

Université Libre de Bruxelles – Faculté des Sciences
IIHE – Interuniversity Institute for High Energies (ULB – VUB)

**Development of the new trigger and data acquisition
system for the CMS forward muon spectrometer
upgrade**

Erik Verhagen

*A thesis submitted in partial fulfillment of the requirements for the degree
of Doctor of Science*

Supervisor:

Gilles De Lentdecker (ULB)

Jury:

Kael Hanson (Président, ULB)

Juanan Aguilar (Secrétaire, ULB)

Frédéric Robert (ULB)

Michael Tytgat (Universiteit Gent)

Nick Van Remortel (Universiteit Antwerpen)

*« Soumettre un nouveau à des épreuves
d'initiation, avant de l'admettre au sein
d'une société scolaire ou universitaire »*

Définition du mot *bizutage*, Larousse.

Acknowledgements

First, I would like to thank Gilles, Kael and Daniel Bertrand for making possible this four-years work contract as an electronics engineer at the Inter-university Institute for High Energies in Brussels. Thanks also to the warm welcome of Yifan in the newly created R&D group and the kindness of Danielle who helped me with the heavy administrative matters of the beginning. These key people made my initial move to Belgium in 2010 a good idea for the future.

Secondly, a number of co-students deserve special thanks too, especially Thomas L. and Alexandre, who helped until the last days with clever inputs to my work and friendly comments. The rest of the R&D team belongs also in this list, Thierry and Florian who sometimes succeeded to convince me with obscure physics explanations, and Patrizia, Valérie and Ryo for sharing their opinions during our informal group meetings.

Thanks also to the people of the *Friday beer* after-works, who helped to keep my spirit high, Thomas R. and Ludivine, Laurent T., Aidan, Thomas M. and David, not forgetting Tomislav to whom I wish especially all the best.

Last but not least, I will not forget the lovely lady who made my stay in Brussels enjoyable for the last four years, and who always believed in me, even when the mood was low. I can now tell that the strongest point I acquired during this work is perseverance, and this is thanks to her.

Revision History

Revision Date	Changes
1.0 23/01/2015	Initial submission for private defense
<i>(partial)</i> 13/02/2015	<ul style="list-style-type: none"> – Old Chapter 3.3 « Front-end signal processing » moved to Chapter 4.3 « Front-end electronics ». This chapter describes a number of preliminary studies performed in answer to the specifications set by CMS on the front-end electronics. – New Chapter 3.3 on the physics performance of the proposed GE1/1 system describing : <ul style="list-style-type: none"> – the goals (the physics behind) – the muon trigger performance – the muon reconstruction performance
<i>(partial)</i> 18/02/2015	<ul style="list-style-type: none"> – New section 6.2.3 summarizing the results of the setup in the beam tests conducted at CERN last December – Added list of acronyms at the end
<i>(partial)</i> 23/02/2015	<ul style="list-style-type: none"> – Corrected typos and some incorrect English construction in Revision 1.0. Many thanks to Kael for spotting and reporting most of them. – Broadened most local end-of-chapter conclusions to give a more consistent link between each individual chapters – Added/replaced some figures in Chapter Four with more recent and explicit ones. – Added a last paragraph in the introduction listing the contribution of my very self in this thesis, to specify where “we” should be read as “me”. – Applied some cosmetics (tale-style images in chapter intros) and added acknowledgements
2.0 26/02/2015	Final submission for public defense

Contents

INTRODUCTION	1
CHAPTER 1: The LHC and CMS: The collider and the experiment	3
1.1 Cultural background in particle physics	4
1.1.1 Short contextual history	4
1.1.2 Introduction to particle accelerators	4
1.1.3 Collider experiments	6
1.2 The LHC	8
1.2.1 Context	8
1.2.2 Key parameters of the LHC	9
1.2.3 Upgrade plan	10
1.3 The Compact Muon Solenoid (CMS) detector	12
1.3.1 Facts and figures	12
1.3.2 Barrel	13
1.3.2.1 Silicon tracker.....	13
1.3.2.2 Electromagnetic Calorimeter (ECAL).....	15
1.3.2.3 Hadron Calorimeter (HCAL).....	17
1.3.2.4 Superconducting solenoid.....	18
1.3.2.5 Muon chambers.....	19
1.3.3 End caps	20
1.3.3.1 Muon chambers.....	20
1.3.3.2 Forward Calorimeter.....	21
1.3.4 Upgrades	22
1.4 Conclusion	23
CHAPTER 2: The CMS trigger and data acquisition system	25
2.1 Basic DAQ design features	26
2.2 Overview of the CMS trigger system	28
2.2.1 Level-1 trigger (LV1)	28
2.2.2 High Level Trigger (HLT)	30
2.2.3 Muon system L1 trigger	30
2.3 The CMS Data Acquisition System architecture	33
2.3.1 Event readout interface	33
2.3.2 Event builder and filter	34
2.3.3 Control and monitoring	35

2.4	Hardware implementation	37
2.4.1	Front-end electronics	37
2.4.2	ATCA standard	38
2.4.3	μ TCA architecture	40
2.5	Conclusion	43
CHAPTER 3: GEM detectors and the GE1/1 system		45
3.1	Muon system technologies comparison in CMS	46
3.1.1	Drift Tubes	46
3.1.2	Cathode Strip Chambers	46
3.1.3	Resistive Plate Chambers	48
3.2	The GEM detector technology	50
3.2.1	The GEM foil	50
3.2.2	Triple-GEM detectors	52
3.2.3	The GE1/1 station	55
3.3	Physics performance	59
3.3.1	Goals	59
3.3.2	Muon trigger performance	60
3.3.3	Muon reconstruction performance	63
3.4	Conclusion	66
CHAPTER 4: GE1/1 read-out system architecture		67
4.1	Constraints and requirements	68
4.1.1	Genesis of the project	68
4.1.2	Performance requirements	69
4.1.3	Mechanical considerations	70
4.1.4	Integration in the CSC trigger logic	72
4.1.5	Technical limitations	73
4.2	System overview for LS2	74
4.3	Front-end electronics	76
4.3.1	Specifications	76
4.3.2	gDSP option	77
4.3.3	VFAT3 option	79
4.3.4	Common mode reduction algorithm	81
4.3.5	Time resolution improvement	85
4.3.6	The GEM Electronic Board (GEB)	87
4.3.7	The opto-hybrid	88
4.4	Back-end electronics	90
4.4.1	μ TCA based architecture	90
4.4.2	The MP7 processing boards	91
4.4.3	The 13th Advanced Mezzanine Card (AMC13)	92
4.5	System description for slice test	94
4.5.1	Motivations	94
4.5.2	Front-end electronics	95

4.5.3	Off-detector electronics	96
4.6	Conclusion	97
CHAPTER 5: μTCA developments		99
5.1	μ TCA infrastructure	100
5.1.1	The μ TCA form factor	100
5.1.2	Module Management Controller (MMC)	100
5.1.3	e-Keying and fabric interface	101
5.1.4	Software integration	101
5.2	Module Management Controller (MMC) testbed	103
5.2.1	Motivation	103
5.2.2	Overview	103
5.2.3	AMC edge connector	105
5.2.3.1	Physical characteristics.....	105
5.2.3.2	Power rails and presence signal.....	105
5.2.3.3	AMC enable signal.....	106
5.2.3.4	Geographical addressing and IPMI.....	106
5.2.3.5	Fabric interfaces.....	107
5.2.4	Power	108
5.2.4.1	Input rails.....	108
5.2.4.2	DC/DC converters.....	108
5.2.4.3	Additional 3.3 Volt payload power.....	111
5.2.4.4	Low Drop-Out (LDO) regulators.....	112
5.2.5	Module Management Controller (MMC)	112
5.2.5.1	Description.....	112
5.2.5.2	Micro-controller.....	112
5.2.5.3	AMC.0 compliant signals.....	113
5.2.5.4	AMC.0 compliant EEPROM.....	114
5.2.5.5	Additional FLASH memory.....	114
5.2.5.6	JTAG scheme.....	114
5.2.5.7	User FPGA reconfiguration.....	115
5.2.6	Monitoring and on-board network	115
5.2.7	Operation and debug interface	116
5.2.8	Fabric interface mezzanine	117
5.3	Remote firmware upgrade	118
5.3.1	Motivation	118
5.3.2	Firmware	118
5.3.3	Software	121
5.4	Conclusion	122
CHAPTER 6: Cosmic test-bench		123
6.1	Experimental setup	124
6.1.1	Triple-GEM prototype	124
6.1.2	Trigger system	126
6.1.3	DAQ electronics	127
6.1.3.1	System overview.....	127

6.1.3.2 Front-end electronics	128
6.1.3.3 Read-out electronics GLIB (μ TCA).....	130
6.1.3.4 DAQ software and run control	131
6.2 Results with the μ TCA based read-out system	133
6.2.1 Performance of the DAQ electronics	133
6.2.2 Muon detection	134
6.2.3 Beam test results	136
CONCLUSIONS	143
Bibliography	145
Table of Acronyms	149

INTRODUCTION

In experimental research, the most visible part of a scientific project is, by far, its scientific production. More specifically, the particularity of experimental particle physics, is the size of the projects nowadays chosen to broaden the scientific activity as much as possible in a field where the number of open questions remains predominant. In this sense, the particle physics community requires constantly bigger, more complex and higher in precision experimental facilities, making more room for non-scientific side activities, such as solving engineering challenges. This work was achieved in the context of such an experiment.

The general idea behind this work is to study a proof of concept, in the field of instrumentation and detector developments for a large scale particle physics experiment. The aim is thus to study the feasibility and evaluate the implications of installing a novel type of detector as a future upgrade of this experiment. All along this document, the emphasis is put on a technical study and technological solutions from an engineering point of view, keeping in mind the objectives and requirements imposed by the final user, namely the particle physicists. It is important to remind the reader that this study started literally from scratch, reflecting the numerous steps needed in a large design process including its successes and its dead ends.

A first chapter will describe the frame in which this work was inserted. The experiment is the Compact Muon Solenoid (CMS) collaboration, one of the experiments for which the European Organization for Nuclear Research (CERN) provides beam and experimental facilities. These two entities are described in this first chapter, after having given to the reader a short introduction on the aim and motivations of research in particle physics. During this chapter, the attention is slowly focused towards the end-cap muon spectrometer, which is the part of the CMS detector of interest in this work.

The second chapter is oriented towards the infrastructure of the CMS detector, and

more specifically, the data acquisition system. This is where the integration of a new type of detector, unforeseen in the original design of CMS, will have the biggest implications. As a matter of fact, the data acquisition system is the most constrained sub-system of the entire experimental facility. Incorporating a new sub-detector and its own data acquisition system inside this already existing infrastructure is a challenge. Understanding the architecture, the constraints and the possible extensions is the aim of this second chapter.


The studied new detector type is explained in Chapter Three. A comparison of the existing types of detectors for the detection of muons in the forward region of CMS is given first, followed by a description of the explanation of the physics goals and performance specifications. A more detailed upgrade path is given afterwards, in a chapter listing the implementation details of the proposed upgrade. A broad technical description of the components, from the front-end to the back-end parts of the data acquisition chain is given in this chapter.

The two last chapters are dedicated to the developments made to finalize the proof of concept. First, to fit into the currently planned evolution of the data acquisition system of the entire experiment, a number of developments were made according to a new chosen standard for the DAQ electronics. Moving towards this new architecture standard, called μTCA , a challenge for engineers in electronics. This is why a detailed study to gain experience and understand the advantages is undertaken in this work. As a conclusion to this study, the last chapter presents a mock-up system of a detector prototype and the associated DAQ chain to be implemented, built in our lab and making use of cosmic muons to generate events. Preliminary results of a very similar system tested with a muon beam at CERN are given last. Although this experimental setup is only a proof of concept, attention is given to handle the different implementation challenges as close as possible to the final system.

These very promising results and the design of the architecture are of course the result of a group work, in the name of the Brussels R&D group of the Inter-university Institute for High Energies (ULB-VUB). A number of contributions, however, are the fruit of the work of the author himself. The study of the μTCA standard for example, as well as all the hardware, firmware and software developments performed in this innovative standard such as described in Chapter Five, constitute the first main personal contribution. The numerical simulations of the front-end electronics signal processing functions, as described in Chapter Four are also the author's contribution, and the biggest part of the hardware and firmware developments described in Chapter Six as well. Finally, the author's knowledge in electronics, microelectronics, software developments, computing, networking and systems engineering were essential to set up an adequate engineering environment, to hold a consultative role during the planning phase of the numerous system developments on which the Brussels R&D group participated and finally, were also useful for contributing to the design of the entire GE1/1 DAQ system for the CMS muon spectrometer upgrade.

CHAPTER ONE

The LHC and CMS: The collider and the experiment



The idea that everything around us is made out of smaller elementary constituents is as old as Democritus (460 - 370 B.C.), who postulated that *the universe consists of empty space and an (almost) infinite number of invisible particles* which differ from each other in *form, position, and arrangement*. For him, all matter is made of indivisible particles, and he even gave these the name of atoms. The reason why it took more than two millennia to actually discover something that could confirm Democritus thoughts (Joseph Thompson with the discovery of the electron in 1898) is closely related to the slow evolution of technology until the 20th century. These

two thousand years saw the most brilliant minds in the history of theoretical physics. But it is only with the experimental discovery of the first elementary building bricks of matter at the beginning of last century that physicists were able to found the bases of what became a new branch of physics, namely particle physics.

This chapter is an illustration of what is mentioned above. Based on very accurate models, a set of orthogonal and well defined building bricks (later referred to as Standard Model in this chapter) and a strong sense of intuition, some new particles have been predicted before being discovered. This discovering phase is what the particle accelerators are designed for. After a short reminder of the basics needed to understand the principles behind these new discovery machines, we will focus on the technologies used to create the best experimental conditions (the LHC) and understand how we make observations (the CMS experiment) in this particular branch of physics.

1.1 Cultural background in particle physics

1.1.1 Short contextual history

Scientific discoveries are resulting from overcoming the obstacles to the broadening of our understanding of the Universe. Through the ages, the obstacles were first cultural (ancient beliefs until the Roman Empire, followed by the omnipotence of the church in Europe) until the second half of the middle ages, limits of the technological advancement (overcame with the industrial revolutions), and geopolitical in the last century. The next most probable obstacles to the coming scientific breakthroughs will be financial, as the number of funded research fields sharing the same budget exploded in the last forty years (eg. space, health, environment, social sciences). The field of particle physics is a good example, since the origins of the world and the nature of the fundamental forces were long attributed to divine instances, the experimental part of this discipline is strongly dependent of the state-of-the-art in the field of engineering and finally, the most brilliant minds in this field were brought together by a number of major events in the history of the 20th century.

After the second world war, which boosted the domination of the United States in large scale experimental nuclear physics, interest in bringing European science back to the foreground grew amongst the European scientific community. Following the post-war trend of creating international peace-keeping organizations, the idea of creating a uniting, pan-European and independent entity to promote non-military research in nuclear physics was proposed by Louis de Broglie in 1949. The main advantage of such a structure, of course, would be to rationalize the prohibitive costs of building a world-class research facility from scratch. Between 1951 and 1954 the “Conseil Européen pour la Recherche Nucléaire” (CERN), which was composed of eleven European states, settled the bases of the future entity. Finally, on the 29th of September 1954, the council was dissolved, and the European Organization of the same name was born.

CERN is located on the Swiss-French boarder close to Geneva. It enjoys a yearly budget of about one billion Swiss Francs and employs around two thousand workmen, technicians and engineers, which makes it currently the world largest laboratory in particle physics. In this sense, it is a service provider, aiming at providing facilities, raw matter (particle beams) and expertise to particle physics experiments. These experiments are usually not run by CERN itself but by worldwide scientific collaborations. They are regularly mentioned in the press for achieving major breakthroughs in the domain of particle physics of course, amongst which it would be unfair not to mention the W and Z bosons (UA1 experiment, 1983) and the now famous Higgs boson (joint ATLAS and CMS experiment [1] [2], 2012).

1.1.2 Introduction to particle accelerators

In experimental particle physics, we can distinguish two types of experiments

allowing the study of the existence and the properties of particles. The first type is the non-accelerator based experiment. These experiments are usually associated with astrophysical particle physics and cosmology, since the incoming bodies have been accelerated by astronomical events such as supernovae, black holes and possibly Gamma-ray bursts. The IceCube Neutrino Observatory for example, is recording the photons emitted by the interaction of neutrinos in a cubic kilometer of the Antarctic icecap under the Geographic South Pole. The second type of experiment is called accelerator-based, and is dependent on the presence of the services and facilities provided by a laboratory. The main difference between these two types of experiments is the deterministic nature of the events. In the first case, no-one knows when an interaction will happen, whereas in an accelerator based experiment, the presence of an upcoming event is known in advance. This will have consequences on the architecture of the data acquisition system. The current work was done in the context of an accelerator-based experiment, the services and experimental facilities being provided by CERN and its Large Hadron Collider (LHC).

The fundamental relation behind every particle accelerator is known as the Lorentz force law:

$$\vec{F} = q(\underbrace{\vec{E}}_{\text{electrostatic}} + \underbrace{\vec{v} \times \vec{B}}_{\text{magnetic}}) \quad (1)$$

where \vec{F} is the force experienced by a particle holding a charge q moving at a velocity \vec{v} , inside an electrical field \vec{E} and a magnetic field \vec{B} . Accelerating and guiding the particles to reach high energy particle beams is thus a matter of cleverly combining electromagnetic fields. A good indicator of the performance of an accelerator is the event rate, in other words the amount of occurrences of the researched process per seconds ($f_{process}$), given by the formula:

$$f_{process} = L \sigma_{process} \quad (2)$$

where $\sigma_{process}$ the interaction cross section of the process (expressed in barns¹) and L the luminosity. The latest quantity, expressed in $\text{cm}^{-2}\text{s}^{-1}$, is fully dependent on the configuration of the accelerator and the shape and density of the particle beam. In the case of a circular colliding storage ring such as the LHC for example, two beams circulating in opposite direction at f turns per second and containing N particles each, with a Gaussian shaped profile yielding a transverse size of $4\pi s_x s_z$ the expression of the luminosity L is:

$$L = \frac{fN^2}{4\pi s_x s_z} \quad (3)$$

From equation (3) we can see that the luminosity, and by definition the corresponding event rate, can be increased either by squeezing the beam ($s_x s_z$), increasing the revolution frequency (f), or increasing the number of particles per beam (N).

¹ One barn = 10^{-24} cm^2

A last parameter sometimes used to characterize the amount of data produced by an accelerator is the integrated luminosity over time:

$$L_{\text{int}} = \int L \cdot dt \quad (4)$$

This quantity, expressed in the inverse of a cross-section² is also used to quantify the effective operation time, in order to estimate the aging of the different components of an accelerator or experiment. In the case of the LHC for example, this measure is linked to a deposited radiation dose, which will limit the reliability of some front-end measurement electronics.

1.1.3 Collider experiments

Two main types of accelerator-based experiments exist. First, the fixed target experiment. When the particle masses may be neglected compared to the beam energy, the energy in the center of mass reference frame, usually noted \sqrt{s} , is given by $\sqrt{s} = \sqrt{2E_{\text{beam}} m_t c^2}$, where E_{beam} is the energy of the incident beam particle and m_t the mass of the target particle. It thus grows with the square root of the energy of the incoming particle. The second type is the collider experiment where much higher energies can be reached, since its energy released in the center of mass reference frame is given by the relation:

$$\sqrt{s} = 2 E_{\text{beam}} \quad (5)$$

The challenge in this technology, however, is to make two particles coming from opposite directions at the speed of light to interact without missing each other.

As opposed to the fixed target experiment where the resulting momentum after impact is in the forward direction because of kinematics constraints, in collider experiments the products of the interaction are strongly dependent of the process underlying the interaction and therefore cylindrically symmetric detectors with sensitivity down to small angles are required. This allows the measurement of four quantities on the final states:

- Particle type
- Spatial position and timing
- Momentum
- Energy

This cylinder is usually composed of two main regions. On one hand the barrel around the interaction point, and the end caps closing the barrel on the other hand.

By convention, the system of Cartesian coordinates (X, Y, Z) centered on the

² 1/nb or nb⁻¹ - nanobarn⁻¹; 1/pb or pb⁻¹ - picobarn⁻¹; 1/fb or fb⁻¹ - femtobarn⁻¹

interaction point is set as follows:

- X points towards the inside of the ring
- Y points to the sky
- Z is in the same direction of the beam line.

The XY plane is thus called the transverse plan. To better describe the momentum angle based physics inside a collider experiment, however, a set of polar coordinates (Z, θ, Φ) is more appropriate, as shown in Figure 1. Here Z is still the axis along the beam line, θ is the polar (longitudinal) and Φ the azimuthal (transverse) angles with Z. In addition, to equally distribute the non-symmetric nature of the final states along the θ angle, a new Lorentz invariant called pseudo-rapidity for highly relativistic particles is defined by³:

$$\eta = -\ln\left[\tan\left(\frac{\theta}{2}\right)\right] \quad (6)$$

Below is a view of the reference frame transformation.




Figure 1: Relation between pseudo-rapidity η and angle θ .

³ Strictly speaking it is the rapidity, defined as $y = \ln\sqrt{\frac{E+P_z c}{E-P_z c}}$, which is Lorentz invariant. For highly relativistic particles, $y \approx \eta$

1.2 The LHC

1.2.1 Context

While the Large Electron-Positron (LEP) collider was still under construction at CERN in the mid-eighties, the particle physics community was already looking ahead to the next steps of the Standard Model exploration. After discussions inside the particle physics community and the letters of intent, in 1992, of the ATLAS and CMS collaborations [1] and [2], it was agreed that the next generation of particle accelerator would be a hadron collider and would be built at CERN. The choice of the location was strongly influenced by the cost savings to be made if reusing the existing infrastructure present at CERN [3], namely the LEP tunnel and the existing accelerators which we will later call the LHC injection chain. In December 1994, the CERN council gave its approval to the LHC project, and once the LEP was decommissioned and dismantled in year 2000 after 11 years of fruitful operation, the 27 km circumference tunnel, a hundred meters under the French and Swiss countryside was freed (see Figure 2) and the construction of the most complex and advanced experiment facility could start, in 2001.




Figure 2: Overall view of the LHC experiments.

The construction and commissioning of the LHC took ten years to complete. During these years many major engineering challenges were overcome, a number of civil engineering master pieces were built and a total of 6.5 billion Euro were spent [4]. A magnet quench incident in September 2008 delayed the startup by more than a year, but finally, on March the 30th, 2010, two proton beams collided at 7 TeV in the center of mass reference frame, marking the start of a new era in high energy physics.

1.2.2 Key parameters of the LHC

The LHC injection complex providing the beam for each run is a chain of four pre-existing accelerators, as shown in Figure 3. The base unit is a bunch of about 100,000 million protons, produced by a bottle of hydrogen and a duoplasmatron. A first 50 MeV/c linear accelerator (LINAC2) followed by two synchrotrons, namely CERN's 1959 Proton Synchrotron (25 GeV/c) and the Super Proton Synchrotron (450 GeV/c) are used to inject a grand total of 2808 bunches inside each ring, according to the 25 ns scheme described in [5]. This 25 ns gap between every single bunch is an important number to keep in mind. It is the origin of the 40 MHz global LHC wide bunch crossing clock which will give the synchronization signal for all the instrumentation of the LHC and the experiments. The revolution frequency of the beam around the LHC is 11,245 turns per second. The nominal momentum after the acceleration phase inside the LHC is 7 TeV/c per beam at a luminosity of $10^{34} \text{ cm}^{-2}\text{s}^{-1}$.

Each run has a nominal duration of 10 hours, during which the beams must remain stable and perfectly under control for about four hundred million revolutions. This is achieved with advanced beam and instrumentation monitoring systems, combined with beam control devices such as beam scrapers and collimators. Because of the destructive amount of kinetic energy stored inside the beam, losing its control would have dramatic consequences on the infrastructure of the LHC itself. This is why an extremely reliable beam dump system was foreseen, capable of dumping the 350 MJ of each beam at once. This system is also used at the end of each run.




Figure 3: LHC injection chain [5].

1.2.3 Upgrade plan

Many uncertainties on the reliability of all the systems arose after the magnet quench incident of September 2008. All the magnets had been trained to reach 9 Tesla individually before installation [6]. After the incident, all the installed magnet circuits, however, were only trained to 6.5 Tesla. This limit came from a larger than expected number of quenches, and gave the limit of 3.5 TeV to the beam energy with half luminosity for a first period of operation. The decision was subsequently taken to define a first 10-year phase of operation (phase 1, 2010-2020), dedicated to carefully reach the nominal performance of the LHC, in 3 successive steps separated by 2 long shutdown (LS) periods. Beyond 2020, the phase 2 is then meant to extend the performance beyond the nominal values originally defined in the specifications, namely by doubling the instantaneous luminosity of the beams to $10^{35} \text{ cm}^{-2}\text{s}^{-1}$. Table 1 summarizes the different operation and maintenance periods.

The first period of phase 1, between 2010 and 2012 (7-8 TeV collision, luminosity of $0.5 * 10^{34} \text{ cm}^{-2}\text{s}^{-1}$) was meant to produce enough collisions, in order to perform a thorough Higg's search in the missing mass range between CERN's LEP experiment ($114 \text{ GeV}/c^2$) and Fermilab's Tevatron accelerator ($600 \text{ GeV}/c^2$). This was a success, since 30 fb^{-1} of integrated luminosity [7] were produced in the LHC and allowed the discovery of the famous missing particle in the Standard Model. After a first long shutdown (LS1) of two years, needed mainly to rework the interconnections of all the bending magnets, the planned performance of 14 TeV collisions at the center of mass

reference frame and a nominal luminosity of $10^{34} \text{ cm}^{-2}\text{s}^{-1}$ will be achieved, at the beginning of 2015. Two years of operation at this level are foreseen, before a new long shutdown where some LHC systems will be modified for operation with double of nominal luminosity. These two last years of phase 1 operation are of particular interest in this work, since this is the target period for the systems design described in the next chapters.

After 2021, a phase-2 is scheduled with a luminosity increase of a factor of ten, to reach an ambitious $10^{35} \text{ cm}^{-2}\text{s}^{-1}$. Since only the beam current is increasing, the major challenges are not on the LHC side, but on the experiments. In most of the maintenance periods of phase 1, the detectors undergo small improvements, but no major structural changes. For phase 2, however, the entire detectors and associated data acquisition systems need to be upgraded, to be able to handle a factor ten increase of events per bunch crossings. On the facilities side, provided by CERN, this upgrade will see the commissioning of a new injection complex, replacing the aging PS and LINAC, and upgrading significantly the SPS.

Period	Energy	Luminosity
2010-2012	7-8 TeV	$0.5 * 10^{34} \text{ cm}^{-2}\text{s}^{-1}$
Long Shutdown 1 (LS1)		
2015-2017	13-14 TeV	$10^{34} \text{ cm}^{-2}\text{s}^{-1}$
Long Shutdown 2 (LS2)		
2019-2021	14 TeV	$2 * 10^{34} \text{ cm}^{-2}\text{s}^{-1}$
Long Shutdown 3 (LS3)		
2022-2030 (?)	14 TeV	$10^{35} \text{ cm}^{-2}\text{s}^{-1}$

Table 1: Summary of the foreseen energy and luminosity upgrades of the LHC [3].

1.3 The Compact Muon Solenoid (CMS) detector

1.3.1 Facts and figures

Four large experiments have been built along the LHC, one around each of the interaction points. The Compact Muon Solenoid (CMS) is one of them. It is with ATLAS, one of the two general purpose detectors. Physically, it is 21 meters long, 15 meters in diameter, is built of about 100 million individual pieces and weighs 21,000 tons. Furthermore, the CMS collaboration is made of over 4,300 active members of 182 institutes from 42 different countries around the world⁴.

As most of the collider experiments, CMS is built of the usual main components, as shown in Figure 4. In the barrel region of CMS, a silicon tracker is the first layer around the interaction point. It determines precisely the momentum and scattering angle of the charged particles produced in the collisions. Next, electromagnetic and hadron calorimeters define the type of the scattered particles and their energy. All this is possible with the help of the magnetic field produced by a solenoid in the next layer. Finally, the outside of the barrel and both end caps are built of muon chambers. The latest are part of the subject in the current work.




Figure 4: Representation of the different parts of the CMS detector.

⁴ As of this writing

1.3.2 Barrel

This description of the CMS detector starts with the barrel, from the origin at the interaction point (IP) to the outer layers.

1.3.2.1 Silicon tracker

The very first layer around the beam line is the silicon tracker. Its role is to track the momenta and the scattering angles of the charged particles leaving the interaction point. Two types of semiconductor detectors are used for this purpose, the pixel and the strip detectors. The advantage of these technologies are multiple. First, the spatial resolution which is of the order of 20 microns for the pixel detectors in CMS. This allows very accurate tracking capabilities and a momentum reconstruction resolution of 1% close to the IP. The second main advantage of silicon detectors is the excellent time resolution, since the recovery time after an event is smaller than the bunch crossing frequency. The trade off for this type of detector is the price. They are usually installed where it is strictly necessary, usually at the very center of the detector.

A pixel detector is made of a PN junction based detection layer, stacked on top of a readout layer. Reverse biasing the PN junction creates an electric field. The interactions produced by a charged particle traveling through the dense silicon create electron-hole pairs, each drifting to the electrode, where the readout front-end array converts the produced current into a signal.




Figure 5: View of a CMS silicon pixel detector stack.

Figure 5 shows the layout of a small part of the CMS silicon pixel detector. The size of each detection layer area is $150 \mu\text{m}^2$, and each of this pixel requires its associated readout front-end. As a result, this exceptional density is at the cost of a complex individual wiring scheme. This is usually the limiting factor to the pixel density of a tracker.

Three layers of these pixel detectors are placed around the interaction point to allow an accurate track reconstruction. The next layer of the tracker are made of silicon strip detectors. The detection principle is the same, but the shape of the detection area becomes a strip. This reduces the amount of expensive pixels as the surface of the cylinder increases, but enables detection only in one single dimension. The spatial

resolution is decreasing accordingly to a minimum of 50 microns on the outer layer.




Figure 6: View of the tracker layout. In blue the pixel detectors, in red the strip detector.

Figure 6 shows the overall geometry of the CMS silicon tracker. The pixel detector is made of three layers in the tracker barrel and two disks in the tracker end cap. The silicon strip detector is made of ten layers in the tracker barrel (TIB + TOB), three layers in the disks (TID) and nine disks in the tracker end cap (TEC). The trajectory of the charged particles is influenced by the strong magnetic field generated in the solenoid. The bending radius, in meters, is given by the relation:

$$R = \frac{P_T}{0.3 B} \quad (7)$$

where P_T is the transverse momentum of the scattered particle in GeV/c, and B the strength of the magnetic field in Tesla. Measuring the bending direction of the scattered particle gives us its charge and the bending radius reveals the momentum. This, of course, requires an accurate correlation search between the different layers of the silicon tracker.

The performance of the silicon tracker for muons in the 1, 10 and 100 GeV/c range, is given in Figure 7.

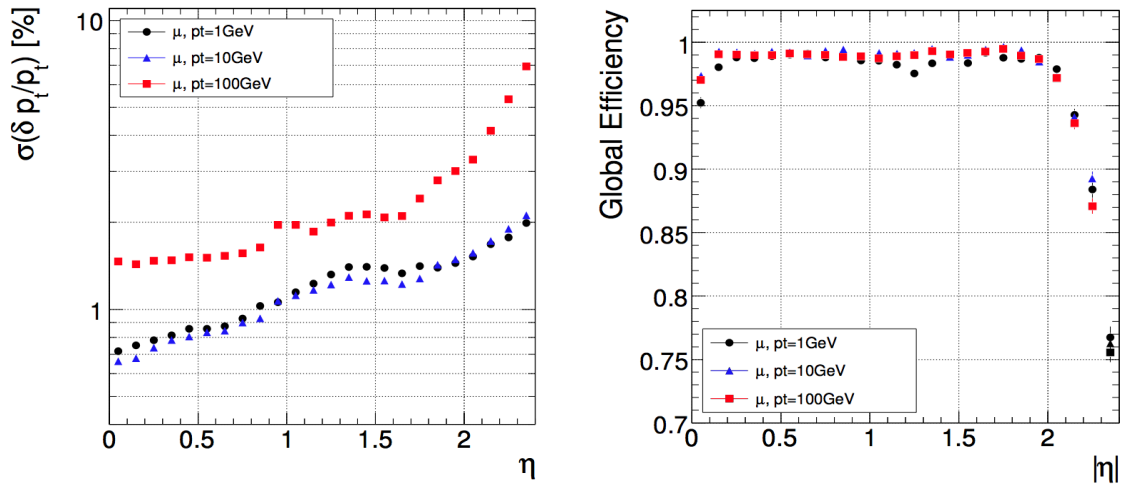


Figure 7: Transverse momentum resolution (left) and tracker efficiency (right) with respect to the pseudo-rapidity.

As expected from the layout given in Figure 6, the efficiency and the resolution are strongly dependent of the pseudo-rapidity, with a remarkable level of precision for $|\eta| < 1.5$, and a rising imprecision beyond this value. This is due mainly because of the higher material density in this region, producing more scatterings, combined with a decreasing strip density.

1.3.2.2 Electromagnetic Calorimeter (ECAL)

Immediately next to the silicon tracker is the electromagnetic calorimeter. The role of a calorimeter is to absorb and measure the energy of electrons, positrons and photons. It makes use of the cascades of secondary particles that incident particles produce via repeated bremsstrahlung and pair production processes inside the material of the calorimeter.

In CMS, for instance, the ECAL is built out of 75,000 PbWO_4 crystal scintillators, each coupled to a photomultiplier tube to amplify and collect the emitted photons. In this type of crystal, the radiation length⁵ of an electron is in the order of a centimeter and the Molière radius⁶ is about 2.2cm [8] Accordingly, each crystal has an area of 2.2cm^2 for a length of 23cm (see Figure 8).

⁵ The radiation length (λ_R) is the distance traveled by an electron before interacting

⁶ The Molière radius approximates the size of a contained electromagnetic shower




Figure 8: One CMS ECAL crystal (left). CMS ECAL crystals setup in one half of an end cap (right) [9].

The energy resolution of the electromagnetic calorimeter is given by the relation:

$$\left(\frac{\rho_E}{E}\right)^2 = \left(\frac{a}{\sqrt{E}}\right)^2 + (b)^2 + \left(\frac{c}{E}\right)^2 \quad (8)$$

where a is the stochastic term representing the number of particles produced in the cascade, b is a calibration constant and depends on the quality of the crystal, and c is a noise factor induced by the electronics. As we can see, the behavior of a calorimeter differs significantly with respect to a silicon tracker in the sense that the resolution improves with the energy. One of the technical challenges for this system is the evolution of the b factor over time, these crystals tending to become opaque with the ambient radiation in the center of the detector. A systematic recalibration of this sub detector will thus be performed during the operation time of CMS. The energy resolution for the CMS ECAL at commissioning time was given by the factors :

$$\left(\frac{\rho_E}{E}\right)^2 = \left(\frac{2.85\%}{\sqrt{E}}\right)^2 + (0.3\%)^2 + \left(\frac{12\%}{E}\right)^2 \quad (9)$$

which is illustrated below:




Figure 9: Energy resolution of the CMS electromagnetic calorimeter [9].

