
Upgrade of the Global Muon Trigger at the CMS Experiment

Von der

Fakultät für Mathematik, Informatik und Naturwissenschaften der
RWTH AACHEN UNIVERSITY

*zur Erlangung des akademischen Grades eines
Doktors der Naturwissenschaften genehmigte Dissertation*

vorgelegt von

Diplom-Physiker
Joschka Phillip LINGEMANN
aus Siegen

Berichter:

Universitätsprofessor
Dr. Achim STAHL
Universitätsdozent
Dr. Claudia-Elisabeth WULZ

Tag der mündlichen Prüfung: 14. September 2016

Diese Dissertation ist auf den Internetseiten
der Universitätsbibliothek online verfügbar.



Contents

Contents	ii
1 Introduction	1
2 Physics and Experiments at the Large Hadron Collider	3
2.1 Theoretical Motivation	3
2.1.1 Overview of the Standard Model of Particle Physics	3
2.1.2 Proton-Proton Collisions	5
2.2 The Large Hadron Collider	7
2.2.1 The LHC Accelerator Chain	7
2.2.2 Luminosity	9
2.2.3 Experiments at the LHC	9
2.2.4 Performance during run-1 and perspectives	10
2.3 The CMS Experiment	11
2.3.1 Magnet	11
2.3.2 Inner Detector	12
2.3.3 Electromagnetic Calorimeter	13
2.3.4 Hadron Calorimeter	14
2.3.5 Muon Detector System	16
2.3.6 The CMS Trigger and Data Acquisition System	18
2.3.7 CMS Detector Upgrades during Long Shut-down 1	19
2.4 CMS Software	19
3 Overview of the L1-Trigger and its Upgrade	21
3.1 Motivation	21
3.2 Common Upgrade Electronics	22
3.3 The Calorimeter Trigger	23
3.3.1 Legacy Calorimeter Trigger	23
3.3.2 Upgraded Calorimeter Trigger	24
3.4 The Muon Trigger	26
3.4.1 Legacy Muon Trigger	26
3.4.2 Upgraded Muon Trigger	28
3.5 The Global Trigger	30
3.5.1 Legacy Global Trigger	31
3.5.2 Upgraded Global Trigger	31
4 Trigger Algorithm Performance Studies	33
4.1 Monte Carlo Simulated Event Samples	33
4.2 Event Selection	33

4.2.1	Efficiency Event Sample	33
4.2.2	Rate Sample	34
4.3	Rate and efficiency estimation	34
4.4	Track-Finder performance	35
4.4.1	Status of the track-finding emulation	35
4.4.2	Categorisation of Track-Finder Muons	36
4.4.3	Performance of the standalone Track-Finders	36
4.4.4	Rate Suppression with Quality Selection	40
4.4.5	Comparison with Legacy System	43
4.5	Isolation in the Upgraded GMT	46
4.5.1	Available Calorimeter Information	47
4.5.2	Tower Calibration and Muon Extrapolation	47
4.5.3	Energy Distribution	50
4.5.4	Estimation of Performance	53
4.5.5	Application in the level-1 trigger	61
5	The Upgraded Global Muon Trigger	65
5.1	Input Data	65
5.1.1	Muon Data	65
5.1.2	Calorimeter Data	66
5.2	Algorithm Blocks	67
5.2.1	Sorting and Cancellation Unit	68
5.2.2	Isolation Unit	69
5.3	Gateware	70
6	Online Software	73
6.1	Common Software Components of the Upgraded L1-Trigger	73
6.2	Upgraded GMT Configuration	74
6.3	Upgraded GMT Monitoring	75
7	Emulator and Hardware-Tests	77
7.1	CMSSW Level-1-Trigger Emulation	77
7.2	The Upgraded Global Muon Trigger Emulator	77
7.2.1	Implementation	77
7.2.2	Data Formats	78
7.3	Evaluation of Hardware Response	78
7.3.1	Test Sample Generation	78
7.3.2	Hardware Tests	79
7.3.3	Results	80
8	Conclusion	83
	Bibliography	85
	Abbreviations	89
	List of Figures	90

1 Introduction

The Compact Muon Solenoid (CMS) is one of the two general purpose detectors at the Large Hadron Collider (LHC). As such, it has a broad physics program that spans from precision measurements to searches for exotic particles predicted by extensions of the standard model of particle physics. Given the rate of collisions in the LHC and technical limitations, not all collision events can be recorded. To avoid losing information that is of interest for the physics programme, an elaborate triggering system is employed that reduces the rate of events to be recorded to manageable levels. To that end, a staged scheme is deployed where the first stage is based on custom built hardware that discards events at the time scale of subsequent LHC collisions, and the second stage is based on more complex filtering algorithms that run on a computer farm.

After a successful first physics data taking period from 2009-2013 at 3.4 and 4 TeV collision energy, the LHC started a 2-year programme of maintenance and upgrades to bring the machine back with almost the design collision energy of 6.5 TeV and higher luminosity. With these changed running conditions, the experiments' triggering schemes have to be revised. To provide the same acceptance for events that are of interest, while increasing the fraction that can safely be rejected, new algorithms have to be developed and the overall triggering strategy has to be reconsidered in the level-1 muon trigger. Information from all three muon detectors of CMS are used in the level-1 trigger. Combining this information at an early stage to ensure the maximum amount of information is available for track reconstruction, is one essential change in the upgrade.

This thesis presents the concept of and performance studies motivating the algorithm of the Micro-Global Muon Trigger (μ GMT). The μ GMT is a central part of the muon trigger; it is the first stage at which muon track candidates from the whole muon detector are available. This means that duplicate tracks have to be identified and handled to avoid double counting. The final trigger stage, the Micro-Global Trigger (μ GT), is capable to calculate topological relations between different trigger objects, but object multiplicities have to be low enough to keep combinatorics in check. Therefore, the μ GMT algorithm foresees sorting the muon candidates according to their momentum and the quality of the momentum measurement.

Lastly, the μ GMT combines calorimeter with muon track information. As the bandwidth available for information exchange and computing resources are a limiting factor in the level-1 trigger, the calorimeter information is received with reduced granularity. Algorithms have to be, therefore, carefully considered and a negative influence of this reduced granularity has to be minimised. With the calorimeter information, an additional

handle is introduced that allows to suppress the trigger rate further while maintaining a high signal efficiency.

Verifying the hardware implementation against an independent software implementation of the same algorithms is an important part during development and also data taking. This software implementation also allows to change aspects of the existing or even to develop new algorithms without the need to disrupt data taking. Also monitoring and control software is an essential tool to ensure successful physics data taking. For example, many parts of the algorithm can be adjusted through modifiable look up tables. The contents of these can be changed through the control software and are checked as part of the monitoring.

Structure

The present thesis is composed of eight chapters, this introduction being the first. The second chapter gives an overview of the standard model of particle physics and the experimental facilities at the Conseil Européen pour la Recherche Nucléaire (CERN) with an emphasis on the CMS experiment. In chapter three, the level-1 trigger before the upgrade is compared to the components that are foreseen, also giving a general overview of the technical capabilities of the hardware components. The fourth chapter contains the presentation of studies that estimate the performance of trigger algorithms at different stages in the muon trigger. As the μ GMT algorithm receives muon candidates from the track-finders, their performance presents the basis for estimating the quality of any μ GMT algorithm. Therefore, the individual track-finders are first studied independently. Then aspects in the sort algorithm of the μ GMT are considered that could be used to increase the number of rejected events while keeping a high signal event efficiency. The chapter is closed by studies that show that the usage of calorimeter information can improve the performance of the muon trigger system. In chapter five, the μ GMT algorithm that currently is implemented in hardware and emulation software is discussed. The sixth chapter introduces the control and monitoring software for the μ GMT. This software was used to verify and test the hardware implementation of the previously described algorithms and the results of these tests are described in chapter seven. The text is then concluded and a short outlook is given.

2 Physics and Experiments at the Large Hadron Collider

This chapter outlines the basics of the Standard Model of particle physics and introduces the experimental facilities that are relevant for this work. Particularly, the LHC and its pre-accelerator chain are described and the run-parameters of the machine and their influence on the luminosity are introduced. In the last part of the chapter, the CMS experiment and its individual detector components and the utilised software are described.

2.1 Theoretical Motivation

In section 2.1.1 the particle content of the Standard Model of particle physics is described. Section 2.1.2 introduces terms that are specific for proton-proton collisions and highlights the most important muon production processes. The following is largely taken from refs. [1, 2].

2.1.1 Overview of the Standard Model of Particle Physics

The Standard Model of particle physics describes the current understanding of the constituents of matter and their interactions. Both the constituents and the interactions can be represented by particles. The constituents are fermions that carry a spin of $1/2$, the mediators of interactions on the other hand are bosons that carry integer spin of 0 or 1.

The fermions can be further split into quarks and leptons which come in three different generations. The three electrically charged leptons are accompanied by neutral partners, the neutrinos. These additional leptons are massless in the Standard Model. However, the discovery of neutrino oscillations necessitated the introduction non-zero masses for at least two neutrinos [3].

Table 2.1 shows the bosons of the Standard Model, the symbols used as abbreviations throughout this text, the spin and electric charge they carry, as well as the interactions they mediate. The last column also indicates which interactions they are subject to.

The massless photon is the mediator of the electromagnetic interaction between particles that carry an electric charge. The electromagnetic interaction is described by the relativistic quantum field theory called Quantum Electro Dynamics (QED).

TABLE 2.1 Bosons (force mediator particles and the Higgs boson) of the Standard Model.

Name	Symbol	Spin	El. charge (e)	Interaction	Subject to
Photon	γ	1	0	EM	-
W boson	W^\pm	1	± 1	weak	weak, EM
Z boson	Z	1	0	weak	weak
Gluon	g	1	0	strong	strong (self)
Higgs boson	H	0	0	-	weak, Higgs (self)

TABLE 2.2 Fermions in the Standard Model.

	Name	Symbol	El. charge (e)	Subject to	Generation
Leptons	Electron	e	-1	EM, weak	1
	Muon	μ	-1	EM, weak	2
	Tau	τ	-1	EM, weak	3
	Electron neutrino	ν_e	0	weak	1
	Muon neutrino	ν_μ	0	weak	2
	Tau neutrino	ν_τ	0	weak	3
	Quarks	Up	u	2/3	EM, weak, strong
Down		d	-1/3	EM, weak, strong	1
Charm		c	2/3	EM, weak, strong	2
Strange		s	-1/3	EM, weak, strong	2
Top		t	2/3	EM, weak, strong	3
Bottom		b	-1/3	EM, weak, strong	3

The W bosons, i.e. the W^- boson and its anti-particle, the W^+ boson, together with the Z boson mediate the weak interaction. Their relatively high masses of $m_W = 80.385 \pm 0.015 \text{ GeV}^1$ and $m_Z = 91.1876 \pm 0.0021 \text{ GeV}$ [4] indicate that the reach of this interaction is small. Since the W bosons carry electric charge, they also participate in the electromagnetic interaction.

Gluons are the force carrying particles of the strong interaction, described by Quantum Chromo Dynamics (QCD). They are colour octets, 8 linearly independent states with different colour charges, colour being the charge of the strong interaction. The gluons themselves carry two colour charges resulting in self-interactions. Due to this self-interaction, quarks as well as gluons can not be observed as free particles, an effect called colour-confinement.

The Higgs boson was the last particle of the Standard Model to be observed. In July 2012, both the ATLAS and CMS collaborations published the observation of a new particle with a mass around 125 GeV that is compatible with a Standard Model Higgs boson [5, 6]. In the frame of the Standard Model, through the mechanism of electroweak symmetry breaking, the Higgs field gives rise to the W and Z boson masses. Fermions gain masses through Yukawa couplings to the Higgs field.

¹ In this work natural units are used which are defined with $\hbar = c = 1$. The typically used unit of energy in this framework is eV. This results in energies and masses being measured in eV, while time and lengths have the unit 1/eV.

Table 2.2 shows all fermions of the Standard Model and the symbols used throughout this text. The electric charge and the interactions they are subject to are given as well as their generations. The quarks can be divided in up-type quarks, up (u), charm (c) and top (t) quark of charge $2/3$ in units of the electron charge (e) and down-type quarks down (d), strange (s) and bottom (b) of charge $-1/3$. The corresponding antiparticles carry opposite electric charge. All quarks participate in weak, electromagnetic and strong interactions. The three generations of charged leptons electron (e), muon (μ) and tau (τ) all carry the charge -1 in units of e , while their antiparticle carry positive charge. The corresponding neutrinos do not carry any electric charge. Charged leptons participate both in electromagnetic and weak interactions, whereas neutrinos only participate in the weak interaction.

All other particles observed can then be described as bound states of these elemental particles. These bound states of quarks are called hadrons and can be grouped into baryons (three bound quarks) and mesons (one quark and one anti-quark). These bound states are bound by the strong interaction.

While many predictions of the Standard Model have been confirmed by experiments to great precision, some open questions remain. For example, driven by the observed matter distribution in the universe, additional non-luminous matter, Dark Matter, has to be introduced which has no corresponding particle in the Standard Model. The explanation of the absence of antimatter needs additional CP violating processes. Another important example is the hierarchy problem, the extreme difference between the coupling constants of gravity and the weak force, which could be explained by Super Symmetry or Extra Dimensions. These and other theories that go beyond the Standard Model predict particles of high masses or that are difficult to detect [7].

2.1.2 Proton-Proton Collisions

At the LHC, protons are collided with protons. Cross sections, however, are calculated at parton level, i.e. for gluon and quark interactions. Therefore, the energy of the partons within the proton need to be known. The parton distribution functions (PDFs) are needed for this estimate. A PDF can be written as $f_i(x, Q^2)$ with $i = u, d, s, c, b, g$, the momentum transfer Q^2 and the momentum fraction x carried by the parton. The structure of hadrons is studied in deep inelastic collisions of leptons and hadrons. Here, the momentum of the lepton before and after the collision, gives the transferred four-momentum q , where $Q^2 = -q^2$, and allows to measure the structure function $F(x, Q^2)$ of the hadron, the sum over all PDFs. The PDFs can be interpreted as the probability to find a parton carrying the momentum fraction x given a momentum transfer.

The number of events N is given by the luminosity L of a collider and the cross section σ of a process:

$$N = \sigma L \tag{2.1}$$

Fig. 2.1 shows the cross sections of various processes as a function of the collider centre-of-mass energy. As one can see from the figure, the total cross section at the LHC is dominated by jet production. Standard Model physics like Z or W boson production are already a factor of 10^{-6} smaller and the Higgs boson production cross section is




FIGURE 2.1 Cross sections σ for various physics processes at hadron colliders as a function of the collision energy [8].

again much smaller. Furthermore, one can see that the total cross section grows with increasing centre-of-mass energy.

Simulating collision events. – The hard interaction for simulated events are chosen based on the matrix element calculation [9]. The final state particles are then fed to a parton shower simulation, an approximation of further splittings of quarks and light quarks into two partons in the soft and collinear limit. Overlaps between the two procedures are then removed with matching methods. In the following hadronisation step, phenomenological models are used to group the resulting partons into hadrons.

In addition to this hard interaction, the so-called underlying event is simulated: The coloured remnants of the two colliding protons are also considered in the hadronisation step. Also further interactions of additional partons from the two protons are considered and lead to multiple parton-parton interactions. Lastly, pile-up collisions are simulated that arise from additional proton-proton collisions that occurred simultaneously. For this last step usually all possible production processes at the collision energy, so-called minimum bias events, are superimposed with the hard process.

2.1.2.1 Muon Production

In proton-proton collisions, muons can be produced through various intermediate states. The momentum of the muon is highly dependent on this intermediate state. Di-muons are mainly produced in Drell-Yan processes that include the production of Z bosons and




FIGURE 2.2 Exemplary Feynman diagrams of processes that have muons in the final state. On the left and right, W and Z boson production with subsequent muonic decay. In the middle the decay of a B meson in the spectator view.

photons or light neutral meson states, such as J/Ψ or ρ mesons. Single muons on the other hand are produced through the decay of W bosons or charged hadron states such as B mesons or other quarkonia.

For low energetic muons, the production through the light hadron states is dominant. When going to higher momenta and, therefore, higher intermediate masses, the production through Z and W bosons becomes more important. At even higher muon momenta other sources dominate, such as top-antitop quark production with subsequent decay $t \rightarrow Wb \rightarrow \mu\nu b$.

2.2 The Large Hadron Collider

The LHC [10] at the European centre for particle and nuclear research (CERN) is a circular hadron collider designed to provide the highest luminosity and collision energy to date. The largest fraction of the LHC physics program is devoted to discover new physics beyond the standard model and measure the properties of the Higgs boson. Therefore, the machine is designed to provide proton-proton collisions with a centre-of-mass collision energy of up to $\sqrt{s} = 14$ TeV and an instantaneous luminosity of $\mathcal{L} = 10^{34} \text{ cm}^{-2}\text{s}^{-1}$. Additionally, it is possible to inject ionised lead, allowing collisions at an instantaneous luminosity of $\mathcal{L} = 10^{27} \text{ cm}^{-2}\text{s}^{-1}$. Nominally, collisions occur at 40 MHz or in other words every 25 ns.

In section 2.2.1 the accelerator complex is described and in section 2.2.2 the luminosity in terms of the machine parameters is discussed and follows ref. [10]. In section 2.2.3 an overview of the experiments at the LHC is given. Lastly, section 2.2.4 shows the overall performance during the first proton-proton collisions run in 2009-2012 and the perspective for the second running period of the LHC (run-2).

2.2.1 The LHC Accelerator Chain

Before reaching the final energy in the LHC, the protons have to be accelerated in several steps. The complete injection chain is shown in Fig. 2.3. Protons are injected into the machine in so-called bunches, packets of protons with high density. In the following the acceleration scheme is described for the nominal proton-proton operation at a collision frequency of 40 MHz.

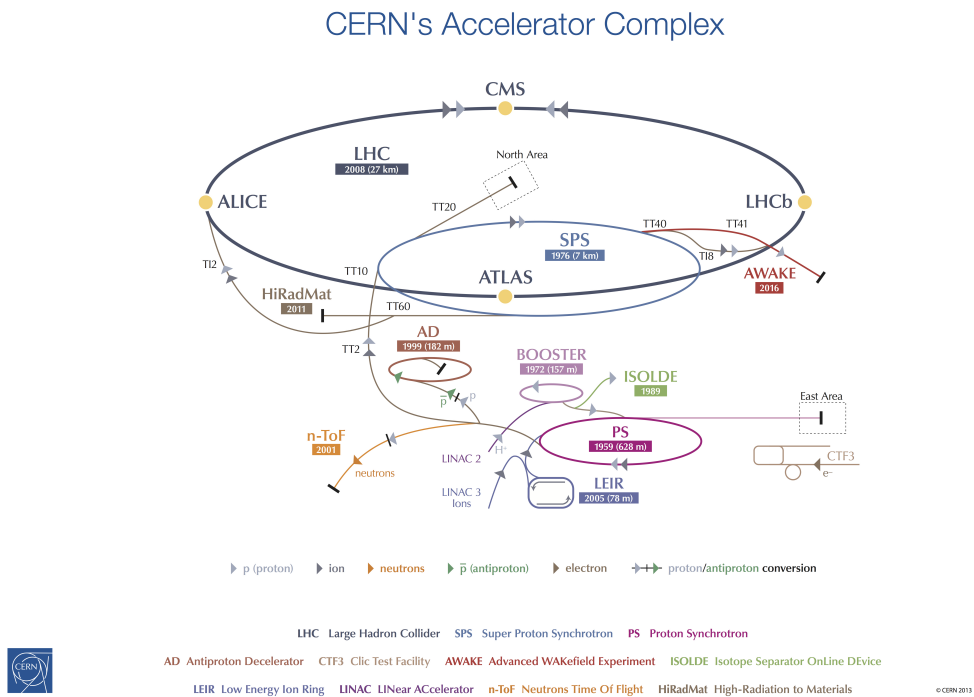


FIGURE 2.3 The LHC accelerator complex, shown are also the experimental areas provided with proton or ion beams along the accelerator chain [11].

The proton supply chain starts at the linear accelerator Linac2 which increases the proton energy up to 50 MeV. The bunches are then injected into the Booster which accelerates them further to 1.4 GeV and injects into the Proton Synchrotron (PS). Here, each of the proton bunches is split into 12, before being supplied to the Super Proton Synchrotron (SPS). One fill of the SPS consists of a total of 72 bunches, 16 PS fills, which are then supplied to the LHC. One such SPS fill is referred to as a bunch train. A total of 39 such trains is required to complete the LHC fill of 2808 bunches. The maximum energy of a circular collider is limited by the maximum magnetic field in the steering dipoles. The LHC is instrumented with superconducting dipole magnets based on NbTi Rutherford cables with a maximum field of 8.33 T. To operate the magnets at these high fields heat depositions and, therefore, beam losses have to be minimised. This means that tight controls on magnetic field errors have to be in place. The maximum energy achievable in the LHC with this magnetic field is 7 TeV per beam.

Fig. 2.3 shows also that the injection chain allows extraction at several energy steps to supply proton and ion beams to various experiments along the supply chain. These include the Antiproton Decelerator (AD), the Compact Linear Collider Test Facility (CTF3), the Isotope Separator OnLine DEvice (ISOLDE), the Neutrons Time of Flight (n-ToF), and the High-Radiation to Materials (HiRadMat) experiments. Descriptions of the physics programs and design of these experiments may be found in refs. [12–16]. The Low Energy Ion Ring (LEIR) is part of the lead ion injection chain.

2.2.2 Luminosity

Beam energy aside, the luminosity of an accelerator is the most important variable for a successful physics program, cf. eq. 2.1 and Fig. 2.1. The instantaneous luminosity can be calculated from a number of machine parameters,

$$\mathcal{L} = \frac{N_b^2 n_b f_{\text{rev}} \gamma_r}{4\pi \epsilon_n \beta^*} F, \quad (2.2)$$

namely the particles per bunch N_b , the bunches per beam n_b , the revolution frequency f_{rev} , the relative gamma factor γ_r , the normalised transverse emittance ϵ_n , the beta function at the interaction point β^* , and the geometric luminosity reduction factor due to the crossing angle F . This factor is given by

$$F = \left(1 + \left(\frac{\theta_c \sigma_z}{2\sigma^*} \right)^2 \right)^{\frac{1}{2}}, \quad (2.3)$$

with the full crossing angle at the interaction point θ_c , the root-mean-square (RMS) of the bunch length σ_z , and the transverse RMS of the beam size at the interaction point σ^* . Equations 2.2 and 2.3 require round beams, $\sigma_z \ll \beta$.

From these equations one can see that in order to achieve the high design luminosity of $\mathcal{L} = 10^{34} \text{ cm}^{-2}\text{s}^{-1}$, high beam intensities have to be achieved. These high densities can not be achieved with antiprotons. To allow the proton-proton collisions, two separate magnetic dipole fields are needed to steer the particles through the machine. The beams are separated in individual beam pipes and brought together at the interaction points. Specialised steering dipoles and focusing magnets are installed at each interaction point to bring highly focused beams to collision.

Over the course of a run the instantaneous luminosity decays due to beam losses. Beam losses arise from the collisions themselves, interactions with rest gas in the beam pipe and intrabeam scattering. The integrated luminosity L is, therefore, given by

$$L = \mathcal{L}_0 \tau_L \left(1 - e^{-T_{\text{run}}/\tau_L} \right), \quad (2.4)$$

with the initial instantaneous luminosity \mathcal{L}_0 , the luminosity lifetime τ_L and the integrated run time T_{run} . The integrated luminosity is optimised by taking the degradation of the instantaneous luminosity and the turn-around time of an LHC refill into account. This results in an optimal run time after which the beams are dumped in a controlled manner and the LHC is refilled.

2.2.3 Experiments at the LHC

Five experiments are provided with collision events by the LHC. Two general purpose experiments with physics goals spanning from measurements of Standard Model parameters to new physics at the TeV-scale: A Toroidal LHC ApparatuS (ATLAS) and the CMS experiment at opposite sides of the LHC ring. The CMS experiment will be described in greater detail in section 2.3 and a description of the ATLAS detector may be




FIGURE 2.4 The LHC performance during first running period of the LHC (run-1). On the left the peak instantaneous luminosity (and rate of collisions) is shown for running in 2010 (green, multiplied by 20), 2011 (blue) and 2012 (red). The right hand figure shows integrated luminosity (and number of collisions) for the 3 years in the same colour coding [22].

found in ref. [17]. Both ATLAS and CMS are accompanied by smaller detectors that measure the total proton-proton cross section and elastic and diffractive scattering. The TOTAL cross section, Elastic scattering and diffraction dissociation Measurement at the LHC (TOTEM) experiment [18] is installed on either side of CMS and the LHCf (LHC forward) experiment [19] on either side of ATLAS. A Large Ion Collider Experiment (ALICE) is a specialised detector for heavy ion collisions and its goals are focused on physics of strongly interacting particles, especially the quark-gluon plasma at extreme values of energy density [20]. The LHCb experiment is dedicated to measurements of B-decays and CP-violation [21].

2.2.4 Performance during run-1 and perspectives

During 2010 first proton-proton collisions at a beam energy of 3.5 TeV were performed, preceded by a first test run at a lower energy of 1.2 TeV in December 2009 [22]. In 2010 the number of bunches was steadily increased from 2 in the very beginning up to 109 at the end of the proton run, corresponding to small luminosities, cf. Fig. 2.4. In 2011 the machine was re-commissioned after an annual winter shut-down and afterwards the number of bunches was further increased from 32 to 1380 bunches with a bunch spacing of 50 ns. At the very end of the proton run, a test fill was performed with 25 ns bunch spacing. A total integrated luminosity of 5.1 fb^{-1} was provided both to ATLAS and CMS. In 2012 the beam energy was increased to 4 TeV and the decision was taken to maximise the integrated luminosity to allow for a maximal dataset for Higgs-boson studies [5, 6]. The default number of bunches per beam in 2012 was 1374 with a bunch spacing of 50 ns. The target intensity at run start was 1.6×10^{11} protons. In a very successful run with more than 30% of the time spent in stable beams, a total integrated luminosity of 23.3 fb^{-1} could be delivered to ATLAS and CMS, cf. Fig. 2.4. To maximise the proton-proton collision data collected, the heavy ion run was pushed back to January and February 2013, which concluded the first running period of the LHC.

After February 2013 the first Long Shut-down 1 (LS-1) of the LHC started with a program of maintenance for the machine [23]. The LS-1 was concluded in the beginning of 2015 with first beams circling the collider in March. The aim for run-2 is to further

TABLE 2.3 Expected parameters for run-2 of the LHC and the resulting luminosities. From left to right: bunch spacing, number of bunches n_b , the transverse emittance ϵ_n , protons per beam N_b , the average number of pile-up events PU , peak instantaneous luminosity \mathcal{L}_{peak} , and integrated luminosity L per year [23].

Spacing (ns)	n_b	ϵ_n (μm)	N_b (10^{11})	PU	\mathcal{L}_{peak} ($10^{34} \text{ cm}^{-2}\text{s}^{-1}$)	L (fb^{-1}/yr)
50	1404	2.0	1.4	61	1.35	40
25	2808	3.0	1.2	30	1.30	38
50	1404	2.0	1.7	87	1.90	56
25	2520	1.3	0.7	23	1.00	29
25	2592	1.4	1.15	63	2.30	70

increase the beam energy, starting with 6.5 TeV and possibly going to the design energy of 7 TeV. The bunch spacing can either be 50 ns or 25 ns. Table 2.3 shows possible configurations and the resulting instantaneous luminosity. The number of pile-up events were already a concern during run-1 and would increase further if the machine were to continue running at 50 ns, as the table indicates. The maximum peak instantaneous luminosity could go as high as $2.3 \times 10^{34} \text{ cm}^{-2}\text{s}^{-1}$, leading to an integrated luminosity of 70 fb^{-1} delivered to ATLAS and CMS per year.

2.3 The CMS Experiment

In this section the different components of the CMS detector will be described. The first section will review the magnetic field configuration. The remaining parts of the experiment will be discussed in the order of their respective positions, starting with the tracker system, or inner detector, then going outward to the electromagnetic and hadron calorimeters. Lastly the outermost component, the muon system, will be described and the chapter is closed by a short description of trigger and data-acquisition systems. The following is largely taken from refs. [24–26].

The right-handed CMS coordinate system has its origin within the experiment at the nominal collision point. The y -axis points vertically upward and the x -axis radially toward the centre of the LHC. Thus, the z -axis points along the beam axis. The azimuthal angle ϕ is measured from the x -axis in the x - y -plane and the polar angle θ from the z -axis. The pseudo-rapidity is defined as $\eta = -\ln \tan(\theta/2)$. Momentum and energy transverse to the beam, p_T and E_T , respectively, are computed from the corresponding x and y components. An imbalance of energy in the transverse plane is labelled as missing transverse energy, E_T^{miss} .

2.3.1 Magnet

The magnetic field is optimised to allow for a good momentum measurement for muons in the outermost part of the detector and charged particles in general in the inner tracker. The stronger the magnetic field, the stronger the bending power for charged particles and, thus, the precision of measurements of their momentum. Therefore, the experiment




FIGURE 2.5 True to scale perspective view of the CMS detector with the interaction point in the centre and all individual detector components labelled [24].




FIGURE 2.6 An r - z cross section of the inner detector which consists of the Pixel, the Tracker Inner Barrel (TIB), the Tracker Outer Barrel (TOB), the Tracker Inner Disk (TID) and the Tracker End-Cap (TEC) detectors [25].

is instrumented with a superconducting magnet with a field strength of 3.8 T. The bore of the magnetic coil encompasses the tracking system and both the electromagnetic and the hadron calorimeter. The return field is strong enough to saturate the iron return yoke which is interleaved with four muon stations.

