



ELSEVIER

Contents lists available at ScienceDirect

Nuclear Instruments and Methods in Physics Research A

journal homepage: www.elsevier.com/locate/nima

Upgrade of the CMS muon trigger system in the barrel region



Dinyar Rabady^{a,*}, Janos Ero^a, Giannis Flouris^b, Jonathan Fulcher^c, Nikitas Loukas^b, Evangelos Paradis^b, Thomas Reis^c, Hannes Sakulin^c, Claudia-Elisabeth Wulz^c

^a Institute of High Energy Physics Vienna (HEPHY), Nikolsdorfer Gasse 18, 1050 Wien, Austria

^b University of Ioannina, 45110 Ioannina, Greece

^c CERN, 1211 Geneva 23, Switzerland

ARTICLE INFO

Article history:

Received 25 March 2016

Received in revised form

16 May 2016

Accepted 17 May 2016

Available online 19 May 2016

Keywords:

CERN

CMS

Trigger

Muon

Electronics

ABSTRACT

To maintain the excellent performance shown during the LHC's Run-1 the Level-1 Trigger of the Compact Muon Solenoid experiment underwent a significant upgrade. One part of this upgrade is the re-organization of the muon trigger path from a subsystem-centric view in which hits in the drift tubes (DT), the cathode strip chambers (CSC), and the resistive plate chambers (RPC) were treated separately in dedicated track-finding systems to one in which complementary detector systems for a given region (barrel, overlap, and endcap) are merged at the track-finding level. This fundamental restructuring of the muon trigger system required the development of a system to receive track candidates from the track-finding layer, remove potential duplicate tracks, and forward the best candidates to the global decision layer.

An overview will be given of the new track-finder system for the barrel region, the Barrel Muon Track Finder (BMTF), as well as the cancel-out and sorting layer: the upgraded Global Muon Trigger (μ GMT). Both the BMTF and μ GMT have been implemented in a Xilinx Virtex-7 card utilizing the microTCA architecture. While the BMTF improves on the proven and well-tested algorithms used in the Drift Tube Track Finder during Run-1, the μ GMT is an almost complete re-development due to the re-organization of the underlying systems from track-finders for a specific detector to regional track finders covering a given area of the whole detector. Additionally the μ GMT calculates a muon's isolation using energy information received from the calorimeter trigger. This information is added to the muon objects forwarded to the global decision layer, the so-called Global Trigger.

© 2016 The Authors. Published by Elsevier B.V. This is an open access article under the CC BY license (<http://creativecommons.org/licenses/by/4.0/>).

1. Context

The Level-1 Trigger of the Compact Muon Solenoid (CMS) experiment is responsible for reducing the event-rate from the nominal LHC bunch-crossing frequency of 40 MHz to 100 kHz.

The general operating principle for the Level-1 Trigger is to find local features of physics objects in early stages of the trigger chain and successively combine these to regional physics objects. They are subsequently received by a global stage where they are sorted before being sent to the Global Trigger (GT). The GT can trigger a read-out decision based on programmable algorithms. These algorithms work on full physics objects such as muons and jets and can include topological conditions.

1.1. The legacy trigger

The legacy Level-1 Trigger was based on VME technology

* Corresponding author.

E-mail address: dinyar.rabady@cern.ch (D. Rabady).

utilizing mainly Field-Programmable Gate-Arrays (FPGAs) as well as Application-Specific Integrated Circuits with galvanic parallel links for inter-card communication. It was synchronized to the Large Hadron Collider (LHC) clock of 40 MHz and worked in a fully pipelined mode.

In the legacy muon trigger (see Fig. 1) hit information from two of the muon systems in CMS (cathode strip chambers (CSC) or drift tubes (DT)) were sent to a local stage where they were separately combined into track stubs within a single muon station. These track stubs were used to form muon tracks by dedicated track-finders (cathode strip chamber track-finder (CSCTF) and drift tube track-finder (DTTF)).

The DTTF received track segments for the measurement of the ϕ coordinate separately from those for the measurement of the η coordinate. It used these to construct tracks in ϕ and η independently, finally merging them into a complete track candidate.

