Commissioning of the CMS tracker and preparing for early physics at the LHC

Matthew Wingham

Blackett Laboratory
Department of Physics
Imperial College London

A thesis submitted for the degree of
Doctor of Philosophy
of Imperial College London
and the Diploma of Imperial College.

July 2008

Abstract

The Compact Muon Solenoid (CMS) experiment is a general purpose detector at the Large Hadron Collider. It has been designed and optimised to discover the Higgs boson and physics beyond the Standard Model. An early discovery of the Higgs boson is the collaboration's top priority and will require a good understanding of both the detector and the physics of the background processes, with a small integrated luminosity. This principle has been the driving force behind the work presented in this thesis.

The Silicon Strip Tracker (SST) sits at the heart of the CMS detector. The development of core algorithms to commission the SST are reviewed and the process of live commissioning at the Tracker Integration Facility is described. A crowning success of this study is the calibration of 1.6 M channels and their synchronisation to a cosmic muon trigger to within 1 ns.

The SST is expected to produce five times more zero-suppressed data than any other CMS sub-detector. As such its efficient handling within High-Level Trigger algorithms is paramount. The performance of the online hit reconstruction software is profiled, the inefficiencies are characterised and a new scheme to focus on physics regions-of-interest only is proposed. As an example of its success, when running the single τ trigger path over $H^{\pm} \rightarrow \tau^{\pm}\nu_{\tau}$ events, hit reconstruction times were reduced from 838 $\pm$ 5 ms to only 5.13 $\pm$ 0.05 ms without any loss in tracking efficiency. The new software is now the tracker community's permanent online solution and is expected to become the offline solution in the near future.

$b\bar{b}Z^0$ production at the LHC is of great interest, primarily due to its status as a background to a supersymmetric Higgs boson production process. The preparation for a cross section measurement with 100 pb^{-1} of data (expected by the end of 2009) is made. The prominent backgrounds are identified and a signal selection strategy is developed and optimised using Monte Carlo. This study demonstrates that a cross section measurement with this amount of data is feasible. Finally, a method to estimate background from data is tested.

Acknowledgements

Working towards my Ph.D has been the most rewarding and enlightening experience of my life. From day one I have found the Imperial College HEP group a welcoming and challenging place to work. I am very proud to have spend the past four years researching here. The thesis presented is the product of four years hard graft and there are many friends and colleagues who have made it possible. I will attempt to list some of the more important ones below.

Firstly, I would like the thank STFC for providing me with the necessary funding. Geoff Hall also deserves a special thank you for supervising me and reading through this document. I owe much to Rob Bainbridge who took me under his wing in my second and third years, teaching me about the CMS tracker and C++. In fact Chapters 4 and 5 wouldn't have left the starting blocks without him. Also, Sasha Nikitenko gave much guidance through my physics analysis in Chapter 6 for which I am grateful. Much of this work has been collaborative, especially when I was based at CERN. I owe special thanks to Jonny Fulcher, Nick Cripps, Giacomo Bruno, Domenico Giordano and Puneeth Kalavase for their efforts.

Many other people have offered me support throughout the years. There are too many to list in full, but they include all my Ph.D contemporaries; Maiko Takahashi for all her useful analysis tips and being a great friend to share this dusty office with; Seb Greder and Freya Blekman for their ROOT expertise; David Colling and Costas Foudas for their help in my first year and general physics advice; Mark Raymond for putting up with my complete ineptitude when I first arrived and anyone who's given me software advice.

Finally, I would like to thank Mum, Dad, Elise, Chris and my partner Claire. They have each been an inspiration which has kept me motivated through the highs and lows. This thesis is for you.

Contents

List of Figures

List of Tables

Chapter 1

The standard model of particle physics

This chapter begins by describing the mathematical framework within which the Standard Model of particle physics is constructed, quantum field theory. This deals with the time evolution and interaction of canonically quantised fields. Crucially, through this framework, testable calculations can be performed. With the tools in place, the Standard Model can then be constructed by considering symmetry arguments. To include mass in the theory, however, this symmetry will need to be broken. Finally, some unsatisfactory features of the Standard Model, namely inconsistencies upon a second order perturbative expansion, will be addressed. One possible solution to this problem is a Standard Model extension named supersymmetry.

1.1 Quantum field theory

Proposed by Paul Dirac in the late 1920s, Quantum Field Theory (QFT) develops quantum mechanics into a theory of canonically quantised fields operating on a vacuum. This corrects several limitations to quantum mechanics such as the possibility of particle creation/annihilation and the negative probabilities introduced when incorporating special relativity.

