TECHNISCHE
UNIVERSITÄT
WIEN

VIENNA
UNIVERSITY OF
TECHNOLOGY

DIPLOMARBEIT

An Integrated Testing Facility for the Global Trigger of the CMS Experiment at CERN

AUSGEFÜHRT AM

INSTITUT FÜR HOCHENERGIEPHYSIK

DER ÖSTERREICHISCHEN AKADEMIE DER WISSENSCHAFTEN

UND DER BETREUUNG VON

DIPL.-ING. DR. VASILE MIHAI GHETE

UNTER DER ANLEITUNG VON

UNIV.DOZ. DIPL.-ING. DR.TECHN. CLAUDIA-ELISABETH WULZ

INSTITUT FÜR HOCHENERGIEPHYSIK

DURCH

THOMAS THEMEL
MATRIKELNUMMER: 0525052
RUBENSGASSE 9/12 - 1040 WIEN

WIEN, MAI 2010   UNTERSCHRIFT DES STUDENTEN   UNTERSCHRIFT DER BETREUERIN

The Global Trigger is part of the Level-1 Trigger of the CMS experiment at CERN, with the task to find the most interesting events corresponding to a rate of 100 kHz from the basic Large Hadron Collider interaction rate of 40 MHz. It is expected to render a decision within $3.2\,\mu s$, which necessitates an implementation using custom hardware. The implementation makes heavy use of Field Programmable Gate Array (FPGA) technology to reconcile the performance requirements with the need for flexibility.

The complexity of the Global Trigger system (13 boards with 51 FPGA chips) makes it vulnerable to a multitude of errors, from electrical errors such as bad solder joints or plug contacts up to logical errors in the implementation of the firmware and the configuration software. The goal of the work described in this thesis was to provide an integrated system that allows users to easily determine whether the system is working correctly and assists experts in tracking down the internal causes of such errors within the context of the Level-1 Trigger Online Control infrastructure.

The operation of the testing system was validated on the basis of errors uncovered in the physics data collected during data taking on cosmic radiation in July and August 2009. The results demonstrate the importance of an easily available testing infrastructure and its integration into the experiment's operating procedure.

Contents

1 Introduction

While the Standard Model of elementary particle physics has withstood all experimental tests since its inception, there is ample reason to believe that it is an effective theory that arises as a special case of a more fundamental description of nature. It is theoretically unwieldy, with a large number of free parameters. Glaringly, it fails to provide an explanation for gravitation, and cosmological observations suggest that large parts of the energy of the universe are not accounted for by any form of matter known in the Standard Model. While many theories exist to extend the Standard Model in theoretically more appealing directions, the final arbiter of their value remains the comparison with experimental data. To that end, the Large Hadron Collider (LHC), currently being commissioned at CERN, is designed to explore the landscape of particle physics at energies an order of magnitude higher than previous machines. Chapter 2 reviews some of the expected features of this landscape from an experimental point of view, both for the Standard Model and some candidate theories beyond the Standard Model. The LHC machine and the Compact Muon Solenoid (CMS) experiment, one of its two large general purpose detectors, form the experimental context of this work. They are described in Ch. 3.

A major technological challenge in the analysis of particle collisions is event selection. As early as 1955, the ability of experiments to generate events outstripped the ability of humans to analyze them, causing Luis Alvarez to complain that "one day of bubble chamber operation could keep a group of 'cloud chamber analysts' busy for a year" [1]. This launched particle physics experiments on a trajectory of ever greater automation and computerization, culminating in the current situation at LHC. When the design luminosity of $\mathcal{L} = 1.0 \times 10^{34}\,\mathrm{cm}^{-2}\mathrm{s}^{-1}$ is reached, the experiments will observe an event (usually containing multiple proton-proton interactions) for every 25 ns interval. With the measurement capabilities of the CMS experiment, this amounts to a data volume of about 40 TB/s, far beyond any storage capabilities currently feasible. In addition, much of this data deluge consists of well-known processes and is thus of little interest for the physics program. Since actual storage capability is limited to about 100 events per second, a rejection factor of almost 10^{-6} has to be achieved in the event selection.

The CMS experiment's approach to this problem relies on two stages. First, all detector data is preserved in buffer memories while a dedicated electronics system makes a preliminary decision to accept or reject the event, based on coarse measurements. This decision has to be rendered within 3.2 µs due to limited buffer space and is expected to reduce the event rate by a factor of about 10^{-3}. This system is the CMS Level-1 Trigger, described in detail in Ch. 4. In a second stage, the events accepted by the Level-1 Trigger are examined by the High-Level Trigger (HLT), implemented in software running on general purpose CPUs, which achieves another 10^{-3} rejection

factor. The selected events are then sent to permanent storage by the Data Acquisition (DAQ) system.

The strict time constraints faced by the Level-1 Trigger exclude an implementation based on general-purpose computers. Instead, the system is based on custom-built electronics, utilizing programmable logic in the form of FPGAs where possible. Since these designs are specific to the CMS experiments, they are produced in small quantities and thus cannot benefit from the economies of scale that would make the complex quality assurance of industrial production feasible. Combined with the large complexity, challenging performance requirements and sometimes harsh environment, these circumstances make the trigger electronics prone to defects which could degrade its function. Since the trigger is a crucial ingredient of the experiment, a tool to check the correctness of the trigger during data taking was necessary. Chapter 5 details such a tool, the major new contribution of this work.

Finally, Ch. 6 discusses some application examples and illustrates the value of the work in the context of the experiment.

1.1 Acknowledgments

2 Physics at the TeV scale

The LHC physics program aims to explore physics at energies about an order of magnitude higher than those achieved in previous accelerators. Proton-proton collisions at the LHC will attain a nominal center-of-mass energy of 14 TeV, compared to the previous record of 1.96 TeV in proton-antiproton collisions at the Tevatron. The following sections review some of the expected physics processes from Standard Model physics to possible new physics beyond the Standard Model.

2.1 Standard Model Physics

2.1.1 QCD and Jet Studies

Since the LHC is a hadron collider, it provides a very rich environment for the study of QCD phenomena. In the dominant hard scattering mode the incoming protons dissociate and produce at least two primary jets from the partons involved in the event. The expected dijet invariant mass distribution[1] is well understood theoretically and has been measured over a wide range of energies and beam intensities at previous hadron colliders. Differences from the expected spectrum (see Fig. 2.1) would be an indicator for possible new physics beyond the Standard Model. The final results should produce an improved understanding of the strong coupling α_S and the parton distribution functions.

2.1.2 Underlying Event Studies

The "underlying event" consists of all parts of a hard scattering event besides those triggered by the primary parton interactions. Typically, a hard scattering event leads to two high-p_T jets from the scattered partons and two jets travelling roughly in the direction of the beam (the remnants of the original hadrons). Effects from the remnants and from multiple parton scattering influence the development of jets and need to be corrected for when matching jet observables to hard scattering events. Different theoretical models exist and substantially disagree on a number of observables, such as the total cross section, charged particle multiplicity or average p_T contributions from the underlying event.

Reference measurements are needed to distinguish between these models. The basic concept exploits the described event topology to define a transverse region relative to the jets from hard scattering. This should be relatively free of products from the hard

---

[1]The invariant mass of a dijet event is defined as the invariant mass of the two highest-energy jets in the event, calculated as $m = \sqrt{(E_1 + E_2)^2 - (p_1 + p_2)^2}$.

Figure 2.1: Expected dijet mass spectrum from pp collisions at $\sqrt{s} = 14\,\text{TeV}$ [2].

scattering process and thus dominated by the underlying event. A similar observation is possible using muon pairs generated through the Drell-Yan process, where a cleaner signal is expected for the price of a much-reduced cross section.

2.1.3 Inclusive b Quark Production

At LHC design parameters, b quarks will be produced at a rate of about 10^6 per second, providing an improved understanding through better statistics and extending the verified reach to higher energies. One of the interesting observables is the distribution of the b quark cross section over the p_T of the produced B hadrons. While the observed shape of this distribution agrees with theoretical predictions, multiple experiments have consistently seen an overall larger cross section than that predicted from perturbative QCD calculations.

Analysis of b quark decays requires an experimental distinction between jets caused by b quark decays and those from light quarks or gluons, called b-tagging. This identification relies on several features of the b decays, primarily the comparatively high lifetime of B hadrons of about 1.5 ps ($c\tau \approx 450\,\mu\text{m}$). This allows a high-resolution tracking detector to assign the tracks originating from this decay to a common secondary vertex. In addition, several other features of B hadron decays are distinctive, such as the higher number of charged particles and higher invariant mass as compared to charm hadron decays. Combining all these variables into a single discriminator allows one to flexibly trade b-tagging efficiency for suppression of lighter decays and achieves predicted suppression rates of two orders of magnitude at an efficiency on the order of of 60%.

2.1.4 B Hadrons and CP Violation

While B hadrons have been studied before, mainly at dedicated B physics experiments such as BaBar or Belle, the LHC environment will provide new opportunities for the

study of CP violation parameters, mainly because of the very high b$\bar{\text{b}}$ production cross section. Studies of J/Ψ, B$_s$, B$_c$ and quarkonia states to measure masses, cross-sections and branching ratios are envisaged by the general purpose experiments ATLAS and CMS. However, the experimental requirements of a B physics program differ from those of the general physics program sufficiently to make a dedicated experiment (LHCb, [3]) useful. LHCb is more suited to a B physics program because it includes a silicon strip detector located very close to the interaction point, Cherenkov detectors for particle identification and a beam defocusing system to intentionally lower the luminosity and avoid pileup.

One area of interest to the general purpose experiments is the study of rare decays like B$_s \to \mu\mu$, which has a Standard Model branching ratio of only 3.85×10^{-9}, but could be strongly enhanced by contributions from new physics.

2.1.5 Diffractive and Forward Physics

Diffractive scattering refers to interactions where the involved hadrons interact through an object with vacuum quantum numbers, i.e. the outgoing states have the same quantum numbers as the incoming ones. The naming of this phenomenon derives from the shape of the elastic proton-proton cross section with respect to momentum transfer, which shows diffraction-like dips for high energies (see Fig. 2.2). At LHC energies, diffractive scattering also includes a large non-elastic component where one or both of the incoming protons dissociate. Due to the small energy loss of the incoming hadrons, the final states are concentrated in the forward regions not directly covered by the general purpose experiments. Specialized forward experiments [4,5] have been developed to instrument the forward regions, allowing the general purpose experiments to distinguish diffractive events through detection of the outgoing products. Expected early results include precise luminosity measurements and improvements of existing cosmic ray simulations (since the interactions of PeV-scale cosmic protons with the atmosphere have a similar center-of-mass energy as LHC collisions).

2.1.6 Heavy Ions

Lattice QCD calculations predict a new state of matter at energy densities $> 1\,\text{GeV/fm}^3$, the quark-gluon plasma (QGP). Colliding lead ions at per-nucleon energies similar to those for protons could produce these states in a comparatively large volume of about 100 fm^3. At high energies, hadrons consist mostly of a dense system of gluons containing very small fractional momenta. Nonlinear gluon-gluon interactions dominate the evolution of the system. Numerous interesting connections to other areas of physics such as early-universe cosmology, the modelling of high-energy cosmic ray interactions or even the Anti-de-Sitter/Conformal Field Theory (AdS/CFT) duality, which predicts QCD observables at finite temperature from black hole thermodynamics.

An easily verifiable prediction is suppression of quarkonia production in QGP environments. This has been observed for the J/Ψ at RHIC and should be detectable for the heavier Υ at the LHC, principally by looking for suitable decays to muon pairs.

Figure 2.2: Elastic proton-proton cross section over momentum transfer measured at various energies.

2.1.7 Top Quark

The goal of top physics at the LHC is an improved measurement of all the properties of the top quark, such as its production cross section, mass and spin. Top quarks will predominantly be produced as $t\bar{t}$ pairs and decay almost exclusively as $t \to W + b$. Experimental signatures of top events differ through the different decay modes of the resulting W pair.

With a branching ratio of about 11%, the W pair will decay into a pair of leptons. Selecting events with lepton pairs and two b-jets provides a very clean channel where the main background comes from directly produced Z bosons decaying to two leptons. This background can be suppressed effectively by simply instituting a kinematic cut to remove lepton pairs whose invariant mass is consistent with a Z decay (see Fig. 2.3). Further selection criteria include the requirement of significant missing energy (from the $W \to \ell\bar{\nu}$ decay) and a lower limit for the transverse momentum of the leptons.

Similar requirements apply to the semileptonic and fully leptonic decay modes (with a branching ration of about 45% each), adapted to the different particle content. Semileptonic events have four jets of which two are b-tagged and an isolated lepton while fully hadronic decays will have at least six jets, two of which are b-tagged. Since the influence of QCD backgrounds worsens with increasing number of jets, analysis methods for these signatures have to be more involved, using neural networks to determine kinematic cuts from simulation results.

6

Figure 2.3: Left: Dilepton mass distribution in top pair decays and indication of kinematic regions that are cut to suppress the Z background. Right: final top mass distribution after event selection and kinematic fit.

Mass Measurement

Measurement of the top mass from dileptonic decays suffers from the principal problem that part of the original energy is lost through the neutrinos in the decay. This can be compensated for by combinatorial effects, requiring the mass of both top quarks to be equal and imposing m_W constraints. The known Standard Model neutrino spectrum further reduces the uncertainty. Fitting the kinematics of each event to the most likely top mass produces the distribution shown in Fig. 2.3.

The semileptonic channel is considered the "golden channel" for mass measurement since it captures the entire energy from the decay of one top quark in three jets, one of which must be b tagged. Measuring the top mass in the fully hadronic channel once again requires a more complex analysis.

Spin Correlations in Top Quark Pairs

The large width of the top quark ($\Gamma \approx 1.4 \, \text{GeV}$, corresponding to a lifetime of about $5 \times 10^{-25} \, \text{s}$) makes the effects of both hadronization and spin flips from gluon emission negligible. This means that spin correlations between the top quarks in a pair should survive until the decay and subsequently be detectable in the angular distribution of decay products.

2.1.8 Electroweak Physics

Production of W and Z bosons are very well understood theoretically. The cross sections for central $W \to \ell\nu$ and $Z \to \ell\ell$ processes at the LHC design energy are $10 \, \text{nb}$ and $1 \, \text{nb}$ respectively, producing about 20 and 2 such decays per second at a luminosity of

7

Figure 2.4: Transverse mass distribution of W $\to \mu\nu$ events and background.

$2 \times 10^{33}\,\mathrm{cm}^{-2}\mathrm{s}^{-1}$. This combination will make these processes very useful during early running for purposes such as precise luminosity monitoring and detector calibration.

Z $\to \ell\ell$ decays can be identified by looking for events where the most energetic dilepton pair has an invariant mass consistent with the Z mass, while the identification of W $\to \ell\nu$ decays is more demanding due to the missing energy involved. A useful discriminatory variable is the transverse mass $m_T = \sqrt{2 p_T^\ell p_T^\nu (1 - \cos(\Delta\phi))}$, which separates events from a W decay from QCD events. Here, p_T is the transverse momentum of the lepton and the neutrino, respectively, $\Delta\phi$ is the azimuthal separation between the lepton and the neutrino. In addition, simulations show that it is useful to compute the missing transverse energy from hard objects only, avoiding a bias introduced by low energy objects which are randomly distributed across the detector. Figure 2.4 illustrates the usefulness of the transverse mass as a discriminant.

Drell-Yan Production of Muon Pairs

The Drell-Yan process produces lepton pairs from the annihilation of a quark and an antiquark in a hadron collisions. In the Standard Model, this involves either a virtual photon or a virtual Z boson. At energies around the Z mass, the Z contribution dominates, and at higher energies the photon and Z contributions mix. The total cross section for this process can be measured by selecting events with a single muon pair. In addition, a forward-backward asymmetry is expected in this process. Due to the symmetric nature of the initial state in pp collisions, the directions of the incoming partons are not known. However, they can be recovered by making use of the fact that the antiquark involved in the process is necessarily a "sea quark" carrying lower fractional momentum so that the momentum of the original quark can be approximated by the momentum of the resulting dimuon pair.

8

Figure 2.5: Expected (left) and unexpected (right) triple gauge boson interactions in the Standard Model.

Measurement of the W Mass

A precision measurement of the W mass is interesting not only as a direct check of the Standard Model, but also because it is sensitive to corrections from new physics. While statistical errors will be negligible due to the high rates, systematic errors inherent in the hadron collision process make a direct measurement of the W mass at the LHC incompetitive. The most promising methods for determining the W mass comes from an approach that uses Z boson observables to calculate those of the W boson, thus trading statistical uncertainty from the lower frequency of Z events for improved systematic error, as most of the experimental and theoretical uncertainties cancel in the comparison of W and Z distributions.

Multiboson Production

In the Standard Model, the neutral gauge bosons do not couple to each other, while they do couple to the charged W bosons (see Fig. 2.5). This should be reflected in the production rates of gauge boson pairs. The cross section for WZ production at the LHC is expected to be about 50 pb and that for ZZ production about 20 pb. Besides their direct relevance in the search for new physics, the ZZ states are also a irreducible background in searches for the Standard Model Higgs boson.

