

The New Global Muon Trigger of the CMS Experiment

Jonathan Fulcher, Joschka Lingemann, Dinyar Rabady, Thomas Reis, and Hannes Sakulin

Abstract—For the 2016 physics data runs, the L1 trigger system of the compact muon solenoid (CMS) experiment underwent a major upgrade to cope with the increasing instantaneous luminosity of the CERN LHC whilst maintaining a high event selection efficiency for the CMS physics program. Most subsystem specific trigger processor boards were replaced with powerful general purpose processor boards, conforming to the MicroTCA standard, whose tasks are performed by firmware on an field-programmable gate array of the Xilinx Virtex 7 family. Furthermore, the muon trigger system moved from a subsystem centered approach, where each of the three muon detector systems provides muon candidates to the global muon trigger (GMT), to a region-based system, where muon track finders (TFs) combine information from the subsystems to generate muon candidates in three detector regions that are then sent to the upgraded GMT. The upgraded GMT receives up to 108 muons from the processors of the muon TFs in the barrel, overlap, and endcap detector regions. The muons are sorted in two steps and duplicates are identified for removal. The first step treats muons from different processors of a TF in one detector region. Muons from TFs in different detector regions are compared in the second step. An isolation variable is calculated, using energy sums from the calorimeter trigger and added to each of the best eight muons before they are sent to the upgraded global trigger (GT) where the final trigger decision is made. The upgraded GMT algorithm is implemented on a general purpose processor board that uses optical links at 10 Gb/s to receive the input data from the muon TFs and the calorimeter energy sums, and to send the selected muon candidates to the upgraded GT.

Index Terms—Compact muon solenoid (CMS), LHC, muon, trigger, upgrade.

I. INTRODUCTION

AFTER having been upgraded during its first long shutdown in 2013 and 2014 the CERN LHC started its second running period (Run-II) in 2015. For Run-II, the pp collision energy was increased from the previous 8 to 13 TeV in the center-of-mass. In addition, the number of pp interactions per bunch crossing (pileup) will reach an average of 50, thus doubling the value of the first LHC running period (Run-I) that ended before the long shutdown. The instantaneous luminosity approached the design specification

Manuscript received June 24, 2016; revised November 22, 2016; accepted January 15, 2017. Date of publication March 21, 2017; date of current version June 26, 2017.

J. Fulcher, J. Lingemann, T. Reis, and H. Sakulin are with the European Organization for Nuclear Research, CERN, 1211 Geneva, Switzerland (e-mail: thomas.reis@cern.ch).

D. Rabady is with the Institute of High Energy Physics, Austrian Academy of Sciences, 1050 Vienna, Austria.

Color versions of one or more of the figures in this paper are available online at <http://ieeexplore.ieee.org>.

Digital Object Identifier 10.1109/TNS.2017.2663442

of $1 \times 10^{34} \text{ cm}^{-2} \text{ s}^{-1}$ in Run-I and is expected to reach $2 \times 10^{34} \text{ cm}^{-2} \text{ s}^{-1}$ by the end of Run-II. With such LHC performance parameters the trigger rate of the first level (L1) of the compact muon solenoid (CMS) trigger system would have increased by around a factor of six, with respect to the time before the long shutdown, were the trigger thresholds kept the same [1]. However, it is not possible to increase the trigger rate as the read-out rate is limited by the detector electronics. Thus, the trigger thresholds would have needed to be increased to maintain the trigger rate, which would have had an undesired impact on the CMS physics program.

The decision was therefore made to upgrade the L1 trigger system to cope with the Run-II conditions and maintain the trigger performance for high mass searches as well as for physics at the electroweak scale. In order to achieve the desired performance, the trigger algorithms needed to be improved to reject more background events while keeping a high efficiency for interesting physics signatures. For triggers involving the muon system, two areas for possible improvements were identified. First, the better measurement of the muon p_T spectrum, especially in the difficult forward detector regions where the magnetic field is inhomogeneous and where also the highest particle rate occurs. The goal was to reduce the tails of the muon resolution spectrum to have less of the many low p_T muons being reconstructed as high p_T muons passing the higher trigger thresholds. Second, the additional use of information from the calorimeters to identify isolated muons that are produced in many processes of interest for the CMS physics program. If a muon is promptly produced in the interaction point, it is usually isolated from other particles and deposits only a small amount of energy in the calorimeter. Contrary to this, if the muon is produced within a jet of other particles via an in-flight decay, those other particles leave a significant energy deposit in the calorimeters, allowing the separation of the two different cases.