Free fermion and boson fields are described by the Dirac (Equation 1.1) and Klein-Gordon (Equation 1.2) Lagrangian densities respectively. Here, γ^μ are the Dirac matrices and m is the particle mass. Instead of the quantum mechanical interpretation of these equations, the wavefunction, ϕ, is now a free field operating on a

vacuum state. It is *quantised* by imposing commutation relations on the field and its conjugate momentum, $\pi(x) = \frac{d\mathcal{L}}{d\dot{\phi}}$.

$$\mathcal{L}_{Dirac} = \bar{\phi}(x)(i\gamma^{\mu}\partial_{\mu} - m)\phi(x) \tag{1.1}$$

$$\mathcal{L}_{KG} = \frac{1}{2}\partial^{\mu}\phi(x)\partial_{\mu}\phi(x) - \frac{1}{2}m^2\phi^2(x) \tag{1.2}$$

The evolution of a field with time can be described by a unitary transformation of the time independent field (or *Schrödinger field*), performed by an operator $e^{-i\hat{H}t}$. Here $\hat{H}$ is the Hamiltonian which can be expressed as a sum of the free field and interaction Hamiltonians, $\hat{H}_0 + \hat{H}_{int}$. To decouple these effects, the evolution is often considered in two separate steps. The evolution with $\hat{H}_0$ is known as the *Dirac field*, the evolution of the Dirac field with $\hat{H}_{int}$ is known as the *Heisenberg field*. The physics of any scattering process is contained in this second step, described by Equation 1.3, where $U(t) = e^{-i\hat{H}_{int}t}$ and ϕ_H and ϕ_D are the Heisenberg and Dirac fields respectively.

$$\phi_H(t) = U^{-1}(t)\,\phi_D(t)\,U(t) \tag{1.3}$$

It is simple to show that $U(t)$ evolves in the manner shown by Equation 1.4.

$$i\frac{d}{dt}U(t) = \hat{H}_{int}(t)\,U(t) \tag{1.4}$$

With the boundary condition $U(t \to -\infty) = 1$, Equation 1.4 yields the solution shown by Equation 1.5.

$$U(t) = Te^{-i\int_{-\infty}^{t}\hat{H}_{int}(t')dt'} \tag{1.5}$$

The *scattering* or S-matrix, which describes the full transformation from outbound Dirac to Heisenberg states, is defined by Equation 1.6.

$$S = U(t \to \infty) \tag{1.6}$$

The transition amplitude from an inbound state at $t \to -\infty$, i, to an outbound state at $t \to +\infty$, f, is described by the *S-matrix element*, S_{fi}, defined in Equation 1.7.

$$S_{fi} = < f, \phi_H(t \to +\infty) | i, \phi_H(t \to -\infty) > \tag{1.7}$$

All perturbative processes in particle physics can be described through the evaluation of the S-matrix element. The *LSZ reduction* form of Equation 1.7, proposed by Lehmann, Symanzik and Zimmermann [1] is shown in Equation 1.8 for an n-particle scattering scenario. The first term describes momentum and energy conservation within the process, the second describes the kinetic energy of each field and the final

term is the *Matrix Element* (ME). $\mathcal{M}_{fi}$ is fully expanded in Equation 1.9 in terms of the Heisenberg fields involved in the interaction and the vacuum state, $|0>$. Here, T indicates the product of fields that follow is time-ordered. $\mathcal{M}_{fi}$ has the form of an n-point Green's function.

$$\begin{aligned} S_{fi} &= i^4 \int d^4x_1 ... d^4x_n e^{-ip_1x_1 ... - ip_nx_n} \\ &\quad \times (\Box_{x_1} + m^2) ... (\Box_{x_n} + m^2) \\ &\quad \times \mathcal{M}_{fi} \end{aligned} \tag{1.8}$$

$$\mathcal{M}_{fi} = < 0|T\phi_H(x_1) ... \phi_H(x_n)|0> \tag{1.9}$$

To perform calculations, $\mathcal{M}_{fi}$ is required in terms of ϕ_D and therefore the form of the interaction Hamiltonian is necessary. As an example, for the interaction of scalar fields the only theory producing finite solutions involves an interaction Hamiltonian of the form shown in Equation 1.10. Here the coupling constant, λ, describes the strength of the interaction between scalar fields.

$$H_{int} = \frac{\lambda}{4!} \int d^3x \phi_D^4(x) \tag{1.10}$$

With this in mind, Equations 1.3, 1.5 and 1.9 can be used to show Equation 1.11.

