universität wien

DIPLOMARBEIT

Titel der Diplomarbeit

„ A hardware abstraction layer for the MicroTCA-based Global Trigger for the CMS experiment at CERN “

Verfasser

Dinyar Sebastian Rabady

angestrebter akademischer Grad

Magister der Naturwissenschaften (Mag. rer. nat.)

Wien, 2012

Studienkennzahl lt. Studienblatt:	A 411
Studienrichtung lt. Studienblatt:	Diplomstudium Physik UniStG
Betreuerin / Betreuer:	Univ.-Doz. Dr. Dipl.-Ing. Claudia-Elisabeth Wulz

Abstract

The Large Hadron Collider (LHC) based at CERN near Geneva collides proton bunches at a rate of 40 MHz. Each collision produces approximately 1 MB of data in the Compact Muon Solenoid (CMS) detector.

In order to reduce this event rate to a more manageable amount, a complex filter system was developed. The first stage of this filter is the so-called Level-1 trigger. This system reduces the incoming event rate to 100 kHz which can then be analyzed and filtered further in a massive computing farm.

The LHC is scheduled to be upgraded to provide collisions with even more particles involved thus making an upgrade of the Level-1 trigger necessary.

This thesis is concerned with the upgrade plans of the Global Trigger (GT) project and ultimately lead to the development of a hardware abstraction layer (HAL) which can provide remote register-level access via Ethernet as well as abstract items to represent the information stored in the registers.

Finally a study of the Global Muon Trigger (GMT) efficiency is presented.

Kurzfassung

Der Large Hadron Collider (LHC) am CERN bei Genf produziert mit einer Frequenz von 40 MHz Teilchenkollisionen. Jede dieser Kollisionen benötigt nachdem sie vom Compact Muon Solenoid (CMS) Detektor aufgezeichnet worden ist etwa 1 MB an Speicher. Um diese enorme Menge an Daten zu reduzieren wurde ein komplexes Filter-System entwickelt. Die erste Stufe dieses Systems nimmt der Level-1 Trigger ein, der die Rate an aufgezeichneten Kollisionen auf 100 kHz reduziert. Diese Kollisionen können anschließend von einer groen Rechenfarm analysiert und weiter gefiltert werden.

Der LHC wird in absehbarer Zukunft ausgebaut werden um Kollisionen mit noch mehr beteiligten Teilchen zu produzieren, was eine Verbesserung des Level-1 Triggers notwendig macht.

Diese Arbeit beschäftigt sich mit den Plänen zu diesem Ausbau innerhalb des Global Trigger (GT) Projekts und führte schlussendlich zur Entwicklung einer hardware abstraction layer (HAL), die entfernten Zugriff auf Hardware-Register über Ethernet erlaubt wie auch abstrakte Elemente zur Verfügung stellt um die Information in den Registern zu repräsentieren.

Abschließend wird eine Studie über die Effizienz des Global Muon Trigger (GMT) präsentiert die zu Verbesserungen für die Datennahme ab dem Jahr 2011 geführt hat.

Contents

Chapter 1

Introduction

Although the Standard Model of particle physics has so-far withstood experimental tests up to very high precision, there is reason to believe it to be an effective theory, embedded in a more fundamental description of nature: It is an ungainly theory with a comparatively large number of free parameters and astronomical observations have suggested that it cannot account for a large part of the universe's energy. Furthermore it fails to describe gravity, one of the four fundamental forces of nature. Several improvements and successors for the Standard Model have been offered, which need to be tested experimentally. Moreover the currently favored mechanism for the leptons to acquire mass requires the production of the so-called Higgs boson, which has until now not been observed.

The Large Hadron Collider (LHC), a particle accelerator based at CERN, near Geneva was constructed to help probe the limits of the current understanding of high-energy physics. It is expected to find signs of new physics as well as possibly the Higgs boson. An overview of the physics work done in 2011 at LHC with an emphasis on the results from the Compact Muon Solenoid (CMS) is given in chapter 2.

The LHC itself is presented in chapter 3 along with the experiments hosted at its interaction points. One of these experiments, CMS, is presented in more depth in chapter 4. As the LHC collides proton bunches with a frequency of 40 MHz it is necessary to implement an event filtering system. The first stage of this system – the Level-1 trigger – is introduced in chapter 5.

In the future LHC will be upgraded to be significantly more powerful, necessitating an improved trigger system. Chapter 6 introduces the technology designated to be used in this upgrade and discusses the advantages as well as possible drawbacks and required modifications of the trigger control and monitoring systems introduced through this change.

One of these modifications is the introduction of a new hardware abstraction layer (HAL) to act as the interface between current control and

monitoring software and the new hardware. A proposal for such a HAL in the upgraded system is shown in chapter 7. Such a HAL will allow engineers to monitor, debug, and maintain the Global Trigger (GT) hardware remotely. It furthermore translates values stored in registers to abstract entities such as "Prescale Factor 3" which can be used by physicists when accessing the hardware. Chapter 7 additionally presents an analysis of this HAL's read and write performance as well as its stability during runs of length comparable to runs of data taking at CERN. Moreover an outlook regarding future developments for this system is given.

Finally a study of the Global Muon Trigger (GMT) efficiency is presented in chapter 8. This study resulted in improvements concerning the GMT's quality assignment being introduced for the 2011 CMS run.

Chapter 2

Physics at the TeV scale

The LHC has been designed primarily as a "discovery machine". Precision measurements, as performed for example at LEP, the former electron-positron collider at CERN, will, however, also be possible due to the large expected integrated luminosities. Events will be less clean than at an e^+e^-machine, however, as collisions occur between partons whose nature and exact energy are a priori undefined. Constituents not involved in the hard scattering lead to spectator jets, which complicate the data analysis. The unparalleled collision energy and luminosity provided by the LHC enable physicists to explore a new energy regime, making it possible to validate the Standard Model even further and to search for signals of new theories such as supersymmetry. The LHC accelerator will be described in more detail in chapter 3.

In the following an overview of the possible physics program at the LHC as well as results from the years 2010 and 2011 with an emphasis on CMS will be given. The topics are ordered according to the organization of the CMS analysis groups.

2.1 QCD physics

The LHC, being a hadron collider, offers a rich environment for studies of quantum chromodynamics (QCD), the theory of the strong interaction governing the nuclear force. Due to the LHC's high center-of-mass-energy and luminosity theoretically well-understood quantities can be studied with very high accuracy.

Among these is the di-jet invariant mass distribution resulting from two hard-scattered protons. Typically a hard-scattering event leads to a least two primary jets from the involved partons. The di-jet invariant mass, the invariant mass of the two highest-energy jets involved in the event, is then defined as $m = \sqrt{(E_1 + E_2)^2 - (p_1 + p_2)^2}$, with $(E_{1,2}, \overrightarrow{p_{1,2}})$ being the 4-vector of the two leading jets. This distribution has been measured at various

colliders for a wide range of energies and luminosities, but could indicate new physics if differences from the expected Standard Model spectrum were found.

Furthermore, CMS's first prominent and unexpected discovery was made in the area of QCD physics during proton-proton collisions in 2010, as presented in [1]. In this result azimuthal correlations of two charged particles with $2 < |\Delta\eta| < 4.8$ were examined which lead to the discovery of a long-range ridge-like structure at the near-side – i.e. $\Delta\phi \approx 0$ when considering high-multiplicity events (figure 2.1). This structure resembles features observed in heavy-ion experiments but had thus far not been seen in proton-proton collisions.

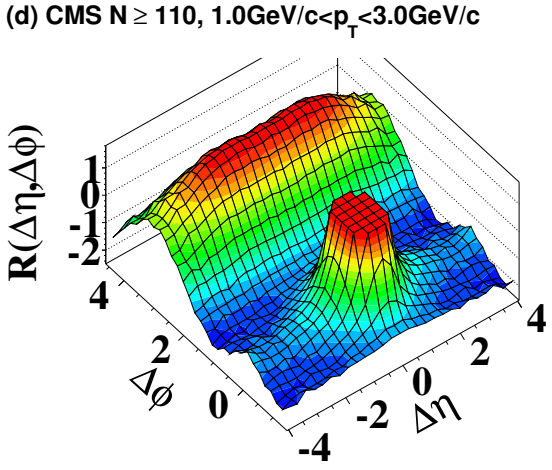

Figure 2.1: A ridge-like structure at the near-side can clearly be seen for a wide range of differences in pseudo-rapidity when investigating azimuthal correlations for two charged particles in high-multiplicity environments. [1]

2.2 Forward physics

Forward physics has traditionally been the study of diffractive and elastic scattering events. As will be described in chapter 4, CMS is very well suited for such studies due to its calorimetry offering almost hermetic coverage. Tagging energetic jets with high rapidities is also an excellent tool to search for new particles through central exclusive production.

At the moment dedicated runs with the LHC's quadrupole magnets in a low-luminosity configuration are used to investigate diffractive events, as pileup (i.e. multiple individual collisions in one bunch crossing) has to be avoided in order to be able to observe e.g. gaps in pseudo-rapidity. It is

however planned to use both high- and low-luminosity bunches in the same accelerator fill in the future.

A phenomenon CMS is well-suited to investigate is jet production from hadron-hadron collisions in the forward regions due to the above mentioned hermetic coverage for the calorimeters. This has been investigated in [2] which finds good agreement of the inclusive single-jet forward cross-section with perturbative QCD (pQCD) as implemented in several event generators while discovering deviations from the measured spectrum in next-to-leading-order (NLO) calculations as well as in all event generators when calculating cross-sections in models for multijet production.

2.3 B physics and Quarkonia

While the LHC hosts a dedicated B physics experiment, LHCb, which has been specifically designed to be able to use a large number of B decay triggers at low thresholds (chapter 3), it is still possible for general purpose experiments such as CMS to contribute to this area competitively.

Studies of J/Ψ, B_s, B_c, and B_d to measure their masses, branching-ratios, and cross-sections are currently being undertaken at CMS. Especially the decay channels for $B_s \to \mu^+\mu^-$ and $B_d \to \mu^+\mu^-$ are strongly suppressed, but very sensitive to new physics. As a study of these channels requires a superior di-muon resolution and an efficient muon trigger, CMS is very well suited for this task and has already contributed significantly to an improved rate limit for the $B_s \to \mu^+\mu^-$ decay, its individual measurement being only slightly higher than LHCb's. [3]

Conversely CMS will not be able to compete with LHCb when investigating hadronic decay modes, as the required low-transverse momentum triggers are only possible in a lower luminosity environment than is required for Higgs or supersymmetry searches.

Just as for B physics a good di-muon resolution is required for the study of Quarkonia – bound states of same-flavor quark anti-quark pairs. Conventionally only $q_c\bar{q_c}$ and $q_b\bar{q_b}$ are labeled as Quarkonia, because bound states made up of lower mass quarks (up, down, and charm) are not observed as pure states in nature due to mixing of mass eigenstates. The Quarkonium state made up of top quarks does not exist since it requires more time to form than the top quark's lifetime.

2.4 Electroweak physics

Electroweak physics is a principal topic in CMS. Apart from its importance for understanding the nature of electroweak symmetry breaking (section 2.6), it can contribute to discoveries, as several quantities such as the W boson's mass are susceptible to new physics. It is also useful for detector

calibration and high-precision luminosity measurements, in particular during early running of the experiment. Other quantities which can be used to test the Standard Model are the W and Z production cross sections along with their ratios. The production of W and Z bosons is dominated by the Drell-Yan process which works by annihilation of a quark anti-quark pair producing a neutral gauge boson (either γ or Z) in the case of $u\bar{u}$ or $d\bar{d}$ or a W boson for $u\bar{d}$ and $d\bar{u}$. These can in turn decay to two leptons. [4]

A study of production cross sections of W and Z cross sections is presented in [5]. The results shown are dominated by systematic errors mainly originating from uncertainties regarding the integrated luminosity which cancels out for the measurement of the ratios. All measurements are in agreement with the Standard Model.

A measurement of weak mixing angles in Drell-Yan processes can uncover evidence of new physics as well. The annihilation process involved here could potentially uncover an additional neutral gauge boson or show deviations of elementary fermion couplings to the neutral electroweak bosons from predictions of the Standard Model. A study presented in [6] however reports results consistent with Standard Model calculations as well.

2.5 Top physics

Studying properties of the top quark such as its exact mass, production cross section, and spin is an important part of CMS's physics program. Before the LHC era these properties have only been examined at the Tevatron, a $p\bar{p}$ collider. Top quarks are mainly produced in $t\bar{t}$ pairs via the strong interaction in pp collisions. While the $t\bar{t}$ production at Tevatron was dominated by quark antiquark annihilation, at LHC top quark pairs are mainly produced through gluon fusion processes. For this reason measurements at the LHC allow new tests for the current understanding of the $t\bar{t}$ production measurement.

The top quark mainly decays through $t \rightarrow Wb$. Depending on whether the emerging W bosons decay to leptons or to quarks, which subsequently hadronize to form jets , top quark pair events can be categorized into the all-hadronic, the di-lepton, or the lepton+jets channels. Furthermore, the mass of the top quark constitutes an important parameter in the Standard Model and also indirectly allows limits to be set for the Higgs boson's mass.

In [7] a study of the $t\bar{t}$ production cross section in the di-lepton channel as well as a top quark mass measurement are presented. The production cross section is reported as $\sigma_{t\bar{t}} = 168 \pm 18(\text{stat.}) \pm 14(\text{sys.}) \pm 7(\text{lumi.})$ pb which is in good agreement with Standard model predictions. The top mass was measured as $175.5 \pm 4.7(\text{stat.}) \pm 4.6(\text{sys.})$ GeV/c^2 after one year of data taking which agrees with the current best value of $173.3 \pm 0.6(\text{stat.}) \pm 0.9(\text{sys.})$ GeV/c^2 that was measured at Tevatron. A similar study using the lepton+jets

channel is presented in [8]. Here the reported production cross section is $\sigma_{t\bar{t}} = 150 \pm 9(\text{stat.}) \pm 17(\text{sys.}) \pm 6(\text{lumi.})$ pb which was combined with the earlier di-lepton result to obtain a cross section of $\sigma_{t\bar{t}} = 154 \pm 17(\text{stat.} + \text{sys.}) \pm 6(\text{lumi.})$ pb.

The top quark production cross section itself was measured in [9] to be $83.6 \pm 29.8(\text{stat} + \text{syst}) \pm 3.3(\text{lumi})$ pb.

2.6 Standard Model Higgs boson

The Higgs boson is the last remaining particle to be discovered for the Standard Model to be complete. It is the particle predicted by the simplest mechanism for electroweak symmetry breaking. Electroweak symmetry breaking in turn is the mechanism by which the fermions as well as the W and Z bosons acquire their masses. As such it was one of the main driving forces in design considerations for the CMS detector.

The different search channels' sensitivity depend on several factors such as the best attainable mass resolution for the Higgs boson, the sensitivity to its resulting decay products, as well as the branching ratio for the Higgs boson and background from other processes at the relevant Higgs mass hypothesis. For Higgs masses below 120 GeV the best sensitivity is attained in the $H \to \gamma\gamma$ channel, which, while only produced with a small branching ratio, is obscured by only moderate background at this energy regime. At these energies the decays $H \to \tau\tau$ and $H \to bb$ are difficult to detect due to their large background. From 120 GeV to 200 GeV the best sensitivity can be achieved by searching in the $Z \to WW$ decay channel, while the branching ratio for $H \to ZZ$ is large above 200 GeV and thus searches for $H \to ZZ \to 2\ell 2\nu$ and $H \to ZZ \to 4\ell$ are used in this energy regime.

In both the $H \to ZZ \to 4\ell$ and $H \to \gamma\gamma$ decay channels, superior mass resolution can be attained through the reconstruction of the four-lepton or di-photon final states. For this reason CMS was equipped with a high resolution electromagnetic calorimeter as well as a powerful muon system and muon trigger, as will be explained in further detail in chapter 4.

Recent results (as reported in [10]) indicate a slight excess signal over the expected Standard Model background at a mass of 124 GeV. The local significance for this excess is 3.1σ which translates to a global significance of 1.5σ to 2.1σ depending on the examined mass range (figure 2.2).

2.7 Supersymmetry

Although the Standard Model has not been seriously challenged experimentally so far it is expected that a new theory will be required to reconcile several perceived inconsistencies in the current model, chief among them the failure to include a description of gravity. A theory superseding the

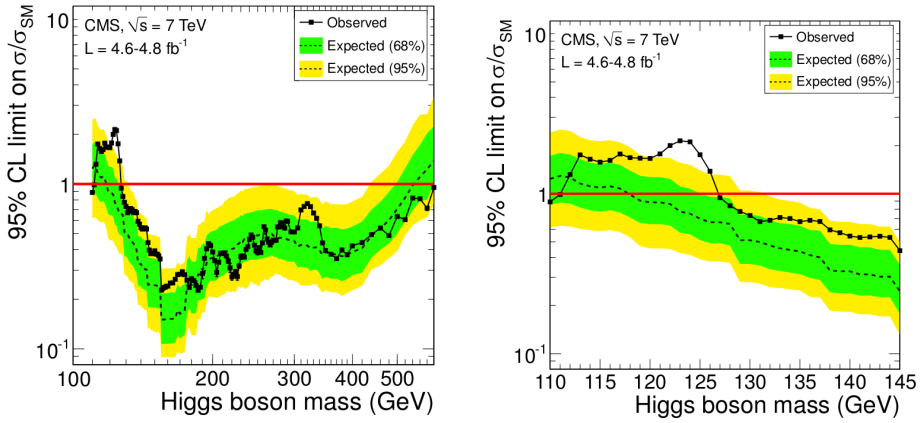

(a) The 95 % CL upper limits for the mass range from 110 GeV to 600 GeV.

(b) The 95 % CL upper limits for the mass range from 110 GeV to 145 GeV.

Figure 2.2: The 95 % CL upper limits for the Standard Model Higgs boson hypothesis mass. The dashed line indicates the expected median results for the expected Standard Model background, the green and yellow bands indicate the ranges that are expected to contain 68 % and 95 % respectively of all results in this model. The solid line represents the observed values.[10]

current Standard Model would ideally also be considered "more fundamental" if it could explain the arbitrary parameters currently required in the Standard Model.

The most studied candidates for such a theory can be grouped under the concept of supersymmetry (SUSY). SUSY extends the current symmetries of quantum field theory with an additional one relating bosonic and fermionic fields with each other. An attractive feature of SUSY is that it could solve both the hierarchy problem and unify the gauge couplings of the Standard Model at an intermediate energy scale. However, to incorporate SUSY in the current Standard Model requires a doubling of the particle content as every fermion would need a bosonic superpartner and vice versa. None of the currently known particles are superpartners of each other, which means that SUSY must be broken at some energy scale.

In order for SUSY to be an attractive extension of the Standard Model the mass differences between particles and their superpartners have to be on the order of TeV which would allow evidence of SUSY to be discovered at the LHC.

Studies presented in [11], [12], and [13] among others have so far not found evidence for supersymmetric partners of current particles. However, the parameter space for SUSY models could be further constrained through these results.

2.8 Exotica

The physics program of the general purpose experiments at the LHC also includes searches for so-called "exotic physics". Large extra dimensions or quark compositeness are examples of topics in the field of exotica.

Large extra dimensions arise in several theories by embedding the four-dimensional space-time of standard quantum field theories in a higher dimensional setting. Evidence for these large extra dimensions could appear through an excess of events at large di-muon or di-electron invariant masses. A search for such an excess is presented in [14], which however reports results consistent with Standard Model expectations.

Another feature predicted by some theories with extra dimensions is the production of microscopic black holes at energies accessible to the LHC. Events containing these would manifest themselves in a decay lacking a preferred direction or particle type. A search for black holes was carried out in [15] which could set the most restrictive limit on black hole production reached thus far at hadron colliders.

Other theories extending the current Standard Model propose that quarks are bound states of more fundamental entities. These models of quark compositeness could explain the number of quark generations, quark charges, and quark masses. Examining di-jet angular distributions in pp collisions can provide clues to a composite nature of quarks. Such a search has been performed in [16], yielding results consistent with QCD predictions.

Chapter 3

The Large Hadron Collider

The Large Hadron Collider (LHC) is the largest and most powerful particle accelerator built in the history of mankind. Most of the year it provides proton-proton collisions (p-p). About one month, usually at the end of the year, is reserved for heavy ion runs, in which ionized lead atoms are made to collide [17]. For 2012 it is also planned to attempt collisions of protons with lead ions enabling physicists to gain more knowledge about the behavior of the quarks and gluons in lead using the proton as a probe. [18]

The LHC is designated a "discovery machine" as its primary goal is to complete the standard model by discovering the Higgs Boson and finding new physics such as SUSY. While lepton colliders provide cleaner events than hadron machines, as the colliding particles themselves are fundamental, the fact that up to now only electrons can be used for these accelerators introduces difficulties when increasing the collider's center-of-mass energy. This is due to synchrotron radiation which scales with $\frac{(E/m)^4}{\rho^2}$ and already reached 18 MW at LEP2 as it operated at a center-of-mass energy of 200 GeV. [19] Above ρ stands for the collider's bending radius and m stands for the particle's mass. This means that more massive particles greatly reduce the synchrotron radiation introduced through higher energies. As the proton is heavier than the electron by a factor ~ 2000 it is a natural candidate to be used in discovery machines. Historically hadron colliders have been used a discovery machines while lepton accelerators were used to explore previously discovered physics in more detail. The development of particle colliders with time can be seen in figure 3.1.

The LHC accelerator is contained in an almost circular tunnel on average 100 m below the surface and measuring 27.6 km in circumference. Two beams are injected into the ring in opposite directions, providing collisions at four interaction points. At each such point an experiment is housed. The LHC is the last stage in an accelerator chain, which begins with the linear accelerator LINAC2 as well as the first circular accelerator Proton Synchrotron Booster in the case of proton beams, and the linear accelera-

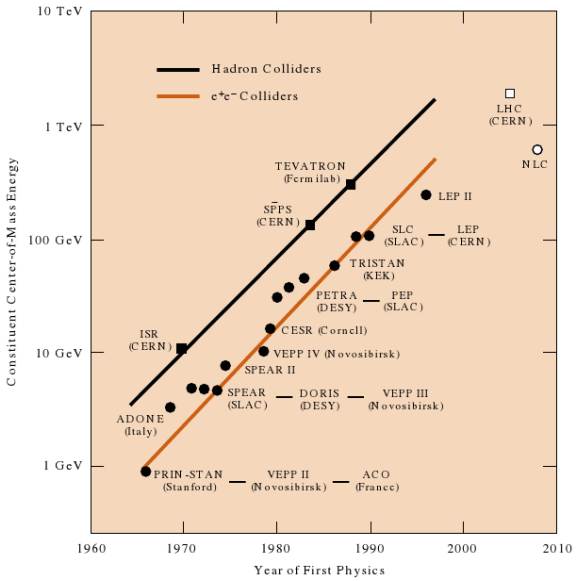

Figure 3.1: The development of particle colliders with time. As can be seen the center-of-mass energy increased almost exponentially until recently. Also to be seen is that hadron colliders have continually provided significantly higher center-of-mass energies than their lepton colliding counterparts at any given time, although not all the energy is available for hard collisions [20].

tor LINAC3 together with the Low Energy Ion Ring (LEIR) for heavy ion runs. The beam is then fed into the Proton Synchrotron (PS) which was CERN's first ring collider and which still provides beams to experiments in the so-called East Area apart from being used as the pre-accelerator for the Super Proton Synchrotron (SPS) and thus ultimately for the LHC. The SPS is the last step in the pre-accelerator chain of the LHC, but it provides beams to other experiments as well, such as the CERN Neutrinos to Gran Sasso (CNGS) neutrino beam for the Gran Sasso underground laboratory or fixed-target experiments in the so-called North Area of CERN. Two transfer lines then connect the SPS to the LHC (figure 3.2).