1.3.2.3 Hadron Calorimeter (HCAL)

Hadrons are another byproduct of proton-proton interactions. A hadron calorimeter relies on the strong interaction force between the incoming hadrons and a very dense medium, called an absorber, to produce cascades. Since the energy range of the interacting hadrons is broad, the radiation length is significantly longer than in the case of electromagnetic interactions. The HCAL thus requires a much larger amount of absorber material, such as steel, copper and gold. In this sense, the hadron calorimeter at CMS is built of an alternation of 16 layers of absorber and plastic scintillators in between, for a total thickness of about a meter in both the barrel and the end caps as seen in Figure 10. This corresponds to more than 5 times the nuclear interaction length⁷.

⁷ Nuclear interaction length λ_a is the mean distance traveled by a hadron before undergoing an inelastic nuclear interaction with the medium

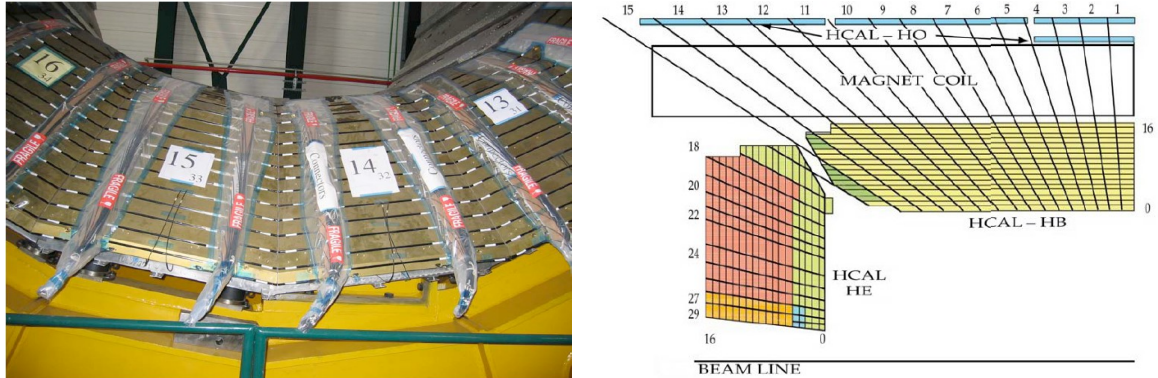


Figure 10: View of the hadron calorimeter (left) and detail of the segmentation (right).

Here also, the effects of the ambient radiation will impair the overall performance of the sub detector. A complex network of optical fibers is laid out along the scintillators to inject calibration pulses. This in order to measure and adjust the b parameter in equation (8) during off-line analyses.

1.3.2.4 Superconducting solenoid

The strong magnetic field needed to measure the transverse momentum of the charged scattered particles is created by a 12.5 meter long solenoid coil around the hadron calorimeter. The advantage of a solenoid coil is the uniformity of its created field in the center, as shown in Figure 11. But creating a 3.8 T capable magnet of this scale is, however, a challenge. To achieve this, the entire spool is enclosed in a cryostat container and cooled down to 4.5 °K to reach below the superconducting temperature threshold of the *niobium-titanium* windings. The power supply required to energize the magnet is rated at a power of 520 kW.

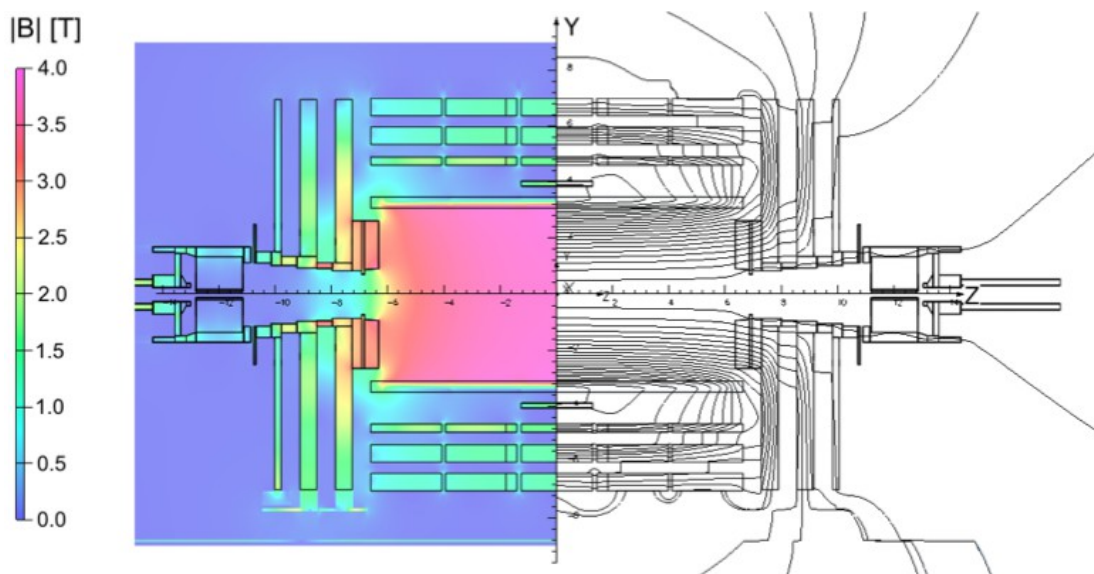


Figure 11: Simulation of the magnetic field produced by the solenoid[10].

The term “compact” in the acronym of CMS comes from the design of this magnet. Measuring the transverse momentum of a charged particle always requires a uniform magnetic field. At high energies, however, the size of detectors grow, making it expensive and technically difficult to use solenoids. To overcome this problem, the ATLAS collaboration built a toroidal magnetic field around the outermost sub detectors of the experiment. CMS, on the other hand, chose to deeply build-in the magnet inside the barrel in order to limit its size, hence the “compact”. This has also as consequence to increase the spatial resolution of the inner sub detectors. The trade-off of this design is that the remaining, and usually voluminous muon chambers, are on the outside of the solenoid. In the case of CMS, this external sub detector takes advantage of the returning (opposite direction) magnetic field lines which are canalized by a steel yoke around the solenoid. Most of the impressive weight of CMS come from this yoke of 10,000 tons on its own. Figure 12 shows the steel yoke components of CMS.



Figure 12: View of the yoke elements around the solenoid [9].

1.3.2.5 Muon chambers

The dense materials present in the hadron calorimeter and the solenoid were chosen with an absorption interaction length⁸ of about 17 cm, which is short enough to absorb the entire hadron energy. In addition, since the energy loss by bremsstrahlung depends on the particle mass as shown in equation:

⁸ The nuclear absorption length is the scattering distance of a hadron in a medium

$$\left\langle \frac{-dE}{dx} \right\rangle \propto \frac{E}{m^2} \quad (10)$$

only muons and weakly interacting particles such as neutrinos are still present beyond the coil. This is thus where the muon chambers are located, assembled inside the layers of the magnet yoke, where an opposing magnetic field exists.




Figure 13: View of the barrel muon stations.

The muon system is the biggest sub detector in terms of surface, with a total covered area of 25.000 m² of detector planes. To make the design cost effective and robust, the choice was made to install gas detectors in the barrel region. Since this document is mainly focusing on the muon system, a complete description of the current state-of-the-art muon chambers will be given in Chapter Three.

1.3.3 End caps

The second part of the detector is composed by the two end caps, which are essentially built of muon chambers.

1.3.3.1 Muon chambers

The end cap muon system is the subject of this work. The barrel muon chambers cover a pseudo-rapidity of $|\eta| < 1.2$. Beyond this angle, the muon flux is caught by the end cap muon system. Two technologies currently coexist in this system, namely Cathode Strip Chambers (CSC) and Resistive Plate Chambers (RPC). Both technologies are described in Chapter Three of this document.

The layout consists in consecutive layers along the Z direction, in order to allow precise momentum measurements by finding coincidences between hits on the trajectory of a muon. For redundancy reasons, the CSC and RPC chambers are both installed in parallel in the $0.9 < |\eta| < 1.6$ region. Initially it was planned to have RPC detectors up to $|\eta| < 2.4$. However, for economic reasons and because it was expected that the muon setup would perform well without RPC's in the range

$1.5 < |\eta| < 2.4$ during the first years of LHC running when the luminosity would be much lower than the nominal value, it has been decided to stage the installation of RPC's in the highest η region. The space for these chambers was nevertheless foreseen in the end caps, and left empty for the initial startup up in 2010. These empty spots are of importance in this work as we will see in Chapter 3. Below is a view of the end caps. MB stands for the muon chambers in the barrel region and ME locates the muon system locations in the end caps.




Figure 14: view of the end cap muon system of CMS [11].

1.3.3.2 Forward Calorimeter

About eleven meters away from the interaction point, measuring between a pseudo-rapidity of 3.0 and 5.0, is the forward calorimeter. Its focus is on hadron products of the collisions, namely in the form of high density jets. The cascades are created by heavy steel blocks, with embedded fibers to detect Cherenkov light. The main research topic of this sub detector was Higgs productions at high mass. Since this channel is now excluded by the discovery of 2012, the main goal of this calorimeter is now to improve the accuracy of the missing transverse energy measurement.

1.3.4 Upgrades

In the coming years, following the LHC upgrade plans, CMS will improve some of its sub detectors in order to adapt the experiment to the foreseen luminosity upgrade, complete several missing functions on the initially installed setup, and to refocus on some aspects of the physics program. For instance, now that a low energy (125 GeV/c²) Higgs scenario has been validated, the upgrades of the LHC and the detectors will allow to focusing on the in-depth study of this particle's properties.

CMS was not built to operate at, nor to withstand the effects of, a luminosity upgrade such as planned at the end of LHC Phase 1 (2019-2020) and phase 2 (beyond 2021), see Table 1. The silicon tracker needs to increase its resolution to avoid significant signal losses at higher interaction densities. In addition, for phase 2, the tracker will have to participate in the trigger decision, which is not the case for the current tracker (see section 2.2.1). This task involves a significant R&D program and won't be ready before the last long shutdown of phase 1 in 2017-2019. For this luminosity upgrade, the entire trigger system needs to be redesigned in order to handle the foreseen average of 20 collisions per bunch crossing and the associated amount of generated events pile up. But one major upgrade is of importance in this work, namely on the end cap muon system.

As was explained in the previous section, some spaces originally foreseen for RPC detectors were left empty at high pseudo-rapidity in the end cap. A particularity of the RPC technology is the degradation of the efficiency as the rate increase. This will become a limitation for the luminosity upgrade of the LHC and makes now the RPC a non-viable solution to fill the spaces initially foreseen for them. This work is part of a feasibility study on the implementation of another technology to replace the initially foreseen RPC chambers at $|\eta| > 1.6$ in the end caps. This novel technology is extensively described in Chapter 3.

A second major topic in the CMS upgrade plans in relation with this work is the improvement of the data acquisition system (DAQ) electronics. The current infrastructure standards used to host the read out electronics is of an older design and will certainly not be able to withstand the dramatic throughput increase associated with the high luminosity upgrades. Another, more modern, infrastructure standard was chosen inside CMS, and the design of the new muon systems based on this new standard will be the topic of Chapter Four.

1.4 Conclusion

We have seen in this chapter that experimental particle physics is a field where large collaborations building over-sized detectors are often required to make discoveries nowadays. CERN is one of the few laboratories hosting these experiments in the world, and CMS is a good example of their size and growing complexity. Throughout this chapter, we learned about different sub-detectors of CMS, to get the big picture of all the sub-components which may be of importance during this study. Narrowing down to the basic principles and key parameters of each of this detector layer allowed us to grasp the challenges of physics processes measurements in the context of collider experiments.

Particle interactions in these experiments are very complex events to analyze, due to the high number of different elementary elements released, and the large energy range window to focus on, in order to see these elements. Many detection layers are usually assembled around the interaction point to identify every constituents. To form a complete picture of an event, all the data from each of these sub-detectors are assembled into a main, detector-wide trigger and data acquisition system. This is the topic of the next chapter, specifically for the CMS experiment.

CHAPTER TWO

The CMS trigger and data acquisition system



As we have seen in the previous chapter, CMS is made out of numerous sub detectors of different technologies, representing millions of heterogeneous signal channels. All these technologies have their specificities, calibration parameters and timing constraints. Nevertheless, the signal of each individual channel has to be conditioned, read out, centralized and transmitted to an event builder which will condense the signals to form an event dataset ready for off-line analysis. This entire chain is called the Data Acquisition system, further referred to as DAQ. In large scale experiments such as in the field of collider physics, the DAQ system requires the use of a broad range of technologies, amongst which electronics, computing and networking are the backbone. Although this work is focusing on the electronics developments of the muon system DAQ for the next upgrade of CMS, a broad understanding of the CMS wide DAQ is needed and will be explained in this chapter.

A first section will give a quick explanation of the different concepts and definitions in the field of DAQ systems, and in particular the strengths and important points of the trigger system. In CMS, 75 million channels are producing data at every bunch crossing. Due to technical limitations on the bandwidth and processing power of the data storage stages, a number of filtering layers have to be applied to this massive data stream without losing any information on already low cross-section processes. This is the role of the trigger system. A second section will thus give an overview of the trigger system of CMS, followed by a broader view on the architecture of the data path in this detector. A last section will focus on the hardware implementation of these systems, with an emphasis on a new standard chosen for the coming upgrades.

2.1 Basic DAQ design features

In the case of collider experiments, the event rate is always given by the bunch crossing frequency, for example 40 MHz in the case of the LHC experiments. This rate is usually high and reading out every sensor channel at this rate is beyond the technical and financially affordable capabilities of today's networking and storage technologies. To limit the amount of data processed, only the hit channels are recorded. Given the fact that the multiplicity⁹ depends on the number of individual interactions at each bunch crossing, the initial scaling of the DAQ architecture is function of the luminosity of the accelerator, which is known during the design phase of a general purpose experiment such as CMS. Concerning the individual sub detector systems, the data volume is function of the event rate such as defined in Chapter One, the area covered by the detector and the granularity of the readout electronics.

Besides transmitting only the hit channels data from the front-end electronics, some advanced filtering is usually performed in order to limit the required bandwidth of the DAQ infrastructure. Paradoxically, two parallel data paths from the front-end to the central DAQ system usually coexist to achieve this filtering. First, a fixed and low latency trigger data path is providing an interruption-type signal to the DAQ instrumentation when a channel shows a hit signal. This allows a first level of fast decision filtering, based on coincidence patterns in regions of interest. Based on this trigger data, if the fast decision algorithm running inside the DAQ system validates the existence of a detector-wide event, the channel data, also called full granularity or tracking data, is read out over the second unconstrained latency channel. Inside the front-end electronics, the event data is usually stored in a circular buffer memory area, and tagged by bunch crossing number to allow its effective retrieval by the DAQ when requested.

The processing power as well as the communication channels bandwidth, both giving the latency constraints of this first level triggering system, are critical. This is explained by the size of the detector. In a fairly large detector, such as modern collider experiments, the flight time of a muon to reach the forward detection region can be longer than the bunch crossing period. This results in a first level of events pileup in the CMS detector. In addition, the trigger decision time which depends on the length of the trigger data path, is typically much longer than the time between two consecutive bunch crossings. The circular buffer memory used to store the successive events inside the front-end chips is usually limited in size. As a result, it is important to define a latency budget for the trigger path, counted in number of bunch crossings and which should never exceed the storage depth of the buffers inside the front-end electronics. This latency budgeting is one of the most important points when designing a DAQ system, and will come back later in this work.

Beyond the experiment specific trigger system, the rest of the data acquisition system for processing events data is less demanding in terms of bandwidth and latency. The event builders used to assemble the event data of all the relevant parts of the detector

⁹ The multiplicity is the number of simultaneously hit channels per bunch crossing

into one coherent dataset, for example, are software engines relying on general purpose IT infrastructure. This enables a high level of scalability and tends to ease the development and integration of auxiliaries such as experiment control interfaces for operators and automation or supervision functions for slow control and monitoring. Although these functions are not visible from a physics point of view, they are considered as being part of the DAQ system, since most of modern physics experiments contain large pieces of industrial infrastructure, such as water cooling, gas distribution and high voltage power supplies.

2.2 Overview of the CMS trigger system

At a collision rate of 40 MHz and taking into account every readout channel, the aggregated data rate produced by CMS is about 100 Gbit/s [12], which is far more than any storage technology can handle so far. Building an efficient trigger system architecture to significantly reduce this data stream without losing events was the complex task of the trigger design group, especially given the performance specifications. For example, where the ATLAS detector has three levels of triggering to reduce the data volume for each event, the CMS detector has only two:

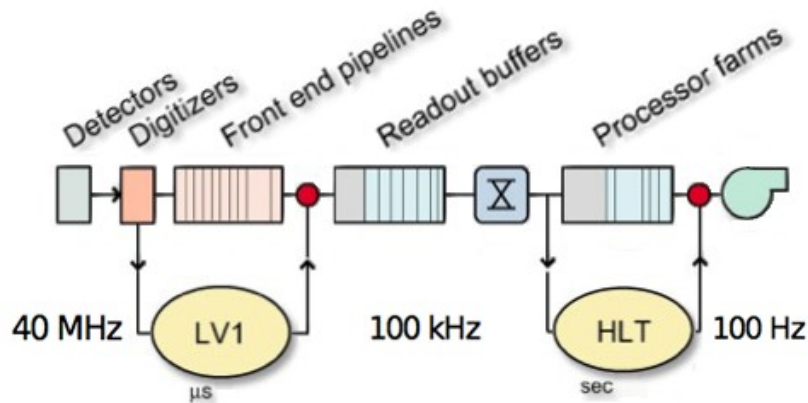


Figure 15: CMS trigger system block diagram with the data rates between each level

The main advantage of having fewer trigger levels is the direct availability of much more unfiltered data in the case an event has been accepted. The event reconstruction is thus more accurate and the precision higher. The required processing power to achieve such a flat trigger function in a delay of no more than several bunch crossings, on the other hand, is much more of a challenge. This section will first focus on the two layers of the trigger system of CMS, and subsequently on the muon system trigger mechanism and its hardware implementation.

2.2.1 Level-1 trigger (LV1)

For each bunch crossing, the sensors of the entire detector are read out and recorded locally in memory buffers inside the front-end electronics. Inside each of this piece of hardware, a local threshold-based decision trigger sends out a signal to a local sector trigger concentrator if the collision produced a signal with a sufficient magnitude to possibly be a valid hit. Several layers of concentrators, aiming at recognizing local patterns and coincidences, are stacked on top of each other to reach the uppermost global trigger system. This is depicted in Figure 16. If a valid trigger pattern is found, a Level-1 accept signal (produced by the Global Trigger Processor) is sent back to the entire detector, containing the bunch crossing number and a command to retrieve the corresponding data. Since most of the front-end electronics is highly optimized for

fast processing as well as for radiation tolerance, the amount of memory in the event data buffers is limited. On the muon system read-out chips, for example, a maximum of 128 events can be stored in the circular buffer before the memory location is overwritten. At a 25 ns bunch crossing period, this gives a total latency of 3.2 μ s, during which the entire triggering chain is solicited, up and downwards.

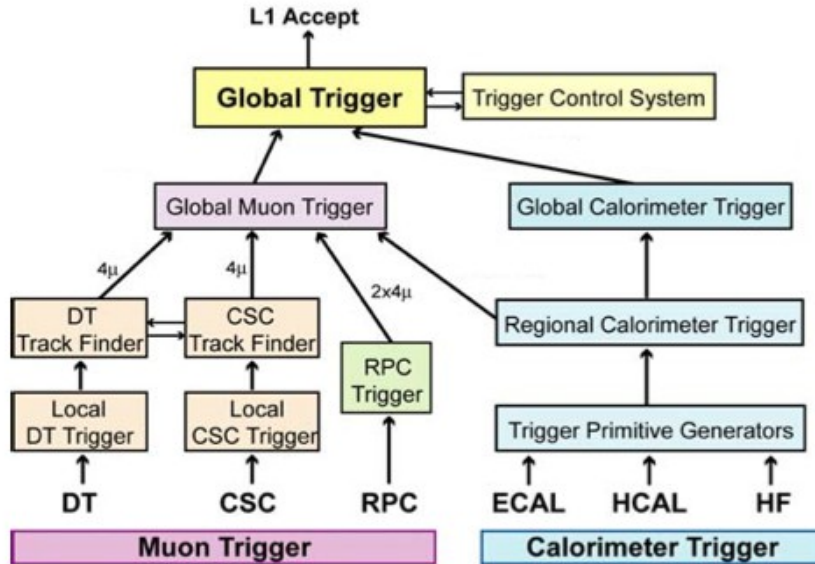


Figure 16: CMS trigger data decision flow [12]

It is needless to say that each individual step in the decision flow should not exceed one or two bunch crossings, in order not to add up to the latency of every communication channel. Packeting/unpacking, modulating/demodulation and propagating an optical signal over a fiber, for example, can be done in a fixed-latency manner. However, it can take several tens of bunch crossing slots over a long distance link, especially between the front-end electronics and the global trigger systems located in the service caverns of CMS. To minimize this layered processing time, the algorithms are parallelized as much as possible. This is possible with the help of modern high density programmable logic, such as Field Programmable Gate Arrays (FPGA).

The advantages of this technology are multiple. Close to the detection areas, the radiation levels usually require the use of dedicated Application Specific Integrated Circuits (ASIC), of which we will see an example later with the VFAT2. But further away from these active zones, concentrating the numerous links not only requires high channel density electronics, but also scalability to control the development costs and adaptability. This last point was the key point promoting the use of FPGA, since the pattern recognition and coincidence search algorithms were likely to evolve with the recalibration and the better knowledge of the detector. The lucky consequence of this is that today, these trigger algorithms which are likely to evolve and be adapted for the coming updates, will require minimal hardware intervention, since a remote upgrade of the firmware can take care of this.

2.2.2 High Level Trigger (HLT)

The drastic data reduction operation achieved by the first level of the CMS trigger system allows a more flexible second layer of filtering at an incoming rate of 100 kHz. Processing events at this rate is possible with modern data center technologies, allowing a streamed processing time per event of the order of a second. This second level is called the High Level Trigger and is performed by software algorithms running on a computer farm of about a thousand processing nodes. The resulting data, when accepted by the triggering software, is finally stored in a distributed storage vault for later off-line scientific analysis at a rate of hundred events per second. Something to keep in mind is that the reduced dataset collected for each event is about 1.5 Mbyte, resulting in several hundreds of megabytes per second of data to be stored during each LHC run.

The technology used for this sub-system is commercial off-the-shelf IT infrastructure. Only the software architecture is designed and built in-house. An extensive description of this architecture and implementation is given in [13], but is beyond the scope of this work.

2.2.3 Muon system L1 trigger

As we can see on Figure 16, the muon system local triggering logic is a big part of the entire L1 trigger system of CMS. The main reason for this is the diversity of sources and resolutions in the muon system, as well as the complex timing scheme caused by the size of this sub-detector. The good spatial resolution of the drift tubes (DT) and the cathode strip chamber (CSC) for example allow these to be used as transverse momentum trigger in the magnetic field of the barrel area. Each of these technologies have thus their dedicated but similar local and regional threshold trigger and track finder electronics. The resistive plate chambers (RPC), on the other hand, are not dependent of any form of drift, and as a consequence, have a much better timing resolution (~ 1 ns). This will allow very accurate bunch crossing (BX) identification, requiring its dedicated electronics.

In the barrel region, the DT local trigger is defined by a coincidence detection on the cathodes of four consecutive modules. In addition to finding the position of the muon track in the DT modules, computing the drift time between layer pairs makes it possible to find the azimuthal angle Φ of the muon, as depicted in Figure 17. The disposition of the DT modules is done in such a manner that coincidences of drift times are comparable and easy to pair. The number of pairs will define the quality of the track. This simple disambiguation mechanism is fast and efficient, and provides a good quality local trigger to the track finder.




Figure 17: Muon track position and angle measurement in DT [9]

Similarly, the CSC local trigger provides an accurate η measurement in the end cap region by finding coincidences amongst hit cathode strips. The only difference being that, to decrease the processing time, strips are grouped with a fast OR logic operation by clusters of 5 to 16 strips inside the electronics. On top of the local track reconstruction, the track finder algorithms are identical for the DT and CSC systems. A powerful track extrapolation method using the incoming angle Φ is used to send the most plausible tracks to the Global Muon Trigger (GMT). This method is based on lookup tables (LUT) where the simulated curvature parameters of the best fitting tracks are stored and retrieved. The advantage of LUT is the very fast access time, and thus the low latency cost on the entire L1 trigger system.




Figure 18: Extrapolating track reconstruction algorithm [9]

However, due to the presence of noise, the hit position uncertainty in each muon detection layer and the large size of the muon spectrometer, the track finder algorithm

often computes more than one muon candidate. A voting scheme based on the best match between measurements and extrapolated tracks retains a maximum of four candidates and send these to the GMT. Given the bending and the position inside the detector, the transverse momentum is estimated and provides an extra decision variable for the global L1-trigger system.

Concerning the RPC, several layers of these chambers are spread all over the detector to provide a very accurate trigger window to the rest of the muon system. Unlike the DT and the CSC, no local coincidence is processed for the RPC. To find the transverse momentum of an incident muon, a region-wide RPC trigger system concentrates the strip signals. A pattern comparison algorithm (PAC) is applied, based on lookup tables produced by simulations performed beforehand. Here again, a number of possible tracks are resulting after the calculation. The two most suitable tracks from the barrel are added to the two best tracks in the end caps and send to the GMT.

2.3 The CMS Data Acquisition System architecture

In addition to the already complex dual stage trigger system, the rest of the DAQ is also impressive. The data readout stage, for example, is meant to retrieve the tracking data from millions of channels, sectored in more than 650 sources at a time, once a L1-accept trigger is issued by the Global Trigger Processor. This operation is performed asynchronously, to make the best usage of the front-end electronics processing resources and communication channels. In the case of CMS, all the retrieved data is processed in an event builder, composed of a farm of standard personal computers. Once the event datasets are built, an event filter coupled to the HLT is applying a last level of data reduction, and dispatches the datasets to the different storage vaults with an acceptable speed for today's storage technology. Finally, a parallel monitoring and control system is providing an interface to the detector operators of the entire DAQ process.

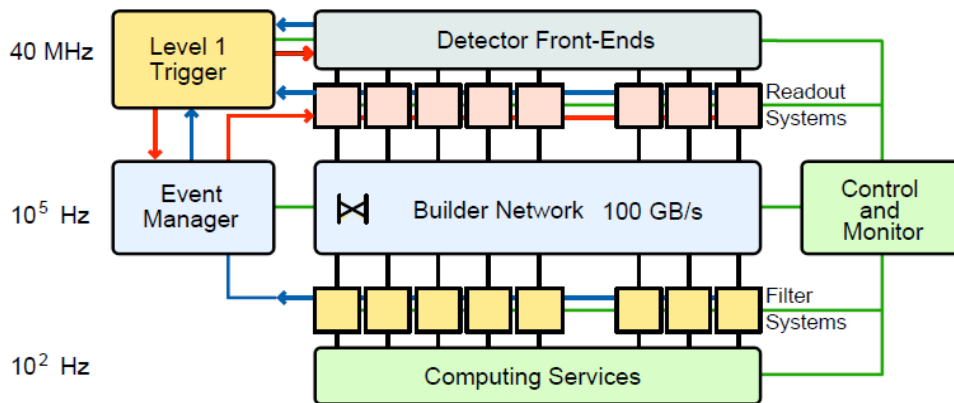


Figure 19: The CMS DAQ architecture [9] with the data rates at each stage

2.3.1 Event readout interface

Each sub-detector is assigned a number of Front-end Readout Links (FRL) to the central DAQ system according to its granularity. These FRL are optical link inputs with a defined communication protocol and data format. The readout and the concatenation of the individual sensor channels is left to the responsibility of sub-detector collaboration. The FRL electronics is composed of a sub-detector specific Front-End Driver (FED) board, which is pushing the region data towards the central DAQ system over the FRL once an L1-trigger is issued. Below is a table summarizing the number of FED allocated to each sub-detector.

Sub-detector	Front-end channels	Number of allocated FED
Tracker (total)	~ 67 Million	536
ECAL	75848	54
HCAL	9072	32
Muon DT	~ 500'000	10
Muon CSC	192'000	8
Muon RPC	195'000	3

Table 2: FED link repartition over the different sub-detector of CMS [9]

On the DAQ side of the 200 meter optical cable located at the surface of the CMS cavern, a Readout Unit (RU) is forwarding the data to the event builder network sequentially. The latency is not a constraint in this transfer, but the data integrity must be guaranteed at a speed of 2.5 Gbit/s.

Something to keep in mind with the front-end part of this event transfer system is the proximity with the detectors. Some concentration points, for example, are located in the balconies surrounding the barrel or in the end caps, and are thus exposed to high magnetic fields and levels of radiation which may exceed the commercial grade standards of most electronic components.

2.3.2 Event builder and filter

As we can see in Figure 19, an event manager is responsible for initiating a region-wide tracking data transfer amongst all the sub-detectors concerned by an event. This event manager relies on the L1-trigger signal received from the Global Trigger Processor. The event building process shown in the schematic view of Figure 20, is divided in two stages. First, the FED builder assembles RU-originating data in blocks of 72 fragments. These fragments are buffered in physically separated memory blocks, corresponding to different FRL regions, called slices. The second stage, called the RU builder (RB), merges these super fragments into 1.5 Mbyte event datasets at a rate of 12.5 kHz. To reach the target 100 Hz output rate, a last filter is applied to the data. This corresponds to the HLT described in the previous section, and is performed by the event filter. This unit first ensures the consistency of the incoming datasets for data quality assurance, then takes the final decision whether the event is rejected or stored in the computing service for later off-line analysis. Finally, the event filter routes the events datasets to the different storage databases, tagged with a unique identification number to allow easy off-line event reconstruction.

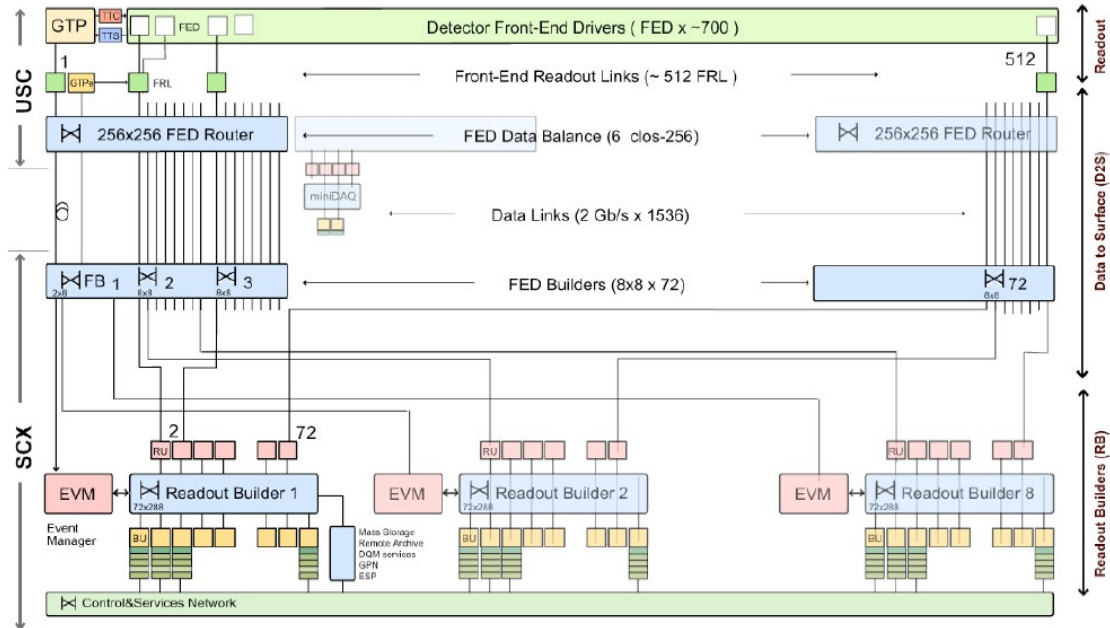


Figure 20: schematic view of the event builder [9]

2.3.3 Control and monitoring

The detector-wide on-line software framework for CMS is called xDAQ [9]. It is a rather complex sub-system which we won't describe in details here. It is bringing together three main components, namely the Run Control and Monitoring system, the totality of the DAQ sub-systems and a Detector Control System, responsible for the slow control and the auxiliaries. This is shown in Figure 21 below.

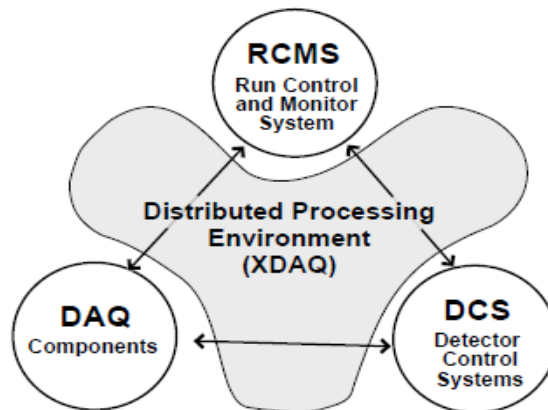


Figure 21: CMS on-line monitor and control systems architecture [9]

The level of complexity of the interconnections and interaction between all these

systems, soberly named as Distributed Processing Environment in the block diagram, is such that it is relying on its own IT infrastructure. Detector operators have access to both the RCMS and the DCS control interfaces from the CMS operator room located in the surface building. The RCMS interface is used for run control, and for good operation and data quality checks. The DCS is handling the infrastructure components such as gas distribution, cooling, power supplies and more.

Although the completeness of the core architecture of this xDAQ framework makes it a complex component inside CMS, a unified protocol allows for easy integration of additional sub-detectors when needed. This interface is called the Hardware Abstraction Layer, and makes it easily possible for xDAQ to communicate with the DAQ electronics of a new coming sub-detector easily, without the sub-detector collaboration to be expert in xDAQ.

2.4 Hardware implementation

2.4.1 Front-end electronics

Two types of technologies are present in the front-end electronics. On the forefront, close to the detectors, and usually enclosing the analog shaping and conditioning circuitry, Application Specific Integrated Circuits (ASIC) are used. These chips are designed by the collaborations to meet the specificities of their sub-detector outputs. In addition, a number of constraints imposed by CMS beforehand need to be met such as power consumption and dissipation or radiation tolerance. ASIC is the ideal technology for this part of the data acquisition chain, since no other technology is able to meet these constraints. The cost of the design process is however not at the reach of small collaborations, this is why several sub-detector collaborations usually tend to federate their efforts in order to build a chip that suites everyone's needs.

A bit further down the data acquisition chain, where the data streams of many front-end chips merge together, FPGA are preferred. This technology, sold as a commercial product, is inherently more sensitive to radiation, hence the distance. In addition, the power density of these chips is usually high, making the powering and cooling impractical in areas close to the interaction point. The main advantages, on the other hand, are the low development costs and effort, the front-end agnostic interfacing possibilities because of a high input / output pin count, and the unlimited remote reconfiguration possibilities. This last feature is, for example, extremely useful for the trigger path, as we will see in Chapter Four. Combined to the high logic density of modern FPGA, it allows increasing the granularity inside sub-detectors without putting more stress on the communication channels to the higher level trigger decision systems.