$$\mathcal{M}_{fi}(x_1 ... x_n) = \frac{\sum_{r=0}^{\infty} \left(-\frac{i\lambda}{4!}\right)^r \frac{1}{r!} < 0|T\phi_D(x_1) ... \phi_D(x_n) \left(\int d^4y : \phi_D^4(y) :\right)^r |0>}{\sum_{r=0}^{\infty} \left(-\frac{i\lambda}{4!}\right)^r \frac{1}{r!} < 0|T \left(\int d^4y : \phi_D^4(y) :\right)^r |0>} \tag{1.11}$$

The time ordered terms in the numerator and denominator can be expanded, through Wick's theorem, into a sum over 2-point Green's functions. Each 2-point Green's function, or *Feynman propagator*, describes a field propagating between two space-time points. This network of bounded states can be elegantly represented through the use of *Feynman diagrams*.

For precise matrix element calculations, solutions to Equation 1.11 with $r > 1$ are required. These can introduce *radiative corrections* (closed loops) in the corresponding Feynman diagram. Each radiative correction to a field demands a corresponding correction to the $r = 1$ *tree* level measurables. Radiative corrections, in turn, introduce diverging field solutions in the region of high momentum and are thus termed ultraviolet (UV) singularities. They can be regulated, however, by imposing an energy scale for the theory, Λ_{UV}.

For further details on QFT some good introductory texts are [2] and [3].

			charge (e)	spin	interaction
ν_e	ν_μ	ν_τ	0	$\frac{1}{2}$	weak
e	μ	τ	-1	$\frac{1}{2}$	EM,weak
u_i	c_i	t_i	$+\frac{2}{3}$	$\frac{1}{2}$	strong,EM,weak
d_i	s_i	b_i	$-\frac{1}{3}$	$\frac{1}{2}$	strong,EM,weak
	H		0	0	-

Table 1.1: The Standard Model fermions and Higgs boson. The index, i, on each quark runs from 0 to 2 and represents its colour.

interaction	mediator	charge (e)	spin	mass (GeV/c^2)
weak	W^+	+1	1	80.4
	W^-	-1	1	80.4
	Z^0	0	1	91.2
EM	γ	0	1	0
strong	8 gluons	0	1	0

Table 1.2: The Standard Model interaction mediators.

1.2 The Standard Model

The Standard Model (SM) is a QFT that describes the elementary particles making up all matter and their interactions (except gravity). To date, it has proven highly successful under experimental scrutiny.

Within the SM there are three types of elementary particle: leptons, quarks and mediators. The leptons and quarks are summarised in Table 1.1. They are all spin $\frac{1}{2}$ fermions, each with a corresponding anti-particle (not shown), and fall into three distinct generations. The one exception is the Higgs boson which has spin 0 and no anti-particle. Quarks are distinguishable from leptons in that they interact via the strong force and come with a *red*, *blue* or *green* strong charge or *colour*. The mediators are summarised in Table 1.2. They are twelve spin 1 bosons, consisting of the photon, γ; the $W^\pm$ and Z^0 and eight gluons which mediate the electromagnetic, weak and strong forces respectively. In total, the SM describes sixty-one elementary particles: thirty-six quarks, twelve leptons, twelve mediators and one Higgs.

At the heart of the SM is the principle of *local gauge invariance* which states that the laws of physics are invariant under certain local phase transformations. This

section will construct the SM Lagrangian density by enforcing this symmetry. The first theory to emerge in this way was quantum electrodynamics.

1.2.1 Quantum electrodynamics

Quantum ElectroDynamics (QED) was developed by a number of physicists in 1929 to describe electromagnetic interactions in terms of particle exchange [4]. It has the structure of an abelian gauge theory, being invariant under local transformations of the U(1) group. Here, Dirac fields transform in the manner shown in Equation 1.12, where $\theta(x)$ is the phase.

$$\phi \rightarrow e^{i\theta(x)}\phi \tag{1.12}$$

Since the free Dirac equation (Equation 1.1) is not invariant under transformations of this form, new terms need to be introduced. This is achieved by exchanging its derivative with a *covariant derivative*, defined by Equation 1.13.

$$D_\mu = \partial_\mu + ig_{em}A_\mu \tag{1.13}$$

Here, g_{em} is known as the *electromagnetic coupling constant*. A new compensating field, A_μ, is necessary to balance the gauge freedom of ϕ. For this reason it is commonly referred to as the *gauge field*. Its corresponding gauge freedom is shown in Equation 1.14.

$$A_\mu \rightarrow A_\mu + \partial_\mu\theta(x) \tag{1.14}$$

The enforcement of QED as invariant under local U(1) transformations has introduced a new field, A_μ, which mediates the electromagnetic force. Interpreting this as the electromagnetic vector potential means the electromagnetic interaction between two particles can be understood in terms of photon exchange.