The LHC is designed to ultimately provide experiments with p-p collisions at a luminosity of $1 \times 10^{34}\,\mathrm{cm}^{-2}\,\mathrm{s}^{-1}$ and a center-of-mass energy of 14 GeV, which means that protons in each beam would be accelerated to 7 GeV. However, until 2011 the LHC was operated at a center-of-mass energy of 7 GeV which was increased to 8 GeV in 2012. Similarly, the luminosity provided by the LHC has been ramped up continuously since the start-up. In 2011 the instantaneous luminosity increased from less than $1 \times 10^{30}\,\mathrm{cm}^{-2}\,\mathrm{s}^{-1}$ to about $3.5 \times 10^{33}\,\mathrm{cm}^{-2}\,\mathrm{s}^{-1}$ as can be seen in figure 3.3a.

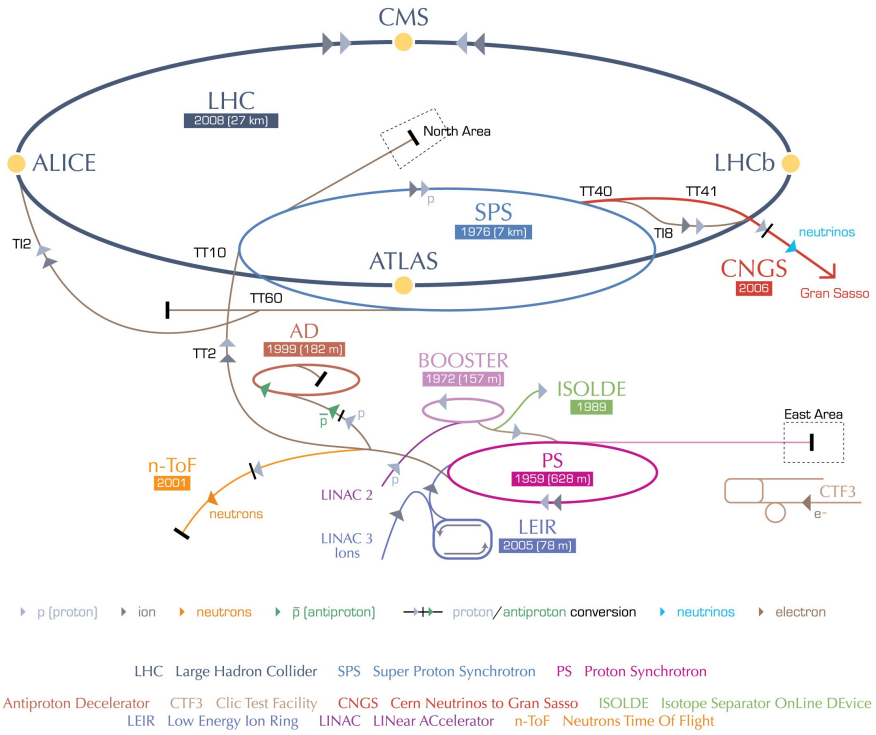

Figure 3.2: The CERN accelerator complex. As can be seen the LHC is fed with protons or heavy ions by the Super Proton Synchrotron (SPS). The SPS itself provides beams to several non-LHC experiments, among them CERN Neutrinos to Gran Sasso (CNGS). The SPS receives accelerated protons from the Proton Synchrotron (PS). PS was CERN's first ring accelerator and it continues to supply smaller experiments with protons. Its proton fill is accelerated by the PS Booster that in turn is filled by the linear accelerator LINAC2. Its heavy ion fill stems from the linear accelerator LINAC3 via the Low Energy Ion Ring (LEIR). (©CERN, 2008)

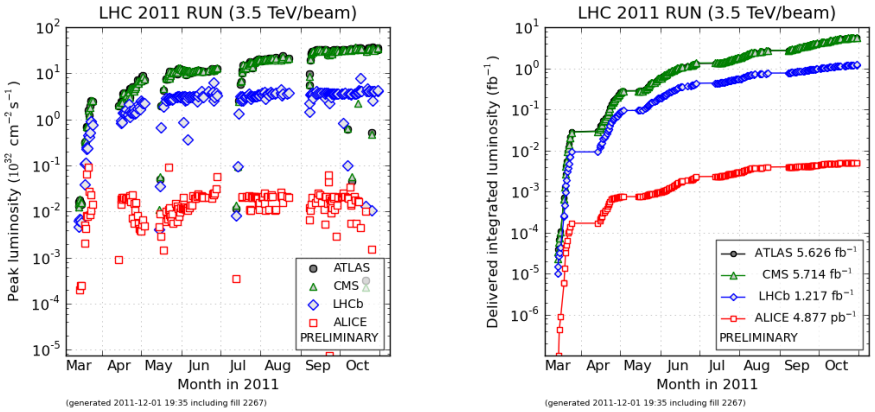

(a) The instantaneous luminosity could be increased over several magnitudes during the LHC's run in 2011.

(b) Similarly to the instantaneous luminosity the delivered integrated luminosity grew massively during 2011. The flat sections indicate technical stops of the LHC.

Figure 3.3: Luminosity plots for the LHC's 2011 run. [21]

3.1 Experiments at the LHC

To make full use of the collisions provided by the LHC, four main experiments along with smaller ones sharing their interaction point were designed. While ATLAS and CMS are both general purpose experiments, designed to discover any new physics and to find the Higgs boson, ALICE was constructed mainly to be used in heavy ion runs, and LHCb is an experiment specializing in studies of heavy flavour physics, in particular beauty, and matter-antimatter asymmetry.

3.1.1 ALICE

A Large Ion Collider Experiment (ALICE) has been built specifically to cope with the conditions produced during heavy ion runs. Therefore it is equipped with an advanced Time Projection Chamber (TPC) which can track a very high number of particles while only introducing a minimum of dense matter into their paths.

This is very useful as individual heavy ion collisions occur at low energies compared to p-p collisions while the frequency of collisions is much smaller mitigating the problem of long drift times in a TPC. For this reason the instantaneous luminosity during p-p runs must be reduced at ALICE's interaction point. [22]

3.1.2 ATLAS

A Toroidal LHC Apparatus (ATLAS) is the largest of the four LHC experiments, measuring 25 m in diameter. A distinguishing feature are its toroid magnets.

It has the same physics program as CMS. The two detectors were built with an emphasis on different detector technologies and magnetic field configurations such that they complement each other. Thus discoveries could be confirmed more or less independently by cross-checking results obtained with the two detectors.

LHCf

LHC forward (LHCf) is designed as a forward experiment sharing ATLAS's interaction point. It detects particles moving almost parallel to the beamline, witht the goal to simulate cosmic rays in laboratory conditions. It is used to measure properties of the production of neutral pions during collisions. This data can then be compared with shower models used to estimate the energy of extremely high energy cosmic rays. [23]

3.1.3 CMS

In contrast to ATLAS, Compact Muon Solenoid (CMS) is very heavy, weighing 12 500 t, with a diameter of 15 m.

Prominent features are its extremely strong magnetic field of 4 T produced by a superconducting solenoid as well as its crystal electromagnetic calorimeter and advanced muon system.

CMS will be described in detail in chapter 4.

TOTEM

Similar to LHCf for ATLAS, TOTal cross section, Elastic scattering and diffraction dissociation Measurement at the LHC (TOTEM) is a forward experiment sharing CMS's interaction point. TOTEM's main goal is the investigation of elastic scattering and diffraction events of protons in the interaction point. This could improve current knowledge on the proton's structure and also give better measurements of the p-p cross-section. [24]

3.1.4 LHCb

The LHC beauty (LHCb) detector has an unusual form for a detector at a collider experiment in being highly asymmetric with respect to the interaction point. This design decision was influenced by the fact that $b\bar{b}$ pairs are boosted and thus predominantly produced in a forward cone which allows the experiment to only be sensitive on one side of the interaction point.

LHCb's most prominent feature is its highly sophisticated Vertex Locator (VELO) which can be moved to only 7 mm distance from the proton beam once beam conditions are declared stable. This allows LHCb to obtain very high resolution measurements for the vertex position. This value is very important for identifying decays from b quarks through the detection of secondary vertices. [25]

MoEDAL

The Monopole and Exotics Detector At the LHC (MoEDAL) experiment is housed inside LHCb's experimental cavern. Its purpose is the direct search for magnetic monopoles and highly ionizing Stable Massive Particles (SMPs).

MoEDAL is placed around LHCb's VELO and features so-called Nuclear Track Detectors (NTDs). These NTDs show characteristic damage when being traversed by highly ionizing particles. Current predictions suggest that the NTDs will be replaced every year due to these changes produced by known particles. [26]

Chapter 4

The Compact Muon Solenoid

4.1 Physics goals of CMS

The Compact Muon Solenoid (CMS) is designed as a general-purpose experiment, with an initial emphasis on searches for the Higgs boson. This is reflected in its superior calorimetry and muon systems which are a result of many Higgs decay channels containing muons, electrons or photons. Since the beginning of data taking in early 2010 a diverse range of studies has been undertaken, from the search for the Higgs boson or SUSY, studies of QCD including Quarkonia to examine the properties of b and c $q\bar{q}$ states, to searches for exotic physics such as SMPs. Apart from the Higgs boson search, the main goal of CMS is to find how the standard model must be extended in order to solve the hierarchy problem and to accommodate for example the dark matter of the universe or gravity.

4.2 Geometry

CMS's dominating feature is the large superconducting solenoid capable of producing a 4 T magnet field in the interaction zone. It has a diameter of 5.9 m making it possible to contain both the inner tracking system and most of the calorimetry.

The magnetic field lines run parallel to the beam axis in the central region. The flux is closed by using a return yoke made from iron, which contains four muon stations made of resistive plate chamber (RPC) detectors as well as drift tubes (DTs) in the barrel or cathode strip chambers (CSCs) in the endcaps.

A schematic of CMS is given in figure 4.1.

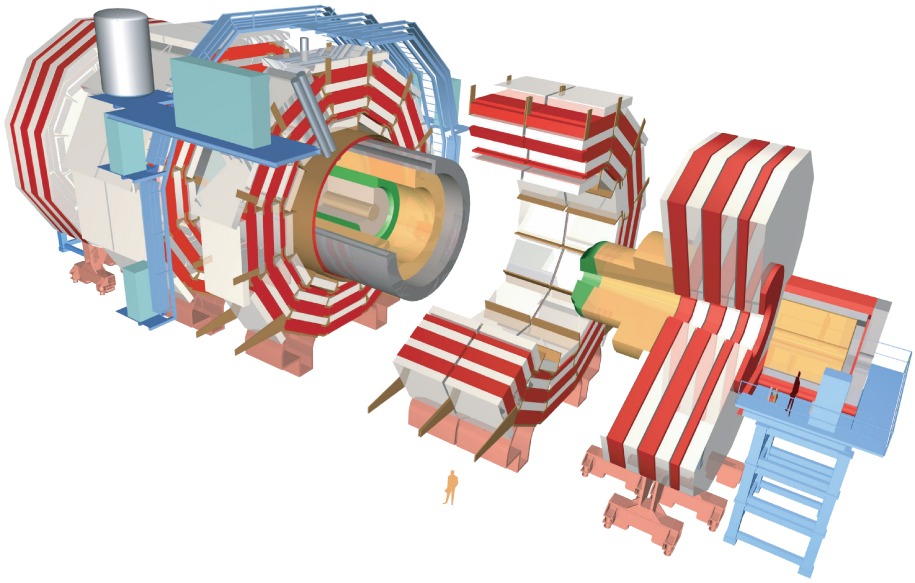

Figure 4.1: Exploded view of the CMS detector. In gray the superconducting solenoid can be seen, surrounded by the red return yoke and white muon stations. Contained by the solenoid are the hadron (orange) and electromagnetic (green) calorimeters, as well as the inner tracking system.

4.3 Coordinate system

CMS's coordinate system is centered at the nominal interaction point of the LHC particle beams, with the x-axis pointing in the direction of LHC's center, the y-axis pointing upwards. The z-axis then points in beam direction.

Due to CMS's cylindrical shape it is useful to define an azimuthal angle ϕ which is zeroed at the x-axis and a polar angle θ which is zeroed at the z-axis. However, it is common to use the so-called pseudo-rapidity instead of the polar angle for high-energy physics experiments. This is useful as the pseudo-rapidity is invariant under Lorenz transformations while being independent from the mass of the considered particle. The pseudo-rapidity is defined as

$$\eta = -\ln\left(\tan\left(\frac{\theta}{2}\right)\right) \tag{4.1}$$

which is approximately the relativistic rapidity in z-direction

$$\xi = \frac{1}{2}\ln\left(\frac{|\mathrm{p}| + \mathrm{p_z}}{|\mathrm{p}| - \mathrm{p_z}}\right) \tag{4.2}$$

for either small masses or big momenta for which $\mathrm{E} \approx \mathrm{p}$ holds.

4.4 Inner tracking

CMS's inner tracker is designed to provide highly accurate reconstructions of particle tracks in order to be able to measure momenta and to distinguish primary and secondary vertices. To do so requires a minimum resolution of about 10 µm. The tracker is built around the interaction point with a length of 5.8 m and a diameter of 2.5 m. Due to the extremely high occupancy of 1000 particles per bunch crossing at LHC's design luminosity a pixel tracker is used in the inner layer. The pixel tracker provides superior resolution and can cope with a very high track multiplicity environment, but requires a staggering 66 million read-out channels. The outer layers of the tracker are constructed with silicon strip modules which provide lower resolution, but require far fewer read-out channels (figure 4.2).

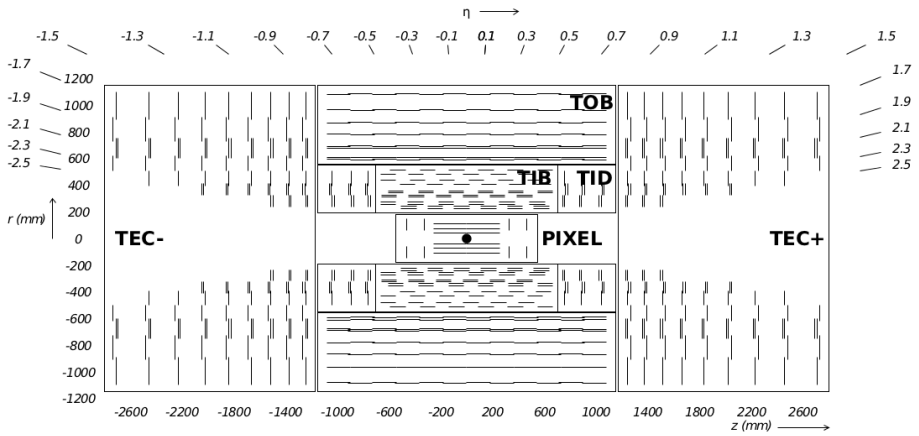

Figure 4.2: The inner tracker of CMS. Each line indicates a detector module. Double lines represent so-called "stereo" modules capable of measuring two spacial coordinates. [27]

A considerable challenge for the silicon tracker's construction was the significant heat dissipation from the detector elements which has to be removed from inside the detector. Simultaneously the amount of material within the tracker had be kept to a minimum in order not to cause significant momentum loss for tracked particles.

4.4.1 The pixel detector

The pixel detector forms the barrel tracker's three inner layers with radii of 4.4 cm, 7.3 cm, and 10.2 cm respectively, and the endcap tracker's two inner layers at a distance from the nominal interaction point of 34.5 cm and 46.5 cm respectively. It consists of a total of 66 million pixels each $100\,\mu m^2 \times 150\,\mu m^2$ in size. Its total size is therefore about $1\,m^2$. The pixel detectors cover a pseudo-rapidity range of $0 < |\eta| < 2.5$.

Each pixel is implemented as a reverse-biased diode which detects a particle by the ionization trail it causes. In the barrel the magnetic field induces a Lorentz angle of 23° which increases charge sharing between neighboring pixels. This improves the ability to obtain sub-pixel resolution and brings the pixel tracker's final resolution to 10 µm in the r and ϕ direction and 20 µm in the z direction. The pixel detectors in the endcaps are installed at an angle of 20° to obtain similar results.

4.4.2 The silicon strip tracker

The silicon strip tracker consists of 10 individual detector layers. The strip size varies between 81 µm×10 cm and 183 µm×25 cm with resolution generally decreasing with distance from the interaction point. It contains 6 million strips on 15148 modules which adds up to an active silicon area of 198 m^2. This makes CMS's silicon strip detector the largest detector of its type.

Some of the modules are implemented as double sided – so-called "stereo" – modules which means two modules glued to each other back-to-back at an angle of 100 mrad. In this way a spatial z-resolution up to 230 µm can be obtained. The resolution in r and ϕ reaches 23 µm to 34 µm in the inside layers of the silicon strip detector.

4.5 Calorimetry

Calorimeters measure the energy of both charged and neutral particles. While electromagnetically interacting particles lose their energy comparatively fast and can therefore be stopped in the first calorimeter layer – called the electromagnetic calorimeter (ECAL) – hadrons have a far smaller cross-section for interaction at the expected energies. Therefore a further layer – the hadronic calorimeter (HCAL) – is built around the ECAL.

As one of the prominent decay channels for the Higgs boson is to two photons, care was taken to construct CMS's ECAL with very high energy and angular resolution.

4.5.1 Electromagnetic calorimeter

As explained before the ECAL's primary benchmark in construction was the $H \rightarrow \gamma\gamma$ channel. This decay channel produces a sharp peak over the background for Higgs masses of up to 140 GeV which means that good energy resolution is a must. Furthermore, good angular precision as well as accurate pion/photon separation are required.

The ECAL is a homogeneous hermetic crystal calorimeter, consisting of 61 200 lead tungstate ($PbWO_4$ crystals in the barrel and 7324 in the endcaps. The barrel covers a pseudo-rapidity region of $|\eta| < 1.479$ which the endcaps

extend to $|\eta| < 3$. Each crystal in the barrel measures $22\,\mathrm{mm}^2 \times 22\,\mathrm{mm}^2$ on the inner face and $26\,\mathrm{mm}^2 \times 26\,\mathrm{mm}^2$ on the outer face. The crystals themselves almost perfectly face the interaction point save for an intentional offset in both ϕ and η which reduces the chance of cracks along probable particle tracks. The endcap ECAL system is constructed similarly with an inner face area of $28.62\,\mathrm{mm}^2 \times 28.62\,\mathrm{mm}^2$ and an outer face area of $30\,\mathrm{mm}^2 \times 30\,\mathrm{mm}^2$. The crystals here are similarly offset. A schematic of CMS's ECAL system is given in figure 4.3.

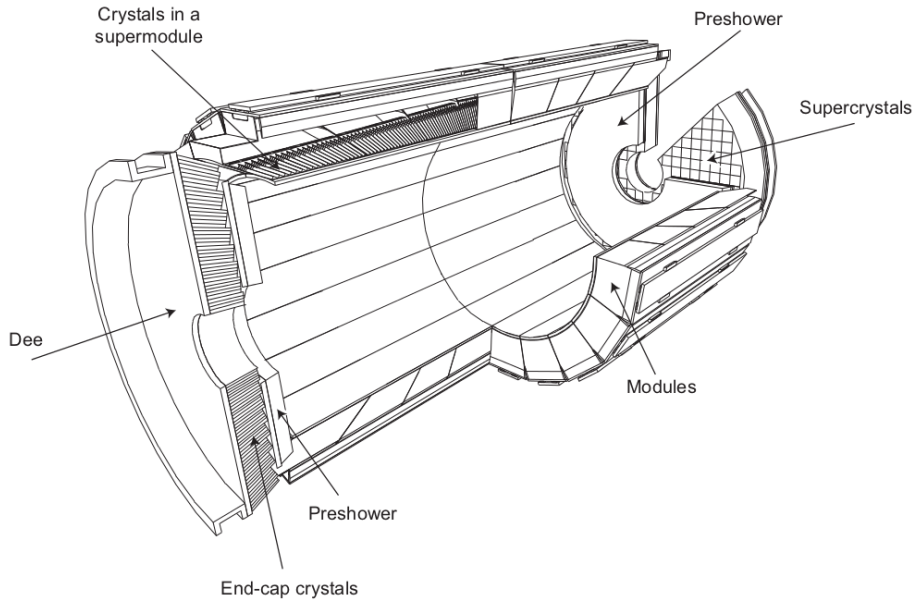

Figure 4.3: The ECAL system in CMS. [28]

In front of the endcaps a preshower detector is placed to identify neutral pions and minimum ionizing particles. It furthermore allows to determine the position of photons and electrons with high granularity.

Lead tungstate was chosen as scintillation material as it offers small light decay time, small Moliere radius, and small radiation length. Its comparatively small light light yield can be compensated by modern photo diodes. A further difficulty is introduced as its light yield is temperature-dependent requiring constant cooling to 18 °C.

4.5.2 Hadron calorimeter

The CMS HCAL enables physicists to measure hadron jets with good precision, thereby making it possible to calculate quantities such as the total missing energy necessary to indirectly detect neutrinos or other weakly interacting exotic particles.

It is built as a sampling calorimeter consisting of several parts. The

barrel part (HB) is located between the ECAL and the solenoid and covers a pseudo-rapidity area of $|\eta| < 1.3$. It consists of 17 layers of plastic scintillators punctuated by layers of brass absorber plates. As the space between the ECAL and the solenoid is not sufficient to completely absorb the most energetic particles, the magnet itself is used as absorption material by the outer barrel part (HO) which consists of one additional layer of scintillators.

The pseudo-rapidity coverage is extended by the endcap calorimeter (HE) to $|\eta| < 3$ using a similar design as for the barrel. For even higher pseudo-rapidities a different design must be used for the hadron calorimeter as the radiation environment becomes very harsh. For this reason the forward hadron calorimeter (HF) uses a steel absorber structure equipped with quartz fibers. This produces Cherenkov light which can then be measured. The HF detector extends coverage to $|\eta| < 5$.

Figure 4.4 shows a longitudinal view of the CMS's HCAL system.

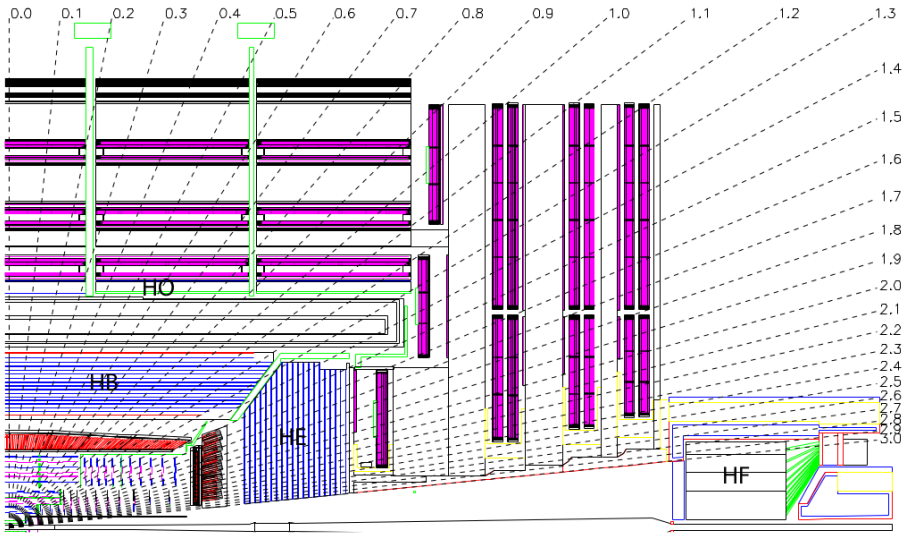

Figure 4.4: The HCAL system for CMS consists of the barrel part (HB), the outer barrel part (HO), the endcap calorimeter (HE), and the forward calorimeter (HF). This provides coverage up to very high pseudo-rapidities. [28]

4.6 Forward detectors

The very forward detectors work in conjunction with TOTEM (section 3.1.3) in order to provide very high coverage in η.