To achieve these improvements, more sophisticated algorithms were developed that require more powerful hardware, which has become available in the time since the design of the legacy trigger system. In addition to this, high-speed links allowed an increase in number of bits sent for each muon, thus making it possible to use finer grained and linear scales for the variables. Furthermore, the higher throughput made it also possible to process a higher number of muons in the system.

In [2], the status and plans for the upgrade of the global muon trigger (GMT) during the development phase are discussed. In this paper, the installed system is described and details on the commissioning phase are given. Furthermore,

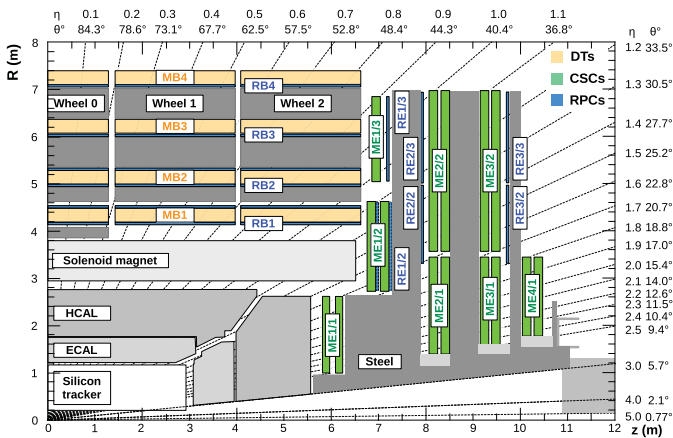


Fig. 1. Quadrant of the CMS detector. The silicon tracker, electronic calorimeter, and hadron calorimeter are located inside the superconducting solenoid. The muon detectors are located between the layers of the steel yoke in the barrel and the endcaps. The DT detector names begin with muon barrel and the CSC names begin with muon endcap. The RPC names begin with RPC barrel and RPC endcap in the barrel and endcap regions, respectively [3].

the performance of the upgraded trigger measured from early 2016 collision data is studied.

A. CMS Detector

The CMS experiment is one of the general purpose experiments at the LHC and located at interaction point 5 of the LHC. Fig. 1 shows a schematic view of a quadrant of the detector that consists of a barrel part and an endcap on each side.

Muons are measured in the pseudorapidity range $|\eta| < 2.4$, with detection planes made using three technologies: drift tubes (DTs), cathode strip chambers (CSCs), and resistive plate chambers (RPCs). The DTs in the barrel cover $|\eta|$ values up to 1.2 and the coverage of the CSCs in the endcaps goes from 0.9 to 2.4. The RPCs are located in the barrel and the endcaps up to $|\eta| = 1.6$. The CMS L1 trigger system, composed of custom hardware processors, uses information from the calorimeters and muon detectors to select the most interesting events in a fixed time interval of $3.8 \mu\text{s}$. The L1 trigger selects interesting events for the physics analysis at a rate of under 100 kHz, which is the maximal read-out rate for some of the front end electronics. This equates to a rate reduction by a factor of 400 compared to the bunch crossing frequency of 40 MHz. The high-level trigger (HLT) processor farm processes the full events built by the central data acquisition (DAQ) system and performs a more sophisticated event selection, further decreasing the event rate from around 100 kHz to less than 1 kHz, before data storage. A more detailed description of the CMS detector, together with a definition of the coordinate system used and the relevant kinematic variables, can be found in [4].

II. UPGRADED MUON TRIGGER

A comparison of the legacy muon trigger system [5] with the upgrade muon trigger is shown in Fig. 2. The upgraded muon trigger system uses a layout based on the detector geometry as opposed to the subdetector-based setup of the

legacy muon trigger. The information from the three muon subsystems is combined at an earlier stage than in the legacy trigger and muon tracks are identified in three detector regions. The track finder (TF) of the barrel region uses information from the DTs and the RPCs and covers the region $|\eta| \leq 0.83$. Information from all three muon subsystems is used in the region $0.83 < |\eta| \leq 1.24$ partially being in the detector barrel and the endcaps. For $|\eta|$ above 1.24 the CSCs and the RPCs are used to build muon tracks in the endcap region.

The TFs are segmented in ϕ in twelve wedges of 30° for the barrel muon TF (BMTF) and six sectors of 60° per endcap for the overlap muon TF (OMTF), and the endcap muon TF (EMTF). The processor from a wedge or sector finds also muons in a ϕ region in the adjacent wedges or sectors, in order to provide overlapping coverage.