In this picture, the free gauge field can be described by the Lagrangian density for the electromagnetic field, $\mathcal{L}_{em}$ (itself U(1) invariant). This is shown in Equations 1.15 and 1.16. Here, $F^{\mu\nu}$ is the electromagnetic field tensor.

$$\mathcal{L}_{gauge} = -\frac{1}{4}F_{\mu\nu}F^{\mu\nu} \tag{1.15}$$

$$F_{\mu\nu} = \partial_\mu A_\nu - \partial_\nu A_\mu \tag{1.16}$$

1.2.2 Electroweak theory

The unification of the electromagnetic and weak interactions into a single gauge invariant theory was first performed by Glashow, Weinberg and Salam [5],[6],[7] in the 1960s. It is found that a non-abelian group, SU(2), describes the weak interaction hence a joint theory of electromagnetism and weak interactions obeys $SU(2)_L \times U(1)_Y$* gauge invariance.

Equation 1.17 shows the local field transformations within this theory. The $U(1)_Y$ group generator, Y, is termed *hypercharge*.

$$\phi \rightarrow e^{i\theta^a(x)T^a} e^{iY\omega(x)}\phi \tag{1.17}$$

A new covariant derivative is required to maintain invariance of the theory under these transformations. The form shown in 1.18 achieves this and invokes three new fields, W^a_μ, which are indeed required for weak interactions. g' and g are the U(1) and SU(2) gauge field coupling constants respectively.

$$D_\mu = \partial_\mu - ig^{'}YA_\mu - igW^a_\mu T^a \tag{1.18}$$

As for QED, the electroweak gauge fields can undergo their own gauge transformations to compensate phase changes in ϕ. This gauge freedom is described by Equation 1.19.

$$W^i_\mu \rightarrow W^i_\mu + \frac{1}{2}\epsilon^{ijk}\theta_\mu W^k_\mu \tag{1.19}$$

The free weak gauge field can be described in covariant form by Equations 1.20 and 1.21. Importantly, the last term in Equation 1.21 does not appear in the corresponding equation for the electromagnetic field (Equation 1.16) and is a direct consequence of the non-abelian nature of the SU(2) group. From QFT it can be inferred that the weak gauge field is therefore self-interacting.

$$\mathcal{L}_{gauge} = \frac{1}{4}Tr\left(T^i\mathbb{F}^i_{\mu\nu}T_i\mathbb{F}^{\mu\nu}_i\right) \tag{1.20}$$

$$\mathbb{F}^i_{\mu\nu} = \partial_\mu W^i_\nu - \partial_\nu W^i_\mu + g\epsilon^{ijk}W^j_\mu W^k_\nu \tag{1.21}$$

The physical electroweak fields are defined as linear combinations of the electromagnetic and weak gauge fields, as described by Equations 1.22, 1.23 and 1.24. Here we see how the SU(2) and U(1) gauge theories are *unified* (or at least entangled).

$$W^\pm_\mu = \frac{1}{\sqrt{2}}(W^1_\mu \mp iW^2_\mu) \tag{1.22}$$

*A field, ϕ can be broken into a sum of its left- and right-handed chiral components, $\phi_L + \phi_R$, where $\phi_L = 1/2(1-\gamma_5)\phi$ and $\phi_R = 1/2(1+\gamma_5)\phi$. The subscript L refers to the coupling of SU(2) gauge fields to left-handed chirality states only (from experimental observation). The subscript Y highlights this is not the same U(1) group used in QED, but the U(1) of hypercharge.

$$Z_\mu = A_\mu \sin\theta_\mathrm{w} - W_\mu^3 \cos\theta_\mathrm{w} \tag{1.23}$$

$$\gamma_\mu = A_\mu \cos\theta_\mathrm{w} + W_\mu^3 \sin\theta_\mathrm{w} \tag{1.24}$$

θ_w is the weak mixing angle (Weinberg angle) defined by Equation 1.25.

$$\frac{g'}{g} = \tan\theta_\mathrm{w} \tag{1.25}$$

Due to the parity violating nature of the weak interaction, only left-handed chirality fermion states interact with the SU(2) gauge field, A_μ^a. To achieve this mathematically, the left-handed component of each field, ϕ_L^i, is represented by an SU(2) doublet and the right-handed component, ϕ_R, by a singlet. The full set of electroweak fermion doublets are listed in Table 1.1. They are paired off according to their generation and quark/lepton status. Each fermion has a corresponding right-handed singlet.