A last main component of the DAQ system of CMS is the Versa Module Europa (VME) infrastructure. It centralizes and processes all the sensor data as well as the trigger data streams (eg. L1-trigger, RU buffers) before being sent to the surface. The VME standard was chosen to host the DAQ electronics from the beginning of CMS because it was the most widespread standard for particle physics experiments, and its low complexity enabled fast initial developments by a rather small collaboration at that time. One remarkable fact, however, is that it quickly became clear to all the LHC partners, including CERN and the other experiments that the existing VME standard would not allow a high enough level of reliability nor sustain the tremendous data rates to be produced by the LHC. This is why, the choice was made to adopt an extended version of the VME bus implementation, named VME64x, becoming quickly the de facto standard for all LHC instrumentation as well as inside CMS. The two main improvements of this implementation are an extra 160-pin data connector for increasing the throughput and a plug-and-play feature to allow quicker replacement when a module is breaking. The crates are located close to the detector to reduce the communication latency, namely in the balconies and in the service cavern.

2.4.2 ATCA standard

With the coming luminosity upgrades of the LHC, the data throughput will increase dramatically. More hit channels per bunch crossing will produce more local triggers, more local coincidences and track candidates will be computed, more sensor data will have to be retrieved and more events will need to be built and filtered. It is clear that the current DAQ systems relying on the VME standard will be unable to handle the rate increase. Aware of this future limitation, a working group was set up inside CMS to evaluate a better suited type of infrastructure for hosting the cavern electronics. Three main requirements for the new data acquisition system architecture have to be fulfilled. First, no hard limit on the possible data throughput should be reached for any of the coming upgrades. This involves a high level of scalability. Secondly, the focus should be set on reliability. This can be achieved by implementing fault tolerance and recovery features. Lastly, integrating the infrastructure in the control and monitoring systems should help to prevent failures and to reduce maintenance time. This is called serviceability. When looking at these requirements, one can easily recognize the usual challenges found in the telecommunication industry. This is why the working group quickly moved towards carrier-grade infrastructure standards.

In the field of computing and telecommunications, the PCI Industrial Computer Manufacturers Group (PICMG) is a consortium of over four hundred major equipment manufacturers developing together open specifications for high performance telecommunication and industrial computing applications. This consortium is at the origin of some of the most widespread standards in the industry, such as PCI, Compact PCI and PCI Express. In December 2002, PICMG finalized a first set of specifications for the next generation of high bandwidth and high availability carrier grade telecommunication equipment, called the Advanced Telecommunication Computing Architecture (ATCA). The form factor of this standard is similar to the VME crates, but rather than being a computer I/O bus, it is meant to be a piece of protocol agnostic infrastructure with hot-swapping capabilities.




Figure 22: General view of an ATCA crate [14]

The ATCA specification only defines requirements on the level of physical dimensions, power distribution, interface connectors and platform management. The modules to be inserted in the crate are dependent on the needs of the user. A wide variety of modules ranging from high density and speed I/O boards to processing server blades are allowed. With high availability in mind, the interconnection backplane is highly redundant, offering several topology variants, such as a full mesh connecting each of the 13 slots to any of the others or a dual star connecting all the slots to two central switching points with the possibility of fail-over redundancy. As the architecture is protocol agnostic, any type of fabric interconnect can be used and even coexist on the differential pairs of the backplane. The most popular fabric interfaces used in ATCA crates are Gigabit Ethernet, XAUI, Serial RapidIO, Infiniband and Packet Routing Switch (PRS).

Every Field Replaceable Unit (FRU) composing the system such as the blade modules but also the power supplies, the fan trays and even the backplane is featured with an Intelligent Platform Management Interface (IPMI) end-point. Besides performing local module monitoring and control, this little piece of intelligent hardware is also enclosing its module specific information, such as the power requirements or the defined fabric interface type, or the topology in the case of the backplane. This information is stored in a FRU record. A redundant pair of Intelligent Platform Management Controllers (IPMC) which is located in the crate, reads out the FRU records when a new module is inserted, and acts as a shelf manager. The goal of this internal infrastructure management system is to avoid fabric interface mismatches or inappropriate power matchings between the loads and the supplies. In addition, this monitoring data is available to the outside over a standard Out-of-Band Management¹⁰ protocol in order to perform fast failure recovery and predictive failure analysis, reducing the downtime in the case of an outage. In short, this new standard was designed for high data throughput, and as an architecture showing no single point of failure to achieve an availability of 99.999 % which represents less than ten minutes downtime in a year.

To reduce the ownership costs associated with this high end standard, the idea came to the PICMG consortium to develop a sub-standard of ATCA, called AdvancedMC (AMC) allowing the insertion of FRU mezzanine board onto ATCA carrier blades. Figure 23 shows a sketch of the relationship. Initially designed to ease the prototyping process of these AMC boards, a second sub-standard of ATCA called μ TCA was released in 2006 [15]. This lightweight type of ATCA interface quickly became popular in the science community, due to the fact that the complex ATCA shelf management is handled outside of the FRU, on a dedicated, commercially available μ TCA Carrier Hub (MCH) FRU.

¹⁰ The Out-of-Band Management is a new trend, allowing remote control and monitoring of a device over a dedicated network channel, and from a dedicated piece of hardware independently from the device itself.




Figure 23: ATCA to AdvancedMC relationship [14]

Because of the high level of availability and performance inherited from the ATCA standard, the μ TCA form factor was chosen to become the reference architecture for hosting the future cavern electronics in CMS.

2.4.3 μ TCA architecture

Like its parent standard, μ TCA is fully protocol agnostic. The standard specifies a number of power requirements, mechanical dimensions and a local simplified platform management scheme. Two physical form factors are defined, a single width and a double width. Both are represented on Figure 23, center and right. Again, there is not strictly defined backplane topology. Full mesh, dual star, and custom user-defined topology are available. The most common topology however, especially in science experiments, is the dual star architecture where all the differential pairs are doubled and routed to two distinct star points, where a redundant set of fabric switches may be inserted. This forbids the existence of a single point of failure. As we will see in Chapter Four, this handy centralization feature will have another function in the case of CMS.




Figure 24: Schematic of a μ TCA crate [14] enclosing Cooling Units (CU), Power Units (PU) and an Rear Transfer Module area (RTM).

From a more functional point of view, we find many similarities with the ATCA standard. Twenty ports composed of two opposite unidirectional differential pairs each are able to fit any possible high speed fabric interface. The first four ports are however bundled in a common options zone. Port 0 and port 1 are dedicated Gigabit Ethernet links and ports 2 and 3 are dedicated SAS/SATA links for connection with a storage medium. Note that these two ports are often routed only between specific locations on the backplane in order to accommodate one or two CPU boards with maximum two disks without having to implement any complex SAS/SATA switch into the fabric switches of the MCH. Ports 4 to 7 and 8 to 11 are often agglomerated in two redundant four-lines communication channels.

CONNECTOR REGION		AMC PORT #	SIGNAL CONVENTIONS			NON-REDUNDANT MCH FABRIC #	REDUNDANT MCH #/ FABRIC #
Basic Side	Common Options	0	AMC.2 1000BASE-BX			A	1/A
		1	AMC.2 1000BASE-BX				2/A
		2	AMC.3 SATA/SAS			B	1/B
		3	AMC.3 SATA/SAS			C	2/B
	Fat Pipe	4				D	1/D
		5	AMC.1x4 PCI-Express	AMC.4x4 SRIO	AMC.2 10GBase-BX4	E	1/E
		6				F	1/F
7	G	1/G					
Extended Side	Extended Fat Pipe	8					2/D
		9		AMC.4x4 SRIO	AMC.2 10GBase-BX4		2/E
		10					2/F
		11					2/G
	Extended Options	12					
		13					
		14					
		15					
		16					
		17					
		18					
		19					
		20					

Figure 25: Ports definition of the μ TCA standard [16]

Since everything is user configurable due to the protocol agnostic nature of the μ TCA standard, several clock network options are available as well. A first option usually referred to as the telecommunication option has a dual interleaved distribution of the four available clock lines of the backplane. Each of the MCH distributes two out of four clock signals to all the slots in such a way that each FRU receives four clock signals in total. In this configuration, an MCH and a clock receiver on the FRU may fail without any clock distribution interruption. The second option is called the fabric interface option, and does not allow any redundancy. In return, a global clock signal is distributed from the first MCH slot to the entire crate on the same dedicated clock lines. This kind of global clock signal is required on some synchronous fabric interfaces such as PCI Express (PCIe), which makes the use of generic CPU boards easier.

A standard size μ TCA crate usually provides twelve slots for user FRU modules, and a dual thirteenth slot for the redundant MCH complex. As for blades in ATCA, each FRU holds an IPMI end-point called the Module Management Controller (MMC), performing local board monitoring and control functions as well as providing the content of its FRU record upon request by the shelf controller. In the case of the μ TCA standard, however, this shelf controller function is located on the MCH rather than in the crate, as it was the case in ATCA. Developing its own FRU modules is thus simplified, as long as a genuine commercial MCH is present inside the crate. This development activity is thoroughly detailed in Chapter Five of this work.

To summarize on the μ TCA standard, the block diagram below gives a good overview of the different components and their interactions.

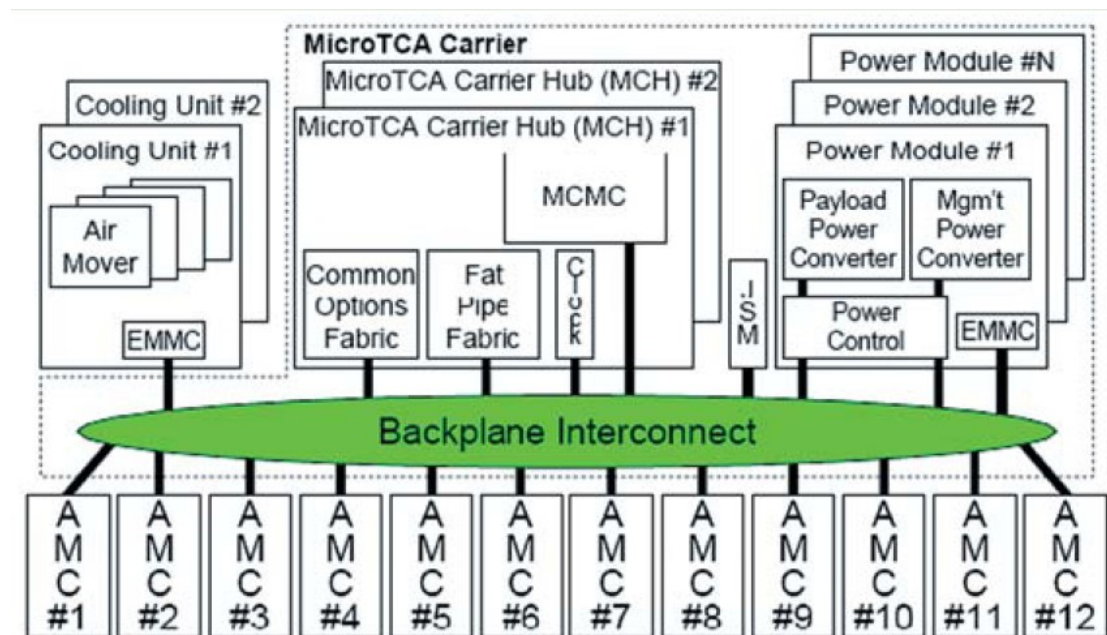


Figure 26: Block diagram of the different components composing a μ TCA crate [16]

Note that each FRU can feature a number of sensors (eg. temperature, voltage, noise) which can be used locally by the MMC to perform monitoring and alarming, but can also be transferred to the μ TCA Carrier Management Controller (MCMC), ie. the IPMI master, in order for this data to be inserted to the experiment-wide monitoring and control infrastructure.

2.5 Conclusion

After having explained the general concepts of a trigger and DAQ system for particle physics experiments in a first section, this chapter focused on CMS-specific features. A good understanding of this architecture is essential since the proposed upgrade of the muon spectrometer will have substantial implications on it. More specifically, the experiment wide muon system will undergo a major upgrade during the Long Shutdown 2, to include every channel in the trigger signal generation and the track reconstruction algorithms. The role of the Cathode Strip Chambers in this upgrade is essential, as these are currently the only detectors installed in the higher pseudo-rapidity region of the end-caps. Reducing the track mis-reconstructions and increasing the trigger efficiency of this sub-detector are thus of importance, as we will see later during this study. This has of course implications in the hardware implementation of the DAQ electronics, where no bottlenecks in the data stream are allowed. This is why the new electronics, for example, will rely on the new μ TCA standard, which was never used before inside the CMS experiment.

The next chapter will focus on the front-end part of the upgrade, and in particular the novel detection technology planned to be used for the forward muon spectrometer of CMS. This technology is the main topic of the Brussels R&D working group. Our choices to address the challenging specification are of course motivated by a number of physics goals and performance requirements, these are detailed in the next chapter as well. These considerations will set the base of our design process, as will be detailed in the second half of this work.

CHAPTER THREE

GEM detectors and the GE1/1 system



The successive luminosity increases scheduled with the coming upgrades of the LHC will have a significant impact on the number of collisions per bunch crossing. The muon system, for instance, is not designed to handle the foreseen particle rate. The main limitation comes from the Resistive Plate Chambers (RPC), especially in the $|\eta| > 1.6$ region of the end caps. Although this detector technology offers a time resolution of about a nanosecond, the recovery time after being hit doesn't allow rates beyond the kilohertz /cm² scale [9]. In 2010, a small collaboration was set up inside CMS to study the feasibility of using Gas Electron Multiplier (GEM)

detectors instead of RPC at higher pseudo-rapidity angles in the end caps. This technology offers the advantage of a much higher rate capability as we will see in this chapter, meeting thus the simulated collision rate requirements of up to 10 kHz per cm², as foreseen in the phase 2 of the LHC upgrades plan.

After a quick comparison of the different gas detectors used in CMS, this chapter will focus on the GEM technology, and in particular the Triple-GEM detector modules elaborated by the GEM collaboration. These modules are to be placed in some locations initially reserved for RPC detectors. This will be followed by a description of the expected physics performance of the system. Many simulations have been carried out to show the improvements brought by the proposed GE1/1 system to the muon trigger and track reconstruction sub-systems of CMS.

3.1 Muon system technologies comparison in CMS

3.1.1 Drift Tubes

Gas detectors offer the advantage of being relatively inexpensive per covered surface inside of an experiment. In this sense, drift tubes (DT) are ideal for very large surfaces thanks to their low level of complexity. A DT is a rectangular shaped chamber filled with a gas mixture specifically chosen to trigger electromagnetic interactions of muons ranging in the 1 GeV to 100 GeV/c momentum levels. An anode wire is stretched between two cathode strips, in the length dimension of the rectangle as shown in Figure 27.




Figure 27: Schematic view of a CMS drift chamber cell [9]

In CMS, each DT cell is 2.4 m long for a width of 42 mm and a height of 13 mm, and is filled with a mixture of 85% Argon and 15% CO₂ [9]. The 1.5 kV/cm electric field applied between the strips and the anode is enough to guide the electron-ion pairs produced by an incoming muon ionizing the gas mixture. Two electrode strips placed on the large edge of the chambers ensure that the drift path is squeezed enough to create avalanches near the anode wire. The amplification thus created is called the gas gain and is of the order of 10^5 in the DT cells of CMS. The time resolution for these modules is of the order of 3 ns, but due to the large dimensions of each cell, the main limitation comes from the 380 ns drift time [9], limiting effectively the rates capabilities. For this reason, DT are only used in the barrel region where the expected muon flux is not likely to exceed ~ 10 Hz/cm².

3.1.2 Cathode Strip Chambers

In the end caps, the cathode strip chambers are based on the principles of the proportional multi-wire chambers invented by Georges Charpak in the sixties. Only the shape and size are adapted to fit in the trapezoidal slots imposed by the layout of

an end cap of CMS. As we can see on Figure 28, a CSC module is built of about 1000 wires, each spaced by 3.2 mm in the case of CMS, tangentially stretched between two layers of copper strips planes. As for the DT, a high electrical potential is applied between the wires (anode) and the cathode strips. The electron-ion pairs created by the ionization of the gas when a muon crosses the chamber are pulled towards each of the electrodes.




Figure 28: Schematic view of a CSC (left), and explanation of detector operation principle (right) [9]

The main advantage of the CSC compared to the DT is a better spatial resolution. Since both the electrons and the ions are absorbed by the perpendicularly arranged wires and strips, reading out both planes enables an accurate muon track position measurement in two dimensions. Furthermore, in the case of CMS, each CSC module is built of a stack of seven strip panels and six wire planes. As a result of this construction, the spatial resolution ranges from 33 to 80 microns, depending on the pseudo-rapidity angle for which the strip is located on the trapezoidal module.

In order to produce a measurable number of electrons after the avalanche, the high voltage applied to the electrodes is set to 3.5 kV which produces a gas gain of the order of $7 \cdot 10^4$. At this operating point, the CSC sub-detector can withstand a rate of up to 2-3 kHz /cm² [9]. This is the reason why this technology is well suited for the higher $|\eta|$ regions of the end-caps. The disposition of the CSC is shown in Figure 29.

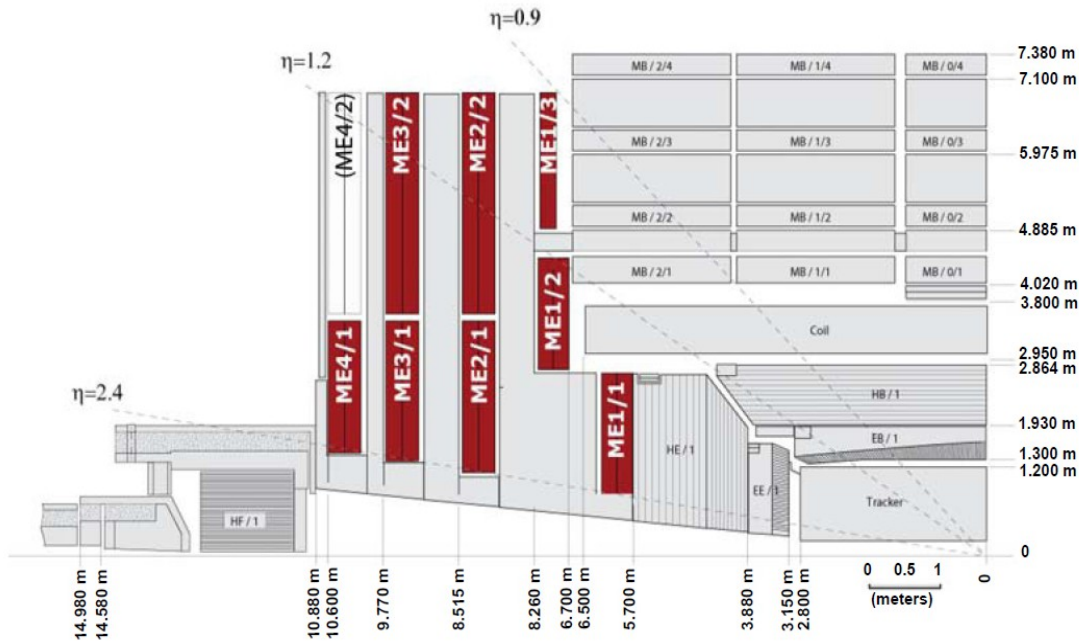


Figure 29: Disposition of the CSC modules in the end caps [9]

A word on the used nomenclature is maybe needed here. In the end-caps, four muon stations are present, called MEx (for Muon End-cap), x grows with the distance to the interaction point. Each station is divided in rings around the beam pipe, and is referenced by MEx/ y , with y growing radially. The first station is built of three rings (ME1/1, ME1/2 and ME1/3) whereas the other stations only contain two rings. Note that the ME4/2 are not populated with CSC modules during the early years of CMS operation [9].

3.1.3 Resistive Plate Chambers

Unlike the two previously mentioned detectors, Resistive Plate Chambers are not relying on a drift and a charge collection mechanism in the electrodes. Instead, very localized variations in the electric field between the electrodes are measured after the creation of a conductive channel between the electrodes. To achieve this, two highly resistive Bakelite plates coated on the outside with a conductive graphite-based paint are stacked on each other with a two millimeter gas gap in between. A high voltage potential is applied between the coating layers, creating an intense electric field of the order of 50 kV/cm in the narrow gas gap. The structure of a double RPC module, like the ones installed in CMS to increase the detection efficiency is shown in Figure 30.




Figure 30: Schematic view of a double RPC module [17]

The ionized gas channel created by a muon is sufficient to trigger the formation of a conductive channel, either in the form of a streamer (spark) or an avalanche if the gas mixture contains a high dielectric strength component such as sulfur hexafluoride (SF_6) [18]. The advantage of an RPC module operating in avalanche mode is a shorter recovery time (several hundreds of nanoseconds) due to the limited number of charges to be evacuated across the resistive Bakelite plates. Still, this recovery time remains the main limitation of the RPC technology, since the highest achievable rate doesn't exceed more than one kHz/cm^2 . In addition, because no predominant gas gain factor is present, the signal amplification is done by the read-out electronics itself, which adds up to the complexity of the front-end chips design. On the other hand, the independence from any drift time of the signal build-up inside the gas mixture provides an excellent time resolution, which is of the order of a nanosecond in the CMS RPC sub-detector.

3.2 The GEM detector technology

3.2.1 The GEM foil

The first studies and developments of the Gas Electron Multiplier (GEM) technology were made in the nineties at CERN, as a part of a broader research effort in micro-pattern gas detectors. The aim was to combine, in a cost effective way, the flexible spatial resolution given by a user-definable electrodes density and a reasonable time performance offered by short drift distances and fast recovery times in gas detectors. Electrodes formed by micro-patterns such as tiny strips, holes or grooves seem to be the ideal solution to match these performance requirements. The main challenge, however, is the production of large quantities of these fine-pitch patterns with a high level of quality assurance, especially on large surface detectors.

In the case of GEM foils, the chosen patterns are $70\ \mu\text{m}$ holes in a $50\ \mu\text{m}$ thin kapton foil, plated on each side by a $5\ \mu\text{m}$ copper layer. Each hole is distant from its neighbor by a distance of $140\ \mu\text{m}$, forming a honeycomb structure [19], as seen in Figure 31. Concerning the industrial production obstacles, these are mostly overcome by using modern printed circuit board (PCB) etching processes, namely photo-lithography and chemical abrasion.




Figure 31: Microscope view of the surface of a GEM foil [19].

Various gas mixtures exist to create the aimed gas gain of a single foil. Argon is commonly used in combination with CO_2 . Some CF_4 may be added to allow an increase of the electron drift velocity. The GEM foil is activated by applying a ~ 400 V potential between the two copper layers, creating the intense electric field (several kV/cm) required to trigger an avalanche reaction when an electron is trapped by the field lines of a hole. Figure 32 shows the shape of the field lines inside the holes of a GEM foil.




Figure 32: Representation of the electric field lines created inside a hole of a GEM foil [20].

A GEM detector, as explained in the next section, consists of a number of these GEM foils stacked up in a closed module containing the gas. A pair of electrodes on the top and bottom of the entire module provide the electric field required to initiate a drift of the electron and ions, similarly to DT or CSC. The multiple GEM foils create a gain factor, the drifting electrons are collected inside the anode strips of the drift field. Figure 33 shows the effective gain comparison of GEM modules built out of single, double and Triple-GEM foils. Limiting the high voltage supply not only simplifies the implementation, increasing the reliability accordingly, but it also reduces the likelihood of potentially destructive discharges happening between the copper layers.




Figure 33: Comparison between the Simple, Double and Triple-GEM detectors of the effective gain and discharge probability wrt. to applied voltage [21]

3.2.2 Triple-GEM detectors

To achieve the highest output signal level with a reasonable voltage supply per GEM foil, the Triple-GEM option was chosen for the CMS muon system upgrade in the end caps. The first gap between the drift cathode and the first GEM foil is called the Drift region. It is higher (3 mm in the case of CMS) in order to offer a longer energy deposition path ($-\langle \frac{dE}{dx} \rangle$) to the incoming muon and therefore a high detection efficiency (> 98 %). This is depicted in Figure 34.




Figure 34: Cutout view of a Triple-GEM detector, showing the amplification process [22]

The space between the two next layers are both called transfer gaps. This is where the electrons produced by the avalanche drift towards the next GEM foil. The last gap before the copper anode strips is called the collection or induction gap. Electrons drifting in this gap induce a signal on the anodes. The anode plane is usually made of conventional glass-reinforced epoxy laminate (FR4) printed circuit boards, produced at low cost. The only critical fact concerning the design of these anode planes is the length and impedance control required on the strips to avoid signal distortion before read-out.

An important point to mention here concerns the energy deposition of an incoming muon outside of the foreseen drift gap. The muon-gas interactions happening in the two transfer gaps, for instance, will produce exactly the same avalanche processes as if happened inside the legitimate drift gap. This is shown in Figure 35, which represents the electrons induced current on the anodes as a function of time. The vertical red lines show “in time” the position of the different gaps of the Triple-GEM structure. In this example, the gas mixture is Ar:CO₂ (70:30) with an electron drift velocity of 0.075 cm/ns




Figure 35: Simulation result on the anode strips of the GEM response to a perpendicularly incoming 1 GeV muon [23]. The different peaks correspond to muon-gas interactions in the successive gaps of the GEM module

This simulation signal was chosen to show the possible magnitude of an electron burst created in the transfer 1 gap of a Triple-GEM detector. Since this gap is closer to the anodes, the drift time is shorter, hence a signal appearing before the genuine drift gap signal. To overcome this problem, an integrating stage is present at the input of the read-out electronics. This will shape the signal to a uniform response of which we know the exact properties, allowing a fine time reconstruction based on the collected charge rather than the instantaneous signal magnitude.

A number of signal simulations have been performed to estimate the size and shape of the charge deposition on the anode plane [23]. The results are shown in Figure 36 for a Triple-GEM detector filled with a 70:30 Argon:CO₂ gas mixture.




Figure 36: Shape of the charge deposition on the anodes [23].

This plot shows the spatial distribution of the electron hits on the surface of the anode plane. The shape is Gaussian, with an RMS of the order of 0.016 cm in both the x and y directions. In a gas detector, the diffusion process of moving electrons is dependent of the drift length (L) in centimeters and the diffusion coefficient (D) according to this law:

$$\rho = D\sqrt{L} \tag{11}$$

which gives an RMS value of 0.02 cm for a drift length of 7 mm and a diffusion coefficient of $0.025 \sqrt{\text{cm}}$ (with a Ar:CO₂ mixture of 70:30 and an electric field of 3 kV)

3.2.3 The GE1/1 station

The GEM detectors to be placed during the coming upgrades of CMS have to fit in the slots originally foreseen for the RPC modules at a high pseudo-rapidity angle in the end caps. The two first locations concerned for the first upgrades are given the names GE1/1 and GE1/2 as shown in red in Figure 37 below.

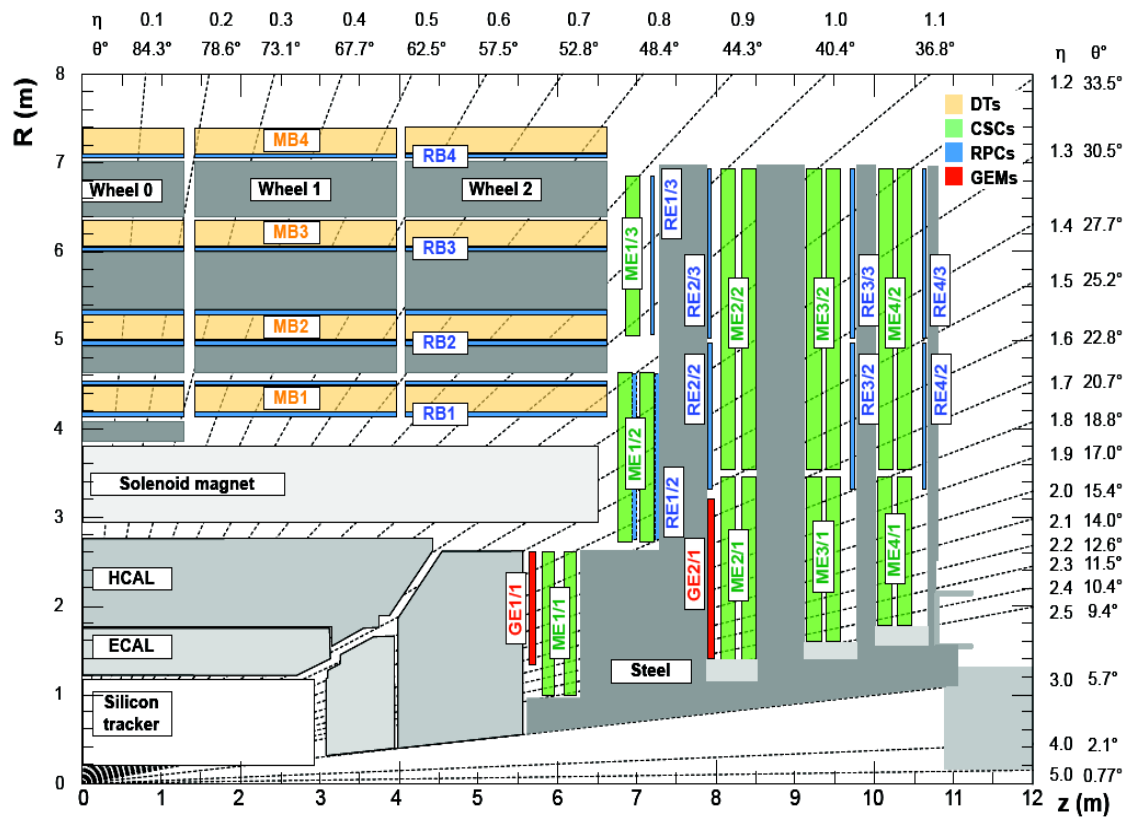


Figure 37: In red, the location of the first to be installed GEM detectors

Each GEM detector covers an Φ angle of 10° , explaining the trapezoidal shape of the modules, depicted in Figure 38.




Figure 38: Structure of a CMS GEM chamber

One of the main challenges during the design of these chambers was the integration of the high voltage DC/DC converters and the front-end read-out chips (called VFAT) between the gas tubing and the electronics water cooling, since GEM detectors are significantly thicker than RPC detectors due to the four drift gaps. The copper strips forming the anodes in the CMS Triple-GEM detectors have a typical pitch of 800 μm . Given the RMS of the Gaussian electron distribution after the drift, often only one strip is hit, and therefore the spatial resolution is given by the relation :

$$\sigma_{\phi} = \frac{\text{pitch}}{\sqrt{12}} \approx 230 \mu\text{m} \quad (12)$$

However, the spatial resolution can be increased to reach 100 μm by applying center-of-gravity search algorithms on clusters of hit strips, since we can simulate precisely the profile of a drifting electron burst. This topic is a work under progress inside the GEM for CMS collaboration. But this gives a fair estimate of the expected spatial resolution of the CMS GEM detector.


The main advantage of the GEM technology over RPC is the limited total charge collection time, not exceeding 100 ns in the case of the modules designed for CMS as we can see in the Table 3 for an Ar:CO₂ (70:30) gas mixture.

Gap	Timing
Collection	0 – 14 ns
Transfer 2	14 – 42 ns
Transfer 1	42 – 56 ns
Drift	56 – 98 ns

Table 3: Timing of the signals arriving on the anode strips and originating from the different gaps inside a Triple-GEM detector [24]

The resulting maximum achievable rate for this detector is thus of the order of 1 MHz/cm², which fulfills the requirements of the future luminosity upgrades of the LHC. The time resolution and efficiency are acceptable as well, as shown in the plots of Figure 39. These plots were obtained with measurements on a small 10x10 cm Triple-GEM prototype. We can see that the best achievable performance on this prototype is tending to a promising 5 ns in time resolution for a detection efficiency of 98% with Ar:CO₂:CF₄ (45:15:40) gas mixture. A gas mixture of Ar:CO₂ has a smaller electron drift velocity which degrades slightly the ultimate time resolution to $\sim 7\text{ns}$, which is still acceptable for CMS. Note that these time resolution measurements are not including any algorithmic improvements as will be explained further in this chapter.

Standard GEM Timing Performance



Standard GEM Efficiency

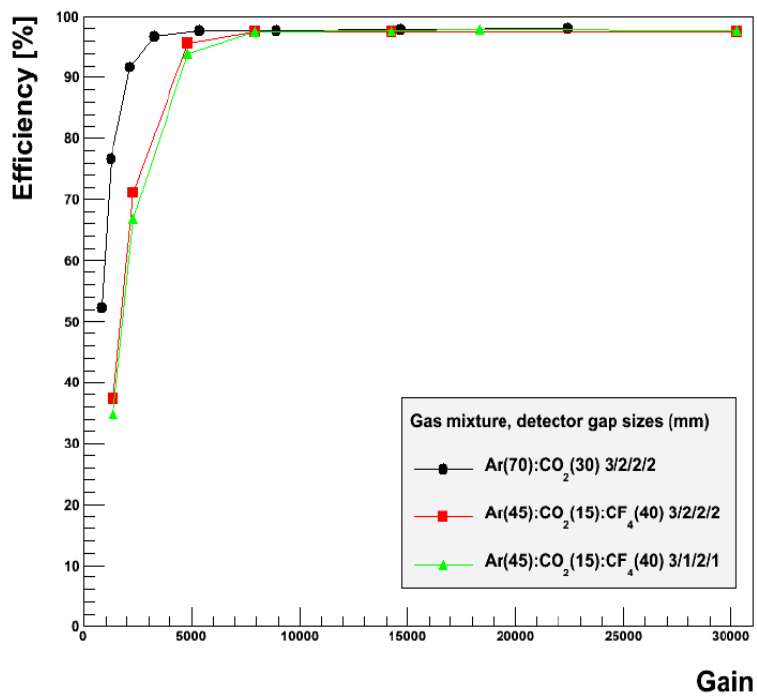


Figure 39: Measured performance of a Triple-GEM detector in terms of time resolution (top) and efficiency (bottom) for different gas mixtures [24]

3.3 Physics performance

3.3.1 Goals

The main goal of the GE1/1 system is to preserve the performance of the muon trigger system in the $1.6 < |\eta| < 2.2$ region for an LHC luminosity of $2 \times 10^{34} \text{ cm}^{-2} \text{ s}^{-1}$ after Long Shutdown 2. This is important on a physics perspective, since this area covers no less than a quarter of the entire CMS acceptance, as we saw on Figure 37. Currently no redundancy is available in this challenging region, as opposed to the RPC detectors backing up trigger and tracking capabilities of the Drift Tubes in the barrel and the Cathode Strip Chambers in the outer rings of the end-caps.

As we can see on Figure 40, with the current setup in the $|\eta| > 1.6$ region (CSC only) and for a luminosity of $2 \times 10^{34} \text{ cm}^{-2} \text{ s}^{-1}$, maintaining the cut on the transverse energy at $p_T > 15 \text{ GeV}$ will result in a trigger rate of 10 kHz, which is similar to the current single muon rate on the entire experiment.