The fully leptonic decays of these states provide clean signatures through the requirement for opposite-sign lepton pairs that fit the Z mass. Discarding events that also include highly energetic jets allows a further suppression of the hadronic background from $t\bar{t}$ and $Zb\bar{b}$ decays.

2.1.9 Standard Model Higgs Boson

One of the cornerstones of the Standard Model is the hitherto unobserved Higgs boson; through the Higgs mechanism (spontaneous symmetry breaking in the Higgs sector), the electroweak gauge bosons W and Z and the fermions acquire masses through interaction with the Higgs field. The search for this particle has been one of the key considerations in the design of the general-purpose LHC experiments. While the mass of the Higgs boson is a free parameter of the Standard Model and there has not yet been a conclusive observation of a candidate particle, there are several indirect constraints on its mass. Lower bounds were determined by the LEP experiments

Figure 2.6: Decay branching ratios (left) and production cross section (right) of a Standard Model Higgs boson at the LHC [7].

(indicating $m_H \geq 114.4\,\text{GeV}$). Tantalizingly, when trying to extract the value of the Higgs mass from its influence on precision electroweak measurements, the greatest likelihood is likewise found at $m_H = 114^{+69}_{-45}\,\text{GeV}$. Multiple theoretical considerations provide upper bounds of about $m_H < 700\,\text{GeV}$, making a discovery at the LHC appear likely. Recent results from the Tevatron experiments also seem to exclude a region around $m_H = 160$ GeV [6].

Higgs Boson Production

Figure 2.6 shows the dominant production cross sections for a Standard Model Higgs boson at the LHC.

Gluon fusion produces a Higgs boson from a QCD process involving a virtual heavy quark. This is a favored process because the coupling between the Higgs field and the fermionic fields of the quarks in the colliding protons is proportional to the mass of the quarks. As the colliding protons predominantly contain very light quarks, the direct production of a Higgs boson from their interaction is unlikely. However, approximately half of the colliding particles' momentum is actually contained in gluons, which can decay to quark pairs. With copious available energy (as in the LHC design case), this process also contributes for the heavy quarks with their stronger coupling to the Higgs field, leading to a strong contribution for Higgs production for the whole m_H range [8]. **Vector boson fusion** produces a Higgs boson from the interaction of two vector bosons as shown in Fig. 2.7. Its cross section is about an order of magnitude smaller than that of the gluon fusion process. **Associated production with vector bosons** occurs when two incoming quarks annihilate to produce a vector boson that is energetic enough to "radiate" a Higgs boson. **Associated production with heavy quarks** can principally be caused by incoming quarks or gluons, but is dominated by gluons at LHC energies due to the considerations mentioned above. The only final states with relevant cross sections are those that produce a $t\bar{t}$ or a $b\bar{b}$ pair in addition

Figure 2.7: Feynman diagrams for Higgs production processes at the LHC - (a) gluon fusion, (b) vector boson fusion, (c) associated production with heavy quarks, (d) associated production with vector bosons.

to the Higgs boson. While the $Hb\bar{b}$ cross-section is actually larger than the $Ht\bar{t}$ cross-section for most of the m_H range where these channels are relevant, the $Hb\bar{b}$ final state is hard to isolate experimentally and so is usually considered irrelevant for initial discovery.

Decay Channels and Experimental Signatures

Figure 2.6 shows the branching ratios for Standard Model Higgs boson decays. For low masses ($m_H < 135\,\mathrm{GeV}$), the dominant decay channels are into a pairs of b quarks or τ leptons, with a combined BR of about 93%. These channels will be unobservable because there is a huge background of similar final states from QCD processes. This leaves the $H \to \gamma\gamma$ process as the most promising channel in this mass region, even though its branching ratio is only on the order of 10^{-3}. Decays into pairs of W or Z bosons dominate intermediate to high mass ranges. Contributions from $t\bar{t}$ become significant for $m_H > 300\,\mathrm{GeV}$, but these suffer from similar detection problems as the $b\bar{b}$ states and will not be considered further. The region of $2M_W < m_H < 2M_Z$ must be considered separately because Z pair production is significantly less likely there than W pair production.

$H \to \gamma\gamma$ **Analysis**   A discovery through the $H \to \gamma\gamma$ channel is experimentally challenging because the backgrounds from direct photon production are much larger than the expected Higgs signal (by a factor of about 10^6). Nevertheless, at $m_H < 130\,\mathrm{GeV}$, it remains the most promising option because the vector boson decays are very rare.

Backgrounds come both from "irreducible" sources which produce two photons

Figure 2.8: Dilepton invariant mass (left) and missing transverse energy (right) distribution for signal events for $m_H = 185\,\mathrm{Gev}$ [2].

directly and from "reducible" sources caused by mismeasurement of other particles as photons. While it is possible to reduce the signal-to-background ratio somewhat through kinematic selection and isolation requirements, the final sample will still have significantly more background events than signal and relies very strongly on an accurate prediction of the background.

$H \rightarrow WW \rightarrow \mu\mu\nu\nu$ **Analysis** The purely leptonic final states that result from $H \rightarrow WW$ decays with a BR of about 4% can be selected very efficiently by looking for appropriate muon pairs with high p_T. Backgrounds mostly come from $t\bar{t}$ and $b\bar{b}$ decays, which can be effectively suppressed by discarding events with high-E_T jets, and Drell-Yan produced Z bosons, which can be suppressed by requiring significant missing energy in the event (see Fig. 2.8).

$H \rightarrow ZZ \rightarrow 4\ell$ **Analysis** Finally, the $H \rightarrow ZZ$ decay with purely leptonic end states offers a number of options where the background is primarily through direct ZZ production, which is well-understood and of comparable magnitude to the signal. The most distinctive such signature is the $H \rightarrow ZZ \rightarrow 4\mu$ decay, which is often referred to as "gold-plated".

The event selection here requires the existence of two opposite sign muon pairs, each consistent with the Z mass. The backgrounds to this signature caused by b quark decays can be suppressed through isolation requirements, leaving direct ZZ production as an irreducible background, which is most relevant for $m_H \approx 2m_Z$. However, the signal clearly dominates this background in the relevant mass range (see Fig. 2.9).

A similar analysis is possible for states for states with one or two electron pairs instead of muon pairs. In addition to the discovery potential, this channel can be used

Figure 2.9: Signal and backgrounds for simulated H → ZZ → 4μ decays in CMS [2].

to evaluate the Higgs mass, decay width and production cross section.

2.2 Beyond the Standard Model

While the Standard Model has not been experimentally contradicted in all accessible situations (or can be extended to accommodate them), its structure presents a tantalizing mix between features forced by the mathematical structure (such as the gauge structure of interactions) and seemingly arbitrary values (such as the many free parameters, the fermion content or the choice of gauge groups). The history of unification that led up to the Standard Model leads many physicists to believe that there should be a more fundamental description of nature, hopefully also including a description of gravitation. This section reviews a number of candidate theories that have been studied for their experimental consequences at LHC energies.

2.2.1 Supersymmetry

The most studied candidate theories for such an extension are based on the concept of supersymmetry, which extends the existing symmetries of quantum field theory by a symmetry relating bosonic and fermionic fields with each other. Its theoretically appealing features include a solution to the hierarchy problem and the prospect of unifying the gauge couplings of the Standard Model at an intermediate energy scale. Trying to produce a realistic theory incorporating supersymmetry in the Standard Model requires a doubling of the particle content, since none of the existing particles can be superpartners of each other. In addition, supersymmetry must obviously be broken at some energy scale because the superpartners have so far not been observed. Some indicators, such as the hoped-for gauge coupling unification, require the mass

differences between particles and their superpartners to be on the order of 1 TeV, making it a likely candidate for discovery at the LHC. There are multiple theories on the mechanism of supersymmetry breaking which differ in their prediction of the superparticle spectra, but under realistic conditions, their phenomenology does not differ greatly.

In the mSUGRA framework, the specifics of SUSY breaking can be described by five parameters - the fundamental boson mass m_0, the fundamental fermion mass $m_{1/2}$, the trilinear coupling A_0, the sign of the Higgs vacuum expectation value $\text{sgn}(\mu)$ and the ratio of the vacuum expectation values of the two Higgs doublets $\tan(\beta)$. By examining the dependence of the particle spectra on these parameters, it is possible to further reduce these to a number of points in the parameter space that are representative of regions with similar low-energy phenomenology.

Experimental Signatures

Baryon and lepton number violation are severely constrained experimentally, but generally present in a supersymmetric theory. A way to rule them out is imposition of an extra multiplicative quantum number "R-parity", $P_R = -1^{3(B-L)-2S}$. All particles of the Standard Model have $P_R = 1$, while all the superpartners have $P_R = -1$. An immediate consequence of this requirement is that superpartners are produced in pairs from initial states containing only Standard Model particles. It also forces the lightest supersymmetric particle (LSP) to be stable since it can not decay to Standard Model particles. A typical SUSY event then will be characterized by large missing energy from the escape of two undetected LSPs and a "positive" signature from the decay chains leading down to these states. Depending on the decay chains, multiple ways to isolate SUSY events have been studied and generally found to be feasible for an early discovery.

2.2.2 Extra Dimensions

Multiple attempts have been made to unify quantum theory and gravitation by embedding the four-dimensional space-time of standard quantum field theory in a higher-dimensional settings. These theories usually predict heavy bosonic particles that show up as resonances in the dijet/dilepton spectra.

A much-publicized consequence of such models is the production of black holes in large extra dimension models. Such states are expected to decay thermally via Hawking radiation within $\approx 10^{-27}$ s. A special feature of these decays is their "democratic" nature, consisting of all existing particle types and lacking a preferred direction. Since the total energy in the decay has to exceed the threshold set by the higher-dimensional Schwarzschild radius, one can select events where the total invariant mass of the decay products is greater than a TeV-scale threshold. Backgrounds can be significantly reduced by selecting highly "spherical" events, which are more likely to result from Hawking radiation.

2.2.3 Contact Interactions

Contact interactions are a generic framework describing the effects of hypothesized particles with an energy larger than the TeV scale, analogous to the Fermi current-current description of the weak interaction. Predicted effects include shifts in the high-energy part of the dimuon spectrum and changes to the geometric distribution of dijet events.

2.2.4 Alternative Electroweak Symmetry Breaking

Technicolor

Technicolor models provide an alternative mechanism of electroweak symmetry breaking by invoking a new QCD-like interaction that provides dynamical masses for electroweak gauge bosons and fermions. These models predict "technimesons" (the spin-zero π_{TC} and spin-one ρ_{TC}, a_{TC} and ω_{TC}), condensates of the fundamental particles with TeV-scale masses, which should be detectable as resonances. The best modes for discovery at the LHC are expected to be $\rho_{TC} \to WZ \to 3\ell\nu$, $a_{TC} \to \gamma W \to \ell\nu\gamma$ and $\omega_{TC} \to \gamma Z \to \ell\ell\gamma$ decays, which are fully leptonic and thus benefit from reduced QCD background and better mass resolution.

Little Higgs

"Little Higgs" models propose a solution to the hierarchy problem of the Standard Model by proposing new TeV-scale particle to cancel the quadratically divergent Higgs mass corrections at one-loop order only. The resulting effective theory is then valid only up to about 10 TeV only, but does not require distasteful fine-tuning. The predicted particle spectrum depends on the detailed realization of these cancellations. In the "Littlest Higgs" model, for example, these additional particles are heavy copies of the electroweak gauge bosons and heavy quarks as well as three scalar electroweak particles with charges 2, 1 and 0. Detection of the doubly charged boson as well as the heavy gauge bosons and top quarks has been studied.

2.2.5 Same-Sign Top Pairs

In the Standard Model, there is no mechanism for exclusive production of same-sign top quark pairs from hadron collisions. While processes that produce same-sign top pairs inclusively certainly exist, they contain a multitude of weak interaction vertices and thus have negligible cross sections. The detection of same-sign top pairs would therefore be a strong indication of new physics. One of the candidate explanations would be the existence of flavor changing neutral currents.

Detection of such events can follow the lines of the search for ordinary leptonic top pair searches with a simple change of sign requirements.

3 Experimental Setup

3.1 Large Hadron Collider

Ever since the discovery of cathode rays leading to J. J. Thomson's identification of the electron, further insights into the structure of matter have been provided by the study of electromagnetically accelerated particles. In the long development of particle accelerators, several different limitations had to be overcome to arrive at current basic design. The energy transferred to a charged particle is proportional to the traversed voltage differential, but increasing the voltage faced technical limits at the MV scale. It is therefore necessary to use multiple stages of acceleration to achieve higher energies. The cost of such devices scales linearly with energy, limiting the energy reach. To achieve higher energies, it is necessary to reuse the acceleration mechanism. This requires bending the particle beam in a magnetic field. Early cyclotron designs used homogeneous magnetic fields to return the particle beam to the accelerating field. Here, the fact that the bending radius increases throughout the acceleration requires the bending field to be present over a prohibitively large area. This is overcome by the synchrotron design, where the bending magnetic field is adjusted to produce a constant bending radius, requiring a magnetic field only along the path of the beam.

Once the kinetic energy of accelerated particles becomes large compared to their rest mass, colliding accelerated particles with a fixed target becomes inefficient since momentum conservation constrains the amount of energy available for reactions to scaling with $\sqrt{E}$. In contrast, when two particles with opposite momenta collide, the total conserved momentum is zero and the entire energy of both particles is available. For this reason, colliders are now generally used to achieve the highest energies.

The Large Hadron Collider currently being commissioned at CERN is the latest in a series of "discovery machines" intended to extend the frontiers of high energy physics. Historically, all machines in a similar role have been hadron colliders (see Fig. 3.1). While lepton colliders provide a cleaner signal, a synchrotron design using electrons can not easily be extended to higher energy regions, principally due to synchrotron radiation losses, which reached a staggering $13\,\mathrm{MW}$ at a center-of-mass energy of $94\,\mathrm{GeV}$ in LEP and scales with $(E/m)^4/\rho^2$, where m is the rest mass of the particle and ρ is the bending radius [10]. This problem is alleviated by using heavier particles (as in hadron colliders or proposed muon colliders) or abandoning the circular geometry in favor of linear colliders.

The LHC provides two counter-rotating beams of either protons or heavy ions. Due to space limitations in the LEP tunnel, the design unites both of the beams in a single ring. To achieve its design goal of providing a center-of-mass energy of $14\,\mathrm{TeV}$ in p-p collisions, the bending dipole magnets need to produce a magnetic field of

Figure 3.1: A "Livingston plot" illustrating the development of accelerator technology [9].

8.3 T. Such high fields were not achievable in the previous generation of accelerator magnets primarily due to the critical current density at their operating temperature of 4.8 K. An improvement is effected in the LHC dipoles by using superfluid helium at a temperature of 1.9 K for cooling.

3.1.1 Layout and Parameters

The general geometry of the LHC lattice is imposed by the characteristics of the preexisting LEP tunnel that houses it. The total layout is split into eight arcs interrupted by so called "long straight sections" (LSS). Each of these LSS is located at an access point to the tunnel. While there is only a single accelerator ring, the two beams are kept in separate vacuum pipes, with crossings occurring at four insertion regions (IR) where the experiments are located: ATLAS at IR1, ALICE at IR2, CMS at IR5 and LHCb at IR8. IR4 houses the RF acceleration and feedback systems for the beams. IR3 houses the momentum cleaning assembly, which strips off leading or trailing particles that do not fit the acceleration envelope provided by the RF system. IR7 contains the betatron cleaning system, which removes particles with large transverse oscillations from the beam [11].

Injector Chain

The LHC is only the final stage of a chain of proton accelerators at CERN. Initially, hydrogen gas is ionized and the protons are accelerated to about 75 keV in a duoplasmatron ion source. A linear RFQ accelerator, LINAC2, then accelerates them to about 50 MeV. Further acceleration stages proceed through three synchrotrons, the Proton Synchrotron Booster (PSB, up to 1.4 GeV), the Proton Synchrotron (PS, 25 GeV) and the Super Proton Synchrotron (SPS, 450 GeV).

Lead ions for heavy ion operation are produced in the Electron Cyclotron Resonance (ECR) ion source, where they are stripped down to a Pb^{27+} state on average. Then, a separate linear accelerator (LINAC3) accelerates them up to an energy of about 4.2 MeV per nucleon before they pass through a carbon foil stripper that increases their ionization state to Pb^{54+} into LEIR, the Low Energy Ion Ring. Because the number of ions produced per source pulse is low, multiple pulses are stacked and conditioned in LEIR. After acceleration up to 72 MeV/n, they pass into the PS and are accelerated to an energy of 1.4 GeV/n. They pass afterwards to the SPS through a second-stage stripper that produces a fully ionized Pb^{82+} state and are accelerated up to 177 GeV/n before passing into the LHC proper.