The full electroweak Lagrangian density in terms of the gauge fields is shown in 1.26. It describes all electroweak interactions, the most relevant to this thesis being the interaction of the Z^0 field. The latter, a linear combination of A_μ and W_μ^3, couples to the left and right-handed fermion fields with coupling strength $\frac{1}{2}gT^3\cos\theta_w + \frac{1}{2}g'Y\sin\theta_w$ and $\frac{1}{2}g'Y\sin\theta_w$ respectively.

$$\begin{aligned}\mathcal{L}_{EW} \;=\; & i\bar{\phi}_{Li}\gamma^\mu(\partial_\mu + \tfrac{i}{2}gT^aW_\mu^a + \tfrac{i}{2}g'YA_\mu)\phi_L^i \\ & + i\bar{\phi}_R\gamma^\mu(\partial_\mu + \tfrac{i}{2}g'YA_\mu)\phi_R \\ & - \tfrac{1}{4}\mathbb{F}_{\mu\nu}\mathbb{F}^{\mu\nu} - \tfrac{1}{4}F_{\mu\nu}F^{\mu\nu}\end{aligned} \tag{1.26}$$

No mass terms exist in the electroweak Lagrangian so far. Neither a fermion mass term of the form $m\bar{\phi}\phi = m\left(\bar{\phi}_R\phi_L + \bar{\phi}_L\phi_R\right)$, nor a boson mass term of $m^2A_\mu A^\mu$ is invariant under the electroweak gauge. Though experiment supports this prediction for the photon, all other SM particles do require a mass. For this reason electroweak symmetry must be broken.

1.2.3 Quantum chromodynamics

Quantum ChromoDynamics (QCD) is the theory of strong interactions. It is formulated in an analogous way to electroweak theory, by enforcing local gauge invariance on the free Dirac Lagrangian (Equation 1.1). However, the $\mathrm{SU(2)}_L \times \mathrm{U(1)}_Y$ group is now replaced by the non-abelian group SU(3), such that the local transformation of the fermion fields takes the form shown in Equation 1.27. T_a are the eight *Gell-Mann* matrices representing the generators of the SU(3) group.

$$\phi \rightarrow e^{i\theta_a(x)T_a}\phi \tag{1.27}$$

The covariant derivative required to maintain the invariance of the theory is shown in Equation 1.28. Here, G^a_μ corresponds to 8 new gauge fields which mediate the strong interaction via gluon exchange.

$$D_\mu = \partial_\mu + ig_{strong}T_aG^a_\mu \tag{1.28}$$

The gauge freedom of G^a_μ required to compensate the local phase transformations of ϕ is shown in Equation 1.29.

$$G^a_\mu \rightarrow G^a_\mu - \frac{1}{\phi}\partial_\mu\theta_a - f_{abc}\theta_aG^c_\mu \tag{1.29}$$

Here, f_{abc} is the matrix of *structure constants* which is defined by Equation 1.30 in terms of the commutation relations between SU(3) generators.

$$[T_a, T_b] = f_{abc}T_c \tag{1.30}$$

The gauge invariant Lagrangian density for the free gauge field is described by Equations 1.31 and 1.32. As for the weak fields, the strong field self-interacts due to the non-abelian nature of the SU(3) group.

$$\mathcal{L}_{gauge} = \frac{1}{4}\mathbb{G}^a_{\mu\nu}\mathbb{G}^{\mu\nu}_a \tag{1.31}$$

$$\mathbb{G}^a_{\mu\nu} = \partial_\mu G^a_\nu - \partial_\nu G^a_\nu - g_{strong}G^b_\mu G^c_\nu \tag{1.32}$$

Each quark field, ϕ_i, is now a triplet where the index i runs from 0 to 2 and corresponds to the quark colour. Since the electroweak singlets and doublets do not transform under SU(3) they do not interact strongly. The full QCD Lagrangian density, shown in Equation 1.33, describes all strong interactions.

$$\mathcal{L}_{QCD} = i\bar{\phi}^i\gamma^\mu\partial_\mu\phi_i - m\bar{\phi}^i\phi_i - g_{strong}\bar{\phi}^i\gamma^\mu T_aG^a_\mu\phi_i - \frac{1}{4}\mathbb{G}^a_{\mu\nu}\mathbb{G}^{\mu\nu}_a \tag{1.33}$$

1.2.4 The Higgs mechanism

The Higgs mechanism, originally proposed by Peter Higgs in the 1960s [8], provides a way of introducing a mass term for the weak bosons whilst leaving the photon massless.