Figure 40: Simulated trigger rate as a function of pseudo-rapidity [25]

According to this simulation, integrating a redundancy in the trigger system (green curve) under the form of a new muon station (MS1/1 in the plot) in this high particles rate and low magnetic field area, will temper the increase of the trigger rate by limiting the inaccuracies of the transverse momentum measurements.

In terms of physics, one of the main objectives of these upgrades is to study the Higgs particle more in details, focusing on less common decay channels than in the very first run of CMS. Some examples are: in the standard model Higgs sector, the channel $h \rightarrow \tau\tau \rightarrow \mu+X$ or more exotic processes such as resonant boson pair productions or electroweak baryogenesis or SUSY scenarios resulting in low momentum leptons. Keeping the transverse momentum threshold low on the muon trigger system increases the sensitivity to these processes [25]. On the tracking side, adding a redundancy to the CSC will maintain the standalone muon reconstruction capabilities well enough over the years as the detector is aging, which is otherwise impairing its performance. This is essential to initiate the study of new physics, especially in scenarios predicting new long living particles decaying in muon pairs.

3.3.2 Muon trigger performance

During LS2, the muon trigger system of CMS will undergo a major upgrade. It will include every muon sub-detector in its track momentum fit. The goal is to minimize the influence of the background signal, mainly composed of soft muons, on the trigger decision. GE1/1 will help with this in the higher eta region. Currently, the transverse momentum of an incoming muon is measured by associating the different stubs produced in the successive CSC chambers. The degrading influence of a parasitic scattering soft muon is thus high. Additionally, due to the curvature of the magnetic field lines as we move away from the solenoid, the first muon station provides the most significant input, as shown in Figure 41.




Figure 41: Azimuthal bending angle of a simulated 10 GeV muon wrt. normal vector to the CSC chamber [25]

The simulation of the azimuthal bending angle of a 10 GeV/c muon shows that the




Figure 43: Muon track segment reconstruction efficiency as a function of pseudo-rapidity [25]

This simulation shows that adding the redundancy on the CSC sub-detector increases the efficiency in the pseudo-rapidity area covered by the GE1/1, smoothing out completely the chambers overlap around $|\eta|=2.1$.

The overall trigger performance after LS1, when the LHC will operate at an instantaneous luminosity of $2 \times 10^{34} \text{cm}^{-2}\text{s}^{-1}$, is given in Figure 44. This simulation shows the benefits of adding the GE1/1 layer to the CSC stubs in order to compute the bending angle of the muons inside the end-caps. Reducing the overall trigger rate according to this plot will allow lowering the momentum thresholds and in turn increase the acceptance of rare physics signatures of the Higgs boson such as di-muon, tri-muon, muon + hadronic tau, etc. The complete description of the simulation parameter are described in [25].




Figure 44: Level-1 muon trigger rate as a function of the momentum

3.3.3 Muon reconstruction performance

CMS aims at a high reconstruction efficiency and a low misidentification rate of muon reconstruction for the next upgrades. The best way to achieve this is by keeping the matching windows as small as possible, even with an increasing number of tracks resulting from the luminosity upgrades. Here again, the bending of the tracks in the magnetic field are maximum at the output of the magnet. In addition, this is also where multiple scatterings are minimum. So this is where a first muon station would have its best place. This would increase the performance of the muon reconstruction performance for single track muons, but also for physics scenarios involving long lived particles, of which the identification depends on the quality of standalone muon reconstruction.

The proposed GE1/1 reconstruction system relies on the readout of the hit channels data from the front-end chips, to form clusters. An important point, which is addressed in this work, is thus to know the cluster size, since this will give the effective spatial resolution. A center of gravity algorithm computes the hit position, and is forwarded to the global muon reconstruction system of the event builder, which in turn will determine the exact momentum. Figure 45 shows the simulated hit resolution in the $R\phi$ -coordinates for two different pseudo-rapidity values.




Figure 45: Distribution of the difference between real and reconstructed hit [25]

The two different positions in pseudo-rapidity are respectively on the top and the bottom parts of the chamber. We see that the effective spatial resolution ranges between 0.029 cm and 0.051 cm for incoming muons of 200 GeV/c. This large range is due to the variable pitch in the radially distributed strips. The influence of the multiple scattering which dominates the spatial resolution degradation at low transverse momentum is shown in Figure 46.




Figure 46: Simulation of the influence of the muon momentum on the spatial resolution

The Figure shows the RMS of the multiple scattering displacement as a function of the muon transverse momentum for GE1/1 and all the other forward muon stations, evaluated at $\eta = 2.0$. The multiple scattering in GE1/1 is typically a factor 2 smaller than in ME2/1, showing once more the interest to add a detection layer at this location in CMS.


3.4 Conclusion

This chapter gave an overview of the innovative gas detector technology proposed for the upgrade of the end-cap muon spectrometer of CMS. The advantages of using GEM detectors compared to the existing technologies were detailed in the two first sections. Compared to drift tubes and resistive plate chambers, a higher particle rate can be reached with the triple-GEM technology. In addition, this technology provides an excellent redundancy for the existing CSC sub-detector in the high pseudo-rapidity region of the end-caps. The goals and performance improvements of installing this extra GE1/1 layer of detectors was detailed in a third section. The benefits on the trigger system was explained first, as the momentum measurement of the CSC chambers will become increasingly unreliable during the successive upgrades, mainly because of a degraded magnetic field in this region, and because of the small lever arm currently offered by the CSC system only. Secondly, for the muon track reconstruction system, the effective spatial resolution of the proposed GE1/1 system is meant to keep the matching window small for single track muons, increasing the reconstruction efficiency and moderating the misidentification rate.

The next step is to think and design the GE1/1 read-out system. This is the subject of the next chapter. A first section will list the specification of the DAQ system, and what was originally planned. This will give us the start point of the study which is, as we will see, mainly driven by technical constraints. Next, we will focus on every individual components of the acquisition chain, from the front-end to the off-detector processing electronics. This will be illustrated by the development of a downscaled version of the entire system, called the slice test, for installation and testing during the first coming technical stop in 2016.

CHAPTER FOUR

GE1/1 read-out system architecture



To perceive the technical challenges of integrating Triple-GEM detectors in the end-caps of CMS, a detailed view of all the components and their interactions is preferable. This is the aim of the current chapter. As we will see, numerous evolutions of the concept took place over the years, mainly thanks to the growing interest for the Triple-GEM technology amongst the CMS community. The Cathode Strip Chambers (CSC) sub-detector collaboration, for example, expressed its interest to benefit from the fast trigger capabilities of the GE1/1 installation, in order to improve their own timing resolution. From a Triple-GEM front-end electronics point of view, this is a major change in the architecture, as we will see in this chapter. To accommodate the successive architecture improvements, the required R&D effort grew to a point where it is not possible to install the entire system at once for the long shutdown 2 (LS2) in 2019. A first engineering step will thus be achieved during the technical stop of end 2016, in order to test the different developments made so far. This step is called the slice test, as it incorporates only a limited number of modules covering a small slice of the full GE1/1 geometry.

In order to understand the chosen components for the GE1/1 system to be installed during LS2, a number of technological considerations and limits need to be understood first. This is the subject of a first section in this chapter. The full LS2 system will then be described in details, from the front-end part to the off-detector electronics successively. Finally, a last section will describe the architecture of the slice test system and the work already achieved to meet the integration deadline of January 2017.

4.1 Constraints and requirements

4.1.1 Genesis of the project

Before the Brussels R&D group joined in 2011, the CMS GEM collaboration had focused on the physics improvements of installing triple-GEM detectors instead of RPC detectors in the high pseudo-rapidity region of the end-caps. Preliminary plans of the development of a front-end read-out chip were existing, these were in the form of two possible options, namely the gDSP and the VFAT3 chips, both sharing the development path of the digital part. The gDSP option included an analog to digital converter per channel, whereas the VFAT3 was a simpler binary hit counter. This is described in section 4.3 of this chapter. On the off-detector electronics side, the choice was to develop some trigger and tracking data concentrator boards, and reuse the RPC PAttern Comparator (PAC) electronics to generate a trigger signal for the CMS muon trigger. This solution is depicted in Figure 47.

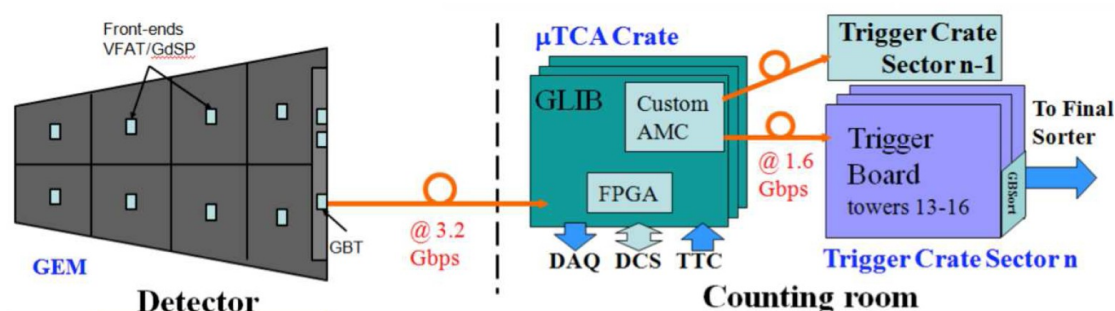


Figure 47: RPC PAC trigger based off-detector electronics

In this proposal, the concentration would already have relied on the μ TCA infrastructure. A number of CERN developed GLIB boards, described in section 4.5 of this chapter, would bring together the trigger and tracking data from each detector over the CERN designed GBT optical link. A custom mezzanine board sitting on top of these GLIB boards would translate the incoming trigger data to a format understood by the RPC PAC trigger boards.

At least four reasons made this proposal inadequate. First, the segmentation appeared not to fulfill the requirements, resulting in an increase of the number of front-end chips to three rows of eight chips. The trigger data rate, in turn, would render the single optical link insufficient. Secondly the architecture presents an inherent paradox : the bandwidth of the new μ TCA-based electronics would have to be bridled to fit the current VME-based RPC PAC electronics. Thirdly, informal discussions started already in 2011 within the CMS collaboration about the future upgrade plans. By the end of 2012, the orientation of these discussions clearly showed that the trigger system of CMS would undergo a major upgrade after 2020. Because of the rate increase, this proposed solution would quickly lack the required flexibility and performance to adapt to these upgrades. The last reason which gave the motivation to

redesign a complete new DAQ architecture came from the cathode strip chambers sub-detector collaboration in 2011. As explained in the previous chapter, providing the CSC Trigger system with a hit signal would increase the global muon trigger efficiency. This implied extracting trigger data from the DAQ chain of the Triple-GEM sub-detector and sending these to the CSC Trigger Mother Boards (TMB). The required latency to perform this transfer is of the order of only a few bunch crossing clock cycles, this is why the trigger data could only originate from the front-end chips directly, adding extra optical links to the design. This CSC trigger system integration is explained in details in section 4.1.4 of this chapter, with the solution of integrating an opto-hybrid data concentrator (described in section 4.3.7) at the bottom of each detector. It is a good example of the difficulty we had to come up with a final system proposal in this constantly moving target environment.

A number of other constraints and successively added requirements became visible as the project moved forward. The performance specifications refined little by little to reach the final list given in the next section. A number of mechanical constraints became problematic as the performance requirements pushed towards the use of more optical links and high voltage cables. This is, for example, why a GEM Electronics Board (GEB) needed to be designed (described in section 4.3.6 of this chapter). As more and more of these structural changes were decided to suit the successive evolutions of the upgrade calendar and the requirements, some more uncertainties were identified. For example, building large Triple-GEM detectors accommodating many front-end chips according to the decided segmentation may induce creation and fluctuation of a common mode over the readout strips. Additionally, using large time constant signal integrators to collect the charges of the GEM process induces a degradation of the time resolution of the detectors. These two problematics were also addressed in section 4.3.

4.1.2 Performance requirements

Since the Triple-GEM detectors are foreseen to fill the vacant RPC module slots in the end-caps, their level of performance is strictly imposed as better. The resulting list of requirements, extracted from the latest Technical Design Review (under approval process at the time of writing) [25] is given below :

- Maximum geometric acceptance within the given CMS envelope.
- Rate capability of 10 kHz/cm² or better.
- Single-chamber efficiency of 97% or better for detecting minimum ionizing particles.
- Angular resolution at trigger level of 300 urad or better in the azimuthal direction.
- Timing resolution of 10 ns or better for a single chamber.
- Gain uniformity of 15% or better across a chamber and between chambers.
- No gain loss due to aging effects after 200 mC/cm² of integrated charge.

As shown in Section 3.2.3, Triple-GEM detectors easily reach the efficiency plateau of 97% with both Ar:CO₂ and Ar:CO₂:CF₄ gas mixtures and with a safety margin. In the GE1/1, the detector efficiency will still be improved by collating two Triple-GEM chamber back to back in a so-called super-module. Combining the individual efficiency of both chambers with a logic OR can increase the super-module efficiency to a level of 99.9 % or more. In addition to the gain in spatial resolution, the super-module structure also increases the timing resolution to a far better level than the imposed 10 ns, since the timing information of the double Triple-GEM detectors can be processed independently.

The spatial resolution of 300 μrad or better in the azimuthal direction, imposed to reach the trigger performance described in section 3.3.2, is the baseline to calculate the pitch of the read-out strips. Considering a Gaussian distribution of the electrons over the anode strips and a binary read-out electronics, the resulting resolution is $300\mu\text{rad} \cdot \sqrt{12} = 1040\mu\text{rad}$. This corresponds to a resolution of 0.8 mm in the azimuthal direction or a pitch of 2.7 mm, at the outer radius of the GE1/1 chambers, which is 2.6 meters away from the beam line. Consequently, each Triple-GEM detector is split into three sectors in Φ , each read out by 128 anode strips. At the outer radius of the GE1/1 chamber, this results in a pitch of 1.2 mm. At the trigger level, the strips are ‘OR’ed by groups of 2 adjacent strips resulting in a pitch well below the requested 2.7 mm. The segmentation in η is given by other technological constraints described at the end of this section.

4.1.3 Mechanical considerations

A number of constraints regarding the size and the shape of the Triple-GEM modules had to be taken into account to ensure the feasibility of the design of the GE1/1 system. First, as we already explained, the modules have to fit in the physical locations initially foreseen for RPC modules, each covering 10° in Φ and a pseudo-rapidity ranging from $1.55 < |\eta| < 2.18$ in the end-caps. The general geometry is thus trapezoidal, with a shorter edge of 22 cm, a long edge of 45.5 cm and a height of 99 cm. Figure 48 shows a picture of an installation slot for a Triple-GEM detector. The red box indicates the narrow space available for each detector.




Figure 48: View of an installation slot for a GE1/1 Triple-GEM detector

Beyond these purely mechanical constraints, the biggest challenge to overcome was a conceptual problem, namely the stretching of thin foils over a large surface. In the case of the Triple-GEM detector modules foreseen for CMS, the gap configuration is 3/1/2/1 [25], which means that two foils carrying a high voltage difference are distant of only three millimeters at most over a 0.25 m² surface. To guarantee the uniformity of the electric field without using spacers, a sophisticated stretching technique was required. The solution was found under the form of an embedded nut holding a free sliding stretcher, which in turn is clamping the three GEM foils. The structure is shown in Figure 49.




Figure 49: Schematic view of the GEM foil stretching technique [25]

4.1.4 Integration in the CSC trigger logic

Reaching the operation limits of a certain detector technology as the luminosity of the LHC will increase is not only a problem for the RPC modules. As we have seen in Chapter 3.3, the Cathode Strip Chambers also have their specific limitations, mainly under the form of ghost interactions when two muons hit the detector at the same moment. The perpendicular wires and strips collecting the electron and the ion bursts create an ambiguity on the position of the real hits, creating the so-called ghost hits. This problem will become more important after LS2, and this is why the CSC sub-detector collaboration expressed its interest to join the discussion table. Including the Triple-GEM trigger data to the CSC local trigger system could be a solution to remove this ambiguity leading to the ghost hits.

The CSC trigger path integration led to a major alteration of the original design. Initially, only three optical Gigabit transceivers (GBT) were foreseen per Triple-GEM detector, as shown in Figure 50. These were foreseen to send the trigger and the tracking data only to the off-detector processing electronics over optical fibers. By design, the latency required to transfer the trigger data from the Triple-GEM detector to the CSC electronics in order to complete its trigger information has to be short (<20 LHC clock cycles). To make the data path as short as possible, per this requirement, a direct optical link is preferable between the Triple-GEM front-end electronics and the CSC trigger system entry point, named the Trigger Mother Board (TMB). Note that it already takes ~ 20 LHC clock cycles to transmit the data over optical fibers from the detector to the service cavern, without any data processing.

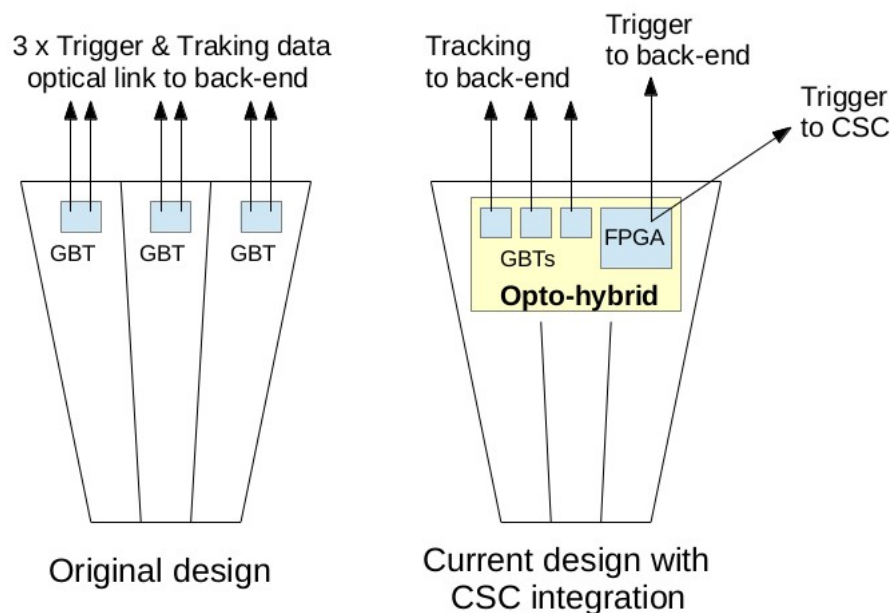


Figure 50: Evolution of the design before (left) and after (right) the integration of the CSC trigger path

The main change in the general architecture is the number and the destinations of the

trigger data paths. In the latest design, the trigger data of the entire GE1/1 module is concatenated inside an additional FPGA. This is possible because the volume of trigger data to be transferred is low. Only the latency is critical on this link. From this FPGA, one single optical link is foreseen for transferring the trigger data to the back-end electronics. And in addition to this, a second optical link with an exact copy of the trigger data is leaving towards the CSC TMB for their own use.

To build this architecture, a major constraint needs to be addressed. The design and the use of an additional opto-hybrid board, as depicted in Figure 50, containing an FPGA is tricky in this area of the detector. This is mainly because of the ambient radiation level. This is the reason why, during the initial phase of the design, the collaboration choose to use the Gigabit Transceiver (GBT) chip set designed by CERN, because it includes a high level of radiation tolerance in its specifications. This chip set is specifically designed to offer a fast and reliable data transmission solution for experiments where the radiation levels exceed the ratings of commercially available products. This is not the case of the components present on the opto-hybrid, and especially not the power converters and the FPGA. There is however, no other choice than using an opto-hybrid board, since the TMB is only compatible with an older optical transmission protocol, called the Gigabit Optical Link (GOL). This is an earlier design from the CERN microelectronics group as well, and incompatible with the GBT protocol. An FPGA is thus required to translate the trigger data from the VFAT3 chips towards the CSC trigger system.

4.1.5 Technical limitations

Several other technical limitations have modeled the design over the years. The transmission over optical links, for example, with the choice of the GBT chip set, gave the final segmentation in η . This chip set is featuring a collection of ten E-port links from which data can be concentrated to be sent over a bidirectional optical cable at the payload data rate of 3.2 Gbit/s. The final segmentation containing three slices in Φ will thus never exceed $3 \times 10 = 30$ front-end chips (VFAT3).

Another limitation concerning the overall architecture of the data read-out electronics is the interface to the rest of the central DAQ system of CMS. Two elements are essential in this interface. First, the data path used to send and receive the trigger and tracking data. These links are optical fibers requiring a well specified transmission rate and data format. Secondly, the CMS wide clocking scheme used to transmit the bunch crossing synchronization signal as well as the Level-1 Accept trigger signals and some fast control commands. Here again, the specifications impose very strict latency constraints on this communication channel. These two data paths are constraining on the overall architecture because they have to be distributed in a latency controlled fashion over the entire read-out electronics of the Triple-GEM project.

4.2 System overview for LS2

The final segmentation of the Triple-GEM detectors chosen for Long Shutdown 2 is set to three columns of eight VFAT3 chips, covering the entire surface with the required spatial resolution. This is illustrated in Figure 51. To fit these 24 front-end chips onto a Triple-GEM module, a special printed circuit board (PCB) needs to be designed in order to avoid the use of cables. This board, called the GEM Electronics Board (GEB) is roughly of the same size as a Triple-GEM module, and is fitted with connectors for the read-out strips on one side, and connector for the multiple VFAT3 chips on the opposite side. In addition, according to the plan of providing trigger data to the CSC detectors, a last connector on the large edge of the GEB is present to accommodate the opto-hybrid board. Instead of cables, the PCB is designed with impedance controlled copper lines converging to the opto-hybrid connector. Furthermore, the GEB ensures that the power network for the VFAT3 chips is uniformly spread over the surface, since the PCB is designed with multiple inner copper layers.

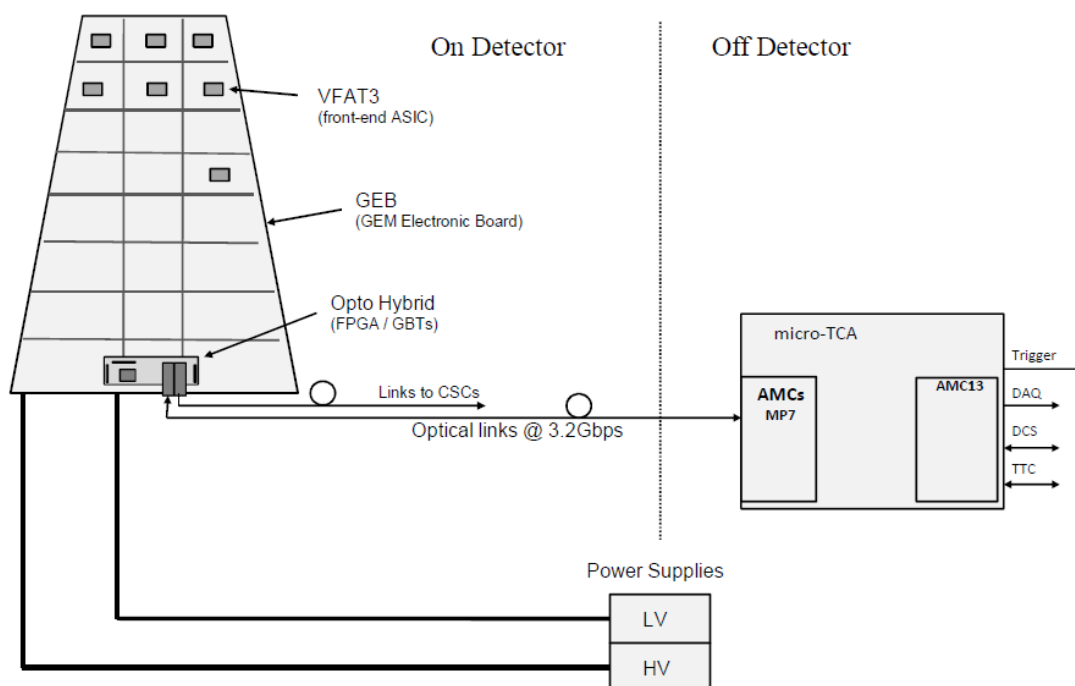


Figure 51: Overview of the entire GE1/1 system as planned for installation during LS2

The next component in the DAQ chain is the opto-hybrid board. It is fixed to the GEB board with a high density connector to fit all the communication lines coming from the VFAT3 chips. This board includes the GBT chip sets which are providing high bandwidth bidirectional optical link facilities for data transfer to the off-detector electronics. A large scale FPGA is also present on this board, to translate the trigger information towards the GOL protocol for shipment to the CSC trigger system.

On the other side of the optical link to the off-detector electronics, is the μ TCA architecture based read-out instrumentation. These twelve-slot crates are stacked inside the USC55 service cavern and contain a collection of MP7 (Master Processor) [26] read-out boards concentrating the incoming optical fibers. In order to shorten the design phase and optimize the developments within the small collaboration, it was decided to maximize the use of existing components. This led to the choice of this specific read-out board, developed by the Imperial College of London. It can accommodate up to 72 optical links and offers a cutting edge Xilinx Virtex-7 family FPGA for data processing at a data rate of up to 10 Gigabits per second. With these remarkable I/O and logic densities, only eight MP7 boards could be enough to handle the entire GE1/1 detector.

As we have seen in Chapter Two, the μ TCA standard crates used in CMS have a dual star topology, which means that two redundant sites exist for the central crate controllers (MCH). In the case of CMS, the choice was made to give up the failsafe feature offered by the redundancy, and use the second MCH site for a central data collection board, called the AMC13. This board is the main interface to the central DAQ system of CMS. It is specifically designed to receive and distribute the Trigger Throttling & Clock (TTC) signals, as well as to concentrate and forward the tracking data and slow control (DCS) over dedicated optical fibers. This board was designed by the CMS group of the Boston University, as a standard μ TCA interface to the central DAQ of CMS, and matched perfectly the needs of the GEM upgrade project. Reusing this development was an effective way to gain some R&D effort as well.

A number of auxiliaries are also part of the design, namely the power distribution network and the gas system. The power generation is depicted in Figure 51 as two distinct high and low voltage supplies, also located in the service cavern. The high voltage (15 kV @ 1 mA) is required for powering the resistive dividers producing the different voltage potentials inside the gaps of the Triple-GEM modules. The low voltage circuit is required for powering the front-end chips and the opto-hybrid board. Concerning the gas distribution, this is very similar to the CSC sub-detector, as the gas mixture is similar. A mix of Argon, CO₂ and CF₄, at the proportions of 45 %, 15 % and 40 % is brewed inside the surface building next to the gas cylinders storage lot, then pumped towards the service cavern and finally to the distribution racks in the balconies. This represents about 300 meters of copper tubing. The choice of using copper is imposed by the presence of CF₄, a highly hydrophilic gas, which otherwise could create hydrofluoric acid and damage the modules.

4.3 Front-end electronics

4.3.1 Specifications

The front-end chip is in charge of reading out the electron collection on the strips of the anode plane. Here is a list of the main specifications as set and approved by the CMS collaboration [25]:

- 128 channel chip
- Read positive and negative charge from the sensor (negative for Triple-GEM detectors)
- Provide tracking and trigger information
- Trigger information : Minimum fixed latency with granularity of 2 channels
- Tracking information : Full granularity after Level-1 Accept.
- Level-1 Accept capability: Level-1 Accept latency up to 20 μ s and Level-1 Accept maximum rate of 1 MHz
- Time resolution of less than 7.5 ns (with detector)
- Integrated calibration and monitoring functions
- Interface to and from the GBT at 320 Mbit/s
- Radiation resistant up to 100MRads (up to 1MRad needed for the muon application)
- Robust against single event effects

The R&D working group in elementary particles of the Inter-university Institute for High Energies in Brussels joined the GEM collaboration in year 2011. In these early days, two options were retained for the design of the front-end read-out chips. First option was to reuse and improve the design of the VFAT2 chip [27]. This design is used in the TOTEM experiment of the LHC, to read out 128 channels of various detector technologies, including GEM chambers. This is a binary chip, only providing an information of hit channels at each bunch crossing. The second option was a more ambitious mixed signal chip called gDSP, including an analog to digital converter (ADC) for each of the 128 input channels, coupled to some digital signal processing cores to perform a very first stage of data reduction and quality check before shipping towards the central DAQ systems of CMS. To allow an efficient design process, both architectures were studied thoroughly in parallel and are described in this section. They had blocs in common. These were studied independently of which design would be chosen at the end. The analog input stage, for example, was developed independently of the rest of the chip.

4.3.2 gDSP option

The gDSP (for GEM Digital Signal Processor) project was an attempt to realize a common dream of the ideal DAQ electronics for particle physics. It featured a high channel density with up to 128 inputs, an accurate full swing signal conversion for each channel provided by 128 parallel ADC blocs, a high speed serialized bidirectional communication link for data retrieval and chip programmability, and a low power consumption and dissipation as well as a small footprint despite the chip complexity and the radiation tolerance. The block diagram is shown in Figure 52 and will be referred to all along this section.

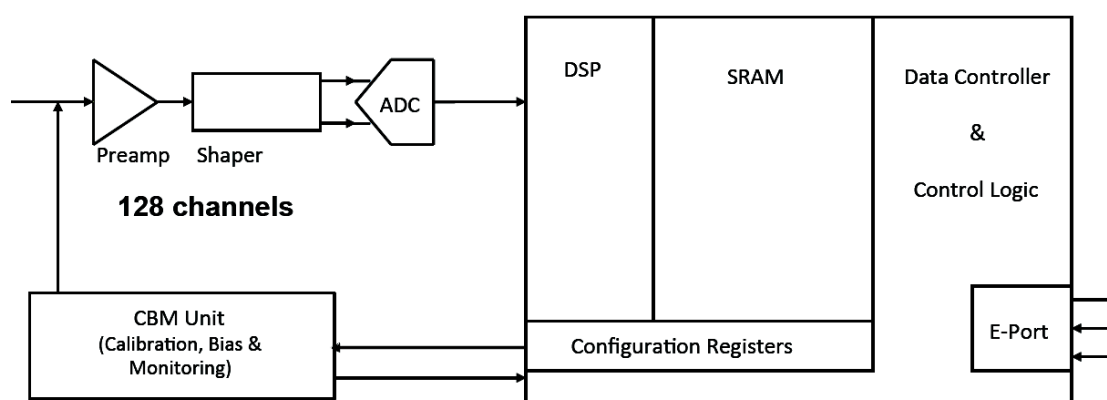


Figure 52: Architecture of the gDSP front-end [28]

This architecture shows a number of standard blocs for on-detector front-end chips. On the left, a number of analog channel signal processing circuits, composed of a low noise preamplifier and a shaper (integrator) are counting the amount of charge transferred from the detector strips. An ADC digitizes the signal and transfers the data to a number of digital signal processing functions. A couple of Static Random Access Memory (SRAM) arrays store the values in a circular buffer type memory, and a controller handles the data retrieval and sending over a communication interface called E-Port in the case of this chip.

The initial study of this architecture was launched by Paul Aspell at CERN on the promising outlooks offered by Figure 53.




Figure 53: 2011 state-of-the-art of the ADC power consumption as a function of the bit resolution [29]

In this plot we see the evolution of the power, normalized with respect to the conversion frequency, of different documented ADC designs, sorted per Equivalent Number Of Bit (ENOB) level. Let's remind that the ENOB is obtained by the following relation:

$$ENOB = \frac{SNDR - 1.76}{6.02} \quad (13)$$

The SNDR is the signal on noise and distortion rate, the 1.76 comes from $10 \log_{10}(3/2)$ which is the quantization error of an ideal ADC, and 6.02 converts decibels into to bits with $20 \log_{10} 2$. According to this plot, the state-of-the-art of ADC design is moving towards an efficiency, also called Figure of Merit (FOM), of 100 fJ per conversion. For a 9-bit ADC clocked at 40 MHz, this corresponds to a power consumption of:

$$Power = FOM \cdot 2 \cdot f_s \cdot 2^{ENOB} \approx 4 \text{ mW per channel} \quad (14)$$

Although this is a lot for a chip, it looks possible, since the biggest part of the power budget of such a mixed signal chip is always given to the signal conversion part. The overall would not exceed the 800 mW for the 128 channels. Unfortunately, the main technology used for designing chips at CERN, namely the 0.13 μm node from a world class foundry is not able to provide the low power consumption described in Figure 53. The FOM described in the articles of the study are reached with deeper sub-micron technologies, below the 90 nm, to which CERN has no access at a reasonable

cost in 2012. This is a first reason why the design of this mixed signal chip stopped in 2012, the second reason being the non-negligible probability that no finalized ADC would be ready before LS2.

The second main feature of the gDSP option was the very early data processing functions implemented on-chip, to perform a first stage of data reduction and correction before being sent out. Four main stages were foreseen in the original design:

- Baseline Correction Filter 1 (BC1), which remove systematic artifacts coming from the electronics for example
- Tail cancellation (TC), which compensates distortions induced by the analog front end
- Baseline Correction Filter 2 (BC2), which cancels low frequency variations of the baseline
- Zero suppression (ZS), which reduces the data to be stored when the programmable threshold of a valid hit is not exceeded

The reason why these four functions were chosen is that they seem essential for a mixed signal front-end chip and were already successfully implemented in another front-end chip, the S-Altro, designed in the same group at CERN for another type of gas detector [30].

In addition to these four general purpose digital signal processing functions, the addition of several other blocs was contemplated, more specifically designed for large GEM detectors. First, a Time Over Threshold (TOT) block between BCF2 and ZC, to improve the timing resolution by compensating the time-walk effect, and a common mode rejection algorithm between BCF1 and TC, to cancel the effects on the strips of fluctuating electromagnetic fields. These two blocs are detailed later on in this chapter. In addition, a last block was under consideration just before the project was canceled, namely a cluster reconstruction algorithm to quickly find a center-of-gravity spot when multiple channels are hit, improving the spatial resolution consequently.

4.3.3 VFAT3 option

VFAT3 is a more realistic design, given the current state-of-the-art of mixed signal integrated circuits design. Furthermore, as it is an evolution of an existing VFAT2 front-end chip, some concepts and elements can be reused. This choice will lead to a considerable time gain in the design process which is not to be underestimated knowing the tight schedule imposed by the LHC upgrades calendar. Figure 54 shows a block diagram of the VFAT3 architecture.

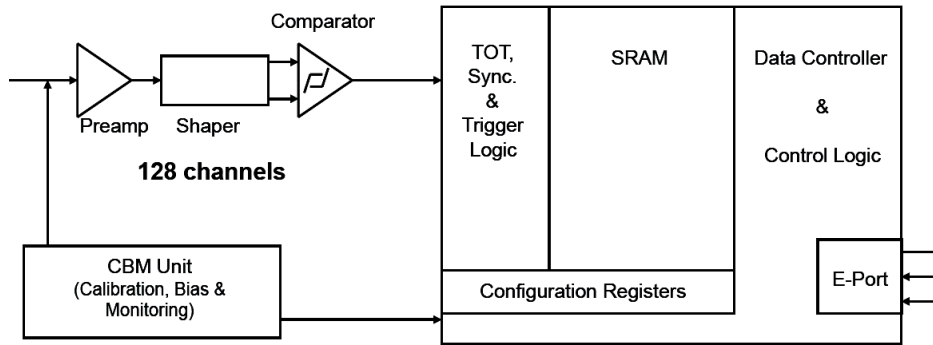


Figure 54: Architecture of the VFAT3 front-end [28]

As mentioned earlier, the analog input stage, composed of a preamplifier and charge sensitive amplifier (third order CRRC shaper) are in common with the gDSP design. But instead of feeding a per-channel ADC, a monostable comparator outputs a binary hit flag when the collected charge inside the shaper reaches a programmable threshold. The monostable is reset only when the input level drops below the same threshold at a rising clock signal.