Due to the needs of RF acceleration and experiment timing, the particles in the LHC travel in bunches. The fundamental LHC clock has a rate of 40 MHz, which provides a natural value for the bunch separation time of 25 ns. This creates space for 3564 bunches on each beam, but practically envisioned operation schemes need to respect further operational constraints such as the switching times of the various injection kicker magnets along the way. The expected bunch dispositions for nominal operation with protons and heavy ions are depicted in Fig. 3.2 and Fig. 3.3, yielding

Figure 3.2: Beam disposition at the LHC for proton operation [12].

a final number of 2808 proton bunches or 591 ion bunches.

Beam Life Cycle

After receiving a filling from the SPS, the LHC runs a "ramp up" program to accelerate the beam up to the final energy, which takes approximately 20 minutes. Once the final energy is reached, the beam is kept in circulation while the luminosity steadily decays, principally through the effects of the collisions, but also through interaction with rest gas in the beam pipe. This decay is modeled as an exponential decay of luminosity $L(t) = L_0 e^{-t/\tau_L}$ with $\tau_L = 15\,\mathrm{h}$. Once the luminosity becomes too low, the beam is aborted by directing it into one of the beam dumps. A controlled "ramp down" of the magnet system also takes about 20 minutes. Including the time spent in system checks and the LHC injection chain, the "turnaround time" between initiation of a beam abort and the availability of a new beam at full energy has been estimated to be 70 minutes in the optimal case. This suggests a beam life of 5.5 hours to optimize integrated luminosity. A more pessimistic scenario assumes a sixfold increase in the turnaround time, which increases the optimal beam life to 12 hours.

3.1.2 Proton-Proton Operation

In proton-proton operation, the LHC collides two proton beams at an expected center-of-mass energy of $\sqrt{s} = 14\,\mathrm{TeV}$, with a design luminosity of $\mathcal{L} = 1.0 \times 10^{34}\,\mathrm{cm}^{-2}\mathrm{s}^{-1}$.
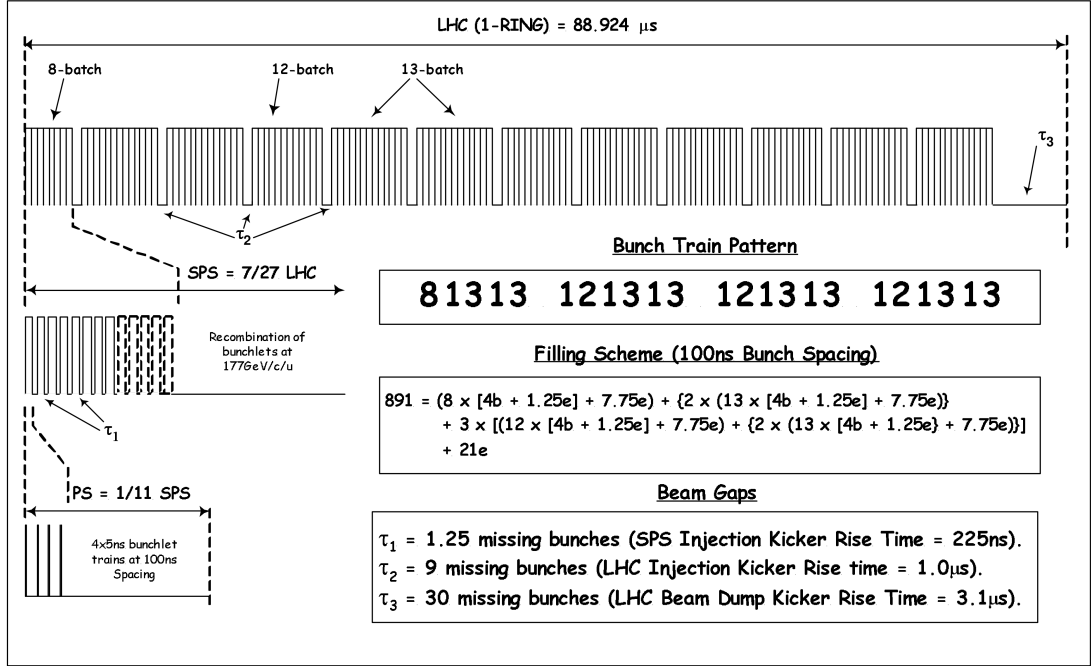
Figure 3.3: Beam disposition at the LHC for nominal ion operation [12].

3.1.3 Heavy Ion Operation

In addition to the standard p-p operation, the LHC is also designed to accelerate heavy ions to a final energy of about 2.76 TeV/nucleon. While the luminosity in heavy ion mode is only $\mathcal{L} = 2 \times 10^{27}\,\text{cm}^{-2}\text{s}^{-1}$, the cross section for heavy ion collisions is considerably larger than for p-p interactions. The primary consumer of the LHC's heavy ion operation is the dedicated ALICE experiment; in addition, both CMS and ATLAS also have heavy ion research programs [13] [14].

3.2 The CMS Experiment

3.2.1 Overview

CMS (**C**ompact **M**uon **S**olenoid) is one of the two large general purpose detectors at the LHC. While its initial design was mostly guided by the search for the Higgs boson, it has since been amended to include a number of different physics goals from the LHC program. The detector is constructed around its central structural feature, a superconducting solenoid magnet with an inner diameter of 5.9 m. This large bore allows placement of both tracking detectors and calorimeters inside the magnet. The stray field from the solenoid is concentrated in a surrounding return yoke constructed of iron. Its thickness (1.5 m of iron) allows four layers of muon stations to be integrated.

Due to the extremely high event rate (10^9 p-p collisions/second at LHC design

luminosity) and limited storage bandwidth (only about 100 full events/second), CMS also requires an extraordinarily powerful trigger and Data Acquisition (DAQ) system to select interesting events.

3.2.2 Coordinate System

The coordinate system used CMS is centered at the nominal interaction point, with the y axis pointing straight up and the x axis pointing at the center of the LHC arc. This leaves the z axis pointing in the direction of the clockwise LHC beam. Since the detector has an approximate cylinder symmetry, it is also useful to define an azimuthal angle φ in the x/y plane with the x axis at $\varphi = 0$. The corresponding polar angle θ is measured from the positive z axis. Instead of θ, the pseudorapidity $\eta = -\ln(\tan(\theta/2))$ is usually used. As the name suggests, it is a good approximation for the relativistic rapidity $\xi = \frac{1}{2}\ln(\frac{E+p_Z}{E-p_Z})$ in the z direction. This can be seen by substituting $p_z = |p|\sin(\theta)$ and performing some trigonometric substitutions in η to arrive at

$$\eta = \frac{1}{2}\ln(\frac{|p|+p_Z}{|p|-p_Z}) \tag{3.1}$$

If the momentum is large or the mass is small, $E \approx p$ obtains and the approximation holds. This allows one to approximately keep the useful properties of rapidity (such as invariance of rapidity differences under Lorentz boosts) while having a general coordinate (independent on the mass of the particle under consideration) that is comparatively simple to calculate.

3.2.3 Inner Tracker

The inner tracker is the innermost component of the CMS detector, surrounding the interaction point in a cylindrical volume that is $5.8\,\mathrm{m}$ long and and has a diameter of $2.6\,\mathrm{m}$. The CMS solenoid provides a homogeneous $4\,\mathrm{T}$ magnetic field in this region. The tracker is designed to measure the tracks of charged particles to very high precision and allow identification of secondary interaction vertices. To identify such vertices from b and τ decays requires a spatial resolution of $10 - 20\,\mu\mathrm{m}$. At LHC design luminosity, the tracker will be traversed by about 1000 charged particles per bunch crossing. To combat occupancy effects in this scenario requires high granularity, so that the innermost layers of the tracker consist of silicon pixel detectors. Since the particle flux per unit area through the outer layers is much lower, these can be built out of silicon microstrips. High granularity and fast readout requirements combine to cause significant power dissipation from the on-detector electronics, which needs to be removed by an efficient cooling mechanism. Furthermore, it is desirable to minimize the amount of material in this part of the detector to avoid unwanted interactions.

The Pixel Detector

The pixel detector consists of the innermost three layers in the barrel tracker (at mean radii of $4.4\,\mathrm{cm}$, $7.3\,\mathrm{cm}$ and $10.2\,\mathrm{cm}$) and the innermost two layers of the tracker

endcaps (at $z = \pm 34.5\,\text{cm}$ and $z = \pm 46.5\,\text{cm}$). Figure 3.4 shows the entire setup. It provides multiple hit coverage over a pseudorapidity range of $0 < |\eta| < 2.5$. The total sensor area is about $1\,\text{m}^2$ and contains 66 million $100 \times 150\,\text{μm}^2$ pixel cells. Each of these pixel cells is basically a reverse-biased diode, where charged particles are detected through the ionization currents they induce. Since the barrel pixels' bias voltage differential is in the radial direction, the high magnetic field causes a Lorentz angle of almost $23°$. This increases charge sharing between neighboring pixels, which is very useful to obtain sub-pixel location information. The spatial resolution of the pixel detectors thus reaches about $10\,\text{μm}$ in the r and ϕ measurements and $20\,\text{μm}$ for the z direction. In the endcaps, the pixel detector modules are installed at an angle of $20°$ to the magnetic field to induce similar charge sharing effects.

The Silicon Strip Tracker

The silicon strip tracker (SST) is the world's largest detector of this kind. It contains 15148 microstrip modules, covering an active silicon area of $198\,\text{m}^2$ in 15148 modules and about 6 million strips. The strip size varies between $81\,\text{μm} \times 10\,\text{cm}$ and $183\,\text{microns} \times 25\,\text{cm}$, generally increasing with distance from the interaction point. In the barrel region, there are four layers of inner barrel detectors (TIB) and six layer of outer barrel detectors (TOB). The distinction between these is mainly due to the fact that the distribution of the strip detector modules attempts to minimize shallow crossing angles. Thus, the inner barrel is shorter than the outer barrel and closed by three "mini endcaps", the inner discs (TID). The TOB is closed by nine layers of endcaps (TEC). In order to obtain two-dimensional hit information from the strip detectors as well, part of the modules are so-called "stereo" modules that consist of two modules glued together at an angle of $100\,\text{mrad}$ [15]. This achieves a spatial z-resolution of $230\,\text{μm}$ in the TIB and $530\,\text{μm}$ in the TOB, respectively; the r-ϕ single-point resolution is $23 - 34\,\text{μm}$ in TIB and $35 - 53\,\text{μm}$ in TOB. The layout ensures that at least nine hits are recorded in the range of $0 < |\eta| < 2.4$, with at least four of those from double-sided modules. Figure 3.5 shows the distribution of tracker modules around the interaction point.

3.2.4 Electromagnetic Calorimeter

The CMS electromagnetic calorimeter (ECAL) is a homogeneous crystal calorimeter consisting of 61200 lead tungstate ($PbWO_4$) scintillator crystals in the barrel and two endcaps of 7324 crystals each. It measures the energy of particles whose primary interaction is through the induction of electromagnetic showers (photons and electrons). The primary design benchmark was observation of the $H \rightarrow \gamma\gamma$ decay, which requires very good ECAL energy resolution.

$PbWO_4$ **Crystals**

An optimal scintillator medium is characterized by fast light decay time, high light yield and both small Moliére radius and radiation length. $PbWO_4$ compares favorable

Figure 3.4: Layout of the CMS pixel detectors [16].

Figure 3.5: Layout of a quarter of the complete inner tracker [16]. Modules shown in red measure only one coordinate, while modules shown in blue are either pixel detectors or stereoscopic strip modules and measure both coordinates.

	NaI(Tl)	BGO	CSI	BaF$_2$	CeF$_3$	PbWO$_4$
Density [g/cm^3]	3.67	7.13	4.51	4.88	6.16	8.28
Radiation length [cm]	2.59	1.12	1.85	2.06	1.68	0.89
Interaction length [cm]	41.4	21.8	37.0	29.9	26.2	22.4
Molière radius [cm]	4.80	2.33	3.50	3.39	2.63	2.19
Light decay time [ns]	230	60 300	16	0.9 630	8 25	5 (39%) 15 (60%) 100 (1%)
Refractive index	1.85	2.15	1.80	1.49	1.62	2.30
Maximum of emission [nm]	410	480	315	210 310	300 340	440
Temperature coefficient [%/°C]	~0	–1.6	–0.6	–2/0	0.14	–2
Relative light output	100	18	20	20/4	8	1.3

Table 3.1: Comparison of scintillator materials used in high-energy physics [17].

to other candidate materials (see Table 3.1) in all these respects except for a comparatively low light yield. However, this disadvantage can be compensated by modern photodetectors. Further complications are introduced by the temperature dependence of the light output, which requires cooling to maintain the standard operating temperature of 18° C throughout the instrument and the radiation environment. The direct effect of radiation on the crystals leads to the formation of color centers, which reduce the transparency of the crystal. However, since PbWO$_4$ is optically transparent, this type of degradation can be measured through controlled illumination with laser light during operation and compensated for in reconstruction.

Geometry

The ECAL's geometry is illustrated in Fig. 3.6. The barrel part of the ECAL (EB) covers the range of $0 < |\eta| < 1.479$. The granularity is 360-fold in ϕ and (2×85)-fold in η. Each crystal has a tapered pyramid shape, measuring about $22 \times 22 \, \text{mm}^2$ (matching the Moliére radius) on the inner face and $26 \times 26 \, \text{mm}^2$ on the outer face, with a length of $230 \, \text{mm}$ or $25.8 \, \text{X}_0$, chosen to limit longitudinal shower leakage to acceptable levels. The placement of the crystals is almost projective with regard to the interaction point, but slightly offset in both η and ϕ to avoid leaving cracks along likely particle paths.

The endcaps extend the range to $1.479 < |\eta| < 3$. Crystals on the endcaps have a front face area of $28.62 \times 28.62 \, \text{mm}^2$, a rear face area of $30 \times 30 \, \text{mm}^2$ and a length of $220 \, \text{mm}$ or $24.7 \, \text{X}_0$. Again, hermeticity is improved by focusing the front faces of the crystals on a point about $1 \, \text{m}$ behind the nominal interaction point.

Figure 3.6: Transverse section illustrating the geometry of the CMS ECAL [16].

Preshower Detector

One of the major reducible backgrounds in the $H \to \gamma\gamma$ decay comes from the misidentification of photon pairs from the decay of high-energy neutral pions as single high-energy photons. Especially in the endcap region, the spatial resolution of the crystal calorimeter is insufficient to make this distinction. Thus, a special preshower detector (ES) is inserted before the endcap calorimeters. It consists of two layers, each comprised by a lead radiator to induce showering and a silicon strip detector with 1.9 mm strip pitch that measure the energy deposited and the lateral shower profile. The thickness of the preshower absorbers (about $3\,X_0$) was chosen as a compromise between the requirement to reliably induce showering for photons of the relevant energy range versus the negative impact of the preshower detector's less precise energy measurement on the total resolution of the crystal calorimeter [18].

ECAL Performance

The relative energy resolution of the ECAL is usually described through the width Gaussian distribution with contributions from three sources:

$$\left(\frac{\sigma}{E}\right)^2 = \left(\frac{S}{E}\right)^2 + \left(\frac{N}{\sqrt{E}}\right)^2 + C^2 \tag{3.2}$$

S is the stochastic term, N the noise and C the constant term. Figure 3.7 shows some values obtained from beam tests.

Figure 3.7: Energy resolution of a CMS ECAL supermodule obtained from beam tests
[16].

3.2.5 Hadron Calorimeter

The hadron calorimeter (HCAL) consists of a number of sampling calorimeter sub-systems, designed to measure the total energy of hadronic particles. The barrel part (HB) is situated between the barrel part of the ECAL and the cryostat housing the solenoid magnet. It covers the region of $0 < |\eta| < 1.3$ with 17 layers of plastic scin-tillators interspersed with brass absorber plates. The outermost absorber plates are made of stainless steel for structural integrity reasons. Granularity is 108-fold in ϕ and 32-fold in η. Since the total thickness of the HCAL is constrained by the radius of the magnet cryostat, it is not sufficient to fully absorb highly energetic particles. To compensate, an outer barrel (HO) part was added that contains additional scintillator layers located on the outside of the cryostat, using the magnet's inert material as an absorber.

The endcaps (HE) extends the coverage up to $|\eta| < 3$ with a very similar design. In higher $|\eta|$ regions, the radiation environment becomes even harsher. Since extend-ing geometric coverage to these regions is crucial for the determination of missing transverse energy, a separate forward calorimeter (HF) has been included in CMS. It uses a steel absorber structure instrumented with quartz fibers where charged shower particles produce Cherenkov light and extends coverage up to $|\eta| < 5$.

HCAL Performance

The HCAL's performance is usually measured by the jet energy resolution and the missing transverse energy resolution. The relative jet energy resolution from simulation

Figure 3.8: HCAL jet energy resolution for the three subsystems (HB in red, HE in blue, HF in magenta) [16].

results is illustrate in Fig. 3.8. The missing transverse energy resolution has been determined as $\sigma(E_T^{miss}) \approx 1.0\sqrt{\Sigma E_T}$ in QCD dijet simulations where the average missing transverse energy is $E_T^{miss} \approx 1.25\sqrt{\Sigma E_T}$.

3.2.6 Forward Detectors

In the very forward region, CMS will be augmented by the CASTOR and ZDC sub-detectors and work in conjunction with the TOTEM experiment [4].