Consider the complex field doublet shown by Equation 1.34, also known as the *Higgs doublet*.

$$\Phi = \begin{pmatrix} \phi^+ \\ \phi^o \end{pmatrix} = \begin{pmatrix} \phi_1 + i\phi_2 \\ \phi_3 + i\phi_4 \end{pmatrix} \tag{1.34}$$

The Lagrangian density for this field is shown by Equation 1.35. Here D_μ is the electroweak covariant derivative defined by Equation 1.18. The first term describes the field's kinetic energy, the second is the field's vacuum-potential which we assume has the form described by Equation 1.36.

$$\mathcal{L}_{Higgs} = (D_\mu \Phi)^\dagger (D^\mu \Phi) - V_{Higgs}\left(\Phi^\dagger \Phi\right) \tag{1.35}$$

$$V_{Higgs} = \mu^2 \Phi^\dagger \Phi + \lambda \left(\Phi^\dagger \Phi\right)^2 \tag{1.36}$$

For a Higgs doublet $\Phi = \begin{pmatrix} 0 \\ 0 \end{pmatrix}$, $\mathcal{L}_{Higgs}$ is invariant under $\mathrm{SU}(2)_L \times \mathrm{U}(1)_Y$ gauge transformations. However, this state is not favoured energetically. The vacuum expectation value for Φ will occur at the potential minimum where $\frac{\partial V}{\partial\left(\Phi^\dagger \Phi\right)} = 0$. This state is shown by Equations 1.37 and 1.38. Unlike for other SM fields it is non-zero and is what breaks electroweak gauge invariance. This is sometimes referred to as *spontaneous symmetry breaking* and is a fundamental concept in the physics of phase transitions.

$$\Phi_{vacuum} = \frac{1}{\sqrt{2}} \begin{pmatrix} 0 \\ v \end{pmatrix} \tag{1.37}$$

$$v = \sqrt{\frac{-\mu^2}{\lambda}} \tag{1.38}$$

The proposed Higgs doublet appears to predict four new scalar bosons. The *Higgs boson*, H^0, corresponds to the field which breaks the electroweak gauge symmetry. This is a massive field due to a mass term introduced with V_{Higgs}.

Three *Goldstone bosons* [9] also seem to exist, corresponding to the three remaining degrees of freedom in the Higgs doublet. These are massless as their vacuum potential is zero. Since no massless bosons beyond the photon have been observed by experiment however, they are instead interpreted as mass terms for the bosons mediating the weak interaction.

To show this explicitly, we must first consider the general form of the Higgs doublet after symmetry breaking. This is shown in Equation 1.39, where a term for the excitations in the Higgs boson field, η, has been included.

$$\Phi = \frac{1}{\sqrt{2}} \begin{pmatrix} 0 \\ \eta + v \end{pmatrix} \tag{1.39}$$

If Equations 1.38 and 1.39 are substituted into the expression for $\mathcal{L}_{Higgs}$, and the covariant derivative is expanded (with hypercharge, $Y = 1$), the Higgs Lagrangian density after symmetry breaking is obtained. This is shown in terms of the physical

weak fields in Equation 1.40. Here, the first line describes the free Higgs boson field and the first terms of lines two, three and four are the mass terms for the H^0, $W^{\pm}$ and Z^0 fields respectively. The higher order terms of line two predict 3 and 4 point Higgs boson self-interactions and the second and third terms of lines three and four predict Higgs interactions with the weak boson fields.

$$\begin{aligned}\mathcal{L}_{Higgs} &= \tfrac{1}{2}\partial_\mu\eta\partial^\mu\eta \\ &\quad -2\mu^2\eta^2 + h.o.t. \\ &\quad +\tfrac{g^2\mu^2}{4\lambda}W^{+\mu}W^-_\mu - \tfrac{g^2\mu}{2\sqrt{\lambda}}\eta W^{+\mu}W^-_\mu + \tfrac{g^2}{4}\eta^2 W^{+\mu}W^-_\mu \\ &\quad +\tfrac{g^2\mu^2}{8\lambda\cos\theta_{\mathrm{w}}{}^2}Z^{0\mu}Z^0_\mu - \tfrac{g^2\mu}{4\sqrt{\lambda}\cos\theta_{\mathrm{w}}{}^2}\eta Z^{0\mu}Z^0_\mu + \tfrac{g^2}{8\cos\theta_{\mathrm{w}}{}^2}\eta^2 Z^{0\mu}Z^0_\mu\end{aligned} \tag{1.40}$$

The mass terms in Equation 1.40 yield the expressions for the SM boson masses shown in Equations 1.41 ,1.42 and 1.43.

