3 The overall upgraded trigger infrastructure

The calorimeter primitive data format is similar to that of the Run I system. However, the communication from the ECAL was upgraded from electrical to optical by retrofitting the ECAL Trigger Concentrator Cards with Optical Synchronization and Link Boards (oSLBs)A total of 576 oSLBs synchronize the ECAL trigger primitives and concentrate them onto 4.8 Gbps links. A complete new  $\mu$ TCA-based system was installed to produce the HCAL and HF trigger primitives. The  $\mu$ HTRimplement a total of  $576^2$  6.4 Gbps links. A total of 1152 links feed into a patch panel connected to the CTP7. Further down, the time multiplexing route from Layer-1 to Layer-2 is provided by a Molex FlexPlane patch panel. 72-to-72 12-fiber MPO cables are routed in three enclosures. This novel technology may be useful in future LHC electronics systems because it can massively simplify and shrink fiber installations at no extra cost. The final trigger decision at Level-1 is performed at the  $\mu$ GT level [7]. The  $\mu$ GT uses an MP7 board to host a total of 280 trigger algorithms<sup>3</sup> and a dedicated board called AMC502 deals with the final “OR” of all algorithms. A

---

<sup>2</sup>504 links are needed for HCAL and 72 for HF. The HF granularity is increased to 12 divisions in positive and negative  $\eta$

<sup>3</sup>The total number of algorithms available is to be compared with the 128-bits of the Run I system



**Figure 2.** The overall upgraded trigger infrastructure. The 3  $\mu$ TCA crates of the Layer-1 equipped with CTP7 boards (left) and the Layer-2 MP7 boards housed in a single crate (top).

multiple board  $\mu$ GT system has been recently implemented and is running successfully [7]. It is used to provide extra selection algorithms.

### 3 The improved Level-1 Calorimeter trigger algorithms

The algorithms of the upgraded Level-1 trigger system have been designed to exploit the full trigger tower granularity and the global calorimeter view provided by the TMT architecture. The goal was to improve single physics object trigger performance in terms of energy and position resolution to maintain low thresholds for physics. These improvements can be achieved by introducing novel reconstruction techniques at the firmware level that are inspired from the offline algorithms. Well-reconstructed single trigger objects can then be used to compute complex quantities and other correlated variables at the  $\mu$ GT level. Trigger decisions based on the invariant mass; object acoplanarity or other topological variables may be introduced at the Level-1. These quantities are expected to enhance significantly the trigger selectivity.

The Level-1 trigger is a hardware system that is a fully synchronous system and therefore operates with a fixed latency. In comparison with HLT, algorithms designed for the Level-1 cannot be iterative. A dynamic clustering has been designed to reconstruct precisely lepton signatures in the calorimeter instead of using a sliding window with a fixed size. The advantage of a dynamical technique is the construction of basic clusters that are combined to reconstruct hadronic  $\tau$  lepton candidates. An optimum-sized window is used to build particle jet candidates directly from trigger towers. Another challenge addressed by the new Level-1 system is the online determination of the pile-up energy without the



information from the tracking<sup>4</sup>. In order to maintain the performance, the pile-up density ought to be subtracted from the measured energy of each calorimeter object. The improved algorithms provide enough handles to adapt the selection technique according to the CMS physics program. A summary of the  $e/\gamma$ ,  $\tau$  and jet algorithm performance measured with 2016 collision data will be presented in what follows.