CASTOR

CASTOR (Centauro and Strange Object Research) is a sampling calorimeter, similar in design to the HF, but covering even higher rapidity ranges ($5.5 < |\eta| < 6.2$). It is installed surrounding the beam pipe, between the CMS endcaps and the HF in $|z|$. It contains quartz plates to generate Cherenkov light from passing particles and tungsten plates as absorbers. The electromagnetic component of the shower is absorbed in a separate front section with thinner layering, while the hadronic component, mainly from neutral pions, is absorbed over the rest of the detector. In heavy ion collisions, CASTOR is expected to record total and electromagnetic energy deposits of about 150 TeV with a resolution on the order of 1%.

ZDC

The Zero Degree Calorimeters (ZDC) is another sampling calorimeter, located at a distance of 140 m from the interaction point, between the two LHC beam lines. Since it is preceded by the LHC beam separator, only neutral particles will reach the detector,

which covers the rapidity range of $|\eta| > 8$ (equivalent to deflection angles $< 400\,\mu\text{rad}$). Projected use cases include triggering for diffractive and heavy ion physics as well as beam optics measurements.

3.2.7 Solenoid

The CMS solenoid provides a magnetic field parallel to the beam axis with an induction of $B = 4\,\text{T}$ at the interaction point. This high magnetic field is required to provide enough bending power to achieve the requirements for the muon momentum resolution. The conductor consists of 4 layers of superconducting NbTi cabling directly bonded with aluminium alloy to achieve the necessary structural stability. To maintain the superconducting state, the 220 t cold mass must be enclosed in a stainless steel vacuum tank and kept below its critical temperature using liquid helium cooling. The nominal operating current is 19.1 kA.

The magnetic flux from the solenoid is returned on the outside through a 1.5 m thick steel yoke that also contains the CMS muon system and is completely saturated at nominal conditions. The yoke consists of 11 elements (6 endcap disks and 5 barrel wheels) that can be moved separately. The central barrel wheel holds the magnet coil and the cryostat. At 10,000 t, this yoke contains almost half of the total mass of CMS.

3.2.8 Muon System

The detection of muons is of central importance to CMS because several expected discovery channels have final states including muons. The mass resolution attainable for muons is superior to that of other leptonic final states because they are much less affected by material interactions in the tracker. The three main tasks of the muon system are identification of muons, measuring of the muon momentum and triggering.

In the barrel, the mainstay of the muon system is formed by drift tubes, which are simply gas-based ionization detectors segmented into tubes for reliability reasons. However, since drift tubes suffer from comparatively slow response, the drift tubes are augmented by resistive plate chamber detectors, which provide better time resolution and thus more reliable bunch crossing identification, but with coarser position resolution.

In the endcaps, a different system is needed because of the higher magnetic field, higher muon rate and greater background. Here, the principal detectors are cathode strip chambers, which measure charges induces on strip cathodes by nearby ionization avalanches.

Drift Tubes

The basic drift tube cell consists of an anode wire with a length of $2 - 3\,\text{m}$, surrounded by a $42 \times 13\,\text{mm}$ tube whose surface contains several cathodes, optimized to attain linear position response over the entire cell volume. The tube is filled with a mixture of about 85% Ar and 15% CO_2. A number of identically oriented cells are organized into layers, which are in turn accumulated into superlayers consisting of four layers

staggered by a half-cell each. A drift tube chamber then contains two or three such superlayers, with wires oriented either in parallel to the beam direction to measure the ϕ coordinate or orthogonal to it to obtain a z measurement.

The single wire resolution for drift tube measurements is about 250 μm, yielding a combined r-ϕ resolution of about 100 μm for an entire chamber.

Cathode Strip Chambers

In the endcaps, drift tubes are not practical because muon and background rates are much higher. In addition, the magnetic field is inhomogeneous, which severely complicates the drift tube's measurements. For this reason, the principal muon detectors there are cathode strip chambers (CSC). A CSC is trapezoidal in shape and consists of six gas gaps, each featuring strip cathodes running radially with a width of $3 - 16\,\mathrm{mm}$ and anode wires running roughly perpendicular to the cathode strips with a spacing of $3.16\,\mathrm{mm}$. The signals from ionization avalanches hitting the anode wires and the mirror charges induced on the cathode strips provide two separate readouts on each layer. The anode readout allows precise determination of the hit timing, while the exact spatial information is determined from the distribution of the mirror charges induced on the surrounding cathode strips.

Resistive Plate Chambers

Both the drift tubes in the barrel and the cathode strip chambers in the endcaps are augmented by resistive plate chamber (RPC) detectors, principally because RPCs provide very good time resolution and acceptable position resolution at comparatively low cost. In an RPC, ionization in a gas gap between two highly resistive electrodes creates an initial charge density. Due to the high resistivity, this can not be immediately compensated by the electrodes but introduces a mirror charge on a separate pickup electrode, which is the useful output signal of a RPC [19]. In CMS, each layer of RPCs consists of two separate gas gaps around a set of central readout strips. This eases operation because the signal charge induced is induced by the sum of the ionization charges in both gaps, so that each can operate at a lower gain.

In the barrel, the readout strips are oriented parallel to the beam axis and vary in size so as to cover about $5/16°$ in ϕ. In the endcaps, readout strips run radially with the same ϕ granularity.

Performance of the Muon System

The determinant of the muon system's performance is its ability to precisely determine the muon's transverse momentum p_T through the bending angle of the track. Since centrally produced muons pass through the tracker as well as the muon system, the final measurement combines results from both subsystems. With low p_T, the muon system's resolution degrades because of multiple scattering in the detector materials that traversed before reaching the muon system, and the tracker measurement is much more accurate. At high p_T, however, effects from energy loss and multiple scatterings

Figure 3.9: The muon momentum resolution using the tracker only, the muon system only or both ("full system") [16].

become less important so that the muon chambers' accuracy is limited by the position resolution. For such tracks, the tracker's p_T resolution degrades through the limited bending power of the small radius, and the resolution of the combined measurements is much better than the individual resolutions. Figure 3.9 illustrates this behavior.

3.2.9 Trigger and Data Acquisition

The high-level organization of the CMS data acquisition path is shown in Fig. 3.10. Since the entire detector has about 55 million readout channels and collisions occur at a rate of 40 MHz in nominal LHC operation, it is clearly impossible to store all data for later analysis. Even though the amount of data can be reduced to about 1 MB per event using zero suppression, the available storage bandwidth is still only on the order of 100 events per second. In order to resolve this problem, CMS employs a two-level approach. While digitized data is held in local front-end systems (FES) on the subdetectors for up to 3.2 μs (128 bunch crossings), the Level-1 Trigger subsystem has to decide whether an event is potentially interesting for physics analysis on the basis of coarse-grained information from the muon and calorimetry subsystems. Since most of this time is spent in signal transitions, the processing time for this decision is limited to about 1 μs and makes an implementation using custom hardware and FPGAs necessary. The implementation of the Level-1 Trigger will be described in more detail in Ch. 4. By design, the Level-1 Trigger should produce a maximum output rate of 100 kHz. Once the Level-1 Trigger decides to accept an event, the

acceptance signal is distributed to the front-end driver (FED) modules, which copy the data from the pipelined buffer into the data acquisition system (DAQ).

The remaining reduction of the event rate is then performed by the high-level trigger (HLT) Event Filter, which consists of a large number of commercial processors running the HLT algorithms. Since HLT algorithms need to access the entire data of the event under consideration, which is stored on over 600 FED components all over the detector, an intermediate system is needed to assemble the events and deliver them to processing nodes. The main challenges for the Event Builder subsystem are the large number of connected components and the high event rate, leading to the requirement to provide a sustained transfer rate of 100 GB/s. Assembly of an event's data proceeds in stages, first concentrating up to two FED units in one front-end readout link (FRL). Data from the FRLs is concentrated into 72 so-called super-fragments of about 16 KB each using a Myrinet switching fabric. These super-fragments are stored in memory buffers on readout units (RU), which are standard PCs connected both to the Myrinet network and a copper-based ethernet network that connects them with the processing units. These processing units consist of standard PCs running a number of software components - the builder units (BU) assemble an entire event from the super-fragments in different RUs, the filter units (FU) run the HLT algorithms and forward the event data over a separate SAN link for storage if accepted.

In addition to the central task of event selection, the DAQ system also allows the execution of additional analysis modules in the Event Filter that perform quality and integrity checks on the processed data, called data quality monitoring (DQM). These provide quick feedback and allow detection of various detector problems without waiting for the process of offline reconstruction.

3.2.10 CMS Online Software

Online software refers to software that is directly required for the operation of the detector. The top-level subsystems are the data acquisition system (DAQ), the run control and monitoring system (RCMS) and the detector control system (DCS). Most of the CMS online software is based on an internally-developed C++ middleware framework named XDAQ, which provides a common basis for the applications running in a clustered computing environment.

Run Control and Monitoring System

RCMS provides a central user interface to the experiment. It provides a common control and monitoring infrastructure for the detector itself (through the DCS) and the DAQ system. Since users need to access this system from all over the world, its primary user interface is web-based. Since its tasks include setting up the software side of the DAQ system, the overall number of components under its control is very high, with about 3000 PCs and 10000 applications involved.

The run control system is organized as a hierarchy of Function Managers (FM), with the central system acting as the root of the tree. The lower-level function man-

Figure 3.10: Schematic representation of the data acquisition path in CMS [20].

agers are provided by the subsystems and expose a standardized web service interface, allowing independent implementation choices. For control purposes, each FM exposes a finite state machine (FSM, see Fig. 3.11) with the central system initiating state transitions. In addition, the RCMS system provides a central configuration database, a log collection facility and a utility for the definition of DAQ configurations.

Detector Control System

The purpose of the detector control system is to ensure correct operation of the experiment with respect to the underlying infrastructure, such as cooling, gas systems or power supplies. It is based upon a commercial SCADA system. The DCS keeps track of almost a million detector parameters and can respond to alarms by initiating automatic recovery actions or notifying human experts. During data taking, it is controlled by RCMS through a web service bridge.

3.2.11 CMS Offline Software

After data from the detector has passed the selection process and has been committed to mass storage, it undergoes standardized processing to reconstruct objects amenable to physics analysis from the raw detector data. To facilitate further analysis, the experiment provides a common C++ framework called CMSSW. The basic concepts in its Event Data Model (EDM) are the *Event*, which contain data associated with a single bunch crossing and *Conditions*, which are comparatively stable external influences such

Figure 3.11: Generic state machine for a DAQ or detector component from the view of the run control system.

as alignment parameters, detector geometry or configuration settings. Each Condition is also associated with an interval of validity (IOV) which describes which events it applies to.

CMSSW provides a Python-based configuration language which allows users to build a processing path out of several types of modules. This configuration defines the selection and order of modules to run, as well as the values for the modules' configuration parameters. Figure 3.12 contains a schematic overview of the components involved in a CMSSW process. Modules fall into one of the following categories:

- Input modules provide event data from an external source, such as a data file or a network stream.

- Producers extend the event data by adding additional records, such as reconstructed track data computed from raw detector hits.

- Filters are used in a configured path to select a subset of events based on their contents, for example applying a quality cut in a physics analysis.

- Analyzers do not modify the contents of the event, but have read-only access to its data. They can be used to create histograms or reports of the events analyzed.

- Output modules store the events after processing.

Framework modules use the event content as their only communication interface. A record inside an event, also called a *product*, is identified by its data type and a number of labels. In addition, the framework stores provenance information that allows a later user to determine the modules and configuration involved in its generation.

Event Data Processing

The initial software view of the detector readout consists of a list of binary streams read out from the various FED units in the system. Starting with the HLT, later

processing requires a more accessible interface. Subsystem-specific unpacker modules therefore transform the raw data formats into logical objects that can be used by reconstruction algorithms to determine physics objects. Since the HLT runs under strict time constraints, it uses regional reconstruction and optimized reconstruction techniques that may be less accurate than full reconstruction for its decision. The event written to permanent storage after acceptance consists of the initial detector readout plus HLT information; the Level-1 Trigger information is part of the readout as the Level-1 Trigger is a hardware system. This is expected to amount to 1-2 MB/event.

Based on this information, the event stream is split into a number of primary datasets which are processed separately from then on. This allows priorities to be placed on streams with more interesting content. Offline processing then recreates the digitized hit information (digis) and performs fine-grained reconstruction to the level of physics objects. Since reconstruction procedures are expected to improve over time, however, the raw data is kept in the event as well. The size of the reconstructed information is about 250 kb/event.

This implies that the event size is dominated by raw data, which is not used in direct physics analysis. For this reason, the event data is divided into a number of different data tiers. RAW data contains the detector readout and HLT results, RECO data is the much smaller reconstruction information, and another level called AOD (Analysis Object Datasets) contains a subset of commonly used RECO objects with a size of only about 50 KB/event.

Simulation Processes

Monte Carlo simulation is an essential tool for both the development of the detector hardware and the physics program to be executed. The CMSSW framework includes interfaces to several standard particle generators. Generated processes can then be run through either a full detector simulation based on Geant4 or a fast detector simulation that sacrifices accuracy of the model for processing speed. The result of this simulation is a set of detector hits, roughly equivalent to the analog detector outputs. From these the simulation produces digitized hits by simulating the front-end electronics. These can then either be converted to the readout format to produce RAW data compatible with that from detector readout or directly processed like digis unpacked from recorded raw data.
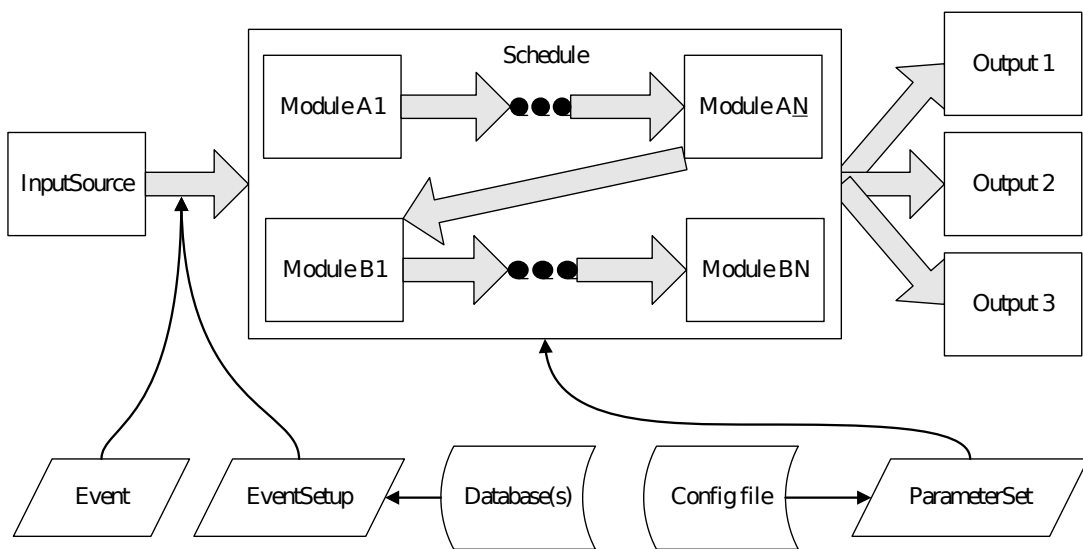
Figure 3.12: Components of a CMSSW framework process [16].

4 Level-1 Trigger

The Level-1 Trigger is the first stage of the event selection process. It has to render a decision to retain or discard an event on every bunch crossing. While this decision is made, all data gathered has to be retained in the detector. The amount of memory available in the tracker and preshower detectors limits this retention period to with 128 bunch crossings ($3.2\mu s$). Given the detector technologies used (with DT drift times of up to 400 ns) and the size of the detector and counting room (signal propagation times add up to $1\,\mu s$), the time for actual decision logic calculations is reduced even further. To produce a decision on each bunch crossing, the trigger system needs to run as a pipeline, with multiple bunch crossings in various stages of processing at the same time. Because it is not possible to process information from the preshower and tracker detectors quickly enough, the trigger only considers physics objects from the calorimeter and muon subsystems. Figure 4.1 shows a high-level view of the subsystems of the Level-1 Trigger. Both the calorimeter and the muon trigger have a similar processing flow, starting with the generation of so-called trigger primitives, which are determined locally on detector elements. These primitives are then combined into object candidates by regional triggers, and a global subsystem trigger stage selects the most interesting of these candidates. The final stage is the Global Trigger (GT), combining information from the subsystems and implementing the event selection criteria derived from physics requirements.

4.1 Calorimeter Trigger

The calorimeter trigger identifies the most energetic electrons, photons and jets in an event. In addition, it calculates a number of global properties, such as the missing transverse energy (due to undetected particles). It encompasses information from the ECAL, HCAL and HF subdetectors. For trigger purposes, the calorimeters are divided into 52×72 ($\eta \times \phi$) so-called trigger towers (see Fig. 4.2), imposing a common geometry on both the ECAL and HCAL subdetectors. The HF calorimeter does not share geometric coverage with the other subdetectors and contains an additional 4×18 trigger towers in each endcap.