2.3.2 Inner Detector

The inner detector or tracker is designed to allow a good resolution for momentum measurements and identification of charged particles which should also allow tagging of both τ s and b -jets. Therefore, the innermost detector is installed as closely as possible to the interaction region. In Fig. 2.6 a schematic view of the tracker volume is depicted. Given the density of the particle flux three regions can be defined in the tracker volume.

- ▷ Inner region, $r \approx 10$ cm: the particle flux is highest the instrumentation has the highest granularity with pixels covering an area of $100 \times 150 \mu\text{m}^2$ each.

- ▷ Intermediate region, $20 < r < 55$ cm: the flow is low enough and micro-strip technology with a cell size of $10 \text{ cm} \times 80 \text{ } \mu\text{m}$ is sufficient.
- ▷ Outer region, $r > 55$ cm: the lowest particle density allows larger pitch silicon strips with a size of $25 \text{ cm} \times 180 \text{ } \mu\text{m}$.

With this instrumentation the occupancy for proton proton collisions ranges from 10^{-4} (inner region) to 1-3% (intermediate and outer regions) per pixel per bunch crossing. In lead-lead collisions, the particle flow is much denser and the occupancy ranges between 1% and 20% for inner region and outer region, respectively.

The Pixel detector, located in the inner region, is divided into barrel and two end-caps. The barrel consists of three layers, located at the radii of 4.4 cm, 7.3 cm and 10.2 cm, with a total of 768 pixel modules. The end-caps consist of 2 disks each, which cover the region of $6 \text{ cm} < r < 15 \text{ cm}$ and are located at $|z| = 34.5 \text{ cm}$ and 46.5 cm . The Pixel detector end-caps consist of a total of 672 pixel modules. The pixel detector allows measurements with a spatial resolution of $10 \text{ } \mu\text{m}$ in r - ϕ and $20 \text{ } \mu\text{m}$ in z .

The Tracker Inner Barrel (TIB) and Tracker Inner Disk (TID) are situated in the intermediate region, instrumenting the barrel and end-caps, respectively. The TIB consists of 4 layers, each $320 \text{ } \mu\text{m}$ thick and with a pitch between 80 and $120 \text{ } \mu\text{m}$. The first two layers consist of stereo modules, two strip modules installed with slightly different orientation to allow measurements both in r - ϕ and r - z . In the TIB, position measurements can be done with a point resolution of 23-24 μm in r - ϕ and $230 \text{ } \mu\text{m}$ in z . The TID consists of 3 small disks where the strips are arranged in rings pointing towards the beam-pipe, necessitating a changing pitch. The first two disks are again stereo modules.

The outer region contains the Tracker Outer Barrel (TOB) and Tracker End-Cap (TEC). The TOB has 6 layers with a higher thickness of $500 \text{ } \mu\text{m}$ and pitch between $120 \text{ } \mu\text{m}$ and $180 \text{ } \mu\text{m}$. Again, the first two layers consist of stereo modules with a stereo-angle of 100 mrad. The position resolution is 35-52 μm in r - ϕ and $530 \text{ } \mu\text{m}$ in z . The TEC consists of 9 disks, where the first 5 rings are equipped with stereo modules. The arrangement is the same as for the TID.

The tracker consists of a total of 15400 modules and operates at a temperature of $-20 \text{ } ^\circ\text{C}$.

2.3.3 Electromagnetic Calorimeter

The Electromagnetic Calorimeter (ECAL), shown in Fig. 2.7, was designed to measure electromagnetic showers with a good energy resolution to allow good di-photon and di-electron mass resolutions. Lead tungstate (PbWO_4) crystals were chosen as scintillators because of their short radiation lengths ($X_0 = 0.89 \text{ cm}$) and Molière radii ($R_M = 2.2 \text{ cm}$). The detector is homogeneous and hermetic with a total of 61200 crystals in the barrel and 7324 in each end-cap. The scintillators are relatively fast, with 80% of the light being emitted within 25 ns, and radiation hard up to 10 Mrad. Their disadvantage is a small light yield necessitating photo detectors with an intrinsic gain. Silicon Avalanche Photo Diodes (APD) were chosen for the barrel and Vacuum Photo Triodes (VPT) in




FIGURE 2.7 A perspective illustration of the Electromagnetic Calorimeter, showing the supermodules and modules composing the barrel, the Dee holding the crystals in the end-caps and the pre-shower devices placed just in front of the end-caps [25].

the end-caps. Both APDs and crystals are sensitive to changes in temperature and variations are kept within 0.1 °C.

The ECAL Barrel (EB) is constructed of 36 so-called supermodules of half barrel length, corresponding to a total coverage of $|\eta| < 1.479$, at a distance to the vertex of $r = 129$ cm. Each supermodule consists of 4 modules. Within these, the crystals are quasi projective (with a tilt of 3° relative to the line to the nominal interaction point) each covering 0.0174 in $\Delta\phi$ and $\Delta\eta$ corresponding to a front face of 22×22 mm². Each crystal is 230 mm long which translates to 25.8 radiation lengths.

The ECAL End-caps (EE) are situated at $|z| = 314$ cm and have a coverage of $1.479 < |\eta| < 3.0$. The Dees, semi-circular aluminium plates support structures for 5×5 crystals, called supercrystals, in an x - y grid. The front face of each crystal is 28.6×28.6 mm² and they are 220 mm = $24.7X_0$ long. In front of the EE a pre-shower device is placed which consists of two planes of silicon strip detectors with a pitch of 1.9 mm and two lead absorber disks. The first lead absorber is $2X_0$ thick and placed in front of the first silicon detector and the second ($3X_0$) between the two silicon planes.

2.3.4 Hadron Calorimeter

The Hadron Calorimeter (HCAL), depicted in Fig. 2.8, is constructed to be as hermetic as possible to allow a measurement of E_T^{miss} . At the same time the goal is to minimise the non-gaussian tails in energy resolution of measurements of hadron showers. Since the HCAL is enclosed by the magnet coil, the available volume for the detector is limited. The chosen design, therefore, maximises the absorber material in terms of radiation lengths while exploiting an active material that is relatively small: Brass (as absorber material) is interleaved with plastic scintillators. A tile/fibre technology was chosen where the plastic scintillator tiles are read out by wavelength shifting fibres connected to Hybrid Photo Diodes (HPD). Furthermore, the system is complemented with scintillators situated between magnet coil and muon system increasing the effective size

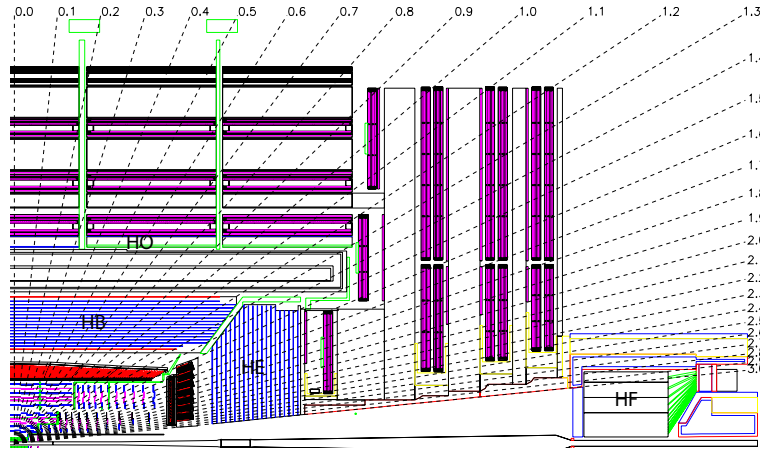


FIGURE 2.8 A r - z cross section of one quadrant of the Hadron Calorimeter (HCAL), highlighted are the Barrel HCAL (HB) and End-cap HCAL (HE) in blue closest to the interaction point, Outer HCAL (HO) in green at the top and the Forward HCAL (HF) on the right hand side [25].

of the HCAL to $10X_0$. Additionally, the Forward rapidity Hadron calorimeter (HF) at $|z| = 11.2$ m increases the coverage to $|\eta| < 5.0$. In the HF steel is used as an absorber and particles emit Cerenkov light in quartz fibres which are read out by photo multipliers.

The Barrel Hadron calorimeter (HB), covering $|\eta| < 1.4$, is segmented into 72×32 towers in $\phi \times \eta$, corresponding to a size of $\Delta\phi \times \Delta\eta = 0.087 \times 0.087$. Each tower consists of 15 5 mm thick brass plates interleaved with scintillator plates of 3.7 mm thickness, except for the first plate having a thickness of 9 mm.

The Outer Hadron calorimeter (HO) consists of the aforementioned scintillators placed between magnet coil and muon system. They follow the geometry of the barrel muon detector, the Drift Tube (DT). They are grouped in 60° sectors in ϕ and five rings in η . The central ring is instrumented with 2 scintillator layers covering both sides of an iron absorber, while the others consist of a single layer of scintillator.

The End-cap Hadron calorimeter (HE) covers $1.3 < |\eta| < 3.0$. It is also segmented into towers, 13 in η and 72 in ϕ . The ϕ -segmentation is $\Delta\phi = 0.087$ (or 5°) for the first five towers (at low η) and $\Delta\phi = 0.174$ (or 10°) for the last 8. The size in η varies from 0.087 for the first five towers to 0.35 for the last tower at highest η .

Finally, the HF covers the region of $3.0 < |\eta| < 5.0$. As mentioned before, the HF consists of steel and quartz fibres embedded therein. The fibres are arranged in a 5 mm-grid in parallel to the beam line and are inserted into grooves in the steel absorber. The HF is segmented into 13 towers in η with a size of $\Delta\eta = 0.175$, except for the tower at lowest η , where the size is 0.1. In ϕ the towers have a size of 10° with the exception of the highest in η that has double size.

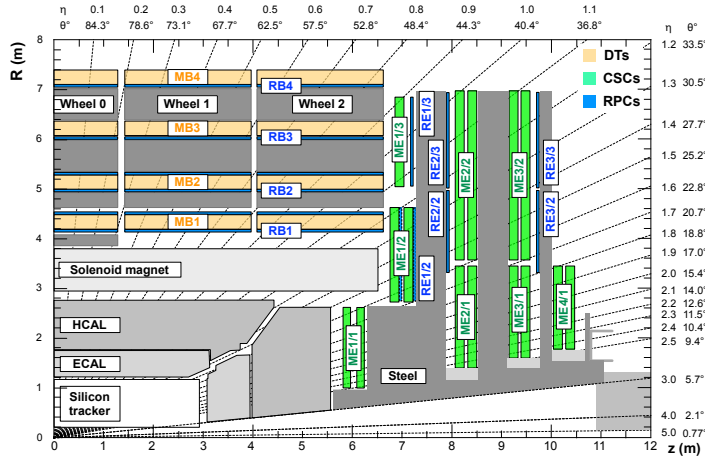


FIGURE 2.9 An r - z cross section of one quarter of the muon system with the z -axis pointing left to right and r increasing bottom to top, the interaction point is at the bottom left. The shown detectors correspond to the situation during run-1 of the LHC. Shown are the drift tube (DT) chambers in the four barrel stations (MB1-4) in orange, the cathode strip chambers (CSCs) in the four end-cap disks (ME1-4) in green and the complementing Resistive Plate Chambers (RPCs), installed both in barrel and end-cap, indicated in blue [27].

2.3.5 Muon Detector System

The aim for the design of the muon detector system includes good identification and momentum resolution for muons, a good di-muon mass resolution and reliable charge measurements up to muon momenta of 1 TeV. The p_T measurement in the muon system essentially relies on the bending at the magnet coil assuming the interaction point as the origin. Muons produced in the centre of CMS are measured both in the muon system and the tracker: For low p_T muons the overall momentum resolution is dominated by the tracker while for high $p_T > 200$ GeV the measurements can be improved by combining the two systems.

The choice of 3 different gaseous detectors was driven by the large surface areas to be covered and the different radiation requirements. In the barrel region, $|\eta| < 1.2$, where the neutron induced background, the total muon rate, and residual magnetic field are low, Drift Tube (DT) chambers were chosen. In the end-caps, up to $|\eta| = 2.4$, where the muon rate and neutron induced background as well as the residual magnetic field are high, Cathode Strip Chamber (CSC) are deployed. Both systems are complemented with Resistive Plate Chambers (RPCs) that were planned to instrument the region of $|\eta| < 2.1$. Only the chambers up to $|\eta| < 1.6$ were installed for run-1. While the RPCs have a fast response and good time resolutions, the position resolution of both CSCs and DTs is better. Fig. 2.9 shows a quarter of the muon system, it illustrates that the Muon Barrel (MB) is divided into 4 stations (labelled MB1 to MB4) arranged in cylinders between the return yoke. Each cylinder is split into 5 wheels. In the Muon End-caps (MEs), the system is split in 4 disks (labelled ME1 to ME4) perpendicular to the beam line. While ME1 consists of 3 rings (ME1/1, ME1/2 and ME1/3), the other 3 disks are divided in 2 rings each. The rings are numbered inside-out, so ME1/1 refers to the innermost ring of the first end-cap disk.




FIGURE 2.10 Layout of the Drift Tube (DT) chambers in the CMS barrel region, showing one of the 5 wheels. The numbers label the twelve sectors of the DT system. The bold lines indicate the RPC chambers. [26].

The DT stations are positioned at $r = 4.0, 4.9, 5.9,$ and 7.0 m and consist of 12 sectors, each covering 30° in ϕ . The chambers are staggered in ϕ to ensure that a high- p_T muon produced at a station boundary passes through at least three of the four stations. Fig. 2.10 shows a view of the DTs along the z -axis of CMS, it illustrates the staggered installation and shows the common labelling of the 12 sectors. In MB1-MB3 eight planes of DTs measure in the r - ϕ plane, four on each side of four z -measuring planes. In contrast, MB4 does not contain a z -measuring plane. The maximum drift length is 2 cm, the point resolution is $200 \mu\text{m}$ and the precision in ϕ is $100 \mu\text{m}$ with 1 mrad direction resolution.

The CSCs are trapezoidal in shape and designed with 6 gas gaps. In each gap a plane of radial cathode strips is complemented with a plane of anode wires which is perpendicular to the strips. The chambers overlap in ϕ to avoid gaps in the acceptance, except for ME1/3, the third outermost ring in the first disk. Each ring consists of 36 chambers except for the inner ring of disks 2-4 (ME2/1, ME3/1 and ME4/1) in which 18 chambers were installed. While ME4 was proposed to consist of 2 rings as ME2 and ME3, only the inner ring was constructed for run-1. The ionisation and resulting electron avalanche signal are used for the position measurement. The resulting spatial resolution is $200 \mu\text{m}$ ($100 \mu\text{m}$ in ME1/1) and the angular resolution in ϕ is 10 mrad.

In the barrel RPCs are installed on either side of the DT chambers for MB1 and MB2 while in MB3 and MB4 only one RPC is attached to the side closer to the interaction point. For the forward RPCs in the end-caps, only stations covering the first 3 disks and the outermost rings of the CSC (ME1/2, ME1/3, ME2/2 and ME3/2) were installed for run-1. RPCs are gaseous parallel-plate detectors that are designed for triggering and therefore are optimised for a very good time resolution of around 1 ns. The spatial resolution on the other hand is worse than that of CSCs and DTs. To achieve high rate capabilities for the CMS trigger, the RPC system is operated in avalanche mode and has the capability to cope with a rate of up to $1 \text{ kHz}/\text{cm}^2$.

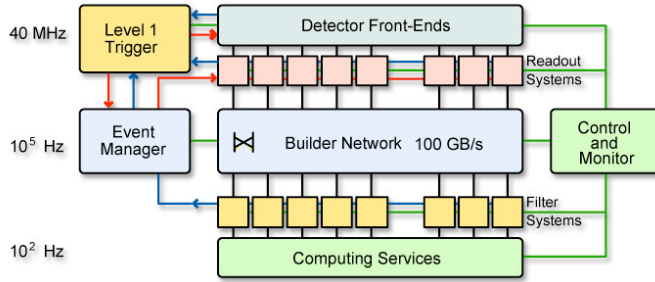


FIGURE 2.11 The CMS Data Acquisition system consisting of level-1 trigger, detector front-ends, the builder network, event manager, monitoring, and the computing services running the high level trigger. On the left hand side the average event rate at each selection stage is shown [24].

2.3.6 The CMS Trigger and Data Acquisition System

The CMS Data Acquisition system (DAQ), shown in Fig. 2.11, consists of the detector front-end electronics, the readout network and an on-line event filter system that runs the High Level Trigger (HLT) software. Along with the HLT, the Level 1-Trigger (L1-trigger) electronics form the CMS trigger.

The L1-trigger is located in another cavern, separated by 8 m of concrete to shield it from radiation. The L1-trigger is allocated a total of $3.2 \mu\text{s}$ to take the decision whether an event should be kept or discarded. This time is limited by the buffer size of the front-end electronics of the detector. It includes the time it takes to transfer data from the detector front-end to the trigger, the execution of the decision logic and the transmission of the decision back to the front-end electronics. Buffers in the front-end electronics hold the event data during this decision time. Taking the time into account that it takes the signals to arrive at the trigger and the transmission back, the decision has to be taken in $1 \mu\text{s}$. The L1-trigger is therefore implemented in custom hardware, including Field Programmable Gate-Arrays (FPGAs), Application Specific Integrated Circuits (ASICs), Programmable Logic Devices (PLDs) and discrete logic components, such as random access memories that contain Look-Up Tables (LUTs). The L1-trigger uses detector information from the calorimeters and all three muon systems, however, only at reduced granularity and resolution. The decision is based on objects representing muons, electrons/photons, E_T^{miss} , jets, or energy sums, that are propagated through the system. The design of the DAQ system and detector front-end electronics allow a maximum L1-trigger rate of 100 kHz, based on the average transfer time of the full event to the read-out system. Taking into account the design bunch crossing rate of 40 MHz and this upper limit on the L1-trigger rate, the system has to ensure a suppression factor of roughly 400, while ensuring a high efficiency close to 100% for interesting physics events.

If an L1-trigger is received, the front-end or front-end drivers may execute additional logic, such as further signal processing, zero-suppression and other data-compression, and places the data in memory for access by the DAQ system. The event builder collects all event fragments from about 700 front-end drivers and forwards them to one of the processors in the computer farm that runs the HLT software. The HLT is designed to discard events as early as possible to allow an efficient reduction of the rate from 100 kHz

to about 100 Hz. Therefore, events are reconstructed as needed: First only calorimeter and muon information is used, followed by tracker information and finally the full event information. Based on the final HLT decision events are then written to disk.

2.3.7 CMS Detector Upgrades during Long Shut-down 1

Hadron Calorimeter Upgrade

Most components of the hadron calorimeter, the HB, the HE and the HO are read out by HPDs, as mentioned in section 2.3.4. During operation in run-1 it was observed that these photo-detectors were not performing as well as expected. The signal-to-noise ratio turned out to be not good enough to allow muon identification in the HO and also lepton isolation both in HB and HE suffered. Since the initial design, Silicon Photo Multipliers (SiPMs) became available which consist of pixel arrays of Avalanche Photodiodes which are operated in Geiger mode. The new technology has a gain that is a factor of 50-500 better than observed in the HPDs. The signal-to-noise ratio can be improved by a factor of eight. Installing this new technology necessitates also changes in the read-out electronics [28]. The SiPMs were installed during LS-1 and the new read-out electronics based on the Micro Telecommunication Computing Architecture (μ TCA) [29] standard was deployed for HF.

Completion of the Muon Detector Installation

Both CSC and RPC detectors, as described in section 2.3.5, were not fully installed before the start of LHC operations. The time during LS-1 was used to complete the installation. Specifically, the additional CSC ring ME4/2 and the corresponding fourth RPC station were added. The extension of the RPC coverage to the full design value of $|\eta| < 2.1$ is foreseen for the next long shut-down of the LHC [28, 30]. Comparing figs. 2.9 and 2.12 shows which chambers were installed during LS-1.

2.4 CMS Software

The CMS software (CMSSW) provides the basic framework that is used for many applications and spans different domains: It is used during data-taking in the High Level Trigger and data-quality monitoring, calibrations and alignment of the detector is performed and then events are reconstructed. Additionally, the framework allows the necessary interfaces to produce simulated events, including both event generation itself and the subsequent simulation of the CMS detector.

In CMSSW modules, written in C++, are defined and may be chained to process events. These modules are steered through Python-based configuration files. Each event is passed through the chain of modules. There is no communication between modules, allowing to test them separately and to chain them in any order. The only way to pass information from one module to the next is to add data products to the event content.

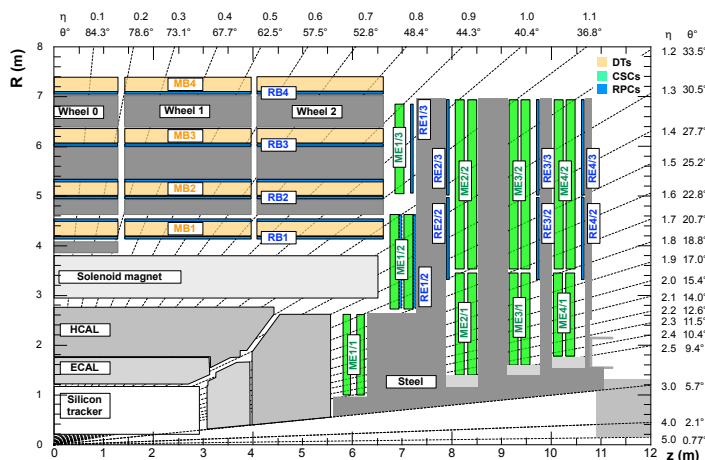


FIGURE 2.12 An r - z cross section of one quarter of the muon system with the z -axis pointing left to right and r increasing bottom to top, the interaction point is at the bottom left. The shown detectors correspond to the situation during run-2 of the LHC. Shown are the drift tube (DT) chambers in the four barrel stations (MB1-4) in orange, the cathode strip chambers (CSCs) in the four end-cap disks (ME1-4) in green and the complementing Resistive Plate Chambers (RPCs), installed both in barrel and end-cap, indicated in blue [31].

Specific modules exist for filtering, analysing or extending the event content. For storing data, CMSSW uses the I/O system provided by the data analysis framework ROOT [32].

The CMS detector simulation is based on GEANT4 [33] and includes a simulation of the full detector structure and material [34]. The software allows a combination of measured data in the simulation. For example measured noise can be introduced to improve the compatibility of the simulation with data events.

3 Overview of the L1-Trigger and its Upgrade

This chapter gives a more detailed description of the L1-trigger and its upgrade. The motivation of the upgrade and the trigger itself will be followed by a short overview of the hardware components that are used in the upgrade. Then, the individual parts of the trigger and their upgrades will be discussed in three sections: The calorimeter, the muon trigger, and the global trigger. The information is largely taken from [35–38] updated with information from private communication and working meetings between the participants [39].

3.1 Motivation

The design of the detector front-end electronics and DAQ system limit the maximum rate of accepted events by the L1-trigger. The design specifications of the LHC with a instantaneous luminosity of $10^{34} \text{ cm}^{-2}\text{s}^{-1}$ result in roughly a total interaction rate of 10⁶ kHz, cf. fig. 2.1. This takes into account that several interactions might occur within one single bunch crossing (pile-up). Reducing this number to the maximum accept rate of 100 kHz while keeping the data that contain interesting physics events at a high efficiency is the task of the L1-trigger.

In section 2.1.2 the cross section dependence on the collider energy was discussed. Taking this into account as well as the instantaneous luminosity expected for run-2, as discussed in section 2.2.4, the number of events will increase by a factor of six compared to run-1 [36]. These figures go beyond the design specifications of the original L1-trigger system. Coping with this new environment without the loss of interesting physics events, the algorithms and, therefore, electronics of the existing L1-trigger have to be replaced.

The L1-trigger upgrade is deployed in 2 stages: Already for the LHC start-up in 2015, the stage-1 upgrade was installed. It includes parts of the fully upgraded calorimeter trigger. During 2015 the installation of the remaining components of the upgrade followed and commissioning has begun during the summer until the end of 2015. In the following, the trigger used during run-1 will be referred to as the legacy trigger.

3.2 Common Upgrade Electronics

The first level of the CMS trigger is hardware based and different technologies are in use. The logical components of the trigger operate in a pipelined mode that is synchronised with the LHC clock. A special property of the CMS L1-trigger, compared to for example the ATLAS L1-trigger, is that information is propagated up to the final trigger component in the form of Trigger Primitives representing physics objects, such as muons, electrons, jets or energy sums. These Trigger Primitives are encoded in bit words¹ where the properties of the object such as η , ϕ , p_T are digitised and occupy different bit-fields.

In the legacy system, the data are sent between the different modules mostly through galvanic parallel links. The upgraded trigger on the other hand utilises serial optical links. These links have the advantage to allow higher link speeds and higher cabling density because of their relatively small size. The disadvantage is that the data has to be first de-serialised for the algorithms that still operate on Trigger Primitives. Then after processing they have to be serialised again to be sent to the next stage. This takes time that would otherwise be available for the processing itself. Most upgraded components in the trigger chain operate the links with the same protocol that allows transmission at 10 Gb/s which is realised through sending 32 bit words with a clock speed of 240 MHz, i.e. 6×32 bit words per bunch crossing. Additionally, control words including a cyclic redundancy check (CRC) sum that ensure the health of the link are sent during the orbit gap, when no collisions occur. The components that receive data from on-detector electronics operate the input links with a 3 Gb/s protocol that is compatible with the legacy electronics.

Core components of the trigger chain are the processing boards that, for example, determine the p_T of muons or take the final trigger decision. Both in the legacy and upgraded trigger, FPGAs are utilised that can be programmed flexibly with so-called Gateware. The Gateware, typically written in a Hardware Description Language (HDL), configures the logic blocks and the interconnects to perform a given logical function. For the upgraded trigger, three different processing boards were designed, the Master Processor 7 (MP7), the Muon Track Finder 7 (MTF7), and the Calorimeter Trigger Processor 7 (CTP7) [40, 41]. In all these cards, latest generation Xilinx Virtex-7 FPGAs are used. In addition to the main processing card, the MTF7 has a mezzanine card that holds a large Random Access Memory (RAM) which can be accessed from the FPGA and holds large LUTs needed in some algorithms. Both MTF7 and CTP7 have an additional smaller FPGA that is used for infrastructure like loading new Gateware images on the processing FPGA.

The trigger relies on the same crate and form-factor specification, μ TCA, as the HF upgrade. The technology is ideal for a triggering system as it is optimised for high availability through redundant power supplies, cooling, and crate-management. The back plane of the μ TCA crate is utilised for the readout of the trigger. One of the redundant crate-management components, Micro-TCA Control Hub (MCH), is replaced with a custom hardware component that provides an optical link for connection to the

¹The digitisation needed for this encoding leads to a loss of information. The precision of a given quantity has an influence on the resolution of the encoded properties.

TABLE 3.1 Sizes and positions of trigger towers in η and ϕ . The towers are symmetric for negative η . For trigger towers 21-28 the corresponding HCAL towers are artificially split in two in ϕ to regain the same granularity.

i_η	1-20	21	22	23	24	25	26	27	28
η_{max}	$0.087 \cdot i_\eta$	1.83	1.93	2.043	2.172	2.322	2.5	2.65	3.0
$\Delta\eta$	0.087	0.09	0.1	0.113	0.129	0.15	0.178	0.15	0.35
$\Delta\phi$	0.087	$\frac{0.174}{2}$	$\frac{0.174}{2}$	$\frac{0.174}{2}$	$\frac{0.174}{2}$	$\frac{0.174}{2}$	$\frac{0.174}{2}$	$\frac{0.174}{2}$	$\frac{0.174}{2}$

DAQ with a bandwidth of 5 Gb/s or 10 Gb/s. This component, the Advanced Mezzanine Card-13 (AMC-13), is designed and produced specifically for these first CMS upgrades. It also distributes the clock, synchronisation and control signals it receives from the Trigger Control and Distribution System (TCDS) to all boards within the crate. Apart from timing, the AMC-13 in concert with the TCDS enforces the trigger throttling rules which ensure that the maximum read-out rate is not surpassed.

3.3 The Calorimeter Trigger

For the CMS calorimeter trigger the HCAL and ECAL is segmented in a total of 2448, 1584 and 144 trigger towers in the barrel, end-cap and forward calorimeters, respectively. The trigger towers follow the geometry of the physical HCAL towers in the central region of the detector, up to $|\eta| < 1.83$. Here, the segmentation is constant with $\Delta\eta \times \Delta\phi = 0.087 \times 0.087$. In order to retain the same size and number of trigger towers in ϕ , the HCAL tower energy for $|\eta| > 1.83$ is split evenly between two trigger towers, as indicated in table 3.1. The η segmentation also starts to increase and sizes vary between $\Delta\eta = 0.09$ and 0.35 for the end-cap. Fig. 3.1 shows a quadrant of the CMS detector where the sizes of the trigger towers are indicated with dashed lines. The towers are associated with indices enumerating their positions in ϕ and η which are used in later stages of the trigger as the representation of angular information. In the forward calorimeters a fine segmentation is not needed, as this information is not used in electron / photon triggers and the tower size increases further to $\Delta\eta \times \Delta\phi = 0.5 \times 0.5$.

3.3.1 Legacy Calorimeter Trigger

The data flow of the legacy calorimeter trigger is displayed in fig. 3.2(a). The first module of the chain, the Regional Calorimeter Trigger (RCT) receives information from the trigger primitive generator that digitises the ECAL, HCAL, and HF energy information and assigns these energies to a given bunch-crossing. The energy information used throughout the trigger is the sum of ECAL and HCAL transverse energy segmented in towers. These are transmitted to the RCT in an 8-bit non-linear scale. Specifically, that means that energy is transmitted in an 8-bit word per tower where the 2^8 possible values are each associated with floating point transverse energies.