$$m_{W^\pm} = \sqrt{\frac{g^2\mu^2}{4\lambda}} \tag{1.41}$$

$$m_{Z^0} = \frac{m_{W^\pm}}{\cos\theta_{\mathrm{w}}} \tag{1.42}$$

$$m_{H^0} = -2\mu^2 \tag{1.43}$$

Extra terms are required in the SM Lagrangian density to account for quark and lepton masses. The *Yukawa* Lagrangian density, $\mathcal{L}_{Yukawa}$, shown in Equation 1.44 is chosen such that it is gauge invariant. The two terms on the right-hand side give mass to the down type leptons and quarks and up type quarks respectively. Here, ϵ_{ij} is the two dimensional antisymmetric tensor, ϕ is a fermion field, q is a quark field and G_{f_d} and G_{f_u} are field-specific constants.

$$\mathcal{L}_{Yukawa} = G_{f_d}\left(\bar{\phi}^i_L\Phi_i\phi_{Rd} + h.c.\right) + G_{f_u}\left(\epsilon_{ij}\bar{q}^i_L\Phi^{*j}q_{Ru} + h.c.\right) \tag{1.44}$$

The form of $\mathcal{L}_{Yukawa}$ after electroweak symmetry breaking is found by substituting 1.39 into Equation 1.44. This is shown in Equation 1.45.

$$\mathcal{L}_{Yukawa} = -G_{f_d}\sqrt{\frac{\mu^2}{2\lambda}}\phi_d\bar{\phi}_d - G_{f_u}\sqrt{\frac{\mu^2}{2\lambda}}q_u\bar{q}_u + \frac{G_{f_d}}{\sqrt{2}}\eta\phi_d\bar{\phi}_d + \frac{G_{f_u}}{\sqrt{2}}\eta q_u\bar{q}_u \tag{1.45}$$

Terms one and two are clearly the ϕ_d and q_u masses. The general form for the fermion masses in terms of the V_{Higgs} parameters and G_f, is given in Equation 1.46. The third and fourth terms predict the Higgs boson field coupling to the fermion fields. The coupling strength is proportional to G_f and hence the fermion mass.

$$m_f = -G_f\sqrt{\frac{\mu^2}{2\lambda}} \tag{1.46}$$

1.2.5 The Standard Model Lagrangian

The SM can therefore be described as a broken SU(3) × SU(2)$_L$ × U(1)$_Y$ symmetry. The complete Lagrangian density, $\mathcal{L}_{SM}$ (Equation 1.47), is the sum of all contributions from the Electroweak, QCD, Higgs and Yukawa sectors, defined by Equations 1.26, 1.33, 1.35 and 1.44 respectively.

$$\mathcal{L}_{SM} = \mathcal{L}_{EW} + \mathcal{L}_{QCD} + \mathcal{L}_{Higgs} + \mathcal{L}_{Yukawa} \tag{1.47}$$

1.2.6 Constraints on the Higgs boson mass

Though Equation 1.43 shows the Higgs boson mass is entirely dependent on a free parameter, μ, there are theoretical arguments to constrain it. To maintain a perturbative weak interaction and unitarity within gauge boson scattering processes it can be shown that the relation $m_H < 1$ TeV/c^2 must be satisfied [10].

The tightest limits, however, come from experimental data. From direct searches, the Large Electron Positron collider (LEP), at CERN, reported m$_H > 114.4$ GeV/c^2, at a 95% confidence level [11]. $W^{\pm}$ and Z^0 boson mass measurements can be used to further constrain the Higgs mass indirectly. Figure 1.1 shows the dominant radiative corrections to $W^{\pm}$ and Z^0 free fields. Precise measurement of the $W^{\pm}$, Z^0, t and b masses can therefore provide evidence for the Higgs boson mass itself.

Figure 1.2 shows the variation of the Higgs boson mass with the $W^{\pm}$ and t masses which have the largest experimental uncertainties of 80.398 ± 0.025 GeV/c^2 [12] and 170.9 ± 1.8 GeV/c^2 [13] respectively. t mass measurements suggest that a SM Higgs boson would not be observed within a mass range significantly above the 114.4 GeV/c^2 lower limit.

Figure 1.3 shows a $\Delta\chi^2$ fit of the SM to existing electroweak data as a function of m$_H$. It gives an upper mass limit of 144 GeV/c^2 at a 95% confidence level. When including the 114.4 GeV/c^2 lower limit from direct searches, then this upper limit increases to 182 GeV/c^2.

1.2.7 Radiative corrections to the Higgs field

A free Higgs boson will itself undergo radiative corrections. Figure 1.4 (left) shows an example of a second order correction from a fermion loop. The corresponding correction to m_H , Δm_H, is shown by Equation 1.48 [14].