The parallel trigger and tracking data path described in section 2.2 and 2.3 is taking its origin at the comparators output. The 128-channel data with its corresponding bunch crossing number is stored in a circular buffer memory, depicted as SRAM1 in the VFAT3 block diagram of Figure 55 below. At the same time, a fast OR operation between clusters of 8 channels is computed inside the low latency Trigger Unit of the chip, and the result is sent to the muon trigger system of CMS. If a level-1 trigger is issued by the global trigger system, this data is transferred to a smaller event memory, referred to as SRAM2 in the diagram. The transfer of these events over the CERN-specific E-Port links is initiated later by the read-out out electronics described in Chapter Four.

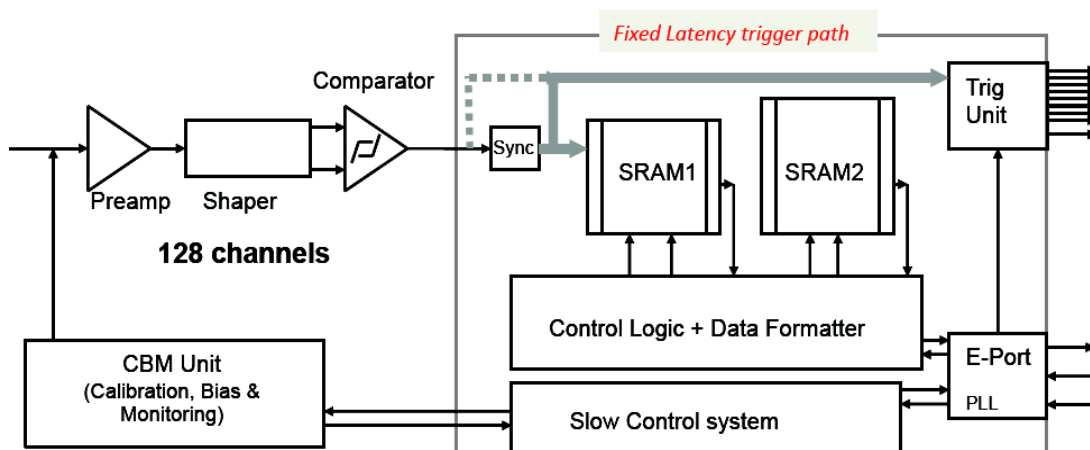


Figure 55: VFAT3 block diagram [25]

The Slow Control system, linked to the Calibration, Bias and Monitoring (CBM) unit provides an access to the analog input circuitry, in order to configure its gain and the peaking time of the shaper. This block will allow an operator to tune the efficiency of the GEM detectors during machine development runs.

4.3.4 Common mode reduction algorithm

It is common in the industry to see interference problems appearing between different elements of a system as it grows in size and complexity. Good care must always be taken during the definition of the specifications of the different sub-systems and their integration plan. This statement counts for large scale particle physics experiments as well. The entire muon system of CMS is composed of many individual gas detectors, all including a high voltage component. In the case of a GEM detector for example, long copper strips carrying charges inside an electromagnetic environment can generate voltage fluctuations on the front-end electronics input. Fortunately, these field effects are local to every GEM chamber and the generated perturbation will most likely affect all the channels of a front-end chip the same way and at the same time. The observable result will thus be a fluctuating common mode signal. As the design of the analog inputs is already very complex due to the gain and the shaper peaking time configuration possibility, it was decided to perform the common mode rejection digitally in the case of the gDSP option. The algorithm studied during the current work is described in this section.

The general idea of a digital common mode rejection algorithm is shown in Figure 56. A first unit reads out the data of the 128 channels and based on this data, evaluates the actual level of common mode on the channels. Knowing this value, a subtraction is performed on all the channels. The easiest way to evaluate the common mode level present at the input is to calculate the mean value over the channels measured by each ADC when no hit is recorded on any of the channels.

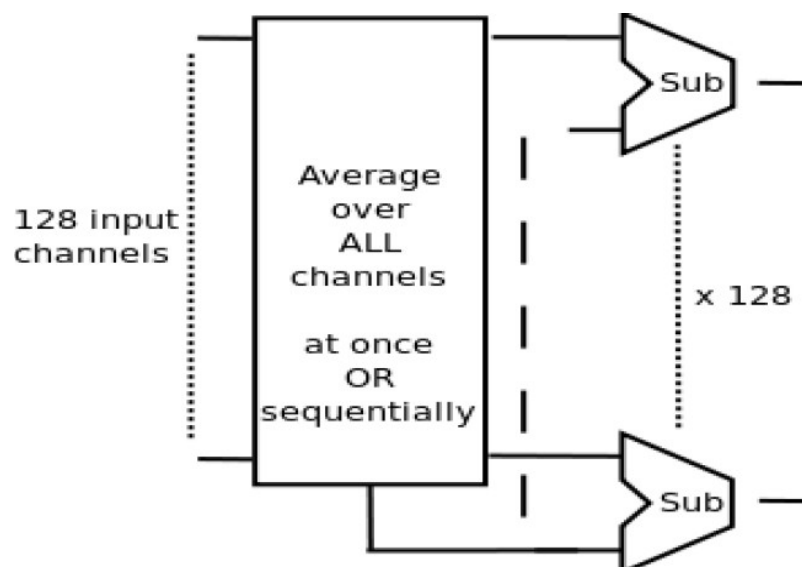


Figure 56: Principle of a common mode rejection unit

This mean value calculation is the most resource intensive part of the unit. Performing the operation on all the channels at the same time involves generating a 128-input adder. Each input being a 9-bit value, a total of 1'143 one-bit adders have to be inferred, not counting each individual 9-bit subtracters needed per channel. To meet the power budget for the digital part of this chip as well as to save some die space, the solution was imagined to pipeline the channel processing with a Moore-type finite state machine, shown in Figure 57

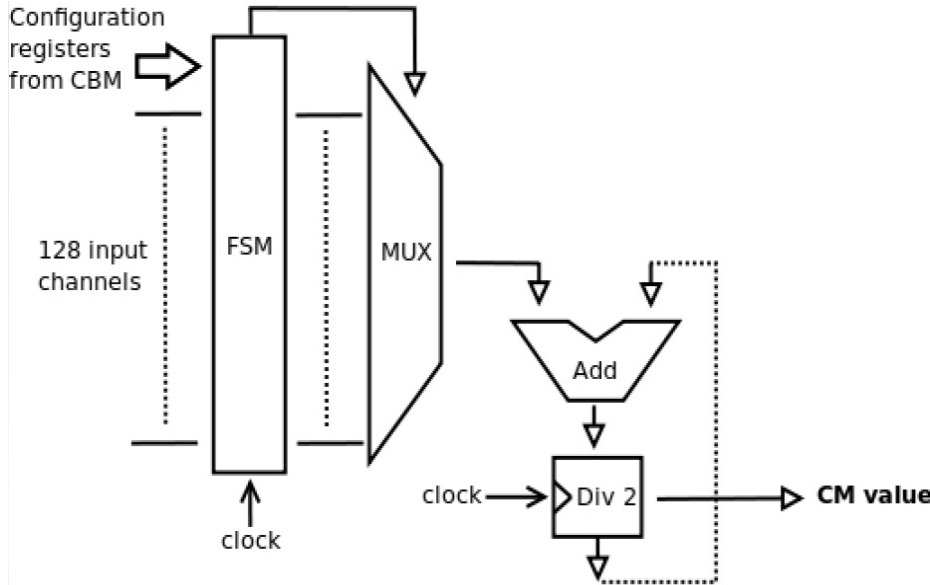


Figure 57: Block diagram of a pipelined common mode rejection unit

Sequentially, each channel value is added to the previously calculated average and divided by two. After a short learning sequence, and provided any hit channel is excluded from the calculation by means of a threshold for example, an accurate common mode measurement is obtained. Some system level simulations were performed to validate the concept. The input to the model is a Gaussian distributed hit, projected on an anode plane of 128 channels, and spread over 1024 samples in time. The general equation of this signal is of the general form :

$$y(x,t) = A \cdot \left[\sum_{\text{samples}} \sum_{\text{channels}} e^{-\frac{(x-64)^2}{2\sigma_x^2}} \cdot e^{-\frac{(t-512)^2}{2\sigma_t^2}} \right] + CM(t) + n(t) \quad (15)$$

The parameters of the Gaussian distribution are borrowed from the GEM signal simulations of Chapter 3.2.2, and the two parasitic parameters, the common mode function $CM(t)$ and the noise $n(t)$ are the variables to be eliminated by our algorithm. Below is a representation of a hit in the middle of the 128 channels, to which a large static common mode component was added, as well as a common mode fluctuation on top of it.





Figure 58: Result of the common mode subtraction on a hit + $CM(t)$ signal

As we can see, the static common mode component as well as the time variations disappeared completely, after an initial learning time of a few samples. For the sake of testing the robustness of the algorithm, it was tested on a signal where not only a time

variation was present in the common mode, but also a spatial component and a high level of noise. The results are shown below :

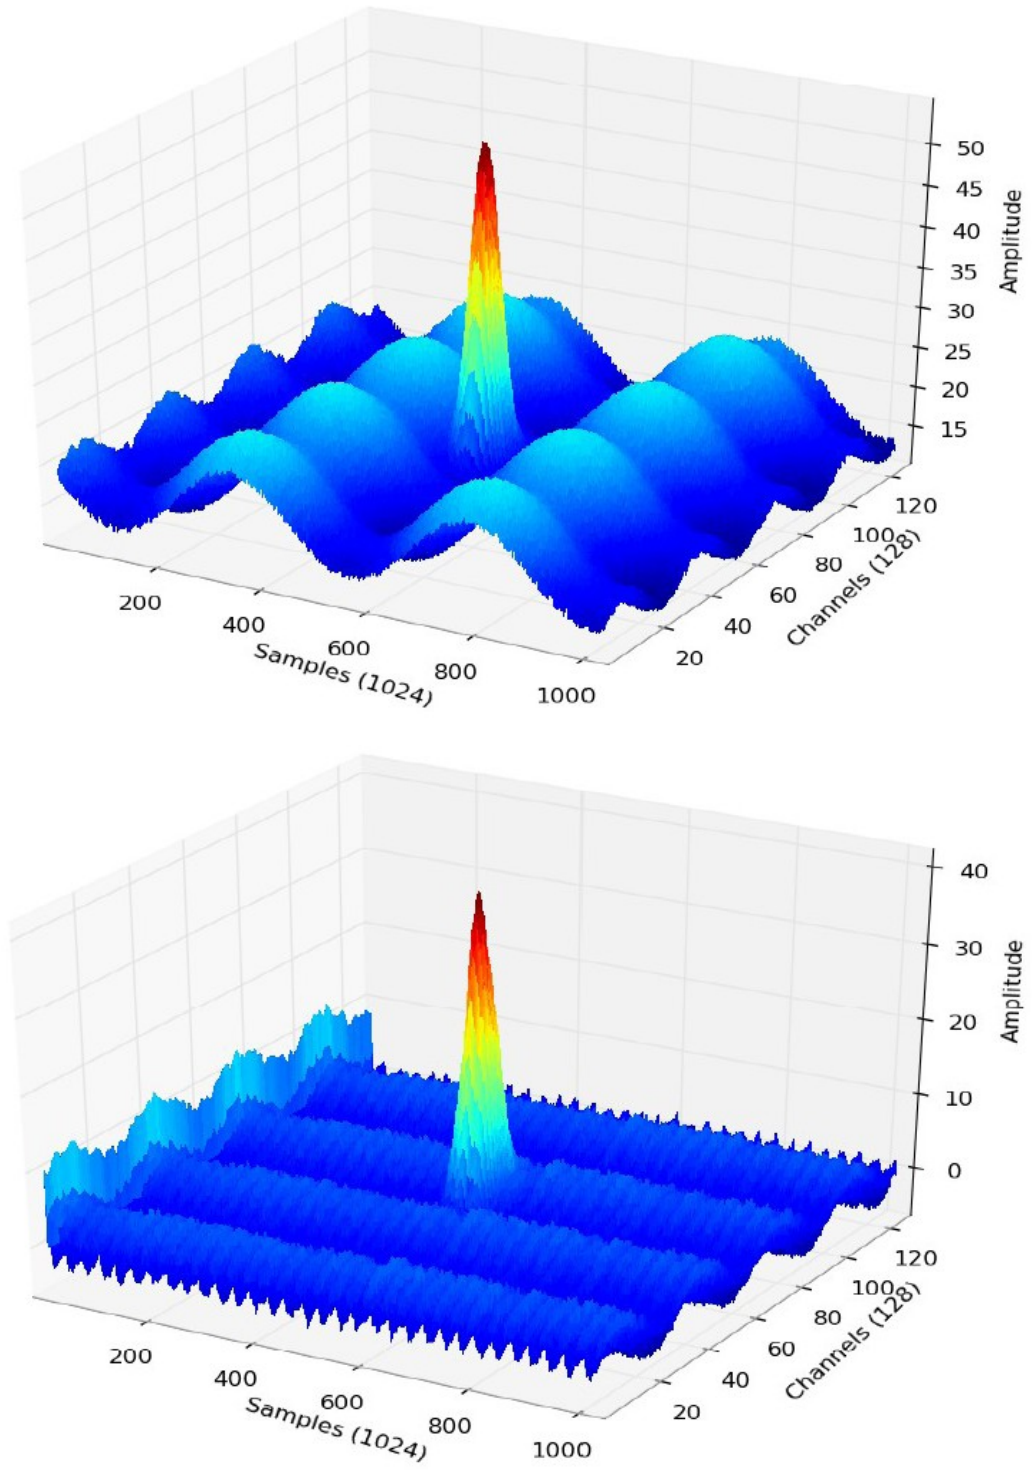


Figure 59: Result of the common mode subtraction algorithm on a hit + $CM(t) + CM(x) + \text{noise}$

Even on the most challenging signal, the residual variations are negligible compared to the input, and the static component of the common mode was always removed successfully.

After the decision was taken to focus on the VFAT3 architecture, this study was abandoned, since eliminating common mode is as easy as tuning the comparator threshold in a binary chip.

4.3.5 Time resolution improvement

The Triple-GEM detector sub-system will provide an input to the CMS muon trigger system over the dedicated trigger data path. It is thus of importance to have the best possible time resolution and the electronics can play a role to gain some precious nanoseconds. This detector performance improvement is described in this section.

As we can see in Figure 60, the magnitude of the signal can be reconstructed by counting the time between two successive threshold crossings in the case of a VFAT3 scenario.

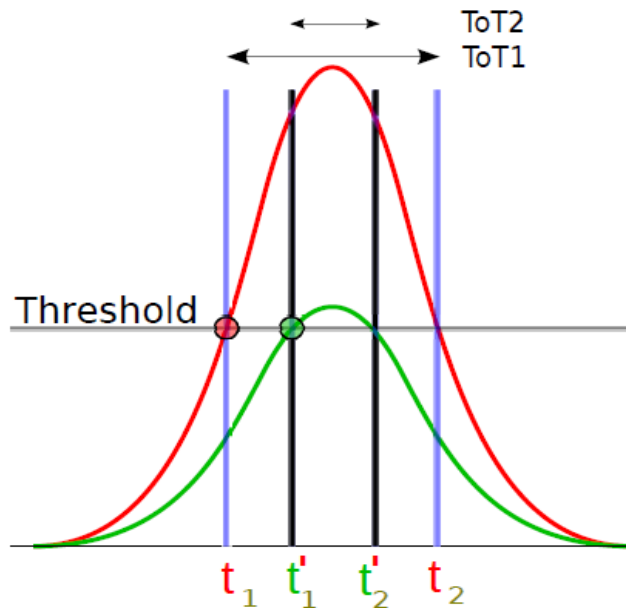


Figure 60: Time over threshold based magnitude measurement

An uncertainty arises from this measurement, however, on the time resolution, due to the time-walk difference between high and low amplitude signals (noted t_1 and t_1' in Figure 60). As a 320 MHz clock signal is present on the chip for the serial communication interface, a period of 3.12 ns is too large to precisely measure the time-walk, leading to an imprecision on the signal start time. The solution to this problem is called the Time over Threshold (ToT) method, developed during the

current work. Knowing the exact signal shape from simulations performed beforehand, it is possible to associate a time over threshold (ToT) count to a signal magnitude value, and thus a time-walk value. This can be done in a very fast way and on the fly with a lookup table (LUT). After performing numerous simulations, it appeared that a big contribution to the time resolution is given by the peaking time of the shaper. This is why the designer of the analog input block changed the shaper transfer function from a second order $h(t)=\left(\frac{t}{\tau}\right)^2 e^{(-2t/\tau)}$ to a third order $h(t)=\left(\frac{t}{\tau}\right)^3 e^{(-3t/\tau)}$. The result of the time over threshold method for several peaking times of both the old and the new VFAT3 shapers are shown in Figure 61.




Figure 61: Time resolution for different shaper peaking times in a Ar:CO₂:CF₄ Triple-GEM detector [31]

This figure shows the time resolution achieved by the electronics as a function of the peaking time of the shaper. Above a peaking time of 50 ns, the time resolution drops down below the 5 ns, which was the initial goal of this relatively simple Time over Threshold algorithm. The data simulated with the program Garfield used for this study comes from data used to perform this simulation comes from the same dataset as used in Chapter 3.2.2

The design process is ongoing, a first release of the analog inputs block was submitted at the beginning of 2014. This first engineering run is under characterization. The final design of the entire chip is scheduled to be ready at the end of 2015.

4.3.6 The GEM Electronic Board (GEB)

The GEB board was designed by the University of Lappeenranta (Finland) to ease the integration of 24 VFAT3 chips on each Triple-GEM detector. The main motivation of designing this board is to avoid the impractical use of cables in the narrow slots foreseen for the GEM modules. The GEB board integrates the power network for the front-end chips as well as the high voltage lines to the Triple-GEM chambers. It also contains the differential pairs for the E-links. These are the 320 Mbit/s communication channels required between the front-end chips and the Gigabit Transceivers (GBT) chip sets used to ship the data over optical links. The GEB also provides the Master clock (MCLK) distribution to ensure the VFAT3 chips are synchronized. This is shown in the schematic views of Figure 62 below.

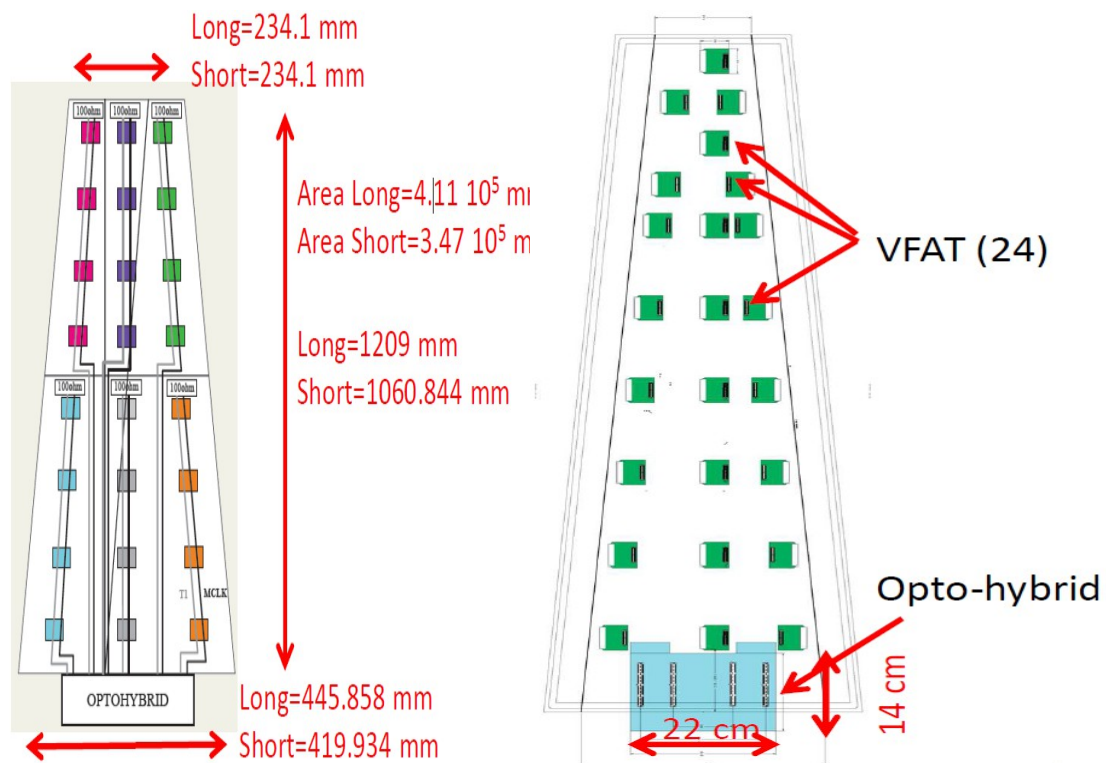


Figure 62: Figure 5: schematic view of the GEB (left) and VFAT3 implantation floor plan (right)

Two versions of this GEB board were produced so far, optimizing at each step the noise decoupling on the power lines and the impedance control of the E-link differential pairs. One of the main critical parts on these boards are the high speed and high density connectors used. The final choice for these connectors are high pin density Panasonic connectors for the 128 copper strips. Figure 63 shows a picture of the first version of the GEB, with the connectors in place, as well as three front-end hybrids in their final position.




Figure 63: Picture of the first GEB version

4.3.7 The opto-hybrid

The opto-hybrid board for the final system has a size of 14.0 x 22.0 cm, and is plugged into the GEB on its long edge (the furthest away from the beam pipe). Here again, the choice of the high density connectors was critical, since impedance mismatches on differential lines can impair proper communication at a speed of 320 Mbit/s. The final choice is to use Samtec QSE-080 connectors. Several iterations were necessary to adjust the performances of the board according to the specifications. Figure 64 shows the block diagram of the last version of the design, which is focusing on the trigger data path and its latency optimization.

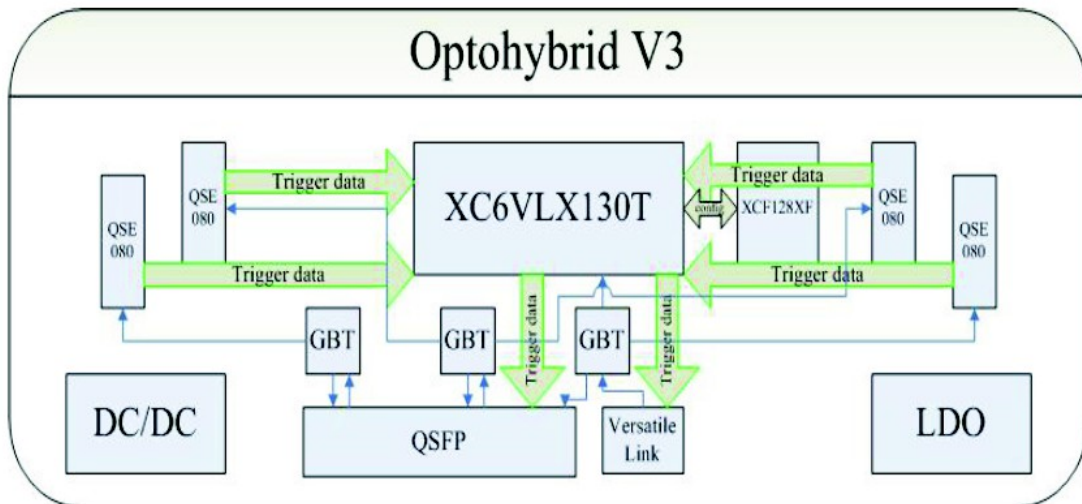


Figure 64: Block diagram of the third revision of the opto-hybrid board [25]

A powerful Xilinx Virtex-6 FPGA (model XC6VLX130T) is centralizing the tracking data requests and shipments, as well as the trigger data fanout to the off-detector electronics and the CSC Trigger Mother Boards. The one-way trigger path are shown in the form of green arrows. Two identical copies of this data is sent outwards, one for the CSC trigger electronics over a CSC compatible protocol, and the second copy towards the GEM trigger electronics over the GBT protocol. The tracking data is sent only over the GBT protocol, the chosen form factor to accommodate the optical link is the dense QSFP standard. Figure 63 Shows the resulting electronics board, where the central FPGA, the optical module cages (on the left) and the Samtec connectors (on top) are clearly visible.




Figure 65: Picture of the last version of the opto-hybrid board [25]

4.4 Back-end electronics

4.4.1 μ TCA based architecture

To centralize the trigger and the tracking data streams of the entire GE1/1 systems of the two end-caps into one single point, an infrastructure based on μ TCA seems to be the best option. This is thanks to this standard's inherent reliability coupled to the tremendous backplane throughput capabilities in a relatively small form factor. The conclusions of the expert group in charge of the evaluation of this standard for the CMS DAQ systems do recommend the use of a dual star topology with a redundant clock network for the backplane. This combination is commonly called the telecommunications option since it is the most commonly sold combination to telecommunication operators. Here is the list of recommendations from a CMS DAQ point of view [32]:

- Commercial MCH1 for crate management, GbE communication, and other user features as desired
- Custom MCH2 providing:
 - LHC 40.08 MHz low-jitter clock distribution
 - Fixed-latency controls distribution (aka TTC)
 - DAQ functionality; readout of data from AMCs
 - Buffer management communications for TTS-like functions as well as possible selective readout control
- Approved crates with the following features:
 - 12 full-height double-width AMC slots preferred
 - Two standard (single-width) MCH slots
 - Approved power modules with 48V bulk input
 - Vertical airflow for cooling
- Backplane with the following interconnections:
 - Dual-star routing of Fabrics A, B, D, E, F, G to MCH1 and MCH2
 - Dual-star routing of CLK1 to TCLKA (MCH1) and CLK1 to FCLKA (AMC13 in MCH2 site)

The choice was made to use VT892 crates from the Vadatech manufacturer, since these fit perfectly in the specifications, and are also used in other sub-detectors of CMS. Figure 66 shows an example of this crate model, with twelve slots, two redundant power modules and two installed MCH units in the central dual star points. In addition, two fan trays on top and bottom of the AMC slots ensure a vertical air cooling flow. These crates were chosen because of their long reliability track record.



Figure 66: Picture of the standard uTCA crate chosen for the DAQ upgrades of CMS [14]

4.4.2 The MP7 processing boards

The MP7 design is an advanced μ TCA data concentrator board, coupled to a powerful processing module. It can accommodate up to 72 bidirectional optical links in its latest revision thanks to 12 Avago miniPOD modules.




Figure 67: Picture of the MP7 board [26]

One of the main advantage of combining the μ TCA infrastructure with FPGA chips is the resulting unlimited flexibility in the communication interfaces. In addition to the 72 high speed transceivers used for the optical fibers, the Xilinx Virtex-7 is also able to distribute the latency constrained trigger data and voluminous tracking data to the backplane over dedicated differential pairs. The reason for using this board is that concentrating the trigger signals of a large part of the detector into one point enables the upfront processing of trigger pattern recognition algorithms.

4.4.3 *The 13th Advanced Mezzanine Card (AMC13)*

The AMC13 is the result of a visionary project initiated by the University of Boston even before the μ TCA standard was chosen to replace the aging VME infrastructure inside CMS. Not only it encloses all the μ TCA.0 specifications for an AMC board, but it also acts as a central DAQ data collecting board to be seated inside an MCH slot of the μ TCA crate. This last point represents the main innovation of this board.

The first revision of the AMC13 was available in 2010, and was already built with the idea of distributing an experiment-wide synchronization clock, as well as distributing a trigger signal to the read-out boards of a rack. The main features list is given below [33]:

- Mounts in redundant MCH slot of dual-star MicroTCA crate
- Occupies Tongue 1 and Tongue 2 slots, optionally Tongue 3
- Receives and decodes legacy TTC fiber (with very low jitter)
- Distributes 40.08 MHz LHC clock on MicroTCA CLK1 (M-LVDS)
- Distributes L1A and fast timing on Fabric B
- Collects DAQ data from AMCs using Fabric A

A block diagram of the AMC13 is given in Figure 68. Tongue 1 is the main first level PCB in an MCH design. Four bidirectional optical fiber transceivers are visible on the left side of this tongue 1 board. These constitute the main interface links with the central DAQ system of CMS, with two DAQ data and slow control links, a Trigger Throttling Clock and Signal link and a spare link for later use.

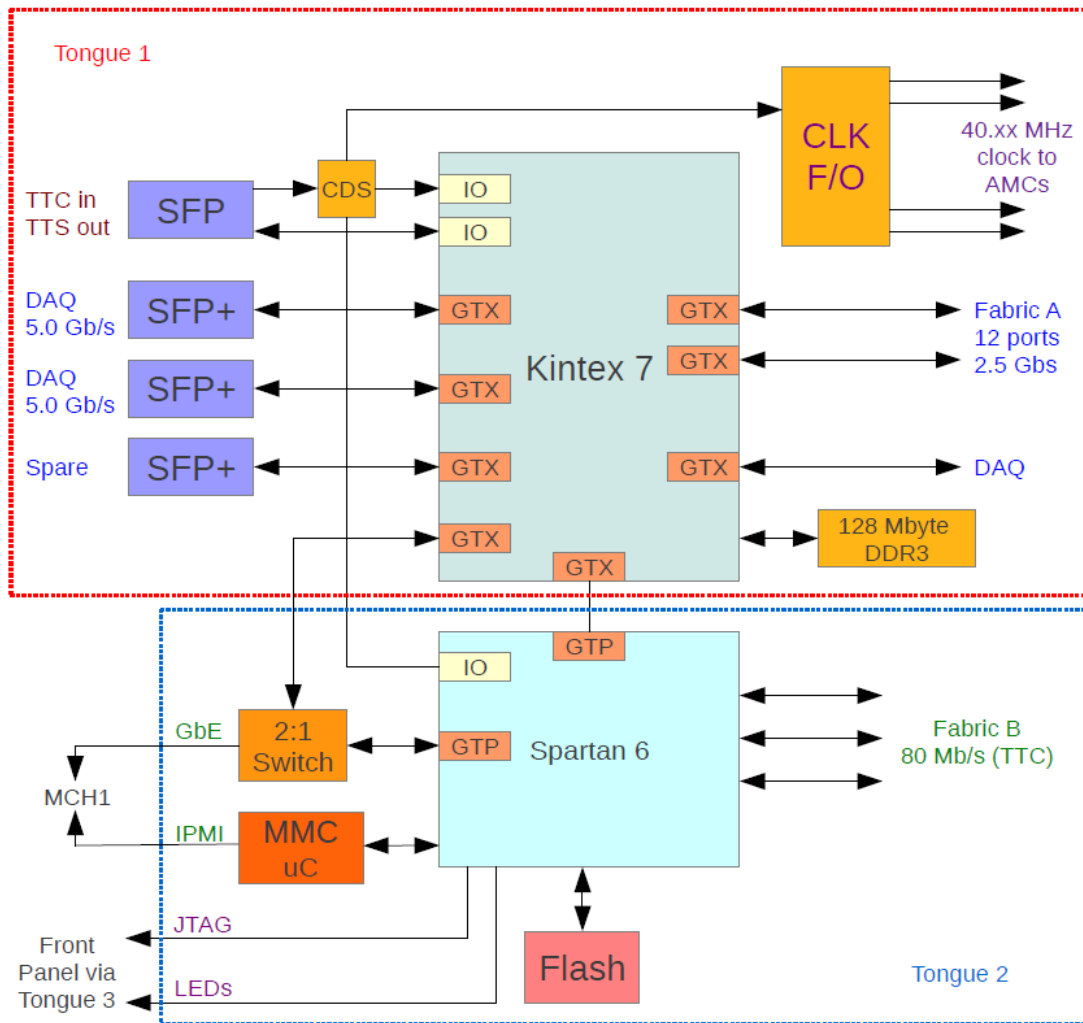


Figure 68: AMC13 (latest revision) block diagram [34]

On the right side of the tongue 1 are located the crate distribution fanout ports, mainly in the form of I/O type differential pairs, plus a group of twelve Gigabit Ethernet links to collect the DAQ data from the entire crate (twelve read-out AMCs). The TTC (clock + LV-1 Accept) signal distribution is ensured by tongue 2, over the unused Fabric B ports of the CMS standard μ TCA crates.

4.5 System description for slice test

4.5.1 Motivations

As a proof of concept, the management board of CMS approved the installation of two or four GE1/1 super-chambers (double back-to-back Triple-GEM detectors) during the end-of-year technical stop of 2016. This installation is called the slice test. The idea is for the CMS GEM collaboration to show the technical advancements and readiness before the Long Shutdown 2 in 2019 and the installation of the full GE1/1 system. These modules will cover a total of 20° or 40° in Φ and will be placed as shown in Figure 69 below.




Figure 69: Position of the two super-chambers for the slice test [25]

As some on- and off-detector components will not be ready yet for this short deadline, some changes to the DAQ electronics are foreseen. This development variant of the global design is specific for the slice test and is detailed in the current section.

4.5.2 Front-end electronics

The first main difference between the slice test electronics and the final system will be the front-end chip. Although the development of the VFAT3 is actively ongoing, the chances to see a fully finished and characterized chip before the of end 2016 are not guaranteed. To mitigate the risk of having a missing critical part in the DAQ chain, the choice was made to build a prototype based on the former VFAT2 chip for the slice test. The block diagram of the VFAT2 chip is shown in [27].

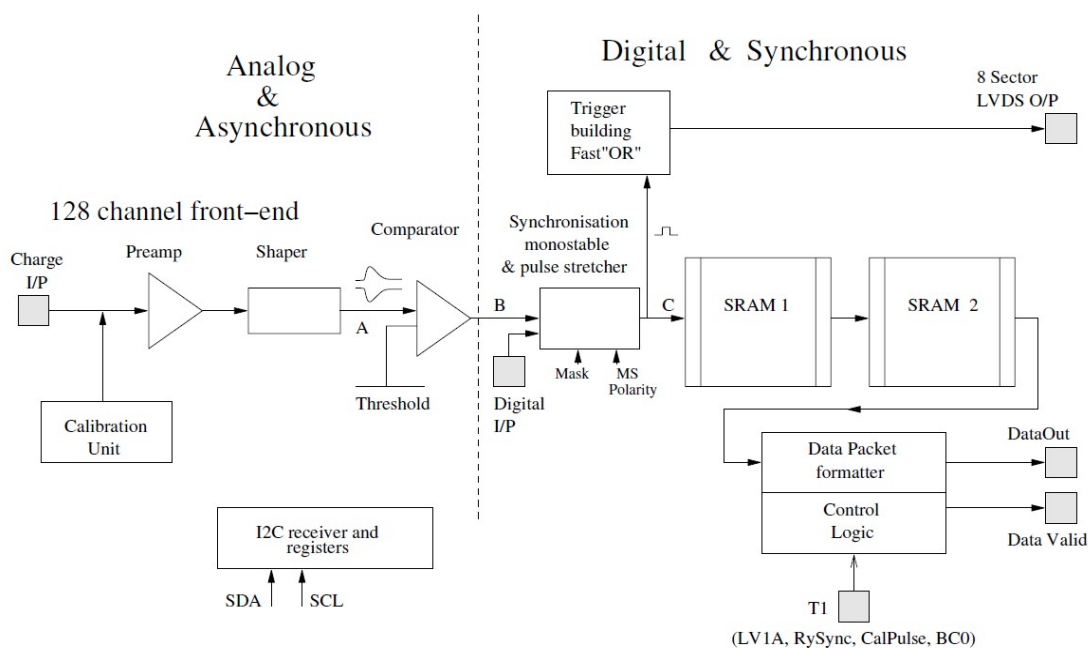


Figure 70: Block diagram of the VFAT2 architecture [27]

The main differences with the final VFAT3 chip are the analog input circuitry, and the digital output format. On the left side of the block diagram, the parameters for the preamplifier and the shaper are fixed in the case of the VFAT2 and the digitizing circuit is a monostable comparator. The VFAT3 will be designed with programmable gain and peaking time parameters, and the digitizing entity will be a Common Fraction Discriminator. The digital output, on the right side of the figure, is composed of differential pairs. This was initially chosen to ease the interfacing of this chip with the CERN designed GOL chip, also used on the CSC trigger system. The VFAT3, on the other hand, chose for the new GBT based transmission protocol to be used as the standard technology for the upgrades at CERN.