The RCT jet and energy sum algorithms operate on RCT-regions, sums of 4-by-4 trigger towers. The RCT is split into 18 processor crates that receive and handle data from the end-cap and barrel as well as forward calorimeters. Each crate holds 7 receiver,

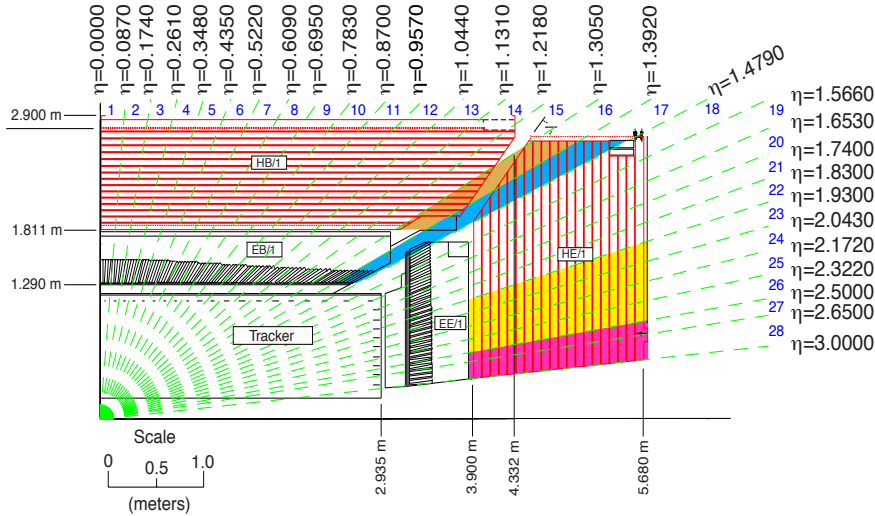


FIGURE 3.1 An r - z projection of one quadrant of CMS, showing the tracker, barrel ECAL (EB), end-cap ECAL (EE), barrel HCAL (HB), and the end-cap HCAL (HE). The green dashed lines and blue numbers indicate the trigger towers [35].

7 electron identifier and 1 jet / summary card. The RCT algorithms running in the receiver and electron identifier cards calculate the energy sums for the RCT-regions and identify electron candidates. The jet / summary card performs jet-finding and τ reconstruction algorithms. It also calculates larger granularity energy sums using the pre-summed regional energies. The HF energy is handled separately and only the energy look-up is performed within the RCT.

The Global Calorimeter Trigger (GCT) receives the 4 highest isolated and non-isolated electrons from each regional crate and the 9×4 highest central jets, forward jets and τ -candidates each from the clustering processor. The candidates are sorted according to their energies and location. Furthermore, the energies are summed to calculate the E_T^{miss} and total energy sums. The four highest candidates for each collection are then sent to the Global Trigger (GT).

3.3.2 Upgraded Calorimeter Trigger

The data-flow of the fully upgraded calorimeter trigger is illustrated in fig. 3.2(b). The ASICs used in the RCT and the Virtex-II FPGAs of the GCT are replaced by 36 CTP7 cards forming the layer-1 and 10 MP7 cards forming the layer-2. To facilitate the transmission via optical fibres, new hardware components, the optical Serial Link Boards (oSLBs), replace the legacy SLBs of the ECAL. While the upgrade of the HF replaced the back-end electronics with μ TCA optical boards, for end-cap and barrel of the HCAL the legacy optical connections can be re-used.

For the upgrade, the operating scheme of the calorimeter trigger is revised: Instead of the traditional processing, the trigger operates in a so-called time multiplexed mode. The data are not, as previously, processed locally and then combined to global objects, instead all data for one bunch-crossing are accumulated in one single FPGA. To accomplish this, the data are pre-processed by the 36 layer-1 processors: each receives

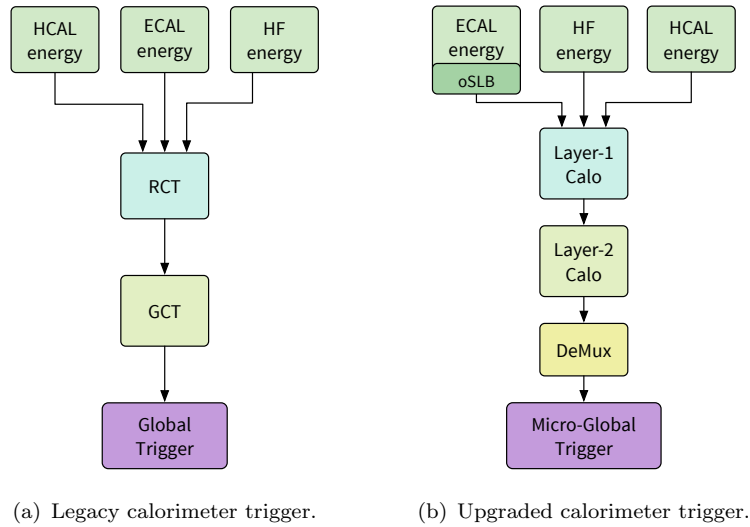
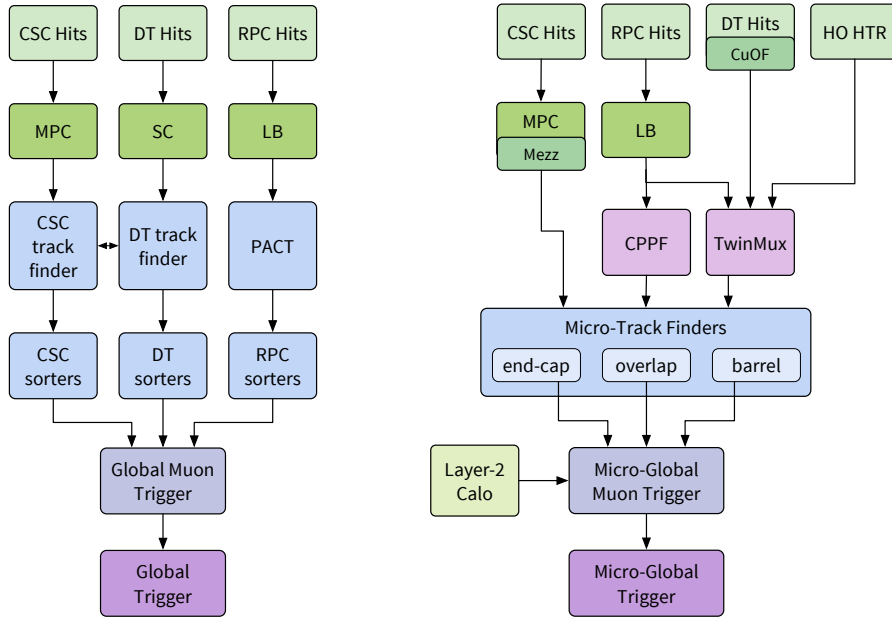


FIGURE 3.2 Illustration of the calorimeter trigger data flow: In the legacy trigger (a) energies from the calorimeters (HCAL, ECAL and HF) are pre-processed by the Regional Calorimeter Trigger (RCT), sorted in the Global Calorimeter Trigger (GCT) and sent to the Global Trigger. In the upgraded trigger (b) new processing boards operate in two layers, with pre-processing being done in the first layer and trigger primitive reconstruction in the second, the additional de-multiplexer unit (DeMux) prepares the data for the upgraded Micro-Global Trigger.

tower information for all towers in positive or negative η and a 4 tower wide strip in ϕ . Then all data from these 36 cards, associated with bunch crossing N , are transmitted to one single layer-2 processor. Given the bandwidth, this transmission has to occur over several consecutive bunch crossings. Therefore, the data for bunch-crossing $N + 1$ are transmitted to the next layer-2 processor. This architecture allows the accumulation of full-granularity data in one single FPGA for jet-finding, electron identification and isolation as well as τ reconstruction and the calculation of energy sums. The reconstruction algorithms are then adapted to be performed on the consecutively arriving strips of calorimeter data and are sorted on the fly during the reconstruction to avoid a big sort-block at the end of the algorithm.

The calorimeter trigger objects are finally sent to the μ GT. However, the time multiplexed operating mode requires another processing step done in the de-multiplexer card (DeMux), where the objects associated with a given bunch crossing are collected and then forwarded to the μ GT.

For the data taking period in 2015, already some of the layer-2 electronics have been installed. Data from the RCT are received and the additional logical resources are exploited to implement more sophisticated τ and electron isolation algorithms. Also, the pile-up mitigation algorithms were improved.



(a) Legacy muon trigger.

(b) Upgraded muon trigger.

FIGURE 3.3 Legacy and upgraded muon trigger chain: In the legacy, Cathode Strip Chamber (CSC), Drift Tube (DT) and Resistive Plate Chamber (RPC) information is available, while in the upgrade also Outer Hadron calorimeter (HO) hit information can be processed. The hit information of the Cathode Strip Chambers is pre-processed in the Muon Port Card (MPC). In the legacy the muon candidates are reconstructed in components dedicated to each detector, the Cathode Strip Chamber Track-Finder (CSCTF), Drift Tube Track-Finder (DTTF) (and Sector Collector (SC)) and PAttern Comparator Trigger (PACT) (and Link Boards (LB)). The candidates are first sorted locally (per sub-detector) and then merged and globally sorted in the Global Muon Trigger (GMT). In the upgrade the candidates are reconstructed track-finders that are defined by η acceptances of the barrel (Barrel Muon Track-Finder (BMTF)), end-cap (End-cap Muon Track-Finder (EMTF)) and the overlap of the two (Overlap Muon Track-Finder (OMTF)). Information of Outer Hadron calorimeter, Drift Tube and barrel Resistive Plate Chamber is combined in the TwinMux and provided to the BMTF. The Concentrator Pre-Processor Fan-out receives end-cap Resistive Plate Chamber data and sends it both to EMTF and OMTF.

3.4 The Muon Trigger

All three muon systems DT, CSC, and RPC participate in the muon trigger of CMS. While the spatial precision of CSC and DT allow good momentum and angular measurements, the timing precision of RPCs can be utilised to assign the muon candidates to the correct bunch-crossing.

3.4.1 Legacy Muon Trigger

The redundancy of the muon detectors is exploited to ensure the availability of a muon trigger even if one of the detectors should be malfunctioning. The data-flow is represented by the flow chart in fig. 3.3(a). Dedicated systems are deployed that use the information from the DT, CSC, and RPC: The Drift Tube Track-Finder (DTTF), the Cathode

Strip Chamber Track-Finder (CSCTF) and the PAttern Comparator Trigger (PACT), respectively. The local information from the 3 systems is collected in the GMT, where the best 4 candidates are chosen and sent to the GT.

DT Track-Finder – The DT hits are processed locally in on-detector electronics to provide track segments that indicate the position in ϕ , the bending angle ϕ_b and a quality code (the number of layers used) to the track finder. The DT track segments are received from the Sector Collector (SC) which synchronises the data for all 4 chambers of a sector, where one sector is a 30° wedge of one wheel with the exception of the central wheel which is split into negative and positive η . The data are then sent to the DTTF processors, one processor per sector, where, depending on the track segment origin, an extrapolation is performed either inward (from MB4 to MB3) or outward (from MB1 and MB2 to any possible outer chamber). The extrapolation is based on programmable windows in ϕ and the bending angle determined in the on-detector electronics. Each processor receives the segments for its associated sector and the two neighbouring sectors in ϕ , as well as three in η . Additionally, the DTTF processors for wheels ± 2 receive track segments from the CSCTF to ensure full coverage. The p_T is then assigned based on the $\Delta\phi$ of the innermost track segments that were combined. A quality code is calculated based on the number of track segments that could be matched and the angular position is computed at MB2. The position in η is assigned in a dedicated η -track finder logic that is based on pre-defined patterns. For each wedge, a synonym for the 30° ϕ sectors, a dedicated sorter module receives 2 candidates from each of the 6 associated processors. The best two are selected based on the assigned quality code and p_T and sent to the DT sorter, resulting in 24 candidates. Again the 4 candidates with highest p_T and quality code are selected and forwarded to the GMT.

CSC Track-Finder – The CSC information is pre-processed in on-detector electronics that reconstruct track segments for anode and cathode hits, separately. The track segments provide a radial and azimuthal position, an approximate bending angle, as well as timing information and an η coordinate. The Muon Port Card (MPC) receives up to two track segments from nine chambers. The best three are transmitted via optical fibres to the CSCTF. An extrapolation is performed on pairs of track segments checking whether the pair is compatible with a muon originating from the beam spot. Any combination is tested with the exception of a direct extrapolation from ME1 to ME4. The extrapolation is performed separately for ϕ and η . In a final step p_T , η , ϕ , charge sign and quality are assigned. The p_T is similarly assigned like in the DTTF, using the $\Delta\phi$ between the track segments. Unlike the DTTF, the CSCTF is segmented into 60° sectors in ϕ , resulting in a total of 12 processors for positive and negative end-caps. The CSC sorter receives up to two muons from each processor and sends the best 4 candidates to the GMT.

Pattern Comparator – In contrast to the DTTF and CSCTF, the PACT uses the hit information from the RPC directly without pre-processing. The hits are synchronised in the Link Boards (LB) and forwarded to the PACT. In a first step, hits are

clustered and the median position of each cluster is used in subsequent steps. A pattern recognition algorithm based on pre-programmed hit patterns that are compatible with muons originating from the vertex is used to identify muon candidates. The hit information from the RPC is compared with these and charge, p_T , η and ϕ are assigned accordingly. In addition, a requirement of a minimum number of participating hits in a given pattern is imposed. The number of hits is also encoded in a quality code. Each of the 84 processing boards can send up to 4 candidates to the RPC sorter. The four best candidates for barrel and end-cap are selected independently and sent to the GMT.

Global Muon Trigger – The final component of the muon trigger, the GMT, merges, sorts, and cancels the received candidates to forward the best four to the GT. The GMT merges the RPC muons from the barrel with the DDTF candidates and the RPC muons from the end-cap with the CSCTF candidates. The p_T assignment can be configured based on the η position and the quality or p_T of the muons. Generally, in 2012 the minimum p_T was assigned to candidates to reduce the rate as much as possible. Furthermore, the GMT cancels muons that are found both in the barrel and in the end-cap. This cancellation is based on the reported ϕ and η coordinates. Finally, the best four muons are selected based on a rank that is a function of p_T and quality code. The GMT is configured to demote unmatched muons of certain qualities that are within a given η region. An isolation algorithm was also implemented. Per RCT-region a MIP bit and isolation bit could be received: The first would indicate whether the deposited energy would be in agreement with a minimal ionising particle and the latter whether the energy in the region was above a configurable threshold. However, the connections between calorimeter trigger and GMT were never commissioned and the algorithm, therefore, not used as the performance at this granularity of calorimeter information did not show promising results.

3.4.2 Upgraded Muon Trigger

The overall approach to the muon trigger was revised for the upgrade, cf. fig. 3.3(b). Instead of having the redundancy of the muon detectors mirrored in the trigger chain, the new approach combines the information as early as possible to exploit the advantages of all three systems to a maximum degree. Also, the upgrade of the HO in LS-1 allows to use this additional muon detector in the trigger. Instead of per-detector triggers, the upgrade approaches the muon identification by deploying dedicated trigger components for the regions that are naturally defined by the detector acceptances. One track-finder will be working in the barrel, $|\eta| \leq 0.83$, the Barrel Muon Track-Finder (BMTF) taking HO, DT and RPC information into account. Another will have specialised algorithms for the difficult transition region between barrel and end-cap, $0.83 < |\eta| < 1.24$, the Overlap Muon Track-Finder (OMTF) that takes DT, RPC and CSC information and utilises a pattern recognition algorithm. In the end-caps, $|\eta| \geq 1.24$, the End-cap Muon Track-Finder (EMTF) will use both RPC and CSC information to identify and measure muons. Again, there will be one component that concentrates all the muon information and selects the best candidates, the Micro-Global Muon Trigger (μ GMT).

Furthermore, the higher available bandwidth allows to transmit the muon properties with a higher precision. For example, the available bandwidth for p_T , η , and ϕ has almost doubled. Overall, the muon object size of previously 32 bits has increased to 64 bits.

Barrel Muon Track-Finder – The BMTF uses hit information from the HO, the DT and the barrel RPC to reconstruct muon tracks: The hit information is combined in the Twin Multiplexer (TwinMux) where similar track segments as previously in the on-detector electronics are generated and sent to the next stages of the trigger at 10 Gb/s. While HO and RPC measurements cannot provide a bending angle ϕ_b they are used as extrapolation targets for the track-finding.

The track-finding algorithm itself will be similar to the one used in the legacy DTTF: An extrapolation window in ϕ is used to associate track-segments from the different detectors and chambers. The p_T is assigned based on the azimuthal angle difference of the associated hits and potentially the ϕ_b of the innermost DT chamber.

The BMTF consists of 12 MP7 processing boards, each processing the information from one wedge. In the upgraded scheme, the wedge sorter is absorbed in the BMTF logic and up to 36 muon candidates are sent to the μ GMT via optical fibre.

Overlap Muon Track-Finder – The OMTF uses hit information from all three muon systems: DTs, RPCs, and CSCs. The CSC information is sent via mezzanine cards installed on the legacy MPCs, which allow transmission at 3.2 Gb/s. The DT and RPC barrel information is received through the TwinMux. The RPC end-cap information is concentrated and distributed with the Concentrator Pre-Processor Fan-out (CPPF). This new processor may also be used to do clustering of the RPC hit information as well as concentrating the information and forward it with 10 Gb/s. Alternatively, the end-cap RPC can also be received directly from the legacy link boards at the lower transmission speed of 3.2 Gb/s.

The track finding algorithm uses all information available in a pattern recognition algorithm. Pre defined patterns, representing muons with different p_T are generated and compared with the detector information. The positions of track segments and hits that can be matched to a pattern are then compared: The probability that a $\Delta\phi$ between hits or track segments is compatible with the pattern is used to define a matching quality. The best matches are then selected based on this quality. This means that muon candidates are selected that have the most hits and where the hits are most compatible with the pattern. As the patterns are defined for different p_T bins, this already yields a p_T assignment.

After the compatible patterns are found, those that have hits or track segments in common are cancelled. In case of these duplicates, only the one with the higher quality is kept.

The OMTF consists of 12 processing boards, 6 for each side of the detector, where the information is split into 60° ϕ sectors. Each processor sends up to 3 muon candidates to the μ GMT.

End-cap Muon Track-Finder – The EMTF receives hit information from the CSCs and RPCs. Exactly like in the OMTF, the RPC information is received from the CPPF or directly from the RPC link boards and the CSC information from the mezzanine cards on the legacy MPCs. In contrast to the legacy track finding, all 18 CSC track stubs are analysed in parallel. A pattern recognition algorithm is used to reconstruct track candidates.

The pattern recognition works in zones, an artificial ϕ - η -segmentation of the end-cap that is introduced as an ordering mechanism for the reconstructed tracks: Each 60° sector is divided into 4 such zones that cover approximately equal angles in θ and ϕ [39]. The hits within one zone are compared with pre-defined patterns to create track candidates. Each pattern can be assigned a quality code based on the number of hits and the type of pattern used. In case a given hit gives rise to several track candidates, these ghosts are cancelled based on the quality: For each station-2 ϕ hit position, the highest quality pattern containing that position is chosen. Additionally, the pattern is compared to the neighbouring patterns.

The found tracks are associated with all hits in the window of the pattern and all $\Delta\phi$ and $\Delta\eta$ between the hits are calculated. These are used for p_T assignment which is based on a multivariate method: Using simulated events containing muon candidates, a Boosted Decision Tree with the same input variables as are available on the processing board is trained to derive the known p_T of these muons. The result of this regression is saved in a LUT which is loaded into the large memory of the MTF7-mezzanine and used in the processing card to assign a p_T .

Finally, all tracks in a sector are compared to each other and in case of shared hits, the one with highest quality is chosen. From the remaining candidates the 3 highest quality tracks are sent to the μ GMT. The μ GMT, therefore, again receives 36 candidates in total.

Micro-Global Muon Trigger – The μ GMT receives a total of 108 muons from the track-finders. In contrast to the legacy system, the muon sub-detector data are already combined at this stage. The μ GMT sorts the received candidates, removes duplicates and the 8 best muon candidates are forwarded to the next trigger stage. Additionally, the μ GMT receives calorimeter information from the calorimeter trigger to calculate an absolute and relative muon isolation. Details of the μ GMT algorithm will be discussed later.

3.5 The Global Trigger

The final component of the trigger chain is the GT or in the case of the upgraded system the μ GT. This central processor receives jets, electrons (photons), τ -jets and energy-sums from the calorimeter trigger and muons from the muon trigger. A distinction between photons and electrons in the L1-trigger is not possible since the tracker information is not available. Based on these objects the Global Trigger forms the final decision whether an event is to be kept.

3.5.1 Legacy Global Trigger

The legacy GT included the processor that calculated the trigger decision based on the trigger primitives, the Trigger Control System (TCS) and dedicated receiver cards for calorimeter trigger data and so-called technical triggers. The TCS forms the final level-1 accept signal based on the trigger decision and the trigger rules, which, for example, prohibit a trigger signal in two consecutive bunch-crossings. Also, if the detector is not ready for a new trigger decision, they are suppressed through the Trigger Throttling System (TTS). The final accept signal is transmitted to all detector sub-systems via the Trigger, Timing and Control system (TTC). The sub-systems then move the data associated with the given bunch-crossing to de-randomising memories where they are retrieved by the DAQ. The TCS limits the accept rate to avoid overflows in the readout chain.

The GT processor receives 4 candidates each of the following classes: muons, isolated electrons (photons), non-isolated electrons (photons), central jets, forward jets and isolated τ -jets. As mentioned previously, these objects are already sorted based on their quality and transverse energy or momentum. Additionally, the GT receives direction and amount of missing transverse energy, the total transverse energy and eight numbers of jets passing programmable E_T thresholds. As the GT receives the trigger objects, and therefore also position of the trigger objects information, the decision may be based on the event topology such as angular distances between various objects.

Up to 128 different trigger algorithms can be calculated in the GT which may require combinations of objects, a given event topology or simply objects with energy or momentum higher than a given threshold. The GT trigger decision is based on a logical or of these 128 algorithms and the technical triggers (at least one of the requirements has to be fulfilled).

3.5.2 Upgraded Global Trigger

The μ GT receives similar objects as the legacy system but with increased bandwidth, higher precision and, thus, larger objects can be sent. Additionally, the number of objects per class has been increased to eight with changes in the object definitions as described previously.

The μ GT is implemented in at least one MP7. The number may increase if more algorithms should be needed to ensure efficient data taking. Per processor, 512 algorithms are foreseen. The logic is similar to the old but several changes to the infrastructure were implemented to make changing trigger algorithms easier. Additionally, the topological conditions that are available were extended, for example an invariant mass can now be calculated. Increasing the number of objects by a factor of two also enables more combinatorial requirements in the μ GT, for example low p_T muons may be used to tag jets as potentially originating from B mesons.

The TCS is split off the Global Trigger in the upgrade and forms the independent Trigger Control and Distribution System (TCDS). Therefore, in contrast to the legacy, the μ GT does not calculate the final accept decision. Instead it only calculates the logical

or of the different algorithms. The TCDS ensures the detector front-end readiness and suppresses triggers if needed. It also enforces the aforementioned trigger rules. Finally, the limitation of the accept rate to avoid overflows is ensured. The final accept, is then only issued by the TCDS if these conditions are fulfilled and the μ GT formed a positive trigger decision. Apart from these triggers based on the μ GT decision, the TCDS triggers calibration sequences of the calorimeters and is able to issue random triggers. The TCDS is already in use since the start-up of the LHC after LS-1 in the beginning of 2015.

4 Trigger Algorithm Performance Studies

This chapter gives an overview of studies evaluating the performance of the upgraded muon trigger, mainly the μ GMT. First, Monte Carlo simulation samples that were used to derive the results shown are introduced and the method to derive efficiency and rate estimations from these is explained. This is followed by an analysis of the track-finder performance which is the input for any μ GMT algorithm. Then, methods are explored that could be applied to enhance this standalone performance in the μ GMT algorithm: Demoting muon candidates in specific η -regions that have a low quality and the introduction of an isolation algorithm. For the latter, several approaches are presented and their performance is evaluated by deriving a rate and efficiency estimation at various working points.

4.1 Monte Carlo Simulated Event Samples

Two simulated samples are used for the results shown in the following. To derive the level-1 trigger efficiency muons from a $Z/\gamma^* \rightarrow \mu\mu$ simulation sample with $\sqrt{s} = 13$ TeV are considered. The events are generated with MADGRAPH [42] and passed to PYTHIA [43] for showering and hadronisation. A total of 30000 events were produced centrally and are used. For estimation of the rate at which muon candidates are triggered with a given algorithm, samples of simulated minimum bias events are used. Here, 13 million deep inelastic collision events at $\sqrt{s} = 13$ TeV without specific requirements on the hard scattering process are generated with PYTHIA. Both minimum bias and $Z/\gamma^* \rightarrow \mu\mu$ simulation samples are generated with an average pile-up of 20. Additionally, the influence of increased pile-up on potential isolation algorithms was checked with a smaller minimum bias sample of 3 million simulated events and 30000 $Z/\gamma^* \rightarrow \mu\mu$ events produced with an average pile-up of 40.

4.2 Event Selection

4.2.1 Efficiency Event Sample

A generator level selection is performed where exactly two muons are required. Furthermore, these two muons are required to have a minimum distance of $\Delta R > 0.2$ to




FIGURE 4.1 Distributions of p_T (left) and η (right) of generated (red) and level-1 muons (black) for the minimum bias (dashed) and $Z \rightarrow \mu\mu$ (solid).

avoid double counting due to the level-1 muon-to-generator matching algorithm that is applied.

At level-1, different selection parameters are applied that are detailed in the relevant text sections below. Usually, at least a p_T threshold is used, i.e. muon candidates are required to have a minimum p_T^{\min} . In that case, only generated muons that have a minimum $p_T > p_T^{\min} + 8$ GeV are considered to determine the plateau efficiency independently of the turn-on effect due to the finite p_T resolution of the level-1 measurement. The latter is considered separately in the corresponding efficiency distribution as a function of the generated muon p_T , where no such requirement is applied.

4.2.2 Rate Sample

For determination of the total rate of the level-1 muon trigger system, the generated event is required to have an L1 muon of any p_T in either the legacy or upgrade system. This minimal requirement allows to compare the two systems. No requirement is placed on generated muons.

Fig. 4.1 shows the p_T and η distributions of generated muons and the level-1 muons prior to any selection for both minimum bias and $Z/\gamma^* \rightarrow \mu\mu$ sample. One can see that the p_T assignment of the level-1 trigger tends to overestimate the muon p_T . For the barrel region, the η assignment is coarse, while for end-cap and overlap it follows the generated distribution. The rate sample is dominated by very low p_T muons.

4.3 Rate and efficiency estimation

Rate – The figure of interest for rate estimation is the number of events per second. This can be calculated by assuming an instantaneous luminosity \mathcal{L} . The total inelastic

cross section at 13 TeV is assumed to be $\sigma_{inel} = 100$ mb [8]. Given these numbers, the number of events per second is calculated with their product. With the total number of generated minimum bias events N , a sample weight can be calculated as

$$w = \frac{\mathcal{L}}{L} = \frac{\mathcal{L}\sigma_{inel}}{N}. \quad (4.1)$$

An instantaneous luminosity of $\mathcal{L} = 1.5 \times 10^{34}$ Hz/cm² is assumed for all distributions shown in this section, cf. table 2.3.

Efficiency – For any generator level muon from the Z boson decay, a match is searched for among the level-1 trigger muon candidates. To determine whether a muon was found by the trigger, the following criterion is applied

$$\Delta R = \sqrt{\Delta\phi^2 + \Delta\eta^2} = \sqrt{(\phi^\mu - \phi_{L1}^\mu)^2 + (\eta^\mu - \eta_{L1}^\mu)^2} < 0.5, \quad (4.2)$$

where η^μ is the pseudorapidity of the generator muon and η_{L1}^μ is the assigned value in the level-1 system, correspondingly for ϕ . The efficiency then can be calculated from the number of muons that have a match N_{match} among the level-1 candidates and all generator muons N_{tot} :

$$\epsilon = \frac{N_{match}}{N_{tot}}. \quad (4.3)$$

Furthermore, it can be instructional to determine the loss of efficiency given an additional selection criterion introduced in the level-1 trigger system. In that case the relative efficiency is calculated as the ratio of the efficiency with and without the criterion applied:

$$\epsilon_{rel} = \frac{\epsilon_{new}}{\epsilon} = \frac{N_{match}^{new}/N_{tot}}{N_{match}/N_{tot}} = \frac{N_{match}^{new}}{N_{match}}. \quad (4.4)$$

4.4 Track-Finder performance

Before evaluating the performance of any new μ GMT algorithm, the quality of its inputs needs to be evaluated. In contrast to the legacy GMT, the parameter assignments of the track-finders is used as is and the quality of this assignment has even more importance for the final performance of the muon trigger.

4.4.1 Status of the track-finding emulation

At the time of writing, the emulation of the upgraded track-finder algorithms is still under development. One important issue is that with the exception of the OMTF, RPC hits are not yet used in the track reconstruction. In its final form, the BMTF will use RPC hit clusters in the same way it is using DT track-segments. The hit information from both muon detectors is combined in the TwinMux. However, no emulation of the TwinMUX is available and, therefore, it is not clear how exactly the combination will be done in the final version of the algorithm. The initial idea is, as outlined in section 3.4.2, to use the hit clusters only as targets in the track-finding algorithm because no bending

angle can be assigned. The EMTF follows a similar approach, where a station-1 CSC hit is required and in the pattern assignment RPC hits are only used for other stations, cf. 3.4.2.