The best muons from the regional TFs are selected in the GMT and sent, together with information from the calorimeter about the muon isolation, to the global trigger (GT), where the final trigger decision is calculated by up to 512 trigger algorithms.

Since the start of the 2016 LHC run the concentrator preprocessor and fan-out (CPPF) card, which combines CSC and RPC hits, has not been available and the RPC information is not yet forwarded to the EMTF. Until the CPPF is commissioned, the OMTF and the EMTF receive the CSC hits directly from the mezzanine of the muon port card and the OMTF gets RPC hits from the link board.

A. Common Hardware

Contrary to the legacy system that used the Versa Module Eurocard standard and many parallel galvanic cables for the interconnects, the upgraded trigger is based on the MicroTCA technology [6] and uses multi Gb/s serial optical links for the data transfer between the modules. The MicroTCA crate provides a high bandwidth backplane, system monitoring capabilities, and redundant power modules. The number of distinct electronics boards could be greatly reduced as many components are based on common hardware designs. In the upgraded trigger the BMTF, GMT, GT, and also parts of the calorimeter trigger use the same processor card, whereas the OMTF and the EMTF also share a common design. With this concept many firmware and control software components, e.g., data read-out and link monitoring, can be reused by several systems and, therefore, the workload for development and maintenance is reduced.

An advanced mezzanine card (AMC), named the AMC13 [7], situated at the position of the omitted second MicroTCA carrier hub card, provides the fast control signals from the trigger control and distribution system to the trigger AMCs over the backplane. In case of the arrival of a trigger signal, the trigger AMCs send their data over the backplane to the AMC13 that also provides the connection to the central DAQ system of CMS via 10 Gb/s optical links.

III. GLOBAL MUON TRIGGER

The muon tracks built by the regional TFs are sent to the GMT via 36 optical links operating at 10 Gb/s. Each link

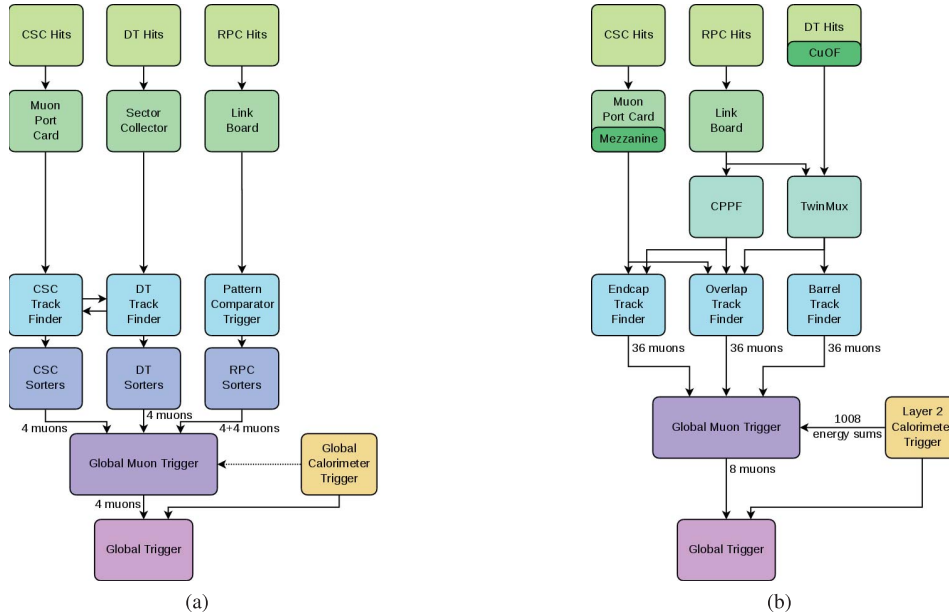


Fig. 2. Block diagrams of the (a) legacy muon trigger and (b) upgraded muon trigger. For the legacy system hits from the CSCs, DTs, and RPCs are sent to the CSC TF, DT TF, and pattern comparator trigger via the muon port card, sector collector, and link board, respectively. The tracks from the three systems are sorted before being sent to the GMT. The GMT also receives calorimeter data but this was not used in Run-I. The best four muons from the GMT are sent to the GT. For the upgrade system hits from the CSCs are sent to the EMTF and the OMTF via a mezzanine on the muon port card. Endcap RPC hits are sent via the link board to the CPPF card and barrel RPC hits are sent to the TwinMux concentrator card. DT hits are sent to the TwinMux card via a copper to optical fiber mezzanine. The EMTF receives RPC hits via the CPPF card. In addition to the CSC hits the OMTF receives DT hits and RPC hits via the CPPF and the TwinMux, which also provides DT and RPC hits to the BMTF. Muons from the three TFs are sent to the GMT, which sorts the muons, performs duplicate removal, adds isolation information from the calorimeter trigger, and sends the best eight muons to the GT.