4.6.1 Castor

Centauro and Strange Object Research (Castor) is a sampling calorimeter similar in construction to HF. It consists of a quartz scintillator and tungsten absorber plates. The Cherenkov light produced in the quartz scintillators is then collected by photomultipliers through air-core light guides.

Castor is used mainly for nucleus-nucleus physics in both the heavy ion and p-p programs as well as diffractive physics for proton runs. It is positioned behind HF, covering a pseudo-rapidity coverage of $5.2 < |\eta| < 6.6$ (figure 4.5).

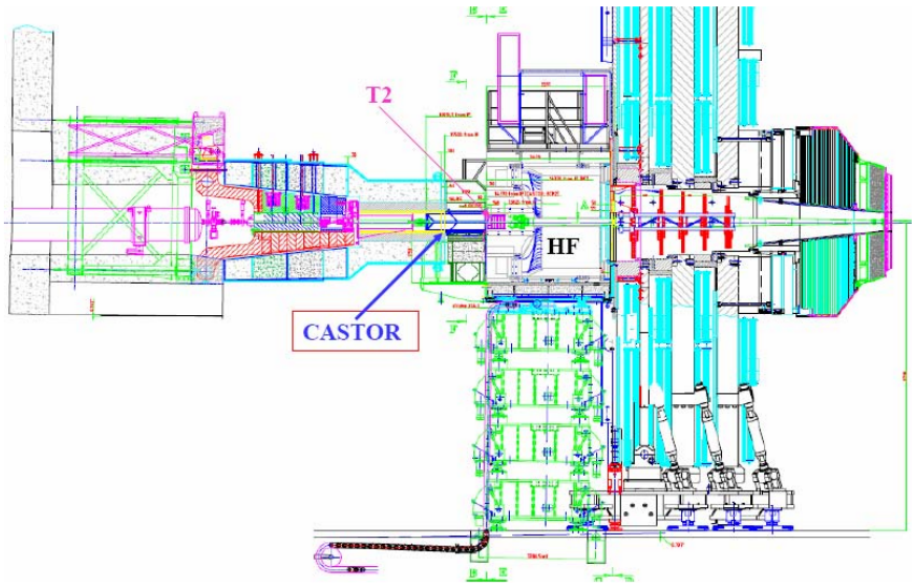

Figure 4.5: The very forward detector Castor is placed right behind HF, extending CMS's calorimetry pseudo-rapidity coverage significantly. [28]

4.6.2 Zero-degree calorimeters

The zero-degree calorimeters are positioned at both sides of CMS at a distance of 140 m from the nominal interaction point, covering a pseudo-rapidity region of $|\eta| < 8.3$.

They are placed behind the LHC beam separator which means that only neutral particles reach it. The ZDC is used mainly to complement CMS's very forward region for heavy ion as well as pp diffractive studies.

4.7 The superconducting solenoid

The solenoid used in CMS provides a magnet field, parallel to the beam-line with an induction of 4 T at the interaction point. Its notable features

are its size of 6 m in diameter and 12.5 m in length, leading to a weight of 220 t, furthermore its 4-layer winding made from a stabilized, reinforced NbTi conductor which provides a very high energy density of $11.6\,\mathrm{kJ\,kg^{-1}}$ leading to a deformation of 0.15 % during energizing.

4.8 The muon system

As explained earlier the muon system is of central importance to CMS's physics goals. This follows from the relative ease of muon detection as well as the possibility to measure a decayed particle's mass with high degree if its final state only consists of muons. The high mass resolution follows from muons only minimally interacting with the material in the tracking and calorimetry systems. Thus the $H \to ZZ \to 4\ell$ decay channel is considered gold-plated when all the leptons in it are muons.

The main requirements for the muon system are consequently muon identification, accurate measurement of their momenta, high confidence in charge assignment, as well as high radiation tolerance and a sophisticated triggering system.

The muon system consists of DTs complemented by RPCs in the barrel and CSCs together with RPCs in the endcaps. Drift tubes were chosen in the barrel due to their comparatively low cost and high position resolution. They suffer, however, from low timing resolution due to comparatively high drift times of about 400 ns. Resistive plate chambers conversely offer low position granularity while touting a time resolution of ~ 1 ns. The endcap muon system suffers from a higher magnetic field, as well as high muon rate and background which means that drift tubes are not viable anymore. For this reason cathode strip chambers are used in the endcaps.

4.8.1 Drift Tubes

Drift tubes in CMS are contained in muon stations placed in four layers in the barrel. Each individual drift tube is 2.5 m long, its face measuring 42 mm × 13 mm. They are filled with an Ar/CO_2 mixture at atmospheric pressure and stacked in layers of four which comprise a so-called superlayer (SL). Three SLs then form a muon station.

The drift tubes themselves are oriented either in parallel to the beam direction in order to measure the ϕ direction or orthogonal to it for better granularity in the z-direction. The single-wire resolution for a drift tube is 250 µm which results in a resolution of 100 µm for an entire station. The drift tube system's coverage is $|\eta| < 1.2$.

4.8.2 Cathode Strip Chambers

The cathode strip chamber system in CMS covers a pseudo-rapidity region of $0.9 < |\eta| < 2.4$. They are used as the high background and muon flux in the endcaps would overwhelm a drift tube system. Furthermore, the inhomogeneous magnetic field there would make drift calculations very complicated.

The cathode strips are used for precision measurements in the r-ϕ plane while the anode wires lying roughly perpendicular can measure η and offer beam crossing identification for individual muons.

4.8.3 Resistive Plate Chambers

The RPC detectors are dedicated exclusively to the trigger system as their timing resolution of 1 ns allows unambiguous bunch crossing identification. Apart from timing information the trigger also receives a coarse measurement of the particle's transverse momentum and spatial coordinates making it possible to match it with measurements taken by the DT and CSCs subsystem, thereby improving the trigger decision.

The RPC system is used both in the endcaps and barrel covering a pseudo-rapidity of $|\eta| < 1.6$

4.9 Trigger and data acquisition

Given the LHC's nominal bunch crossing frequency of 40 MHz at a luminosity of $1 \times 10^{34}\,\mathrm{cm}^{-2}\,\mathrm{s}^{-1}$ which would translate to a data rate of 40 TB/s if every event were recorded, it is paramount to be able to promptly and efficiently select interesting events while discarding the rest.

CMS uses a two-stage trigger system to accomplish this task. The Level-1 trigger (L1T), built from custom hardware provides a first selection within 3.2 µs, thereby reducing the event rate to 100 kHz. During this time the data under consideration is stored in a ring buffer on local detector front-end systems. The Level-1 trigger reaches a read-out decision on the basis of coarse calorimeter and muon system information. A more detailed explanation of the L1 trigger is given in chapter 5.

The second stage is provided by the High-Level trigger (HLT) Event Filter which is implemented as a computing farm, analyzing data taken online using a C++ framework. The HLT operates on the complete event data which is stored on over 600 front-end driver modules on the detector. For this reason a large builder network is necessary to construct an event from this multitude of sources. The HLT then outputs events to be recorded on tape at a rate of about 300 Hz. A visualization of this pipeline is provided in figure 4.6.

In addition to the task of filtering the Level-1 output stream, the HLT also provides on-line data-quality monitoring (DQM) allowing operators to

identify and diagnose problems during runs without having to wait for later offline analysis.

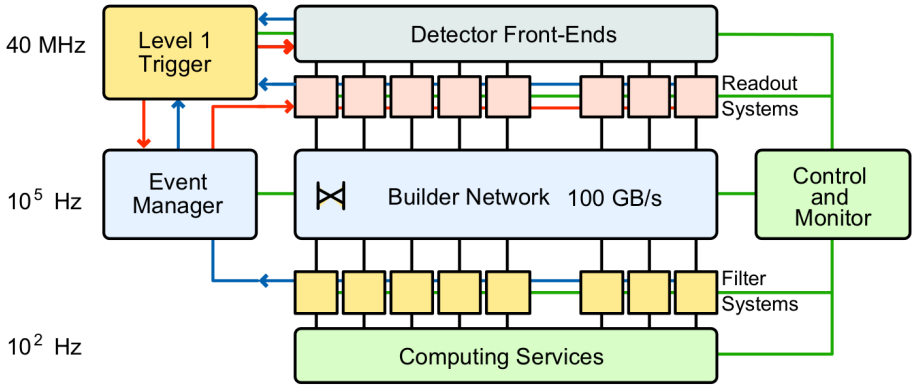

Figure 4.6: The CMS read-out chain. The Level-1 trigger reduces the incoming event rate of 40 MHz provided by LHC to about 100 kHz which can be managed by the software filter systems (HLT). In this stage the event rate can be reduced even further to about 300 Hz which are then stored on tape. [28]

Chapter 5

The CMS Level-1 trigger

As the LHC supplies a very high collision rate and luminosity to the experiments it is impossible to store every event. Storing the collision data for one event requires at least 1 MB of storage space. At the nominal collision rate of 40 MHz this would add up to at least 40 TB/s consisting largely of uninteresting data.

Conversely, the cross-sections for interesting events vary over orders of magnitude. While the production of $b\bar{b}$ has a cross-section on the order of mb, the cross-sections for the production of the Higgs boson in the relevant mass region is on the order of fb (figure 5.1).

These conditions require a very flexible trigger system that can be programmed to accept different relative quantities of different events. For this reason the L1 trigger was designed to work with so-called physics objects, listed in the following:

- Muons
- Electrons or photons
- Jets
  clusters of energy in the calorimeter
- τ jets
  narrow jets corresponding to the decay of a τ lepton
- Energy sums
  - E_T^{miss}
    missing transverse energy
  - E_T^{total}
    total transverse energy
  - H_T^{miss}
    missing hadronic jet energy

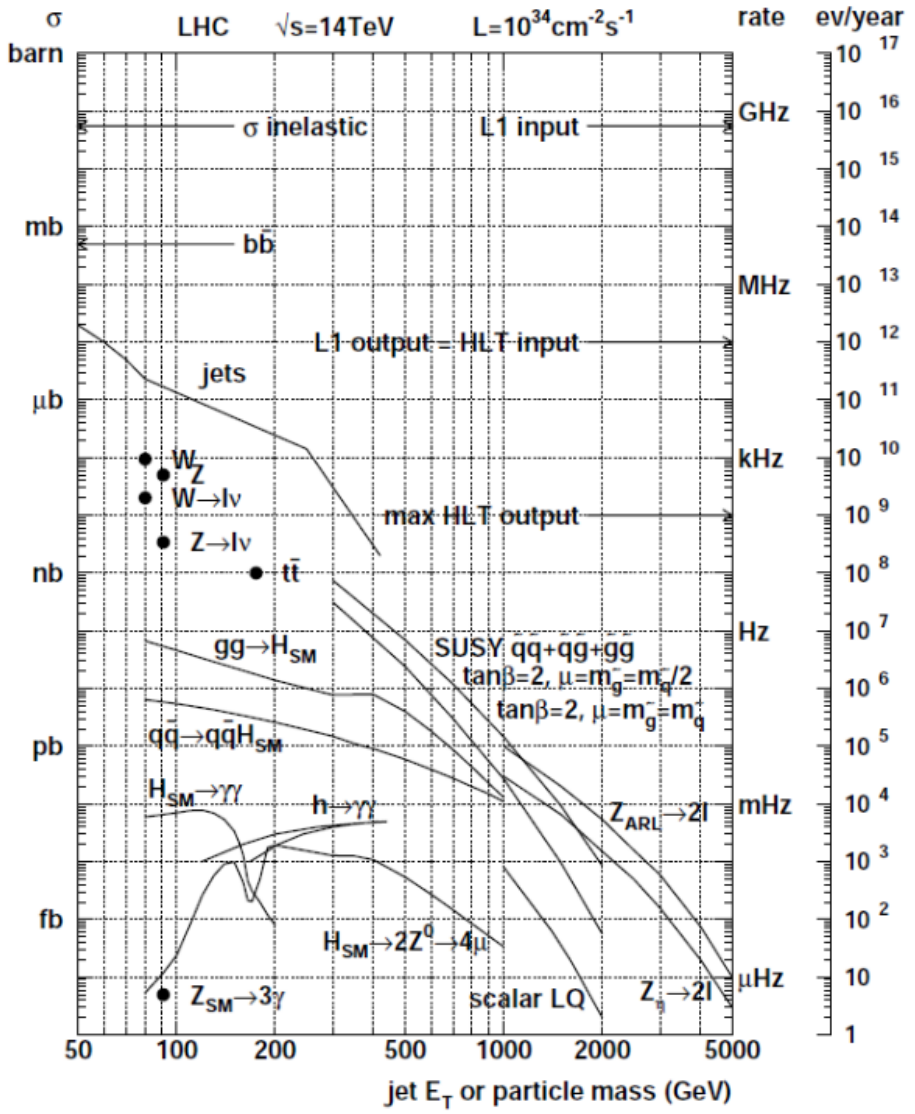

Figure 5.1: The cross-sections for phenomena to be studied with CMS vary widely over orders of magnitude. On the right hand side the interaction rates at nominal luminosity are given. [29]

– H_T^{total}
total hadronic jet energy

Furthermore, the L1 trigger is equipped to provide prescaling on an algorithm basis, where a prescale factor n indicates that only every n^{th} event is selected.

These features make it possible to formulate trigger algorithms in a natural way – e.g. to record every event containing 4 muons, or to record only 1 in every 100 events containing 2 muons but no jet.

As the Level-1 trigger is the first stage in the read-out chain the data is only read out of the detector when the read-out decision is made, until which the event information has to be stored in buffers inside the detector. The maximal latency for the read-out decision to be made is determined by the tracker and preshower detector buffers which contain space for 128 bunch crossings. This translates to a time of 3.2 µs. The actual processing time is reduced even further when considering delays introduced due to the drift time in the drift tube chambers or the signal propagation time due to the distance between the detector and the counting room, where a large part of the trigger electronics is located. This signal propagation time sums up to about 1 µs.

5.1 The structure of the L1 trigger

Due to the requirement of speed for the Level-1 trigger it works in a "divide and conquer" fashion: Relatively simple algorithms are run in parallel on custom hardware (field-programmable gate arrays (FPGAs) and application-specific integrated circuits (ASICs)).

Furthermore, the calorimeters as well as the muon system are handled separately until the highest level within the L1 trigger. The final accept or reject decision is then made in the Global Trigger, taking into consideration the results of the calorimeter and muon trigger sub-systems.

Calorimeter and muon trigger systems are built in a similar fashion. Initially so-called trigger primitives are found – i.e. track segments or energy clusters – which are then merged in regional triggers to produce trigger objects. Finally these objects are merged and ranked in the Global Calorimeter and Global Muon Triggers, which forward the best candidates to the Global Trigger (figure 5.2).

5.1.1 The calorimeter trigger

The calorimeter trigger is designed to identify the four most energetic physics objects from each of the following categories.

- isolated electron or photon candidates in $|\eta| < 2.5$

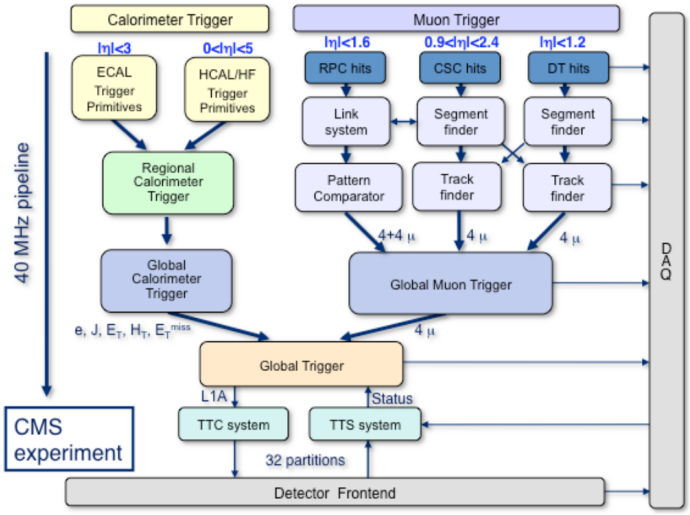

Figure 5.2: The Level-1 trigger in CMS is built in a highly modular way. In both the muon and calorimeter triggers initially trigger primitives are built. These are forwarded to regional triggers, which construct the final trigger objects and assign values for transverse momentum or energy and quality to them. Subsequently the Global Muon and Global Calorimeter Triggers send the best of these objects on to the Global Trigger which decides whether to initiate a detector read-out. [30]

- non-isolated electron or photon candidates in $|\eta| < 2.5$
- jets in the central region, $|\eta| < 3$
- jets in the forward region, $3 < |\eta| < 5$
- τ-like jets in $|\eta| < 2.5$

Furthermore it should measure these objects' transverse energy and position, identify the bunch crossing the event originated from, and calculate the total missing energy as well as the total transverse energy in the calorimeters.

Theses tasks must be performed at an input frequency of 40 MHz in a pipelined-mode without any dead-time.

The calorimeter trigger is logically split into several steps as can be seen in figure 5.4. The Trigger Primitives Generator (TPG) sub-system which is integrated with the calorimeter read-out, is the first step in the pipelined calorimeter trigger system. Its main tasks are to assign a bunch crossing to each detector pulse and calculate the transverse energy sum for each so-called trigger tower.

Trigger towers are defined (η, ϕ)-regions used for electron/photon identification. They are comprised of 5×5 crystals in the ECAL. The size of these

towers was chosen in a trade-off between background-suppression which is improved with smaller tower sizes and cost issues which increase with a larger number of trigger channels.

A group of 4×4 trigger towers is then combined in a calorimeter region which forms the basis for jet and energy triggers which are calculated in the Regional Calorimeter Trigger (RCT). In the HF region a calorimeter region consists of only one trigger tower (see figure 5.3).

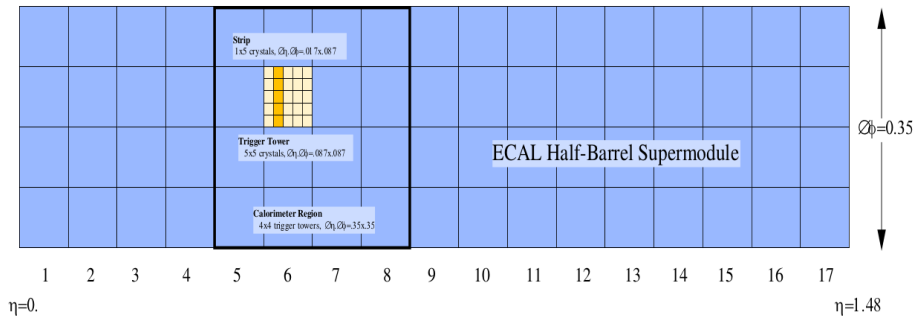

Figure 5.3: Layout of trigger towers and calorimeter regions in the barrel of CMS. [29]

The RCT receives the energy deposits for each of the trigger towers from the TPGs. It then builds τ leptons, jets, and isolated as well as non-isolated electron/photon candidates from them. It consists of 20 processing crates of which 18 are used for the calorimeters in the barrel and endcap, covering a region of $|\eta| < 3$. One of the remaining crates covers both HF calorimeters which extends coverage for missing transverse energy and jet detection to $|\eta| < 5$. The final crate combines information from the remaining 19 crates to find jets and τ leptons. Furthermore it sums the transverse energies calculated in the other crates to provide energy sums for various ϕ-regions.

Each of the 18 crates sends its four highest ranked isolated as well as non-isolated electron/photon candidates to the Global Calorimeter Trigger (GCT). The crate finding jets and τ leptons transmits the information for its highest energy 9×4 τ leptons and central as well as forward jets as well as the transverse energy sums for 18 ϕ-regions to the GCT.

In the GCT the trigger objects received from the RCT are sorted. Following this the best four electron/photon, jet, and τ lepton objects are forwarded to the Global Trigger. Furthermore the number of jets, the global transverse energy as well as the total missing energy vector are calculated and sent to the Global Trigger as well.

The GCT is also involved in luminosity monitoring by providing transverse energy sums to the Luminosity Monitor. [29, 28]

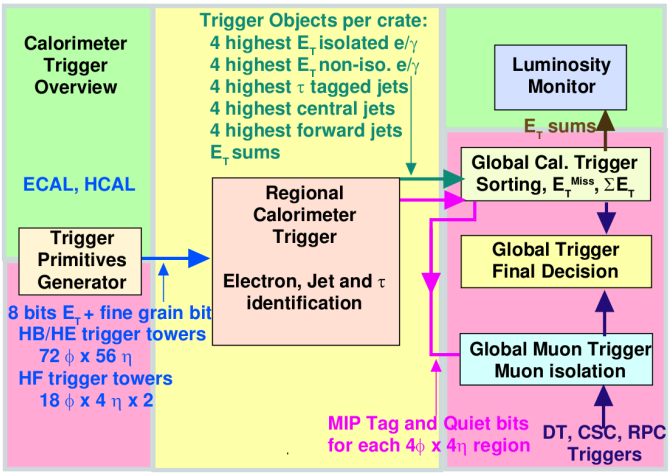

Figure 5.4: Flow diagram for the operation of the calorimeter trigger. So-called trigger primitives are constructed by the Trigger Primitives Generator which are then forwarded to the Regional Calorimeter Trigger. Here the actual trigger objects are built and then forwarded to the Global Calorimeter Trigger which sorts these, calculates several global objects and then forwards them to the Global Trigger. The transmission of MIP and Quiet bits as indicated in this figure has not been fully implemented in the current system as explained in section 5.1.2. [29]

5.1.2 The muon trigger

As explained in chapter 4 a major goal in the development of the CMS experiment was the precise measurement of the properties of muons. The muon system consists of a DT sub-system in the barrel and a CSC sub-system in the endcaps as can be seen in figure 5.5. Additionally an RPC sub-system is provided exclusively for triggering purposes. This allows tracks to be unambiguously matched to specific beam crossings.

The Level-1 muon trigger is compartmentalized into several sub-systems:

- local DT and CSC triggers
- regional DT and CSC triggers
- the Pattern Comparator Trigger (PACT)
- the Global Muon Trigger (GMT)

Their data flow is visualized in figure 5.6.

The local triggers for the DT and CSC sub-systems build track segments which are transferred to the regional triggers. Track segments from the overlap regions shown in figure 5.5 are shared between the regional triggers

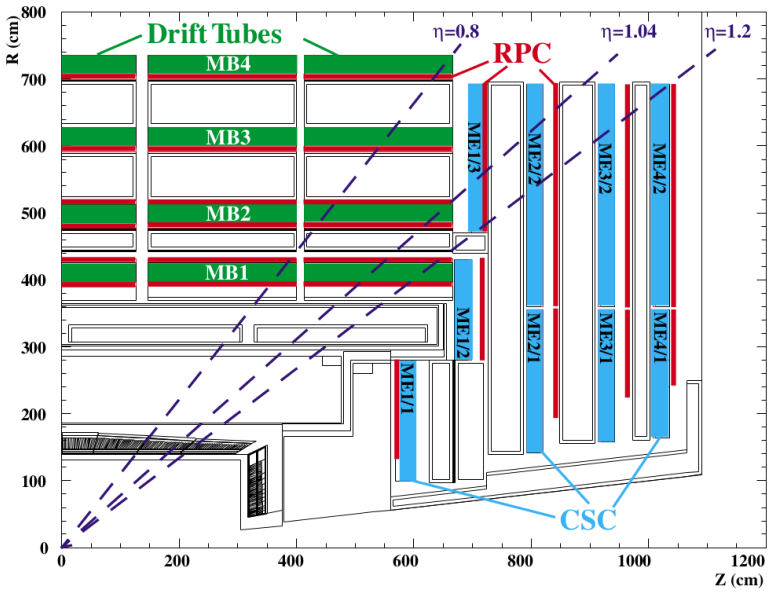

Figure 5.5: The detectors used in the muon trigger in a longitudinal cut. As can be seen there is an overlap region between the DT and CSC systems, both of which are complemented by RPC stations. [31]

in order to be able to suppress ghosts. Additionally coarse data from the RPC detectors in the endcaps can be transferred to the CSC local trigger to help reduce spatial and temporal ambiguities in multi-muon events.