All emulators used for obtaining the results shown in the following agree with the firmware available at the time of writing to either a relatively high percentage of events (90% of the events have the same results for the BMTF emulator / firmware) or at the bit-by-bit level (for OMTF and EMTF). The latter means that, given the same input data, the exact same output data is produced to the available integer precision by hardware and emulator.

4.4.2 Categorisation of Track-Finder Muons

Two different categorisation approaches are used in the following sections. The muons can be distinguished by the track-finder, which means that in comparisons with generated muons these are restricted to the nominal acceptances of each track-finder. This also applies for efficiency calculation where the track-finder muon candidates are analysed. The following acceptances are used:

- ▷ Barrel region of the BMTF: $|\eta| \leq 0.83$,
- ▷ Overlap for the OMTF: $0.83 < |\eta| < 1.24$,
- ▷ End-caps in the EMTF: $|\eta| \geq 1.24$.

The second categorisation used is done based on the quality assigned by the track-finders. While in the final version the quality will be compatible between the different systems, in the available version of the emulation two different scales are used. In case of barrel and overlap track finders, the legacy scale is used that assigns a value between 1 and 7 to each candidate. This value depends on the number of hits used for reconstruction, which hits were used and, in the case of the overlap, a number assigned to each pattern that encodes the compatibility with a prompt muon. The EMTF uses a new scale, a 4 bit value, where each bit represents whether a hit from a given station (one to four) was used in the reconstruction. The disadvantage of this scale is that muons with a higher quality code do not necessarily have a more reliable p_T assignment. E.g. muons that were reconstructed without the first station (corresponding to the least significant bit) have generally a less reliable p_T assignment.

4.4.3 Performance of the standalone Track-Finders

The standalone performance of the track-finders can be evaluated at the input stage of the μ GMT. Rate and efficiency are calculated in the same way as for the μ GMT, but the generator muons are restricted to the nominal acceptance of the track-finder. For the following plots, no quality selection is applied and simply all candidates from the track-finders are considered.

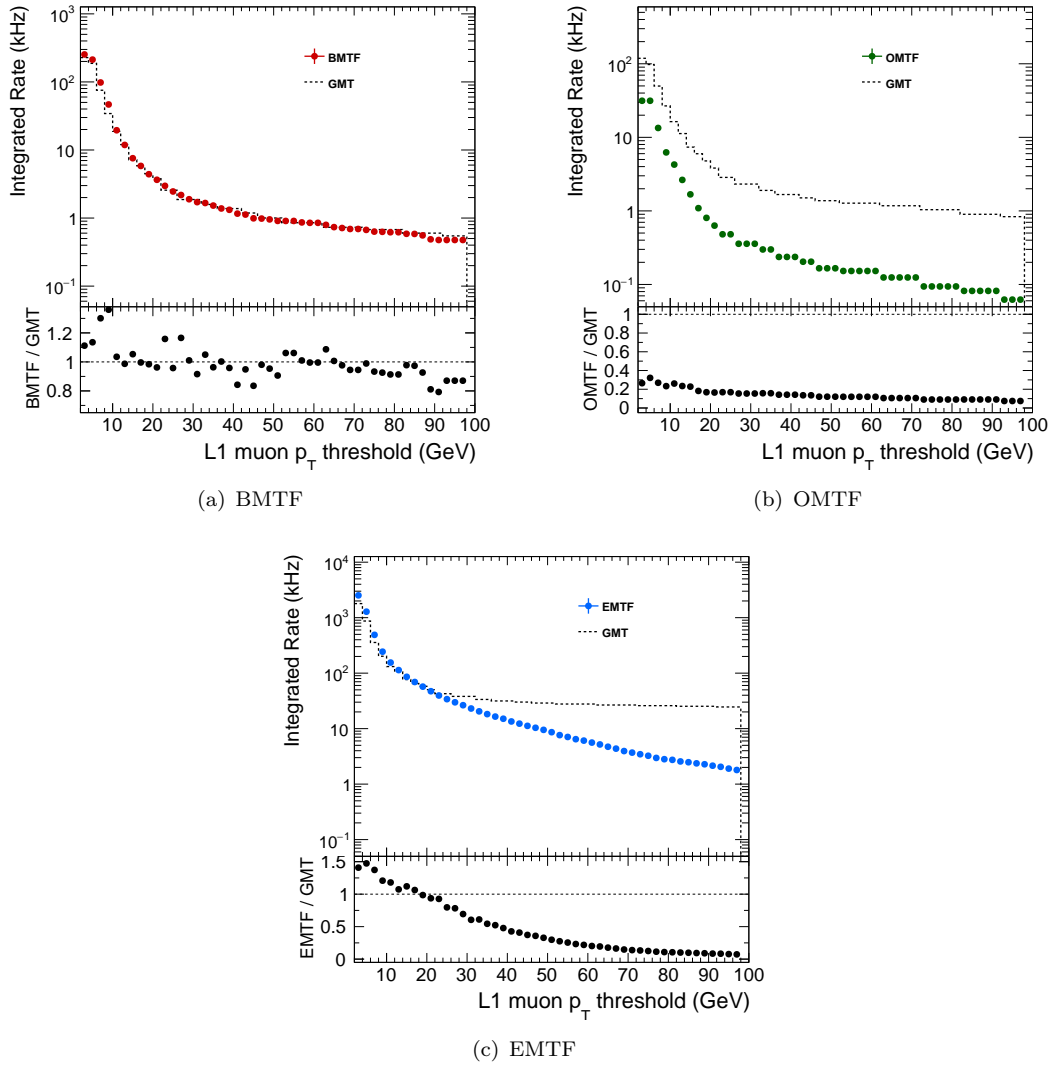


FIGURE 4.2 Rate as a function of p_T threshold. The dashed line represents the legacy GMT rate in the corresponding region. Dots show the estimation for the track-finders: Barrel in (a), overlap in (b) and end-cap in (c).

Rate – as a function of the p_T threshold is shown in fig. 4.2. In fig. 4.2(a) it is seen that the behaviour in the barrel is very similar as in the legacy system. This is to be expected as the legacy track-finding algorithm is used and RPC information is not yet incorporated. The rate is slightly higher because in the legacy GMT candidates are merged and the smaller p_T is assigned. Furthermore, duplicate candidates are cancelled.

Fig. 4.2(b) shows the evolution of rate in the overlap region, demonstrating a significant improvement with respect to the legacy system. The rate can be reduced by up to a factor of 10, depending on the threshold, and a minimum of a factor of five for very low thresholds. For reference, the typical threshold for level 1 muons is 16 GeV for both run 1 and the beginning of run 2, where the rate is reduced to roughly 18% of what was obtained with the legacy system.

The EMTF also has improved performance with respect to the legacy GMT. However, only for relatively high p_T thresholds greater than 21 GeV. Also here, the rate can be reduced by up to a factor of 10, but only for very high thresholds.




FIGURE 4.3 Estimated rate as a function of η with an applied p_T threshold of 16 GeV. The dashed line represents the legacy GMT rate. The stack shows the rate contributions from the track-finders. In red the barrel, green the overlap and blue the end-cap track-finder.

In fig. 4.3 the rate at an applied p_T threshold of 16 GeV is shown as a function of η . Again, the improvements in the overlap region are visible. In the barrel similar performance is seen across the whole region. In the end-caps, the difference between upgraded EMTF and legacy GMT becomes less significant at the edge of the RPC acceptance of $|\eta| < 1.6$. The algorithm of the legacy GMT merges candidates that are reported by two detector sub-systems and, therefore, combines the RPC information with the CSC information. However, the upgraded EMTF emulation does not yet incorporate the hit information from the RPCs. Therefore, this reduced performance of the upgraded algorithm is expected and will likely be improved once the algorithm is extended to make use of RPC information. Beyond this acceptance, however, the rate can be reduced with the new algorithm.

Efficiency – as a function of several muon parameters is shown in fig. 4.4. The track-finders are directly compared in these plots and are shown as differently coloured dots: BMTF in red, OMTF in green, and EMTF in blue. For all of the plots in this figure a threshold of 16 GeV was applied at the track-finder and legacy GMT level.

The turn-on in the barrel region as shown in 4.4(a) has a similar shape as in the legacy system, but the plateau efficiency is roughly 5% lower. The overlap region exhibits a slightly slower turn-on than the legacy system and reaches a plateau of 90%-95%. For the end-cap, the plateau efficiency is similar to the legacy and around 95%. The turn-on shape, however, is again less steep with regards to the legacy system.

Looking at the efficiency as a function of η , cf. fig. 4.4(b), one can directly see the difference in the plateau efficiency is present in the whole barrel region. Again, the introduction of RPC information in the muon reconstruction will likely help to increase the efficiency in this region. Especially at the wheel boundaries around $|\eta| \approx 0.2$, the efficiency deteriorates without RPC hit information. For the OMTF the drop in the plateau efficiency is due to relatively hard acceptance cuts that are applied at the track-finder emulator output. Here, the assigned η of the OMTF muon candidates is required

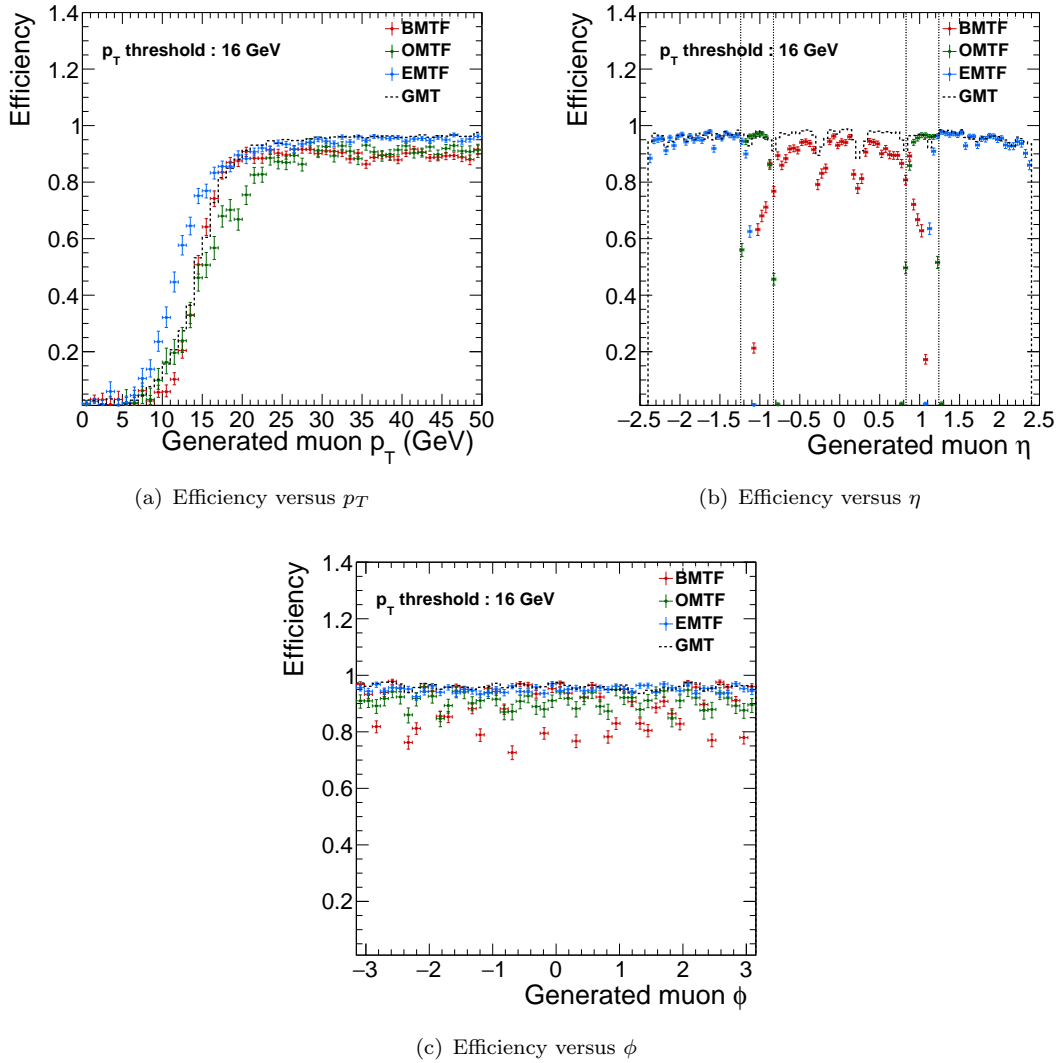


FIGURE 4.4 Efficiency of the track-finders as a function of p_T (a), η (b), and ϕ (c). All plots show the efficiency at a threshold of 16 GeV. The dashed line shows the efficiency of the legacy system and the coloured dots the upgraded track-finders efficiency. The BMTF in red, the OMTF in green, and the EMTF in blue.

to be within the nominal acceptance of the track-finder. This leads to a drop in efficiency already in the last bins within the nominal acceptance of the OMTF. These cuts can be configured to be looser. Otherwise, the overlap region shows similar or better efficiency with respect to the legacy system. In the end-cap region, the overall efficiency is similar to the one observed in the legacy system, with the exception of the highest η bins.

While for OMTF and EMTF the efficiency as a function of ϕ is flat within the uncertainties, cf. fig. 4.4(c), it deteriorates in the barrel at the wedge boundaries. This repeating pattern of decreased efficiency is very similar to the DT only reconstruction efficiency in the DTTF that is observed in the legacy trigger [37]. Therefore, it is expected that efficiency here can be recovered by including RPC hit information.

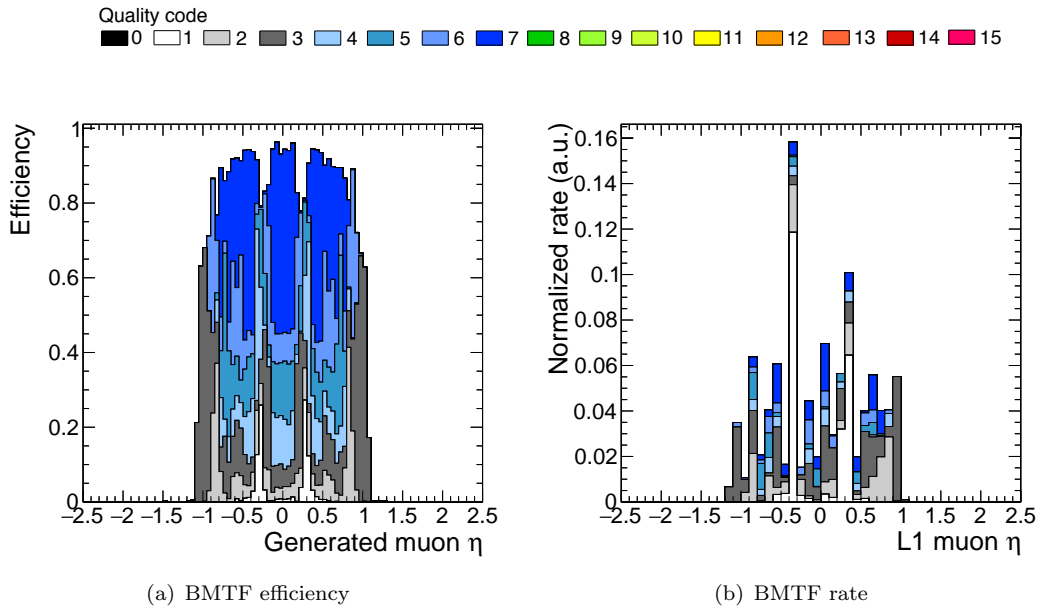


FIGURE 4.5 The efficiency (a) and rate (b) in the BMTF split into the qualities reported by the track-finder. The quality scale in the BMTF ranges currently from 1 to 7, while the μ GMT accepts values between 0 and 15.

4.4.4 Rate Suppression with Quality Selection

The possibilities to reduce the muon trigger rate at the μ GMT are limited. The μ GMT cannot improve on the p_T assignment of the track-finders, but a pre-selection based on the quality reported by the track-finders can be implemented. Categories of muon candidate qualities can be found that contribute only little to the efficiency of muon triggers while they produce large rates. In the hardware, muons from a given track-finder or within a certain η region could be assigned new (lower) quality codes at the μ GMT (or already at the outputs of the track-finders). E.g. given a quality code assigned by the track-finder and the η of the muon candidate, a new quality could be assigned that the μ GMT then uses to determine whether these candidates can be safely ignored or only used in certain trigger algorithms. For example muons of relatively low quality could still be of interest for b-tagging triggers or requirements for di-muon triggers could be less strict than those for single-muon triggers.

In fig. 4.5, efficiency and rate are shown in the barrel region. The muon candidates are split by quality and the resulting histograms are stacked. The colour coding is indicated at the top. The upgraded quality scale ranges from 0 to 15 (a 4 bit scale), but since the BMTF currently re-uses the legacy algorithm it transmits the same scale as the DTTF which ranges from 1 to 7. The quality depends on the stations used for reconstruction, where stations 1 and 2 are most important and otherwise more stations correspond to a higher quality. E.g. a track reconstructed from all stations has quality 7, a track reconstructed from station 1 and 2 and any other station is assigned a quality of 6 and a track from stations 3 and 4 has quality 1, see table 4.1.

Comparing the contribution per quality to the efficiency, fig. 4.5(a), and the rate,

TABLE 4.1 Indicated are the required station hits in a track candidate for a given quality code.

Quality Code	1	2	3	4	5	6	7
Hits required	3-4	2-3	1-2	2-3-4	1-3-4	1-2-3	1-2-3-4
in stations		2-4	1-3			1-2-4	
			1-4				

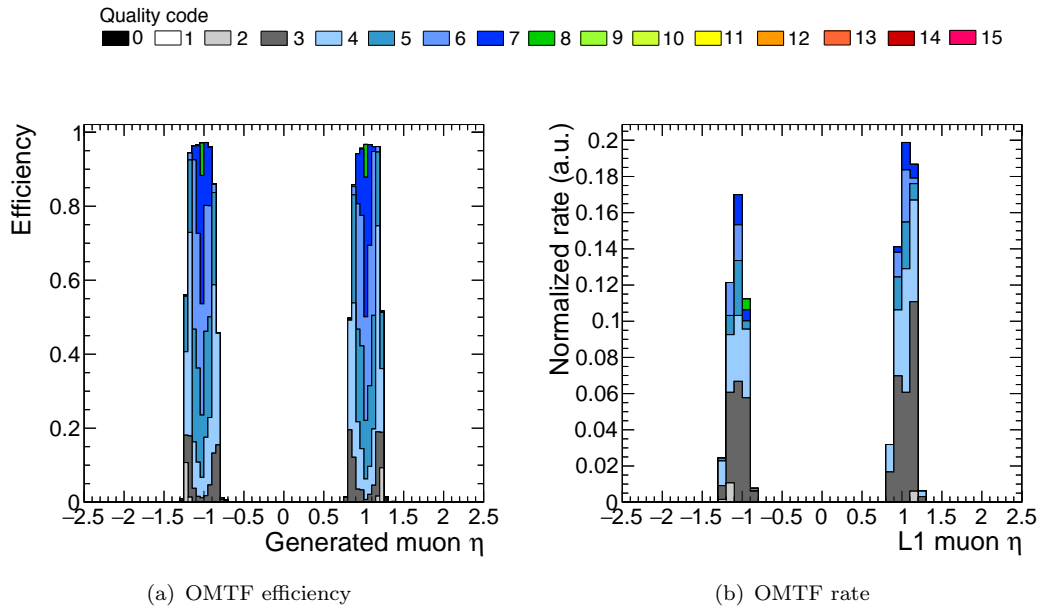


FIGURE 4.6 Efficiency (a) and rate (b) contributions of the different qualities reported for muon candidates in the OMTF. The quality code is the number of hit layers used for a given candidate.

fig. 4.5(b), one sees that there is no quality that only contributes to rate while contributing not at all or little to the efficiency. When comparing the two plots it should be noted that while for the efficiency the generator muon η is used as the x-axis and for the rate the level-1 η . This can not be avoided due to the definition of the efficiency where sometimes no level-1 muon is available and of the rate where generator muon information may not be available.

Fig. 4.6 shows similar distributions for the overlap region. Again, not the full range of quality codes is currently utilised in the track-finder emulation. In the current implementation the number of hit layers is used as quality.

All muon candidates of a given quality contribute significantly to the efficiency in the region where they show contribution to the rate. Therefore, no optimisation based on quality is possible for the OMTF.

In the end-cap region, the EMTF reports a bit-pattern that indicates the stations used in the track reconstruction; Each of the four bits represents the presence of a hit of one station of the CSC. E.g. a quality of 12 has a binary representation of 0b1100 and therefore indicates that stations 1 and 2 were used in the track reconstruction. As the algorithm requires at least two hits not all values between 0 and 15 are expected in the

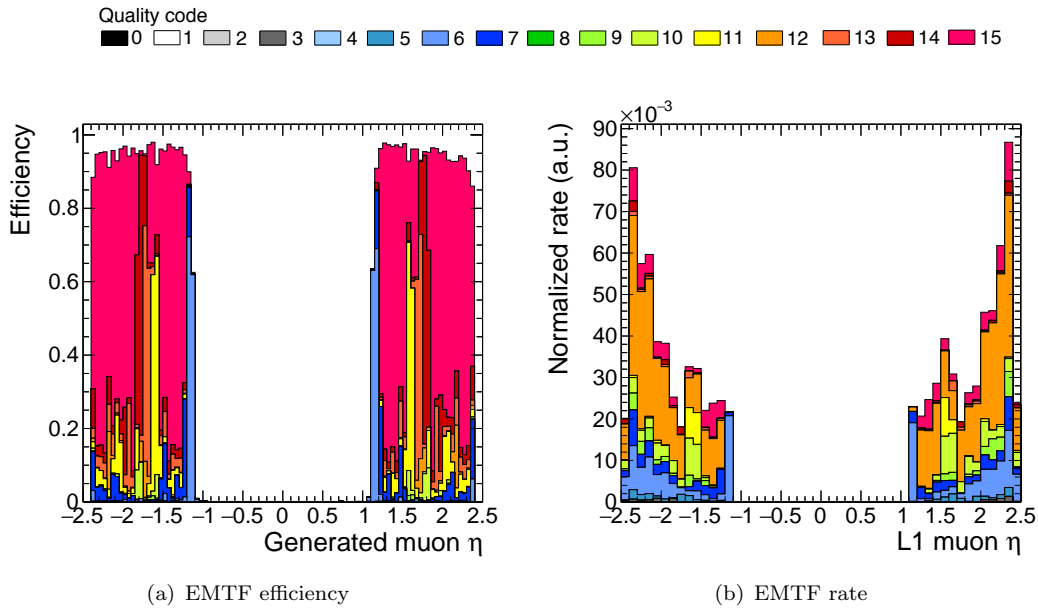


FIGURE 4.7 Efficiency (a) and rate (b) contributions split in quality of muon candidates in the EMTF. The quality code corresponds to the stations that were used in the track reconstruction.

μ GMT. Fig. 4.7 shows the efficiency and rate contribution of each quality code. As expected from the definition of the quality, a wider range of values is seen but still a few are never reported as the track-finder algorithm requires at least two hits for a valid track.

For the EMTF, several muon candidates of certain qualities only contribute little to the overall efficiency. Muon candidates with qualities below 10 with the exception of 7 and 6 (representing patterns using station 2, 3 and 4 and station 2 and 3, respectively) have little to no contribution to the efficiency, cf. 4.7(a). In the rate distribution, shown in fig. 4.7(b), however, they do contribute significantly. Also, muons with quality 6 could be masked if $|\eta| > 1.2$. Another quality that could be masked is code 12, corresponding to tracks with hits in station 1 and 2 only. Qualities 9 and 10 (station 1-4 and station 1-3 candidates) do also contribute more to rate than to efficiency.

Applying this quality selection results in efficiency and rate distributions as shown in figs. 4.8 and 4.9, respectively. Here, the efficiency and rate are shown for muon candidates as available in the μ GMT with no η requirements but a p_T threshold of 16 GeV applied. On the left hand side, the distributions without a quality selection are shown to be compared with the distributions on the right hand side where the selection is applied. Comparing fig. 4.8(a) with fig. 4.8(b), one sees that the selection results in almost identical efficiency with the exception of a few points around $|\eta| \approx 1.7$, where the efficiency is reduced to 80%. However, comparing fig. 4.9(a) with fig. 4.9(b), the resulting rate reduction is up to a factor of 7 in the high η region and roughly a factor of 2 across the whole EMTF acceptance.

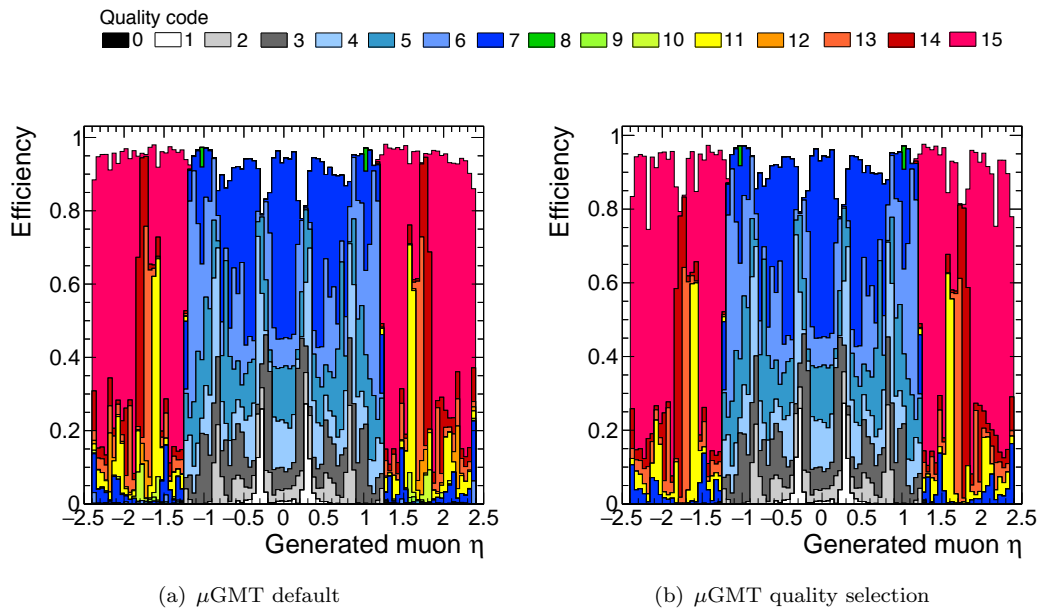


FIGURE 4.8 Efficiency of the μ GMT as a function of η . Without quality selection (a) and selection as described in the text (b) applied. In both cases a threshold of 16 GeV is applied.

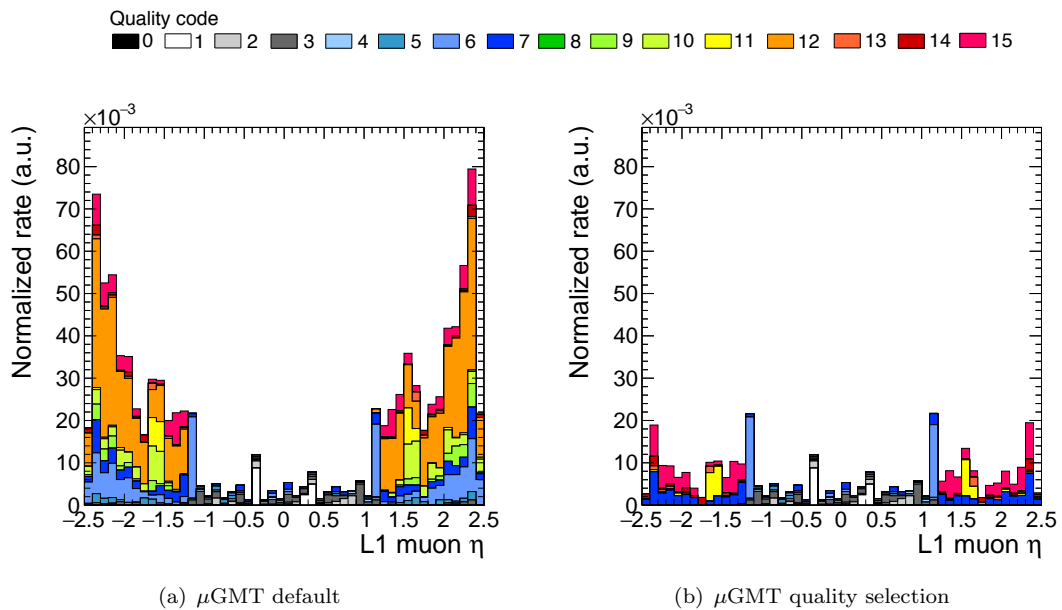


FIGURE 4.9 Rate of the μ GMT as a function of η . Without quality selection (a) and selection as described in the text (b) applied. In both cases a threshold of 16 GeV is applied.

4.4.5 Comparison with Legacy System

The legacy system used in run-1 presents a good benchmark for the upgrade. The aim is to have a similar efficiency but reduce the rate as much as possible. Similar to the previously described quality masking, the legacy GMT had algorithms in place to demote the quality of track-finder muons based on their quality and η . Therefore, a

comparison between the upgraded and legacy GMT is done with and without quality selection applied in both systems.

For the μ GMT the quality selection described in the previous section is used. The GMT uses the configuration as it is used in the beginning of run-2, where CSC TF candidates are masked if the quality code is 2 or if the quality code is 1 and $|\eta| > 1.3$.

The μ GMT emulation performs a cancel-out based on a window $\Delta R < 0.1$. If candidates are found within that distance of one another, the one with lower p_T is selected. This choice effectively reduces the rate at a given threshold while having little effect on the efficiency. This cancel-out is performed between muons from neighbouring track-finder sectors and at the track-finder boundaries (e.g. between BMTF and OMTF).