The GBT chip itself is also an ongoing project, and no formal release date is available to the date of today, as it is still under characterization. Two scenarios were developed to ensure that the slice test is producing data at the end of the technical stop, early 2017. The key component allowing to easily switch between the two solutions is the opto-hybrid board, developed under two concurrent versions. First of all, the VFAT2, originally designed for interfacing with the GOL, is not compatible with the GBT. A

large scale FPGA is mandatory to read out the 24 VFAT2 chips foreseen on the GEB. In the first scenario, the opto-hybrid board will include 3 GBT chip sets, if these are available. The FPGA, featuring by default several 4.8 Gbit/s capable transceivers, will then emulate the E-Port communication protocol from the VFAT3, and translate the incoming and outgoing data packets between the VFAT2 chips and the GBT chip sets. This would be the closest solution to the final system in terms of technology integration. The alternative path, in case of the GBT chip set not being available, is to use this same FPGA to emulate the GBT protocol and use regular small form-factor pluggable (SFP) transceivers to interface the opto-hybrid with the optical fibers. This second option will require more R&D and tuning work on the firmware level, but is meant to test the data throughput and reliability of the optical links, as the shipped data will be identical to the data provided by a regular GBT chip set.

4.5.3 Off-detector electronics

Concerning the off-detector electronics, one of the main components of the design might very well be missing for the integration deadline of the slice test as well, namely the MP7 board from the Imperial College of London. Furthermore, even if the board is available, the time required to confidently develop a working user application on top of this novel product is beyond the deadline of end of 2016. This is why another μ TCA read-out board will be used for the slice test: the CERN proprietary GBT Link Interface Board (GLIB) [35].

When it became more and more clear that VME was not a long term solution for hosting the always growing DAQ applications, CERN decided to follow the trend already initiated by many forward looking high energy physics institutes such as the University of Boston, DESY or SLAC to build up some experience with the μ TCA standard. Since the design of the GBT became one of the leading projects inside the Electronics Systems for Experiments group at CERN, it seemed an ideal occasion to build a testbed for this chip set on the high performance specifications of the μ TCA standard. This led to the definition of a specification list for the GLIB in 2010, and a very first available and usable revision followed in the second half of 2012 [35].

The GLIB features a large Xilinx Virtex-6 FPGA, four SFP cages two on-board mezzanine sites and some additional logic such as SRAM, clock distribution circuitry and signal integrity chips for communication with the backplane. Since this board was initially planned as a testbed to gain experience with the μ TCA standard, it became a fairly complex project. By end of 2013, a workable and reasonably well tested GBT chip set emulator was released by CERN, allowing the testing of protocol reliability and performance with the implementation of error counters on loop back and inter-GLIB optical links.

4.6 Conclusion

After having described in the previous chapter the Tripe-GEM technology to be used for the upgrade of the CMS forward spectrometer and detailed the goals of the upgrade in terms of physics, the proposed architecture of the GE1/1 read-out system was detailed in this chapter. The aim was not only to describe the work achieved during the design process of the proposed solution, but also to explain the choices of the chosen technologies. In a first section, we posed the problematic by explaining the short comings of what was in place before the Brussels R&D group joined the collaboration. Some challenges to overcome include the performance requirements, the mechanical integration, the request from the CSC sud-detector collaboration to get trigger data, and some more technical limitations.

The entire system was described as it will be installed during the Long Shutdown 2 phase of the LHC, from the front-end electronics to the back-end processing instrumentation, not to forget the slice test instrumentation which will be installed earlier to test the integration feasibility of the system beforehand.

The next chapter is dedicated to the hardware developments performed in the frame of the infrastructure upgrade of the back-end electronics towards the μ TCA standard. A description of the key functionalities of this standard are given first, followed by the specifications of a μ TCA board designed as part of this work.

CHAPTER FIVE

μ TCA developments



Particle physics experiments are usually on the forefront of the technological advancement. This is because the limits to the precision, the data quality or the event rates is set by the current technological limits such as bandwidth, scalability or storage density. For the coming upgrades of the LHC, the performance limits of the currently installed VME infrastructure will be reached. The decision was made within CMS to switch to a new standard for any electronics under development, in order to start replacing little by little the aging VME equipment and reach full performance for the final upgrade, beyond 2020. The chosen form factor for the upgrades is the newly born μ TCA standard. The main advantage over VME of μ TCA is that it is considered as a piece of protocol agnostic network infrastructure. The dependency to any specific CPU board supplier as it was the case with VME is thus eliminated. Furthermore, the throughput described in the current version of the standard is far beyond the requirements of CMS. In other words, this standard should not be the limit to any more extensions in the near future.

Of course, these promising performance levels come at a cost. In the case of μ TCA, this is under the form of a much higher level of complexity and a true integration challenge. To give an idea of the implications of the move towards μ TCA, a first section will describe the few important requirements imposed by the standard when developing μ TCA boards. The main goal of this present work is to gain experience with this new standard, and this is why some developments in this field were performed. The design of an MMC testbed board is detailed in the next sections, and a last section will describe the work achieved to allow remote firmware upgrades over standard network infrastructures.

5.1 μ TCA infrastructure

5.1.1 *The μ TCA form factor*

It was already revealed in Chapter Two that replacing the VME infrastructure with μ TCA constitutes a real gain in performance and flexibility, but it is at the cost of a serious effort in R&D. This new standard is complex, especially when the final product is outside of the commercially available telecommunication market, and when it is highly application specific as it is the case for most DAQ applications.

The first most evident difference is the physical size difference. An AMC module is about 25% smaller than a standard VME EuroCard. The consequence is of course a much higher component density and a larger number of routing layers. Designing such a board is a challenge as the current trend for DAQ applications is to increase the application power and the input channels density. The MP7 board from the Imperial College of London is a good example of complexity and power density. The cooling and the temperature monitoring become critical functions in the μ TCA crates.

Another consequence of the performance increase resulting from the current channel density levels is the massive use of differential pairs. This tends to add two steps to the design process of an AMC, namely the impedance calculation of all the pairs and a thorough signal integrity check. These are common requirements when designing long distance high bandwidth communication channels such as in the case of optical fiber links. But more and more, these techniques are also required to evaluate the performance of internal short distance connections, such as the AMC to MCH links over the μ TCA backplane.

5.1.2 *Module Management Controller (MMC)*

This is one of the main innovations of μ TCA over VME. Every piece of the infrastructure (crates, power supplies, fan trays, backplane, etc.) as well as the user applications (AMC boards and extensions) are monitored and managed from a central entity. This allows for a high level of serviceability. For example, alarms can be sent out in the case of a failure. But it also allows for an increasing reliability, with automatic fail-over features if any of the modules show first signs of weakness. Here again however, these functionalities are mandatory for any module to be activated upon insertion, and implementing these requirements is a significant R&D effort.

A good illustration of the complexity of the MMC functions is the powering sequence of a newly inserted AMC. Upon insertion, a slot mating circuit on the AMC gives a signal to the crate controller (the MCH) that a new module is present inside a given slot. The MCH enters in contact with the new AMC over a dedicated I²C pair of lines, also called the IPMI communication channel, to inquire the Field Replaceable Unit (FRU) information memory about the new module's capabilities and requirements, such as power needs or fabric interfaces. If these requirements match the crate

configuration, the MCH allows the AMC to be activated whenever needed. This state is called the Standby mode. To activate the AMC, the user has to push in the hot-swap handle. When this is done, the MCH allows the AMC to power up by telling the power modules to activate the Payload Power (PP) on this specific slot of the backplane. This simple example tells us that every AMC board needs to carry an intelligent board controller for the MMC functionalities, usually under the form of a micro-controller chip, a non-volatile memory component to store the FRU information, a hot-swap handle with an interrupt line to the MMC and several dedicated geographical addressing lines plus an address discovery state machine to tell an MCH the I²C address of the newly inserted AMC.

In addition to handling the infrastructural functions described above, the MMC is also responsible for the on-board AMC monitoring and housekeeping. Several temperature and current sensors are usually distributed over the AMC. These measure the health status of the payload, which can be forwarded to the MCH, where an out-of-band management system is running. This allows a user to retrieve these status informations from a central system such as the slow control (DCS) part of the xDAQ framework. A number of upper and lower threshold values can also be set inside the MMC to define some critical operating points not to exceed. The idea is to make sure the AMC is immediately switched off if any of these bounds are reached during operation, preventing the AMC and its surroundings to be destroyed in the case of a short circuit for example.

5.1.3 e-Keying and fabric interface

The e-Keying is another of these mandatory but complex procedures set in place by the μ TCA standard to avoid any capabilities mismatch between newly inserted AMC board in a running crate. As described in Chapter Two, the backplane can come in a number of variants, depending on the end-user application. But there are always several bidirectional ports on the backplane, dedicated to any of the chosen fabric interface. The used fabric interface, however, must be coherent in the entire crate, which means the MCH must feature a dedicated fabric switch and the backplane has to provide the right topology.

All this is known by the MCH, which reads out its own fabric interface type and the FRU information of the backplane revealing its topology. When a new AMC is inserted inside the crate, the MCH reads out its capabilities over IPMI. Knowing its own configuration, and reading the information about the capabilities of the AMC and the backplane, the Payload Power is granted if and only if a match is found between these three components. This complex procedure makes it impossible to switch ON an incompatible AMC inside a running crate.

5.1.4 Software integration

With VME, the read-out and control of a board inside a crate was possible by

allocating registers in the memory space of a CPU board running an operating system. Additionally, interrupts were also part of the extension of the CPU bus. Writing a device driver was the only requirement. In the case of the μ TCA standard, this is not the case anymore. Depending on the ports and the used fabric interfaces, the solution to communicate with an AMC may vary.

A first option, is to consider the μ TCA crate as a centralized microprocessor system. In this case, a CPU boards inside the crate acts as the root complex of any of the most common fabric interfaces present on the market, such as PCIe, SRIO, XAUI or even USB and SATA. The only limitation to which fabric interface to use is the existence of the corresponding fabric switch inside the MCH. The resulting system will behave as a computer and run an operating system. This option is then similar to VME. The performance will be high since many ports can be aggregated to form high bandwidth trunks, but no reliability feature is present. There is no redundancy possible.

The second option is to consider the μ TCA crate as a cluster of independent systems. The chosen fabric interface should not be a point to point protocol, but rather be a switched network protocol such as Gigabit Ethernet. The MCH acts as a genuine network switch and each AMC owns an address (IP, MAC, etc.), generated from the geographical addressing capability of the crate. The advantage of this solution is the inherent modularity. In addition, if one of the AMC fails, replacing it while the rest of the system is running, reduces the downtime to the level of adaptability of the read-out software. The main challenge in this option is to communicate with the payload function from a central point, over a protocol which is inherently not made for point to point connections in terms of latency and addressing.

To solve this issue, the xDAQ development group proposed to use a Hardware Abstraction Layer (HAL), called IPBus. The role of this library is to emulate the existence of a local computer architecture such as described in the first centralized microprocessor system option, and distribute it over the cluster of independent systems of option two. The resulting address space (seen from the top DAQ computer) is mapped in regions pointing to different physical systems. When an agreement was found on the address space access software methods, the μ TCA Hardware Abstraction Layer (μ HAL), developed by the University of Boston, was born. This is a clever way to solve the problem, especially because the only change with respect to the existing VME system is the address space. All the existing code developed before the existence of μ TCA can thus be ported and reused.

5.2 Module Management Controller (MMC) testbed

5.2.1 Motivation

The goal of building a first Module Management Controller testbed AMC was to illustrate our comprehension of the μ TCA standard. Before even being able to develop a user defined payload, certain requirements need to be fulfilled on the infrastructure point of view, under the form of a μ TCA compliant MMC slave. Mastering the development of this little piece of intelligent hardware is mandatory to develop new μ TCA based DAQ boards. But due to the complexity and the youngness of the standard, the development of MMC slaves remain a profitable activity for the time being, making it difficult to freely build μ TCA based DAQ components.

Some organizations in the particle physics community, including CERN, started an effort to centralize the development of a generic MMC modules. Two weak points appear in this initiative. First, the final products come with a cost [36] in order not to destabilize the existing market. As a consequence, no source code is shared. Secondly, the field of possibilities and capabilities is unlimited with μ TCA, which makes that no generic product can be good enough for everything possible. An MMC slave is AMC specific, this is why it is important to master its development.

Fortunately, we found a reliable partner in the name of DESY (Hamburg), which is involved in the development of μ TCA boards since the release of the standard. As opposed to other institutes, DESY was willing to share its source code, which gave a starting point to the development of home made AMC boards. The next step was to build a testbed, incorporating all the additional innovative features developed in Brussels, and their description is the topic of this section.

5.2.2 Overview

On a mechanical point of view, the MMC testbed is a single height, mid-size module (180.6 x 73.5 mm) with a PCB-style edge connector of 170 pins. The reason for choosing this size was the availability of a μ TCA starter kit from *N.A.T* specifically designed for promoting the development of AMC boards. This starter kit is composed of a five-slot, single height μ TCA crate with a MCH and a power supply. The PCB is a standard four-layer FR4 design, compliant with the specifications in term of thickness, electrical and mechanical characteristics. Below is an overview picture of the MMC testbed AMC.




Figure 71: Picture of the MMC testbed AMC module

The μ TCA edge connector is on the left side of the board, the two DC/DC converters described and tested later in this section can be seen on the bottom left of the picture, the MMC micro-controller is further to the right (labeled MMC MCU on the picture) and the reserved site for a mezzanine board is traced on the top half of the picture. These elements can be found back in the block diagram of Figure 72.

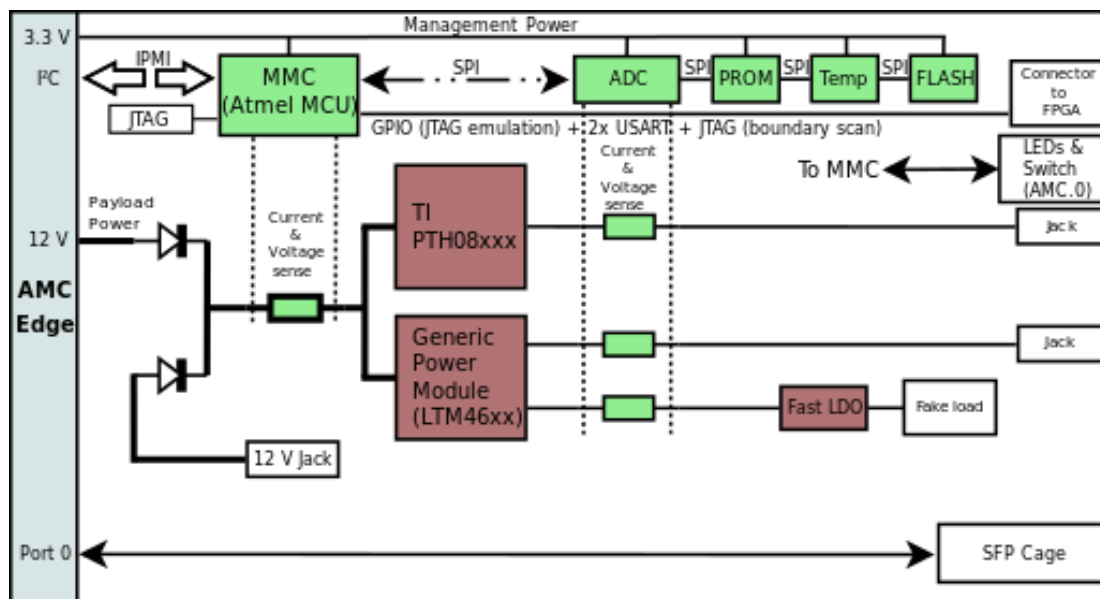


Figure 72: Block diagram of the MMC testbed AMC module

All the elements present inside this block diagram will be detailed in the current section, from left to right.

5.2.3 AMC edge connector

5.2.3.1 Physical characteristics

For this board, use was made of a PCB designed golden fingers connector, in opposition to the widespread Harting AdvancedMC™ plug-style connector. There are two reasons for this. First the price and availability of such a connector for a simple testbed board. In addition, the PCB style connector requires less design effort, the Harting AdvancedMC™ connector being poorly documented. So far we have not seen any of the common drawbacks of using a PCB version of the connector, such as copper pads peeling off or excessive copper wearing. This is a topic we will keep an eye on for future AMC/μTCA board developments. The connector size is AMC.0 compliant, 65.0 mm in width and 7.9 mm in depth. The full 274-pin version was implemented, however only 170 pins were used according to the usual implementation of the AMC backplane connection. For revision Rev A1, no actual gold plating was used on the edge connector, mainly for cost reasons.

All 56 ground (GND) lines have been routed to a common net and extended to a ground plane in one of the inner layers of the PCB. All 8 Payload Power rails have also been distributed across the board on a separate PCB layer. Calculations of the minimum track widths have been made to ensure all power rails comply to the AMC.0 requirements regarding power levels (AMC.0 Chapter 4-2 REQ 4.2 and REQ 4.4b). All payload power nets have a minimum width of 0.4 mm.

5.2.3.2 Power rails and presence signal

Maintenance power rails are provided on a dedicated single net on the bottom edge of the board, where all the components requiring this supply are located. For further AMC developments, components requiring to be tied to this power supply should be placed on the same side, to ease the routing. Here again, minimum routing width of this net is 0.4 mm to fulfill the specifications (AMC.0 Chapter 4-2, REQ 4.6b and REQ 4.7b).

The AMC presence signal is provided to the carrier by a low drop Schottky diode between signals *PS0#* and *PS1#*, as recommended in the specifications (AMC.0 Chapter 3.2.2, paragraph 18 & 19). Figure 73 shows the AMC.0 compliant direction of the presence detection diode.




Figure 73: AMC.0 compliant module presence detection diode

The diode is a low cost NXP BAT17 Schottky barrier diode, in a SOT23 package.

5.2.3.3 AMC enable signal

During the module insertion sequence, the next step after presence detection (*PS0#* and *PS1#*) is the enabling of the module by the crate manager (MCH). This is done by asserting signal *ENABLE#* on the backplane connector. Very common practice (and recommended by the specification manual, AMC.0 Chapter 3.4, Figure 3-3) is to use this signal as a reset for the Module Management Controller (MMC) intelligence (in green on the right side of Figure 74 below).

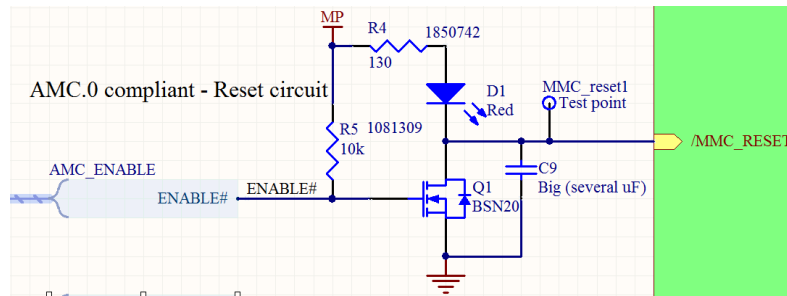


Figure 74: AMC.0 compliant module enable circuit

One detail that should be noted, is that the signal provided by the carrier is in open drain logic, whereas the MMC reset input is a standard CMOS input pin. This is why additional components (transistor and a pull-up resistor) were needed.

5.2.3.4 Geographical addressing and IPMI

Geographical addressing is performed through three tri-state address lines controlled by the crate manager. These three lines are tied to pull-up resistors and to an I/O pin of the MMC micro-controller. This scheme is in compliance with the AMC.0 specification set (AMC.0 Chapter 3.4).




Figure 75: AMC.0 compliant geographical addressing logic

The present address determination is needed by the MMC to broadcast its I²C contact point to the MCH. Since three states are allowed on the backplane side (low, high and open circuit) the address determination is performed by a finite state sequence inside the MMC firmware.

IPMI communication is done over two I²C lines connected to the shelf manager through the crate backplane. The I²C bus physical interface, traditionally composed of a data line (SDA) and a clock line (SCL) requires pull-up resistors to the management power net (called MP) on both lines (AMC.0 Chapter 3.5, REQ 3.34). The correct

values for the pull-up resistors are shown in Figure 76.




Figure 76: AMC.0 compliant I²C circuit for the IPMI signaling

It is worth mentioning that the standard data rate for the I²C protocol inside the IPMI scheme of the μ TCA standard is 100 kbit/s.

5.2.3.5 Fabric interfaces

On the main data path side, also called fabric interfaces, most ports are routed to the mezzanine connector, to allow further developments of home-made user applications. One exception to this is Port 0 (GbE), which is routed to an on-board SFP+ cage on the front of the board. This is to allow to doing some testing on different power supply schemes, since the data rate in these fiber optic modules is strongly dependent of the power supply fluctuations.

The first telecommunication clock (TCLKA) and the fabric clock (FCLK) are also routed to the mezzanine. To ensure some reasonable level of performance, all fabric signals to the mezzanine have equalized routing lengths, to synchronize these interfaces with the fabric clocks. In addition, both fabric clocks have been routed with an equal net length. A detailed summary of the port destinations is given in Table 4.

Port	Type	Destination
0	GbE	SFP on front panel
1	GbE	Mezzanine connector
2	SATA	Mezzanine connector
3	SATA	unrouted
4	PCIe	Mezzanine connector
5	PCIe	Mezzanine connector
6	PCIe	Mezzanine connector
7	PCIe	Mezzanine connector
8 to 15	Various	unrouted

Table 4: Summary of the ports destination

It is also worth mentioning that coupling capacitors are present on the receiving end of the differential pairs as stated in the AMC.0 specification document (AMC.0 Chapter 6.2, REQ 6.5b), and impedance has been controlled to meet the 100 Ω

requirement $\pm 10\%$, on each pair (AMC.0 Chapter 6.2, REQ 6.11).

5.2.4 Power

5.2.4.1 Input rails

The Maintenance Power is provided by the crate at any time. The voltage is 3.3 Volt, and the current is limited to 150 mA. Only components which have to be active even if no payload power is present should be tied to this net since the maximum current of 150 mA should never be exceeded according to the AMC.0 specification set (AMC.0 Chapter 4.2.2, REQ 4.6b). This is to be kept in mind for further AMC developments.

The Payload Power is provided by the crate, after a successful negotiation process with the crate manager. To be able to test our board and perform the initial debugging and MMC firmware developments, an alternate Payload Power path is present. It is in the form of a power-jack connector on the back of the board. This can be plugged to a 12 Volt transformer. The merging of these two power nets is done with OR-diodes, as shown in Figure 80. Current-sensing resistors are present behind the diodes, to achieve a measurement of the total input power for monitoring purposes by the Module Management Controller.

5.2.4.2 DC/DC converters

The AMC edge connector is providing stable and regulated Payload Power of 12 Volt upon request. This voltage is available to the user applications hosted on any AMC type board. To convert this power supply to more usable voltage levels, high efficiency DC to DC converters are required. The purpose of this testbed is also to perform some measurements of two different types of DC to DC converters, in order to make a choice for future developments. The two Devices Under Test are a *Linear Technology* LTM4619 Dual 4A Buck DC/DC converter, and a *Texas Instruments* PTH08T261 3A DC/DC converter. Figure 77 below shows a schematic capture of the LTM4619 module, with the required external components to set the voltage levels.

highly integrated component, requiring very few external components, but is only available in a Lead-less Land Grid Array (LGA) package which is not trivial to solder. The PTH08T261 module is a small circuit in a Dual In-line Package (DIP) format, which is very handy, but it is rather high and might constrain the development of mezzanine boards. Both DC/DC converters have *ENABLE* lines, controlled by the Module Management Controller, and can be switched ON or OFF individually. In the case of the LTM4619 module, two DC/DC converters are in fact integrated in one package, both can be controlled individually through independent *ENABLE* lines.

The results of a ramp-up performance test, in Figure 79 below, show that the response time of the *Texas Instrument* PTH08T261 is much better than the *Linear Technology* LTM4619 module at an 80 % load for both DUT.




Figure 79: Comparison at a 80 % load of the ramp-up times of the LTM4619 (left) and the PTH08T261 (right)

The goal of performing a comparison of the voltage ramp-up steepness over time is primarily to estimate the start-up time of the payload electronics once the MCH enables the Payload Power. This has an importance when using some modern fabric interfaces such as PCIe, since these rely on a bus discovery protocol at the very beginning of the root complex initialization (10 ms power up time according to [37]) Note that these ramp-up times are in accordance with the values given in the datasheets of the respective DUT.

On this board, the PTH08T261 and one of the Buck converters in the LTM4619 module are routed to power the mezzanine (user application). This corresponds to the upper DC/DC converter in Figure 80. A set of jumpers is used to select which source is currently under test. The second DC/DC converter in the LTM4619 is dedicated to provide power to the FPGA core through low noise regulators as shown in Figure 80.

Power distribution scheme

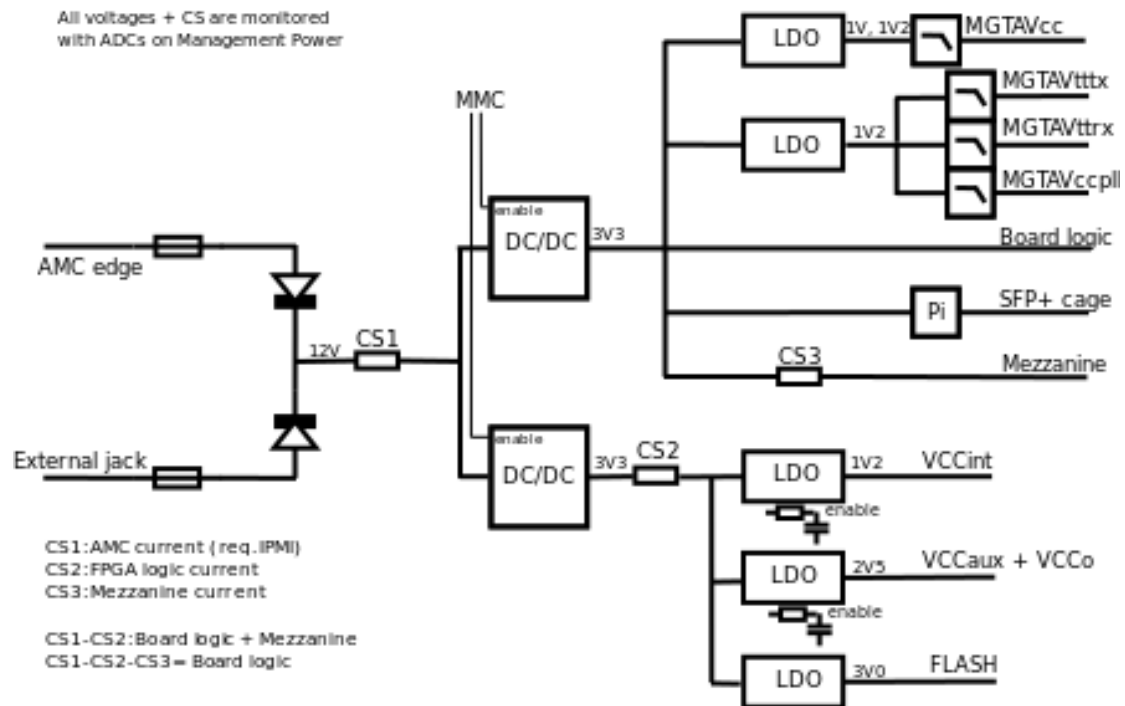


Figure 80: Power distribution network of the MMC testbed AMC module

This figure depicts the power distribution network, as it is intended to be implemented once the test phase is over and the first boards carrying a payload are to be designed. Three current sense (CS) resistors are used to evaluate the repartition of the power flows. In addition, real-time current sensing inside the MMC is mandatory to predict current surges which could damage the DC/DC converters over time.

5.2.4.3 Additional 3.3 Volt payload power

An additional 3.3 Volt power net was initially added only for AMC.0 specification compliance. According to this document (AMC.0 Chapter 6.4, REQ 6.52), the JTAG signal lines shall have 10 k Ω pull-up resistors to a 3.3 Volt power net derived from the Payload Power on the input. To satisfy this requirement, a dedicated voltage regulator was added. In addition, functions that do not need a permanent 3.3 Volt supply (such as the Payload Power monitoring functions) are powered from this regulator. This net, like the Management Power, is reserved to local AMC level functions, and are thus not routed to the mezzanine, and should never be used to power user application components on further developments. The component used to generate this voltage level from the 12 Volts Payload Power is a *Texas Instruments* μ A78m33-Q1 Fixed Voltage Positive Voltage Regulator in a SOT-223 case. This choice was made because it requires very few external component to operate, it can supply up to 500 mA, and it suits the 12 Volts input level provided by the crate.

5.2.4.4 Low Drop-Out (LDO) regulators

To provide a stable voltage to the user application's noise sensitive components (like Gigabit transceivers of an FPGA) a low noise, high bandwidth power supply rejection ratio and low drop-out linear voltage regulator has been added, drawing its current from the second output of the Linear LTM4619. The chip itself is a Texas Instrument TPS7A8001 in a SON-8 package. This component was chosen because of the very good Power Supply Rejection Ratio (PSRR) offered at the frequency the DC to DC converter is operating. The key role of this addition is to filter out the remaining switching noise, to ensure a good level of performance of the downstream high speed communication controllers. On this version of the board, the output voltage is set to 1.0 Volt, but this is easily adjustable by placing an appropriate pair of resistors in the feedback loop.

5.2.5 Module Management Controller (MMC)

5.2.5.1 Description

As already explained in Chapter Two, the Mezzanine Management Controller refers to the intelligent piece of hardware on-board in charge of :

1. Monitoring the general functioning of the board (power, temperature, etc.)
2. Forwarding this information to the crate manager and, alternatively, to a remote management user-interface over the standard IPMI protocol.
3. Negotiating backplane requirements (voltage, bus type, etc.) with the shelf manager

The second point in this list is the point we wanted to enhance in this project. The vision of this project is to let the MMC be the centralized board management entity, with exhaustive communication to the user application, such as a central slow control system. The aim is to be able to forward user application messages (eg. error counters for communication links, DAQ trigger levels and statistics, debug messages, etc.) originating from their processing units (CPUs, DSPs, FPGAs...) to a centralized detector control framework, such as DOOCS (in development at DESY) or the DCS of xDAQ for CMS. The channel for this exchange would be the standard IPMI interface present on the shelf manager, alongside the usual global environmental monitoring data (power supply status, temperature, voltage etc). At the time this project started, none of the similar boards under development in the community, nor any other obscure proprietary boards available on the market provided this functionality.

5.2.5.2 Micro-controller

The processor chosen to implement the Mezzanine Management Controller is the well-known *ATMEL* ATmega 2560, with 256 KByte of FLASH memory to store its firmware. It is one of the biggest model in the 8-bit family of micro-controllers from *ATMEL*. The reason for choosing this model is to never be memory-constrained in the implementation of new functions, since these user application messages forwarding,

for instance, can be unlimited. This micro-controller series also include the required peripherals to fulfill the needs of implementing exhaustive board management features and user application processors communication. To comply with the available Management Power voltage and to ease the integration, version -16AU (up to 16 Mhz clock frequency at 3,3 V in TQFP package) of this chip was identified as being the best choice. An 8 MHz external crystal oscillator is providing the clock. The reset line is tied to a reset circuit controlled by the AMC backplane *ENABLE#* signal. The micro-controller is completely autonomous, does not require external memory or signals to startup, and is ready to operate when the AMC is inserted into the crate.

5.2.5.3 AMC.0 compliant signals

All the signals required by the AMC.0 standard are present on this revision of the board, as described in chapter 3.3 “Additional local Module functionality” of the AMC.0 specifications document.

- BLUE LED is present and tied to an I/O pin of the MMC. This LED can provide feedback to the user on the current Hot Swap state of the module.
- HOT SWAP HANDLE is present. For cost effectiveness on this first prototype, however, it has been replaced by a SPDT switch on the front of the board. The connection of the switch to an I/O pin of the MMC is compliant with what can be seen on Figure 3-3 of the AMC.0 specifications.
- LED1 (mandatory) and other LEDs (optional) are present and tied to I/O pins of the MMC, in a current sink connection as depicted on Figure 3-3 of the AMC.0 specification set. Two bi-color (red/green) LEDs are counted in the optional LEDs
- A watchdog timer is always preset in the MMC.
- A PRESENCE DETECT loop (*PS0#* and *PS1#*) is present and integrates a low voltage drop (Schottky) diode, to notify the carrier that the module has been inserted properly.
- An *ENABLE#* line is present, to start up the MMC functionalities as soon as the Module is inserted. This line is tied to the RESET line of the MMC through a piece of open-drain conversion logic, as advised on Figure 3-4 in the AMC.0 specifications.
- IPMB-L signals (*SCL* and *SDA*) are routed to a dedicated I²C peripheral of the MMC. The mandatory pull-up resistors to the Management Power are in place.
- GEOGRAPHICAL ADDRESSING signals (*GA0*, *GA1*, *GA2*, *P1*) are also routed to the MMC, and pull-ups from *GA0*, *GA1* and *GA2* towards the *P1* port are in place.

5.2.5.4 AMC.0 compliant EEPROM

The AMC.0 specification set states that each module should have its Module Management Controller (MMC) and a dedicated EEPROM to store the FRU information data requested by the Carrier upon Module insertion. The memory chip is present, and is a 2 kbit *Microchip* 25AA02E48T EEPROM with a hard-coded globally unique MAC address. The reason why this chip was chosen is to improve the genericity of this board. High bandwidth protocols such as Ethernet require a unique Media Access Control (MAC) address also called Hardware Address, which is physically identifying the node on a network. These protocols are typically implemented in the user application of this board, since these are making use of the differential pairs (ports) of the backplane. If we leave this identification exclusively to the user application, however, the user application will require individual configuration files (firmwares) containing this different globally unique Hardware Address. Since this board was designed to be as generic as possible, the Hardware Address is on-board and can be provided to the user application to immediately identify itself on the network, and to enable the use of generic non-individual firmware files in the user application.