Each regional trigger for the CSC and DT sub-system transmits information consisting of transverse momentum, spacial coordinates in the detector and a value indicating the quality of measurement to the GMT. The PACT which uses the RPC sub-system finds four muon candidates each for both the barrel and endcaps and similarly sends these candidates along with the same information as provided by the regional triggers to the GMT. The GMT then merges and ranks the muon candidates from the complimentary systems using a measure of quality which is indicative of the confidence to have found a true muon. Finally the four muons with the highest quality are propagated to the Global Trigger (GT).

Local triggers

Local triggering in endcap and barrel regions is very similar. Both systems initially generate trigger primitives within small detector regions whose tracks are then extrapolated and matched to form global track candidates to be forwarded to the track finders.

As can be seen in figure 5.5 a substantial overlap region between the DT and CSC coverage for muons originating from the interaction point

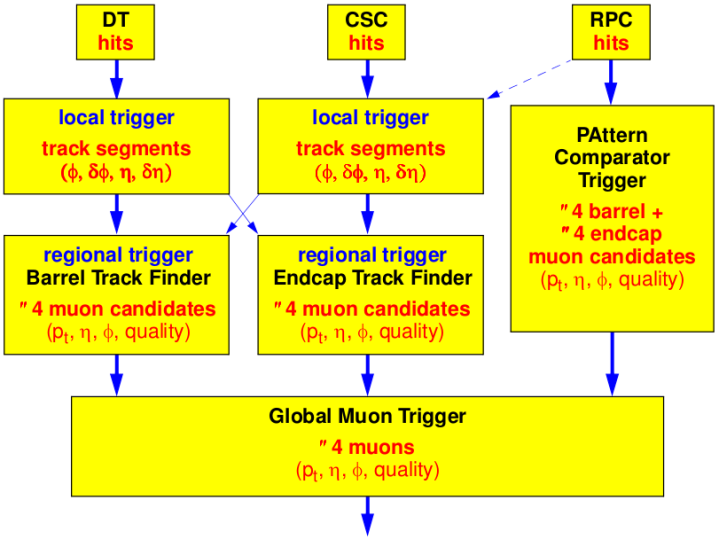

Figure 5.6: Flow diagram for the muon trigger. The three independent detector sub-systems (DT, CSC, and RPC) each are equipped with their own local and regional triggers. The results found in the regional triggers for each sub-system are then forwarded to the GMT which merges and sorts candidates in order to send the four best muon candidates to the GT. [29]

exists. In order to improve the measurements made by the muon system and to mitigate the risk of double-counting muons in both systems a hard boundary at $|\eta| = 1.04$ is defined between the local DT and CSC triggers. This means that track segments found in MB2 muon stations are sent to the CSC Track Finder while track segments found in ME1/3 are sent to the DT Track Finder.

The DT local trigger is comprised of

- Bunch and Track Identifiers (BTIs)
- Track Correlators (TRACOs)
- Trigger Servers (TSs)

Each BTI is connected to nine drift tubes in four layers within one SL. It is designed to search for coincident, aligned hits and extrapolate the so-found track segment. The BTI requires a minimum of three layers with hits, but track segments formed from four hit layers are assigned a higher quality. As each SL only measures either ϕ or η of a measured track segment the two coordinates are separate at this point.

TRACOs correlate the tracks measured in ϕ-type SLs. In this way the angular resolution for the tracks can be increased. The TS for one muon station then receives the information from all the TRACOs in that station and selects the two tracks with the lowest bending angle to forward to the Drift Tube Track Finder as this correlates with high transverse momentum. The BTIs within the SL measuring the muon's pseudo-rapidity directly transmit their measured values to the TS.

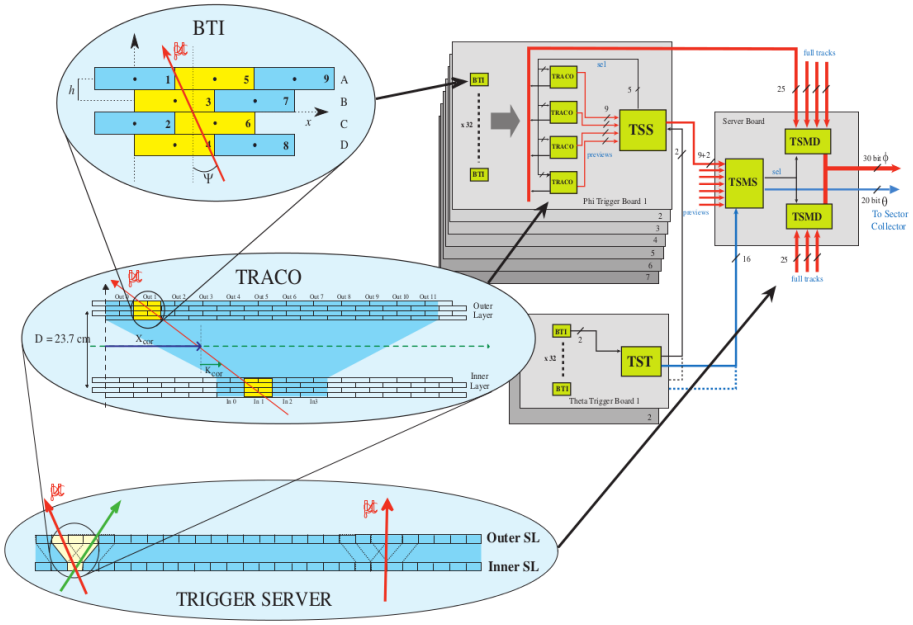

Figure 5.7: Visualization of track segment finding in the DT local trigger. [29]

The CSC local trigger works by finding Local Charged Tracks (LCT) which are comprised of the position, angle, and bunch crossing information for the candidate track. Similarly to the DT local trigger the two coordinates, η and ϕ are measured separately by the anode wires and cathode strips respectively. The anode wire electronics was specifically optimized for bunch crossing identification while the cathode strip electronics was designed to optimally measure ϕ.

The endcap region is especially challenging due to high occupancy and the comparatively high momentum of particles in this region of the detector. To suppress this large background a coincidence of four CSC layers is required. This, paired with the fact that cathode strip signals are digitized with half-strip width allows a resolution of 0.15 strip widths to be reached.

If a track is found using the cathode strips it is then forwarded to the trigger motherboard (TMB) as a Cathode LCT (CLCT). Similarly Anode

Local Charged Tracks are built and sent to the TMB. Here these tracks are correlated and merged to LCTs that are sent to the Muon Port Card (MPC).

The MPC receives LCTs from 9 TMBs and propagates the best two or three track candidates to the CSC Track Finders.

Regional triggers

Both the DT Track Finders (DTTFs) and CSC Track Finders (CSCTFs) are designed to identify muon candidates in the barrel or endcap and determine their transverse momentum p_T, their spatial coordinates η and ϕ as well as assign them a measure for the quality of their measurement.

The candidates are then sorted by their rank which depends on the candidate's transverse momentum and quality in order to send the best four muon candidates each in endcap and barrel to the GMT.

DT Track Finder The DTTF receives track segments from each muon station where they are connected into a full track and assigned a value for their particle's transverse momentum.

To efficiently manage this job the detector is split into wedges of 30° in the ϕ angle which are then again split into 12 sectors in the z direction each (figure 5.8).

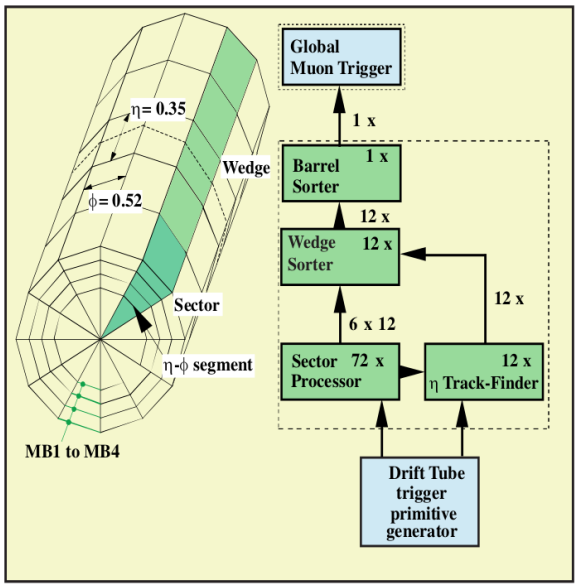

Figure 5.8: The barrel is logically split into wedges which are then again split into 12 sectors each. A sector correlates with the area of one detector wheel, except for the central wheel which is split into two sectors to avoid asymmetries. [29]

Each sector is equipped with a Sector Processor (SP). These are logically split into an Extrapolation Unit (EU), a Track Assembler (TA), and an Assignment Unit (AU). SPs for the outermost wheels receive information from the CSC sub-system as explained before.

The EU matches track segment pairs of different muon stations by extrapolating a hit coordinate for each track segment using its spacial coordinate ϕ and bending angle. It can also use track segments from neighboring sectors as targets for its extrapolation to account for tracks crossing sector boundaries.

The TA then finds the two tracks with the highest rank in a given sector. Here the rank is determined by the number of matching track segments and the extrapolation quality from the EU.

Finally the AU determines the candidate muon's transverse momentum, its spacial coordinates η and ϕ, as well as a value for quality of the track measurement. (See figure 5.9)

The SP then forwards its two highest ranking muon candidates to the Wedge Sorter which selects the two candidates with the highest transverse momentum in the corresponding wedge and propagates them to the Muon Sorter. Here potential duplicate tracks due to sharing of information between sectors by the SPs are merged and all candidates from the complete detector are again sorted, the four tracks with the highest rank being sent to the GMT.

CSC Track Finder Similar to the DTTF the first stage in the regional CSC trigger is handled by SPs that receive LCTs from the MPCs. A significant difference between the DT and CSC regional triggers is introduced due to the high background in the endcaps, produced by low transverse momentum muons and other particles. Furthermore the magnetic field is not axial in the endcaps leading to complicated particle trajectories.

The CSCTF system is split into a total of 12 sectors, each covering 60° in ϕ. Every SP then receives up to 16 muon candidates from the local triggers. The SPs operate very similar to the SPs in the DTTF with the exception that sharing of information between neighboring sectors is not necessary due to the small bending angle in the endcaps.

The SPs then forward the best three muon candidates together with their transverse momentum, spacial coordinates, and quality to the forward Muon Sorter Board. From here the best four muons are propagated to the GMT.

RPC trigger The RPC trigger differs from the DT and CSC triggers in that no local trigger is used. As the RPC sub-system is used solely for triggering its entire readout is directly transferred to the counting room. The RPC system can unambiguously assign tracks to bunch crossings due

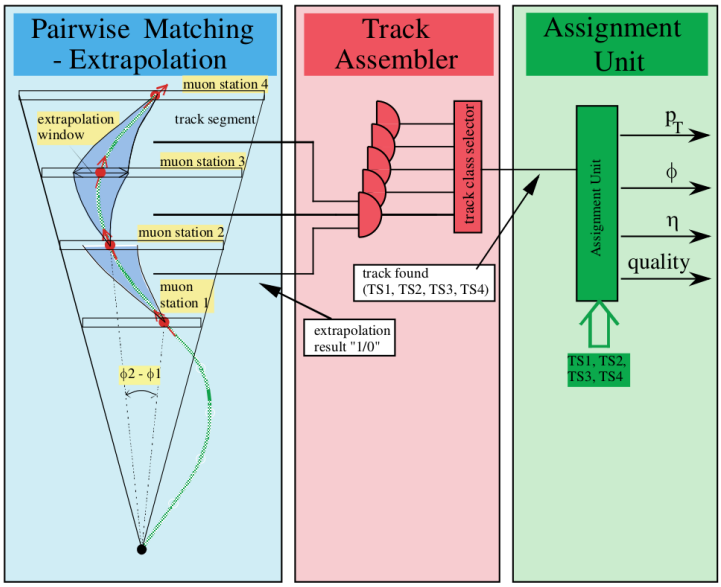

Figure 5.9: The Sector Processor (SP) first matches track segments in the Extrapolation Unit (EU), finds and determines the best two tracks in a given sector using the Track Assembler (TA) and finally calculates the tracks' transverse momentum, spatial coordinates and quality in the Assignment Unit (AU). [29]

to its superior time resolution of 1 ns which makes it ideal for triggering purposes.

In the counting room the information from the RPC sub-system is segmented into 33×144 trigger towers in the η and ϕ directions respectively. A separate PACT is assigned to each trigger tower. It then finds track candidates with constant η values. To do this it requires a minimum of 3 hit layers from 4. Finally the four highest ranked muons from every η-ring are sent to two area sorters. These area sorters find the best four muons for both barrel and endcap regions which are propagated to the GMT.

The Global Muon Trigger

The GMT forms the final stage of the muon trigger. It receives a total of 16 muon candidates from the regional triggers – four each from the DT, and RPC systems in the barrel as well as from the CSC, and RPC systems in the endcaps.

It greatly improves the trigger efficiency by attempting to merge muon candidates that were detected in two complementary sub-systems (i.e. DT+RPC or CSC+RPC). Candidate muons that were confirmed in this way are assigned the highest quality by the GMT.

Furthermore the GMT is equipped with cancel-out units which remove

duplicate candidates originating in the overlap regions from both DT and CSC sub-systems.

The GMT is also equipped with logic to use information indicating quiet regions (which would allow improved background suppression) and signatures of minimum-ionizing particles (MIP, indicative of muon tracks) in the calorimeter that was transmitted by the RCT. This logic is not used currently due to the lack of sufficient studies for this feature.

Finally the GMT ranks the muon candidates according to a quality value determined by the quality propagated from the regional trigger. The highest quality value is reserved for confirmed candidates. The four highest ranking candidates are then propagated to the GT.

5.1.3 The Global Trigger

The GT forms the final part of CMS's Level-1 trigger. Its main task is to decide every 25 ns whether to read out an event from the detector. This decision is based on so-called algorithms which are a collection of different physics conditions, such as momentum thresholds, energy thresholds, muon quality, and spatial conditions for detected objects. Especially the possibility for spatial conditions is a big improvement over prior first level trigger systems. Furthermore the GT uses 64 so-called "external conditions" provided by other systems and 64 "technical triggers" to come to a read-out decision. The results of these arguments are combined in a final OR (FINOR) which determines whether to read-out the detector.

The maximal allowed read-out rate is determined by the HLT that can only cope with an input rate of 100 kHz. To enforce this limit while continuing to be able to trigger on very common events, algorithms can be prescaled to only trigger on every n^{th} occurrence of a certain event. These prescale factors can also be applied to the technical triggers.

Due to the latencies introduced by the preceding trigger systems the GT is allowed approximately 0.25 µs to come to a decision. To keep cable lengths to a minimum the complete GT is housed in a single crate along with the GMT (figure 5.10).

The GT receives the best four muons from the GMT. They are sorted by their rank, which depends on their quality and transverse momentum. The GT also receives the muon's transverse momentum, spatial coordinates η and ϕ, electric charge, and quality from the GMT.

Furthermore the GT receives

- 4 τ-jets
- 4 central and 4 forward jets
- 4 isolated and 4 non-isolated electron/photon objects

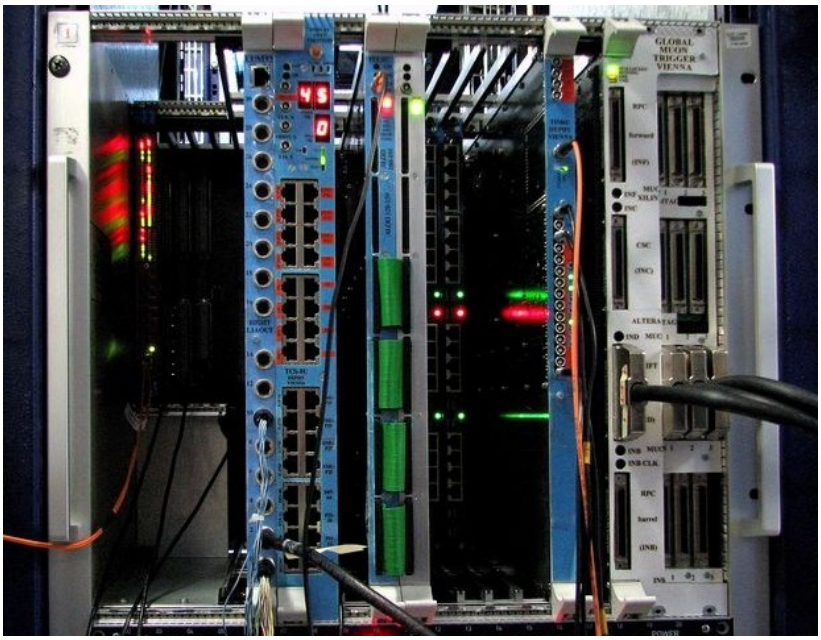

Figure 5.10: Picture of the GT crate. From left to right the cards are: TCS, FDL, GTL, PSB, TIM, GTFE, and GMT.

- the total transverse energy E_T^{total}, the magnitude and direction of the missing transverse energy E_T^{miss}, the total hadronic jet energy H_T^{total}, and the magnitude and direction of the missing hadronic jet energy H_T^{miss}

from the GCT. This information has to be delayed in the GT as the processing in the muon trigger takes longer than in the calorimeter trigger.

Global Trigger hardware

The current GT is housed in one VME crate together with the GMT. All modules inside the crate are connected via the back-plane which is also used for communication with the online control and monitoring software. This software accomplishes its duties by accessing exposed registers and memories on the FPGA chips carried by the VME modules. The use of FPGA technology additionally allows the trigger to be extremely flexible as new firmware can be loaded into the FPGA before a run.

Timing Module The Timing Module (TIM) forwards broadcast information received from the Trigger Timing and Control (TTC) system which provides among others the clock signal from the LHC. This is used to synchronize different parts of the detector electronics to each other and to the LHC clock.

Pipeline Synchronized Buffer Boards The Pipeline Synchronized Buffer (PSB) boards receive data from external systems. Their main functionality is implemented in the SYNC chip. Here the phases of all incoming signals are detected through oversampling. Then each channel is synchronized to the LHC clock by applying a programmable delay to it.

PSB modules are used to synchronize signals from the calorimeter and technical triggers. They are equipped with four InfiniBand connectors, each providing two 1.28 Gbit/s data streams which translates into 8 logical input channels. Optionally the first two of these channels can be also fed by 18 RJ45 ports used for LVDS parallel data at 80 MHz.

Global Trigger Logic board The Global Trigger Logic (GTL) module is the main board of the GT. It uses five FPGA chips of which three are receiver chips which are provided with the 80 MHz input data from the GMT and PSB modules via the back-plane. From these chips the data is then distributed to the two condition chips.

These chips actually apply the conditions and algorithms to the input data. Each chip can process up to 96 algorithms in parallel, but currently only 128 algorithms are used in total.

The decision bits for these 128 algorithms are then sent to the Final Decision Logic (FDL) board via flat-band cables.

Final Decision Logic Board The FDL module receives 128 algorithm decision bits from the GTL as well as 64 technical trigger bits from a dedicated PSB board.

It contains rate counters, prescale factors, and – for technical triggers – Veto Mask Logic for every incoming bit. If a veto bit is set for a specific technical trigger an event is suppressed if the technical trigger bit itself is also set.

The data acquisition (DAQ) system for CMS is divided into eight partitions which can be read out separately if need be. The FDL can produce a separate FINOR bit for each of these partitions forming an 8-bit FINOR word.

Global Trigger Front-End board The GT is linked to CMS's DAQ system via the Global Trigger Front-End (GTFE) board. If a Level-1 Accept (L1A) signal is produced by the Trigger Control System (TCS) board all relevant GT boards (FDL, TCS, PSB, and GMT) forward their trigger records to the GTFE board which merges them to a single record and sends it to the DAQ system.

Furthermore it merges data from FDL and TCS modules in another event record which is transmitted to the DAQ Event Manager (EVM).

5.1.4 Central Trigger Control System

The TCS produces the L1A signal upon receipt of a positive FINOR signal. However it also takes the current state of the detector into account.

This is done on the one hand by local state machines emulating the state of the front-end buffers, which implement several hard rules (at most 1 L1A signal in 3 bunch crossings, 2 in 25, 3 in 100, and 4 in 240). On the other hand the TCS is notified by a "back pressure" system when buffer space runs out in the HLT, prompting the TCS to suppress triggers until the situation has improved.

5.2 The Trigger Supervisor

The L1 trigger electronics are exposed to CMS's central run-control (RCMS) as a single sub-system.

However, the different trigger stages are housed in different crates and built by different teams, which makes a distributed control and monitoring system necessary.

Such as system is provided by the Trigger Supervisor C++ framework, which uses the CMS-wide XDAQ framework as middle-ware. [32]

Internally a Trigger Supervisor-based system is segmented into a "cell" for each sub-system (e.g. the GT cell) which offers both a web-interface for direct control by a human operator and a SOAP-based API. Cells can be highly customized by the sub-system's maintainer.

For central access to the L1 trigger electronics by RCMS the so-called Central Cell is provided.

Configuration data for the complete CMS detector is stored in the Online Master Database system (OMDS), an Oracle database. It can be easily accessed by Trigger Supervisor cells using XDAQ's TStore service. The L1 trigger stores e.g. its different configurations in this database.

5.2.1 Level-1 Trigger Menu

Among the objects stored in OMDS are the so-called L1 Trigger Menus. A Trigger Menu is the collection of all algorithms that run on the GTL's condition chips at any given time. On the chip-level a trigger menu presents itself as a firmware image that was compiled from VHDL code.

To edit a Trigger Menu a separate Trigger Supervisor Cell is used, the Level-1 Trigger Menu Editor.

5.3 Offline Software

For offline analysis of trigger performance or detector studies a L1 trigger emulator is needed. It must be used e.g. to validate design decisions made

for the trigger before implementing them.

The C++ software framework CMSSW provides bit-level emulators of all L1 trigger components which allows a user to e.g. run already taken data samples only through the stages of the trigger relevant for a specific study.

Using these emulators it is also possible to search for errors in data taking by comparing the emulated trigger's results with the results obtained during data taking.

Chapter 6

The Global Trigger upgrade

6.1 Motivation for the GT upgrade

The current Level-1 trigger electronics will reach their performance limit when the LHC passes its design luminosity of $1 \times 10^{34}\,\mathrm{cm}^{-2}\,\mathrm{s}^{-1}$. To continue triggering efficiently will require higher granularity, more sophisticated triggering algorithms and eventually even the inclusion of more sub-detector systems such as the silicon strip tracker. New physics – found by the current system – will call for more flexible trigger menus that will use more trigger algorithms at the same time.

To fulfill these requirements technologies such as the currently used galvanic cables and now-outdated FPGAs and ASICs will need to be replaced by optical links and more powerful chips. The upgraded system will also have to support many more input channels to be used in more complex algorithms. These algorithms then will be combined in a single FPGA, further increasing flexibility compared to the current system that utilizes two FPGAs for this combination step, as explained in chapter 5.