Efficiency – Figs. 4.10 to 4.12 show the trigger efficiency including the μ GMT algorithm emulation in relation to the trigger efficiency of the legacy GMT. The left hand distributions show the efficiency without any quality requirements applied while on the right hand side the above described quality criteria are applied both in μ GMT and GMT. The GMT efficiency is shown as a dashed line, while the μ GMT is represented by a stack split by the individual track-finder contributions. In the lower part of each plot the ratio of μ GMT efficiency over GMT efficiency is drawn.

Comparing the distributions in fig. 4.10(a) with fig. 4.4(b) the effect of the cancellation on the overall efficiency can be seen. While the effect of the cancellation within the track-finders themselves is negligible, a very small effect can be seen at the boundary of OMTF and EMTF. Here, the efficiency drops by a few percent, compared to the standalone efficiency. Overall, the efficiency is at a similar level as for the legacy system with the exception of the drops due to the quality selection and within the BMTF acceptance. This can also be seen in fig. 4.11, where the efficiency is shown as a function of ϕ . The overall efficiency of the upgraded muon trigger is at a similar level as that of the legacy muon trigger, with the exception of the periodic deterioration due to the wedge boundaries of the DT system, which will be recovered by introducing RPC information into the BMTF. Both with and without the quality selection applied, the overall efficiency is similar comparing upgraded with legacy muon trigger.

Looking at the efficiency as a function of p_T in fig. 4.12, one can see that the turn-on is slightly slower in the upgraded system, also visible in the ratio distributions. The overall plateau efficiency is slightly lower in the upgrade (mainly driven by the lower BMTF efficiency) but no trends are visible across the p_T range.

Rate – Fig. 4.13 shows the rate as a function of the p_T threshold. Again, in fig. 4.13(a) no quality selection is applied, while it is applied in fig. 4.13(b). The ratio between legacy GMT and upgrade μ GMT looks very similar for both plots, the overall rate is similar at a p_T threshold of 16 GeV. The dashed line shows that for the legacy system increasing the threshold beyond approximately 25 GeV does not reduce the rate any further. However, for the upgrade increasing the threshold remains to be a viable option up to very high thresholds.

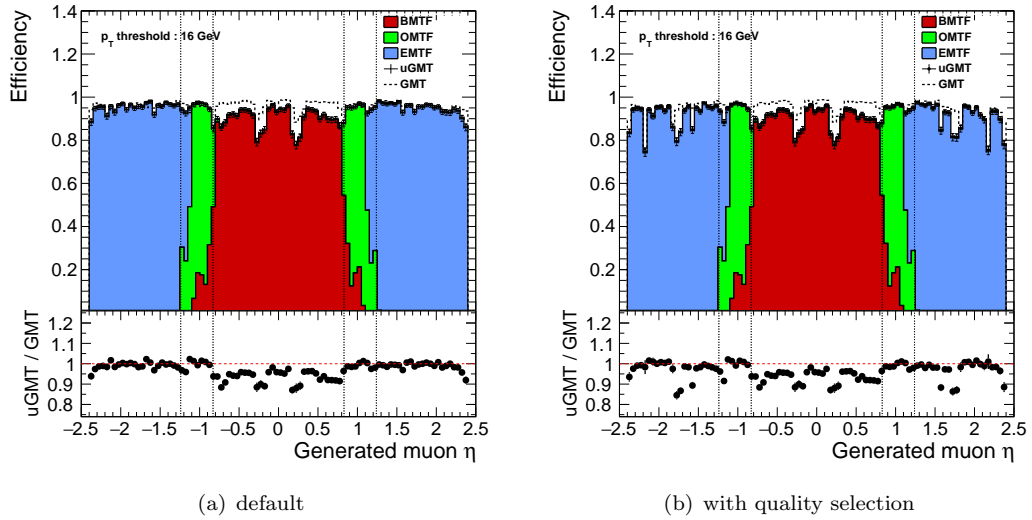


FIGURE 4.10 Efficiency as a function of η , comparing legacy GMT (dashed line) with μ GMT (stack). The μ GMT is stacked in contributions of the individual track-finders (BMTF red, OMTF green, and EMTF blue). In (a) comparison without quality selection and in (b) with quality selection as described in the text.

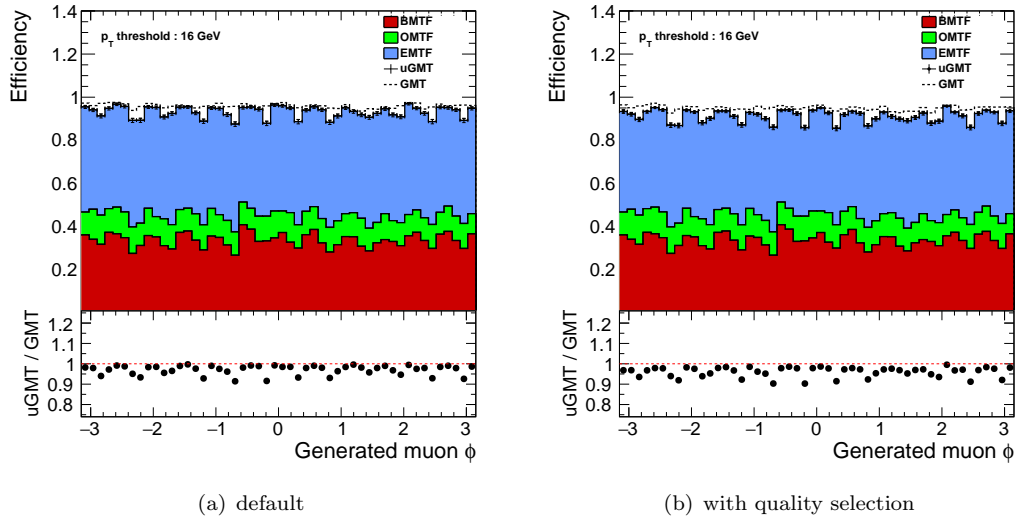


FIGURE 4.11 Efficiency as a function of ϕ , comparing legacy GMT (dashed line) with μ GMT (stack). The μ GMT is stacked in contributions of the individual track-finders (BMTF red, OMTF green, and EMTF blue). In (a) comparison without quality selection and in (b) with quality selection as described in the text.

Overall, both efficiency and rate are slightly lower for thresholds of 16 GeV in the upgrade μ GMT. While efficiency is approximately at 95% of the legacy and the rate is around 90%. However, the efficiency will be increased by the inclusion of RPC information. At the same time the additional hits from the RPCs can improve the p_T assignment and therefore reduce the rate. At higher thresholds, the rate reduction in the upgraded system can go up to a factor of 10 with respect to the legacy muon trigger.

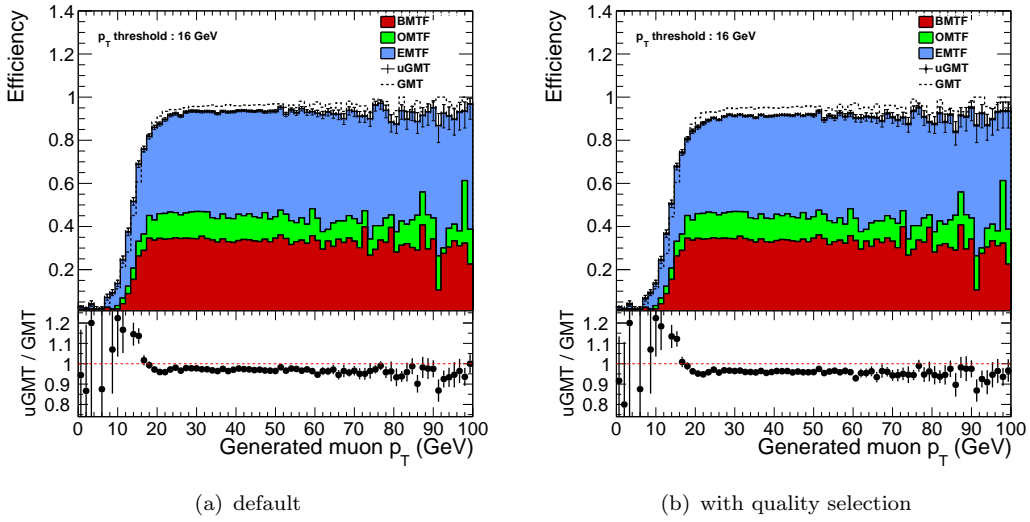


FIGURE 4.12 Turn-on distributions of the legacy GMT (dashed line) and the μGMT (stack). The μGMT is stacked in contributions of the individual track-finders (BMTF red, OMTF green, and EMTF blue). In (a) comparison without quality selection and in (b) with quality selection as described in the text.

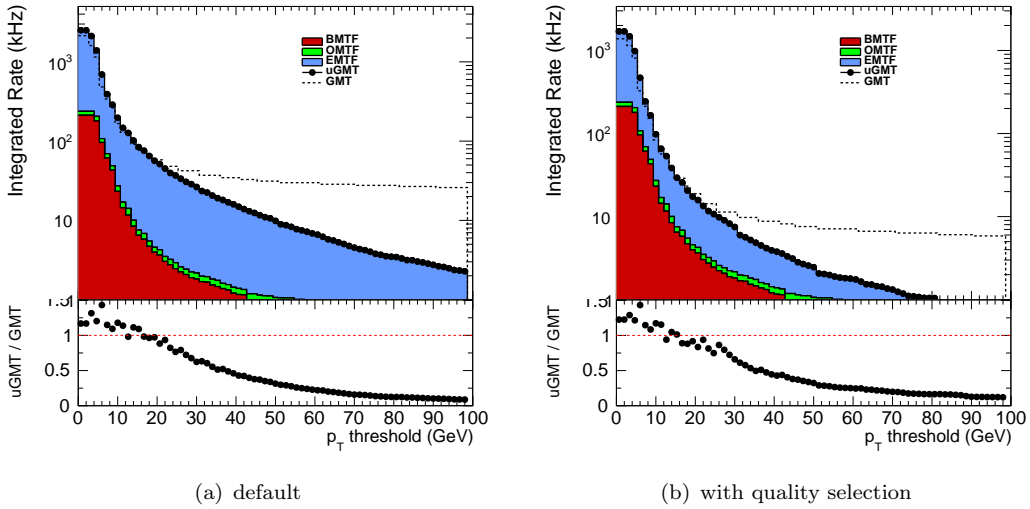


FIGURE 4.13 Rate as a function of p_T threshold for the legacy GMT (dashed line) and the μGMT (stack). The μGMT is stacked in contributions of the individual track-finders (BMTF red, OMTF green, and EMTF blue). In (a) comparison without quality selection and in (b) with quality selection as described in the text.

4.5 Isolation in the Upgraded GMT

The upgraded μGMT will receive energy information from the calorimeter trigger. In the following the available calorimeter information is described and possible usage in the μGMT algorithm discussed. The general strategy is to calculate the sum of energy in an area around a muon candidate and use this to distinguish between prompt isolated muons and those that are produced in jets.




FIGURE 4.14 Ratio of level-1 E_T index to E_T calculated in off-line reconstruction as a function of the tower η index.

4.5.1 Available Calorimeter Information

In the upgraded calorimeter trigger, information is available with tower granularity, as described in section 3.3. More precisely, the sum of measured E_T in ECAL and HCAL is calculated and used in the algorithms. The limiting factor for the granularity and precision of calorimeter energy information received in the μ GMT is the available bandwidth. Given the bandwidth two different strategies are possible:

- ▷ Utilise full granularity and receive pre-calculated cones,
- ▷ Receive at a reduced granularity (“tower sums”) and calculate isolation cone in the μ GMT.

While the former can make use of the full granularity available in the trigger for the cone calculation, the bandwidth limitation results in the fact that these cones can only be received for every other tower or very reduced precision. Receiving the towers directly would allow for more flexibility in the energy calculation but increases the logic resources needed in the μ GMT.

Another factor of limitation due to bandwidth is the scale in which energy sums can be received; Depending on the granularity chosen the energy information would have to be compressed. Within the calorimeter trigger, the tower energies are transmitted with a scale of 0.5 GeV bins.

For the following studies, full tower granularity and the calorimeter trigger energy precision is assumed and possibilities for an implementation in the level-1 trigger are discussed at the end.

4.5.2 Tower Calibration and Muon Extrapolation

The level-1 calorimeter system uses a discrete scale to represent the energy of each tower. To get the actually physical meaningful E_T , this index has to be re-calibrated. For this,




FIGURE 4.15 The result of the tower calibration. On the left, the reconstructed tower energy as a function of the level-1 E_T index can be seen. On the right the corresponding distribution of reconstructed tower energy for the E_T index of 9 is shown.

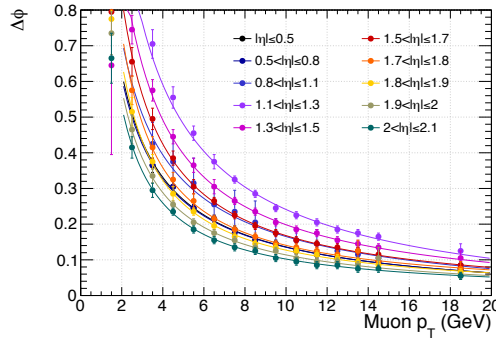


FIGURE 4.16 Extrapolation to the vertex. The differently coloured lines show the fitted functions used for the extrapolation.

the fully calibrated tower energies used in the CMS reconstruction can be compared with the level-1 scale. Fig. 4.14 shows the ratio of level-1 E_T index to the reconstructed E_T for the different tower η indices i_η . One can clearly distinguish between the distribution for towers with $|i_\eta| \leq 21$ (“central”) and those beyond (“forward”). While the central distribution peaks around $E_T \approx 0.5$ GeV, the towers in the forward region tend towards a coarser binning and have a larger spread. The change in calibration coincides with the artificial split of HCAL energy for the trigger towers beyond $|\eta| = 1.83$ (corresponding to $i_\eta > 21$).

In fig. 4.15 the method of re-calibration is illustrated. fig. 4.15(b) shows the distribution of the ratio of reconstructed E_T and the level-1 index for one level-1 index. Similarly, the reconstructed E_T in all level-1 E_T bins is used to determine the physical E_T value; The mean of the distribution is used as a central value and the RMS as the uncertainty. The resulting reconstructed E_T as a function of the level-1 index can be seen in fig. 4.15(a).

The level-1 muon trigger system uses MB2 in the barrel and ME2 in the end-cap as a reference frame for the measurement of the azimuthal angle ϕ . The calorimeter, on the other hand, uses the tower indices as a scale for η and ϕ position. The mapping




FIGURE 4.18 Average tower energy around a muon in case of prompt muons (a) and muons found by level-1 in minimum bias (b). On the x- and y-axis, the tower index relative to the extrapolated muon position is used.

Given the distributions in fig. 4.18, one can already define several key parameters of an isolation algorithm:

- ▷ To contain most of the energy seen in fig. 4.18(b), a cone size of ± 5 towers around the muon is chosen
- ▷ An inner cone should have the size of 3-by-3 towers to contain the typical energy foot-print of a prompt muon.

Using these parameters, the corresponding energy sums can then be calculated as

$$E_T^{\text{out}} = \sum_{|\Delta i_\eta|=2}^5 \sum_{|\Delta i_\phi|=2}^5 E_T^{\text{twr}}(i_\eta, i_\phi), \quad (4.6)$$

$$E_T^{\text{in}} = \sum_{|\Delta i_\eta|=0}^1 \sum_{|\Delta i_\phi|=0}^1 E_T^{\text{twr}}(i_\eta, i_\phi), \quad (4.7)$$

$$E_T^{\text{tot}} = E_T^{\text{out}} + E_T^{\text{in}} = \sum_{|\Delta i_\eta|=0}^5 \sum_{|\Delta i_\phi|=0}^5 E_T^{\text{twr}}(i_\eta, i_\phi), \quad (4.8)$$

where $E_T^{\text{twr}}(i_\eta, i_\phi)$ is the transverse energy of a tower at position i_η, i_ϕ , and the sum indices correspond to the differences $\Delta i_\eta = i_\eta^\mu - i_\eta$ and $\Delta i_\phi = i_\phi^\mu - i_\phi$.

4.5.3 Energy Distribution

Effect of extrapolation – Already in the previous section, it was shown that on the one hand the effect of bending due to the magnetic field is small for the prompt muons from the Z boson decay. The low p_T muons found in the minimum bias sample

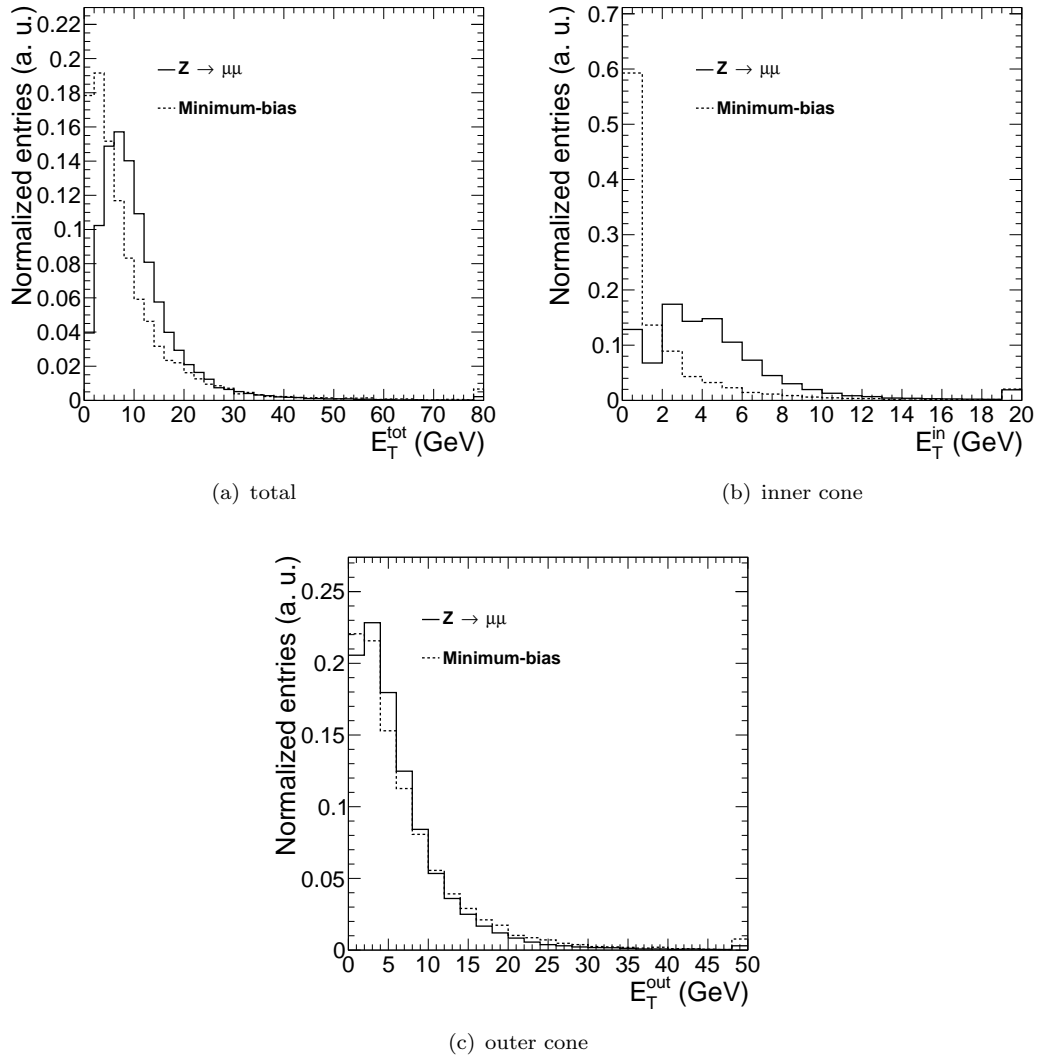


FIGURE 4.19 Sum of tower energies in cones around the muon, shown are the distributions for muons from Z boson decays (solid line) and the muons found by the level-1 trigger in the minimum bias sample (dashed line). All towers within a 11-by-11 square (a) E_T^{tot} , the towers immediately around the muon (b) E_T^{in} and all towers except those immediately around the muon (c) E_T^{out} . Here, the extrapolation is not applied. The last bin represents the overflow.

on the other hand need to be corrected. In figs. 4.19 and 4.20 the resulting energy sum distributions are shown without and with extrapolation, respectively. While the total energy slightly shifts towards higher values around muons from the minimum bias sample, E_T^{in} shifts slightly towards higher values for the prompt muons. However, overall the effect is small.

Generally, the total energy E_T^{tot} , in fig. 4.19(a), shows little promise as the discriminating power between muons from Z bosons and those in the minimum bias sample is very small. Already fig.4.18(b) showed that the average energy is rather small. However, looking at the inner cone energy E_T^{in} , cf. fig. 4.19(b), one can clearly see the muon foot-print from the prompt muons. Also the outer cone energy E_T^{out} , shown in fig. 4.19(c) does not show a lot of discriminating power. But still, after removing the muon energy by

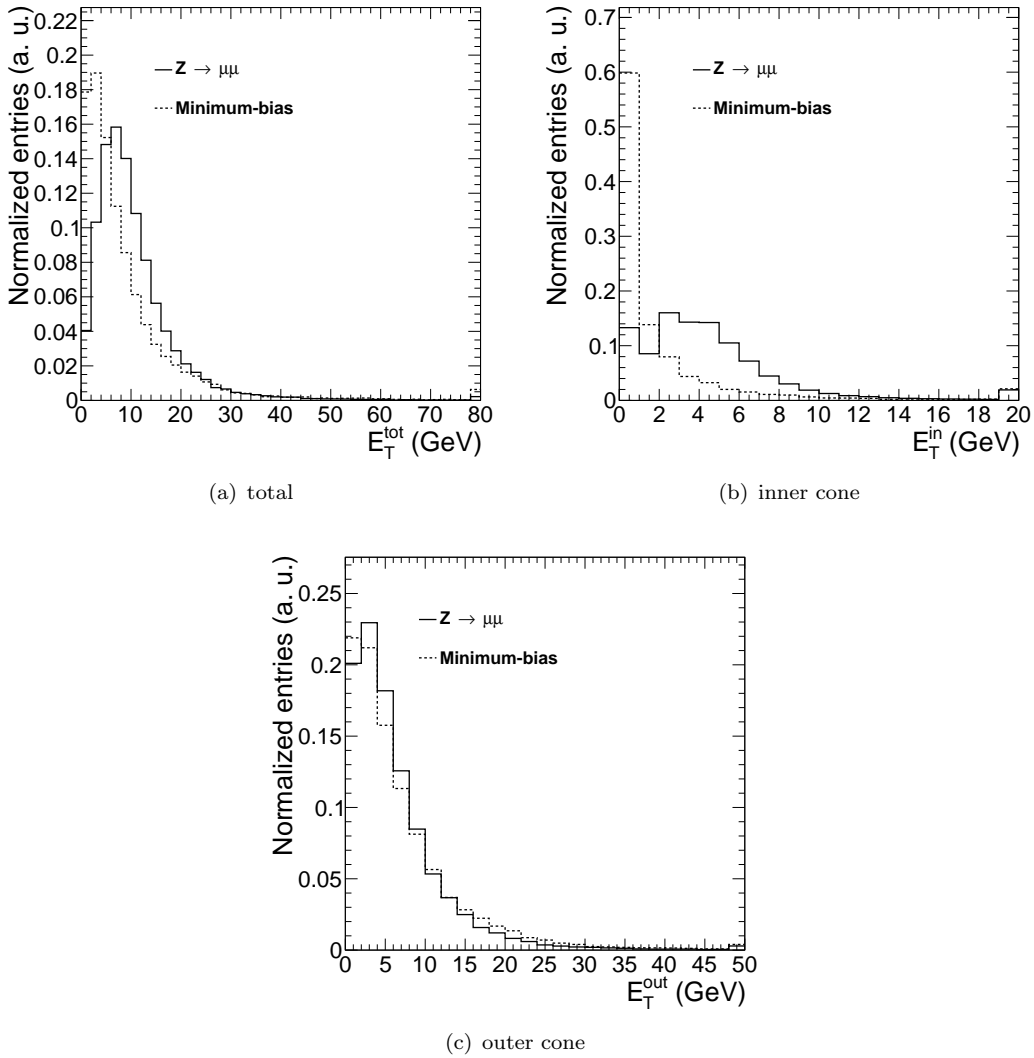


FIGURE 4.20 Sum of tower energies in cones around the extrapolated muon, shown are the distributions for muons from Z boson decays (solid line) and the muons found by the level-1 trigger in the minimum bias sample (dashed line). All towers within a 11-by-11 square (a) E_T^{tot} , the towers immediately around the muon (b) E_T^{in} and all towers except those immediately around the muon (c) E_T^{out} . The last bin represents the overflow.

subtracting the inner cone energy, the energy around muons in the minimum bias sample tends towards higher values with regard to the muons from Z boson decays.

Improving discriminating power – The ratio of energy deposits around the muon and the muon track transverse momentum, is a selection criterion that is used in the particle-flow relative isolation [44]. This ratio could also be calculated in the level-1 trigger. In fig. 4.21(a), the ratio E_T^{out}/p_T^μ is shown, where p_T^μ is the muon transverse momentum reported by the track-finder. The discrimination between muons contributing to the rate and those that contribute to the efficiency is improved by using this ratio. While the ratio for muons from the Z boson decay tends to smaller values, it tends towards larger values for muons from in-flight decays.




FIGURE 4.21 Sum of tower energies, E_T^{out} divided by the level-1 muon p_T in fig. 4.21(a) and divided by the total sum of energy E_T^{tot} in fig. 4.21(b). The dashed line for muons found in the minimum bias sample and the solid line for muons from Z boson decay. In fig. 4.21(a) the last bin serves as an overflow bin.

In fig. 4.18, the different topology of tower energies for the two samples is clearly visible. It is possible to exploit this fact by looking at the ratio of energy in the outer cone to the overall energy $E_T^{\text{out}}/E_T^{\text{tot}}$, cf. fig. 4.21(b). While this ratio is relatively evenly distributed across the whole range $[0, 1]$ for isolated muons from the Z boson decay, it clearly peaks at one for the muons found in the minimum bias sample.

4.5.4 Estimation of Performance

The performance of the isolation variable can be estimated with the receiver operating characteristic (ROC) curve. These curves are defined for binary classifier systems as false positive as a function of true positive at different thresholds of the discriminating variable. In the case of isolation for the level-1 trigger, these can be interpreted as the rate and the signal efficiency. In the following, ROC curves are determined with the resulting rate and efficiency of the μ GMT exploiting the previously illustrated isolation variables relative to the μ GMT not applying any requirements on isolation variables.

Absolute isolation – Since the distributions of E_T^{out} shown in fig. 4.20(c) are very similar for the muons from Z boson decay and minimum bias sample, the resulting ROC curve shows that this variable is not very useful as a discriminator. Essentially, applying a threshold for the outer cone energy results in the same effect a pre-scale would have; Both efficiency and rate are reduced by the same amount, cf. fig. 4.22. The performance of an absolute isolation is independent of the applied p_T threshold.

Relative Isolation – Given the level-1 muon p_T and the tower energies, a relative isolation variable, E_T^{out}/p_T , can be calculated. This introduces an explicit dependency




FIGURE 4.22 Relative efficiency as a function of relative rate for different isolation thresholds applied for the absolute isolation variables. The different colours show the influence of changing the p_T threshold.




FIGURE 4.23 Relative efficiency as a function of relative rate for different isolation thresholds applied to the relative isolation variable. The different colours show the influence of changing the p_T threshold.

on the p_T assignment to the isolation threshold. This effect is also visible in the corresponding ROC curve, cf. fig. 4.23. For higher thresholds, the discrimination power is reduced. However, especially for lower p_T thresholds, 10 and 16 GeV, working points can be defined where a few percent are lost in efficiency while the rate is reduced by 10-20% (depending on the applied p_T threshold). This can be used, for example, in di-muon triggers, where thresholds are typically lower.