5.2.5.5 Additional FLASH memory

In addition, a 1 Mbit FLASH memory chip (*Microchip* 25LC1024-E/SM) was added to the MMC chip in order to store user application related information. The idea is to be able to store a simple user application firmware file to be used in an emergency situation, such as memory loss or alteration due to radiation effects. Alternatively, this FLASH memory could host a Power ON Self Test (POST) firmware image of the user application to be loaded on Payload Power availability, ensuring the user application hardware is ready before the production firmware (user application) image is loaded.

5.2.5.6 JTAG scheme

The on-board JTAG scheme is very simple on this board. Effort was made to ease comprehension and usability of the JTAG functionalities to the user, rather than increasing the amount of individual JTAG channels. There is only one JTAG access point on the board. This access point is used for programming the MMC, any JTAG compliant device on board and the user-application JTAG compliant devices (such as FPGAs) as long as there is a valid IEEE 1149.1 JTAG chain. This access point can be the on-board connector or the JTAG lines on the back plane. These two are totally equivalent, the on-board connector is meant to be skipped in future releases, once the backplane JTAG lines are fully understood and accessible. The key concept of the JTAG functionality of this board, is that the access point (on-board connector or backplane lines) are routed to the MMC chip I/O pins. These pins can be configured dynamically as JTAG pins or as regular I/O pins.

When the Hot Swap handle is pulled out (board in Management Power only mode), the four pins are configured in JTAG mode, which means access is given to the JTAG debug and configuration features of the MMC micro-controller. When the Hot Swap handle is pulled in (Payload Power is ON, and user application is started), the pins are configured as regular I/O pins, and JTAG signals coming onto these pins are forwarded to the user application JTAG chain. From the outside point of view, if the

handle is pulled out, you see the MMC, and if the handle is pulled in, you do not see the MMC but you have full access to the user application JTAG chain.

5.2.5.7 User FPGA reconfiguration

Several I/O pins of the MMC have been dedicated to user application reconfiguration, in the field firmware update or running firmware validity check (with radiation effect tolerance in mind). These services are FPGA oriented, but could be extended to other types of user applications like DSPs or CPUs. Since the Rev A1 revision has no user application, the signals are routed to the Mezzanine connector. The available signals are summarized in Table 5.

Signal	Description
FPGA_DONE	Valid reconfiguration finished
FPGA_INIT_B	Reconfiguration monitoring pin
FPGA_RESET	Reset line for FPGA
FPGA_RECONFIGURE	Reconfiguration strobe
FPGA_CONFIG_MODE_x	Reconfiguration mode select
FPGA_REV_SEL_x	Firmware revision select
FPGA_DOUT	Serial data output
FPGA_DIN	Serial data input
FPGA_CCLK	Serial data clock
FPGA_FUTURE	Reserved for future use

Table 5: Available payload reconfiguration controls

5.2.6 Monitoring and on-board network

The on-board network is one of the most innovative features of this board. The entire AMC is ruled by a Module Management Controller. This entity does the monitoring of the environmental conditions, the operational conditions (voltages, currents, Payload and Management powers) and can take decisions regarding these values, such as sending alarms or switching of parts of the boards. The on-board network is an extension of this control and monitoring feature towards the user application. Two communication lines are dedicated to UART message exchanges between all the UART compliant chips on the board, in the form of a ring. A *transmit* line is tied to the *receive* line of the next component in the ring. A token is traveling around the loop and ensures all the devices are active. This acts as a watchdog for all the components at the same time and is the first goal of the on-board network. A short message passing system can also exist on the on-board network. The MMC is always considered as node 0. It can ask status information (transceiver error counter, temperature, DAQ status...) or set a register (data rate preset, operation mode...) on one of the Nth processing node of the user application in the ring, and send this data to the IPMI engine on the crate controller or the carrier. The idea is to be able to

integrate the user application into the global management environment provided natively by IPMI in the xTCA world.

Monitoring sensors have been placed on the module. One is monitoring the general environment temperature of the MMC. Another has been placed under the *Texas Instruments* DC-DC converter to have some data on the heat dissipation of such devices. A third module is available in the form of a diode, to simulate the die temperature sensing diodes available on most large scale FPGAs, DSPs and CPUs. Voltage dividers are placed onto the Payload Power input of the board, and current sensing resistors are used to monitor the current flowing into the board. These values can assess the instantaneous power consumption of the board, and enable the MMC to take an appropriate decision if the measured value exceeds the maximum threshold defined in the AMC.0 specification set. Several voltage and current sensors are distributed across the board to check if a given part of the board is not showing problems such as short circuits or open connections. To start, the Management Power is monitored, with a reference voltage driven ADC. The 3.3V payload power needed for the JTAG and some small on-board features is monitored too. The voltage and current provided by both DC-DC converters are also under monitoring, as well as the 1.0 Volt *MGVtt* transceiver-specific power after its dedicated Low Drop-Out regulator.

5.2.7 Operation and debug interface

To be able to test all the features of this testbed, an operation and debug interface is available on the front side of the board. The communication protocol is the standard RS-232, with the parameters:

Baudrate	9600
Data bits	8
Stop bits	1
Parity	None
Flow control	None

Table 6: Interface access parameters

When a standard serial console is connected to this interface, the following screen is displayed as soon as <Enter> is pressed:

```
AMC MMC TESTBED Rev. A1 - August 2011  
IIHE Brussels - DAQ Group  
Board debug interface - Use carefully !  
Type "help" for list of available commands  
>_
```

which is the command interpreter of the small home-made operating system running inside the MMC. The "help" command provides a list of the currently implemented callable functions:

```
> help
General command format is "n[0-F] <command> [OPTIONS]"
For n>0, the command is passed to the Nth target on the board network.
Target 0 is the MMC itself. Note: you may omit n0 for local MMC commands.
Available commands are:
led0 [ON | OFF] Switch ON/OFF blue AMC.0 Hot swap handle LED
led1 [ON | OFF] Switch ON/OFF red AMC.0 LED
led2 [RED | GREEN | OFF] Control user LED 1
led3 [RED | GREEN | OFF] Control user LED 2
amc [MAC] Retrieve board hardware (MAC) address
power [lrm1 | lrm2 | pth] [ON | OFF] Control on-board DC-DC converters
temperature [PTH | ADC | DIODE] Display temperature sensor values
voltage [MP | PP | PP3 | PTH | LTM1 | LTM2 | MGT] Display voltage values
current [PP | PTH | LTM1 | LTM2] Display current values
flash [erase | readout] FLASH memory operations
```

5.2.8 Fabric interface mezzanine

In order to validate the entire Module Management Controller functionalities and especially the fabric interface e-Keying, an end-point to the fabric interfaces needed to be developed and tested. This was done under the form of a small mezzanine board featuring an Altera Cyclone IV FPGA chip. Since the small *N.A.T* development crate is including a PCIe friendly MCH module, this protocol was chosen to be implemented.

To achieve this, an FRU record generator was designed (in Python) to create the correct file structure and content summarizing the information present in Table 4. In addition to these fabric interface descriptions, the FRU file contains all the mandatory operational fields (institute name, board name, serial number, power requirements, sensor descriptions), the specific clock information (frequency, modulation and signaling standard) and all the associated checksum characters. Once the file was transferred to the dedicated on-board EEPROM chip and the entire development environment set up on the CPU board inside the crate, our board appeared in the addressable devices on the PCIe bus:

```
root@dev-utca:~# lspci -s 0c:00.0
0c:00.0 Unassigned class [ff00]: Altera Corporation Device 0004 (rev ff)
```

This line states that a new device with a yet unassigned driver has been discovered at address 0c:00.0 on the PCIe bus, containing an Altera FPGA-based PCIe end-point.

5.3 Remote firmware upgrade

5.3.1 Motivation

Modern electronics, and especially DAQ system often rely on field programmable logic systems such as FPGAs. The reason is that FPGAs offer a large potential of processing power in an easily achievable form (allowing quick development time) as well as an incomparable flexibility during run time. The user application processing system (called firmware) can be updated in the field to minimize downtime. This feature is especially popular in DAQ systems for physics experiments, since these systems can now, with the help of FPGAs, follow the evolution and enhancements of the detector itself.

Systems built around FPGAs are usually filled with their user application firmware during the test and calibration phase. The system is then placed in its duty location (counting house, service space or cavern, instrumentation rack) where a firmware update requires either human access, or a local master processing node such as a CPU. This is common for VME, PCI or older architectures, but not for the μ TCA standard which relies on Ethernet networking embedded inside the backplane. Access to the firmware requires a network layer on the processing node, which now can be anywhere on the network, and does not need to be in the crate itself anymore.

This section gives a description of the design and the specifications of a remote firmware upgrade feature meant to be implemented in every future DAQ μ TCA board. This study was performed in 2010, during the initial evaluation of the possibilities offered by the μ TCA standard. More evolved and inherently safer implementations of the feature have been developed by other working groups in the meantime, but this first-of-the-kind study was necessary to prove the feasibility of the concept. No advanced considerations were made on data integrity at this point and neither was it concerning safety. Both are, at this point, relying on the basic mechanisms present on the underlying Media Access Control network protocol (such as cyclic redundancy check, hardware address filtering and handshaking).

5.3.2 Firmware

An FPGA has usually its firmware (the description of the logic core) stored in an external FLASH memory chip. This firmware is loaded into the FPGA at power up, or when a logic reconfiguration is triggered. Updating the firmware is thus as simple as replacing the content of the FLASH memory with a newly updated and synthesized firmware. Several types of FLASH memories are common to store firmwares. In our case, the Xilinx FPGA was featured with a Serial Peripheral Interface (SPI) type memory chip. Implementing this functionality on other types of interfaces such as Byte-wide Peripheral Interface, however, would require minimal design changes.

The general idea is to have a minimal boot-up application, called *golden* or *failsafe*

firmware image at the beginning (address 0) of the FLASH memory. This small application will by default always be loaded into the FPGA at power up and provides access to the FLASH memory for the firmware upgrade. A second firmware, the user application, which is the firmware for which the board was designed is stored further in the FLASH memory array. After normal boot-up and when the DAQ system is ready to start, this user application image overloads the failsafe image inside the FPGA, in a so-called *multiboot* configuration scheme. The remote firmware upgrade functionality is meant to upload this user application image into its location inside the FLASH memory. The uploading is always performed by the failsafe application and over the network.

Figure 81 shows a block diagram of the firmware upgrade functionality to be implemented inside the failsafe image of the FPGA.

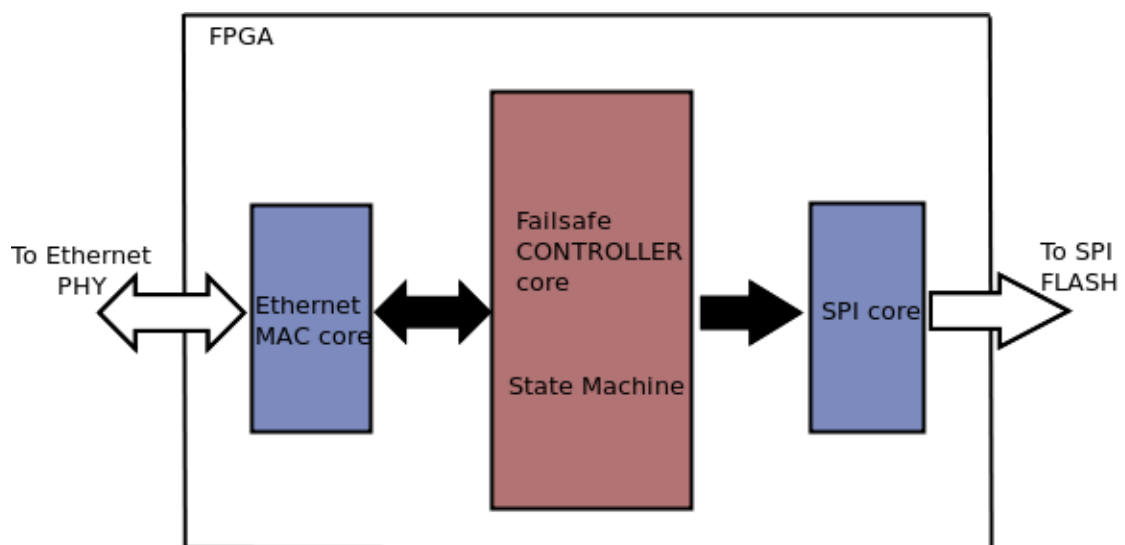


Figure 81: Firmware upgrade functionality, providing write access to the FLASH memory from the network

For the details of the testbed implementation, the Ethernet MAC core is the Ethernet tri-mode project available from OpenCores.org and provided under LGPL license. It implements a 10/100/1000 Mbit/s tri-mode Ethernet MAC conforming to the IEEE 802.3 specifications. The output interface is a variant of the so-called Atlantic Interface promoted by Altera for its SoC interconnection bus. On the SPI side, the core is a home-made state machine offering a configurable line speed to the FLASH chip. A typical upgrade transfer sequence is shown in Figure 82. This sequence is valid for a point to point communication between a master (remote computer or CPU blade inside the μ TCA crate) and the target FPGA board. A multi-point (broadcast) version was also envisaged, but the idea was abandoned for obvious safety and reliability reasons.

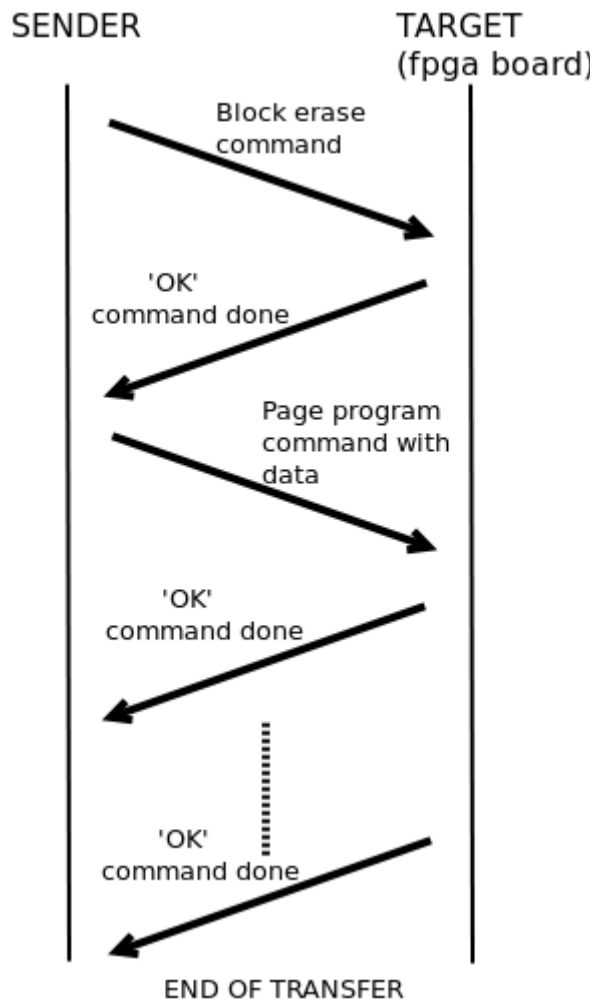


Figure 82: Typical firmware upgrade sequence

A controller state machine is in charge of :

1. Storing the received frame in a local (inside the FPGA) RAM cell
2. Grab only the Ethernet payload data by digesting the header after a CRC check has been applied
3. Reassemble the data in a double-word data packet format
4. Finally, add the synchronization signals for the SPI core.

The access to the FLASH memory requires a certain number of manufacturer-specific commands for write and erase operations. These mnemonics are issued by the master software, and recognized on the controller side as valid operations. This aims at preventing random and potentially dangerous access by sending erroneous commands. In addition to this mnemonic detection, a frame is not interpreted as valid if it does not belong to the right Ethernet Type of Service (ToS, third field in the frame header). This ToS and the mnemonic filters are, at this point, the only higher level safety features in place in this application. Any hardware (MAC) address is

considered on the receiving end and the acknowledgment frames are sent back to the initiating sender.

The SPI core is based on a state machine asserting the Chip Select and clock lines of the memory chip. The first data double-word (32 bits) is stored in a shift register as soon as the entire incoming frame is locked inside the RAM cell. The data is then clocked out on falling clock edges, enabling the FLASH memory chip to capture the data on the rising edges. This polarity is fully configurable, alongside the command set and transfer speed to suite most of the commercially available FLASH memory chips. The transfer speed between the FPGA and the FLASH memory chip, however, does not need to be fast, since the biggest time lag comes from the self-trimmer page write operation inside the memory array, which is forced inside the software.

5.3.3 Software

On the computer side, a small piece of software was written (in Python) to parse and chop the binary firmware file and to provide it to the controller state machine in a valid sequence. First of all, the firmware file must be in the right format (raw binary) to be stored directly in the FLASH memory. This is achieved by applying the following command to the boundary scan bitstream file (.bit) after synthesis:

```
> promgen -spi -w -u 0 <firmware>.bit -o <firmware>.bin -p bin
```

The "-spi" option ensures the data is prepared for an SPI type FLASH memory (no bit swapping), and the "-p bin" option provides the right output format. The output is the resulting <firmware>.bin file, that the programming utility will be able to handle directly.

To program the FPGA over Ethernet, the only information needed is the MAC (hardware) address of the target. The utility retrieves the file size, performs an erasing of the right amount of blocks inside the memory array, and includes an address counter to place the bytes at the right place, starting from an offset which can be specified as a generic parameter. Here again, also the appropriate FLASH commands (block erase and page program) can be changed to cope with different chip manufacturers. Here is the command to launch a program, after the target MAC address has been modified inside the executable:

```
> ./xilinx_ethernet_reconfigure.py <firmware>.bin  
FLASH erase: 100 % [ ##### ] Erase success : 524288 bytes = 8  
blocks  
FLASH program: 100 % [ ##### ] Program success : 464196 bytes  
= 1814 pages
```


5.4 Conclusion

The aim of this chapter was to consolidate our knowledge of the μ TCA environment by designing a fully μ TCA compliant board. As we have seen, it brings together many features required by the μ TCA standard, but also user level additions which can be useful for the development of proper DAQ boards in the future. First, besides test our ability to develop μ TCA boards fully complying with the standard and fitting in this complex IPMI manageability scheme, we tested a number of concepts which would be useful in production DAQ boards. Incorporating an EEPROM containing the globally unique hardware address of the board in the MMC complex is one example. Designing an onboard network to exchange DAQ data between the payload and the MMC for easy slow control integration is another example. Secondly, we gave the proof of the ability to reprogram FPGA firmware remotely. This is becoming common currently, but this was the first study of this kind at the time we started our developments in 2010.

To conclude the testing of the trigger and data acquisition system for the CMS forward muon spectrometer, a cosmic test-bench was built as well. This experiment will make use of incoming cosmic muons to test the Triple-GEM detector response. The associated read-out electronics is making use of most of the components which will actually be used for the final system. This is described in the next chapter and will conclude the proof of concept for a Triple-GEM based muon system.

CHAPTER SIX

Cosmic test-bench



After the analysis and the thorough understanding of all the components of the future DAQ system for the muon spectrometer upgrade of CMS, the best way to test the concepts was to build a real system. The aim is to reproduce the entire DAQ chain, from the energy deposition inside a Triple-GEM module to the off-detector trigger and event dataset generation, including the raw data transfer over optical links and the read-out inside a μ TCA environment. This chapter is a description of the work performed to successfully test the setup, bringing together the fruit of four years of concepts studying and knowledge acquisition. Of course, the final DAQ system

is to be installed in the frame of a particle accelerator experiment, which we can't simulate. But we can make use of a universal and inexpensive source of particles, namely cosmic muons, to calibrate and characterize the developed electronics. The momentum of these particles is typically of the order of a GeV/c. Muons will be detected thanks to the ionization of the gas molecules and the high electron amplification provided by the GEM foils. To end this chapter, some very preliminary results coming from a similar setup placed into a test beam are give as well.

CERN has a long history in the development of gas detectors for collider experiments. In addition, potential applications of GEM foils were found in nuclear medicine recently, which brought a small R&D group at CERN to develop and sell small Triple-GEM detector prototypes at an affordable price. Two of these 10 x 10 cm² modules were purchased for providing a realistic signal at the input of the front-end preamplifier. A description of these detector prototypes as well as the entire experimental setup is given in a first section of this chapter, followed by a complete overview of the results obtained with the setup.

6.1 Experimental setup

6.1.1 Triple-GEM prototype

The small $10 \times 10 \text{ cm}^2$ Triple-GEM prototype is very similar to the large CMS Triple-GEM detector spanning a full 10 degrees sector in Φ as described in the previous chapter. Three GEM foils are stacked in a 3-2-2 mm gap configuration¹¹ and the anode strips can be read out by one front-end chip. As for the full scale prototype, the gap electric fields are provided by an external resistor divider bridge. Figure 83 shows a picture of the prototype, installed in its testbed.




Figure 83: View of the $10 \times 10 \text{ cm}^2$ Triple-GEM prototype detector [38]

On the bottom left part of the picture, the high voltage supply, followed by the resistive divider bridge, providing up to 5 kV on the drift electrodes and typically up to 600 V between the 2 mm transfer gaps. The gas outlet is visible on the bottom right corner of the picture. A computer controlled gas mixture control system was developed, allowing a dynamic adjustment of the gas proportions between runs. On the right side of the prototype, the two strip connectors are visible, temporary covered with channel-merging LEMO modules. These are the strip connectors on which the front-end electronics will be plugged. The copper coating is meant to shield away any noise inside the system. In the center of the module, the lid foil can be seen, evidently closing up a gas filled detector.

¹¹ This gap configuration is slightly different for the CMS standard which is 3/1/2/1 but this has no effect on the conclusions of this chapter

To perform a reliable measurement campaign, this detector needed to be characterized. This was done using an ^{55}Fe 5,9 keV X-ray source and the commercial ORTEC PC172 front-end electronics [38]. Below are the resulting gain and rates plots:




Figure 84: Gain (top) and rates (bottom) measurements of the 10 x 10 cm² Tripe-GEM prototype [38]

As it can be seen in the top plot of Figure 84, the Triple-GEM detector can easily reach very high gains of several thousands. Although no discharge was observed, we preferred to limit the measurement at a high voltage of 4600 Volt, corresponding to a gain of about 14.000. On the bottom plot of Figure 84 we can see that the efficiency plot for the 5.9 keV photons starts at 4.4 kV, corresponding to a gain of ~ 3500 . This also shows that even without pushing the detector to its limits, a comfortable efficiency zone can be reached with this detector technology. Since this study aims at validating the concept of the new DAQ system, we are not measuring the absolute detection efficiency of the detector. These kind of measurements will be performed later by the Brussels team with this new DAQ system but still requires a lot of improvements to the testbed.

6.1.2 Trigger system

To provide a trigger signal to our electronics, a fully characterized cosmic muon detection chain was needed. Two scintillators are placed around the device under test, here our Triple-GEM prototype, as shown in Figure 85. The bottom scintillator is enclosed inside a wooden structure below the Triple-GEM, and is coupled to a pair of photomultiplier tubes (PMT) of which the efficiency plateau is well known, with a value of 95 % at 900 V. The reason for using two PMT on the same scintillator is to exclude false positives due to noise inside the PMT itself. The covered surface of this scintillator is $21 \times 21 \text{ cm}^2$.




Figure 85: Triple-GEM detector prototype on its test bench

The top scintillator is smaller ($9 \times 9 \text{ cm}^2$), to fit the size of the active zone inside the Triple-GEM prototype. Its efficiency plateau was obtained experimentally by increasing the high voltage supply until the ratio of detected muons over the triggered

hits from the bottom PMT was close to the scintillators surface ratio of 20%.

The readout electronics for this trigger setup is done inside a NIM crate. Each PMT is connected to a discriminator module providing a NIM-logic pulse of equal length for each input channel. A global AND is computed on these three inputs, and the result is sent to a NIM to TTL converter. This pulse will provide the trigger signal to the DAQ electronics.

6.1.3 DAQ electronics

6.1.3.1 System overview

The two main components chosen to build this system are as close as possible to the slice test electronics, as described in Chapter Four. The front-end is a VFAT2 chip, the read-out electronics is a GLIB board inside a μ TCA crate. In between, a pair of multi mode optical fibers will transfer the slow control data as well as the fast trigger commands towards the front-end chip on one side, and return the trigger and tracking data on the other side. Figure 86 gives an overview of the entire front-end system.




Figure 86: Overview of the front-end DAQ electronics

In the middle of the picture, the Triple-GEM detector prototype sits on the wooden bottom PMT enclosure. On the left side, the top PMT has its scintillator reaching just above the detector. On the immediate right side of the Triple-GEM we see a VFAT2 hybrid, from which a gray ribbon cable leaves towards the front-end side GLIB board, detailed below.

One component had to be added and specifically developed to emulate the GEB board

with the opto-hybrid concentrator. This component is interfacing the optical cables coming from the read-out electronics to the VFAT2 front-end chip. To speed-up the development of this component, an already existing hardware platform was used, namely a second GLIB board. To avoid confusion with the back-end μ TCA GLIB board, this board is called the front-end GLIB. The hardware is identical, the firmware core inside the Xilinx Virtex-6 FPGA is completely different, since it was forked from the opto-hybrid development tree. The I²C link towards the VFAT2 chip is simplified with respect to the opto-hybrid, to accommodate only one single front-end chip, and the trigger and tracking data paths are not duplicated towards the GEB board.

6.1.3.2 Front-end electronics

Using a common hardware platform such as the GLIB allows some gain in the development time on one side, but requires adjusting the version-specific parts of the project on the other hand. The VFAT2 chip, for example, uses a non-standard communication protocol, which needed to be emulated and ported to the GLIB board. This, however, motivated the use of the GLIB, since a site is reserved on the original design, to accommodate user-specific hardware, under the form of a high density FMC connector. A small FMC mezzanine board was thus designed to forward the signal from the VFAT2 chip to the on-board FPGA. The FMC is shown in Figure 87.




Figure 87: Picture of the FMC mezzanine board

Two LEMO connectors can be seen on the left side, to provide the trigger pulse from the NIM crate into the FPGA. This is the reason why the NIM-logic had to be translated to TTL levels. A resistive divider on the FMC ensures protection to the 2.5 Volt LVCMOS inputs of the FPGA.

The VFAT2 chip has three distinctive communication channels [27]. The slow control path is relying on a I²C bus, and is used to set configuration registers. This is the simplest frame structure in the entire DAQ chain. A read or write request is sent from the FPGA on the GLIB to the slave VFAT2 chip under the form of a valid I²C

transaction, with start, stop and slave acknowledgement bits present. Both a read and write request contain a VFAT2 slave address (3 bits), a read/write flag, a register address (4 bits) and the 8 bits of data to be read or write. The frame structure is given below:

15	13	12	11	8
Slave Address		Read/Write	VFAT register address	
7	Data to read or write (8 bits)			0

The slave address bits correspond to a hard wired value soldered onto the VFAT2 hybrid with pull-up and pull-down resistors. This allows addressing up to seven front-ends on the same I²C bus, which will be useful in the context of a full GEB populated with 24 VFAT2 chips.

On the opposite side of the FPGA, towards the back-end electronics, two high speed optical links are implemented using the built-in multi-gigabit transceivers (GTX) of the FPGA. The first link is simulating a fixed latency communication path for the local trigger signals (the so-called S-bits). In the final design, this link will be reserved for the trigger data to be sent to the CSC trigger electronics and the GEM μ TCA trigger crate. The second link is used to transfer the tracking data containing the channel hit patterns as well as the slow control commands towards and from the front-end.

For both links, the encoding is the standard 8b10b scheme. According to this standard, twelve 10-bit characters coding 8 bit data are reserved as control characters, to indicate start of frames for example. These are used in our system to identify the content of the data transferred on the optical links, between tracking data and slow control. Figure 88 shows a simplified block diagram of the tracking part of the firmware architecture running on the FPGA of the front-end GLIB board.

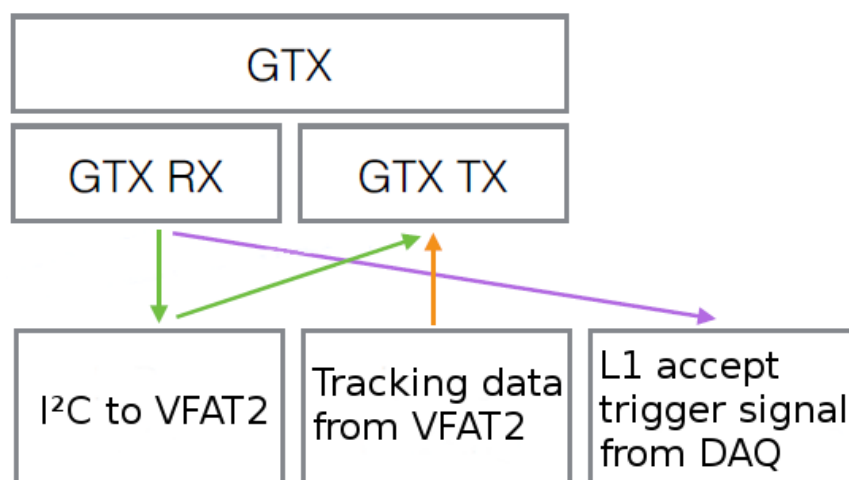


Figure 88: Block diagram of the front-end GLIB side FPGA firmware

The GTX RX and GTX TX block, which are directly connected to the optical link transceivers (GTX) are high speed multiplexers designed to forward the frames to the right destination according to their provenance or control character. The GTX RX entity, for instance, transfers the incoming slow control requests to the I²C bus, and forwards the Level-1 Accept trigger packets. Upon this Level-1 trigger signal, the VFAT2 returns tracking data to the GTX TX entity, which multiplexes these frames with I²C data from a read command on the slow control path.

6.1.3.3 Read-out electronics GLIB (μ TCA)

On the back-end electronics side, the GLIB is used as a transparent interface between the end-user and the front-end electronics. The only visible interface from a user point of view is a set of registers located on an IPBus client (see Chapter Five) on the network. To access the VFAT2 registers and data, and thanks to the abstraction layer provided by IPBus, the user only needs to know the IP address of the GLIB inside the μ TCA crate and the register bank structure. Figure 89 shows the block diagram of the firmware running inside the GLIB on the back-end electronics.

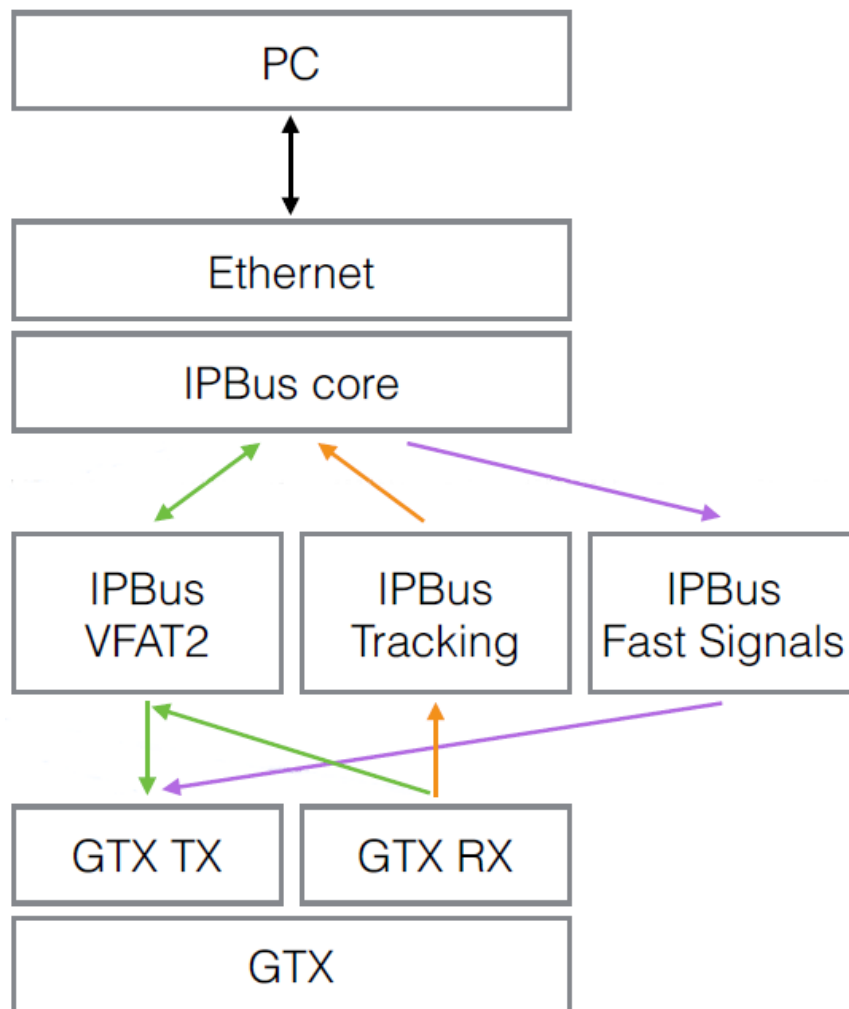


Figure 89: Block diagram of the back-end μ TCA read-out GLIB firmware

The IPBus core distributes the requests and returned data into several register banks, depending on the type of data transmitted. The IPBus VFAT2 registers are the slow control access points. The tracking data is available through the IPBus Tracking block. In addition, a L1-Accept message can be sent to the VFAT2 from the user as well, for testing purposes. A similar pair of RX and TX multiplexers as in the front-end firmware are present here as well, to merge the different requests and responses into a single optical fiber transceiver. This is where the correct 8b10b control characters are encoded.

The frames content is given by the back-end registers, as full copy plus a control character. These frames are forwarded directly to the final destination on the other side of the link, such as the I²C bus. The back-end registers thus contain the entire I²C frame, and no local interpretation of the frame content or translation is done anywhere on the chain. This full transparency was decided to increase the adaptability of this system to the quickly evolving specifications and requirements of the first prototypes. As the system was designed here, the only adaptations required in the content of the requests can be changed in the software on the user-end of the chain. The content of such a register is shown in the table below:

31	27	26	25	24	16
Unused		Error	Valid	VFAT2 Number	
15					0
16-bit VFAT2 frame					

Bits 24 to 16 are reserved for the VFAT2 number. The purpose of this is to implement an addressing method for the 24 front-end chips which will be present on the GEB board for the final design. Inside the back-end electronics firmware, the VFAT2 registers structure is duplicated 24 times. The error and valid bits are simple flags to ensure the request was successful from end to end.

6.1.3.4 DAQ software and run control

In order to operate this experiment, a full run control environment was developed as well, based on technologies present in the most modern experiment interfaces found in the experimental physics community. Here also, a front-end and a back-end component exchange physics data and slow control data. The front-end part, in this context, is a web based graphical user interface, allowing to simply and comprehensively act on the entire detector operation. This entire run control system was developed by the Brussels DAQ R&D group.

The back-end component is relying on a MySQL database and a number of data handlers. These handlers gather data from the experiment. Slow control parameters such as the high voltage and gas flow for example are continuously monitored. In addition, these data handlers are also designed to set or modify these parameters in real-time, from values on which a change occurred inside the database. To apply these changes, the front-end web interface is also constantly in contact with the database,

triggering the handlers whenever a value is updated. The reason to build the entire system on a database is to ease the tracing back of run parameters when doing detector development runs. The entire parameter sets are linked together inside a single run structure. Figure 90 shows an example of the high voltage control and monitoring screen for the experiment.