While the current system utilizes the VME-standard it was decided to base its successor on the newer MicroTCA which offers superior bandwidth on the back-plane, thus opening the possibility of migrating signal ways away from the currently cluttered front-plane. Furthermore it is hoped to be able to equip a future MicroTCA-based Level-1 trigger with mostly commercial of the shelf (COTS) components using one module-type across the Global Trigger project. This would mean that new physics requirements could be fulfilled by simply adding another such commercially available module. Similarly maintenance will be facilitated through such a system. [33]

6.2 MicroTCA as a replacement candidate

Micro Telecommunications Computing Architecture (TCA) (MicroTCA) is a technology born out of the AdvancedTCA (ATCA) platform.

While ATCA is meant to offer "carrier-grade" – i.e. extremely reliable ("five nines" or more) and fast-to-recover ($< 50\,\text{ms}$) – service it is prohibitive in its cost and complexity for the Level-1 trigger project.

A viable alternative is available with the MicroTCA platform. It offers many of the same benefits such as redundancy, system-level management, and hot-swappable field-replaceable units (FRUs) at a lower level of cost and complexity when compared to ATCA.

This makes MicroTCA a viable successor to the current VME-based system. However its versatility as well as its roots as a telecommunications technology introduce complex challenges.

6.2.1 Composition of a MicroTCA system

A MicroTCA system consists of at least one so-called MicroTCA carrier which houses the Advanced Mezzanine Cardss (AMCs), a MicroTCA Carrier Hub (MCH), as well as power systems. All of these systems, except for the AMCs, can be used in a fully redundant way.

Management functions on a carrier level are fulfilled by the Carrier Manager which runs on the primary MCH's MCH Carrier Management Controller (MCMC). All other modules communicate with the MCMC via their own management controller which varies between classes of modules. AMC modules use simple Module Management Controllers (MMC) while power modules and cooling units require Extended MMC (EMMC). [34]

Up to 16 MicroTCA carriers can be held by one MicroTCA shelf which provides cooling to the contained carriers. The carriers and the shelf's cooling system are then managed by a shelf manager which can be implemented on any FRU in one of the carriers (i.e. AMCs, MCHs, power modules or cooling units) or even outside the shelf. [35]

A full MicroTCA system can consist of one or more shelves which are managed by a system manager. The system manager is not strictly required by the MicroTCA standard and can be implemented either as a remote application running on standard PC hardware, or on an AMC or MCH equipped with a CPU. [35]

MicroTCA Carrier Hub

The MCH assumes the functions usually held by the ATCA carrier board, ATCA shelf manager, and ATCA switch board for both cost and space reasons. It is connected to the AMC modules through a double star Ethernet bus, and a star Intelligent Platform Management Interface (IPMI) bus called IPMI local (IPMI-L). The connection to the power modules and cooling units is implemented through a double star IPMI bus called IPMI Channel 0 (IPMI-0). (See figures 6.1 and 6.2.)

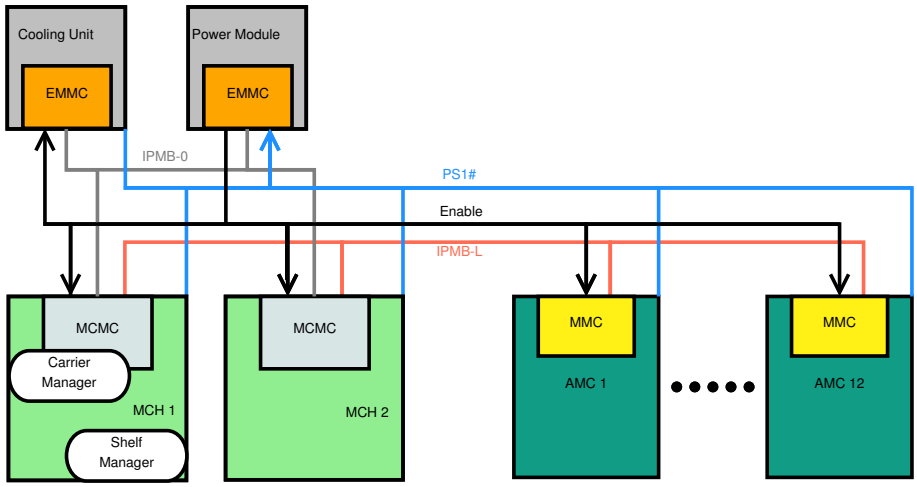

Figure 6.1: Management aspects of a MicroTCA shelf. MCHs, power modules, and cooling units are connected by IMPI-0; MCHs and AMCs are connected by IPMI-L. PS1# is used to assert presence by an FRU, Enable is asserted when the FRU has been enabled.

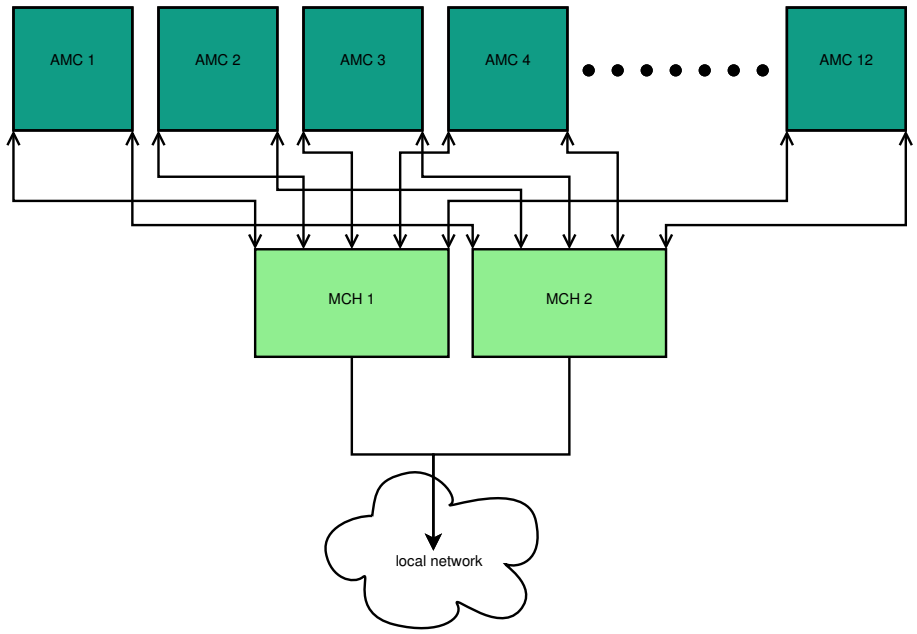

Figure 6.2: The MCHs include an Ethernet switch to implement a double-star Ethernet network. For this reason the signal density near the MCHs is very high.

Through the IPMI bus the MCH can instruct the primary power module to activate and deactivate all components, check temperatures and other sensor readings, as well as control fan speeds.

This tight control enables hot swapping by E-keying. Before an AMC module is powered with 12 V its MMC is enabled with a supply voltage of 3 V. Then some standard parameters are read out of the MMC in order to check compatibility with the crate. Only if compatibility is detected the module is powered fully with 12 V. [36]

The MCH generally acts as the hub for all protocols routed over a MicroTCA system's back-plane. It provides an Ethernet switch and usually routes signals traveling on the so-called fat-pipes of the back-plane between AMC modules. To do so it can be equipped with appropriate fabric tongues depending on the protocol that should be deployed on the fat-pipes.

Advanced Mezzanine Card

The AMC module is the actual workhorse of the MicroTCA platform. It can be equipped with front-panel I/O and a powerful FPGA via an FPGA Mezzanine Card (FMC) and optionally provides a general purpose CPU capable of running a rudimentary GNU/Linux operating system. One MicroTCA crate can hold up to 12 AMC modules.

In the context of the CMS trigger upgrade project AMC modules will be used as carrier cards for FPGAs which then run the trigger algorithms.

6.2.2 Advantages

MicroTCA is an industry standard rising in popularity. This will ensure the long-term availability of both replacement modules and support. It will also lead to a thriving ecosystem of various readily available modules, similar to the current situation with VME. Being a standard also mitigates the drawbacks of vendor lock-in. Soft- and firmware developed for a certain MicroTCA module should be able to be ported with minor changes to a new system if the need arises. [33]

It supports system level health management, is fully redundant, and supplies high bandwidth via the back-plane (~ 6 Gbit/s). Furthermore the ability to hot-swap all FRUs such as the AMC modules as well as cooling and power units together with the possibility to run a fully redundant MicroTCA system will allow the Level-1 Trigger to offer very high availability.

6.3 Problems and challenges

Several problems arise when using MicroTCA as a replacement technology in high-energy physics. This results from MicroTCA being designed for use

in the telecommunications sector. However these challenges can be mostly worked around.

6.3.1 Galvanic input signals

Only the rather high-level components (such as the GMT, GCT and GT) in CMS's triggering system will be upgraded to the MicroTCA standard. This means that a sizable part of the input signals will still be arriving via many galvanic connections. This is especially a problem for the GMT's upgrade as its current front-plane is fully occupied by input cables and a MicroTCA module is significantly smaller than a VME module.

It has been proposed to solve this problem by converting the electrical input signals to optical ones in so-called conversion crates which can be built from cheaper and simpler hardware. (See figure 6.3.) Optical cables will have a smaller footprint and could fit onto the new MicroTCA-based modules. It needs to be investigated whether the so-introduced conversion time will add significantly to the Level-1 trigger's latency. Such a system is currently in the conceptual stage and could be implemented using MicroTCA modules.

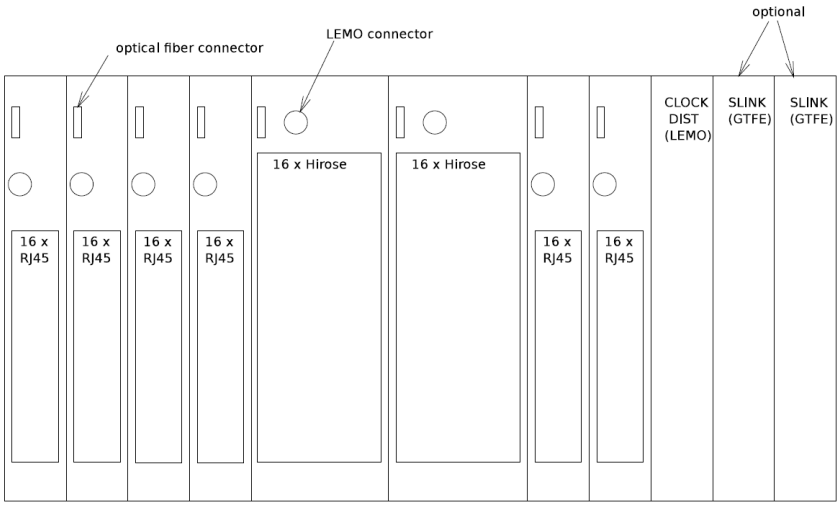

Figure 6.3: Conversion crates can provide a cheap and effective way to offer backwards compatibility by accepting galvanic links whose signal is then converted to optical signals for use in the MicroTCA systems.

6.3.2 Real-time communication via the back-plane

The various trigger modules require a near-real-time communication protocol via the back-plane. Currently all solutions for communication via the back-plane incorporate some protocol with a significant amount of latency.

Working in tandem with the HEPHY Vienna Trigger group Vadadech Inc. has developed a tongue for the MCH that allows developers to deploy a custom, potentially non-latency protocol via the fat-pipes on the back-plane.

6.3.3 System monitoring

A MicroTCA-based system offers comprehensive system monitoring capabilities. Ideally this information should be exposed to the Trigger Supervisor and potentially other monitoring systems through a common interface in order to enable engineers to swiftly debug possible problems.

6.3.4 Memory access to FPGAs on AMC modules

A further challenge introduced when migrating the VME-based system to MicroTCA is a lack of direct access to the memory registers of an FPGA mounted on an AMC module. Such register access is necessary to store various configuration settings before a run as well as for online monitoring of key values such as the trigger accept rate.

While a VME crate is connected via a PCI-bus to a controlling PC, exposing also its FPGA's registers in this way, a MicroTCA system is connected to a node PC via Ethernet. Thus it is necessary to implement a hardware abstraction layer (HAL) via Ethernet in order to be able to realize this memory access.

A proposal for such a HAL is presented in chapter 7.

Chapter 7

The Global Trigger hardware abstraction layer

Control and monitoring systems for the GT routinely access the memory registers located on the FPGAs that implement the GT's logic. This is done in order to load operating parameters into the GT and to read out information about the trigger performance during a specific run. The former VME-based system was connected to its controlling PCs directly via PCI-bus, which made it comparatively easy to access the FPGA's memory registers. Furthermore it lacked system-level health management features which would improve the ability for online debugging and monitoring of the trigger electronics.

A MicroTCA system in contrast offers system-level health management which should be exposed to control and monitoring systems in the CMS Control Room. Furthermore, register-level access is not easily achievable anymore as the only communication to the outside-world offered by the GT MicroTCA system is Ethernet, which introduces significant overhead.

It would also be desirable to represent the memory to be accessed as "items" that are represented as strings, integers, and maps towards the control and monitoring software as opposed to address offsets with the true byte-representation. It was therefore agreed that an abstraction layer to provide this should be part of a future framework.

Such an abstraction layer has been implemented in the so-called GT HAL framework. This framework provides a custom communications protocol for register-level access to the AMC module's FPGA as well as software for both the AMC module and the controlling node PC. In the future it should include an IPMI interface to access the system-health information offered by the MicroTCA system. All this information can then be exposed in a unified interface by the node PC.

7.1 Hardware components

The evaluation MicroTCA-system used to implement the current version of the GT HAL consists of the following major components:

- the MCH UTC001
- the AMC AMC514
- the crate VT891

All these components were manufactured by Vadatech Inc.

As described in chapter 6 the MCH also acts as an Ethernet switch, which means that it was possible to directly communicate with the AMC514 module witout the need for any advanced routing.

The AMC514 module includes a PowerPC (PPC) processor capable of running a rudimentary GNU/Linux operating system. From this system it is possible to access the FPGA's memory registers via PCIe.

7.2 Architecture of the proposed HAL

It was decided to implement the GT HAL framework in a modular way, thereby allowing parallel development of the different components. It consists of the following components:

- amc514d service
  The daemon running on the AMC module responsible for performing read and write operations on the FPGA's memory registers
- AmcInterface
  The abstraction and communication layer for the node PC. Abstract items are defined in an XML file
- GtControl
  A collection of command-line scripts in order to utilize the framework
- GtWebControl
  A web client for access to abstract items stored in the FPGA's memory registers

This implements the software model described in figure 7.1.

As rapid development is a major concern it was decided to implement the software for the node PC in the Python programming language. The AMC module's daemon had to be implemented in C++. Furthermore, usage of free and open technologies already available was preferred as this also contributes to faster development and easier maintenance.

In the following the individual components of the GT HAL framework will be described.

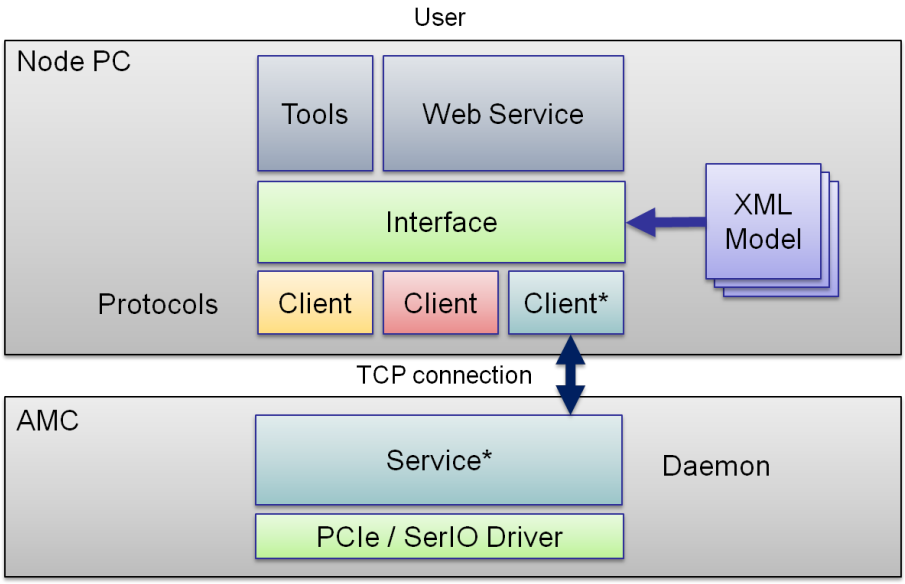

Figure 7.1: Block diagram showing the software model for the currently implemented GT HAL framework. *Web Service* is provided by GtWebControl, *Tools* is implemented in the GtControl package, *Client* as well as *Interface* layers are bundled in AmcInterface, and *Service* is realized by amc514d. Multiple Clients can be instantiated for parallel access to multiple cards.

7.2.1 amc514d service

The amc514d service running on the AMC module accesses the memory registers and exposes these via TCP. The node PC can then access these registers via the AmcInterface introduced in section 7.2.2.

The service internally uses a network library to communicate with the node PC as well as the so-called libmtca library developed in-house for driver access (figure 7.2).

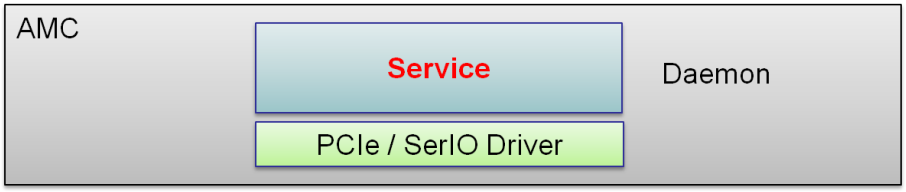

Figure 7.2: The software stack on the AMC module consists of the driver to access the FPGA's memory registers, the amc514d binary, and the TCP stack provided by the GNU/Linux operating system on the PowerPC.

7.2.2 AmcInterface

In order to connect to the AMC daemon the node PC runs a software package which provides a communication as well as an abstraction layer.

In the following the communication layer is referred to as "Client layer" while the abstraction layer is implemented as "Interface layer".

Interface layer

The Interface layer translates "items" that can be defined in an XML file and represent actual objects and values understood by the trigger algorithms such as `RandomTriggers` or `True/False` into hardware addresses in order to potentially access only bit ranges within the individual registers. In this way a 32-bit register can be used to store e.g. up to 32 boolean values. This abstraction layer furthermore provides the possibility of using several registers to store large strings or look-up tables.

The Interface layer can then be used conveniently as shown in listing 7.1.

Listing 7.1: Accessing items is convenient using the abstraction layer provided by AmcInterface. Reading a `value` returns the hexadecimal representation in the actual register, reading `text` returns the representation as defined in the Interface layer. Writing always uses the representation defined in the Interface layer.

```
amc.read('Temp').value() => 0x5634
amc.read('Temp').text() => '34.5℃'
```

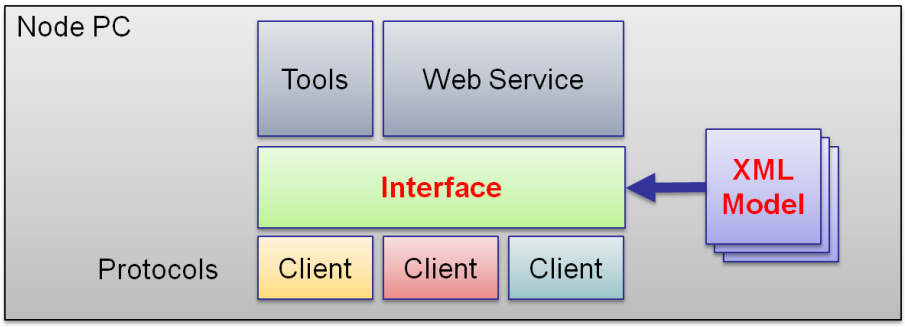

Figure 7.3: The Interface layer is responsible for the mapping of abstract items to hardware addresses as described in an XML document.

```
amc.read('Counter').value() => 255
amc.read('Counter').text() => 'Overflow'
amc.write('Mode', 'Idle') == write('Mode', 0x3)
```

Client layer

The Client layer serializes data structures containing the information provided by the Interface layer (i.e. the addressing information along with a potential value to be written) and then deserializes the incoming TCP stream to data structures which are then passed on to the Interface layer.

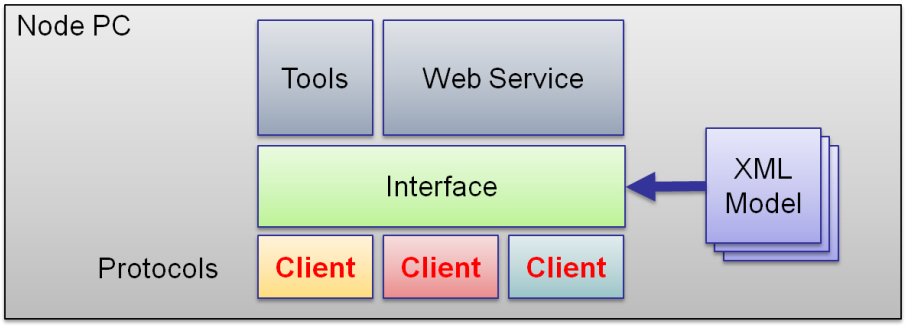

Figure 7.4: The Client layer connects to the amc514d service running on the AMC module using a custom protocol.

It is implemented as a Python class which provides `read` and `write` as public methods. Both of these use the `request` method which in turn calls the appropriate methods to serialize, send, receive, and deserialize the messages from and to the AMC module's amc514d service.

7.2.3 GtControl

For debugging and basic maintenance a suite of scripts was written. Using these it is possible to gain direct access to the FPGA's memory registers or profile the performance of the GT HAL framework.

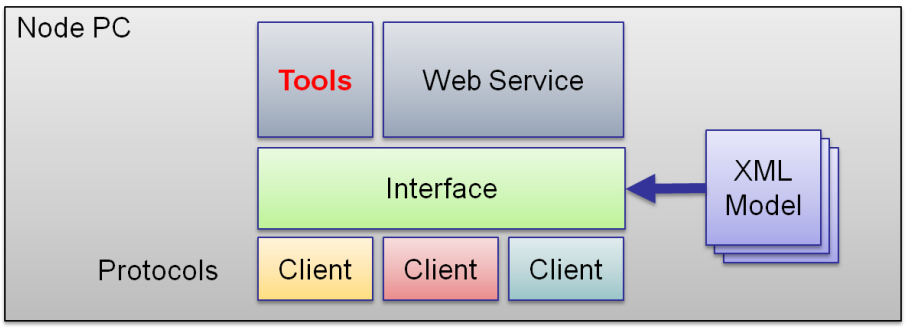

Figure 7.5: The GT HAL framework provides a tool suite to allow easy debugging by engineers.

The profiling tools contained in the GtControl package were heavily used for the performance tests presented in section 7.4.4.

The GtControl suite accesses the hardware without using the Interface layer, thereby allowing low-level memory register access but sacrificing the comfort of using predefined items. To utilize the Interface layer a web client was written. This decision was made as a web client can be easily accessed from any platform and extended more flexibly and rapidly than a native GUI.

7.2.4 GtWebControl

Using the Django framework a web client (figure 7.6) for the GT HAL framework was written and deployed. Using this client it is possible to access the physical items described through the abstraction layer. It can be flexibly extended to support several AMC modules. This is currently being used for day-to-day work on the AMC module in the test crate at HEPHY Vienna.

7.3 Evaluating potential technologies

In order to find possible bottlenecks in the proposed system to access hardware registers via Ethernet, a very basic communications protocol was implemented at the application layer. Following this a more sophisticated protocol could be developed which implemented the lessons previously learned. Thus the Client layer was developed independently from the Interface layer which allowed development of the Client layer to use an exploratory approach while not affecting the Interface layer adversely.