The four best tracks found in both the CSCTF and DTTF systems were subsequently sent to the Global Muon Trigger (GMT). The resistive plate chamber (RPC) system used a pattern-matching approach in the Pattern Comparator (PACT) to reconstruct tracks. It

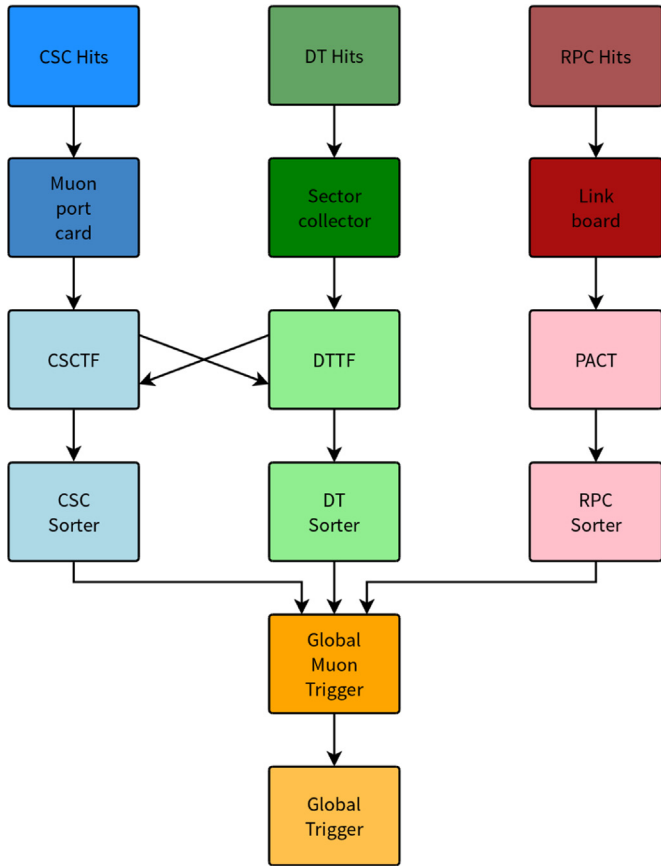


Fig. 1. Block diagram of the legacy CMS Level-1 muon trigger. Hits from the cathode strip chambers (CSC) are sent to the CSC track-finder (CSCTF) via the muon port card. In parallel drift tube (DT) hits are sent to the DT track-finder (DTTF) by the sector collector. Resistive plate chamber (RPC) hits are collected by the link boards and then forwarded to the pattern comparator (PACT). The output of each of these systems is then sorted individually before being sent to the Global Muon Trigger.

then sent the four best muons for both the barrel and endcap region to the GMT.

The GMT matched muon tracks from complementary muon systems. Tracks from RPC PACT were merged with matching CSCTF or DTF tracks according to programmable algorithms that derived an improved track from the two original tracks. Due to the geometry of the CMS detector a muon could be reconstructed in both the barrel as well as the endcap muon systems. The GMT found such tracks and canceled the duplicate. The resulting muons were ranked and subsequently sorted in two stages before the four highest-ranked muons were sent to the GT. The legacy GMT is described in more detail elsewhere [1].

The GT could then combine the muon tracks with information received from the calorimeter trigger in 128 algorithms, of which each could trigger a read-out decision.

A complete description of the legacy Level-1 Trigger is provided in the Technical Design Report [2].

1.2. The upgraded trigger

The LHC's expected instantaneous luminosity during Run-2 exceeds the original design specification. The number of interactions per bunch crossing (pile-up) already surpassed the expected number in the 2012 run.

The Level-1 Trigger is required to support a physics program that both allows searches at the TeV scale and is sensitive to electroweak scale physics. This cannot be achieved by increasing