carries up to three muons from one TF processor, covering a barrel wedge or an overlap or endcap sector. In total, each muon TF can send up to 36 muons resulting in a maximal number of 108 muons as input to the GMT, more than six times the number of muons the legacy GMT received. Compared to the legacy trigger muons, which used 32 b, the input muons in the upgrade system use 62 b in which the following variables are encoded:

- 1) p_T ;
- 2) η coordinate;
- 3) local ϕ coordinate in the wedge or sector;
- 4) charge;
- 5) charge measurement validity bit;
- 6) quality;
- 7) halo bit / η -fine bit;
- 8) track address.

The halo bit is used by the EMTF and indicates muons from beam halo, while the η -fine bit indicates that the BMTF the η assignment has high precision. The track address encodes the track segments used by the TF to build the muon track.

A block diagram of the GMT algorithm is shown in Fig. 3. In the GMT, the muons are sorted in two sorting stages, duplicate muons are removed and up to eight muons are sent to the GT, doubling the number of muons the legacy GT received. Details on the implementation of the algorithm are given in the following sections.

The global ϕ of the output muon is calculated using the local ϕ of the input muon in a wedge or sector and adding an offset depending on which wedge or sector the muon originates from.

Since the GT can consist of up to six processor boards six copies of the GMT output muons are provided on six times four 10 Gb/s optical links.

In addition to the muons, the GMT receives calorimeter information from the calorimeter trigger in order to calculate the isolation of the muons. Around a thousand energy sums over calorimeter towers¹ are received on 28 optical links running at 10 Gb/s.

A. Hardware and Firmware

1) *Hardware*: The GMT algorithm is implemented on the Master Processor 7 (MP7) general purpose data stream processor AMC designed by Imperial College [8]. The main feature of the MP7 is a Xilinx Virtex 7 690 field-programmable gate array. Compared to the legacy system this chip provides a significant increase of resources for the implementation of the algorithm. The chip has about three times the amount of logic and seven times the amount of memory than the legacy system with its ten Virtex II chips [5]. The MP7 has 72 optical input and 72 optical output links that are operated at 10 Gb/s. Inputs and outputs can be connected on the front panel with four multifiber termination push-on connectors holding 36 fibers each.

2) *Firmware*: The current implementation of the algorithm uses half of the available logic and a third of the available block RAM. Parts of the logic run at a 240 MHz clock while others are required to operate at the 40 MHz LHC clock. The use of a higher clock at the input processing

¹The size of a calorimeter tower is $\Delta\eta \times \Delta\phi = 0.087 \times 0.087$.