Applying a p_T threshold of 16 GeV leads to efficiency and rate distributions as shown in figs. 4.24 and 4.25, respectively. The dashed line represents the μ GMT without additional isolation criterion applied, while the data points show the performance with relative isolation. Here, for the isolation threshold a working point is chosen that leads to an overall efficiency loss of 10%. Looking at the efficiency as a function of p_T , shown in figs. 4.24(c) and 4.24(d), the effect of the additional p_T dependence becomes visible; The steepness of the turn-on is reduced when applying a threshold on the relative isolation. This also explains why the efficiency loss as a function of the angular position is so low:

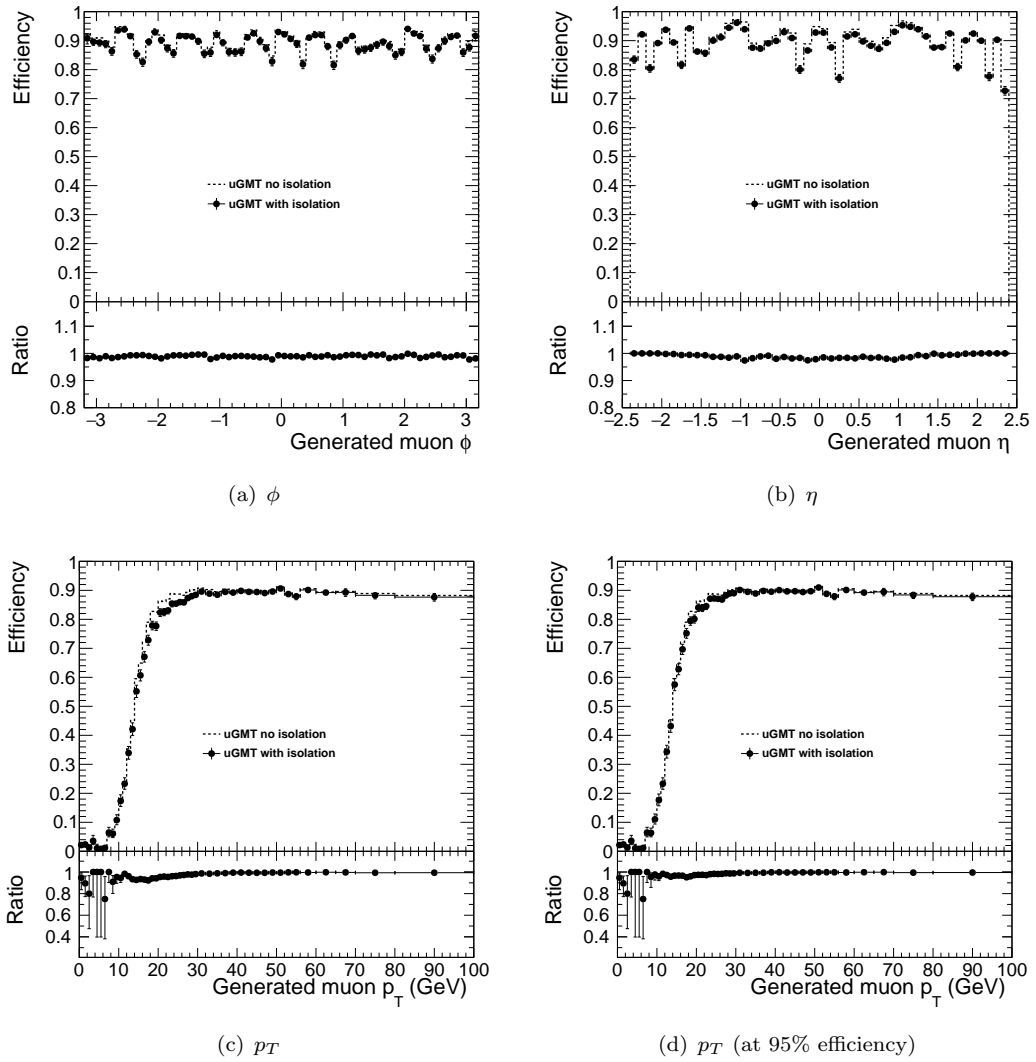


FIGURE 4.24 Efficiency of μ GMT with (points) and without (dashed line) relative isolation applied. The lower halves show the ratio of the two. For figs. (b) to (c) a isolation threshold resulting in 90% overall relative efficiency was applied, while in fig. (d) a 95% relative efficiency working point was chosen.

To disentangle from the turn-on effect, only generated muons with a $p_T > 24$ GeV are considered in figs. 4.24(b) and 4.24(a). Overall, the efficiency loss as a function of η or ϕ is flat.

The rate distribution as a function of level-1 η in fig. 4.25(a) shows that applying a relative isolation threshold slightly increases the already present fluctuations due to detector geometry and the discrete level-1 η scale. The rate versus p_T threshold distribution, shown in fig. 4.25(b), reveals that the application of an isolation threshold that corresponds to a 10% efficiency loss has the same effect as increasing the level-1 p_T threshold by 2 GeV. Increasing the p_T threshold results in an overall efficiency loss of roughly 5% (depending on the thresholds). Fig. 4.26 shows the efficiency reduction as a function of the generated muon p_T for increasing the threshold as well as requiring relative isolation. This shows that the relative isolation has little benefits compared to applying an increased p_T threshold.

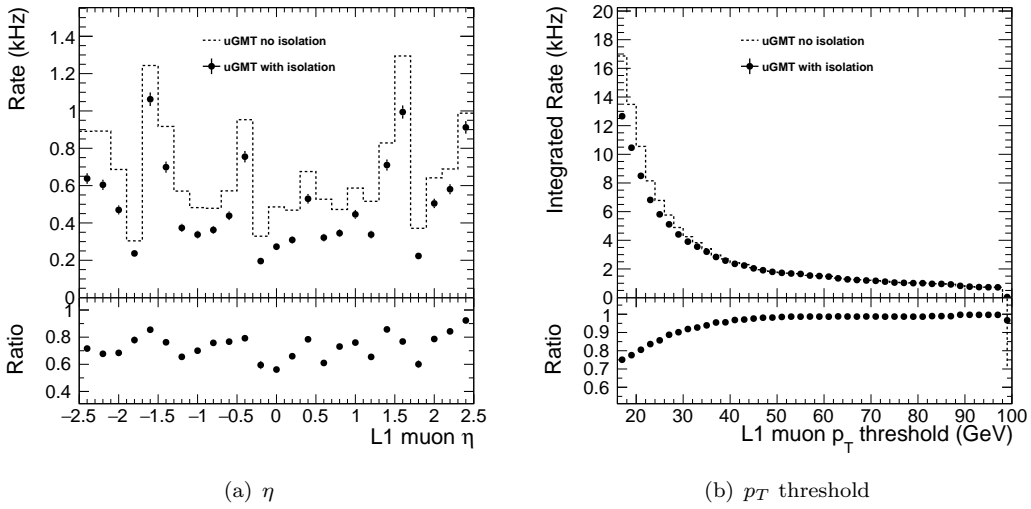


FIGURE 4.25 Rate of μ GMT with (points) and without (dashed line) relative isolation applied. The lower halves show the ratio of the two. A threshold was chosen that results in 90% overall relative efficiency.




FIGURE 4.26 Comparing the efficiency loss as a function of p_T for increasing the p_T threshold by 2 GeV (in red) or applying a threshold with 10% efficiency loss to the relative isolation.

Overall, the relative isolation is not an ideal candidate to reduce rate and maintain a high efficiency. However, it can still be useful in di-muon triggers where low thresholds are applied and in analyses that require a relative isolation on the reconstructed muons which also decreases the acceptance at low muon p_T . A threshold has to be selected that only reduces the efficiency by a few percent to gain any benefit compared to increasing the p_T threshold, compare figs. 4.24(c), 4.24(d) and 4.26.

Muon confirmation – Another approach is to exploit the difference in shape seen for E_T^{in} , cf. fig. 4.20(b). Conversely to the previously used variables, here, a lower threshold will be applied; Only muons that have a minimum energy deposit in E_T^{in} will be selected. The corresponding ROC curve is shown in fig. 4.27. While the variable discriminates very well between muons contributing to the level-1 rate and those that contribute to the

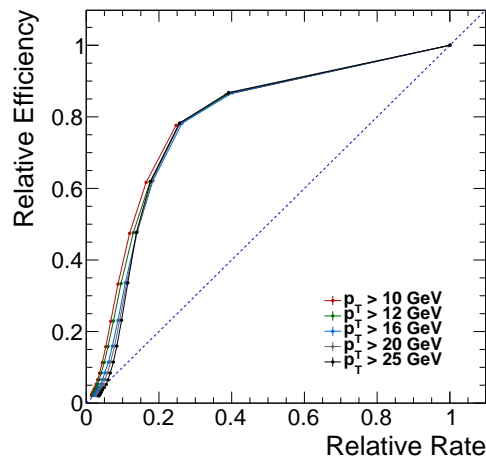


FIGURE 4.27 Relative efficiency as a function of relative rate for different thresholds applied to the inner cone energy. The different colours show the influence of changing the p_T threshold.

efficiency, the downside is the precision of the available energy information resulting in discrete working points, as can be seen in the curve: Each point in the graph represents an increase of the threshold by one unit of the discrete scale, corresponding to 0.5 GeV. Therefore, it will be difficult to tune this variable to a desired working point of a given minimum efficiency. The minimum overall efficiency loss is roughly 15% and increasing the required energy by 0.5 GeV results in steps of 5% and 10%.

When choosing the lowest efficiency loss working point, resulting in a 15% drop, and simultaneously applying a 16 GeV p_T threshold, the efficiency and rate distributions shown in figs. 4.28 and 4.29 are obtained. While the efficiency loss as a function of ϕ is flat, a pronounced η dependency is visible: At the border of the OMTF acceptance, the efficiency begins to drop to 40-70%. This effect can be explained by the change in granularity of the tower energy in this region and will be discussed in further detail later. If this variable were to be used, it could only be applied to muon candidates within $|\eta| < 1.25$.

The rate reduction on the other hand is quite significant. Fig. 4.29(a) shows that the reduction is slightly larger for muons beyond $|\eta| > 1.25$, where the rate drops to 10 to 25%. In the more central region it is reduced to approximately 35% to 40%. The region $|\eta| > 1.25$ contributes a larger part of the overall rate. Therefore, restricting the usage of this variable to the central detector region reduces its performance significantly. As a function of p_T threshold, the reduction is roughly constant. At the 85% efficiency working point, the rate is reduced to 20% to 25% for all thresholds, cf. fig. 4.29(b).

Energy ratio – The last variable described previously is the ratio of the outer cone to the full cone energy, $E_T^{\text{out}}/E_T^{\text{tot}}$, cf. fig. 4.21(b). The resulting ROC curves for different level-1 p_T thresholds are shown in fig. 4.30. Similarly to using the inner cone energy, a minimum efficiency loss has to be accepted. However, this loss is slightly lower, around 10%. This efficiency loss would result in a slightly smaller rate reduction of 50%.

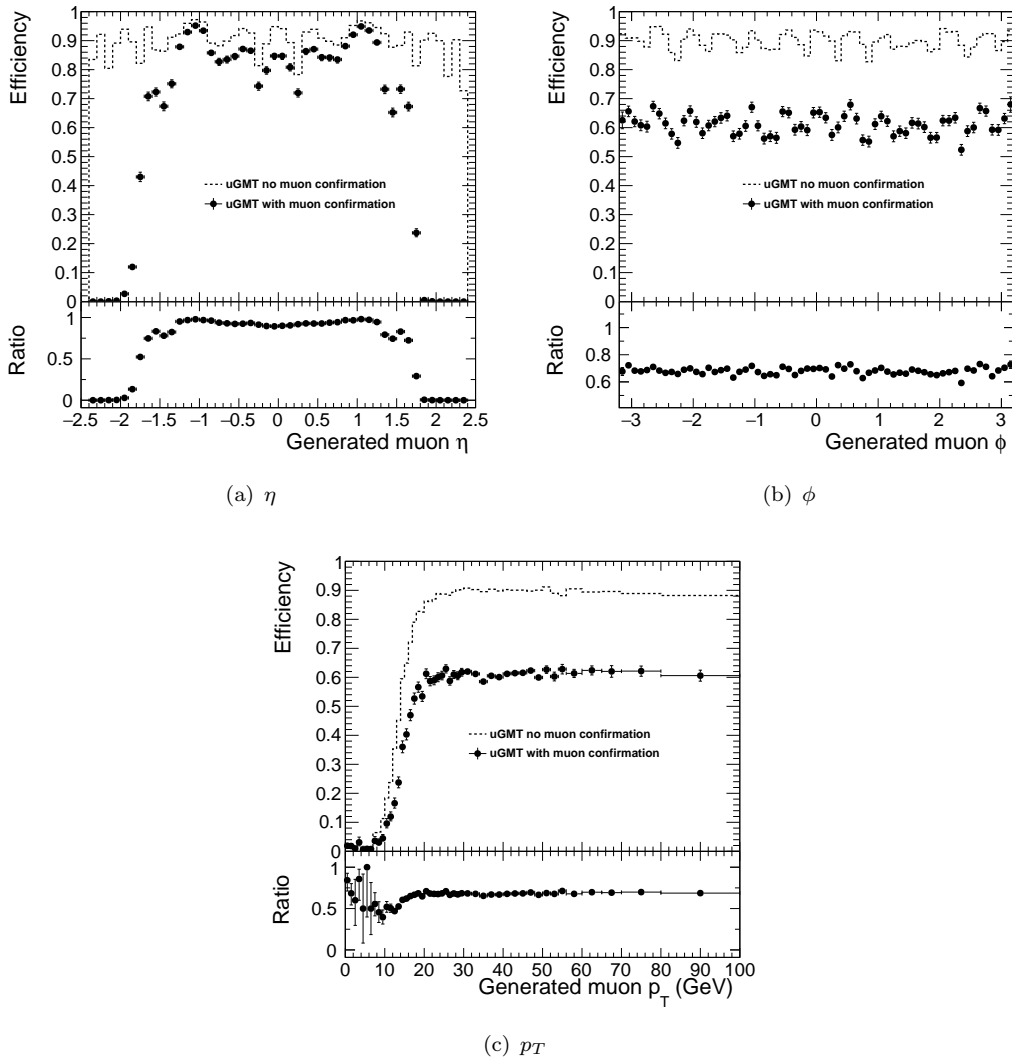


FIGURE 4.28 Efficiency of μ GMT with (points) and without (dashed line) relative isolation applied. A 85% relative efficiency working point was chosen to determine the threshold.

Applying a 16 GeV p_T threshold and a threshold on the cone energy ratio that corresponds to the 90% efficiency working point, results in the rate and efficiency distributions shown in figs. 4.31 and 4.32. Similarly to using E_T^{in} directly, using the ratio has a similar η dependence, cf. fig. 4.31(b). Figs. 4.31(a) and 4.31(c) show that the reduction of efficiency as a function of ϕ and p_T (at the plateau) is flat.

The corresponding rate distributions, fig. 4.32, show that the rate reduction is again flat in L1 p_T threshold. Also similarly to the previously shown isolation variable, the rate reduction as a function of η follows the structure present due to the detector structure.

Energy ratio (with η restriction) – The strong η dependence observed when using the inner cone energy can be explained by looking at the available calorimeter information. While for central towers ($|\eta| < 1.83$), the correspondence between actual tower energy and the E_T index assigned in the level-1 calorimeter trigger system shows a narrow peak around the nominal value of 0.5 GeV, this distribution has a much wider




FIGURE 4.29 Rate of μ GMT with (points) and without (dashed line) applying a threshold to the inner cone energy. The lower plot in each figure shows the ratio of the two. A threshold was chosen that results in 85% overall relative efficiency.




FIGURE 4.30 Relative efficiency as a function of relative rate for different thresholds applied to the ratio of outer cone energy to the full cone. The different colours show the influence of changing the p_T threshold.

spread for towers beyond that, cf. fig. 4.14. The cause of this difference is the lack of information from the HCAL for which the width in ϕ is doubled beyond $i_\eta = 21$.

The decrease can be mitigated by only applying the threshold for muons within a restricted η range. To fully avoid any influence of the towers, a range of $|\eta| < 1.8$ is chosen. The ROC curve derived with this restriction on η before applying the threshold on the energy ratio is shown in fig. 4.33. As can be seen, the efficiency loss is reduced, but also the rate reduction is smaller compared to applying the threshold over the full η range. However, at a loss of 5% in efficiency (the loosest working point), the rate can be reduced by 35%. At a 10% loss, the reduction can be increased to 45%.

For the comparison of efficiency and rate, the loose working point is chosen in fig. 4.34, where also a 16 GeV p_T threshold is applied. As for the previously shown distributions, the efficiency loss in ϕ is flat, cf. fig. 4.34(a). With the η restriction, the loss as a

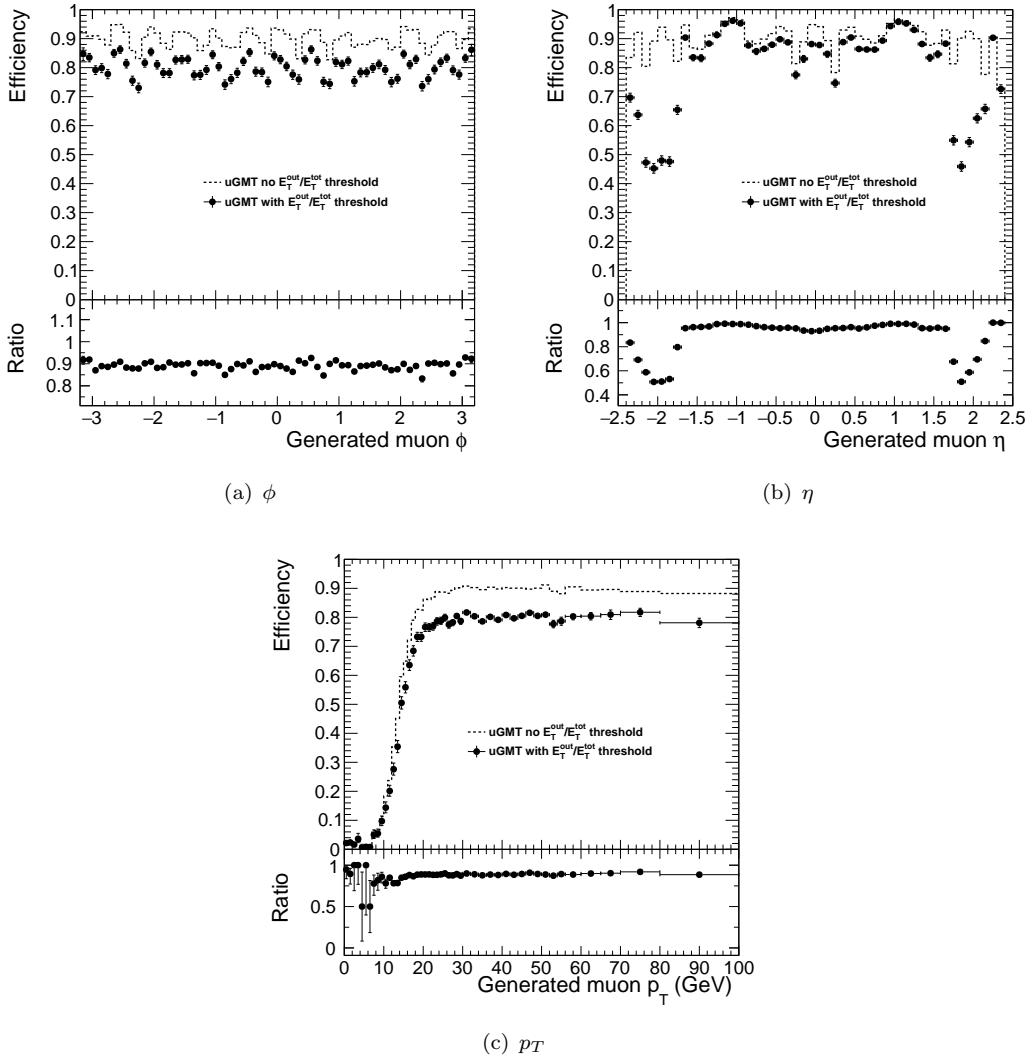


FIGURE 4.31 Efficiency of the μ GMT with applying a threshold to the energy ratio as described in the text, shown with points. The dashed line indicates the efficiency without any isolation criterion. Apart from the isolation also a p_T threshold of 16 GeV is applied.

function of η shown in fig. 4.34(b) exhibits a jump at the boundary of the range that is chosen. However, since a working point was chosen with relatively small efficiency loss, this jump is not very large. The turn-on in fig. 4.34(c), shows that the steepness is not significantly changed and the plateau is reached at the same generator muon p_T . Specifically, the point at which an efficiency of 90% of the plateau is reached does not change, which is the value that defines the p_T threshold value at level-1. The efficiency loss in the plateau region is flat. Also, the turn-on shows that only the efficiency for muons with very low p_T , i.e. below the p_T threshold, is reduced.

The reduction in rate, shown in fig. 4.35, exhibits the same η dependence as seen before. The restriction of applying the threshold to muons within $|\eta| < 1.8$, results also here in an additional jump at the corresponding boundary. As a function of the L1 p_T threshold, the decrease in rate is mostly constant for the whole range with a slight increase for very high thresholds.

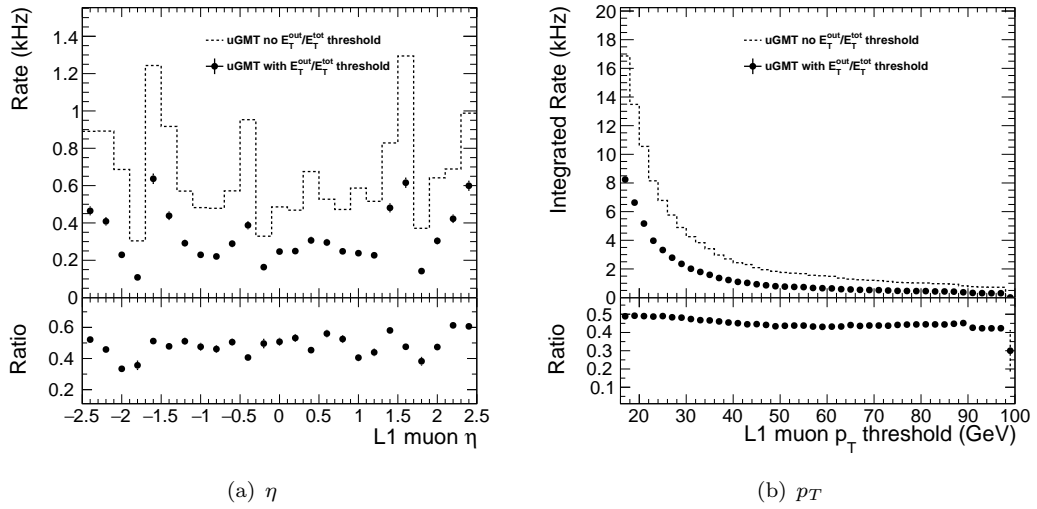


FIGURE 4.32 Rate of the μ GMT without (dashed line) and with (points) applied threshold on the ratio of outer and full cone energy. Additionally a L1 p_T threshold of 16 GeV is applied.




FIGURE 4.33 Relative efficiency as a function of relative rate for different thresholds applied to the energy ratio of outer and full cone. The threshold is only applied for muons within $|\eta| < 1.8$ to retain efficiency. The different colours show the influence of changing the p_T threshold

4.5.5 Application in the level-1 trigger

The performance estimation shown in the previous section indicates that an isolation algorithm as it is commonly used on reconstructed muons cannot be applied successfully at the level-1 trigger stage. Specifically, the absolute isolation has no positive influence while the relative isolation could be used but results in the same efficiency loss and rate reduction as applying a higher p_T threshold. This may seem surprising since similar variables are often used as selection criteria in analyses of CMS data. However, in the level-1 trigger, the precision and granularity of the calorimeter energy information is reduced. Additionally, the p_T assignment in the track-finders is not tuned for the momentum resolution but to have a high efficiency when placing a p_T threshold. Lastly,

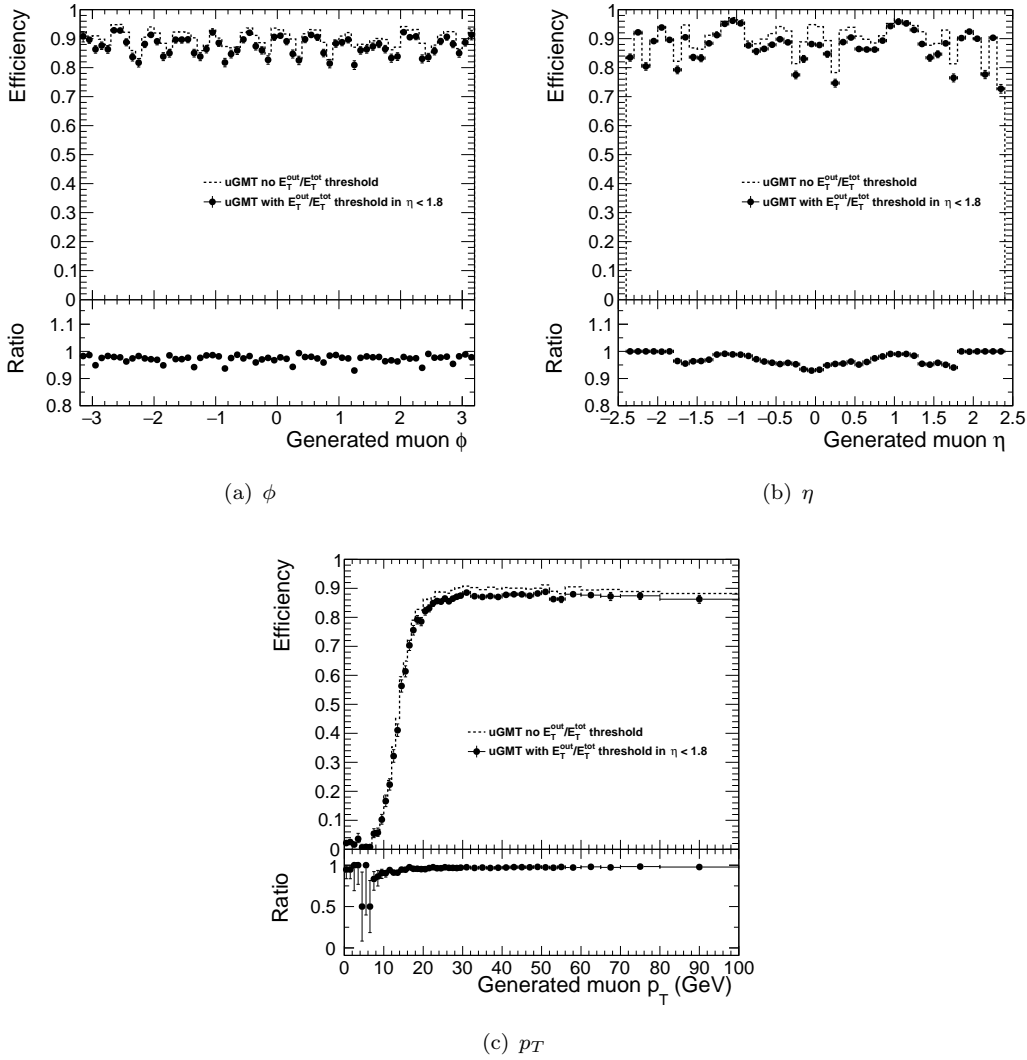


FIGURE 4.34 Efficiency when restricting the application of a threshold on the energy ratio to $|\eta| < 1.8$. Both efficiency with (points) and without (dashed line) applied threshold are shown.

the position resolution of reconstructed muons is much better than of the level-1 muon candidates since tracker information can be used.

However, the confirmation of a muon through the energy deposited in the calorimeter exhibits a good discrimination between non-prompt muons that dominate the level-1 rate and prompt muons from a decaying Z -boson. The down-side of applying such an algorithm is that due to the discrete energy-scale that is available, a chosen working point can not be finely tuned. This can be improved by calculating the ratio of the total cone energy and the outer energy. While the total rate-reduction is slightly lower, the efficiency loss working point can be freely chosen. Due to the reduced ϕ granularity of the tower energy information at high $|\eta|$, such an algorithm would have to be limited to $|\eta| < 1.8$. This results in a rate reduction of 35-45% at an efficiency loss of 2-10%.

The above mentioned algorithms were evaluated assuming full trigger tower granularity being available and also using the full energy scale that is available in the layer-2 calorimeter trigger. This information would not be available in the μ GMT but only in

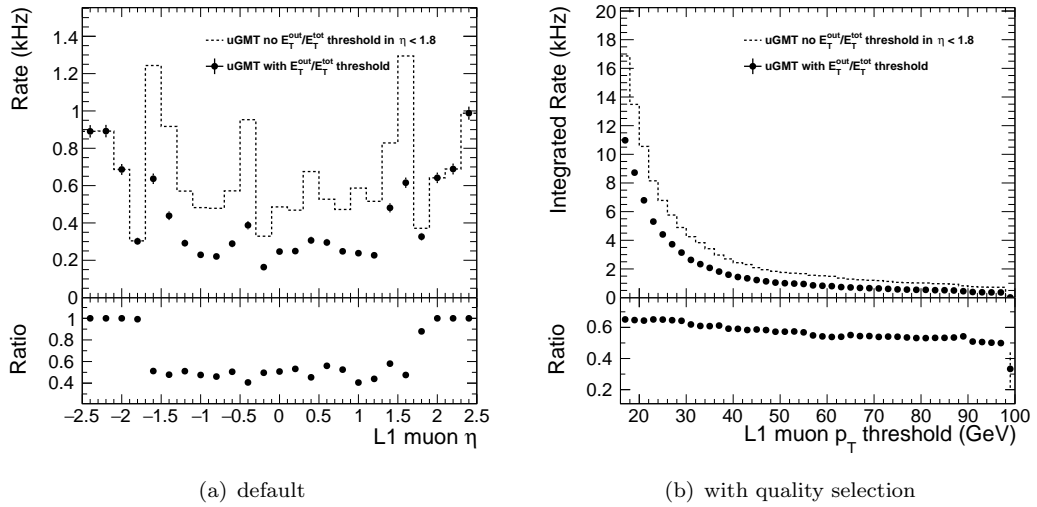


FIGURE 4.35 Rate of the μ GMT with η -restricted (only for muons with $|\eta| < 1.8$) energy ratio threshold. The rate is shown both with (points) and without (dashed line) applied threshold.

the layer-2 calorimeter trigger. Here, instead of pre-calculating the single cone energies, both the outer and full cone could be calculated. Then the ratio of the two energies could be sent to the μ GMT at a reduced energy precision. Alternatively, a threshold could already be applied and the corresponding bit could be sent.

In either case the information would have to be available at tower granularity to reproduce the performance shown in this section. However, since the tower energy information beyond tower 21 is not sufficient for an efficient algorithm, any information beyond that could be discarded. This would reduce the number of towers to 42×72 . For the energy information 28 links are foreseen in the μ GMT. When the threshold is applied and only one bit is transmitted this would fit into the budget of 28 links at 10 Gb/s corresponding to 192 bit per bunch crossing ($42 \times 72 = 3024 < 28 \times 192 = 5376$). Transmitting the ratio itself, however, is not feasible at a viable energy precision.

5 The Upgraded Global Muon Trigger

This chapter gives a more detailed description of the μ GMT algorithm as it is implemented both in the emulation code and gateway. First the input data are described, followed by an overview of the building blocks of the algorithm. The chapter is closed by a short discussion of the gateway implementation.