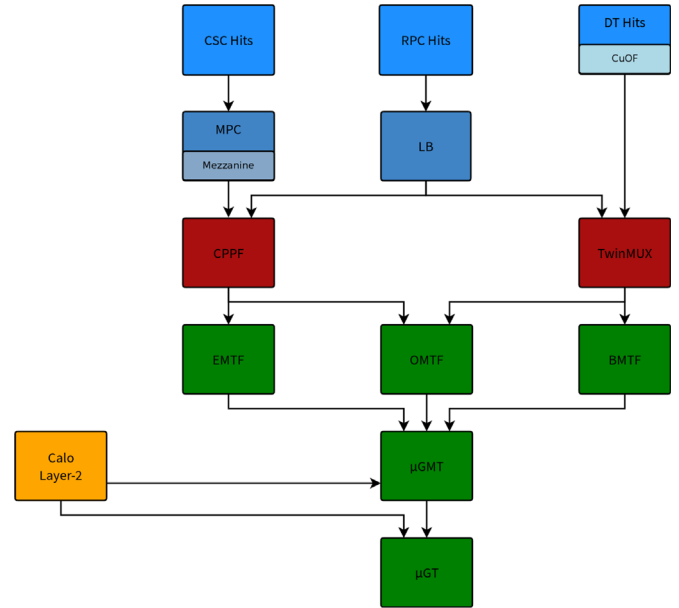


Fig. 2. Block diagram of the upgraded CMS Level-1 muon trigger. Hits from the cathode strip chambers (CSC) are sent via the muon port cards (MPC) to the concentrator and preprocessor fan-out (CPPF) boards that also receive hits from the resistive plate chambers (RPC) in the endcap region via the link boards (LB). Drift tube (DT) hits are sent with a Copper to optical fiber mezzanine (CuOF) to the TwinMUX system which also receives the RPC hits from the barrel region. The track segments are used in the endcap, overlap, and barrel muon track finders (EMTF, OMTF, and BMTF). The resulting tracks are propagated to the Global Muon Trigger (μ GMT) where information from the Calorimeter Layer-2 system is used to assign isolation values. Finally they are sent to the Global Trigger (μ GT).

the Level-1 accept rate as several detectors would require major upgrades in order to accommodate the required read-out rate. Thus an improvement in several areas of the Level-1 Trigger was required. For the muon trigger this implies the introduction of muon isolation as well as an improvement in the muon parameter precision, especially the transverse momentum.

The track-finder architecture was re-structured from detector-based to region-based. In the legacy trigger there was a separate track-finding system for the DT, CSC, and RPC systems, whereas in the upgraded system there is a barrel, overlap, and endcap muon track-finder. Both concerning their hardware as well as their firmware the individual track-finder processors were upgraded, which lead to increases in the track-finders' precision for muon track parameters as well as greater flexibility for future improvements. The inclusion of hit information from the RPC system permits an increase in the quality of reconstructed tracks at an earlier stage than in the legacy system (see Fig. 2).

In order to accommodate the track-finders' increased precision, muon objects were increased in size from 32 bits to 64 bits, which allows the Level-1 Trigger to move to a linear scale for the muon track parameters while remaining flexible for possible later changes (see Table 1).

Table 1
Comparison of bit widths for muon objects in the μ GMT.

	Legacy GMT	Upgraded GMT
ϕ	8 bit	10 bit
η	6 bit	9 bit
p_T	5 bit	9 bit
Quality	3 bit	4 bit
Charge	1 bit	1 bit
Charge valid	1 bit	1 bit
Track addresses	N/A	27 bit

Additionally, the number of muons delivered to the μ GT by the μ GMT was increased from 4 to 8, which allows greater flexibility for the trigger algorithms, such as potentially using lower quality muons for b-tagging.

A comprehensive description of the planned upgrades is given elsewhere [3].

2. Common hardware

For the Level-1 Trigger upgrade the current VME crates were replaced by the microTCA crate system. This system provides system-level health management, redundant power supplies and cooling as well as a high-capacity back plane.

The barrel muon track-finder (BMTF) and μ GMT are implemented in a Xilinx Virtex-7 690 FPGA placed on an Advanced Mezzanine Card (AMC), the Master Processor 7 (MP7), built by Imperial College [4], which provides a significant increase in resources compared to the chips used in the legacy system.

The MP7 supplies 72+72 10 Gb/s channels for receiving and sending data via optical links. Multi-fiber Termination Push-on (MTP) connectors are used for I/O, in which one connector bundles 36 fibers. The backplane is used to receive fast control signals from the Trigger Control and Distribution System [5] and for sending readout information to the data acquisition system upon receipt of a Level-1 Accept signal.

3. Barrel muon track-finder

The BMTF is responsible for reconstructing muon tracks in the barrel area ($|\eta| < 0.83$). It is logically split into twelve 30° “wedges” where each wedge consists of five sectors containing independent drift tube and resistive plate chamber detectors. Data from each wedge are treated by a dedicated track-finder processor, whereas hit information is shared between track-finder processors for neighboring wedges in order to avoid inefficiencies at the overlaps.

The BMTF receives so-called “super-primitives” built by concentrator cards (“TwinMUX”) from DT and RPC hits. These super-primitives describe track segments within a detector module and consist of either the ϕ or η coordinate, the bending angle (if containing the ϕ coordinate), and quality bits indicating the confidence in the measurement. The RPC hits in particular are used to correct mis-measurement of hits that were assigned to a bunch-crossing too early or too late by the DT system.