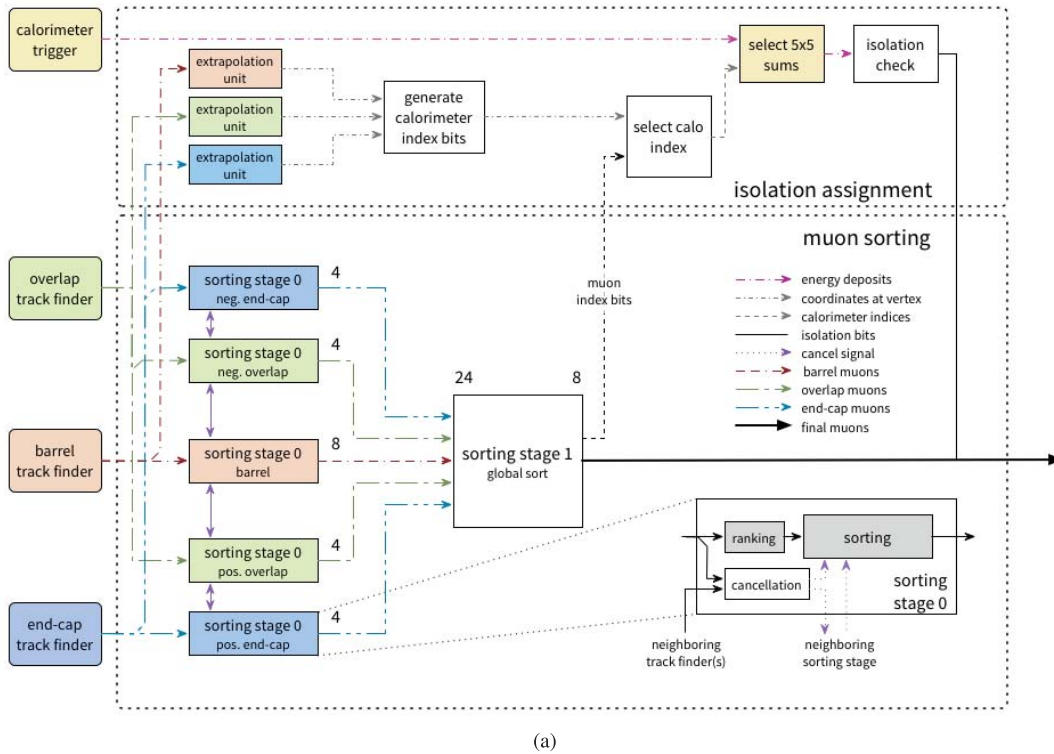


Fig. 3. Block diagram of the GMT algorithm. The top region contains the blocks for the isolation assignment. The bottom region shows the blocks for the muon sorting. The first sorting stage involves the blocks denoted stage 0 and the second sorting stage is named sorting stage 1 in the drawing. The inset on the bottom right shows one regional sorting block of the first sorting stage, which includes the rank assignment and duplicate muon identification.

reduces significantly the amount of memory needed as the same lookup tables can be used in a pipelined process for the successively arriving muons. Area constraints were needed for some blocks in order to achieve timing closure of the firmware. The total latency amounts to 250 ns. Half of this time is needed for the algorithm while the other half is used for deserialization/serialization.

B. Muon Sorting

In the upgrade muon trigger, the sorting of the input muons, which was performed by separate cards in the legacy system, has become a part of the GMT. This saves latency as a deserialization/serialization step can be omitted and the duplicate muon removal can be performed at an earlier stage.

The sorting in the GMT is performed in two stages.

The first sorting stage sorts the muons from the positive and negative side of the OMTF and the EMTF, as well as from the BMTF separately. Before the sorting, a rank is calculated for every incoming muon using the p_T and quality of the muon track. Duplicate muons, found by different TF processors, are identified for later removal in the sorting block. For the OMTF muons and the EMTF muons from each detector side the four with the highest rank are kept. Together with the highest ranked eight BMTF muons this leads to a maximum of 24 muons being sent to the second sorting stage.

The second sorting stage compares the ranks of the muons coming from the first stage and selects the eight with the highest rank to be sent to the GT.

C. Duplicate Muon Removal

Due to the overlap between adjacent wedges or sectors of the TFs a muon traversing the detector in these overlap regions can be found by the TF processors of both sides of the overlap. Since the TF processors do not know about muons found in their neighboring processors, in such cases two TFs will send a muon track to the GMT that stems from the same muon. In addition to this overlap in ϕ , the different regional TFs have also an overlap in η where a muon can be found by the BMTF and the OMTF, or by the OMTF and the EMTF. Such duplicate muons would significantly increase the trigger rate for double muon trigger algorithms and need to be removed while at the same time keeping a high efficiency for events with two muons.

The duplicate removal is performed at the GMT where the muon track information from all regional TFs is gathered for the first time. The following boundaries are considered for duplicate removal:

- 1) adjacent wedges or sectors for each TF;
- 2) BMTF and OMTF on both detector sides;
- 3) OMTF and EMTF on the positive detector side;
- 4) OMTF and EMTF on the negative detector side.

Two different methods are used for the identification of duplicates.

The first method makes use of the track address of the muon, which encodes the track segments used to build the muon track, to find duplicates between BMTF wedges. Two BMTF tracks from adjacent wedges are marked as duplicates if one




Fig. 4. Area in η and ϕ used for the calculation of the calorimeter energy sum. The arrow marks the calculated direction of the muon at the interaction point.

or more track segments are shared between two muon tracks. This method has the advantage that only duplicate muons are found, since two different muons do not have shared track segments.