4.1.1 ECAL Trigger Primitive Generation

In the electromagnetic calorimeter, each trigger tower consists of 5×5 crystals. The analog signals from the scintillators are digitized on four parallel amplifiers with gains of 33, 9, 5 and 1 respectively, forwarding the signal from the lowest non-overflowing amplifier as the result. This is encoded in a pseudo-floating point format (12 bits of

Figure 4.1: Overview of the Level-1 Trigger subsystems [21].

Figure 4.2: Layout of calorimeter trigger towers in η [22].

voltage information and 2 bits specifying the gain) for every 25 ns period. For the Trigger Primitive Generation (TPG), a sliding window of ten such samples around the current bunch crossing is considered. The signals from 25 crystals are linearized and summed to five strip sums before undergoing filtering and time deconvolution to suppress noise and identify the originating bunch crossing. These strip sums are used to provide a fine-grain veto bit on the result when the shower profile is not sufficiently concentrated in one strip to be compatible with an electron or photon. The sum of the five strip energies and the fine-grain veto bit are forwarded to a trigger concentrator which serializes inputs from a number of trigger towers and sends them to the Regional Calorimeter Trigger (RCT).

4.1.2 HCAL Trigger Primitive Generation

HCAL TPG is similar in principle to the process for the ECAL. The output signals from photodiodes are digitized to a pseudo-floating point format (5 bits of significant digits and 2 bits of range). An intermediate processor linearizes these values, provides bunch crossing identification and sums the appropriate inputs for a trigger tower. The fine-grain veto bit of the ECAL TPG is replaced by the minimum ionizing particle (MIP) bit which indicates that the energy deposit is compatible with the passage of a muon. The output to the RCT happens in the same format as for the ECAL.

4.1.3 Regional Calorimeter Trigger

The Regional Calorimeter Trigger (RCT) consists of 19 electronics crates in the CMS counting room. 18 of these crates receive the E_T values from the trigger towers in both HCAL and ECAL, each crate covering one half of the detector in z and $40°$ in ϕ. Each of these identifies the four highest-ranking isolated and non-isolated electron or photon candidates in its region and sends them to the global calorimeter trigger. In addition, it coarse-grains the E_T sums over 4×4 trigger tower regions, assigning MIP (if one of the HCAL trigger towers produced a MIP bit), quiet (if the energy sum is below a threshold value) and τ veto (if the activity in the 4×4 region is not compatible with a τ jet) bits to each region. These are forwarded to the Global Calorimeter Trigger (GCT). The HF trigger regions are forwarded to the GCT without processing.

4.1.4 Global Calorimeter Trigger

The Global Calorimeter Trigger (GCT) receives the e/γ information from the RCT and sorts them in two Leaf Cards. A Concentrator Card receives data from the two Leaf Cards and performs the final sorting. Only the overall four highest-ranked candidates out of 36 regional candidates are selected and forwarded to the GT for each of the particle categories output by the RCT. The calorimeter trigger also looks for clusters in the E_T sums provided by the RCT and produces jet candidates from these. The four candidates with highest E_T for forward jets (from high-$|\eta|$ regions), τ jets (without a τ veto bit in their region) and central jets each are forwarded to the global trigger.

In addition, the GCT calculates a number of global energy sums for use in trigger decisions. Total transverse energy (ETT) is the sum of all transverse energies from the regions. Missing transverse energy (ETM) is calculated by splitting the regional E_T into E_x and E_y and reversing the resulting vector. Total jet transverse energy (HTT) is the scalar E_T sum of all jets above a programmable threshold. Missing jet transverse energy (HTM) is the ETM equivalent with jet transverse energy only. It is also envisioned to send the MIP and quiet bits received with the regional E_T sums to the Global Muon Trigger (GMT), but this is currently not implemented.

4.2 Muon Trigger

The muon trigger uses information from all three types of muon detectors in CMS (DT, CSC and RPC). This approach allows the trigger to combine the information from a muon chamber system (DT in the barrel, CSC in the endcaps) with that of the associated RPCs, which serve as dedicated trigger detectors. Since the RPCs and the muon chambers are susceptible to different types of background, an intelligent combination of the information from different subsystems allows an optimal compromise between efficiency and background suppression to be achieved. This task is performed by the GMT.

4.2.1 Trigger Primitive Generation

The logical organization of the DT and CSC triggers is quite similar. Each system produces localized track segments in per-chamber electronics, which are then correlated into global track candidates through a sorter logic. The most promising candidates from each subsystem are then forwarded to the GMT.

Drift Tube Trigger Local Trigger

Figure 4.3 shows an overview of the drift tube local trigger. The first stage in the identification of drift tube tracks is performed by a Bunch and Track Identifier (BTI), which is connected to nine drift tubes over four layers of a single superlayer (SL). Since each superlayer measures only one coordinate, the ϕ and θ measurements are entirely separate at this point. To produce a track candidate, hits in at least three layers are required, but track candidates with hits in four layers are assigned a higher quality. The result from such a measurement is an approximate hit location and track direction.

Since the ϕ coordinate is usually measured on two superlayers separated by a lever arm, the results from BTIs in the two superlayers are then combined by track correlators (TRACO). A local trigger server (TSϕ) receives input from all the TRACOs in a muon station and forwards the two track candidates to the track finder, preferring low bending angles.

Processing in the θ view is simpler since there is only one BTI measurement per track. A cross connection between the θ and ϕ trigger servers allows improved noise

Figure 4.3: Overview of the DT local trigger [22].

reduction by requiring coincident tracks.

Cathode Strip Chamber Local Trigger

The cathode strip chambers also detect separate track candidates in the two coordinates η and ϕ. While both views produce Locally Charged Track (LCT) segments, the anode readout electronics are optimized to provide precise bunch crossing information by measuring the coincidence of hits on multiple layers, while the cathode readout is optimized to provide a precise measurement of the ϕ coordinate.

In the ϕ view, hit information is generated from cathode strip signals digitized at half-strip precision. To suppress the considerable background in the endcap regions, the hit distribution is then converted to track candidates by requiring coincidence between multiple layers compatible with predetermined muon hit patterns. If a match is found, a cathode local charged track (CLCT) is generated and sent to the trigger motherboard (TMB).

For the θ coordinate, signals from wire groups consisting of about 10 anode wires form the primary hit information. A pattern comparison similar to that for CLTCs is performed and the resulting anode locally charged tracks are sent to the TMB. The TMB then correlates θ and ϕ information and forwards the resulting LCT information to a muon port card (MPC), which receives LCT information from nine TMBs and

Figure 4.4: η boundary between the DT and CSC subsystems [22].

forwards only the best two or three LCT candidates to the CSC Track Finder.

Overlap Region Processing

The geometry of the muon system implies that a region exists where muon tracks originating from the interaction point pass through both the barrrel and endcap regions. To improve the measurements made by the DT and CSC subsystems and avoid duplicate processing of track candidates, it is therefore desirable to share information between these systems. This is done at the local trigger level by defining a sharp coverage boundary between the DT and CSC subsystems at $|\eta| = 1.04$ (see Fig. 4.4). To allow this redefinition of coverage, track segments from the MB2/2 muon stations are sent to the CSC Track Finder and track segments from the ME1/3 region are sent to the DT Track Finder.

Drift Tube Track Finder

The first stage of the Drift Tube Track Finder (DTTF) is formed by 72 sector processors. A sector covers a 30° region in ϕ and usually one detector wheel (see Fig. 4.5).

The central wheel is split in half and contains one sector for each to avoid asymmetries. Each sector processor attempts to match the track segments from the separate muon stations in the sector into muon tracks. To do so, it first attempts to link track segment pairs from neighboring stations by extrapolating the measured location and bending angle to the next station. These segment pairs are then assembled into global tracks and assigned p_T, η and ϕ values as well as a quality indicator. Each sector processor produces up to two candidate tracks.

In the next stage, candidates from the six sectors with identical ϕ are combined in a wedge sorter that cancels duplicates arising from a track crossing neighboring sectors. It outputs two muon track candidates per wedge, which are further reduced in the barrel sorter which correlates information from all 12 wedges and forwards only the four highest-ranking muon candidates to the GMT.

Cathode Strip Chamber Track Finder

While the principal objective of the DTTF and the CSCTF are similar, there are also fundamental differences. First, the CSCs have to cope with much larger backgrounds from low-p_T muons and other particles. Second, the magnetic field in the endcap region is non-axial, making track calculations more complicated.

The CSC Track Finder organizes the system into 12 sectors, covering $60°$ in ϕ and one endcap each. Each sector processor receives up to 16 muon candidates from the local triggers and reconstructs them intro track candidates. Each such candidate is assigned a 7 bit rank consisting of its p_T and quality assignment. The three highest-ranked candidates from each sector are then sent to the CSC sorter, which extracts the overall four highest ranking candidates and forwards them to the GMT.

Resistive Plate Chamber Trigger

RPC Trigger processing differs from that of the other muon subsystems in that it does not include local triggers for each muon station. The entire readout of the RPC system is transfered to the counting room via zero-suppressed readout over optical links. The trigger view of RPC system is segmented into 33×144 trigger towers in η and ϕ. The RPC Trigger contains a separate pattern comparator (PACT) chip for each of these trigger towers which extracts a muon track when there are coincident hits in at least three of the four RPCs in the tower. Track candidates from towers representing a constant-η ring are then sorted and compared to eliminate candidates from one muon track passing through multiple towers, which leaves the 4 highest-ranked candidate tracks from each η ring. These are again concentrated by two area sorters to produce four tracks for each the barrel and the endcap region, which are then sent to the GMT.

4.2.2 Global Muon Trigger

The Global Muon Trigger (GMT) receives 16 muon candidates from the muon subsystems (4 each from DT, CSC, barrel RPC and endcap RPC) and is tasked with extracting the best four overall muon candidates for the global trigger, exploiting knowledge

Figure 4.5: Organization of the DTTF ϕ view [22].

about the characteristics of different subsystems in order to arrive at an optimal compromise between efficiency and background reduction. The primary means of reduction is the matching of duplicate hits from the overlapping RPC and DT subsystems in the barrel or RPC and CSC subsystems in the endcaps, merging the candidates to provide optimum precision. A programmable match logic assigns a match quality based on the differences in η and ϕ. In addition, the ranks assigned by the subsystem triggers are converted to a global muon rank based on p_T, local quality, η and detector type. The GMT also receives MIP/Quiet bit information from the calorimeter trigger and matches this information with the muon tracks. After matching, the four highest-ranked muons are forwarded to the GT.

4.3 Global Trigger

The GT is the final stage in the Level-1 Trigger process. It compares the physics objects provided by the previous stages according to programmable signatures and decides whether to retain or discard an event. For every bunch crossing, it evaluates up to 128 programmable algorithms over the input objects received from the GCT, GMT and up to 64 "external conditions" delivered as binary inputs from other systems. In addition, it receives up to 64 "technical trigger" signals. The results of the physics algorithms and technical triggers are combined in a final OR (FINOR). Prescale factors and trigger masks (inhibiting the result of that trigger) can be applied to each of the

192 trigger (physics algorithm or technical trigger) bits; veto masks can be used for technical triggers.

4.3.1 Decision Logic

The CMS DAQ subsystem can be split into 8 independent partitions interested in different parts of the event stream. The actual FINOR contains the decision for each of these partitions, forming an 8-bit FINOR word.

The trigger decision (per partition) in the global trigger is highly programmable. The basic building blocks of event selection decisions are predefined conditions which are applied to the input objects. Table 4.1 shows the defined condition types. These conditions, along with 64 external conditions (like presence of beam from the LHC Beam Pick-up Timing Experiment (BPTX) detectors or signals from CASTOR, ZDC and Beam Scintillator Counters (BSC)) can be combined using boolean operators into 128 algorithms.

Up to 64 binary signals from external systems (BPTX, HCAL, CSC, DT, RPC, BSC...) are received directly by GT as technical triggers.

The following factors influence the FINOR word:

- Prescaling: each trigger bit (algorithm or technical trigger bit) is assigned a prescale factor, selecting only a given ratio of the triggers of that type. Prescaling is needed to select samples of very common processes without exceeding the target trigger rate. Actually, at the beginning of the run, sets of prescale factors are defined which are applied sequentially in time to compensate for the decrease of the luminosity during the run (see Sec. 3.1.1).

- Masks: for each partition in the FINOR word the trigger can apply a trigger mask on each trigger bit. This mask causes the result of the corresponding trigger bit to be ignored in the final OR.

- Trigger Veto Mask for the technical trigger bits: if a technical trigger fires and the associated veto bit is set, the event is not selected (forcing the FINOR to 0).

4.3.2 Global Trigger Hardware

The GT system consists of a number of VME electronics boards housed in a VME 9U crate. This crate also contains the GMT and the Trigger Control System (TCS) boards. All boards in the crate are connected through the backplane.

Common Features

All the boards that make up the global trigger contain a VME interface, allowing the online software to configure and monitor the operation of the boards through exposed registers and memories. In addition, the custom chips in the global trigger are all based on FPGA technology and can be reprogrammed with new firmware over the VME interface.

Condition Type	Subtype	Parameters
muon	single muon	η, ϕ, p_T ranges, isolation bits, charge
	2-4 muon	η, ϕ, p_T ranges per muon, isolation bits per muon, charge correlation
	double muon correlation	like double muon, $\Delta\eta$ and $\Delta\phi$ ranges
calorimeter	1-4 objects	object type (e/γ, isolated e/γ, central/forward/τ jet), per-object η, ϕ ranges, E_T threshold
	double object correlation	like double object, $\Delta\eta$ and $\Delta\phi$ ranges
energy sum	ETT condition ETM condition HTT condition HTM condition	total E_T threshold missing E_T threshold, ϕ range total jet E_T threshold missing jet E_T threshold, ϕ range
correlation	-	two single object conditions (of different type) with $\Delta\eta$ and $\Delta\phi$ ranges
HF bit count	-	HF subregion, number of HF towers over threshold in subregion
HF ring energy sum	-	HF subregion index, threshold for E_T in subregion
External condition	-	number of external input bit required

Table 4.1: Basic conditions in the Level-1 global trigger.

Pipeline Synchronizing Buffer Boards

Pipeline Synchronizing Buffer (PSB) boards are used to receive data signals from external systems. Each PSB board features four InfiniBand connectors, each delivering two 1.28 Gbps serial data streams. These are assigned to eight logical channels. In addition, there are 16 RJ45 ports for LVDS parallel data at 80 MHz, which can be assigned in two groups to replace the InfiniBand inputs on channels 0 and 1.

The core functionality of the board is implemented in the SYNC chip, which detects the phases of the incoming signals through oversampling. It applies a programmable delay to each channel. It then writes the data to a ring buffer on the board and sends it over the crate backplane as 4×32 bit 80 MHz Gunned Transceiver Logic (GTL+) signals.

The Readout Processor (ROP) chip waits to receive an L1A signal through the backplane and then copies the data from the PSB's ring buffer to the Global Trigger Frontend (GTFE) board.

In a full setup, the GT contains 3 PSB boards for input from the GCT and another to receive the technical trigger signals. In addition, the transmission of the MIP/Quiet bits from the GCT to the GMT uses another three PSB boards. These are usually named in reference to the board slots they are installed in, for example PSB13 for the PSB in slot 13 of the GT crate.

Global Trigger Logic Board

The Global Trigger Logic (GTL) board executes the trigger logic described in 4.3.1. It receives input signals from the PSBs and the GMT over the backplane. It contains three receiver chips (REC1-3) which translate the 80 MHz backplane signals to the board's 40 MHz clock. Output from the REC chips is then duplicated and fed to two condition chips (COND1 and COND2), which calculate up to 96 algorithms each. This would in principle allow up to 192 algorithms, but only 128 signal lines are actually connected to the outputs. The 128 algorithm bits are sent via parallel flat band cables to the Final Decision Logic (FDL).

Final Decision Logic Board

For every clock cycle, the Final Decision Logic (FDL) board receives 128 bits of algorithm information from the GTL over the flat band cables mentioned above and 64 bits of technical trigger signals from a dedicated PSB board. It calculates the FINOR word as described in 4.3.1 and forwards it to the TCS board.

In addition, the FDL also contains a ROP chip that includes data from the FDL board in the readout in the event of an L1A.

Global Trigger Front End Board

The Global Trigger Front End (GTFE) board links the global trigger to the CMS DAQ system. In the event of an L1A, it assembles readout data from the different

46

boards making up the GT and sends the to the DAQ system. There are two ROP chips on the GTFE board, one for the DAQ system and the other connected to the DAQ Event Manager (EVM). While technically similar, their readout records contain different subsets of the trigger information.

Timing Board

The Timing (TIM) board broadcasts clock signals received from the TTC system. This includes the L1A signals distributed through the TCS that initiate the readout process.

4.4 Trigger Control System

The Trigger Control System (TCS) acts as an intermediate between the GT and the rest of the CMS systems. Upon receipt of a positive FINOR verdict, it distributes the L1 accept signal to the different subsystems through the Trigger Timing and Control system (TTC). Even though the trigger subsystem could in principle operate without dead time, the fact that the readout of an entire event takes significantly longer than a single bunch crossing precludes triggering on each bunch crossing in reality. To avoid overloading the DAQ system, the TCS implements a combination of hard limits on L1A, so-called "trigger rules" (no more than 1 L1A in 3 bunch crossings, 2 in 25, 3 in 100, 4 in 240) and a "backpressure" system, where the TCS is notified when buffer space is running out and suppresses triggers until the situation improves.