A given BMTF processor constructs the ϕ - and η -components of a track separately using super-primitives from within a given wedge and from its neighbors in three steps (see Fig. 3):

1. The *extrapolator unit* searches for track segments in adjacent detector modules compatible with a track from a muon originating at the interaction point.
2. The *track assembler unit* receives track segments found by the extrapolator unit. It then creates a track as well as an associated quality value indicating the number of track segments used.
3. The *assignment unit* computes the spatial coordinates, p_T , a quality value, and track addresses using look-up tables (LUT).

The three best muons from each sector are then sent to the Global Muon Trigger by the Wedge Sorter.

4. Upgraded global muon trigger

The upgraded Global Muon Trigger (μ GMT) differs significantly from the legacy system, as it is not necessary to merge muons from the DT and CSC track-finders with those delivered by the RPC system. Additionally, the μ GMT integrates the formerly independent final sorting layers of the track-finders, which in turn requires the canceling out of duplicates at each boundary between two track-finder processors. Finally it computes the isolation of a muon based on the energy deposits in the calorimeter around that muon's track.

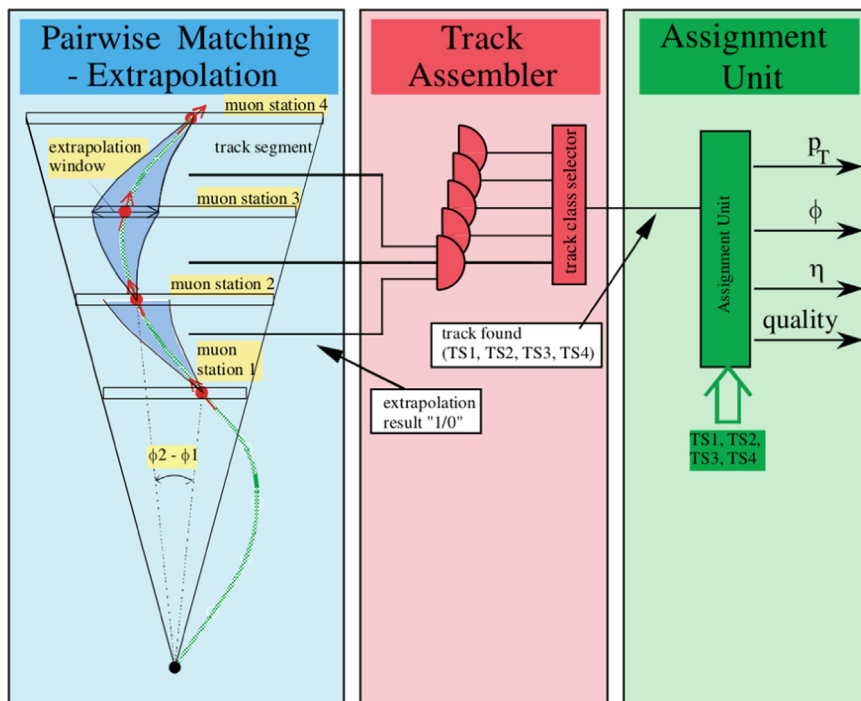


Fig. 3. Illustration of the extrapolation logic in the Barrel Muon Track-Finder.