The second method uses the muon track coordinates and is applied to find duplicates between adjacent sectors in the OMTF and the EMTF and between different regional TFS. Around each muon a region in η and ϕ is defined and, if there is a second muon track found within this region, the two tracks are marked as duplicates. The distance between the two tracks is calculated as $\Delta R = (f_\eta \Delta\eta^2 + f_\phi \Delta\phi^2)^{1/2}$, where $\Delta\eta$ and $\Delta\phi$ are the distances in η and ϕ , respectively, and f_η and f_ϕ are weighting factors for the two distances. In some cases, the BMTF is not able to calculate the η value of a track and it assigns a coarse eta value to the track based on the position of the stations that were used to build it. A bit is reserved in the incoming muon data to indicate if the eta coordinate has been measured with coarse or fine precision and a different weighting factor f_η can be chosen accordingly. Monte Carlo (MC) simulated events are used to determine the optimal size and shape of the regions in which tracks should be marked as duplicates. These regions can be different for the various boundaries.

If two tracks have been identified as duplicates, the one with the lower quality is marked for removal. In case of a tie in quality, the muons coming from the BMTF and the EMTF win over muons from the OMTF and muons from the wedge or sector on the lower ϕ side win over muons from their neighboring wedge or sector, except for the boundary at the ϕ wrap-around where muons from wedge 12 or sector 6 win over muons from wedge 1 or sector 1, respectively.

D. Isolation

For each output muon, the GMT calculates an absolute and a relative isolation bit, which are added to the muon. These

bits indicate if there is a calorimeter energy deposit that can be associated with the muon, which would indicate a nonprompt muon that could have been created within a particle jet. At the GT, these bits can then be used to reject such in-flight muons and achieve a rate reduction for prompt muons.

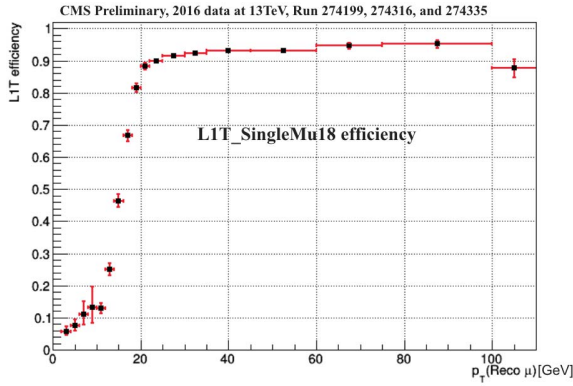
The GMT receives energy sums at a granularity of 2×2 calorimeter towers as shown in Fig. 4 from the calorimeter trigger. The energy sums are calculated over an area of 5×5 of the 2×2 calorimeter tower regions. The 2×2 calorimeter tower granularity is necessary since the input bandwidth is not sufficient to allow for finer segmented energies. In the ϕ direction 36 sums are received and in the η direction 28, resulting in a total of 1008 energy sums with a 5 b precision. In order to calculate the isolation bits the muon η and ϕ coordinates are extrapolated to the nominal interaction point at the center of the CMS detector using the muon p_T . While a muon that is produced in a particle jet follows a curved trajectory induced by the magnetic field of the CMS solenoid, with a curvature depending on the muon p_T , the other particles of the jet follow a more straight trajectory and deposit their energy at a different position in the calorimeter than where the muon is passing. Since muons within particle jets are usually produced close to the interaction point the extrapolation of the muon coordinates to that point is necessary. With the coordinates at the interaction point, the index of the corresponding energy sum is calculated and if the sum is below a predefined threshold the absolute isolation bit is set. The relative isolation bit is set if the energy sum at the index divided by the muon p_T is below a predefined threshold. The two isolation bits are merged with the output muons that are sent to the GT.

The described algorithm for the calculation of the isolation bits is currently implemented in the firmware. However, the input energy sums are not yet received from the calorimeter trigger the isolation bits are not yet used at the GT level. Ongoing studies indicate that a different algorithm for the isolation might result in a better performance, which would justify a future replacement of the algorithm currently implemented.

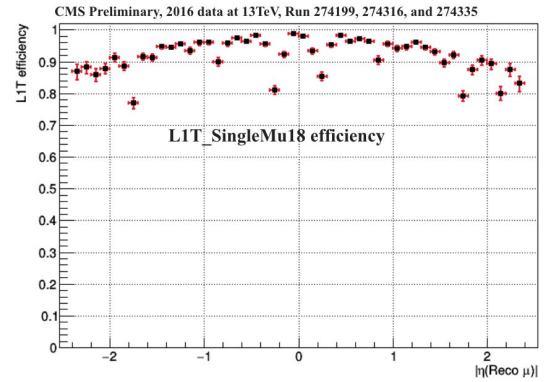
E. Commissioning of the Upgrade Trigger

During the 2015 run of the LHC the legacy trigger system provided the triggers for CMS. For the commissioning of the upgrade trigger system in parallel to the ongoing data taking the signals from the muon system were duplicated using splitters and mezzanine cards. Only a part of the DT detector data was available as input to the upgrade system in case of the BMTF and the OMTF, whereas for the EMTF the complete CSC data were available. During the 2015 heavy ion data, taking the timing of the trigger processors was adjusted and the TF algorithms were checked.