### 3.1 The electron, photon and tau lepton finders

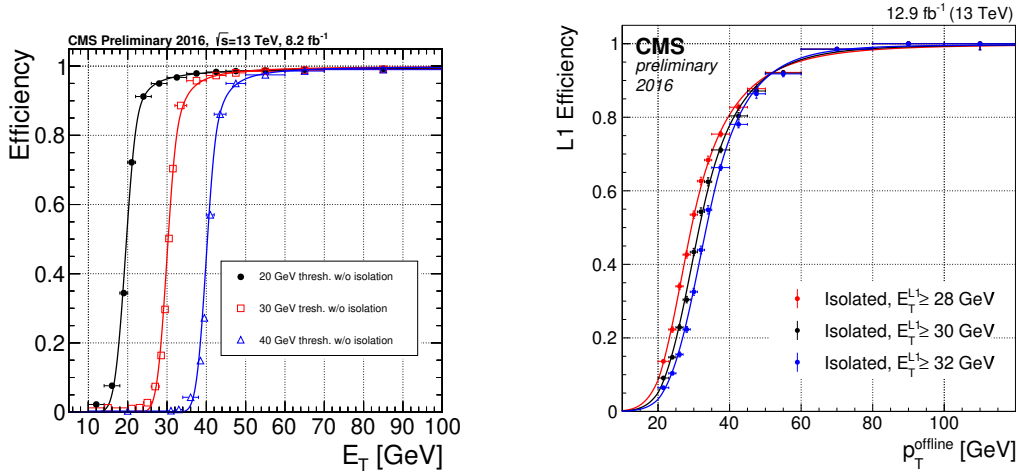
In the CMS detector, electrons (and converted photons) tend to lose some of their energy through radiation when passing through the tracker material. These Bremsstrahlung emissions of photons spread along the  $\phi$  direction due to the bending induced by the magnetic field. In order to obtain a better containment of the electron energy, the reconstruction of the associated cluster should accommodate an extension in the  $\phi$  direction. Inspired from the offline reconstruction, the dynamic clustering was designed to adapt the size of the cluster to precisely match the electron footprint in the calorimeter to optimize the trigger response [9] [10]. As compared to the Run I algorithm [1], the energy resolution is improved by almost 30% in the transition region between the barrel and the endcap where the tracker material is maximum. In the case of hadronically decaying  $\tau$  leptons, many energy clusters associated with each decay product may be produced. The dynamic clustering is used here to reconstruct individual clusters, which can be subsequently merged. The  $\tau$  lepton finder algorithm developed here is far superior to the standalone Run I algorithm that could not reach a 100% efficiency at high energy. Another advantage of the enhanced granularity is the ability to compute the position of the trigger candidate as an energy-weighted average centered on the seed tower. The position resolution is therefore improved by a factor 4 for these objects when compared to the Run I system. An extra collection of isolated candidates can be produced for both the  $e/\gamma$  and  $\tau$  leptons. The isolation energy is evaluated as the energy deposited into ECAL and HCAL in a  $6\times 9$  TTs window around the cluster seed after subtracting the candidate's  $E_T$ . The pile-up energy dependence of the isolation energy is taken care of by introducing a threshold cut that depends on the transverse energy, the position in  $\eta$  and the multiplicity of trigger tower in the event. The cluster shapes produced by the  $e/\gamma$  and  $\tau$  finder algorithms are categorized to provide further background discrimination.

### 3.2 The jets and energy sums finders


The reconstruction of particle jets is based on a  $9\times 9$  TTs sliding window centered around a local maxima. In order to avoid double counting of overlapping jets without efficiency loss, the trigger towers are required to satisfy a set of inequalities [9]. The size of the window is chosen to correspond to the 0.4 cone radius used for the offline anti- $k_T$  reconstruction algorithm. Other global quantities are computed with full calorimeter granularity such as the total  $E_T$ , the missing transverse energy ( $ME_T$ ) and  $H_T$ , the jet-based equivalent of the total  $E_T$  as well as  $MH_T$ , the jet-based equivalent of the  $ME_T$ ). In order to preserve the energy resolution performance, a local pile-up correction technique called “chunky donut” is used to estimate the pile-up energy around the jet candidate on an event-to-event basis.

---

<sup>4</sup>It is planned to introduce the tracking information in the Level-1 trigger for the Phase II CMS detector



**Figure 3.** Trigger efficiency curve for a 20, 30 and 40 GeV electron threshold as a function of the offline reconstructed electron  $E_T$  (left). Trigger efficiency curve for isolated  $\tau$  lepton thresholds of 28, 30 and 32 GeV as a function of the offline reconstructed  $\tau$   $p_T$  (right).



**Figure 4.** Jet trigger efficiency as a function of the offline reconstructed jet  $E_T$  for various Level-1 seed thresholds (left). The single jet trigger efficiency is checked against various average pile-up intervals (right).

### 3.3 The algorithm firmware implementation

The firmware implementation of these algorithms was particularly challenging as all the finder algorithms described above were to fit in a single Xilinx Virtex 7 FPGA. In the TMT approach the data from the calorimeters are reorganized in consecutive rings of 72 TTs in  $\phi$  transmitted from Layer-1 to Layer-2 for each pair of positive and negative  $\eta$  every bunch crossing. An input pipeline is necessary to process the data at the incoming data rate starting upon reception of the first data word. For the 32-bits received on each link, the internal computing frequency achieved is 240 MHz. The structure of the firmware is

organized so that consecutive algorithm steps converge in the core of the chip where the sorting of the trigger candidates takes place. The firmware obtained is compact and easily maintainable. Since the start of the Run II period, the firmware was rebuilt more than 15 times successfully. The total latency of the upgraded system remains under 48 BX in total.

## 4 Commissioning, operations and performance of the calorimeter trigger

The upgraded system has been installed during the long shutdown of the LHC between 2013 and 2015. The integration tests were performed in 2014 and the new system ran in parallel with the intermediate upgrade trigger during the fall of 2015. These first collision data were used to validate the algorithms using a bit-level emulation of the firmware implemented in C++. After a commissioning period using cosmics data taking early 2016, the first collisions were recorded successfully in April the same year. Since the start of operation, the instantaneous luminosity has reached an unprecedented level that required adjustments to be made to the physics selection strategy. Thresholds as well as other parameters such as the lepton isolation and calibration were derived many times over. The LHC has been able to run 65% of the time with stable proton beams declared useful for physics to be compared with 35% in 2015. CMS was able to record collisions at high pace with very little time to perform optimizations. The proton data-taking period was completed on the 28<sup>th</sup> of October 2016 with more than 40 fb<sup>-1</sup> of collision data successfully recorded by the upgraded trigger.