4.5 Online Software

The trigger electronics are represented as a single subsystem to the central CMS run control software (RCMS, see section 3.2.10). Since the actual hardware spans many organizations and electronics crates, however, a distributed control system is necessary. All online software used by the subsystems is based on the Trigger Supervisor(TS) C++ framework, which in turn relies on the CMS-wide XDAQ framework.

The basic autonomous unit of the TS framework is an application called a Cell. Each cell provides a web interface featuring both a browser-based user interface and an API using SOAP. TS cells also provide channels to interact with other cells. For integration with the RCMS system, a single central cell represents the entire trigger subsystem.

Subsystem developers specialize their TS cells through three principal extension mechanisms:

- Commands can be specified which take a number of predefined parameters and return a data stream and/or error and warning information. These can be invoked from a generic user interface, through the SOAP API or from inside the cell C++ code.

- Operations provide support for finite state machines. A prime example is the configuration operation, which usually represents the state of the hardware controlled by the cell to the rest of the system.

- Panels can implement more complex graphical user interfaces based on the AJAX paradigm as implemented in the AjaXell library.

Configuration data for the entire CMS detector is stored in a shared Oracle database called the Online Master Database System (OMDS). The TS framework also provides a standardized way to access data from the database through the XDAQ TStore service.

Due to the relative simplicity of its hardware (only a single VME crate), the GT is controlled by a single cell instance. Since the GMT is located in the same VME crate, the GMT Cell runs on the same computer as the GT Cell.

4.5.1 Level-1 Trigger Menu

The programmable algorithm logic running in the GTL card's COND chips is called the Level-1 Trigger menu. These trigger menus are maintained through a graphical TS application, the Level-1 Trigger Menu Editor. Trigger menu specifications can be stored in the OMDS database or an XML format. To actually implement a new trigger menu in the hardware, the algorithms and conditions need to be converted to VHDL code and processed with the FPGA vendor's compiler to produce a firmware image.

These firmware images could in principle be loaded automatically when changing the GT configuration, but this is not currently done because firmware loading is slow compared to the rest of the configuration process and happens infrequently. Instead, the online software only verifies that the loaded firmware version fits the version expected by the configuration.

4.6 Offline Software

The CMSSW software collection contains bit-level emulators of all the trigger components. This is principally required for two reasons:

- Detector simulation requires a simulated trigger to allow realistic physics studies and for the validation of trigger design decisions.

- Running the emulator on recorded data and comparing the calculated outputs with those actually recorded allows detection of errors in data taking.

Due to the modular design of the CMSSW framework, these emulators usually consist of the following subsystems:

- Condition formats: objects describing configuration information for the subsystem, such as the trigger menu and prescaling factors used.

- ESProducers: components that read the condition database or other input sources and produce the condition format objects used by applications.

- Data formats: objects describing event content, for example the trigger decisions or input values for a certain bunch crossing.

- Unpackers: when read out from the DAQ system, the event data simply consists of a stream of 64 bit words for each subsystem. Unpackers implement the logic to decode the subsystem specific format and produce object representations for the event.

- Analyzers: actual emulator logic is then implemented in the form of CMSSW analyzers that use the input objects generated by previous stages to add objects representing the emulator decision to the event.

Figure 4.6 shows this setup for the L1 GT emulator. The `L1GlobalTrigger` class extracts input objects from the event and performs board-by-board emulation of the hardware to produce the output objects, listed in the top right corner of the diagram.

Figure 4.6: Package diagram of the Level-1 global trigger emulator in CMSSW [21]

5 Pattern Testing for the Level-1 Global Trigger

5.1 Requirements

Errors in the L1 GT system can be divided into two major classes: low-level errors resulting from comparatively trivial electronics conditions such as loose connections, cross-talk from neighboring lines or short circuits, and high-level errors introduced through wrong cabling, incorrect configuration or implementation errors in the firmware. It is common for both of these errors to have insidious effects that do not cause obvious malfunctions, but manifest as subtle biases in the functioning of the trigger. Due to the trigger's crucial importance for the functioning of the CMS experiment, it is important that such errors are identified rapidly and reliably. CMS also employs a data quality monitoring (DQM) subsystem that should detect malfunctions. The DQM system's approach relies on monitoring the output data stream during physics data taking. A more preventive approach that allows verification of the system's correctness before data taking starts (especially during development of new features and after hardware changes) is required.

Of course, the development of a system of the GT's complexity would not have been possible in the absence of verification tools. While such testing was in principle understood and possible before the present work, a test of the entire system required considerable expertise about the internals of the hardware and manual execution of multiple disconnected software programs. To summarize, the following problems plagued the existing system:

- Testing was rare. Due to the large effort involved and expertise required, it never became standard procedure to run a full system test regularly, or after minor changes like cabling adjustments. This was especially true for low-level connection problems.

- Since most actual integration testing was performed by hardware developers, a disconnect between the settings used in tests, usually performed at the development crate in Vienna, and the production setup at CERN developed, making it even more difficult to test actual production hardware.

- As another consequence of this disconnect, the test data sets used did not resemble the inputs during physics data taking. This would improve confidence that a positive test result would imply a system according to specification.

The goal of this project was to remedy these problems through the development of new software components, making testing easily accessible and painless to allow its integration into everyday operation of the experiment.

5.2 Implemented Tests

5.2.1 Function Tests

Function tests are a category of tests designed to verify the overall functioning of the GT/GMT system in conditions close to those found during normal operation. The tests simulate all the inputs to the global trigger and compare the results in the form of the global trigger readout records (as sent to the DAQ/EVM systems from the GTFE) with expected values.

Input Simulation

The function test simulates all input to the GT and the GMT. Inputs that are processed through PSB cards can be simulated by enabling simulation memories on the PSB cards. While the inputs for the GT from the GMT also provide a simulation option, the simulation memory in the GMT sorter chip is limited to a size of 1024 bunch crossings. Since the PSB memories as well as the GMT input memories allow the simulation of an entire LHC orbit (3564 bunch crossings), the simulation starts at the GMT inputs (the muon quadruplets from the four regional muon triggers) and runs the GMT alongside the rest of the system, allowing the use of all 3564 bunch crossings for testing.

Input Generation

The input of a function test requires information for all the inputs to the GT as well as the GMT. In addition to information representing the inputs to the system, it also needs reference outputs to verify that the system is operating correctly. This entire information adds up to 2328 bits for each bunch crossing, detailed in Table 5.1. A portable representation of this information existed before this project in the form of the "pattern file" format - a text file representing this information in 119 columns of hexadecimal numbers, with one line per event. Since this format was already established with the existing software infrastructure and contains a fairly good a human-readable representation of the electronics-level view of the system, it was also adopted as the specification format for function tests of the system.

In order to conduct tests with realistic inputs from either physics data or Monte Carlo simulations, the event data therefore needs to be converted to the pattern file format from the CMSSW format that the datasets are originally produced in. This is done using an analyzer framework module called `L1GtPatternGenerator`, which creates the pattern file from digis-level inputs. Most of the information stored in a pattern file is available from the object representations in the exact format needed, so

the required processing is mostly confined to re-ordering the objects from the internal formats (which feature multiple bunch crossing per event) into the pattern file format (one bunch crossing per event). This is done by organizing the internal data in a two-level index hierarchy (event, relative bunch crossing) that is only collapsed when the processing is finished.

In addition, the following extra conversions are needed:

- The HF energy sums and bit counts are not available in an object format that corresponds to the GT input, required for the pattern file. They need to be repacked in the pattern generator.

- In order to detect the difference between a disconnected cable and an all-zero input ("no muon") on the GMT inputs, bit 8-15 of the regional muon candidates are inverted in the electronic representation. Since this is what is represented in the pattern file, the representation needs to be converted.

The flexibility of the CMSSW framework allows the pattern generation to be used unchanged in all the required use cases.

Pattern Generation Process   The general production procedure starts from RAW data files and uses the predefined unpacker sequences to produce the digi-level input objects. Figure 5.1 shows the framework configuration used in this case. The only processing module involved is the pattern generator itself, with the trigger decision taken from the actual hardware output (in case of data).

Re-running the Global Trigger   To test changes to the global trigger's operation, such as the implementation of a new trigger menu or a change in the hardware rules, the global trigger emulator has to be run alongside pattern generation. Even though a specific trigger configuration will be provided by the conditions attached to the input data, the framework es_prefer mechanism can be used to override specific parts. Figure 5.2 shows the framework configuration used to simulate a new trigger menu with existing data.

GT/GMT Readout

The GTFE board is the interface to the DAQ system for the entire global trigger crate, including the TCS and GMT. Consequently, the readout records sent to DAQ and EVM each consist of a number of board payloads. In addition to the bunch crossing that the L1A occurred in, the readout record also contains data for a number of surrounding bunch crossings, which helps in detecting synchronization problems or events spread over multiple bunch crossings. Depending on board firmware and configuration, data from either three or five bunch crossings are contained in a single readout record. The GTFE board's ROP chips contain a spy memory to capture readout records as they are generated. Since the two readout record encapsulate all

Figure 5.1: CMSSW components required for pattern generation from existing data. The swim lanes indicate CMSSW data tiers.

RAW

DIGI

FEDRawDataCollection

GctRawToDigi

L1GlobalTriggerRawToDigi

L1GctJetCandCollection (forJet)

L1GctJetCandCollection (cenJet)

L1GctJetCandCollection (tauJet)

L1GctEmCandCollection (isoEm)

L1GctEmCandCollection (nonIsoEm)

L1GctEtTotalCollection

L1GctEtMissCollection

L1GctEtHadCollection

L1GctHtMissCollection

L1GctHFBitCountsCollection

L1GctHFRingEtSumsCollection

L1MuGMTCandCollection

L1GlobalTrigger

L1GlobalTriggerReadoutRecord

Figure 5.2: Relevant event records and producers to produce an emulated global trigger decision for use in pattern decision. The L1GlobalTriggerReadoutRecord instance produced at the end of this process replaces the one from the third column in Fig. 5.1 as the input to the pattern generator.

Logical name	Physical input	Bits	Content
TTRIG	PSB9 ch 0/1	64	technical trigger signals (via RJ45)
CA1	PSB13 ch 6/7	64	4 isolated e/γ candidates
CA2	PSB13 ch 4/5	64	non-isolated e/γ candidates
CA3	PSB13 ch 2/3	64	central jet candidates
CA4	PSB13 ch 0/1	64	forward jet candidates
CA5	PSB14 ch 6/7	64	τ jet candidates
CA6	PSB14 ch 4/5	64	E_T sums
CA7	PSB14 ch 2/3	64	HF bit counts/HF energy sums
CA8	PSB14 ch 0/1	64	currently unused
CA9	PSB15 ch 2/3	64	currently unused
CA10	PSB15 ch 0/1	64	external condition data (via RJ45)
MQF4	PSB19 ch 6/7	64	48 MIP/Quiet bits for endcap regions
MQF3	PSB19 ch 4/5	64	48 MIP/Quiet bits for endcap regions
MQB2	PSB19 ch 2/3	64	48 MIP/Quiet bits for barrel regions
MQB1	PSB19 ch 0/1	64	48 MIP/Quiet bits for barrel regions
MQF8	PSB20 ch 6/7	64	48 MIP/Quiet bits for endcap regions
MQF7	PSB20 ch 4/5	64	48 MIP/Quiet bits for endcap regions
MQB6	PSB20 ch 2/3	64	48 MIP/Quiet bits for barrel regions
MQB5	PSB20 ch 0/1	64	48 MIP/Quiet bits for barrel regions
MQF12	PSB21 ch 6/7	64	48 MIP/Quiet bits for endcap regions
MQF11	PSB21 ch 4/5	64	48 MIP/Quiet bits for endcap regions
MQB10	PSB21 ch 2/3	64	48 MIP/Quiet bits for barrel regions
MQB9	PSB21 ch 0/1	64	48 MIP/Quiet bits for barrel regions
RPC forward	GMT INA	128	4 regional muon candidates from forward RPCs
CSC	GMT INB	128	4 regional muon candidates from CSC
DT	GMT INC	128	4 regional muon candidates from DT
RPC barrel	GMT IND	128	4 regional muon candidates from barrel RPCs
GMT muons	-	128	
Algo	-	192	128 algorithm decisions
FINOR	-	24	final decision word

Table 5.1: Description of data required for the function tests in one bunch crossings. Rows with a value of - for the input column indicate values that are produced by the GT/GMT system.

information about an event that is available to the higher levels of the system, they provide a useful probe for testing.

One notable missing piece in the readout record is the algorithm bit information as generated by the GTL. Since the GTL has no separate readout channel, the readout record only contains the algorithm bits after FDL processing, which includes application of the prescale factors. This has some adverse consequences for the reprocessing of data.

Execution Modes

There are two subtypes of function tests which use the same input and rely on similar verification methods, but differ slightly in the method of L1A generation.

GT/GMT Function Test   For the GT/GMT function test, the L1A signal that causes the readout is generated through the test trigger mechanism of the TCS. This allows the generation of the L1A at a fixed bunch crossing within each orbit. Since the readout contains at least three bunch crossings, the entire simulated LHC orbit is examined by stepping the test trigger from bunch crossing 1 to 3563 in increments of 3 and extracting the associated readout records for each.

The output data from a GT/GMT function test thus consists of 1188 DAQ readout records and 1188 EVM readout records.

FINOR Test

The GT/GMT function test is defined to mostly bypass the FDL logic - prescaling and masks are set to trivial values, and the FINOR word is intentionally forced to zero - the L1A signal is generated through the test trigger mechanism. In the FINOR test, the L1A is generated through the normal mechanism involving the FDL and TCS boards and the TTC system. While this extends the test coverage, it has the disadvantage of making the output more unpredictable due to the nature of the GTFE spy memory, which can currently capture only one readout record before it needs to be reset. Since the bunch crossing number within the orbit when the test is started is unpredictable, the resulting readout record contains the first event after the start of the test to generate an L1A with the current settings.

Extending this test to the case where multiple events in the orbit would have triggered an L1A in an efficient manner is difficult since repetition of the test is increasingly likely to simply cause a repeat of a former run. For the time being, the FINOR test thus simply runs once and produces one DAQ and EVM readout record each per execution.

Output Verification

The output readout records are analyzed in two different respects:

Identifier	Description
Comparison1	Bunch crossing number of the GTFE board record and central bunch crossing number of all other boards must be equal.
Comparison2	Bunch crossing number and local bunch crossing number must agree on every board record.
Comparison3	Within consecutive records for a single board, the local bunch crossing number must be consecutive.
Comparison4	All board records must share an identical event number.
Comparison5	The relative bunch crossing numbers for every board record must be correctly ordered (E,F,0,1,2 or F,0,1).
Comparison6	Board identifiers for all records must conform to known boards.
Comparison7	The event number must conform to the central bunch crossing number in every record.
Comparison8	The luminosity segment number of the TCS record must match that of the FDL record.
Comparison9	The orbit number of the TCS record must match that of the TCS record.

Table 5.2: A list of the consistency checks applied to the GT readout records produced in pattern tests.

- **Consistency tests** verify the internal consistency of the readout record in a number of respects. For example, they detect internal synchronization problems or format errors in the readout. Table 5.2 lists all implemented consistency checks.

- **Data tests** compare the data contents of the readout record with the expectations from the emulation. Each of the columns[1] from the pattern file is also present in the readout record and can be checked against the original value for correctness.

5.2.2 Interconnection Test

The interconnection test is optimized to detect low-level electronic problems in the GT crate. It sends preprogrammed signal patterns from the input memories to the GTL. The COND chips on the GTL run in a special bypass mode that disables the algorithm logic and simply mirrors a subset of the received inputs to the output bits. The basic unit for testing is a group of 64 input bits corresponding to two GMT

---

[1]with the current exception of the MIP/Quiet bits, which are not currently sent in the production system.

muons[2] or two PSB channels (named CA1-CA10 from Table 5.1 for the inputs from GMT and MU1MU2 and MU3MU4 for the two muon pairs from the GMT). The resulting algorithm bits are read out from an FDL spy memory.

The patterns are designed to highlight common electronics problems:

- "Running zero" test: sets all input bits but one in each 16 bit input word.

- "Running one" test: clears all input bits but one in each 16 bit input word.

- "Counter" test: run all bytes of the input through sequences from 0 to 255.

- "Inversion" test: performs inversion of subgroups from 1 to 32 bits, which is especially sensitive to problems in the power supply of the tested components.

- "Random number" test: complements the pre-defined patterns, in the hope of revealing anomalies missed through the more systematic tests.

Each input is distributed to both of the COND chips, but since only 128 of the 192 output bits of the GTL reach the FDL, only the lower 32 bits of the output from the COND1 chip can be examined. Comparing the output to the expected input then yields the result. Thus, the result of a full interconnection test is organized in a two-level hierarchy according to the input tested and the pattern used. For each of these parameters, it produces a list of mismatches between input values and results.