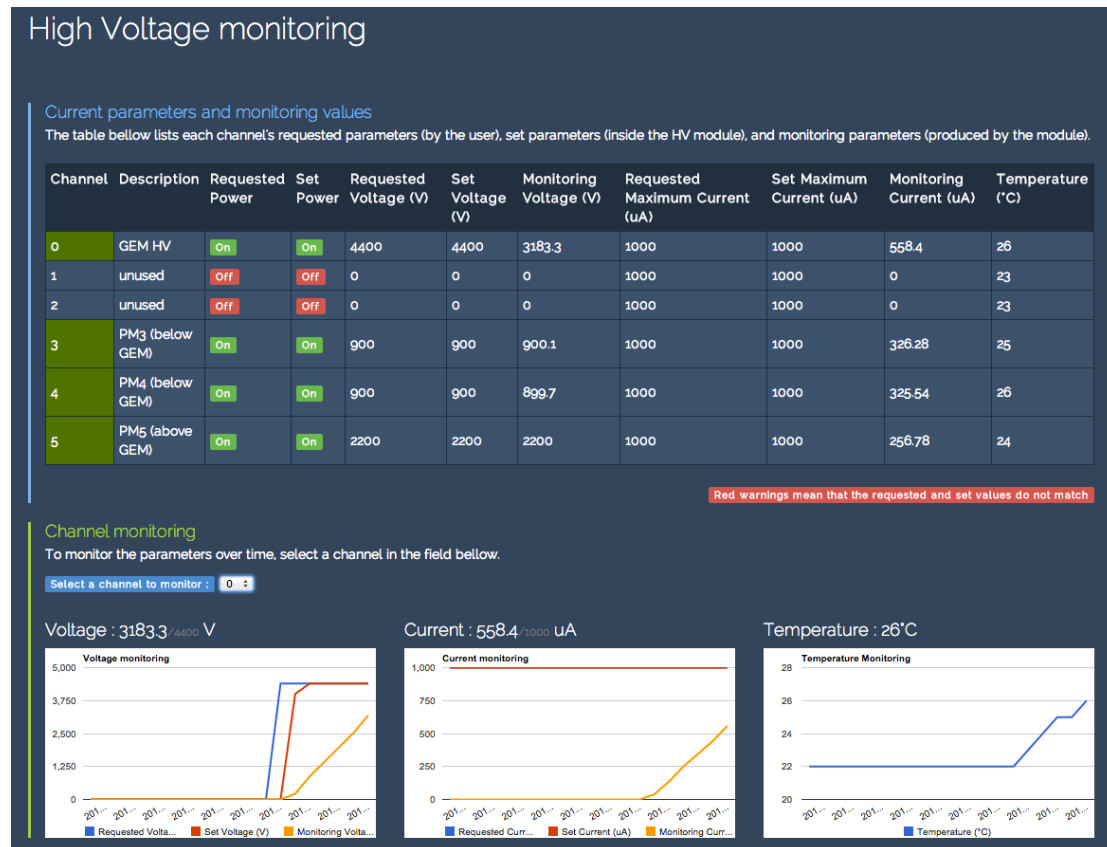


Figure 90: Example of high voltage monitoring screen on the run control web application

In this example screen, we see that four high voltage channels are switched ON. More specifically, the high voltage for the Triple-GEM detector prototype itself, and the high voltage sources for the three PMT surrounding the Triple-GEM. On the bottom, we see the evolution of the voltage, current and temperature of the high voltage sources. These plots were only possible with the use of a global run control database.

Besides the slow control, these data handlers and database system are also getting and setting the front-end (VFAT2) configuration registers, according to parameters entered by the user. The DAQ data upon valid trigger is also stored inside the database, to mimic the CMS physics data sets generation and storage after HLT reduction. These data sets can be visualized inside an event viewer and downloaded for further analysis.

6.2 Results with the μ TCA based read-out system

6.2.1 Performance of the DAQ electronics

Several parameters need to be characterized before connecting the electronics to the detector and being able to use the DAQ chain for scientific results. These parameters will be described in this section.

The VFAT2 front-end chip needs to be characterized as well, with parameters entered before each run. The monostable threshold is one of them. A hit channel is detected at the output of the analog block by a monostable and a threshold. This threshold is set by an internal 8-bit digital-to-analog converter (DAC), of which a copy is also available on the output pins of the hybrid. Setting the output to a given level and reading out this level with an ADC of higher precision located on the FMC board, will enable to accurately measure the linearity of the threshold values. For this, a signed 10-bit ADC ranging from 0 to 2.5 Volts was used for this, Figure 91 shows the resulting linearity curve.




Figure 91: Result plot of the threshold DAC linearity scan

This plot shows that the highest achievable threshold value is ADC count of 512, which is the highest possible value for a signed 10-bit ADC. This level corresponds to the baseline. The lowest possible threshold value is 325 ADC counts, which corresponds to a voltage of -0.45 V.

This voltage is far beyond the noise level fluctuations, but will cut off the events with a lower shaper amplitude, for events which spread over several anode strips for example. Another scan was needed to evaluate the value of the threshold to be just

above the noise level. To perform this scan, a logic OR of all the channels is defined inside the analysis software at the user end, and the threshold level is slowly increased until the number of valid hits coming from the analog block is close to zero. Figure 92 is showing the resulting plot.




Figure 92: Results of the threshold level scans

This plot shows a good immunity to noise starting at a threshold value of 30 DAC counts, which corresponds to a voltage of -52 mV. This becomes our new baseline for the next measurements.

6.2.2 Muon detection

A second parameter to optimize, is the chain latency, and is related to the architecture of the VFAT2 chip itself. Finding the value of this parameter, however, needs to be done on-line, with triggered events. The VFAT2 hybrid is fitted on the anode strips connector on the side of the detector, 4.400 Volts are applied to the detector and a 70:30 mixture of Argon:CO₂ is released at a flow of 30 ml/min.

At each clock cycle, the output of the comparators is saved into a first barrel-shifting memory block. Upon Level-1 Accept, it is recalled and sent to the second longer term memory block, from which it will be forwarded towards the back-end electronics. However, to find the event back in the constantly shifting barrel of the first memory area, a latency depth parameter needs to be provided to the front-end chip. This parameter corresponds to the number of clock cycles (bunch crossings in the case of a collider experiment) issued between the event recording and the Level-1 trigger. This value is constant in a detector, and is the reason why it is important to have only fixed and known latency communication channels in the trigger paths, as described in Chapter Two. A register inside the VFAT2 chip allows us to set this parameter before

the start of a run.

In the case of our cosmic test bench experiment, the latency is given by the length of the trigger electronics and cables since no high level decision mechanism is implemented. In addition to this fixed latency, our setup was built with a fixed 500 ns extra latency, in order to ensure the trigger latency is longer than the entire charge collection inside the Triple-GEM detector prototype. Our first estimations concluded that a 200 ns delay would be added to the 500 ns extra delay, split into a 120 ns signal build-up inside the 9 mm Triple-GEM detector, and an 80 ns delay inside the cabling, the NIM-electronics and the shaper. The procedure was thus to increase the latency counter register little by little between 15 and 40 clock cycles (375 ns to 1000 ns) and record a thousand events for each register value. Figure 94 shows the resulting plot.

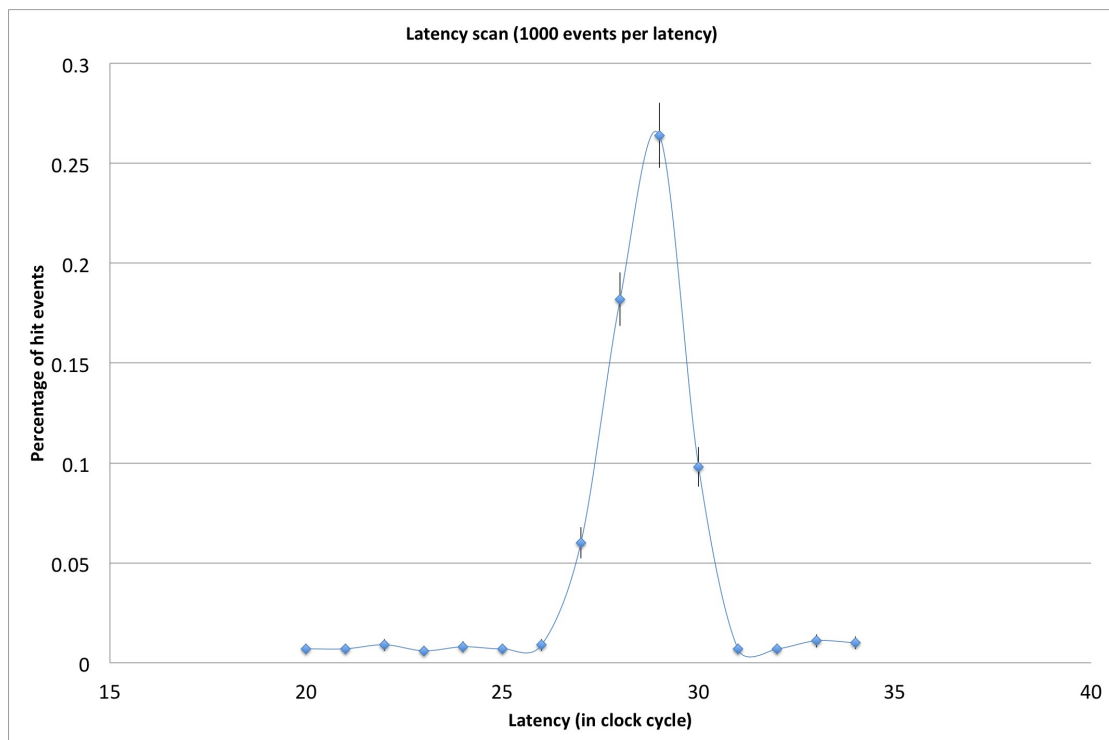


Figure 93: Plot of the latency scan from 375 ns to 1000 ns

From the plot, we can see a peak at 28 clock cycles, which corresponds to 700 ns latency. This value will thus be used for the latency configuration register during chip initialization at the start of each run. The spread of this distribution, however, shows a bad read-out efficiency. This is due to the poor quality of the NIM electronics from another age, used to generate the trigger signal.

The last step of the testing of the entire DAQ chain is the detection of cosmic muons with the Triple-GEM prototype. The resulting event rate given by the trigger appeared to be at 1.5 Hz, of which a majority are actual muon hits for the Triple-GEM prototype as well. These events can be read out by the DAQ electronics over the entire chain. One of the main questions for the study of the Triple-GEM detectors for

the CMS muon spectrometer upgrade was the cluster size, in other words, the number of anode strips hit for each muon. This value gives the spatial resolution capability of a Triple-GEM detector. The result is given in the plot below:




Figure 94: distribution of the number of channels hit per event (cluster size)

As we can see from these 140 events, a majority of hits are visible on one strip only. A number of muons, however, are spreading their signature over several strips, which most probably indicates an incoming angle with respect to the vertical axis. Nevertheless, the cluster size average of 1.7 is close to the value obtained during the latest test beam performed at CERN by the CMS GEM Collaboration with a muon beam as we will see in the next section.

6.2.3 Beam test results

At the end of 2014, the Brussels R&D group has installed and exposed to a particle beam at CERN a GE1/1 detector prototype equipped with the complete new readout chain, including : a GEB board to hold the VFAT2 chips, a 1st version of the opto-hybrid board and the uTCA electronics. A picture of the entire front-end part of the setup is visible in Figure 95. The tests have been performed at the CERN H4 beam line. The proton beam extracted from the SPS (see section 1.2.2) is sent to a target where the incident protons create secondary particles, in our case it produces pions with a momentum of 150 GeV/c. Because of their short lifetime, $2.6 \cdot 10^{-8}$ s, the pions decay and produce muons. In the test beam area, we can switch between pions and muons by acting on the beam collimators. The beam is coming from the left on Figure 95, perpendicular to the detectors.



Figure 95: Picture of the front-end part of the beam test setup

Four other GE1/1 detectors equipped with the former VFAT2-TURBO electronics [39] are installed. Only the third one is equipped with the readout system, recognizable in the picture by the green GEB board holding the front-end chips. On the very bottom of this detector, we see the optical outputs of the opto-hybrid, holding an FPGA to concentrate the incoming VFAT2 trigger and tracking data carried out individually to the off-detector electronics by the two visible orange optical link.

On the other side of the optical links, both data streams are recorded by a GLIB board sitting in a μ TCA crate. This is shown in Figure 96.



Figure 96: Off-detector electronics in a μ TCA crate

The two optical links carrying the trigger and tracking data are read out by the GLIB board. The two brown Lemo cables provide clock synchronization signal and an external trigger pulse to the read-out electronics, as it was the case in the cosmic muon stand. On the right side in the middle of the crate we see the MCH module, offering a star point to all the slots inside the μ TCA crate. The white network cable plugged into this MCH is the interface of this crate to the analysis computer, where data are recorded. On the very left side of the crate, we can see the two redundant power supply units of the crate.

As for the $10 \times 10 \text{ cm}^2$ prototype, the first step to commission this setup was the characterization of the front-end electronics inside this new environment. The comparators threshold scan and the latency scan are shown in Figure 97. On the top plot, we can see that the noise level is higher in the beam test setup than in the muon lab setup. As a result a higher threshold value is measurable. According to this plot, to perform effective measurements, the chosen value to exclude any false positives is a value of 35 on the ADC, which corresponds to a threshold voltage of -61 mV . This is typically a factor 2 larger than the threshold applied to the other GE1/1 detectors under test. The higher noise level could be explained by the fact that the grounding between the GEB board and the detector anode plane is not yet optimized. It is also important to say that the other GE1/1 prototypes were already tested in a similar test

beam a couple of weeks before which resulted in the fact that all their operation parameters were already well tuned and their grounding optimized.

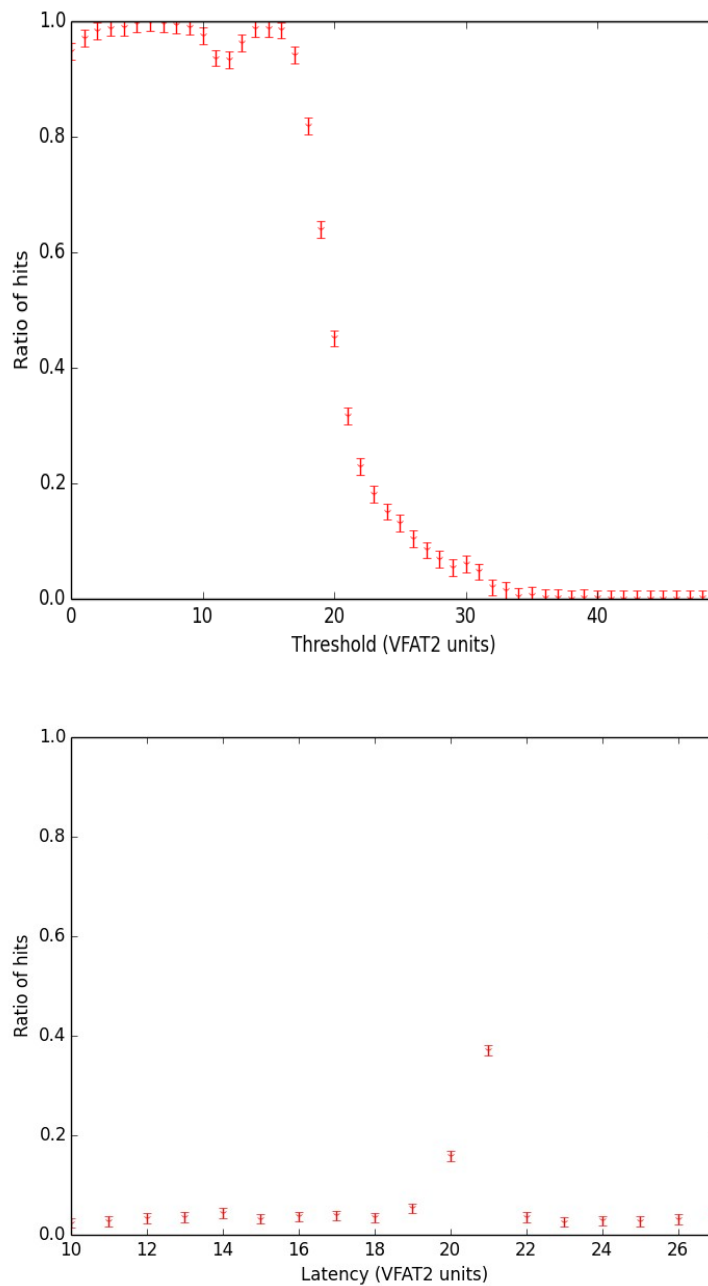


Figure 97: Comparator threshold scan (top) and latency scan (bottom) of the VFAT2

The second key parameter which needed characterization was the latency of the entire trigger chain, giving the depth of the stored hit data inside the circular buffer memory

of the front-end chip (seen on the bottom plot). Compared to the lab setup, where the trigger logic is generated by old NIM electronics, the 40 MHz clock period during which the trigger signal was received is much better defined, at a value of 21 slots. This corresponds to a trigger latency of 525 ns.

Once the DAQ chain was characterized and the parameters configured in the front-end chips and the read-out electronics and software, the data taking could start. Figure 98 shows the cluster size distribution for our Triple-GEM detector in the beam test.




Figure 98: Cluster size distribution inside the GE1/1 Triple-GEM detector

The cluster size distribution is quite different from the one recorded with the cosmic test bench (see Figure 94). The average cluster size amounts to 1.3 while it is about 1.7 in the cosmic test bench. Several factors may explain these differences. First the strip pitch is much smaller in the 10x10 cm² prototype : 0.4 mm while it is at least 0.6 mm in the GE1/1 prototype, depending on the position along the strips. In addition, during the beam tests, the particles are mainly perpendicular to the strips while in the cosmic test bench the muons follow the $\cos^2\theta$ distribution expected for atmospheric muons [40] at the ground. Another effect could be the presence of dead channels on the VFAT2 chips. This can be seen from Figure 99 which shows the number of hits recorded by each strip during a typical *muon run*.




Figure 99: Beam profile measurement showing a number of dead channels

This distribution is often referred to as the *beam profile*. The empty bins in the beam profile histogram indicate probable dead channels. This may be due to the fact that these VFAT2 chips have already served in several test beams and may have been damaged. At the time of this writing, deeper analyses are ongoing to quantify the influence of a dead channel in the spatial resolution, and what can cause such a mechanical or electrical failure.

Nevertheless, the cosmic muon stand experiment followed by the beam test experiment confirmed the validity of the CMS Triple-GEM + GEB + opto-hybrid + μ TCA + GLIB setup for the design of a DAQ system for GE1/1. This concludes the stage of “Proof of Concept” of the new DAQ system and it is now submitted to the CMS collaboration for approval, as part of the Technical Design Review.

CONCLUSIONS

This work is a contribution to the elaboration of the new DAQ system for the Triple-GEM gas detectors for the next upgrade of the muon spectrometer of CMS, starting from the early days of the project to the validation of the proof of concept with the first fully working prototype. As often in the field of research, the design process had its successes but also its dead ends. Every time, these results were nonetheless analyzed and mitigated to keep the design process efficient and to reach the best possible final state. Using this Triple-GEM detector technology implies a number of considerations for the read-out and processing electronics, and this was the main focus of this study. Furthermore, CMS is a complex experiment, based on technologies available at the moment of the Technical Proposal of 1994. Integrating a novel type of detection mode was, at that time, certainly not foreseen, which makes the integration of this new technology a true challenge. To fully understand the implications of upgrading the forward muon spectrometer, a number of questions were answered over the six chapters of this work.

The first question we asked ourselves, after having described the context of the CMS experiment, was the reason for upgrading the forward muon spectrometer of this experiment. As explained, this is related to the discovery, in 2012, of the particle for which this entire experiment was built. Now that the discovery is done, it is time to study its characteristics, and this is why the entire LHC complex will be upgraded. The improvements will essentially affect the accelerator luminosity, increasing the production rate, hence the need of a performance upgrade of the detector.

The second question addressed in this document, was related to what is actually needed to be improved. To define this, an exhaustive tour of the existing muon detection technologies was given, with a focus on their technical parameters, strengths and drawbacks. After this analysis, we demonstrated that none of the existing technologies, as built now, would be adequate to fulfill the challenges of the planned upgrades of the LHC. Worse, the currently installed Resistive Plate Chambers inside

the end-caps are already known for being too limited in detection efficiency.

Subsequently, the next question concerned the technology to be used. This is when a specific working group was set up inside the CMS collaboration, with the task to study the feasibility of using GEM chambers as a solution to this problem. Triple-GEM detector, in particular, have a number of characteristics which make them unsuitable for high granularity tracking applications. But this is not the aim of the end-cap muon spectrometer. The timing resolution of such a detector is better than 8ns which is excellent. The charge collection time ($< 100\text{ns}$), providing the rate capability of $\sim 1\text{ MHz}$, is good enough for the expected flux between the $1.6 < |\eta| < 2.4$ region. More important, the production cost for large Triple-GEM detector compared to other detectors of this performance is pushing towards the use of this technology for the future upgrade.

The GEM foil technology starts to be described in the literature at the end of the nineties. The Triple-GEM evolution of this detection technique is even younger, and has never been used at the scale of an entire detector. This is why the question arose of the technologies involved to read out such a large surface, knowing the constraints on the trigger path and the data volume. In phase with some other sub-detectors of the CMS experiment, the μTCA architecture was chosen to withstand the dramatic increase of the required read-out bandwidth in the electronics. In addition, the data concentration on the Triple-GEM detector itself was a second challenge, knowing the constraints in space, power and radiation.

Knowing the challenges and the technologies to be used, the last question was, of course, does it work? A number of developments were made to prove the concept of a data acquisition chain for gas detectors, based on an architecture which would allow seamless integration in the existing back-end DAQ of CMS. On the front-end side, a Triple-GEM detector prototype was assembled and a complete cosmic muon read-out system was built to show the effective feasibility of muon detection with this technology. Seeing the excellent results of some preliminary measurements, the setup will become a test bench for more characterization work to be done by future students after this current work. The proposed studies could be the influence of the gas composition and concentration to the detector efficiency, or the fine tuning of high voltage parameters and distribution over the gaps. One could also think about a study of the optimal Triple-GEM gap geometry, and the study of the spatial resolution as a function of the incoming cosmic muon angle, etc.

This promising outlook will be completed soon by the upcoming results of the same type of setup, installed in a test-beam facility at CERN by the Brussels R&D group. The main difference with the cosmic muon setup will be the deterministic nature (time, energy) of the incoming particles. Depending on these results, the status of the collaboration for the slice test will be known. In the case of a success, the main effort will be transferred to the development of the next generation front-end chip and the integration into the complex CMS detector operation framework with the imminent submission of the Technical Design Review to the CMS-wide collaboration and shortly afterwards to the LHC committee for final approval.

Bibliography

- [1]: The ATLAS collaboration, Letter of Intent for a General-Purpose pp Experiment at the Large Hadron Collider at CERN , CERN/LHCC 92-4, 1992
- [2]: The CMS collaboration, Letter of intent by the CMS Collaboration for a general purpose detector at LHC, CERN/LHCC 92-3, 1992
- [3]: Lyndon Evans, Philip Bryant, LHC Machine, Journal of Instrumentation, 2008 JINST 3 S08001, 2008
- [4]: CERN - Communication Group, CERN faq -The LHC guide, <https://cds.cern.ch/record/1165534/files/CERN-Brochure-2009-003-Eng.pdf>, 2009
- [5]: Roger Bailey & Paul Collier, Standard Filling Schemes for Various LHC Operation Modes (Revised), LHC-Project Note 323_Revised, 2003
- [6]: E. Todesco, C. Lorin, M. Bajko, Energy of the LHC after the 2013-2014 shutdown, http://www.researchgate.net/profile/Clement_Lorin/publication/257770742_ENERGY_OF_THE_LHC_AFTER_THE_2013-2014_SHUTDOWN/links/00b49525d5effe2d08000000.pdf, 2012
- [7]: Mike Lamont, Status of the LHC, Journal of Physics: Conference Series 347 (2012) 012001, 2012
- [8]: Q. Ingram, The Lead Tungstate Electromagnetic Calorimeter of CMS, PACS: R29.40.Mc; 29.40.Vj, 2007
- [9]: The CMS collaboration, The CMS experiment at the CERN LHC, Journal of Instrumentation, 2008 JINST 3 S08004, 2008
- [10]: The CMS collaboration, Precise Mapping of the Magnetic Field in the CMS Barrel Yoke using Cosmic Rays, Journal of Instrumentation, 2010 JINST 5 T03021, 2010
- [11]: The CMS collaboration, Technical Proposal for the upgrade of the CMS detector through 2020, CMS U1TDR, 2011
- [12]: G. Bauer et al., The data-acquisition system of the CMS experiment at the LHC, Journal of Physics: Conference Series 331 (2011) 022021, 2010
- [13]: The CMS Trigger and Data Acquisition Group, The CMS High Level Trigger, PACS: 13.85.-t, 07.05.Kf, 2005
- [14]: Kay Rehlich, DESY, MicroTCA® Update from a Physics Lab's Point of View, https://portal.slac.stanford.edu/sites/ad_public/events/mtca_llrf_jun_2012/Published_Documents/Introduction%20to%20mTCA.pdf, 2009
- [15]: Pigeon Point Systems, Introduction to μ TCA Hardware Platform Management, MicroTCA Conference 2009, 2009
- [16]: Vollrath Dirksen, NAT GmbH, MicroTCA-FAQ , PRAXIS PROFILINE – MicroTCA – August 2007, 2007
- [17]: Karol Bunkowski, PhD thesis - Optimization, Synchronization, Calibration

- and Diagnostic of the RPC PAC Muon Trigger System for the CMS detector, http://cds.cern.ch/record/1308715/files/TS2010_009.pdf, 2009
- [18]: P. Camarri, R. Cardarelli, A. Di Ciaccio and R. Santonico, Streamer suppression with SF₆ in RPCs operated in avalanche mode, ATLAS Internal Note MUON-NO-226, 1998
- [19]: Serge Duarte Pinto, Gas Electron Multipliers - Research and Development on GEM technology, <http://cds.cern.ch/record/1376429/files/Thesis-2011-Duarte.pdf>, 2011
- [20]: S. Bachmann, A. Bressan, L. Ropelewski and F. Sauli, A. Sharma, D. Mörmann, Charge amplification and transfer processes in the Gas Electron Multiplier, Nuclear Instruments and Methods in Physics Research Section A, Volume 438, Issue 2-3, p. 376-408., 1999
- [21]: S. Bachmann, A. Bressan, M. Capeans, M. Deutel, S. Kappler, B. Ketzer, A. Polouektov, L. Ropelewski, F. Sauli, E. Schulte, L. Shekhtman, A. Sokolov, Discharge studies and prevention in the gas electron multiplier (GEM), Nuclear Instruments and Methods in Physics Research A 479 (2002) 294–308, 2002
- [22]: Laura Franconi, Master thesis - Performance of large-size triple GEM detectors for the CMS Forward Muon Upgrade project, https://twiki.cern.ch/twiki/pub/MPGD/CmsGEMPapers/LauraFranconi_MasterThesis.pdf, 2011
- [23]: Geoffrey Mullier, Master thesis - Étude de détecteurs "Triple-GEMs" pour la mise à niveau du spectromètre à muons de l'expérience CMS pour les phases de haute luminosité du LHC, http://ulb-phys.akegroup.org/Memoire/Memoire_Mullier_Geoffrey_2012.pdf, 2012
- [24]: T.Maerschalk, High-eta Upgrade for CMS : Triple-GEM, Internal communication, 2012
- [25]: GEM for CMS collaboration, CMS Technical design report for the muon endcap upgrade: GE1/1 - The station 1 GEM project, CERN-LHCC-2014-NNN (Draft V01-07), 2014
- [26]: Andrew W. Rose, Greg Iles, Aaron Bundock, Sarah Greenwood – Imperial College, London, User guide for the MP7 Master Processor, Virtex-7, <http://www.hep.ph.ic.ac.uk/mp7/documents/UserGuideForTheMP7.latest.pdf>, 2014
- [27]: P.Aspell, G.Anelli, P.Chalmet, J.Kaplon, K.Kloukinas, H.Mugnier, W.Snoeys, VFAT2 : A front-end system on chip providing fast trigger information, digitized data storage and formatting for the charge sensitive readout of multi-channel silicon and gas particle detectors., Nuclear Science Symposium Conference Record, 2008. NSS '08. IEEE, ISBN:978-1-4244-2714-7, 2006
- [28]: Paul Aspell, GEMs for CMS : Electronics System, <https://indico.cern.ch/event/158341/session/1/contribution/45/material/slides/0.pdf>, 2012
- [29]: Paul Aspell, Future Possibilities for the Salto Architecture, <http://agenda.linearcollider.org/event/5107/material/slides/0?contribId=6>, 2011

- [30]: P. Aspell, M. De Gaspari, H. França, E. García, Super-Altro 16: a Front-End System on Chip for DSP Based Readout of Gaseous Detectors, Nuclear Science Symposium and Medical Imaging Conference (NSS/MIC), 2012 IEEE, ISBN: 978-1-4673-2028-3 , 2012
- [31]: Thierry Maerschalk, Timing Resolution Techniques- TOT and CFD, Internal communication, 2013
- [32]: CMS DAQ group, CMS MicroTCA Crate Specification and AMC Card Requirements, http://joule.bu.edu/~hazen/CMS/AMC13/The%20CMS%20uTCA%20Crate%20v0.9_esh.pdf, 2013
- [33]: Eric Hazen (Boston University), CMS AMC13 Hub Card, http://joule.bu.edu/~hazen/CMS/AMC13/AMC13_Short_Spec-1.pdf, 2012
- [34]: Charlie Hill, AMC13 XG Operational Specification, http://ohm.bu.edu/~chill90/op_specs/AMC13_Operational_Spec_05_13_2013.pdf, 2012
- [35]: Paschalis VICHODIS, First test results with the Gigabit Link Interface Board (GLIB), Journal of Instrumentation, 2011 JINST 6 C12060, 2012
- [36]: Vincent Bobillier et al., MMC mezzanine, Presentation, Third meeting of the xTCA Interest Group, September 2011 / TWEPP Vienna, 2011
- [37]: Altera, FPGA Configuration via Protocol, <http://www.altera.com/literature/wp/wp-01132-stxv-cvpcie.pdf>, 2011
- [38]: Marco Zecchin, Characterization of a Triple-GEM Detector Prototype for the CMS Muon Upgrade with GEM Detectors, http://iihe.ac.be/publications/Zecchin_Marco.pdf, 2014
- [39]: E. Graverini, A Large GEM detector prototype - Test beam results and analysis, , 2010
- [40]: K.A. Olive et al. (Particle Data Group), Cosmic rays chapter of Particle Data Group, Chin. Phys. C, 38, 090001 (2014), 2013

Table of Acronyms

ADC	Analog-to-Digital Converter
AMC	Advanced Mezzanine Card
ASIC	Application Specific Integrated Circuit
ATCA	Advanced Telecommunication Computing Architecture
ATLAS	A Toroidal LHC ApparatuS
BCF _n	Baseline Correction Filter
BX	Bunch Crossing
CBM	Control, Bias and Monitoring
CERN	European Organization for Nuclear Research
CM	Common Mode
CMOS	Complementary Metal-Oxide Semiconductor
CMS	Compact Muon Solenoid
CPU	Central Processing Unit
CSC	Cathode Strip Chamber
CU	Cooling Unit
DAC	Digital-to-Analog Converter
DAQ	Data AcQuisition
DC	Direct Current
DCS	Detector Control System
DESY	Deutsches Elektronen-SYnchrotron
DSP	Digital Signal Processor
DT	Drift Tube

DUT	Device Under Test
ECAL	Electromagnetic CALorimeter
EEPROM	Electrically Erasable Programmable Read-Only Memory
ENOB	Equivalent Number of Bits
FED	Front-End Driver
FMC	FPGA Mezzanine Card
FOM	Figure of Merit
FPGA	Field Programmable Gate Array
FRL	Front-end Readout Link
FRU	Field Replaceable Unit
GBT	GigaBit Transceiver
GEB	GEM Electronics Board
GEM	Gas Electron Multiplier
GLIB	Gigabit Link Interface Board
GMT	Global Muon Trigger
GOL	Gigabit Optical Link
GTX	Gigabit Transceiver
HLT	High Level Trigger
HCAL	Hadron CALorimeter
I ² C	Inter-Integrated Circuit (communication)
IP	Internet Protocol
IPMB	Intelligent Platform Management Bus
IPMC	Intelligent Platform Management Controller
IPMI	Intelligent Platform Management Interface
JTAG	Joint Test Action Group
LDO	Low Drop-Out (regulator)
LED	Light Emitting Diode
LEP	Large Electron-Positron Collider
LHC	Large Hadron Collider
LSn	Long Shutdown
LUT	Look Up Table
LV1	Level-1
LVC MOS	Low Voltage Complementary Metal-Oxide Semiconductor
LVDS	Low Voltage Differential Signaling

MAC	Media Access Control
MCH	μ TCA Carrier Hub
MCMC	μ TCA Carrier Management Controller
MCU	Micro Controller Unit
MMC	Module Management Controller
NIM	Nuclear Instrumentation Module
PAC	PATtern Comparator
PCB	Printed Circuit Board
PCI	Peripheral Component Interconnect
PICMG	PCI Industrial Computer Manufacturers Group
PMT	Photo Multiplier Tube
POST	Power On Self Test
PP	Payload Power
PROM	Programmable Read Only Memory
PRS	Packet Routing Switch
PSRR	Power Supply Rejection Ratio
PSU	Power Supply Unit
PU	Power Unit
QSFP	Quad Small Form-factor Pluggable (optical transceiver module)
RAM	Random Access Memory
RB	Readout unit Buffer
RCSM	Run Control and Monitoring System
RPC	Resistive Plate Chamber
RTM	Rear Transfer Module
RU	Readout Unit
RX	Receiver
SAS	Serial-Attached SCSI (Small Computer System Interface)
SATA	Serial Advanced Technology Attachment
SFP	Small Form-factor Pluggable (optical transceiver module)
SLAC	Stanford Linear Accelerator Center
SOT	Small Outline Transistor
SNDR	Signal over Noise and Distortion Ratio
SPI	Serial Peripheral Interface
SRAM	Static Random Access Memory

TC	Tail Cancellation
TIB	Tracker Inner Barrel
TID	Tracker Inner Disk
TEC	Tracker End-Cap
TMB	Trigger Mother Board
TOB	Tracker Outer Barrel
TOS	Type Of Service
TOT	Time Over Threshold
TOTEM	TOTAL Elastic and diffractive cross section Measurement
TQFP	Thin Quad Flat Pack
TTC	Trigger Throttling System
TTL	Transistor-Transistor Logic
TX	Transmitter
UART	Universal Asynchronous Receiver-Transmitter
UDP	User Datagram Protocol
VME	Versa Module Europa
XAUI	10 Gigabit Attachment Unit Interface
ZS	Zero Suppression
μHAL	μTCA Hardware Abstraction Layer
μTCA	Micro Telecommunication Computing Interface