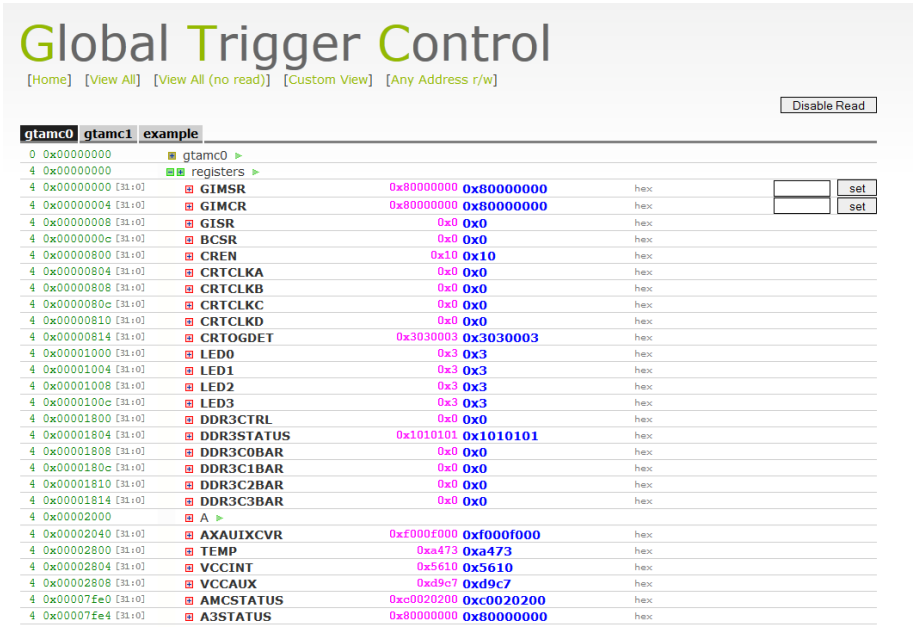

Figure 7.6: The web client GtWebControl offers access to the items as defined in the Interface layer.

When developing the evaluation protocol fast development as well as the possibility for easy testing and debugging were paramount. Therefore it was agreed to use a text-based protocol capable of doing one atomic operation per connection and using TCP at the transport layer. A typical interaction is shown in listing 7.2.

Listing 7.2: An example interaction between Client layer and amc514d service in the text-based protocol as seen from the Client-side.

```
CONNECT
> write [addressing information] [value]
> [reply from service]
DISCONNECT
...
CONNECT
> read [addressing information]
> [reply from service]
DISCONNECT
```

While this implementation is not optimized for speed it allows easy testing and debugging by using a simple telnet client and can be quickly implemented as no low-level byte-stream manipulation is needed.

7.3.1 Evaluated technologies

Local access

A benchmark for the fastest possible access needed to be set so that comparison with the network-bound methods was possible. To this end a simple program was written that directly interfaces with the kernel driver and performs several read and/or write operations on the registers in question. Care was taken so that only read- and writable registers were checked.

Timing was done using `clock_gettime()`, a Unix-specific function providing high-precision time information.

Local access was measured to be about as fast as would be expected for operations on memory registers with an access time of about 1 µs for write access and approximately 4 µs for read accesses. The bandwidth is therefore less than 4 MB/s.

It is notable that a read operation seems to be slower by about a factor 4 than a write operation.

Testing the TCP/IP-based protocol

As the designated AMC module for the GT is equipped with a PPC CPU running the GNU/Linux operating system, a logical first step was to use the TCP protocol, as it offers several beneficial features such as guaranteed transmission of messages and flow control.

The first naive implementation used a TCP server on the AMC module that was not capable of accepting multiple connections at once. This meant that parallel access showed very erratic access times, because a second connection attempt was refused by the amc514d service requiring the Client application to retry sending.

A further problem that surfaced were unacceptably slow transfer times. The reason for this could be traced to the nature of the evaluation protocol described in listing 7.2. TCP requires no unhandled packets to be left on the line which requires the disconnect operation to wait for potentially orphaned packets. As one register access transmitted comparatively small amounts of data, the disconnect operation added significant overhead.

While it would be possible to switch to UDP in order to potentially increase transfer speed, an alternative which does not necessitate the replacement of TCP in both the client as well as the daemon is a persistent connection between the client and daemon. Although this requires a forking or multi-threaded daemon to handle multiple connections the client can be left largely unchanged.

Modifying our system to use a persistent TCP/IP connection during the transfer of a fairly large data block results in an improvement in transfer time of about a factor 10.

The protocol used in this test is described in listing 7.3.

Listing 7.3: An example interaction between Client layer and amc514d service in the persistent TCP protocol as seen from the Client-side.

```
CONNECT
> write [addressing information] [value]
> [reply from service]
...
> read [addressing information]
> [reply from service]
...
DISCONNECT
```

Investigating UDP

UDP speed tests were performed to gauge the feasibility of switching away from a TCP-based protocol. No actual read or write operations were performed on the hardware, but a simple echo service was implemented using UDP. In order to estimate the improvement when using UDP one message passed via this protocol was considered equal to one read or write operation via TCP.

In order to evaluate UDP's possible speed advantage compared to TCP in our setup two very crude tests were done. In both tests the actual memory on the AMC was never accessed.

In the first test a protocol as described in listing 7.2 was used, which means that the client needed to wait for a reply after each of its send operations. This method of transmission did not offer significant improvements over the TCP-based protocol when using a persistent connection. In the second test the client sent messages as fast as possible without waiting for a reply from the server. This second method increases possible transfer speeds by a factor of more than 10 when compared to the challenge-response protocol.

7.3.2 Bandwidth

Computing transfer rates for the different means of access is not straightforward due to the nature of the tests. Seeing as different application level protocols as defined in the OSI reference model [37] were tested it is useful to compare the effective bandwidth (i.e. considering only the information to be written as payload) with the bandwidth obtained when considering any transmitted data as payload as well as with the maximum possible bandwidth supplied by the hardware.

The bottleneck on the network side was a 100 Mbit/s switch, so we would expect a transfer rate of at most 12.5 MB/s. The final access speed is however limited by the transfer rate for direct access, which is less than 4 MB/s as was seen in section 7.3.1.

Effective bandwidth

As shown in listings 7.2 and 7.3 the transmitted text string containing the information to be written was formatted as `write x y z` where `z` is the information to be written. During the tests the information to be written on the AMC module per operation was encoded in one byte of information.

This leads to the following transfer rates:

- TCP "standard": 0.17 kB/s
- TCP persistent: 1.87 kB/s
- UDP "send-rcv": 1.65 kB/s
- UDP "burst": 25 kB/s

Full sending bandwidth

The bandwidth when also considering the overhead incurred through addressing and other meta-information is significantly higher:

- TCP "standard": 2.33 kB/s
- TCP persistent: 26.2 kB/s
- UDP "send-rcv": 23.15 kB/s
- UDP "burst": 350.2 kB/s

7.3.3 Lessons learned

As could thus be seen the effective bandwidth is severely harmed by the overhead of the preliminary protocol. Therefore it was decided to move to a more compressed, binary protocol. It is also possible to first send any write or read requests and only then receive. This will reduce transaction time by a significant fraction as well.

UDP's advantages are obviously its simplicity and speed, but it would introduce several difficulties. Among them are:

- Problematic loss of packets which leads to the need for an implementation of a custom reliable protocol, potentially negating any gains in speed over TCP
- Need for time-stamping of packets or other sequencing as UDP makes no guarantees that packets arrive in order
- Lack of congestion control

This means that using UDP would severely increase the complexity of the proposed software framework whereas TCP has this complexity included in its implementation.

7.4 A Protocol Buffer-based HAL

As the increased application complexity potentially introduced by using UDP was deemed impractical it was decided to implement a binary protocol which could reduce the overhead incurred due to sending ASCII strings. Furthermore the protocol was designed such that successive operations could be packed into one TCP transaction. Additionally minor improvements were made to optimize overall stability and speed.

Some of the improvements made are described in the following.

7.4.1 Local write on the AMC module

As it is only possible to write and read whole registers at the hardware level a write operation in the evaluation protocol required the whole register to be read out by the node PC when it was necessary to write less than the full 32 bits in a register. Then a new value was calculated for the complete register on the node PC that was finally sent to the AMC module.

This was highly inefficient when often writing only single-bit values and furthermore lead to the possibility of race conditions when moving to stacked writes or concurrent access: Two instances of GtControl would access the same register at the same time and read out the same information. Both calculate a different new value for the register in question which is then sent to the AMC module. The value that arrived later would overwrite the changes done by the other instance of the GT HAL framework.

In order to remove the risk of such events it was decided to move this logic to the AMC daemon. The node PC would send a modified write command consisting of the usual addressing information and payload to be written as well as a bitmask specifying the bit range to be written to.

Both the bitmask and the value to be written can efficiently be encoded as integers. On the AMC daemon these can then be combined with the original value of the register by a simple bit operation in order to calculate the new value of the memory register.

7.4.2 Protocol Buffers as binary data format

Having evaluated a number of possible data structures to be used as exchange format it was decided to use Protocol Buffers. [38]

Protocol Buffers is a technology to serialize structured data, initially developed by Google. It is released as free software and is used widely within Google itself, which implies continued development and availability as well as good robustness. The general principle is as follows:

- The structure of the data is defined in a so-called .proto file which is "compiled" once with the `protoc` compiler. This results in a C++, Java or Python class to handle the future data structure

- The data itself is serialized to a binary string which can be streamed across a network via TCP but also presented in a human-readable form by the class handling it.
- Protocol Buffers are meant as a pure data-exchange format without built-in support for remote procedure calls (unlike e.g. XML-SOAP)
- Due to its binary nature Protocol Buffer data is very compact and efficiently parseable compared to e.g. XML

These properties make Protocol Buffers a very good match for our software framework.

In order to enter or extract information from the FPGA's registers five pieces of information need to be known by the amc514d service:

- the kind of operation (read or write)
- addressing information
- a bitmask
- optionally the information to be written

These requirements can be encoded easily in a simple .proto file.

Listing 7.4: .proto file for a possible binary exchange format

```
message gtMsg {
  enum operation {
    READ = 0;
    WRITE = 1;
  }
  required operation op = 1;
  optional unit32 address_offset = 2;
  optional uint32 bitmask = 3;
  optional uint32 msg = 4;
}
```

As can be seen in listing 7.4 the fields in a message are tagged with unique numbers that identify the fields in the message binary format. This enables the user to add additional fields to the format later on flexibly as long as the tags are kept the same. Tags numbered 1 through 15 use one byte to encode, while tags in the range of 16 to 2047 use two bytes. Fields can be marked `required`, `optional`, or `repeated`.

An optional field does not have to be set in a Protocol Buffer message, while the generated classes throw an exception for omitted required fields. Most fields in the data structure for the GT HAL framework were declared

as optional as this allows greater flexibility when changing the protocol at a later date. Sanity checks making sure that the required fields have been set are then implemented on the application-level. Repeated fields can be present multiple times in an individual Protocol Buffer message. These were not used in the GT HAL framework's protocol.

Classes to read and write such messages can then easily be generated from the above definition.

7.4.3 A stacked, asynchronous communications protocol

Additionally to the use of Protocol Buffers it was decided to allow multiple access commands to be grouped into one TCP message. This would reduce the overhead incurred through connection establishment and disconnection observed in section 7.3.1.

The Client layer introduced in section 7.2.2 then actually performs a dual serialization: In a first step the data structure representing an atomic operation such as read or write is serialized to a byte string using the protobuf library. The second step then serializes the individual byte strings into a stream to be sent via TCP. In order to be parsed on the receiving end the strings representing an atomic operation are separated by length fields. The whole TCP stream is finally terminated by a `NULL` field (figure 7.7).

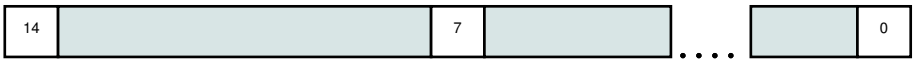

Figure 7.7: The TCP stream consists of several serialized Protocol Buffer messages representing an atomic operation, preceded by fields indicating their length. The final message is terminated by a `NULL` field.

The Client layer can be used to directly access registers on the AMC module, as shown in listing 7.5.

Listing 7.5: Accessing registers directly can be achieved by using the Client layer. Stacked operations are easily achievable by collecting several operations (stored in Python dictionaries) in a list which is then passed to the `read` or `write` command.

```
msg0 = {'addres_offset': 0x7ff4, \
  'value': 0xffffffff, 'bitmask': 0xf}
msg1 = {'address_offset': 0x7ff4, \
  'value': 0x0, 'bitmask': 0xf}
msgList = [msg0, msg1]

Client c
result = c.write(msgList)
```

```
msg2 = {'address_offset': 0x7ff4}
msg3 = {'address_offset': 0x7ff8}
msgList = [msg2, msg3]

Client c
result = c.read(msgList)
```

Similarly the amc514d service needed to be modified to support this protocol.

7.4.4 Evaluating the new system

Bandwidth

Transfer rates for the GT HAL framework showed a marked improvement when compared to those of the primitive protocol introduced in section 7.3. Access times were on the order of 50 µs which translates to 20 000 operations per second or about 80 kB/s. As explained later this access time varies slightly for different block sizes and could be improved further.

Because finding possible bottlenecks in the complete framework was a major concern specialized tools for this task were developed and included in the tools suite provided by the GT HAL framework at an early stage. The AMC daemon was additionally equipped with basic function execution timing capabilities.

Finding bottlenecks on the node PC

Using the main profiling tool various aspects concerning the runtime especially of the node PC's software stack could be investigated (figure 7.8a).

Several bottlenecks were identified. As can be seen in figure 7.8a the major part of the execution time is spent in functions controlled by outside sources such as the Google Protocol Buffer library or the TCP stack. Especially the Protocol Buffer library implemented in pure Python could potentially be swapped for its C++ counterpart to achieve significant performance gains.

Of the remaining functions especially the ones for receiving or handling Protocol Buffer data structures (`_receive`, `_serialize`, and `_proto2dict`) seem to be causing excessive execution times. The function `_proto2dict` is concerned with converting individual Protocol Buffer messages to Python dictionaries and as such should exhibit O(N) behavior. Its run-time might be slightly improved by using more efficient data structures to store the individual replies from the AMC daemon.

The functions `_receive` and `_serialize` make copious use of Python lists' appending and removing methods which present O(N) behavior as

well. These functions would certainly benefit from a faster container class replacing Python lists.

Part of the *Other* and *_receive* categories in the above runtime is spent waiting for the AMC daemon to reply to the node PC's request. For this reason the AMC daemon was profiled.

Bottlenecks on the AMC module

Profiling was not straightforward on the AMC module as its limited resources made the usual profiling tools difficult to use. Thus a simple profiling solution was added directly to amc514d's source code. This was comprised of basic timing capabilities together with output into a log file for later examination.

The profiler was thus tailored for the needs of this project and could measure the timing of tasks as opposed to just functions. As can be seen in figure 7.8b the largest fraction of the total runtime is spent receiving the incoming TCP stream. Querying the hardware constitutes another major operation. It needs to be pointed out however that a hardware query's measured runtime is about 2 µs when divided by the number of read-write operations that were done, which is to be expected for low-level register access. Parsing the individual Protocol Buffer messages uses a minor amount of time as the C++ Protocol Buffer library is implemented in a highly efficient way. It may be possible to reduce the total running time by merging the parsing and hardware querying steps. Instead of storing the parsed Protocol Buffer messages in a data structure and only querying the hardware in a later step in bulk it may be more efficient to query the hardware after every individual message was parsed.

It is worth noting that the absence of a *Sending* category in both node PC and AMC profiling is not by mistake. The send operation returns almost immediately and does not indicate a completed transmission. It is therefore not a useful gauge of network speed.

Long-term testing

Another tool (`long-term-test`) was devised to simulate the access typically seen during normal runs at CERN. This repeatedly reads out and sets various register ranges for several hours. Using this tool it could be shown that the framework works reliable over an extended period of time accessing various registers.

Profiling stacked operations

The reason stacked operations were introduced initially was for possible performance benefits as they would mitigate the overhead of a full TCP packet for one operation. In order to measure the gain in performance

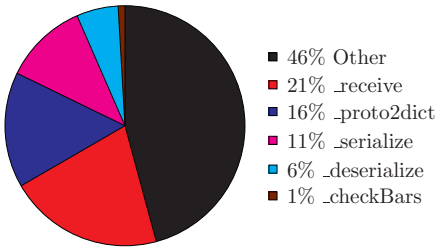

(a) Pie chart showing relative function execution times on the node PC. As can be seen a large fraction of total execution time is taken up by functions controlled by outside developers such as those used for argument parsing, for establishing the TCP connection, or those included in the Google Protocol Buffer library. It would be possible to improve the performance of the Google Protocol Buffer library by moving to a complete C++ implementation. The remaining function's large execution times could however be reduced utilizing different Python data structures.

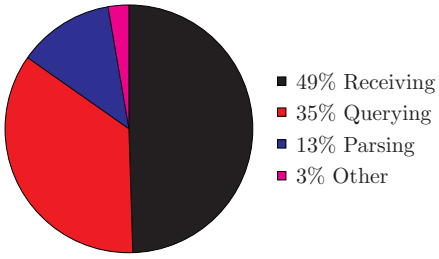

(b) Pie chart showing relative function execution times on the AMC module. As can be seen the largest fraction is taken by receiving the TCP stream and storing the individual Protocol Buffer messages in a `std::vector` data structure. Searching for optimizing strategies in this area seem most promising. Parsing the individual Protocol Buffer messages can be done very efficiently as the protobuf library is highly optimized C++. The speed for querying the hardware is limited by the hardware itself.

Figure 7.8: The relative execution times for different routines in the GT HAL framework split into the software stack on the node PC and AMC module.

over single operations a script was written to perform a certain range of read/write operations per block. It was run for different ranges and step sizes, averaging every data point over 10 runs.

Initial tests produced confusing measurements: While increasing the block size from 1 to 1000 operations resulted in reduced transaction times per operation they seemed to rise when doing the same test up to 1 million operations.

Decreasing the step-size and examining various ranges of operations per block showed that the transaction time per operation falls with growing block size until about 15 000 operations per block, at which point it first flattens out until about 50 000 operations per block where it begins to grow again (figures 7.10a, 7.10b, and 7.10c). When increasing the block size up to 450 000 operations per block the performance decreases even further (figure 7.10d). This is probably due to increased overhead on the client side.

Specifically Python's data structures might begin to show performance problems at this size of data structures – a data structure containing 100 000 access instructions would use at least 500 kB disregarding any overhead required for book-keeping or memory alignment. As the typical memory to be accessed will only contain about 10 000 registers this is no big concern for the GT project.

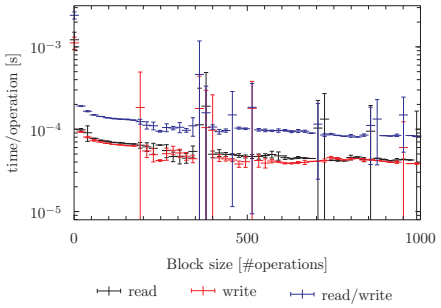

(a) Plot from 1 to 1000 operations per block, 100 points of data taken

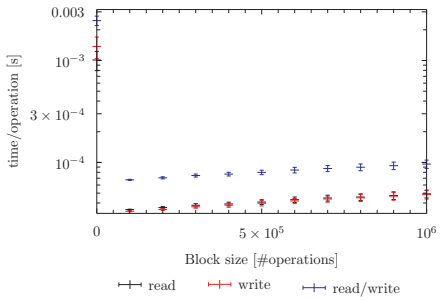

(b) Plot from 1 to 1 million operations per block, 10 points of data taken

Figure 7.9: Initial tests comparing the execution speed per operation for different block sizes showed mixed results: When plotting block sizes from 1 to 1000 operations per block execution times per operation clearly decreased. Conversely when extending this test to 1 million operations per block the execution time per operation rose. It seems that at some point complexity overheads from data structures in the Client layer dominate the total execution time.

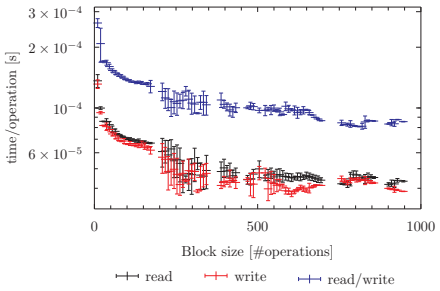

(a) Plot from 1 to 1000 operations per block, 100 points of data taken. In this plot it can be clearly seen how execution time per operation decreases with increasing block size.

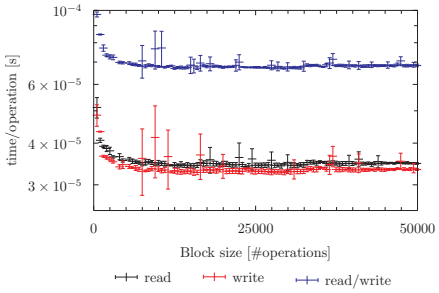

(b) Plot from 1 to 50 000 operations per block, 100 points of data taken. Execution time per operation still decreases with growing block size, but begins to flatten out.

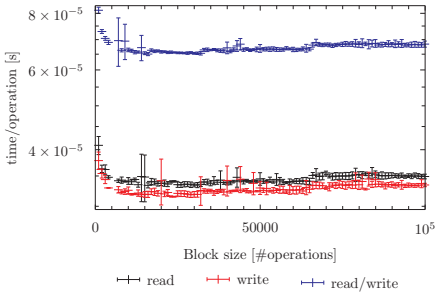

(c) Plot from 1 to 100 000 operations per block, 100 points of data taken. Execution times per block size begin to flatten out at 15 000 operations per block. An increase can be observed from about 50 000 operations per block.

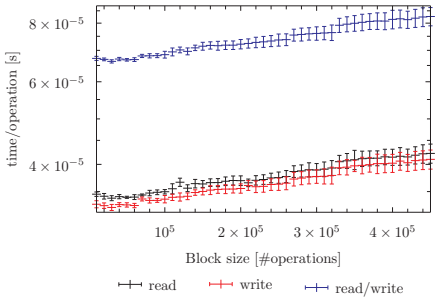

(d) Plot from 1 to 450 000 operations per block, 50 points of data taken. On this scale a dip at about 60 000 operations per block can be seen with the execution time per operation rapidly increasing with growing block size from this point on.

Figure 7.10: Plots using finer step-sizes while measuring the execution time per operation for various block sizes. The data points for 1 operation per block was excluded as were several data points that were identified as outliers. All data points were averaged over 10 runs.

Comparing read and write times

A large discrepancy was observed in section 7.3 between individual read and write access. To determine whether this difference is still observed in the larger framework the daemon was profiled separately for read and write accesses. Results indicate a slight excess in time for writes compared to the time for reading (table 7.1). This is in stark contrast to the results obtained earlier, which indicated a significantly higher time for reading.

Run	Time for read	Time for write
1	1060.23 ms	1208.27 ms
2	1055.25 ms	1207.3 ms
3	1068.42 ms	1231.57 ms
4	1056.37 ms	1206.83 ms
5	1065.6 ms	1213.83 ms
6	1054.56 ms	1210.87 ms
7	1055.38 ms	1215.46 ms
8	1056.45 ms	1211.47 ms
9	1073.78 ms	1226.95 ms
10	1060.11 ms	1207.97 ms
Average:	1060.615 ms	1214.052 ms

Table 7.1: Results for separate profiling of read and write access. While profiling of direct access showed a significantly slower read access as opposed to write access the framework seems to offset this difference. This might be due to higher overhead when writing on the daemon-side. One run consisted of a block of 50 000 operations.