Interpretation of Results

The detailed routing of signals on the GTL board is shown in Fig. 5.3. Each input is routed over the backplane to exactly one of the three receiver chips (see Table 5.3). This layout implies the following conditions that are helpful for analysis of bit errors:

- Errors common to all input particles are likely caused by either the COND chips or the signal routing to the FDL.

- Static errors between the input simulation and the REC chip appear in triplicate: they affect both lines deserialized from a single 80 MHz line, which are again duplicated to two COND chips, and one of the four instances is truncated because only 32 bits from COND1 are sent to the FDL.

- In contrast, errors between the REC and COND chips will manifest only once per COND chip.

These criteria could in principle be used for automated diagnosis of these error modes, but the added complexity did not appear to be justified because the additional information does not help non-experts to deal with the problems and experts prefer to be confronted with a unified view of errors for arbitrarily complex failure mode.

---

[2]Contrary to the function tests, the GMT is not included in these tests because it lacks a similar bypass mode and the more repetitive manner of the bit patterns negates the disadvantage of the short simulation memory.

Receiver Chip	Logical Input
REC1	MU1MU2
	MU2MU3
	CA1
	CA2
REC2	CA3
	CA4
	CA5
	CA6
REC3	CA7
	CA8
	CA9
	CA10

Table 5.3: Assignment of logical GTL inputs to receiver chips.
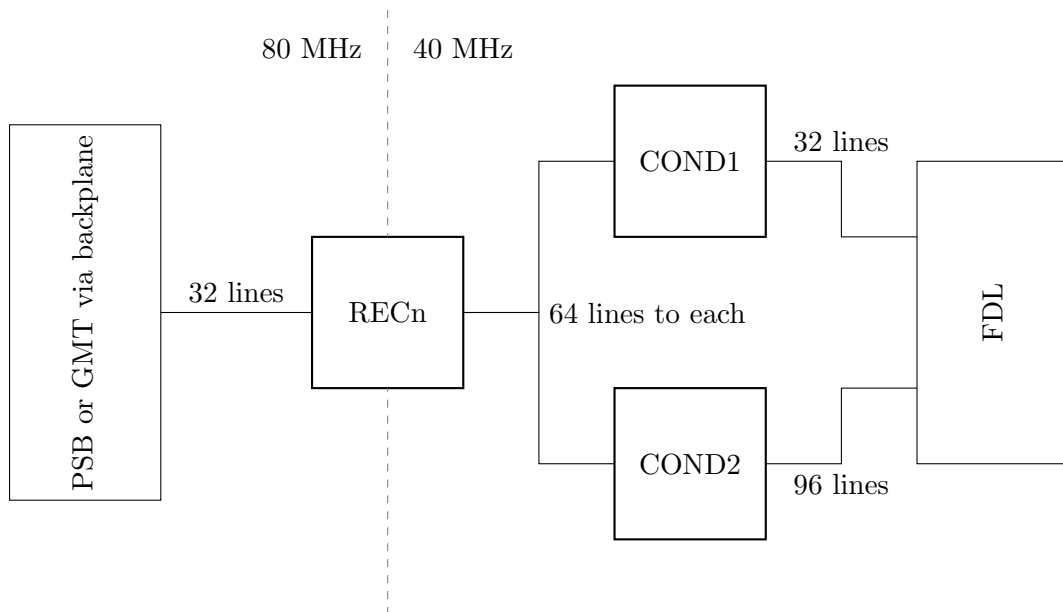
Figure 5.3: Schematic of GTL connections exercised during the interconnection test.

5.3 Software Implementation

The execution of the tests described in the previous section requires access to both the GT and GMT hardware. Since GT and GMT are already represented in the Trigger Supervisor deployment of CMS as two cells running on the machine connected to the GT crate, it was natural to implement the testing system using the same technology. This is especially important for the goal of testing in conditions as close as possible to those of actual operation. Thus, the user interface and logic of the testing process was implemented as another Trigger Supervisor cell called the TestCell, which communicates with the subsystem cells using the TS framework mechanisms, principally by calling commands for hardware access.

5.3.1 User Interaction Concept

One of the key requirements of the TestCell project was to make it very easy to verify correct operation of the system. For this reason, the execution of a test had to be made as foolproof as possible to allow non-experts to run such tests. For the function tests at least, the selection of a correct test configuration is not exactly trivial since the generated inputs need to match the hardware configuration in points such as the scales and trigger menu used. This led to the introduction of the concept of a **test specification**, which contains all the parameters necessary to execute a test. While the creation and maintenance of test specifications is an expert task, the step of **test execution** simply requires the user to select a test specification and press a button in the user interface. The system then executes the tests and provides a summary verdict that makes it easy to identify whether the test has succeeded. If errors have occurred, the application should also help experts analyze the problem. Since the errors are not necessarily reproducible, the application allows users to save a **test result** after execution. These results can then be reviewed by experts to pinpoint the source of the problem.

These three separate concepts result in the basic separation of the user interface into three TS panels - the Test Specification Panel, the Test Execution Panel and Test Results Panel.

5.3.2 Object Model

The object model for the TestCell is designed to be extensible. To this end, the cell logic operates on an abstract object model (shown in Fig. 5.4) where each type of test is represented as a `TestCategory` object. A `TestCategory` provides the central interface needed to specify, execute and analyze tests of that method. The following objects need to be provided for each type of test:

- A `TestSpecification` provides all the input information required to execute a test.

- A `TestOperationBase` instance is a TS operation that implements actual execution of the test. Its finite state machine is shown in Fig. 5.5.

Figure 5.4: Abstract object model for a test category.
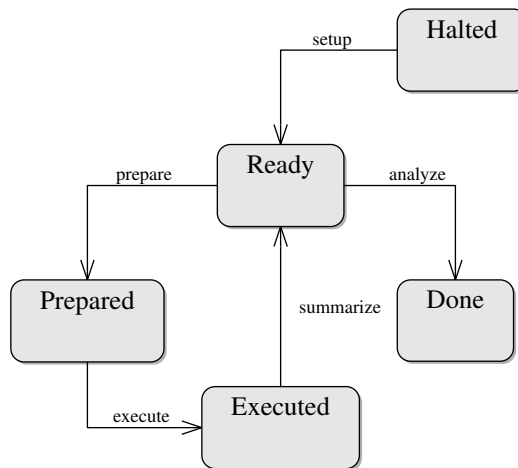
Figure 5.5: Finite State Machine for the generic test operation.
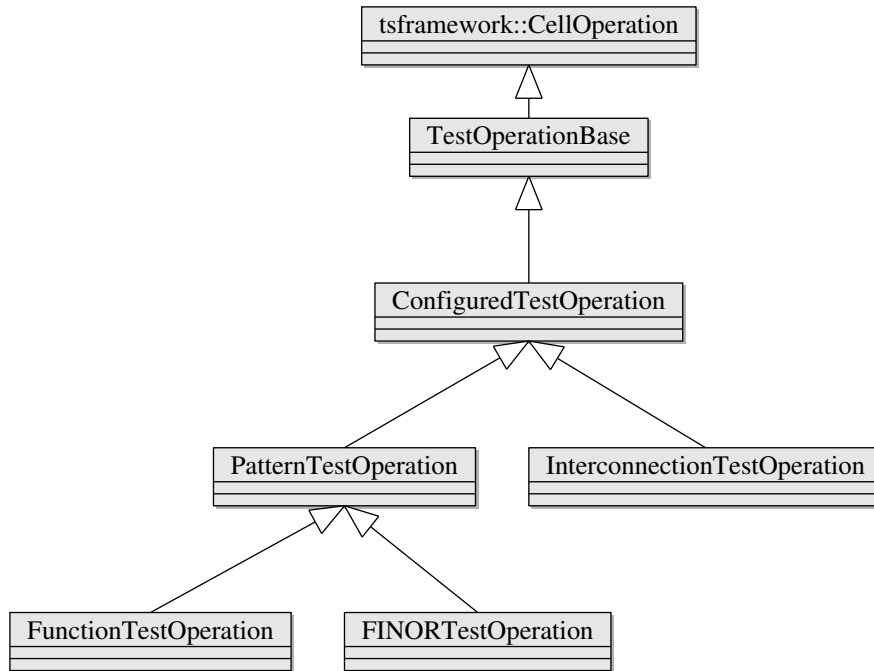
Figure 5.6: Test Operation class hierarchy.

- A `TestResult` contains the entire result information from an executed test instance.

- A `TestSpecificationEditor` is a GUI widget that can be used to create, view or edit `TestSpecification` objects.

- A `TestResultBrowser` is a GUI widget than displays test results.

Each of the three existing tests then provides an implementation for these classes. However, the inheritance hierarchy is not entirely uniform because the functional differences between the pattern test and the FINOR test are rather small. It is therefore not necessary to have separate `TestSpecification` and `TestResult` subclasses, and instead there are common `PatternTestSpecification` and `PatternTestResult` classes for both. In addition, the inheritance hierarchy contains an intermediate layer named `ConfiguredTestOperation` (and so on) for tests that require the hardware to be configured before executing the tests. This is common functionality for all existing tests. Fig. 5.6 shows the inheritance hierarchy for the operation classes, where the entire hierarchy is expressed.

5.3.3 Integration with Existing GT/GMT Cells

The GT and GMT cells already implement commands for hardware configuration from the central database, normally used to set up the hardware from the RCMS

Figure 5.7: Sequence diagram showing the interaction between the different cells for the execution of a GT/GMT function test.

system. Each subsystem maintains its own configuration information, identifying each configuration with a unique key. The global configuration of the trigger system is then defined by a so-called TSC key, which simply aggregates configuration keys for all subsystems involved. The subsystem level of abstraction is also appropriate for use in the TestCell - while testing for data-taking would in principle be better served by directly relying on the global configuration keys in order to eliminate another error source, the creation of a new TSC key requires significant organizational overhead and so is not suitable for the development use cases. Subsystem keys are more easily created because the responsibility is more localized.

Thus, all the existing tests require selection of both a GT and a GMT key from the configuration database. Configuration is performed before every test by sending the respective TS command to the cells (`Configure` for the GT and `GmtConfigureFromDB` for the GMT). Relying on the cells for hardware access also guarantees that tests do not accidentally interfere with data taking since the GT cell refuses to change its configuration while a run is active.

The actual execution of test cycles is implemented in the GT cell, using a different command for each test type (`FinorTest`, `PatternTest`, `InterconnectionTest`). Since simulation data also has to be loaded into the GMT for the function tests, a separate command `GmtLoadPatterns` exists there. After initial experience showed that sending the whole pattern content as command parameters was slow (pattern data for an entire orbit is about 2.5 MB) and all the cells have access to a common file system, it was decided to use the file system for communication of larger data sets (pattern test output and pattern data). This means that the sub-cell commands get file system paths for their input and output files.

When a test is executed in the TestCell, an instance of the associated test operation is created and driven through a predefined set of state transitions. Figure 5.7 shows the interactions in the case of a GT/GMT Function Test.

5.3.4 Persistence

Test specifications and test results need to be stored persistently across cell restarts. Storing this information in the OMDS database was not considered adequate because test results generate large amounts of mostly unstructured data that are better stored in the file system. Consequently, the cell itself implements a simple persistence system called a `DirectoryStorage`. This system manages a collection of named items, each of which contains a single layer of named attributes. This interface is sufficiently general to allow replacement of the underlying storage implementation without affecting the user code.

In the current implementation, each `DirectoryStorage` object manages one underlying file system directory. Items are mapped to subdirectories, and attributes correspond to files in these subdirectories. Since attribute values are stored in text form inside the attribute files, the contents are easily accessible for both human and programmatic editing. The specifics of translating between the `DirectoryStorage` format and the object representation have to be implemented by the appropriate
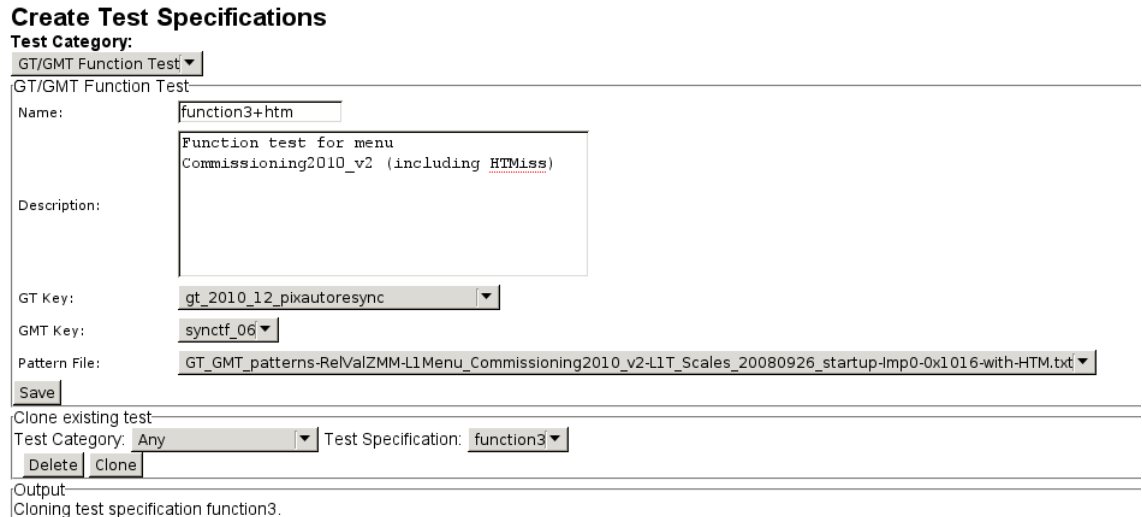
Figure 5.8: Screen shot of the Test Specification Panel, showing the specification for a GT/GMT Function Test.

`TestCategory` instance.

The global interface to the storage system is the `TestDataManager` singleton, accessible through the `CellContext` provided by the TS framework. It maintains separate stores for test specifications, pattern files and test results and manages the conversion of persistence requests at the object level ("save this test specification") to the item/attribute level by forwarding to the appropriate `TestCategory` object.

5.3.5 User Interface

As mentioned in the introduction to this chapter, the user interface is split into three panels according to the separation of use cases for experts and non-experts.

Test Specification Panel

The test specification panel's functionality is limited to the maintenance of the test specification database. The top combo box is used to specify the type of test specification to create. The content of the framed box underneath displays the selected test category's `TestSpecificationEditor` widget, an Interconnection Test in the example. The Save button creates the specification. The second frame contains the elements to delete existing test specifications or create new test specifications based on an existing object. The final frame contains textual feedback about actions performed. Figure 5.8 shows an example of this panel.

Test Execution Panel

The test execution panel is the primary interface for non-expert users. In its initial state, it allows the user to select a test specification, either directly from all selected specifications or filtered to a certain test category. Upon selection of a test specification, the frame underneath the selector presents the parameters for the selected test specification. The "Execute test" button then initiates a test run. The "Test Output" frame then logs the progress of the test while it is running. Upon completion, the "Result" frame receives the summary information from the test. If a test result was obtained, the group of controls next to the "Execute test" button is enabled. "Reset" discards the test result or warnings and returns the panel to the state before test execution. This is required between test executions and prevents accidental information loss through re-execution of the test. The "Details" button enables the detailed error view, as defined by the test category's `TestResultBrowser`. The next two controls are used to save the test result. The label can be edited, but defaults to the name of the test specification with an added timestamp.

Test Results Panel

The test results panel is used to load previously stored test results and display their content. At the top is a test result selector used to choose among the stored test results, optionally filtered according to test category and test specification. Upon loading, the corresponding `TestResultBrowser` is displayed.

Interconnection Test Result Display

The results of an interconnection test are organized according in a hierarchy according to the involved pseudoparticle and test pattern. In the summary view, it displays a short list of pseudoparticles containign failures and the number of patterns containing errors. The detailed view contains a line-by-line summary of the failures, organized in the pseudoparticle-pattern hierarchy. Since test results can become very long, especially for long-running interconnection tests, the output is limited to 200 lines by default and can be scrolled using the arrow buttons.

A simple parser takes the plain text results and highlights the errors. Figure 5.9 shows an excerpt from a typical interconnection test result.

Pattern Test Result Display

The display for pattern test results is organized along the lines of the readout records received from the hardware, since all the tests are defined at the level at the level of the individual readout record level (see Sec. 5.2.1). The result display consists of a tab control with three panes, where one displays error statistics and the other two contain readout record browsers, one for the lists of DAQ and EVM records each.

The readout record browser, shown in Fig. 5.10, provides a navigation interface that allows the user to browse the result collection, either in order or filtered to records

Figure 5.9: ICT result

exhibiting a certain error. The record display contains a textual description of the errors and a readout record display. The readout record is displayed in hex dump format, but contains color highlighting to indicate the boundaries of board records. In addition, clicking on one of the board records displays the fields of that particular record in a table representation. Where applicable, another layer of information is represented by tool tips (for example, GMT records contain table fields for the muon candidates, and these can be decoded to the muon candidate's p_T and other attributes).