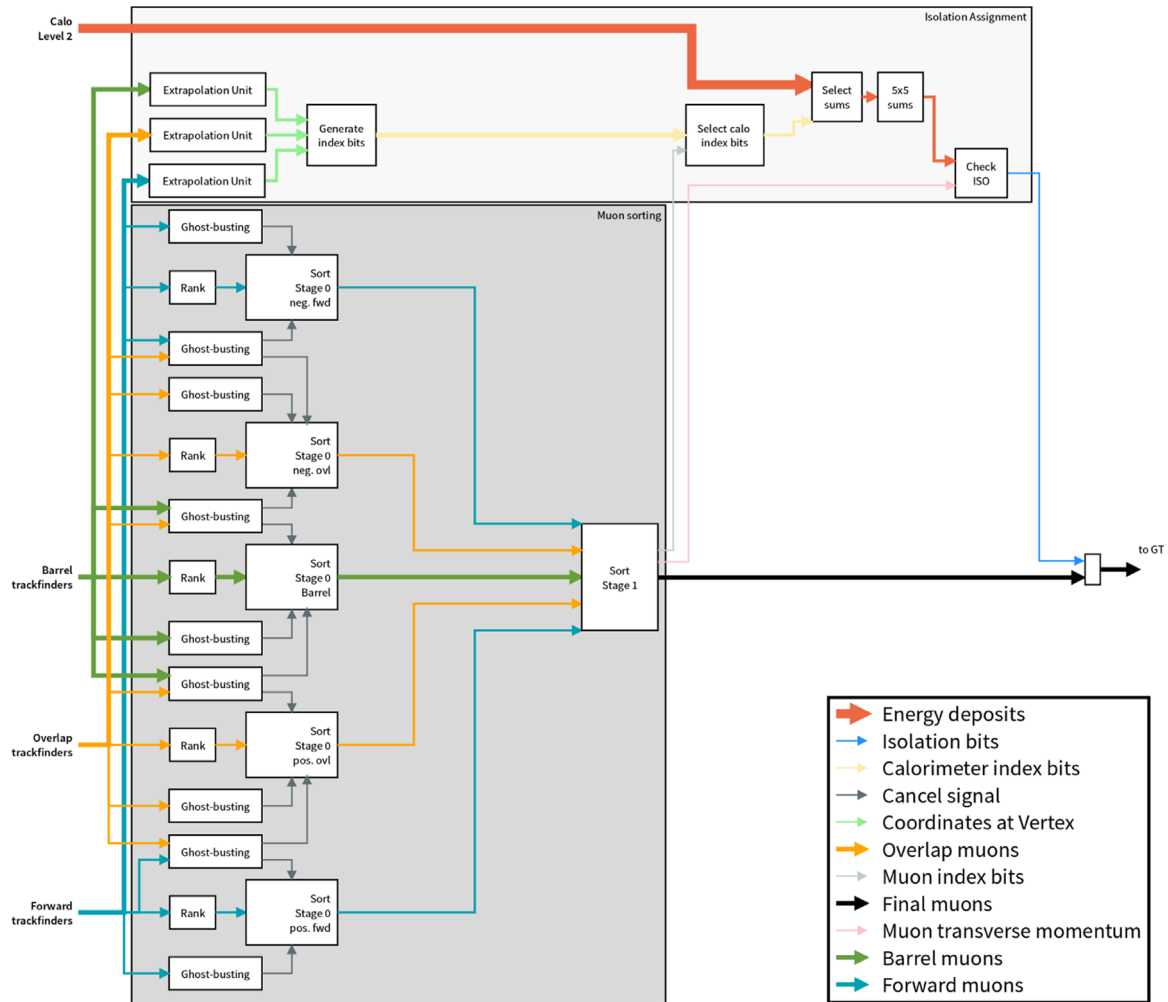


Fig. 4. Block diagram of the upgraded GMT (μ GMT). The μ GMT sorts input muons from the positive and negative sides of the detector in the endcap and overlap region as well as the input muons from the barrel region separately. In parallel ghost-busting takes place as well as an extrapolation for each muon to the vertex. In a second sorter stage muons from all detector regions are sorted again and the eight best are sent to the isolation unit in order to determine their isolation. Finally these eight muons and their isolation values are sent to the Global Trigger.

4.1. Muon sorting

The μ GMT absorbs the track-finders' final sorting stage removing the need for dedicated sorter cards, thereby saving latency otherwise required for de-/serialization at the optical transceivers. As each track-finder consists of 12 processor boards, each with an output of 3 muons, the μ GMT receives 108 input muons at 64 bits. Apart from the latency savings this allows the removal of ghosts between the track-finders at an earlier stage in the processing chain. This is accomplished in two steps (see Fig. 4). In the first stage the muons from each track-finder are sorted separately. Muons from the overlap and endcap track-finders are additionally sorted separately for the positive and negative sides of the detector. Sorting is done according to a rank assigned depending on p_T and the quality of a muon as given by the track-finders. In parallel the muons at track-finder boundaries are matched in order to find possible ghosts. The duplicate muon with the lower quality is then canceled out in the sorter.

The second sorter stage receives four muons each from the positive and negative regions of both the overlap and endcap track-finders, as well as eight muons from the barrel track-finder. The best eight muons out of these 24 are then sent to the μ GT.

4.2. Ghost busting

The removal of duplicate tracks is necessary between the different track-finders, but also for muons from neighboring sectors or wedges in the same track-finder.