These data have been used to study the performance on single object trigger efficiency as well as energy and position resolutions. Figure 3 (left) shows the trigger efficiency curve for the various single electron thresholds as a function of the offline reconstructed electron  $E_T$ . The result was obtained by looking at events selected with a tag-and-probe method from a  $Z \rightarrow ee$  data sample. Figure 3 (right) displays the trigger efficiency curve for isolated  $\tau$  leptons as a function of the offline reconstructed  $p_T$ . A data sample with  $Z \rightarrow \tau\tau \rightarrow \mu\tau$  selected events is used where the  $\mu$  is the tag and the hadronic  $\tau$  is the probe. The efficiency curves for single jet triggers are shown on Figure 4 (left) and for various average pile-up ranges (right). For the jet studies, a single  $\mu$  dataset was prepared. These curves display excellent triggering performance of the upgrade system for single physics objects. The sharpness of the turn-on curves allowed maintaining the low selection thresholds. The  $\tau$  lepton trigger performance is now comparable to other objects. Lepton isolation and jet reconstruction have been proven to be pile-up resilient.

The single electron trigger threshold was kept below  $\sim 35$  GeV for  $1.5 \times 10^{34} \text{cm}^{-2} \text{s}^{-1}$  and the double electron trigger required  $\sim 25$  GeV and  $\sim 12$  GeV thresholds on both legs. A double isolated  $\tau$  lepton trigger was deployed with a threshold that remained under 32 GeV. The extension of the  $\mu$ GT [7] hardware offers a lot more flexibility in the implementation of cross-triggers and topological triggers based on invariant masses for example. The first dedicated VBF trigger algorithms were deployed in the fall of 2016.

## 5 Conclusions

The development, installation and commissioning of the upgraded trigger has been conducted with a very aggressive schedule to be ready in time for the first collisions. This new system has achieved excellent performance throughout the course of the 2016 LHC proton collision period. Thresholds for single objects were maintained low enough for the CMS physics program to be carried out. The flexibility of the system is demonstrated by its capabilities to adapt to higher instantaneous luminosity and harsher conditions by introducing more topological variables. The collaboration is currently investigating the scalability of the existing system towards the upgrade for the Phase II of the LHC. The generic stream-processing engines designed within the Phase I upgrade program are considered as demonstrators for future trigger architecture developments.

## Acknowledgments

We gratefully acknowledge the following sources of funding: CERN, CNRS/IN2P3(France), STFC (United Kingdom), the U.S. Department of Energy (DOE) and the National Science Foundation (NSF). The work of the author of this paper has been partly funded by the P2IO LabEx (ANR-10-LABX-0038) in the framework “Investissements d’Avenir” (ANR-11-IDEX-0003-01) managed by the French National Research Agency (ANR).

## References

- [1] S. Chatrchyan et al. (The CMS Collaboration), *CMS TRIDAS project technical design report, volume 1, the trigger systems*, CMS TDR CERN/LHCC 2000-38, CMS-TDR-006-1
- [2] S. Chatrchyan et al. (The CMS Collaboration), *CMS technical design report for the level-1 trigger upgrade*, CERN-LHCC-2013-011, CMS-TDR-12 (2013)
- [3] B. Kreis (Fermilab) et al. (The CMS Collaboration), *Run 2 Upgrades to the CMS Level-1 Calorimeter Trigger* 2016 JINST **11** C01051 e-Print : arXiv :1511.05855
- [4] M. Baber et al. (The CMS Collaboration), *Development and testing of an upgrade to the CMS level-1 calorimeter trigger*, 2014 JINST **9** C01006.
- [5] A. Svetek et al. (The CMS Collaboration), *Construction, Testing, Installation, Commissioning and Operation of the CMS Calorimeter Trigger Layer-1 CTP7 Cards* Proceedings of the Tweep2015 conference. CMS-CR-2015-317.
- [6] Imperial College London, MP7 homepage, <http://www.hep.ph.ic.ac.uk/mp7>.
- [7] *Design and Performance of the Phase I Upgrade of the CMS Global Trigger*, J. Wittmann et al. (The CMS Collaboration), these proceedings
- [8] G. Larrea et al. (The CMS Collaboration), *IPbus: a flexible Ethernet-based control system for xTCA hardware*, 2015 JINST **10** C02019.
- [9] A. Zabi et al. (The CMS Collaboration), *Triggering on electrons, jets and tau leptons with the CMS upgraded calorimeter trigger for the LHC RUN II*, 2016 JINST **11** C02008
- [10] *The CMS Electron and Photon Trigger for the LHC Run 2*, N. Dev et al. (The CMS Collaboration), these proceedings