Figure 5.10: The pattern test result view's readout record browser, showing details of a readout record with mismatching algorithm decision bits. The lower section displays a GMT board record, and the tooltip shows the details of a regional muon candidate.

6 Test Validation Using Cosmic Events

The TestCell software was deployed to the CMS environment in September 2009, after an initial development period in Vienna. It quickly proved useful in verifying previously known hardware defects and discovering new ones. Since the experiment was at the time taking data from cosmic radiation, the interesting opportunity to look for the consequences of undetected errors exists. This section will discuss some examples and their consequences.

6.1 Data/Emulator Comparison

The primary tool for after-the-fact evaluation of the consequences of electronics errors is the comparison between the results of the hardware - as recorded by the DAQ readout - and a reprocessing of the input values with the GT emulator.

The process setup for the comparison between data and emulator results is very similar to the one used for pattern generation in Sec. 5.2.1. The setup of the CMSSW framework for this purpose is shown in Fig. 6.1. The two readout records obtained from the data file and the emulator are compared by a framework analyzer, which produces a number of histograms on disagreements and errors encountered.

Limits of data/emulator comparison for trigger decisions

While this form of validation is a very powerful tool, it is not perfect. First, there is the obvious bias inherent in looking only at data recorded by the experiment - events that were erroneously missed by the trigger will never make it to the input dataset and will therefore not be noticed. In practice, this is mitigated by the fact that there are usually multiple independent trigger criteria, so that some subsets of the errors will show up because algorithms not affected by the malfunction will still select the event. Second, due to the fact that the readout only contains the algorithm decision word after prescaling has been applied, it is difficult to interpret disagreements on prescaled algorithms. Running the trigger emulator with prescaling disabled allows an error to be flagged when the emulator does not match an emulator and the hardware does, but for the converse, it is impossible to distinguish locally whether there is a malfunction or the prescaling suppressed the bit. Even statistical interpretations are complicated by the interaction of multiple prescalers and trigger masks.

70

Figure 6.1: CMSSW process setup for data/emulator comparison.

Execute Test

Test Category: Any ▼ | Test Specification: Interconnection ▼

Interconnection Test

Name:	Interconnection
Description:	test run with two loops (basic)
GT Key:	gt_2009_35
GMT Key:	synctf_06
Number of loops:	2

Execute test | Reset | Details | Test result label: Interconnection-2010 | Save

Test Output
Executing test specification `Interconnection'.
Configuring GT and GMT...
Warnings from GT configuration:
Loading pattern data...
Executing tests...
Analyzing results...

Results
Ok

Figure 6.2: Result of an Interconnection Test without errors.

6.2 Case Study

The following section reviews the effects of some errors encountered during initial deployment of the TestCell and their traces in the experiment's output. All of the histograms in the following section are based on a sample of 92190 events from a short period of data taking in August 2009.

6.2.1 Baseline Scenario

If the GT is working according to specification, we expect the Function Test, FINOR Test and Interconnection Test to complete without errors. In this case, the information returned by the test is very succinct (see Fig. 6.2 and Fig. 6.3) and only indicates that no errors have occurred. A comparison of recorded data with emulator results should also yield no differences if the comparison is restricted to algorithms that are not prescaled.

6.2.2 Bit Errors in GTL Connections

During data taking, comparing the trigger rates of non-isolated electron triggers with different thresholds to reference values showed an unexpected abundance of highly ranked non-isolated electron triggers. Further investigation of this phenomenon showed that electron candidates arrived correctly at the PSBs from the GCT, but often matched trigger algorithms far above their actual rank. A further analysis revealed a short circuit between the lowest bit of the candidate η and the highest bit of its rank, leading to a significant overvaluation of candidates with odd η. Fig. 6.4 shows the error in a representation of the 16-bit e/γ candidate.

Figure 6.3: Result of a Function Test without errors.

Figure 6.4: Encoding of an e/γ candidate from the GCT. Bit 15 is used for synchronization information. The linked red bits indicate the short circuit encountered.

Figure 6.5: $|\eta|$ distribution of non-isolated e/γ candidates in events from run 112265 matching the `L1_SingleEG5` algorithm.

Since the trigger masks were configured to accept objects with a rank[1] ≥ 10 and the error added 32 to the rank of all candidates, odd-η candidates of arbitrary rank were accepted. Figure 6.5 clearly shows the resulting skewed $|\eta|$ distribution. Comparing the recorded events with the emulator shows a large number of disagreements for electron algorithms, see Fig. 6.6. Figure 6.7 shows the results of an interconnection test with the faulty board.

6.2.3 Missing Muon "Mystery"

The error described in the previous section was traced to a faulty connection on the GTL board, which was then replaced by a spare. A further analysis of the output from the data/emulator comparison revealed some disagreements in the muon algorithms as well (see Fig. 6.8). The number of disagreements for single muon algorithms was small, but the double muon algorithm showed a much higher number of disagreements. Initially, this was puzzling since there were no known problems with the muon part of the system.

At the time, the interconnection test as implemented in the TestCell did not yet exercise the muon inputs. When this functionality was added, it revealed that the entire `MU3MU4` input of the GTL board consistently read as zero, regardless of the inputs sent from the GMT. Upon analysis of the board, this fault was traced to a

---

[1] The numerical value in the name of the trigger algorithm, `L1_SingleEG`, refers to the e/γ E_T, which in this case was converted to hardware units with a linear scale of $E_T = \text{rank}/2\,\text{GeV}$.

Figure 6.6: Disagreements in e/γ algorithms between emulator and data for an analysis of run 112265.

Figure 6.7: Results of the interconnection tests for the `CA2` input exhbiting the error in this section.

Data vs emul: non-matching algorithm decision word for BxInEvent = 0
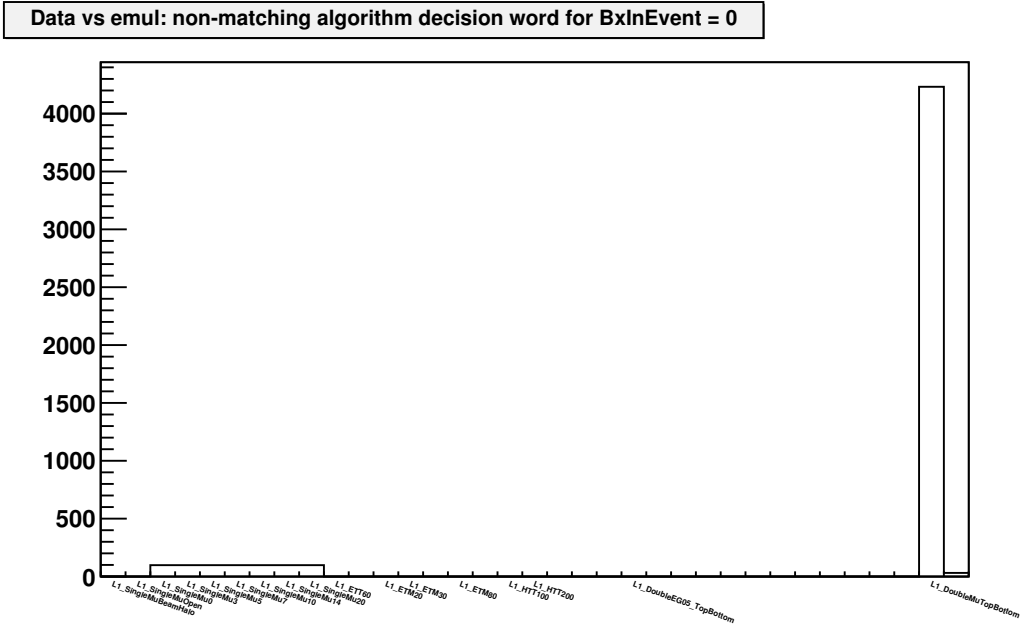
Figure 6.8: Disagreements in muon algorithms between emulator and data for an analysis of run 112265.

disconnected voltage supply pin on the board which left an entire sector of the GTL board without power. As a consequence, at most two muons from the GMT reached the GTL. Fig. 6.9 shows the results of an interconnection test on the affected input.

Since the GMT sorts its output muons according to a calculated rank that generally increases along with p_T, it was comparatively rare for the highest-energetic muon to be sorted to position 3 or 4, which explains the low frequency of the single muon mismatches. Since the double muon trigger did not contain any p_T thresholds, it was far more likely to be affected by this issue. It is particularly interesting that such a rather drastic electronics error can lead to comparatively subtle consequences in the results.

```
Details:
Particle: MU3MU4 ▼  Pattern: RANDOM  ▼
Loop #0, pattern 9, particle: 12 FAILED!

--------------------------------------------- Interconnection Test ---------------------------------------------
Particle: MU3MU4, Pattern: RANDOM
Used COND chip version: V100D

Event |--------------- ALGO in pattern ------------------|   |--------------- ALGO in hardware ----------------|
  #   b191   160 159   128 127    96 95    64 63    32 31     0   b191   160 159   128 127    96 95    64 63    32 31     0
      |     | |     | |     | |     | |     | |     |         |     | |     | |     | |     | |     | |     |
0000  00000000 00000000 01ab84c7 00000000 01ab84c7 01ab84c7   00000000 00000000 00000000 00000000 00000000 00000000
0001  00000000 00000000 00937101 00000000 00937101 00937101   00000000 00000000 00000000 00000000 00000000 00000000
0002  00000000 00000000 02e0e8cb 00000000 02e0e8cb 02e0e8cb   00000000 00000000 00000000 00000000 00000000 00000000
0003  00000000 00000000 02db9750 00000000 02db9750 02db9750   00000000 00000000 00000000 00000000 00000000 00000000
0004  00000000 00000000 007f18ee 00000000 007f18ee 007f18ee   00000000 00000000 00000000 00000000 00000000 00000000
0005  00000000 00000000 0324ca4a 00000000 0324ca4a 0324ca4a   00000000 00000000 00000000 00000000 00000000 00000000
0006  00000000 00000000 00ddf67a 00000000 00ddf67a 00ddf67a   00000000 00000000 00000000 00000000 00000000 00000000
0007  00000000 00000000 00bb6e4a 00000000 00bb6e4a 00bb6e4a   00000000 00000000 00000000 00000000 00000000 00000000
0008  00000000 00000000 00655bc1 00000000 00655bc1 00655bc1   00000000 00000000 00000000 00000000 00000000 00000000
0009  00000000 00000000 01bba1f4 00000000 01bba1f4 01bba1f4   00000000 00000000 00000000 00000000 00000000 00000000
0010  00000000 00000000 035da7e2 00000000 035da7e2 035da7e2   00000000 00000000 00000000 00000000 00000000 00000000
0011  00000000 00000000 01e9b9e2 00000000 01e9b9e2 01e9b9e2   00000000 00000000 00000000 00000000 00000000 00000000
0012  00000000 00000000 0165fbe1 00000000 0165fbe1 0165fbe1   00000000 00000000 00000000 00000000 00000000 00000000
0013  00000000 00000000 038de641 00000000 038de641 038de641   00000000 00000000 00000000 00000000 00000000 00000000
0014  00000000 00000000 0064263d 00000000 0064263d 0064263d   00000000 00000000 00000000 00000000 00000000 00000000
0015  00000000 00000000 01f25512 00000000 01f25512 01f25512   00000000 00000000 00000000 00000000 00000000 00000000
0016  00000000 00000000 0183a1b6 00000000 0183a1b6 0183a1b6   00000000 00000000 00000000 00000000 00000000 00000000
0017  00000000 00000000 03c965f7 00000000 03c965f7 03c965f7   00000000 00000000 00000000 00000000 00000000 00000000
0018  00000000 00000000 00723074 00000000 00723074 00723074   00000000 00000000 00000000 00000000 00000000 00000000
0019  00000000 00000000 013f8a5e 00000000 013f8a5e 013f8a5e   00000000 00000000 00000000 00000000 00000000 00000000
0020  00000000 00000000 0044bacd 00000000 0044bacd 0044bacd   00000000 00000000 00000000 00000000 00000000 00000000
```

Figure 6.9: TestCell output from an interconnection test of the broken muon input of the GTL card described in this section.

7 Conclusion and Outlook

An intergrated testing facility for the Global Trigger of the CMS experiment was presented. After establishing the context of the work in the first chapters, the preceding two chapters describe an implemented testing solution for the CMS Global Trigger and illustrate its application in real-world scenarios. The underlying software has been deployed at CERN and integrated into the experiment's operating procedures.

While the accessibility of testing has been significantly improved, the development has also shown that there is significant opportunity to further improve upon this status. At the time of this writing, the generation of the pattern files described in Sec. 5.2.1 for a specific configuration of the system remains an expert task and needs to be done outside the application. Formalizing the necessary knowledge of the connections between the hardware configuration specified in the online software and the associated configuration for the pattern generation job in the offline software framework so that tests can be performed for arbitrary hardware configurations without expert involvement is the obvious next step. A prerequisite for this step is the consistent availability of this configuration information from a central database, which has only recently been achieved.

Additionally, efforts are underway to create a general framework for interconnection tests between the different trigger subsystems. Integrating this with the current solution would expand the test coverage significantly by including cabling and synchronization issues.

Another interesting direction for future work is suggested by the divergent needs of experiment operation and hardware development - while the test procedures used to verify correct operation of the experiment must be straightforward and leave little room for user error, tests performed during development require the flexibility to override safeguards and change configuration parameters quickly and easily. In the existing solution, priority was given to the former considerations, so that the user interface hides much of the flexibility inherent in the underlying implementation. Developing the current solution into a flexible tool for developers while maintaining the simpler interface for operational tests is an interesting challenge.

Bibliography

[1] Abraham Pais. *Inward Bound*. Oxford University Press, 1986.

[2] The CMS Collaboration. *Detector Performance and Software, Physics Technical Design Report, Volume 1, CERN-LHCC-2006-001, CMS-TDR-008-1*. March 2007.

[3] The LHCb Collaboration. The LHCb detector at the LHC. *Journal of Instrumentation*, 3(08):S08005, 2008.

[4] The TOTEM Collaboration. The TOTEM Experiment at the CERN Large Hadron Collider. *Journal of Instrumentation*, 3(08):S08007, 2008.

[5] The LHCf Collaboration. The LHCf detector at the CERN LHC. *Journal of Instrumentation*, 3(08):S08006, 2008.

[6] The CDF and DZERO Collaborations. Search for Higgs boson production in dilepton and missing energy final states with $5.4 fb^{-1}$ of $p\bar{p}$ collisions at $\sqrt{s} = 1.96$ TeV. *Phys. Rev. Lett.*, 104(6):061804, Feb 2010.

[7] A. Djouadi. The Anatomy of Electroweak Symmetry Breaking Tome 1: The Higgs Boson in the Standard Model. *Physics Reports*, 457(1-4):1–216, 2008.

[8] Frank Wilczek. Decays of Heavy Vector Mesons into Higgs Particles. *Phys. Rev. Lett.*, 39:1304–1306, 1977.

[9] J. Nash. Current and Future Developments in Accelerator Facilities. *IOP 2010*, March 2010.

[10] R. Bailey et al. Synchrotron radiation effect at LEP. *Proceedings of the 6th European Particle Accelerators Conference*, 1998.

[11] O. Bruening et al. *LHC Design Report Volume 1: The LHC Main Ring*. 2004.

[12] R. Bailey and P. Collier. Standard filling schemes for various LHC operation modes. *LHC Project Note 323*, 2003.

[13] P Steinberg and the ATLAS Collaboration. Heavy Ion Physics at the LHC with the ATLAS detector. *Journal of Physics G: Nuclear and Particle Physics*, 34(8):S527–S534, August 2007.

[14] L Sarycheva. Heavy Ion Physics with the CMS Detector at LHC. *Nuclear Physics B - Proceedings Supplements*, 177-178:173–177, March 2008.

[15] Carlo Civinini. The CMS Silicon Strip Tracker. *Nuclear Instruments and Methods in Physics Research A*, 579:726–730, 2007.

[16] The CMS Collaboration. *Physics Performance, Physics Technical Design Report, Volume 2, CERN-LHCC-2006-021, CMS-TDR-008-2.* June 2006.

[17] The CMS Collaboration. *The CMS Electromagnetic Calorimeter Project: Technical Design Report, CERN-LHCC-97-033, CMS-TDR-004.* December 1997.

[18] E Tournefier. The Preshower Detector of CMS at LHC. *Nuclear Instruments and Methods in Physics Research Section A: Accelerators,Spectrometers,Detectors and Associated Equipment*, 461(1-3):355–360, 2001.

[19] The CMS Collaboration. *The CMS Muon Project: Technical Design Report, CERN-LHCC-97-032, CMS-TDR-003.* 1997.

[20] The CMS Collaboration. The CMS Experiment at the CERN LHC. *Journal of Instrumentation*, 3(08):S08004.

[21] V. M. Ghete. The CMS L1 Trigger Emulation Software. Technical Report CMS-CR-2009-123, CERN, Geneva, May 2009.

[22] S. Cittolin et al. *The Trigger Systems, TriDAS Project Technical Design Report, Volume 1, CERN-LHCC-2000-038, CMS-TDR-006-1.* December 2000.

List of Figures

1f8a73a03493467c9789db9c4fdb5d6d06d1657e thesis.tex