5.1 Input Data

The μ GMT receives data both from the calorimeter Layer-2 and the track-finders, as mentioned in section 3.4.2. The data are transmitted on 10 Gb/s optical links with 10b/8b encoding. 32 bit words are sent at a frequency of 240 MHz allowing a total of $6 \times 32 \text{ bit} = 192 \text{ bit}$ per link and bunch-crossing to be transmitted. The MP7 processing board in which the μ GMT is implemented has a total of 72 input channels of which 36 are used for muon data and 28 for calorimeter data.

5.1.1 Muon Data

In the upgraded trigger, muon objects are represented by 64 bit words. A total of 108 Muons are received, 36 muons each from the three track-finders, BMTF, OMTF, and EMTF. Eight properties are transmitted and encoded in the data. Table 5.1 shows the bit-width available for a given property, the scale factor to calculate the physical meaning of the integral number, the possible range of the integral number and gives a comment on special meaning. For p_T the transformation between the integral value and the floating point physical value is simply

$$p_T = (i_{p_T} - 1) \cdot 0.5 \text{ GeV}, \quad (5.1)$$

where i_{p_T} is the integral values that were transmitted. The value $i_{p_T} = 0$ is reserved for invalid muon candidates. The values of η and ϕ are encoded in 2's complement, where the most significant bit is used to specify the sign. The transformation of η is given by

$$\eta = \left(-a_8 2^8 + \sum_{i=0}^7 a_i 2^i \right) f_\eta, \quad (5.2)$$

where a_i are the individual bits of the binary representation of the integral value and f_η is the scale factor given in table 5.1. The ϕ coordinate is transmitted in a track-finder

specific local coordinate system. To calculate ϕ in the global CMS coordinate system, an offset has to be applied that depends on the wedge or sector in which the muon was detected. The transformation is therefore given by

$$\phi = \left(-a_7 2^7 + \sum_{i=0}^6 a_i 2^i \right) \cdot \frac{2\pi}{576} + \phi_{\text{offset}}, \quad (5.3)$$

where a_i are again the individual bits of the integral value. ϕ_{offset} is the offset that is given by the wedge or sector. The offset is always the ϕ coordinate corresponding to the lower edge of a given sector or wedge.

Information about the charge of the muon is transmitted in two bits, the charge sign bit s and the valid charge bit. Through the former, the charge q_μ may be calculated with

$$q_\mu = (-1)^s e, \quad (5.4)$$

with the electric charge of the electron e . The valid-charge bit gives information whether the charge could be determined at the track-finder level: For highly energetic muons, the bending is not sufficient to reliably measure the charge and the valid bit is set to 0.

The quality definition depends on the track-finder but in general depends on the number of hits used to form the track of the muon and the compatibility with a muon originating from the vertex. More details of the current implementation of muon qualities was discussed in section 4.4.4. The halo or η -fine (H/F) bit gives information whether the muon is compatible with a halo muon when the muon was sent from the EMTF or whether the η coordinate could be determined with high precision when the muon is sent by the BMTF. Finally, the 27 bit wide track address allows the identification of the hits and track segments used in the reconstruction of the muon.

Information about the bunch-crossing is sent per link instead of per muon: The least significant bits of the bunch counter BX0/1/2 are sent and in the first bunch-crossing of the orbit, the muon data is marked through the BC0 bit.

5.1.2 Calorimeter Data

While for the muon objects several properties are transmitted, the energy sums transmitted from the calorimeter trigger are simply the energy values. Each value represents the energy sum calculated for an programmable area, with the default of a 4 tower radius. The numbers are encoded in 5 bits, where the highest value is used as an overflow bin. This results in the possible transmission of 6 energy-sums per link per 240 MHz clock cycle or $6 \times 6 = 36$ energy sums per link per bunch-crossing.

The energy information is only needed within the acceptance of the muons, $|\eta| < 2.42$, which would correspond to 2×26 towers in η , cf. table 3.1. In ϕ , all 72 towers are needed. The position information of the tower corresponding to a given energy value does not have to be transmitted explicitly, instead the position in the transmitted data could be used to derive the tower index and with this derive the physical η and ϕ from a corresponding map.

TABLE 5.1 Given in this table are the scale definitions of muon candidates transmitted to the μ GMT: Name of the property, number of bits, the unit (value of the least significant bit), the integer range used and a comment.

Parameter	n_{bits}	Scale factor f	Range	Comment
p_T	9	0.5 GeV	[0, 511]	0: empty candidate
η	9	0.010875	[-240, 239]	2's complement. $value \times f$ gives the centre of the bin.
ϕ	8	$2\pi/576$	[-5, 218]	2's complement. $value \times f$ gives the lower edge of the bin.
charge sign	1		[0, 1]	1: negative, 0: positive
valid charge	1		[0, 1]	1: charge sign is valid, 0: charge sign cannot be determined
quality Q	4		[0, 15]	Definition dependent on track-finder.
H/F	1		[0, 1]	Indicates halo-muon or η -fine bit, depending on track-finder.
track address	27			to be defined
BC0	1		[0, 1]	1: first bunch crossing of orbit, 0: else. Received per link
BX0/1/2	3			3 least significant bits of the bunch-crossing counter. Received per link.
SYNC	1			Indicates synchronisation errors. 1: synchronisation lost, 0: no error. Received per link.

Section 4.5 introduced the muon confirmation and an isolation variable solely based on the ratio of cone energies. Possible scenarios of transmitting the necessary data to the μ GMT were also discussed. Essentially, the layer-2 calorimeter tower would pre-calculate the two needed cone energies and send the ratio to the μ GMT. Alternatively, since no information of the muon is needed, a threshold could already be applied in the layer-2 calorimeter trigger and only a bit per tower transmitted.

5.2 Algorithm Blocks

The main task of the μ GMT is to sort the muon candidates reported by the BMTF, OMTF, and EMTF. Additionally, candidates that potentially originate from the same muon should be identified and the duplicate with lower quality or, alternatively, higher p_T should be discarded. Lastly, an isolation can be calculated from energy deposits that are received from the Layer-2 calorimeter trigger. Fig. 5.1 illustrates the logic blocks that accomplish this and which are described in the following sections. The algorithm can be split into two main blocks, the sorting and cancellation unit and the isolation unit.

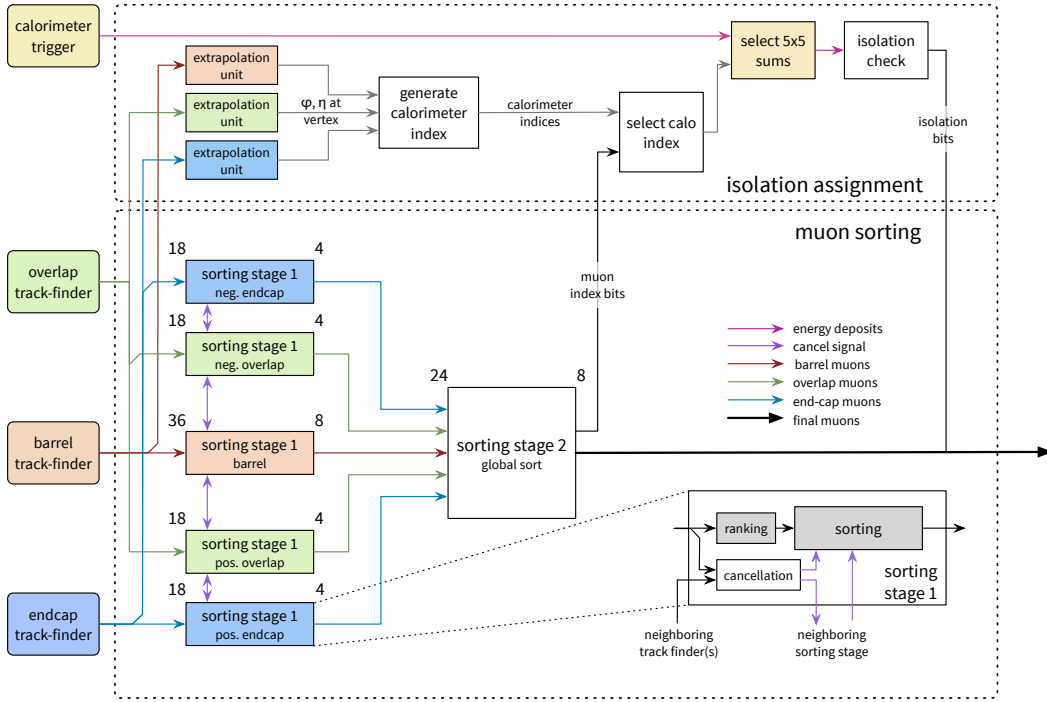


FIGURE 5.1 This block-diagram shows the data flow within the μ GMT algorithm. The data flows left-to-right, the different track-finder muons are indicated with red, green and blue arrows, representing barrel, overlap and end-cap track-finders, respectively. Calorimeter data are shown with magenta arrows. The numbers above the sorting stages indicate the number of muon candidates at start and end of the sort stage. Each sorting stage 1-block is a simplification as indicated in the inset box at the bottom right.

5.2.1 Sorting and Cancellation Unit

Sorting

In the sort algorithm, muons are sorted according to their quality Q and p_T . To streamline this sorting mechanism, a rank r is assigned for each muon:

$$r = f(p_T, Q), \quad (5.5)$$

where $f(p_T, Q)$ is a configurable, arbitrary function of p_T and Q . The algorithm compares each muon i with all other muons j that are relevant in that stage and calculates a weight w_i

$$w_i = \sum_{j=0}^N \theta(r_i - r_j), \quad (5.6)$$

where $\theta(r_i - r_j)$ is the Heavyside step function which gives 1 for $r_i \geq r_j$ and 0 if $r_i < r_j$. Thus, the n muons of highest rank can be selected out of the N muons by taking the muons with the weights $w = N - 1, \dots, N - n$.

Since comparing all 108 muons with each other in the same stage would require more logical resources than are available in the MP7, the sorting is split into two stages. In the first stage muons are sorted individually within barrel, positive overlap, negative

overlap, positive end-cap, and negative end-cap: For the barrel, the 36 input muons are sorted down to 8 muons and for the positive and negative overlap and end-cap regions the 18 down to 4. In the second stage, the remaining 24 muons are sorted down to 8, cf. lower left of Fig. 5.1.

Cancellation

Duplicate muons can arise either from one track-finder reporting it in two neighbouring ϕ sectors, as there is no information exchange between the different track-finder processors, or two neighbouring track-finders each reporting it. The cancellation in the μ GMT algorithm is therefore performed between the neighbouring track-finders and at the boundaries of the track-finder $60^\circ/30^\circ$ sectors in ϕ .

Two possibilities are available to the μ GMT to perform this algorithm and different approaches may be used for different regions:

- ▷ *Coordinates* (η, ϕ) may be used to match muon candidates in a window $\Delta r = \sqrt{\Delta\phi^2 + \Delta\eta^2} < x$.
- ▷ *Track-addresses* that encode the track segments and hits that are associated with the muon could be used to identify muons that have any in common.

For example, in the barrel track-finder region, the 7 bit encoding of track segments used in the legacy system could be used and extended to the RPC hits by using some additional bits. However, for the overlap region, where pattern recognition algorithms are employed, the encoding is not straight-forward. The exact nature of cancellation is still evolving and for the moment the cancellation based on coordinates is implemented and will be used during the start up of the upgraded trigger. Especially in the track-finder transition regions, a matching based on track addresses is unlikely to be functional due to the differences in the algorithms of the track-finders.

Whichever method is chosen to identify muons that may originate from the same muon, they have to be cancelled against each other. This will be based on the quality code or p_T that is assigned to each muon: The one with lower quality or higher p_T will be discarded to decrease the rate while accepting a small reduction in efficiency. The cancellation takes place in parallel to the first sorting stage and the weight of cancelled muons is forced to 0.

5.2.2 Isolation Unit

Calculating the isolation requires the position of the muon at the vertex, as was discussed in section 4.5.3. For all 108 incoming muons (η, ϕ) at the vertex is determined with LUTs that take both η and p_T as input, in other words

$$\Delta\phi = f(\eta, p_T) \quad \text{and} \quad \Delta\eta = f(\eta, p_T). \quad (5.7)$$

The corresponding LUTs are generated independently for barrel, overlap and end-cap. Therefore, the most significant bits that represent the coarsest η information are not necessary. Also, the muon is assumed to have propagated approximately straight, when $p_T > 32$. These two restrictions allow to reduce the size of the LUT significantly.

In a second step these coordinates at the vertex are transformed to tower indices, using the mapping described in section 3.3. Now, for the final 8 muons that were selected in the sorting stage, these tower indices are used to select the energy sums to be associated.

Both absolute and relative isolation bits are possible to be calculated, with

$$b_{abs} = \theta(E_T^{sum} - T_{abs}) \quad b_{rel} = \theta\left(\frac{E_T^{sum}}{p_T^\mu} - T_{rel}\right), \quad (5.8)$$

where E_T^{sum} is the selected calorimeter energy sum, p_T^μ is the transverse momentum of the muon, θ is again the Heaviside step function, while T_{abs} and T_{rel} are programmable thresholds. Additionally, investigations have started on how to implement an algorithm that corresponds to the energy ratio shown in section 4.5.3. One likely implementation could be to calculate the ratios in the calorimeter layer-2 processors. The result could then be transmitted at a tower-by-tower granularity. Alternatively, a bit per tower could be transmitted that indicates whether the ratio in that tower was below or above a programmable threshold. The μ GMT would then only map the muon candidates extrapolated η and ϕ to the corresponding tower index and use the transmitted bit.

5.3 Gateware

The previously described blocks are what is currently implemented as the trigger algorithm of the μ GMT. Alongside these blocks, additional logic is required to de-serialise and serialise the muon data and implement the link protocols. Furthermore, logic is needed to allow reading out the in- and output data of the μ GMT and make it available for DAQ-readout.

The gateware is written in a modular way in VHDL (Very high speed integrated circuit Hardware Description Language). This modular description allows that commonly needed logic, as link synchronisation and protocol, is provided centrally by the MP7 developers. The developers of the actual trigger logic are free to write their algorithms which receive data and control signals from the common logic modules.

While the main algorithm runs in synchronisation with the LHC clock at 40 MHz, the (de-)serialisation stages are implemented with a 240 MHz clock as they are required to receive and send 32 bit words at this frequency. Some look-ups are also performed at this higher clock speed to reduce the needed memory resources. Since rank calculation and extrapolation only depend on the information contained in one 32 bit word, they can be performed in the 240 Hz clock region before de-serialisation. This allows to re-use the same memory blocks 3 times (for the three muons transmitted per link) and therefore reduce the memory needs for these calculations by a factor of three.

Additionally, controls are available that allow synchronisation with sender and receiver cards to ensure that the received muons are all from the same bunch-crossing. Inputs

may also be masked per track-finder by writing to dedicated memory registers in case problems are encountered during a data taking run. Error counters are available to test whether the synchronisation with senders are lost and may be read-out per link. Additional counters allow to estimate the rate at which muons are received from the different track-finders in order to monitor their behaviour.

6 Online Software

In this chapter the existing infrastructure for the online software that controls and monitors the μ GMT during data taking is described. First the common components used in the level-1 trigger are mentioned and then the implemented software specific for the μ GMT is discussed.

6.1 Common Software Components of the Upgraded L1-Trigger

Trigger Supervisor – The trigger supervisor is the control and monitoring software suite of the level-1 trigger. It is divided into so-called cells that control parts of the system at various granularities (individual boards to systems). A central cell delegates commands from the central CMS run-control to the different sub-system cells with SOAP commands. The cells can be accessed and controlled via a web-based user interface that is based on a custom server-side software library AJAXELL [45].

IPbus – The μ TCA technology provides Gigabit Ethernet connections for communication with the hardware. However, no specific protocol is predefined. The CMS collaboration has developed the packet based communication protocol IPbus [46] based on UDP. Via this protocol, read and write access can be established. A software and gateway suite that implements the needed functionality is available and used in almost all components of the upgraded trigger (the only exception is the EMTF, using PCIe). To access the mapped memory within the chip, an address table is defined that connects the memory addresses with human readable node names. Each node can represent a block-RAM, register or FIFO.

Communication from several PCs to the same hardware is possible via the control hub. The control hub is a software application that arbitrates simultaneous access to one device from several control or monitoring applications. Each connected application is handled by a separate process within the control hub. The connections from control applications to the control hub use the TCP communication protocol that has more sophisticated flow control and congestion mitigation algorithms than UDP.

SWATCH – Many of the trigger components share the same hardware, therefore, a control and monitoring software layer was introduced that is common between the

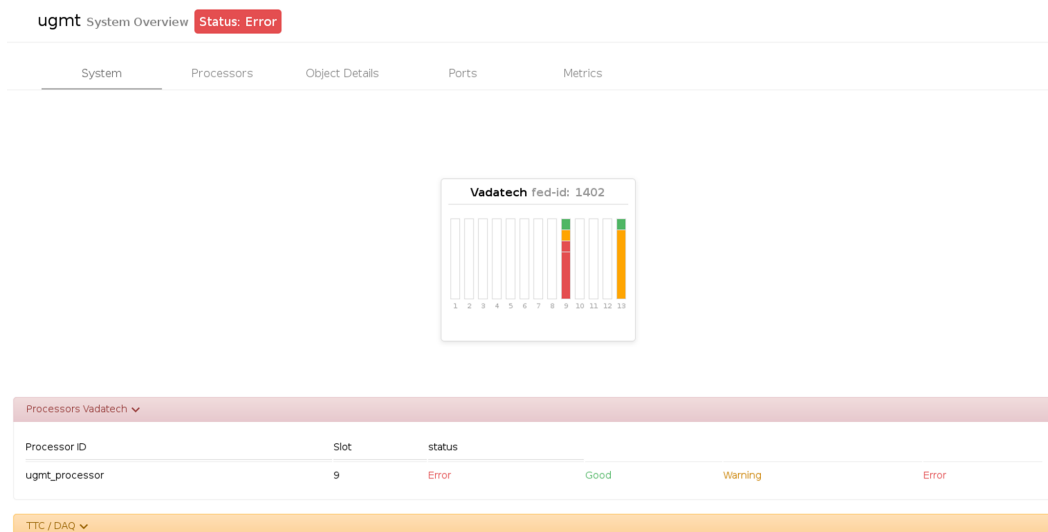


FIGURE 6.1 Demonstration of the system overview user interface developed for the central online software. For this screen-shot, the cell was connected to a demonstrator system with more than one board and crate.

different components. This layer is in-between the actual drivers that interact directly with the trigger processors and the high level software that serves the user interfaces. This SoftWare for Automating conTRol of Common Hardware (SWATCH) unifies the interfaces for the different hardware to make the control and monitoring software as uniform as possible [47].

6.2 Upgraded GMT Configuration

The configuration of the μ GMT for the most part involves loading and verifying LUTs. From the software point of view, these can be seen as a space in memory that is accessible through IPbus. All LUTs are automatically discovered through attributes specified in the address-table. The contents of a LUT can either be served from a local file or, in the future, from a table in the configuration data-base. The LUTs are first written to the card and then read-back to be compared with the original content source. If a mismatch is found an error is flagged in the attached monitoring node.

Apart from LUTs, four configuration registers are exposed to the control software that allow to mask all inputs from one of the four sources of μ GMT inputs: BMTF, OMTF, EMTF and calorimeter trigger. Per-channel masking is also available through the standard control software.

As the μ GMT is based on the MP7 and integrated in the infrastructure gateway, the corresponding standard control software can be used that includes the configuration (and alignment) of input and output ports, as well as the read-out logic.

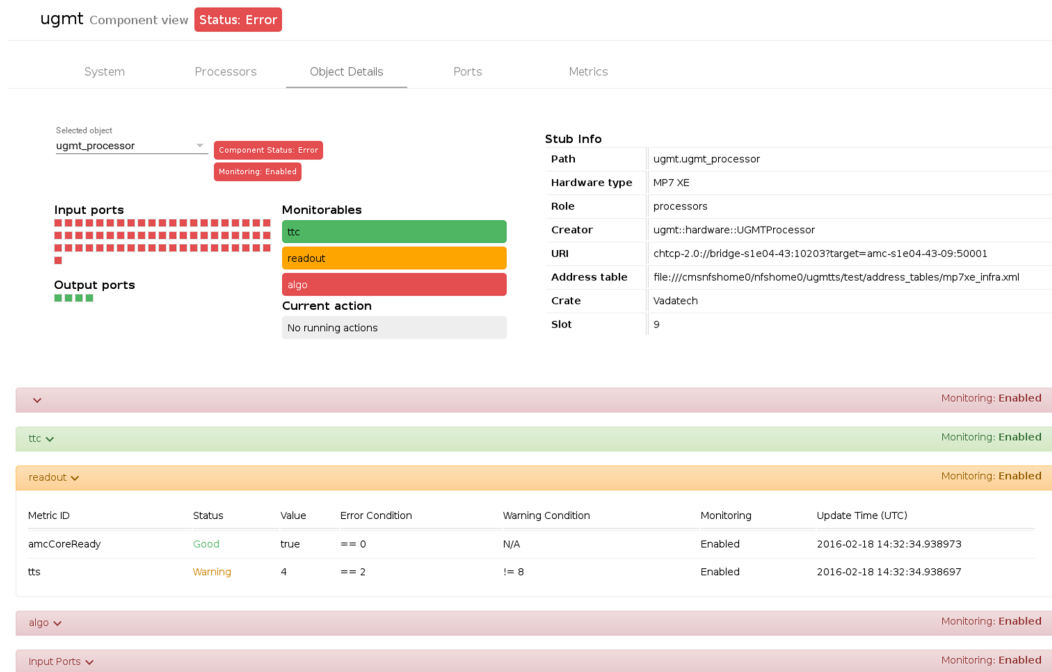


FIGURE 6.2 Demonstration of the object-detail user interface developed for the central online software.

6.3 Upgraded GMT Monitoring

For monitoring purposes, several counters are implemented in the μ GMT and exposed to the control and monitoring software. For one, the bunch counters transmitted by the track finders are compared to the internal counter and any mismatch increases a counter. Similarly, the BCO flag is checked for synchronisation. In addition, the number of muons received from the track finders is counted in the firmware. This can give an estimate on the muon rate and can be used to discover problems with any of the upstream systems.

Because of SWATCH and the common hardware used throughout the upgrade project, many monitoring tasks can also be done in a centralised manner. Link (alignment, CRC check sum, etc.), TTC and DAQ monitoring are of interest to all systems. Also, due to the common interfaces defined in SWATCH, user interfaces for monitoring only need to be developed once and then can be used for all systems.

Displaying monitoring data – Making the monitoring data available to the operations team and publish information about errors is essential to ensure reliable physics data taking. Therefore, panels were designed that display this data in a concise manner. Fig. 6.1 shows a screen-shot of the overview panel. Here, all crates that belong to a system are displayed prominently at the top. The overall status is collected and displayed as a flag at the top of the page. The status of individual processing boards is represented by common colour-coding (red for errors, yellow for warnings and green if no problems are encountered) and displayed in the position where the boards can be found in the corresponding μ TCA crate. Additionally, sub-components of the processor, like the TTC status, DAQ link or overall algorithm status are indicated by blocks within

that display. Here, the same colour coding is used to indicate problems. Below the crate overview, several buttons (also colour coded depending on the component status) are shown that can be clicked to show more detailed monitoring information in case problems are encountered. This allows the user to find sources of problems while also having a comprehensible overview by default.

Fig. 6.2 shows another panel that was developed to indicate the status of one single processing board in more detail. The overall status is displayed near the selection box at the top. Below that, each link is represented as a square that indicates the status with the corresponding colour. Right next to it, the status of typical monitoring items is displayed that can be aggregations of some sub-components. As an example, the algorithm status is the result of determining the status of all its metrics, such as the above mentioned error counters in the μ GMT or the correctness of LUTs. As a last point, any currently running command (such as configuration) is indicated and detailed information about the processor such as the location of the address table or the position in the crate are displayed as a table on the right hand side. Similar to the system overview page, collapsible tables below the graphical overview allow the user to find more information in case problems are encountered.

7 Emulator and Hardware-Tests

In this chapter, the emulation software of the μ GMT is described. Also, a short overview of the data formats for information exchange between the trigger stages is given. The chapter is closed with the description of gateway verification that was performed with the emulator.

7.1 CMSSW Level-1-Trigger Emulation

The emulation of the L1-trigger hardware is done through a collection of CMSSW modules, one per trigger stage. Each produce object-collections which are put into the event-record. These collections are then used as input for the next trigger stage.

The collections needed by a given emulator module can either be produced by the preceding emulator stage or the converted output of the actual hardware. Software is available for all components that convert between hardware readout and CMSSW collection and vice versa. This allows independent tests of the individual hardware components.

In general, the aim of the emulation is to achieve bit-by-bit correctness: Given the same inputs both hardware and emulation should produce the exact same results. In order to achieve this, the objects on which the emulation operates should have the information used in the algorithm available at the same precision as the hardware. Therefore, the properties of, for example, muons are encoded in the same integral scales as in the hardware, cf. section 5.1.

The emulation serves multiple purposes. During development it is used to verify that the gateway functions as intended and provides information for debugging. During data-taking it is used to monitor the functionality of the hardware and, for example, can be used to identify instances of mis-configurations. Moreover, the emulation can be used to implement and test new algorithms or optimisation of the existing algorithm and to test the performance without impact on the recording of physics data.

7.2 The Upgraded Global Muon Trigger Emulator

7.2.1 Implementation

Generally, the implementation of the trigger algorithm of the emulator follows what was described in chapter 5. Since the emulator operates on the aforementioned objects in

the event-record, the serialisation and de-serialisation stages are not emulated.

The data-flow is controlled in three classes, the main emulator class, the isolation unit class and the cancel unit class. To achieve bit-by-bit agreement, LUTs are used throughout the algorithm. These LUT classes have methods that allow to write text-files to disk which can be used to configure the hardware accordingly.

7.2.2 Data Formats

Three interfaces exist for the μ GMT: One to the calorimeter, one to the track-finders, and one to the μ GT. For all these interfaces, data formats are specified that define the information exchanged. As mentioned before, these data formats need to contain the quantities of the muon or the energy values in the same encoding as used in the hardware to allow bit-by-bit agreement. Therefore, all quantities are available in the integral values that correspond to the values exchanged in the hardware. Additionally, the methods are implemented that allow the conversion to the physical quantities that are useful for studying the system performance.

7.3 Evaluation of Hardware Response

Testing the gateway implementation of the algorithm is one of the important use cases of the L1-trigger emulator. For this, VHDL test bench modules were developed that can supply data to individual logic units, such as the isolation, the serialisation, the de-serialisation, the sorting, and so on. The logic is then simulated with dedicated software `Modelsim` by Mentor Graphics. The responses of the gateway can be compared with text-files that contain the response from the emulator. These modules that allow testing of parts of the algorithm are referred to as *test-benches* and the text-files that contain input or expected emulator input are called *test-patterns* in the following.

Furthermore, the functionality has to be tested when the gateway is placed into the FPGA. To allow these tests, special memory blocks are available that can be configured to supply the algorithm logic with simulated input data and read out the results. These buffers can also be used to load text-files that contain the expected results and send them via optical fibre to another processing board. With this functionality, tests can be done that can involve several stages of the trigger or one module alone. The software allows saving the contents of these buffers to disk in text-files, which are referred to in the following as *buffer-dumps*.

7.3.1 Test Sample Generation

In order to be able to produce test-patterns without being dependent on the track-finder emulation, different options were implemented that allow the generation of fake track-finder muon candidates, cf. fig. 7.1. One implementation takes generator level muons from Monte-Carlo simulation which are converted into the expected input data-format. Also, legacy track-finder data can be converted to the upgraded data-formats.

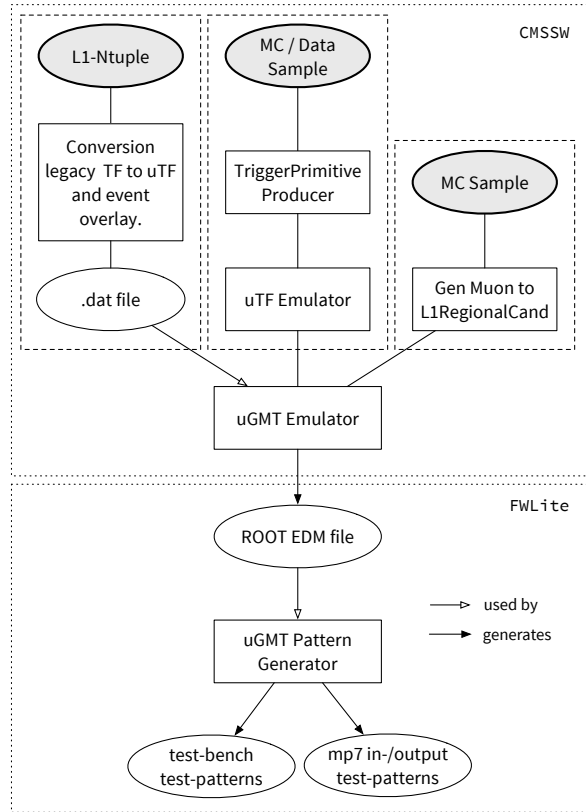


FIGURE 7.1 Work flow of test pattern generation. The patterns to verify the gateway implementation can be derived from several inputs such as Monte Carlo simulation or legacy GMT inputs from 2012 data samples. On these inputs the μ GMT emulator is run and its output is used to derive the test patterns in two different formats.

Optionally, the muon candidates from several bunch-crossings may be placed into one bunch-crossing to simulate the increased number of muons. For this “event overlay”, up to three muons are collected per 60° ϕ -wedge and track-finder η -range. Additionally, the μ GMT emulator can be run in the chain of emulation that also includes the track-finder emulators which then generate the expected input for the μ GMT.