After the end of the year break of CERN, several cosmic runs were taken with all subsystems participating. At first, with the legacy trigger providing the triggers and the upgrade system running in parallel to adjust the timing. During this time, the DT inputs for the legacy muon trigger that were not duplicated were moved to the upgrade system. The switchover to the upgraded trigger system providing triggers for CMS was done in a single day, near the end of February 2016.



(a) Efficiency as a function of reconstructed muon p_T for a single muon trigger with a threshold of 18 GeV. The highest p_T bin contains all muons with $p_T > 100$ GeV.



(b) Efficiency as a function of reconstructed muon η for a single muon trigger with a threshold of 18 GeV and reconstructed muon $p_T > 22$ GeV.

Fig. 5. Efficiency for a single muon trigger with a threshold of 18 GeV. (a) Efficiency as a function of reconstructed muon p_T for a single muon trigger with a threshold of 18 GeV. The highest p_T bin contains all muons with $p_T > 100$ GeV. (b) Efficiency as a function of reconstructed muon η for a single muon trigger with a threshold of 18 GeV and reconstructed muon $p_T > 22$ GeV.

The subsequent weeks until the start of collisions of the LHC were used to test the new system and integrate the configuration in the CMS run control software. In parallel, the online data quality monitoring was set up to provide prompt feedback of the data taken and help spot and diagnose problems.

F. Results of Studies on Data Taken During the LHC Startup 2016

Since the start of regular pp collisions the data recorded are used to study the performance of the upgraded muon trigger. Trigger efficiencies are measured for a single muon trigger with a threshold of 18 GeV using a tag and probe technique with offline reconstructed muons from preselected Drell–Yan events decaying to the muon final state [10]. The data for the study shown here were taken during three runs of several hours duration between May 29, 2016 and June 2, 2016.

The tag muon is required to be a muon reconstructed with the CMS particle flow algorithm [9] with $p_T > 27$ GeV and an isolation variable lower than 0.1, matched to an HLT trigger object within $\Delta R < 0.3$,² that passed the HLT trigger for isolated muons with 24 GeV threshold, which is seeded by the single muon 22 GeV L1 trigger. Furthermore, the tag muon needs to fulfill the medium muon quality criteria, which require the muon to be reconstructed involving the inner tracker hits with a valid fraction greater than 0.8. The medium muon quality criteria also require the muon track to pass at least one of the two selections detailed in the following lines. The reconstructed muon track must have a segment compatibility value passing a tight selection. Muon tracks fulfilling only a looser segment compatibility selection are accepted if the track also includes hits from the muon system and passes several additional reconstruction quality criteria.

The probe muon, for which the efficiency is measured, must fulfill the medium muon criteria as well. Furthermore, the tag and the probe must be in different regions of the detector.

$$^2 \Delta R = \sqrt{\Delta \eta^2 + \Delta \phi^2}$$

For a probe muon with $|\eta| < 1$, the tag muon must be in the EMTF region ($|\eta| > 1.2$), while for a probe with $|\eta| \geq 1$ the tag muon must be in the BMTF region ($|\eta| < 0.8$). This guarantees that the tag and the probe are two different muons. To ensure that the tag muon would have fired the L1 trigger a TF muon required to fulfill the single muon quality criterion and having a $p_T \geq 18$ GeV must be found in the region of the tag.

To calculate the efficiency, a L1 muon from the triggering bunch crossing, coming from the GMT, must be found matching the probe muon within $\Delta R < 0.3$. Fig. 5(a) shows the turn on efficiency as a function of the probe muon p_T . At the threshold value, the efficiency reaches about 86% of the plateau efficiency of about 93%. The efficiency as a function of probe muon η is shown in Fig. 5(b). Compared to the legacy trigger the efficiencies found are slightly lower by 2%–3%. It is expected, however, that the efficiency in the endcap regions up to $|\eta| = 1.6$ will increase once the RPC information is forwarded and used in the EMTF.