Further investigating this result showed that the different execution times between read and write access result from the increased size of the data structures storing write commands. (See table 7.2.) This can be clearly seen as the times to actually query the hardware were very similar with a slightly longer time for reading from registers while tasks such as receiving and parsing the access commands took longer for write access.

7.5 Conclusion

A solution for register-level access via Ethernet was presented. The GT HAL framework offers as custom communications protocol via TCP, it also provides the ability to use abstract items representing various data structures which are stored in the accessed registers.

Furthermore the GT HAL framework includes a web interface to modify these abstract items as well as a suite of tools to access the MicroTCA

Run	Parsing	Query hw	Receiving
1	130 ms	410 ms	483 ms
2	129 ms	405 ms	484 ms
3	129 ms	405 ms	497 ms
4	129 ms	406 ms	485 ms
5	129 ms	406 ms	494 ms
6	129 ms	405 ms	483 ms
7	129 ms	405 ms	485 ms
8	129 ms	405 ms	485 ms
9	129 ms	406 ms	502 ms
10	129 ms	406 ms	488 ms
Avg:	129 ms	406 ms	488 ms

(a) Timing for 50 000 read operations.

Run	Parsing	Query hw	Receiving
1	156 ms	396 ms	628 ms
2	156 ms	396 ms	627 ms
3	156 ms	396 ms	651 ms
4	157 ms	395 ms	627 ms
5	156 ms	396 ms	634 ms
6	156 ms	395 ms	631 ms
7	156 ms	396 ms	635 ms
8	156 ms	397 ms	631 ms
9	156 ms	395 ms	647 ms
10	156 ms	396 ms	628 ms
Avg:	156 ms	396 ms	634 ms

(b) Timing for 50 000 write operations.

Table 7.2: More accurate results for profiling of read and write access. It appears that the main culprits for the increased time a write operation takes are the functions responsible for parsing and receiving. Querying the hardware is almost equally fast for both operations. Data was taken for a block of 50 000 operations.

hardware's registers directly via Ethernet, diagnose problems and profile its performance.

As was expected a performance penalty was incurred through both the abstraction layer and the need for Ethernet access. This penalty reduces the bandwidth from at most 4 MB/s to 80 kB/s. This ratio could be improved as was described in section 7.4.4.

7.5.1 Future plans

Providing an IPMI interface to the GT HAL framework

Adding an IPMI interface to the node PC is one of the next planned steps. In this way the node PC could provide a unified interface for a MicroTCA system to the outside world. (See figure 7.11)

This interface would communicate with an IPMI daemon running on the MCH module via TCP/IP over Ethernet analogous to the current AMC daemon system.

Hardware architecture for the GT HAL framework

While the current design consists of a node PC interacting directly with each AMC module (see figure 7.12a) it is planned to include a so-called Processor

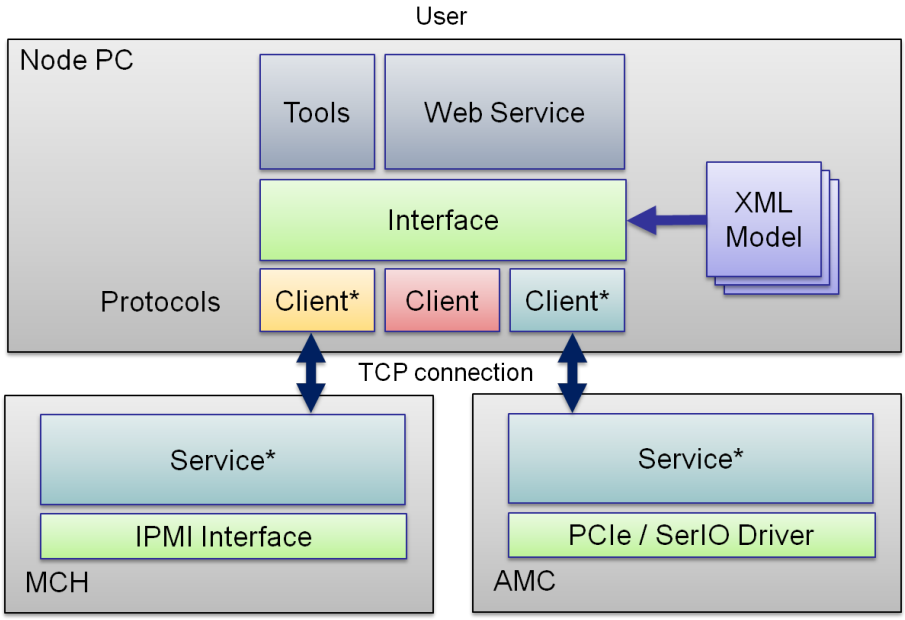

Figure 7.11: The current GT HAL framework, consisting of a daemon running on an AMC module and a hardware abstraction and communication layer running on the node PC would be extended to connect to an IPMI interface running on the MicroTCA system's MCH. It would be implemented analogous to the hardware abstraction layer via a TCP/IP transport layer over Ethernet. In the graphic, a star at the object's name indicates it is currently communicating.

AMC (PrAMC) in the MicroTCA system. A PrAMC includes only a PPC CPU and can then connect to AMC modules via the fat pipes on the back-plane. Therefore the AMC daemon would run on the PrAMC and work as the single point of access for the node PC (figure 7.12b).

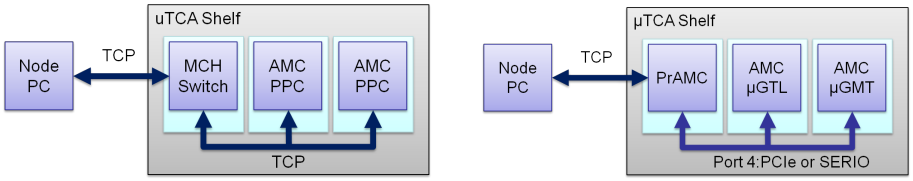

(a) The current system contains an MCH that routes TCP traffic to the AMC modules. The node PC acts as control hub for several modules.

(b) The proposed system containing a PrAMC that connects to the node PC. It then acts as the control hub for several AMC modules.

Figure 7.12: The current evaluation system consists of a node PC, an MCH, and an AMC module. The AMC module houses a PowerPC (PPC) CPU running the daemon to access the FPGA's memory registers. The proposed future system would contain a so-called PrAMC which houses the PowerPC CPU. The AMC modules would only hold the FPGA which can be accessed via the back-plane's fat-pipe.

Improving the framework's performance

An improved access time for memory registers could probably be attained by re-evaluating the choice of data structures in the communication layer on the node PC. Collections of operations are currently implemented as Python lists which could be replaced by Python deques as already now all operations are done on the beginning and end of lists which are O(N) for lists but only O(1) for deques.

If these changes do not bring on the desired improvements a further possibility is to re-implement the current communication layer in C or C++. However, this would mean an increased development and maintenance effort.

Integration into the CMS infrastructure

To integrate the current GT HAL framework into the future CMS network it might be necessary to switch to the so-called µHAL system which consists of a Control Hub on the node PC and currently uses the IPBus protocol over UDP to communicate with the AMC module's FPGA. The GT HAL framework could be adapted to utilize µHAL's Control Hub but continue to use the current Protocol Buffer-based protocol. This could be a viable course of action as maintenance work could so be spread over several teams in the CMS experiment.

Chapter 8

Improving the trigger's efficiency

8.1 Quality assignment in the GMT

As described in chapter 5 the primary role of the trigger is to reduce the rate of recorded events while retaining most interesting data. The trigger internally classifies each muon using quality categories. These are first assigned in the regional triggers. Quality assigned by the regional triggers is determined from the number and configuration of layers in the muon system that detected the respective candidate particle. This is a measure for the accuracy with which the muon's transverse momentum (p_T) could be measured.

The GMT then combines the muon candidates received from the regional triggers (up to 4 each from the DTs, the CSCs, the RPCs in the barrel, and the endcap RPCs) and transmits the best four candidates to the GT. The quality assigned to GMT candidates depends on whether the muon candidate could be confirmed and if it was not confirmed, by which of the sub-systems it was detected.

The redundancy of the CMS muon system helps to reduce the number of ghost candidates and aids background rejection. Candidate muons which are found in two complementary sub-systems (i.e. either DT+RPC or CSC+RPC) are automatically considered confirmed and accepted regardless of their regional quality. Candidates with low regional quality which are found in only one sub-system can be tagged for exclusion from single- and di-muon triggers in certain pseudo-rapidity regions. [39]

The GT decides, based on the estimated transverse momentum of the muon candidate and on the quality assigned to it by the GMT whether the muon candidate can contribute to trigger the CMS data acquisition.

In order to decide which categories of muon candidates should be excluded from single- and di-muon triggers, studies were undertaken during

the GMT design-phase using simulated data. In this work, these results are verified and re-evaluated using data taken in the year 2010 which amounted to a total integrated recorded luminosity of $43.17\,\mathrm{pb}^{-1}$. [40]

8.2 Techniques

Tuning of the GMT quality assignment has been done by examining histograms of both the rate and efficiency of all recorded muons split by their quality and binned in the histograms according to their pseudo-rapidity.

In the following the techniques for determining the trigger rate and efficiency will be discussed. While the technique for determining the trigger rate is fairly straightforward, there are several techniques to find the trigger efficiency depending on the kind of data set used.

8.2.1 Rates

To measure the relative trigger rates per pseudo-rapidity bin, each event in a minimum bias data sample is analyzed whether it contains a muon candidate that caused a trigger. This muon is then examined whether it passed a transverse momentum cut. If these requirements are satisfied, the muon is accounted for, according to its quality and its pseudo-rapidity, in the histogram.

The data for this study was extracted from the so-called NanoDST stream. This stream consists of every tenth event (it is pre-scaled by a factor of ten) that caused a L1A signal to be sent. For storage reasons these events record only limited information since they will not be used for physics analysis.

8.2.2 Efficiency

When determining the trigger efficiency using minimum bias samples a rather straightforward algorithm can be used. Only events with one reconstructed muon are examined. This muon's transverse momentum must pass a certain threshold, and must have been detected within the pseudo-rapidity area that is covered by the muon system.

All such muons are inserted into a normalization histogram in which they are binned according to their pseudo-rapidity. Then each of these muons is checked whether it can be associated to a muon candidate in the GMT. In that case it is inserted into the appropriate histogram depending on its quality. These histograms are then divided by the normalization histogram in order to obtain an efficiency.

While this method is preferred due to its simplicity, it is only possible to use it with minimum bias data. However in 2010 not enough data containing

muons with high transverse momentum were collected with minimum bias triggers to allow a statistically significant analysis using this method.

Tag and Probe

In order to use data that were collected with the help of the GMT for efficiency studies a technique called "Tag and Probe" can be used. This makes it possible to avoid the bias from using events triggered by the muon trigger itself.

The Tag and Probe technique uses events identified as containing mass resonances (such as Z bosons in the case of measurements with muons) to establish the trigger efficiency without introducing bias.

The Z boson events used in this analysis decayed following $Z \rightarrow \mu\bar{\mu}$. As the Z boson's mass is about $90\,GeV/c^2$ most of its decay products will have transverse momenta greater than $20\,GeV/c$.

As the two muons were triggered on independently of each other it is possible to use one of the two as the "tag" and the other as the "probe" which is used to actually determine the trigger efficiency.

The data set used in this analysis was obtained by filtering prompt isolated muons originating from Z and W boson decays.

If one of the muons in events identified as Z decays can be associated to a L1 trigger candidate, this muon can be used as a "tag" while the other will then be used as a "probe". If also the "probe" can be associated to a L1 trigger candidate, the "probe" counts as efficient. If both muons in the event were triggered both represent a "tag" and at the same time count as an efficient "probe".

Additionally both "tag" and "probe" muons are required to pass certain thresholds for their transverse momentum and quality. The "tag" muon has to pass the requirements of the L1_SingleMu7 algorithm. This means a GMT quality greater than 3 and a transverse momentum greater or equal to $7\,GeV/c$. The "probe" muon is required to pass a threshold for both the transverse momentum calculated during offline reconstruction as well as for the transverse momentum calculated in the GMT – these were $20\,GeV/c$ and $14\,GeV/c$ respectively in this study.

The algorithm used to implement tag and probe when measuring the trigger efficiency can be summarized as follows.

Every reconstructed muon is examined whether it satisfies several cuts:

- $|\eta| < 2.4$
  Determines whether the muon was found in the acceptance region of the muon system

- $p_T \geq 20\,GeV/c$
  The p_T – cut on the reconstructed transverse momentum

- global muon?
  Determines whether the muon was found in both the tracker and the muon system

If a muon can be found that satisfies these cuts, other reconstructed muons in the event are examined whether they qualify as a tag.

A muon must then satisfy the following requirements to qualify as a tag:

- detected in the GMT?
  Determines whether the reconstructed muon was found by the GMT during data taking

- $q_{GMT} > 3$
  Quality requirement posited by the L1_SingleMu7 algorithm in the GT

- $p_T \geq 7\,\mathrm{GeV/c}$
  The $p_T -$ cut used by the L1_SingleMu7 algorithm in the GT

If a tag and probe muon are found in the same event a normalization histogram is filled with the probe muon. Then the probe muon is checked whether it can be matched to a muon candidate in the GMT and whether its transverse momentum as determined by the GMT exceeds $14\,\mathrm{GeV/c}$. If this is the case, the histograms for the appropriate quality are filled with the probe muon.

Finally the latter histograms are divided by the normalization histogram to obtain the GMT efficiency.

8.3 Results

As can be seen when comparing figures 8.1a and 8.1b the current cuts in place for single muon triggers greatly reduce the rates at high pseudo-rapidities while leaving the rates at pseudo-rapidities of less than 1.5 comparatively the same.

Of especial noteworthiness is the hot tower at $|\eta| = 1.5$, largely dominated by muon candidates measured in the CSC with quality 2. This region corresponds to the gap between two CSC rings. Following this measurement, the CSCTF algorithm was modified to mitigate extra rates originating from this region since suppressing these candidates at the level of GMT would result in a substantial loss of efficiency.

Other, less discernible regions, such as at $|\eta| = 1.$ and $|\eta| = 0.6$, are a result of discrete binning in regional muon trigger systems.

The GMT efficiency is relatively high at $|\eta| < 2$ but falls off significantly at larger pseudo-rapidities (figure 8.2).

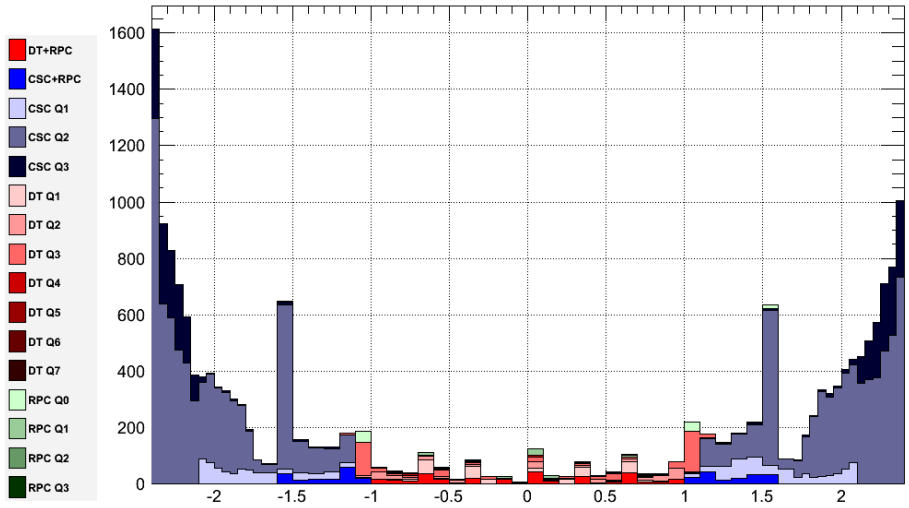

(a) All qualities

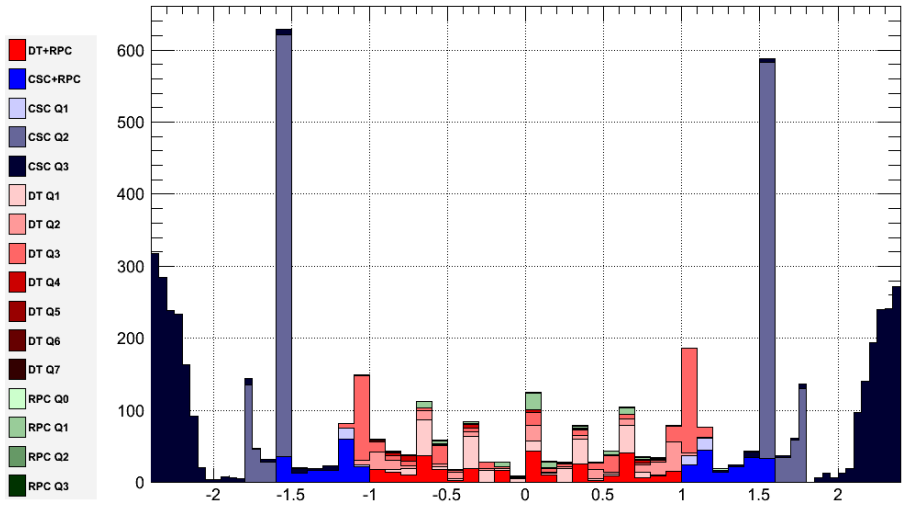

(b) After current quality cuts

Figure 8.1: Rates for the L1_SingleMu14 trigger algorithm split into different quality categories and plotted against the muon's pseudo-rapidity as determined by the GMT. As can be seen most of the rate comes from CSC qualities 2 and 3.

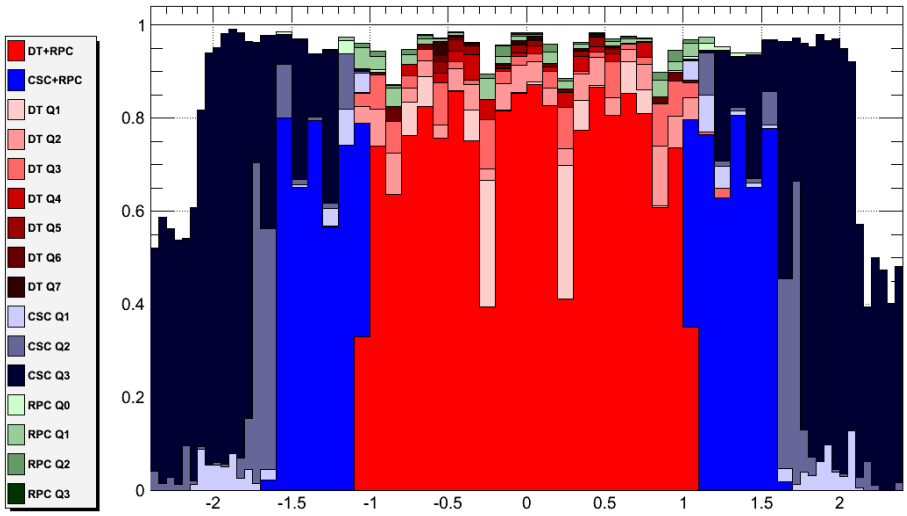

(a) All qualities

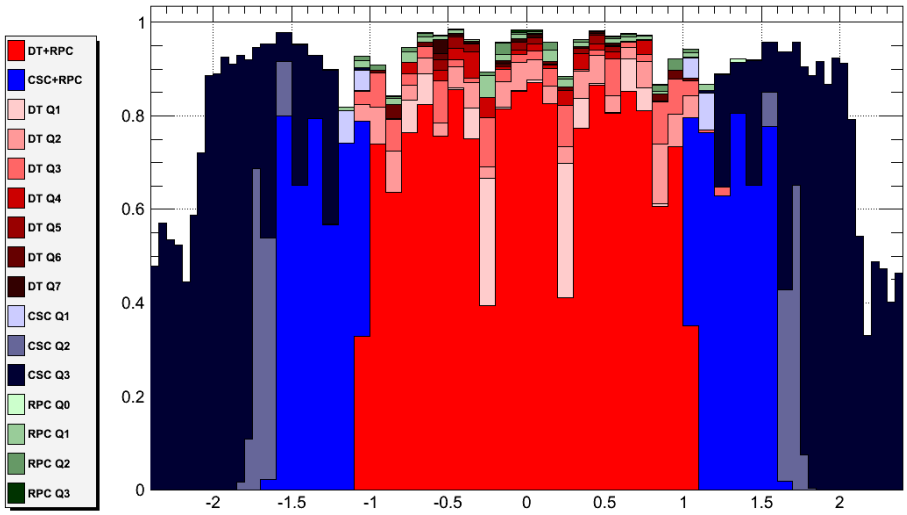

(b) After current quality cuts

Figure 8.2: Efficiencies for different quality categories in the L1_SingleMu14 trigger algorithm plotted against the reconstructed pseudo-rapidity. The data used for this analysis was obtained using a W/Z skim.
It can be seen that the efficiency for large pseudo-rapidities is significantly smaller than towards the center of the detector. This is to be expected as measurements in this region are significantly more difficult and also less valuable to physics analysis.

It is not significantly affected by the current cuts as can be seen when comparing figures 8.2a and 8.2b. However, it seems there might be room for improving the current efficiency without increasing the rates significantly.

When examining promising combinations of quality and pseudo-rapidity to be recovered through changes to the current trigger cuts, it is convenient to examine the relative trigger efficiencies (figure 8.3).

This improves the ability to find combinations of qualities and pseudo-rapidities which could contribute a significant amount to the GMT's efficiency while not increasing the rates substantially.

As can then be seen in figures 8.3a and 8.3b the following quality codes could substantially increase the GMT's efficiency.

- $q_{CSC} = 1$ at $1 \leq |\eta| \leq 2$
- $q_{CSC} = 2$ at $|\eta| = 1.1$
- $q_{RPC} = 0$ at $1 \leq |\eta| \leq 1.1$
- $q_{RPC} = 1$ at $0.8 \leq |\eta| \leq 0.9$

Comparing these results with the rates shown in figure 8.1 it can be seen that removing the suppression for $q_{CSC} = 1$ at high pseudo-rapidities would lead to a significant increase in L1A rates. Removing the suppression on muons found using the RPC sub-system would however only marginally increase the L1A rate.

As a result of these studies several proposals were made for a change in GMT quality suppression.

Formerly some combinations of RPC qualities and η were suppressed in both single and double muon triggers. It was proposed to remove all suppression as the RPC patterns have changed since this decision was made and it seems as though the L1A rates would not be increased substantially.

Similarly, unconfirmed CSC candidates were originally suppressed from single muon triggers in the following cases:

- $q_{CSC} = 1$ in $|\eta| > 1.2$
- $q_{CSC} = 2$ in $|\eta| > 1.8$ and $|\eta| < 1.5$

The above configuration was proposed to be changed to:

- $|\eta| > 1.3$ and $q_{CSC} = 1$
- $|\eta| > 1.8$ and $1.2 < |\eta| < 1.5$ with $q_{CSC} = 1$

After various cross-checks, the proposed GMT quality assignment configuration has been deployed in Spring 2011 and has been used since then, which means for the large majority of data collected by CMS.