In any case, these input data are added to the event record and the emulator is run. This produces the expected output, which in turn is placed in the event-record and saved to disk in ROOT files. In a final and fast step, these files are read to produce the test-patterns that are used in the test-bench and hardware tests. For the former, test-patterns are generated for each independent logical step of the algorithm, including the serialisation at the input to the algorithm and de-serialisation at the output.

7.3.2 Hardware Tests

Through the previously mentioned specialised buffers, it is possible to provide patterns to the algorithm gateway when it is placed in the FPGA. This allows to test the trigger algorithm in the same configuration as it will be run during physics data taking. While the test-benches allow testing the basic implementation of the algorithm in HDL, they do not allow to check whether the timing of the electronic signals is correct. Therefore,




FIGURE 7.2 Example of emulator and hardware comparison. Each bit of the 64 bit words that represent the muons sent to the μ GT is compared and in case of a mismatch highlighted. In this example, errors in bit 32 were found for muons that were ranked as first and second by the μ GMT sorting algorithm.

this additional test in the actual chip is important to verify the functionality of the algorithm.

For this test, the hardware algorithm is provided with the same input patterns as the emulator was. The test patterns are run through the algorithm gateway and the muons after the first sorting stage, the selected final muons and the assigned ranks are read out from the hardware. All of these are then compared and checked to be identical between hardware and emulator, identical meaning that the integral values of all previously described properties match between the two sets.

Comparing the two samples amounts to checking whether the reported 64 bit words are identical for emulation and hardware. However, it can be useful to have a graphical representation of this information. Fig. 7.2 shows such a representation for the final muons. The bins in the histogram indicate whether a mismatch in a given bit was found, in this example errors were found in bit 33 of the first two muons.

7.3.3 Results

The hardware results were checked against the results of the emulator at this bit-by-bit level for all main parts of the algorithm. At the time of writing, the exact nature of the track-address based cancellation algorithm was still under discussion. Therefore, this part of the algorithm was not possible to be tested. Testing was performed on several Monte Carlo samples, including an extended test in which 10000 events are run and no error is encountered. Samples are selected that have muons in the final state, such as top anti-top production, Z -boson production and W -boson production. Additionally, to stress-test the sorting algorithm, muon candidates from several events were merged into one to simulate a full occupancy of the μ GMT inputs.

The implementation of the isolation algorithm using the energy ratio discussed in section 4.5 was still in development and was not yet tested. The default isolation algorithm with absolute and relative isolation and using the extrapolation was tested and found to

agree between emulation and hardware implementation. Lastly, the sorting algorithm and coordinate based cancellation were tested and were also found to agree in all events.

8 Conclusion

With the upgrade of the LHC to reach the design collision energy and higher luminosities, changes to the triggering scheme of the CMS experiment have become necessary. To achieve a similar event rate suppression while keeping signal events with high efficiency, the hardware-based first level trigger is being upgraded.

The three muon detectors (DT, CSC and RPC) of CMS provide information to the level-1 muon trigger. An essential part of this muon trigger is the Global Muon Trigger which combines the information from the three muon track-finder systems of CMS.

This thesis presented the new approach in the upgraded muon trigger, in which the information of the muon systems is combined as early as possible in the track-finders. Another new feature of the upgraded trigger is the possibility of information exchange between the calorimeter and muon trigger systems allowing the introduction of isolation algorithms in the μ GMT. The concept of the Micro-Global Muon Trigger (μ GMT) is therefore very different from its legacy counterpart. The main task is to select the 8 best muon candidates from the up to 108 input muon candidates and assign the isolation.

A crucial part of successful physics data taking is the ability to ensure the correct functioning of the trigger. Monitoring tools were designed and implemented that are used throughout the upgraded level-1 trigger. Additionally, the necessary configuration software for the upgraded Global Muon Trigger was developed. In particular, it allows to efficiently change parts of the algorithm that use look-up tables.

Throughout the level-1 trigger project, emulators are used to test new variations of the existing algorithms without an impact on data taking. The emulators are essential for the trigger operation and commissioning. The μ GMT emulator, implemented in the CMS software framework, was used for tests and verification during the development of the gateway implementation. Furthermore, it was used to estimate the expected performance of the trigger as presented in this thesis. Ways of suppressing the level-1 accept rate by doing a selection based on the quality of input muons was discussed and found to reduce the rate by up to a factor of 3 in some detector regions with a loss of efficiency of a few percent.

A muon isolation algorithm within the level-1 context was presented. It allows to reject muons produced as subsequent states in a hadronic decay. The calorimeter trigger data is used to calculate a ratio of the energy deposited by the muon itself and any deposits around it. The available granularity of calorimeter energy information requires to restrict usage of this isolation variable to $|\eta| < 1.8$. Applying a threshold on the ratio yields a reduction of the event rate by a factor of two in the aforementioned η region. At the

same time, the efficiency loss for identification of muons from Z boson decays can be kept within a few percent. Challenges in the implementation of this algorithm in the level-1 hardware trigger, such as bandwidth limitations, were discussed and solutions presented.

During 2015, commissioning of the μ GMT and the other components of the trigger upgrade started. The signals from the detector electronics were split to be received both by legacy and upgraded trigger. This allowed to start tests of the trigger with measurements of muons passing through CMS.

The commissioning will be finished in the beginning of 2016 and the upgrade will very likely be used as the default trigger system in the remainder of run-2. In the beginning of data taking in 2016 the μ GMT will run without the final isolation algorithm. The decision was taken that the priority for initial commissioning will be the sorting and cancellation algorithms of the μ GMT. However, the necessary connection between calorimeter trigger and μ GMT was installed and tested. This allows the algorithm to be introduced as soon as the implementation of the isolation algorithm in gateway is finalised.

Bibliography

- [1] David J Griffiths. *Introduction to elementary particles; 2nd rev. version*. Physics textbook. Wiley, New York, NY, 2008.
- [2] Francis Halzen and Alan Douglas Martin. *Quarks and leptons: an introductory course in modern particle physics*. Wiley, New York, NY, 1984.
- [3] Norbert Schmitz. *Neutrino physics*. Teubner, Stuttgart, 1997.
- [4] K.A. Olive et al. (Particle Data Group). 2014 Review of Particle Physics. *Chin. Phys. C*, 38(090001 (2014)), Jan 2014.
- [5] CMS Collaboration. Observation of a new boson at a mass of 125 GeV with the CMS experiment at the LHC. *Phys. Lett. B*, 716:30–61, Jul 2012. doi: 10.1016/j.physletb.2012.08.021.
- [6] ATLAS Collaboration. Observation of a new particle in the search for the Standard Model Higgs boson with the ATLAS detector at the LHC. *Phys. Lett. B*, 716:1–29, 2012. doi: 10.1016/j.physletb.2012.08.020.
- [7] P. Nath, B.D. Nelson, H. Davoudiasl, B. Dutta, D. Feldman, et al. The Hunt for New Physics at the Large Hadron Collider. *Nucl. Phys. Proc.* doi: 10.1016/j.nuclphysbps.2010.03.001.
- [8] P. W. J. Stirling. private communication. URL <http://www.hep.ph.ic.ac.uk/~wstirling/plots/plots.html>.
- [9] Andy Buckley, Jonathan Butterworth, Stefan Gieseke, David Grellscheid, and Stefan Hoche. General-purpose event generators for LHC physics. *Phys. Rept.*, 504:145–233, 2011. doi: 10.1016/j.physrep.2011.03.005.
- [10] O. S. Brüning, P. Collier, P. Lebrun, S. Myers, R. Ostojic, et al. *LHC Design Report Volume 1: The Main Ring*. CERN, Geneva, 2004.
- [11] J. Haffner. The CERN accelerator complex. Complexe des accélérateurs du CERN. Oct 2013.
- [12] S. Maury. The Antiproton Decelerator: AD. *Hyperfinde Interactions*, 109:43–52, Aug 1997.
- [13] M. Hori and J. Walz. Physics at CERN’s Antiproton Decelerator. *Prog. Part. Nucl. Phys.*, 72:206–253, 2013. doi: 10.1016/j.ppnp.2013.02.004.

-
- [14] *International Linear Collider Technical Review Committee: Second Report*. SLAC, Feb 2003.
- [15] M. Lindroos and T. Nilsson. *HIE-ISOLDE : the technical options*. CERN, Geneva, 2006.
- [16] H. Gaillard, C. Theis, H. Vincke, I. Efthymiopoulos, S. Evrard, et al. HiRadMat: A New Irradiation Facility for Material Testing at CERN. *Conf. Proc.*, C110904: 1665–1667, 2011.
- [17] ATLAS Collaboration. The ATLAS Experiment at the CERN Large Hadron Collider. *J. Instrum.*, 3:S08003. 437 p, Aug 2008.
- [18] TOTEM Collaboration. The TOTEM Experiment at the CERN Large Hadron Collider. *J. Instrum.*, 3:S08007, Aug 2008.
- [19] LHCf Collaboration. The LHCf detector at the CERN Large Hadron Collider. *J. Instrum.*, 3:S08006, Aug 2008.
- [20] ALICE Collaboration. The ALICE experiment at the CERN LHC. A Large Ion Collider Experiment. *J. Instrum.*, 3:S08002. 259 p, Aug 2008.
- [21] LHCb Collaboration. The LHCb Detector at the LHC. *J. Instrum.*, 3:S08005, Aug 2008.
- [22] R Alemany-Fernandez, E Bravin, L Drosdal, A Gorzawski, V Kain, et al. Operation and Configuration of the LHC in Run 1. Nov 2013.
- [23] W. Herr. Luminosity performance reach after LS1. *Charmony 2012 Workshop on LHC Performance*, Feb 2012. doi: 10.5170/CERN-2012-006.276.
- [24] CMS Collaboration. *CMS Physics: Technical Design Report Volume 1: Detector Performance and Software*. Technical Design Report CMS. CERN, Geneva, 2006.
- [25] CMS Collaboration. The CMS Detector at the CERN LHC. *J. Instrum.*, 3:S08004. 361 p, Aug 2008.
- [26] CMS Collaboration. *The CMS muon project: Technical Design Report*. Technical Design Report CMS. CERN, Geneva, 1997.
- [27] CMS Collaboration. The performance of the CMS muon detector in proton-proton collisions at $\sqrt{s} = 7$ TeV at the LHC. *JINST*, 8:P11002, 2013. doi: 10.1088/1748-0221/8/11/P11002.
- [28] CMS Collaboration. *Technical proposal for the upgrade of the CMS detector through 2020*. Number CERN-LHCC-2011-006. LHCC-P-004 in Technical Design Report CMS. CERN, Geneva, Jun 2011.
- [29] PCI Industrial Computer Manufacturers Group. Micro telecommunications computing architecture base specification: micro TCA. *PICMG*, 2006.
- [30] M. Tytgat, A. Marinov, P. Verwilligen, N. Zaganidis, A. Aleksandrov, et al. The upgrade of the CMS RPC system during the first LHC long shutdown. *JINST*, 8(02):T02002, 2013.

- [31] CMS Collaboration. CSC local efficiency measurements from run II data. Oct 2015.
- [32] F Rademakers and R Brun. ROOT – An object oriented data analysis framework. 389:81–86, Apr 1997. doi: 10.1016/S0168-9002(97)00048-X.
- [33] S. Agostinelli et al. GEANT4: A Simulation toolkit. *Nucl.Instrum.Meth.*, A506: 250–303, 2003. doi: 10.1016/S0168-9002(03)01368-8.
- [34] F. Cossutti. Tuning and optimization of the CMS simulation software. *J.Phys.Conf.Ser.*, 219, 2010. doi: 10.1088/1742-6596/219/3/032005.
- [35] G L Bayatyan, N Grigorian, V G Khachatryan, A T Margarian, A M Sirunyan, et al. *CMS TriDAS project - Technical Design Report, Volume 1, The Trigger Systems*. Number CERN-LHCC-2000-038 in Technical Design Report CMS. CERN-LHCC, Geneva, Dec 2000.
- [36] CMS Collaboration. *CMS Technical Design Report for the Level-1 Trigger Upgrade*. Number CERN-LHCC-2013-011 in Technical Design Report CMS. Geneva, Jun 2013.
- [37] CMS Collaboration. The CMS Trigger. *JINST*, currently in preparation.
- [38] H. Sakulin. *Design and Simulation of the First Level Global Muon Trigger for the CMS Experiment at CERN*. PhD thesis, Nov 2002.
- [39] D. Acosta, P. Bortignon, K. Bunkowski, M. Carver, Janos Erö, et al. private communication and level-1 trigger group meetings.
- [40] K Compton, S Dasu, A Farmahini-Farahani, S Fayer, R Fobes, et al. The MP7 and CTP-6: multi-hundred Gbps processing boards for calorimeter trigger upgrades at CMS. *Journal of Instrumentation*, 7(12):C12024, 2012.
- [41] D Acosta, G Brown, A Carnes, M Carver, D Curry, et al. The CMS Modular Track Finder boards, MTF6 and MTF7. *Journal of Instrumentation*, 8(12):C12034, 2013.
- [42] J. Alwall, R. Frederix, S. Frixione, V. Hirschi, F. Maltoni, et al. The automated computation of tree-level and next-to-leading order differential cross sections, and their matching to parton shower simulations. *JHEP*, 07:079, 2014. doi: 10.1007/JHEP07(2014)079.
- [43] Torbjörn Sjöstrand and Stephen Mrenna and Peter Skands. PYTHIA 6.4 physics and manual. *Journal of High Energy Physics*, 2006(05):026, 2006.
- [44] I. M. de Abril and M. M. de Abril. Enhancing the User Interface of the CMS Level 1 Trigger Online Software with AJAX. *IEEE-NPSS*, 15, 2007. doi: 10.1109/RTC.2007.4382776.
- [45] Serguei Chatrchyan et al. Performance of CMS Muon Reconstruction in pp Collision Events at $\sqrt{s} = 7$ TeV. *JINST*, 7:P10002, 2012. doi: 10.1088/1748-0221/7/10/P10002.
- [46] C. Ghabrous Larrea, J. Brooke, K. Bunkowski, A. Thea, I. A. Carli, and C. Lazaridis. SWATC: common control SW for the uTCA-based upgraded CMS L1 Trigger. *JPCS*, currently in preparation.

- [47] C. Ghabrous Larrea, K. Harder, D. Newbold, D. Sankey, A. Rose, A. Thea, and T. Williams. IPbus: a flexible Ethernet-based control system for xTCA hardware. *JINST*, 10:C02019, 2015.

Abbreviations

μGMT	Micro-Global Muon Trigger
μGT	Micro-Global Trigger
μTCA	Micro Telecommunication Computing Architecture
AMC-13	Advanced Mezzanine Card-13
BMTF	Barrel Muon Track-Finder
CERN	Conseil Européen pour la Recherche Nucléaire
CMS	Compact Muon Solenoid
CPPF	Concentrator Pre-Processor Fan-out
CSC	Cathode Strip Chamber
CSCTF	Cathode Strip Chamber Track-Finder
CTP7	Calorimeter Trigger Processor 7
DAQ	Data AcQuisition system
DT	Drift Tube
DTTF	Drift Tube Track-Finder
ECAL	Electromagnetic Calorimeter
EMTF	End-cap Muon Track-Finder
FPGA	Field Programmable Gate-Array
GCT	Global Calorimeter Trigger
GMT	Global Muon Trigger
GT	Global Trigger
HCAL	Hadron Calorimeter
HDL	Hardware Description Language
HLT	High Level Trigger

HO	Outer Hadron calorimeter
L1-trigger	Level 1-Trigger
LHC	Large Hadron Collider
LS-1	Long Shut-down 1
LUT	Look-Up Table
MB	Muon Barrel
ME	Muon End-cap
MP7	Master Processor 7
MTF7	Muon Track Finder 7
OMTF	Overlap Muon Track-Finder
PACT	PAtern Comparator Trigger
RCT	Regional Calorimeter Trigger
RPC	Resistive Plate Chamber
run-1	first running period of the LHC
run-2	second running period of the LHC
SWATCH	SoftWare for Automating conTrol of Common Hardware
TCDS	Trigger Control and Distribution System

List of Figures

2.1	Cross sections σ for various physics processes at hadron colliders as a function of the collision energy [8].	6
2.2	Exemplary Feynman diagrams of processes that have muons in the final state. On the left and right, W and Z boson production with subsequent muonic decay. In the middle the decay of a B meson in the spectator view.	7
2.3	The LHC accelerator complex, shown are also the experimental areas provided with proton or ion beams along the accelerator chain [11].	8
2.4	The LHC performance during run-1. On the left the peak instantaneous luminosity (and rate of collisions) is shown for running in 2010 (green, multiplied by 20), 2011 (blue) and 2012 (red). The right hand figure shows integrated luminosity (and number of collisions) for the 3 years in the same colour coding [22].	10
2.5	True to scale perspective view of the CMS detector with the interaction point in the centre and all individual detector components labelled [24].	12
2.6	An r - z cross section of the inner detector which consists of the Pixel, the Tracker Inner Barrel (TIB), the Tracker Outer Barrel (TOB), the Tracker Inner Disk (TID) and the Tracker End-Cap (TEC) detectors [25].	12
2.7	A perspective illustration of the Electromagnetic Calorimeter, showing the supermodules and modules composing the barrel, the Dee holding the crystals in the end-caps and the pre-shower devices placed just in front of the end-caps [25].	14
2.8	A r - z cross section of one quadrant of the Hadron Calorimeter (HCAL), highlighted are the Barrel HCAL (HB) and End-cap HCAL (HE) in blue closest to the interaction point, Outer HCAL (HO) in green at the top and the Forward HCAL (HF) on the right hand side [25].	15
2.9	An r - z cross section of one quarter of the muon system with the z -axis pointing left to right and r increasing bottom to top, the interaction point is at the bottom left. The shown detectors correspond to the situation during run-1 of the LHC. Shown are the drift tube (DT) chambers in the four barrel stations (MB1-4) in orange, the cathode strip chambers (CSCs) in the four end-cap disks (ME1-4) in green and the complementing Resistive Plate Chambers (RPCs), installed both in barrel and end-cap, indicated in blue [27].	16
2.10	Layout of the Drift Tube (DT) chambers in the CMS barrel region, showing one of the 5 wheels. The numbers label the twelve sectors of the DT system. The bold lines indicate the RPC chambers. [26].	17
2.11	The CMS Data Acquisition system consisting of level-1 trigger, detector front-ends, the builder network, event manager, monitoring, and the computing services running the high level trigger. On the left hand side the average event rate at each selection stage is shown [24].	18

- 2.12 An r - z cross section of one quarter of the muon system with the z -axis pointing left to right and r increasing bottom to top, the interaction point is at the bottom left. The shown detectors correspond to the situation during run-2 of the LHC. Shown are the drift tube (DT) chambers in the four barrel stations (MB1-4) in orange, the cathode strip chambers (CSCs) in the four end-cap disks (ME1-4) in green and the complementing Resistive Plate Chambers (RPCs), installed both in barrel and end-cap, indicated in blue [31]. 20
- 3.1 An r - z projection of one quadrant of CMS, showing the tracker, barrel ECAL (EB), end-cap ECAL (EE), barrel HCAL (HB), and the end-cap HCAL (HE). The green dashed lines and blue numbers indicate the trigger towers [35]. 24
- 3.2 Illustration of the calorimeter trigger data flow: In the legacy trigger (a) energies from the calorimeters (HCAL, ECAL and HF) are pre-processed by the Regional Calorimeter Trigger (RCT), sorted in the Global Calorimeter Trigger (GCT) and sent to the Global Trigger. In the upgraded trigger (b) new processing boards operate in two layers, with pre-processing being done in the first layer and trigger primitive reconstruction in the second, the additional de-multiplexer unit (DeMux) prepares the data for the upgraded Micro-Global Trigger. 25
- 3.3 Legacy and upgraded muon trigger chain: In the legacy, Cathode Strip Chamber (CSC), Drift Tube (DT) and Resistive Plate Chamber (RPC) information is available, while in the upgrade also Outer Hadron calorimeter (HO) hit information can be processed. The hit information of the Cathode Strip Chambers is pre-processed in the Muon Port Card (MPC). In the legacy the muon candidates are reconstructed in components dedicated to each detector, the Cathode Strip Chamber Track-Finder (CSCTF), Drift Tube Track-Finder (DTTF) (and Sector Collector (SC)) and PAttern Comparator Trigger (PACT) (and Link Boards (LB)). The candidates are first sorted locally (per sub-detector) and then merged and globally sorted in the Global Muon Trigger (GMT). In the upgrade the candidates are reconstructed track-finders that are defined by η acceptances of the barrel (Barrel Muon Track-Finder (BMTF)), end-cap (End-cap Muon Track-Finder (EMTF)) and the overlap of the two (Overlap Muon Track-Finder (OMTF)). Information of Outer Hadron calorimeter, Drift Tube and barrel Resistive Plate Chamber is combined in the TwinMux and provided to the BMTF. The Concentrator Pre-Processor Fan-out receives end-cap Resistive Plate Chamber data and sends it both to EMTF and OMTF. 26
- 4.1 Distributions of p_T (left) and η (right) of generated (red) and level-1 muons (black) for the minimum bias (dashed) and $Z \rightarrow \mu\mu$ (solid). 34
- 4.2 Rate as a function of p_T threshold. The dashed line represents the legacy GMT rate in the corresponding region. Dots show the estimation for the track-finders: Barrel in (a), overlap in (b) and end-cap in (c). 37
- 4.3 Estimated rate as a function of η with an applied p_T threshold of 16 GeV. The dashed line represents the legacy GMT rate. The stack shows the rate contributions from the track-finders. In red the barrel, green the overlap and blue the end-cap track-finder. 38

- 4.4 Efficiency of the track-finders as a function of p_T (a), η (b), and ϕ (c). All plots show the efficiency at a threshold of 16 GeV. The dashed line shows the efficiency of the legacy system and the coloured dots the upgraded track-finders efficiency. The BMTF in red, the OMTF in green, and the EMTF in blue. 39
- 4.5 The efficiency (a) and rate (b) in the BMTF split into the qualities reported by the track-finder. The quality scale in the BMTF ranges currently from 1 to 7, while the μ GMT accepts values between 0 and 15. 40
- 4.6 Efficiency (a) and rate (b) contributions of the different qualities reported for muon candidates in the OMTF. The quality code is the number of hit layers used for a given candidate. 41
- 4.7 Efficiency (a) and rate (b) contributions split in quality of muon candidates in the EMTF. The quality code corresponds to the stations that were used in the track reconstruction. 42
- 4.8 Efficiency of the μ GMT as a function of η . Without quality selection (a) and selection as described in the text (b) applied. In both cases a threshold of 16 GeV is applied. 43
- 4.9 Rate of the μ GMT as a function of η . Without quality selection (a) and selection as described in the text (b) applied. In both cases a threshold of 16 GeV is applied. 43
- 4.10 Efficiency as a function of η , comparing legacy GMT (dashed line) with μ GMT (stack). The μ GMT is stacked in contributions of the individual track-finders (BMTF red, OMTF green, and EMTF blue). In (a) comparison without quality selection and in (b) with quality selection as described in the text. 45
- 4.11 Efficiency as a function of ϕ , comparing legacy GMT (dashed line) with μ GMT (stack). The μ GMT is stacked in contributions of the individual track-finders (BMTF red, OMTF green, and EMTF blue). In (a) comparison without quality selection and in (b) with quality selection as described in the text. 45
- 4.12 Turn-on distributions of the legacy GMT (dashed line) and the μ GMT (stack). The μ GMT is stacked in contributions of the individual track-finders (BMTF red, OMTF green, and EMTF blue). In (a) comparison without quality selection and in (b) with quality selection as described in the text. 46
- 4.13 Rate as a function of p_T threshold for the legacy GMT (dashed line) and the μ GMT (stack). The μ GMT is stacked in contributions of the individual track-finders (BMTF red, OMTF green, and EMTF blue). In (a) comparison without quality selection and in (b) with quality selection as described in the text. 46
- 4.14 Ratio of level-1 E_T index to E_T calculated in off-line reconstruction as a function of the tower η index. 47
- 4.15 The result of the tower calibration. On the left, the reconstructed tower energy as a function of the level-1 E_T index can be seen. On the right the corresponding distribution of reconstructed tower energy for the E_T index of 9 is shown. 48
- 4.16 Extrapolation to the vertex. The differently coloured lines show the fitted functions used for the extrapolation. 48
- 4.17 Average tower energy around a muon in case of prompt muons (a) and muons found by level-1 in minimum bias (b). On the x- and y-axis, the tower index relative to the non-extrapolated muon position is used. 49

4.18	Average tower energy around a muon in case of prompt muons (a) and muons found by level-1 in minimum bias (b). On the x- and y-axis, the tower index relative to the extrapolated muon position is used.	50
4.19	Sum of tower energies in cones around the muon, shown are the distributions for muons from Z boson decays (solid line) and the muons found by the level-1 trigger in the minimum bias sample (dashed line). All towers within a 11-by-11 square (a) E_T^{tot} , the towers immediately around the muon (b) E_T^{in} and all towers except those immediately around the muon (c) E_T^{out} . Here, the extrapolation is not applied. The last bin represents the overflow.	51
4.20	Sum of tower energies in cones around the extrapolated muon, shown are the distributions for muons from Z boson decays (solid line) and the muons found by the level-1 trigger in the minimum bias sample (dashed line). All towers within a 11-by-11 square (a) E_T^{tot} , the towers immediately around the muon (b) E_T^{in} and all towers except those immediately around the muon (c) E_T^{out} . The last bin represents the overflow.	52
4.21	Sum of tower energies, E_T^{out} divided by the level-1 muon p_T in fig. 4.21(a) and divided by the total sum of energy E_T^{tot} in fig. 4.21(b). The dashed line for muons found in the minimum bias sample and the solid line for muons from Z boson decay. In fig. 4.21(a) the last bin serves as an overflow bin.	53
4.22	Relative efficiency as a function of relative rate for different isolation thresholds applied for the absolute isolation variables. The different colours show the influence of changing the p_T threshold.	54
4.23	Relative efficiency as a function of relative rate for different isolation thresholds applied to the relative isolation variable. The different colours show the influence of changing the p_T threshold.	54
4.24	Efficiency of μ GMT with (points) and without (dashed line) relative isolation applied. The lower halves show the ratio of the two. For figs. (b) to (c) a isolation threshold resulting in 90% overall relative efficiency was applied, while in fig. (d) a 95% relative efficiency working point was chosen.	55
4.25	Rate of μ GMT with (points) and without (dashed line) relative isolation applied. The lower halves show the ratio of the two. A threshold was chosen that results in 90% overall relative efficiency.	56
4.26	Comparing the efficiency loss as a function of p_T for increasing the p_T threshold by 2 GeV (in red) or applying a threshold with 10% efficiency loss to the relative isolation.	56
4.27	Relative efficiency as a function of relative rate for different thresholds applied to the inner cone energy. The different colours show the influence of changing the p_T threshold.	57
4.28	Efficiency of μ GMT with (points) and without (dashed line) relative isolation applied. A 85% relative efficiency working point was chosen to determine the threshold.	58
4.29	Rate of μ GMT with (points) and without (dashed line) applying a threshold to the inner cone energy. The lower plot in each figure shows the ratio of the two. A threshold was chosen that results in 85% overall relative efficiency.	59
4.30	Relative efficiency as a function of relative rate for different thresholds applied to the ratio of outer cone energy to the full cone. The different colours show the influence of changing the p_T threshold.	59

- 4.31 Efficiency of the μ GMT with applying a threshold to the energy ratio as described in the text, shown with points. The dashed line indicates the efficiency without any isolation criterion. Apart from the isolation also a p_T threshold of 16 GeV is applied. 60
- 4.32 Rate of the μ GMT without (dashed line) and with (points) applied threshold on the ratio of outer and full cone energy. Additionally a L1 p_T threshold of 16 GeV is applied. 61
- 4.33 Relative efficiency as a function of relative rate for different thresholds applied to the energy ratio of outer and full cone. The threshold is only applied for muons within $|\eta| < 1.8$ to retain efficiency. The different colours show the influence of changing the p_T threshold 61
- 4.34 Efficiency when restricting the application of a threshold on the energy ratio to $|\eta| < 1.8$. Both efficiency with (points) and without (dashed line) applied threshold are shown. 62
- 4.35 Rate of the μ GMT with η -restricted (only for muons with $|\eta| < 1.8$) energy ratio threshold. The rate is shown both with (points) and without (dashed line) applied threshold. 63
- 5.1 This block-diagram shows the data flow within the μ GMT algorithm. The data flows left-to-right, the different track-finder muons are indicated with red, green and blue arrows, representing barrel, overlap and end-cap track-finders, respectively. Calorimeter data are shown with magenta arrows. The numbers above the sorting stages indicate the number of muon candidates at start and end of the sort stage. Each sorting stage 1-block is a simplification as indicated in the inset box at the bottom right. 68
- 6.1 Demonstration of the system overview user interface developed for the central online software. For this screen-shot, the cell was connected to a demonstrator system with more than one board and crate. 74
- 6.2 Demonstration of the object-detail user interface developed for the central online software. 75
- 7.1 Work flow of test pattern generation. The patterns to verify the gateway implementation can be derived from several inputs such as Monte Carlo simulation or legacy GMT inputs from 2012 data samples. On these inputs the μ GMT emulator is run and its output is used to derive the test patterns in two different formats. 79
- 7.2 Example of emulator and hardware comparison. Each bit of the 64 bit words that represent the muons sent to the μ GT is compared and in case of a mismatch highlighted. In this example, errors in bit 32 were found for muons that were ranked as first and second by the μ GMT sorting algorithm. 80