The muon trigger rate has been studied with a zero bias data sample taken with a prescaled trigger that only requires a colliding bunch for triggering. For the single muon trigger with 18 GeV threshold, the rate has been found to be 3.7, 5.1, and 5.7 kHz for a luminosity of $4.7 \times 10^{33} \text{ cm}^{-2} \text{ s}^{-1}$, $6.7 \times 10^{33} \text{ cm}^{-2} \text{ s}^{-1}$, and $7.2 \times 10^{33} \text{ cm}^{-2} \text{ s}^{-1}$, respectively. This is roughly a rate reduction of 30% compared to the legacy trigger system as was estimated from studies with MC simulated samples. The rate reduction improves when going to higher trigger thresholds giving a good handle for tuning in higher instantaneous luminosity conditions.

IV. CONCLUSION

The CMS L1 trigger system was upgraded for the 2016 run of the LHC. The system was implemented using the MicroTCA standard and uses common hardware designs for the components. The muon trigger moved to a layout based on the detector regions, using the available muon subsystems in a region to build muon tracks. The GMT absorbs the regional

sorting stage and performs a ranking and duplicate muon removal. After a second sorting stage, the highest ranked eight muons are sent to the GT. When commissioned, the GMT will add absolute and relative isolation information, calculated with calorimeter energy sums from the calorimeter trigger, to the muons sent to the GT.

During the 2015 heavy ion run and in cosmics data taking before the start of the LHC in 2016 the upgrade trigger system was running in parallel to the legacy system for commissioning. The trigger of CMS switched to the upgrade system in February 2016 before the collision run of the LHC.

The analysis of data taken at the beginning of the 2016 pp run shows that the efficiency reaches about 93%, which is slightly lower than in the legacy system, for a single muon trigger with an 18 GeV threshold. The efficiency is expected to improve once the hits in the RPCs can be used in the endcap region. At the same time, the rate could be reduced compared to the legacy trigger, which is important to cope with the increasing luminosity that the LHC will provide until the end of Run-II and beyond.

Further improvements are expected from tuning the matching areas of the duplicate removal and from using the added isolation information of the muons once the inputs needed for the calculation are provided to the GMT.

REFERENCES

- [1] CMS Collaboration, “CMS technical design report for the level-1 trigger upgrade,” CERN, LHC Experiments Committee, Geneva, Switzerland, Tech. Rep. CERN-LHCC-2013-011 CMS-TDR-12, Aug. 2013.
- [2] J. Manfred, J. Lingemann, D. Rabady, H. Sakulin, and A. Stahl, “Upgrade of the CMS global muon trigger,” *IEEE Trans. Nucl. Sci.*, vol. 62, no. 3, pp. 1104–1109, Jun. 2015.
- [3] A. Tapper and D. Acosta, Eds., “The performance of the CMS muon detector in proton-proton collisions at $\sqrt{s} = 7$ TeV at the LHC,” *J. Instrum.*, vol. 8, no. 11, p. P11002, Nov. 2013.
- [4] S. Chatrchyan *et al.*, “The CMS experiment at the CERN LHC,” *J. Instrum.*, vol. 3, no. 8, p. S08004, Aug. 2008.
- [5] S. Dasu *et al.*, “CMS. The TriDAS project. Technical design report: The trigger systems,” CERN, LHC Experiments Committee, Geneva, Switzerland, Tech. Rep. CERN-LHCC-2000-038, Dec. 2000, vol. 1.
- [6] S. Jamieson. (Sep. 2006). Micro telecommunications computing architecture short form specification. PCI Industrial Computer Manufacturers Group. [Online]. Available: https://www.picmg.org/wp-content/uploads/MicroTCA_Short_Form_Sept_2006.pdf
- [7] E. Hazen, A. Heister, C. Hill, J. Rohlf, S. X. Wu, and D. Zou, “The AMC13XG: A new generation clock/timing/DAQ module for CMS MicroTCA,” *J. Instrum.*, vol. 8, no. 12, p. C12036, Dec. 2013.
- [8] K. Compton *et al.*, “The MP7 and CTP-6: Multi-hundred Gbps processing boards for calorimeter trigger upgrades at CMS,” *J. Instrum.*, vol. 7, no. 12, p. C12024, Dec. 2012.
- [9] CMS Collaboration, “Particle-flow event reconstruction in CMS and performance for jets, Taus, and MET,” CERN, Geneva, Switzerland, Tech. Rep. CMS-PAS-PFT-09-001, Apr. 2009.
- [10] CMS Collaboration, “Pre-LHCP 2016 L1 muon trigger performance,” CERN, Geneva, Switzerland, Tech. Rep. CMS-DP-2016-024, Jun. 2016.