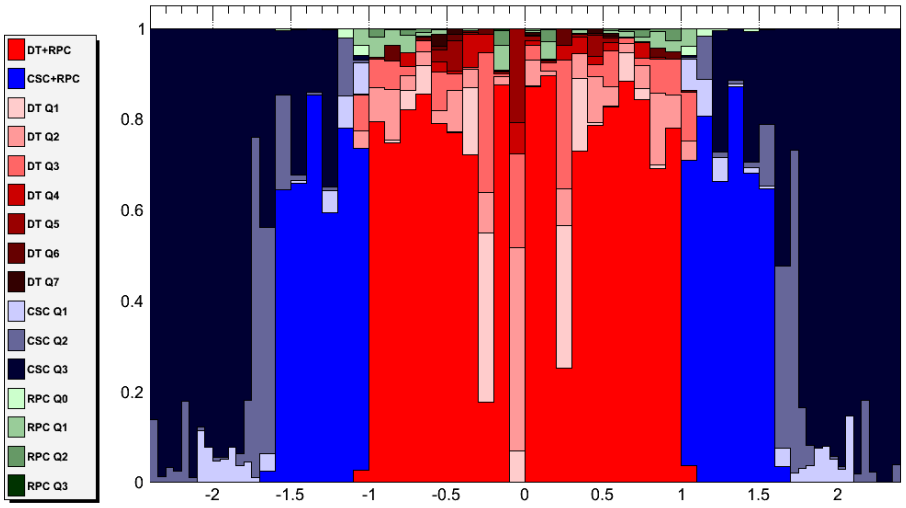

(a) All qualities

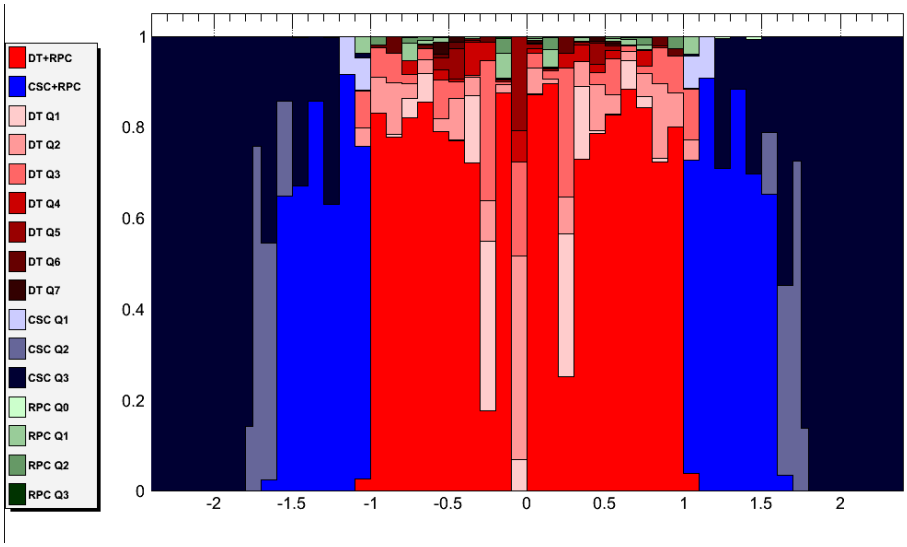

(b) After current quality cuts

Figure 8.3: Relative efficiency contributions for the trigger algorithm L1_SingleMu14 plotted against the pseudo-rapidity as determined by the GMT, split into different quality categories. The data used for this analysis was obtained using a W/Z skim.
Using these plots it can be seen that an increase in efficiency could be obtained by further tweaking the current suppression algorithms for unconfirmed muons, especially for CSC qualities 1 and 2 as well as for all RPC quality categories.

Acknowledgments

I would like to thank my supervisor Claudia-Elisabeth Wulz for enabling me to do this thesis and for giving me the opportunity to come to CERN before and during my thesis. The amazing atmosphere at this laboratory provided me with motivation and inspiration for my work. I am also thankful for lots of corrections and pointers given during the completion of this thesis.

Ivan Mikulec helped me come to grips with the Compact Muon Solenoid (CMS) analysis framework as I was working on the study of the Global Muon Trigger (GMT) efficiency. His patient explanations made the work on this analysis a very rewarding experience. Furthermore I would like to thank him for pointers given regarding the chapter of this thesis concerned with the GMT efficiency.

Babak Rahbaran provided me with lots of advice and guidance while working on the Global Trigger (GT) upgrade as well as offered lots of valuable input regarding the first drafts of this thesis. He also proved to be a driving force behind the completion of the thesis for which I am thankful.

Similarly Bernhard Arnold supported me with helpful tips and code as well as previous experience while working on the hardware abstraction layer (HAL).

I would also like to thank Manfred Jeitler for giving me the initial push to come to CERN and for lots of interesting discussions before and during my thesis as well as Vasile Mihai Ghete for thorough explanations about the CMS offline software framework and good advice as my thesis progressed.

Finally I would like to thank my family for all the support I received during my studies as well as their ongoing patience while I get carried away explaining my work.

Bibliography

[1] V. Khachatryan et al. "Observation of long-range, near-side angular correlations in proton-proton collisions at the LHC". In: *Journal of High Energy Physics* 2010 (9 2010). 10.1007/JHEP09(2010)091, pp. 1–38. ISSN: 1029-8479. URL: http://dx.doi.org/10.1007/JHEP09(2010)091.

[2] Serguei Chatrchyan et al. "Measurement of the inclusive production cross sections for forward jets and for dijet events with one forward and one central jet in pp collisions at $\sqrt{s} = 7\,\text{TeV}$. oai:cds.cern.ch:1421821". In: (Feb. 2012). Comments: Submitted to the Journal of High Energy Physics.

[3] Serguei Chatrchyan et al. "Search for $B_s^0 \to \mu^+\mu^-$ and $B^0 \to \mu^+\mu^-$ Decays in pp Collisions at $\sqrt{s}$=7 TeV". In: *Phys. Rev. Lett.* (19 Nov. 2011), p. 191802. DOI: 10.1103/PhysRevLett.107.191802. URL: http://link.aps.org/doi/10.1103/PhysRevLett.107.191802.

[4] Sidney D. Drell and Tung-Mow Yan. "Massive Lepton-Pair Production in Hadron-Hadron Collisions at High Energies". In: *Phys. Rev. Lett.* 25 (5 Aug. 1970), pp. 316–320. DOI: 10.1103/PhysRevLett.25.316. URL: http://link.aps.org/doi/10.1103/PhysRevLett.25.316.

[5] S. Chatrchyan et al. "Measurement of the inclusive W and Z production cross sections in pp collisions at $\sqrt{s} = 7\,\text{TeV}$ with the CMS experiment". In: *Journal of High Energy Physics* 2011 (10 2011). 10.1007/JHEP10(2011)132, pp. 1–76. ISSN: 1029-8479. URL: http://dx.doi.org/10.1007/JHEP10(2011)132.

[6] S. Chatrchyan et al. "Measurement of the weak mixing angle with the Drell-Yan process in proton-proton collisions at the LHC". In: *Phys. Rev. D* 84 (11 Dec. 2011), p. 112002. DOI: 10.1103/PhysRevD.84.112002. URL: http://link.aps.org/doi/10.1103/PhysRevD.84.112002.

[7] S. Chatrchyan et al. "Measurement of the $t\bar{t}$ production cross section and the top quark mass in the dilepton channel in pp collisions at $\sqrt{s} = 7\,\text{TeV}$". In: *Journal of High Energy Physics* 2011 (7 2011).

10.1007/JHEP07(2011)049, pp. 1–49. ISSN: 1029-8479. URL: http://dx.doi.org/10.1007/JHEP07(2011)049.

[8] S. Chatrchyan et al. "Measurement of the $t\bar{t}$ production cross section in pp collisions at 7 TeV in lepton+jets events using b-quark jet identification". In: *Phys. Rev. D* 84 (9 Nov. 2011), p. 092004. DOI: 10.1103/PhysRevD.84.092004. URL: http://link.aps.org/doi/10.1103/PhysRevD.84.092004.

[9] S. Chatrchyan et al. "Measurement of the t-Channel Single Top Quark Production Cross Section in pp Collisions at $\sqrt{s} = 7$ TeV". In: *Phys. Rev. Lett.* 107 (9 Aug. 2011), p. 091802. DOI: 10.1103/PhysRevLett.107.091802. URL: http://link.aps.org/doi/10.1103/PhysRevLett.107.091802.

[10] Serguei Chatrchyan et al. "Combined results of searches for the standard model Higgs boson in pp collisions at $\sqrt{s} = 7$ TeV. oai:cds.cern.ch:1422614". In: (Feb. 2012). Comments: Submitted to Physics Letters B.

[11] Serguei Chatrchyan et al. "Search for physics beyond the standard model using multilepton signatures in collisions at $\sqrt{s} = 7$ TeV". In: *Physics Letters B* 704.5 (2011), pp. 411 –433. ISSN: 0370-2693. DOI: 10.1016/j.physletb.2011.09.047. URL: http://www.sciencedirect.com/science/article/pii/S0370269311011208.

[12] Serguei Chatrchyan et al. "Search for Supersymmetry at the LHC in Events with Jets and Missing Transverse Energy". In: *Phys. Rev. Lett.* 107 (22 Nov. 2011), p. 221804. DOI: 10.1103/PhysRevLett.107.221804. URL: http://link.aps.org/doi/10.1103/PhysRevLett.107.221804.

[13] Serguei Chatrchyan et al. "Search for supersymmetry in pp collisions at $\sqrt{s} = 7$ TeV in events with a single lepton, jets, and missing transverse momentum". In: *Journal of High Energy Physics* 2011 (8 2011). 10.1007/JHEP08(2011)156, pp. 1–43. ISSN: 1029-8479. URL: http://dx.doi.org/10.1007/JHEP08(2011)156.

[14] Serguei Chatrchyan et al. "Search for large extra dimensions in dimuon and dielectron events in pp collisions at $\sqrt{s} = 7$ TeV". In: (2012). Comments: Submitted to Physics Letters B. arXiv:1202.3827 [hep-ex].

[15] Serguei Chatrchyan et al. "Search for microscopic black holes in pp collisions at sqrt(s) = 7 TeV". In: (2012). Comments: Submitted to the Journal of High Energy Physics. arXiv:1202.6396 [hep-ex].

[16] "Search for quark compositeness in dijet angular distributions from pp collisions at sqrt(s) = 7 TeV". In: (2012). Comments: Submitted to the Journal of High Energy Physics. arXiv:1202.5535 [hep-ex].

[17] JM Jowett et al. *First run of the LHC as a heavy-ion collider. oai:cds.cern.ch:1382049.* Tech. rep. CERN-ATS-2011-143. Geneva: CERN, Sept. 2011.

[18] Amy Dusto. *Accelerator Soup: recipe for heavy-ion collisions in proton-lead test.* URL: `http://alicematters.web.cern.ch/acceleratorsoup` (visited on Mar. 4, 2012).

[19] *Energy Record gives LEP new Discovery Possibilities.* URL: `http://press.web.cern.ch/press/pressreleases/Releases1999/PR09.99ELEP200GEV.html` (visited on June 17, 2012).

[20] Wolfgang K. H. Panofsky. "The Evolution of Particle Accelerators & Colliders". In: *Beamline* (1997).

[21] LHC Programme Coordination. *LHC Luminosity Plots for the 2011 Proton Run.* URL: `http://lpc.web.cern.ch/lpc/lumiplots.htm` (visited on Mar. 4, 2012).

[22] Anna Phan. *LUMI LEVELING: What, Why and How?* URL: `http://www.quantumdiaries.org/2011/04/24/lumi-leveling-what-why-and-how/` (visited on Mar. 5, 2012).

[23] O Adriani et al. *LHCf experiment: Technical Design Report.* Technical Design Report LHCf. Geneva: CERN, 2006.

[24] V Berardi et al. *Total cross-section, elastic scattering and diffraction dissociation at the Large Hadron Collider at CERN: TOTEM Technical Design Report.* Technical Design Report TOTEM. Geneva: CERN, 2004.

[25] P R Barbosa-Marinho et al. *LHCb VELO (VErtex LOcator): Technical Design Report.* Technical Design Report LHCb. Geneva: CERN, 2001.

[26] James Pinfold et al. *Technical Design Report of the MoEDAL Experiment.* Tech. rep. CERN-LHCC-2009-006. MoEDAL-TDR-001. Geneva: CERN, June 2009.

[27] CMS Collaboration. "Commissioning and performance of the CMS silicon strip tracker with cosmic ray muons". In: *Journal of Instrumentation* 5.03 (2010), T03008. URL: `http://stacks.iop.org/1748-0221/5/i=03/a=T03008`.

[28] CMS Collaboration. "The CMS experiment at the CERN LHC. The Compact Muon Solenoid experiment". In: *J. Instrum.* 3 (2008). Also published by CERN Geneva in 2010, S08004. 361 p.

[29] CMS Collaboration. *CMS TriDAS project: Technical Design Report; 1, the trigger systems.* Technical Design Report CMS.

[30] Pamela Renee. *Operation and Performance of the CMS Level-1 Trigger during 7 TeV Collisions.* Tech. rep. CMS-CR-2011-183. Geneva: CERN, Sept. 2011.

[31] Hannes Sakulin and Massimiliano Fierro. *Studies of the Global Muon Trigger Performance.* Tech. rep. CMS-NOTE-2001-003. Geneva: CERN, Jan. 2001.

[32] I Magrans de Arbril, C-E Wulz, and J Varela. "Conceptual design of the CMS trigger supervisor". In: *IEEE Trans. Nucl. Sci.* 53 (2006), pp. 474–483.

[33] B Rahbaran et al. "MicroTCA-based Global Trigger Upgrade project for the CMS experiment at LHC". In: *Journal of Instrumentation* 6.12 (2011), p. C12054. URL: http://stacks.iop.org/1748-0221/6/i=12/a=C12054.

[34] Paul Stevens. *An Integrators Guide to MicroTCA Part I – Platform management overview*. Tech. rep. Advantech Inc. URL: http://buy.advantech.com/resource/Whitepaper/An_Integrator's_Guide_to_MicroTCA_Part%20I_Platform_management_overview.pdf (visited on Apr. 13, 2012).

[35] Peter Marek. *An Integrators Guide to MicroTCA Part III – System management software*. Tech. rep. Advantech Inc. URL: http://buy.advantech.com/resource/Whitepaper/An_Integrator's_Guide_to_MicroTCA_PartIII_System_mngmt_sw.pdf (visited on Apr. 13, 2012).

[36] Vollrath Dirksen. "MicroTCA-FAQ". In: *Praxis Profiline – MicroTCA* (Aug. 2007). URL: http://www.embedded4you.com/index.php?option=com_docman&task=doc_view&gid=5&tmpl=component&format=raw&Itemid=25 (visited on Apr. 14, 2012).

[37] H. Zimmermann. "OSI Reference Model–The ISO Model of Architecture for Open Systems Interconnection". In: *Communications, IEEE Transactions on* 28.4 (1980), pp. 425–432. URL: http://ieeexplore.ieee.org/xpls/abs_all.jsp?arnumber=1094702 (visited on June 16, 2012).

[38] *Protocol Buffers – Google's data interchange format*. URL: http://code.google.com/p/protobuf/ (visited on Apr. 20, 2012).

[39] Hannes Sakulin. "Design and Simulation of the First Level Global Muon Trigger for the CMS Experiment at CERN. oai:cds.cern.ch:1324143". Presented on 2002. PhD thesis. Vienna U.: Vienna U., 2002.

[40] CMS Collaboration. *CMS Luminosity Collision Data 2010*. URL: https://twiki.cern.ch/twiki/bin/view/CMSPublic/LumiPublicResults2010 (visited on Feb. 27, 2012).

Glossary

ALICE
A Large Ion Collider Experiment. 14

AMC
Advanced Mezzanine Cards. 46, 48, 50–52, 54–56, 58, 59, 61, 63–65, 70, 72, 91

application-specific integrated circuit
an integrated circuit customized during manufacturing for a particular use, as opposed to general-purpose usage. 29, 87

ASIC
application-specific integrated circuit. 29, 45

ATCA
AdvancedTelecommunications Computing Architecture (TCA). 45, 46

ATLAS
A Toroidal Large Hadron Collider (LHC) Apparatus. 14, 15

AU
Assignment Unit. 36, 37

BTI
Bunch and Track Identifier. 34

Castor
Centauro and Strange Object Research. 22, 23

CLCT
Cathode Local Charged Tracks (LCT). 35

CMS
Compact Muon Solenoid. i, iii, 1, 2, 15, 17, 18, 20–25, 31, 39, 41, 42, 48, 49, 51, 72, 73, 79, 81

CNGS
CERN Neutrinos to Gran Sasso. 11

COTS
commercial of the shelf. 45

CSC
cathode strip chamber. 17, 24, 25, 31–33, 35–38, 73, 76, 79, 88

CSCTF
cathode strip chamber (CSC) Track Finder. 36, 37, 76

DAQ
data acquisition. 41

DQM
data-quality monitoring. 25

DT
drift tube. 17, 24, 25, 31–33, 35–38, 73, 88

DTTF
drift tube (DT) Track Finder. 36, 37

ECAL
electromagnetic calorimeter. 20, 21, 30

EMMC
Extended Module Management Controllers (MMC). 46

EU
Extrapolation Unit. 36, 37

EVM
Event Manager. 41

FDL
Final Decision Logic. 41

field-programmable gate array
an integrated circuit that can be configured by the user after manufacturing. 29, 88

FINOR
final OR. 39, 41

FMC
field-programmable gate array (FPGA) Mezzanine Card. 48

FPGA
field-programmable gate array. 29, 40, 41, 45, 48, 50–52, 55, 62, 72, 88

FRU
field-replaceable unit. 46, 48, 89

GCT
Global Calorimeter Trigger. 31, 40, 49

GMT
Global Muon Trigger. i, iii, 2, 32, 33, 36–41, 49, 73–76, 79, 81

GT
Global Trigger. i, iii, 1, 33, 39–42, 49, 51, 52, 55, 56, 58, 61, 62, 64, 67, 69, 72, 73, 76, 81

GTFE
Global Trigger Front-End. 41

GTL
Global Trigger Logic. 41, 42

HAL
hardware abstraction layer. i, iii, 1, 50–52, 55, 56, 61, 62, 64, 69, 72, 81

HB
barrel part. 21

HCAL
hadronic calorimeter. 20–22

HE
endcap calorimeter. 22

HF
forward hadron calorimeter. 22, 23, 31

HLT
High-Level trigger. 25, 39, 42

HO
outer barrel part. 21

hot-swappable
possibility to exchange a field-replaceable unit (FRU) while the system is switched on. 46

Intelligent Platform Management Interface
an interface for maintenance and management of computer systems. 46, 89, 90

IPMI
Intelligent Platform Management Interface. 46, 51, 70, 90

IPMI-0
Intelligent Platform Management Interface (IPMI) Channel 0. 46

IPMI-L
IPMI local. 46

L1A
Level-1 Accept. 41, 42, 74, 79

L1T
Level-1 trigger. 25

LCT
Local Charged Tracks. 35–37, 87

LEIR
Low Energy Ion Ring. 11

LHC
Large Hadron Collider. i, iii, 1, 11, 12, 14–18, 23, 25, 27, 40, 45, 87, 90, 92

LHCb
LHC beauty. 15, 16

LHCf
LHC forward. 15

MCH
MicroTCA Carrier Hub. 46, 48, 49, 52, 70, 90

MCMC
MicroTCA Carrier Hub (MCH) Carrier Management Controller. 46

MicroTCA
Micro TCA. 45

MMC
Module Management Controllers. 46, 48, 88

MoEDAL
Monopole and Exotics Detector At the LHC. 16

MPC
Muon Port Card. 35–37

NLO
next-to-leading-order. 5

NTD
Nuclear Track Detector. 16

OMDS
Online Master Database system. 42

PACT
Pattern Comparator Trigger. 32, 33, 38

PPC
PowerPC. 52, 58, 70

pQCD
perturbative quantum chromodynamics (QCD). 5

PrAMC
Processor Advanced Mezzanine Cards (AMC). 70

PS
Proton Synchrotron. 11

PSB
Pipeline Synchronized Buffer. 40, 41

QCD
quantum chromodynamics. 3–5, 9, 91

RCT
Regional Calorimeter Trigger. 31, 39

RPC
resistive plate chamber. 17, 24, 25, 31–33, 37, 38, 73, 79

SL
superlayer. 24, 34

SMP
Stable Massive Particle. 16, 17

SP
Sector Processor. 36, 37

SPS
Super Proton Synchrotron. 11

SUSY
supersymmetry. 8, 11, 17

TA
Track Assembler. 36, 37

TCA
Telecommunications Computing Architecture. 45, 87, 90

TCS
Trigger Control System. 41, 42

TIM
Timing Module. 40

TMB
trigger motherboard. 35, 36

TOTEM
TOTal cross section, Elastic scattering and diffraction dissociation Measurement at the LHC. 15, 22

TPC
Time Projection Chamber. 14

TPG
Trigger Primitives Generator. 30, 31

TRACO
Track Correlator. 34

TS
Trigger Server. 34

TTC
Trigger Timing and Control. 40

VELO
Vertex Locator. 15, 16

Dinyar S. Rabady

Deinhardsteingasse 3/14
1160 Vienna
☎ *+43 676 9185556*
✉ *dinyar.rabady@cern.ch*

Personal data

Date of birth **November 30th, 1986**.
Place of birth **Waidhofen an der Thaya, Austria**.
Citizenship **Austria**.

Educational history

Feb. 9th – Feb. 16th 2011 **International School of Trigger and Data Acquisition 2011**, *Rome, Italy*.
Lectures and practical sessions on Trigger, DAQ, NIM, VMEbus, PCI, FPGAs, Microprocessors, etc.

2010 – July 2012 **Institute of High Energy Physics of the Austrian Academy of Sciences**, *Diploma thesis*.
Thesis: "A hardware abstraction layer for the MicroTCA-based Global Trigger for the CMS experiment at CERN"

July 5th – Sept. 6th 2010 **"HEPHY@CERN" Summer intern**.
Work on the software of the Level-1 trigger control system of the CMS experiment, shifts for the Level-1 trigger system. Attended the lectures offered by CERN

2006 – 2012 **University of Vienna**, *Study of physics*.

1997 – 2005 **Secondary school in Waidhofen an der Thaya**.
School leaving exam: "Feynmandiagrams in Quantumelectrodynamics"

1993 – 1997 **Primary school in Schwarzenau**.

Work experience

March 2012 – **Software developer**, *Catalysts GmbH*.
Work on a fraud detection and protection mailrelay server.

October 2008 – February 2011 **Student tutor in the group Aerosol Physics and Environmental Physics**, *University of Vienna*.
Student tutor in the physics laboratory exercises for biologists.

Talks and posters

Sept. 26th – Sept. 30th 2011 **MicroTCA-based Global Trigger Upgrade project for the CMS experiment at LHC**, *TWEPP '11 in Vienna, Austria*, Poster with Babak Rahbaran, Herbert Bergauer, Bernhard Arnold, and Markus Eichberger.
Overview of HEPHY's electronic group's current plans for the Global Trigger upgrade using μTCA technology.

Sept. 11th – Sept. 15th 2011 **CMS Global Trigger for an upgraded LHC**, *CMS Physics Week in Brussels*, Poster with Manfred Jeitler, Babak Rahbaran, and Claudia-Elisabeth Wulz.
Overview of current plans for the upgrade of the Global Trigger for the CMS experiment to the μTCA standard.

June 17th 2011 **The upgrade of the CMS Trigger Electronics for SLHC**, *Joint Annual Meeting of the Swiss Physical Society and Austrian Physical Society in Lausanne*, Contributed Talk on behalf of the CMS collaboration.
Overview of the current plans for the upgrade of especially the Global Trigger of the CMS experiment.

Conferences attended

Sept. 26th – Sept. 30th 2011 **TWEPP '11**, *Topical Workshop on Electronics in High Energy Physics*, Vienna, Austria.
Member of the local organising team

June 13th – June 17th 2011 **SPS/ÖPG Annual Meeting '11**, *Joint Annual Meeting of the Swiss Physical Society and Austrian Physical Society*, Lausanne, Switzerland.

Languages

German	**Mother tongue**	
English	**Second language**	*Spoken with father and grandmother since birth.*
French	**Basics**	*From 5. to 8. class in secondary school*