The μ GMT thus needs to perform ghost busting between the following areas:

- barrel and positive/negative overlap track-finders,
- positive overlap and positive endcap track-finders,
- negative overlap and negative endcap track-finders,
- neighboring wedges or sectors of each track-finder.

Tracks can be matched either based on their spatial coordinates or based on common track segments used during the tracks' assembly.

- Matching based on spatial coordinates uses a matching window $\Delta R^2 = f_1 \times \Delta \eta^2 + f_2 \times \Delta \phi^2$ where f_1 and f_2 are scaling factors. If the muons are closer than a pre-programmed value for ΔR^2 , the one with the lower quality is marked as a ghost. This matching algorithm requires no dedicated information from the track-finders, however it is less accurate than track-address based cancel-out.
- Ghost busting using the track addresses works by matching the




Fig. 5. Illustration of the area used for the computation of the energy sums required to determine the isolation values.

track-segments used for a muon's track for each station. If a segment is shared between two tracks they are flagged as duplicates of each other. Matching based on track-addresses is more accurate than using spatial coordinates. However, significantly more information is required to perform this kind of ghost busting.

While track addresses are available in the μ GMT for the barrel and endcap track-finders, the overlap track-finder cannot propagate these due to bandwidth constraints. For this reason cancel-out within wedge and sector boundaries for barrel and endcap track-finders is performed using track-addresses while the remaining boundaries use coordinate-based cancel out.

4.3. Muon isolation

The μ GMT calculates both relative and absolute isolation for each muon sent to the upgraded μ GT. This allows algorithms to be used that ignore muons created in jets from in-flight decay.

The μ GMT receives 5 bit energy values for each 2×2 trigger

tower¹ region from the Layer-2 calorimeter trigger. These values are sums of 5×5 of the aforementioned 2×2 trigger towers (see Fig. 5). Muon tracks are extrapolated to the vertex. The value for the energy sum in the calorimeter at that position is compared with a pre-programmed value, to calculate the absolute isolation value, and divided by the muon's p_T value to determine the relative isolation.

The calculated isolation values are then merged with the output muons from the final sorter and propagated to the μ GT.

Currently an alternative to the presented isolation algorithm is under consideration and may be implemented at a later stage.

5. Summary

An overview of the current developments for the upgraded CMS muon trigger in the barrel region has been given. Due to the significantly increased luminosity that will be provided by the LHC during Run 2, the CMS Level-1 Trigger was upgraded. In this upgrade the formerly detector-centric track-finders were replaced by regional ones requiring, in turn, a replacement for the legacy Global Muon Trigger.

Both the Barrel Muon Track Finder and the upgraded Global Muon Trigger were implemented in a Virtex-7 FPGA using the microTCA crate technology. The Barrel Muon Track Finder receives track segments comprised of drift tube and resistive plate chamber hits, which it uses to build tracks in the barrel regions. The upgraded Global Muon Trigger uses the increased input and processing capabilities to provide muon isolation as well as increase both the number and precision of muon objects. Muon sorting in the μ GMT saves latency and allows improved ghost busting.

The system as presented is currently undergoing commissioning in the muon path. The isolation assignment will be commissioned in a next step.

References

- [1] H. Sakulin, A. Taurok, The Level-1 Global Muon Trigger for the CMS Experiment, Technical Report, CMS-CR-2003-040, CERN, Geneva, Oct 2003. URL: <https://cds.cern.ch/record/687846>.
- [2] G. Bayatyan, et al., CMS TriDAS project: Technical Design Report, Volume 1: The Trigger Systems, Technical Design Report CMS. URL: <https://cds.cern.ch/record/706847>.
- [3] A. Tapper, D. Acosta, CMS Technical Design Report for the Level-1 Trigger Upgrade, Technical Report, CERN-LHCC-2013-011. CMS-TDR-12, CERN, Geneva, Jun 2013. URL: <https://cds.cern.ch/record/1556311>.
- [4] K. Compton, et al., The MP7 and CTP-6: Multi-hundred Gbps processing boards for calorimeter trigger upgrades at CMS, JINST 7, 2012, C12024. <http://dx.doi.org/10.1088/1748-0221/7/12/C12024>.
- [5] J.G. Hegeman, et al., The CMS Timing Control and Distribution System, Technical Report, CMS-CR-2015-321, CERN, Geneva, Dec 2015. URL: <https://cds.cern.ch/record/2110277>.

¹ A trigger tower has a size of $\Delta\phi = 5^\circ$ and $\Delta\eta \approx 0.1$.