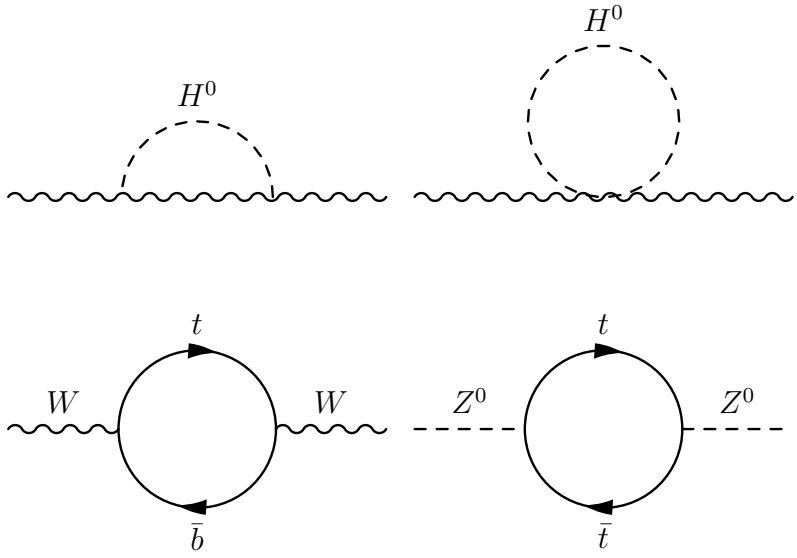

Figure 1.1: Feynman diagrams showing dominant loop corrections to free $W^{\pm}$ and Z^0 boson fields.

$$\Delta m_H^2 = \frac{G_f^2}{16\pi^2}[-2\Lambda_{UV}^2 + 6m_f^2 \ln\left(\frac{\Lambda_{UV}}{m_f}\right)] \tag{1.48}$$

The conclusions drawn from these calculations depend very much on the chosen energy scale, Λ_{UV}. If Λ_{UV} is chosen to be the Planck mass ($M_P = (8\pi G)^{-\frac{1}{2}} = 2.4 \times 10^{18}$ GeV/c^2), then Δm_H is fifteen orders of magnitude greater than the 1 TeV/c^2 theoretical limit. If Λ_{UV} is chosen to be small, new physics above this must still cut the diverging loop integral through some mechanism. This represents a problem for the whole SM mass spectrum, which is defined entirely by the Higgs boson.

Further radiative contributions to m_H can come from the effects of arbitrarily heavy complex scalar particles which may exist beyond the SM (Figure 1.4 right). These would take the form shown in Equation 1.49 [14].

$$\Delta m_H^2 = \frac{G_b^2}{16\pi^2}[\Lambda_{UV}^2 - 2m_b^2 \ln\left(\frac{\Lambda_{UV}}{m_b}\right)] \tag{1.49}$$

A cancellation of the terms in Equations 1.48 and 1.49 could serve to remove the Higgs mass divergences (though higher order diagrams also need to be considered). For the terms to exactly cancel, a symmetry between bosons and fermions would need to exist. One theory with such a feature is *supersymmetry*.

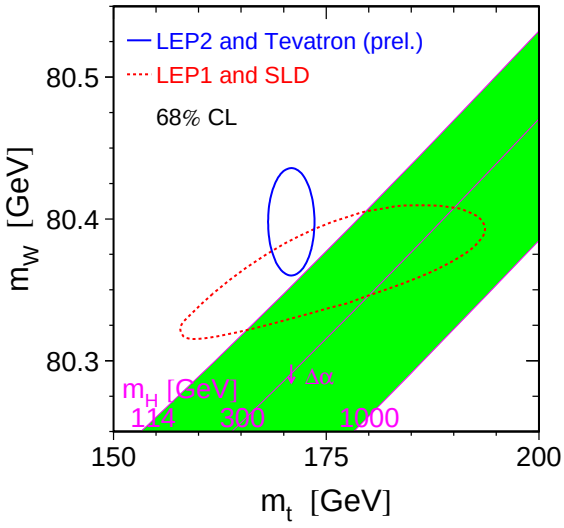

Figure 1.2: Indirect SM Higgs boson mass measurement from $W^{\pm}$ [12] and t masses [13]. The region below 114 GeV/c^2 is excluded from direct LEP searches. t mass measurements at Tevatron suggest the mass would not be observed significantly above this.

This section aimed to provide a brief introduction to the SM tailored to the physics interests of the Large Hadron Collider. A more comprehensive description can be found in any of the following sources [15], [16] and [17].

1.3 Beyond the Standard Model : supersymmetry

SUperSymmetrY (SUSY) is a theory relating fermions with bosons through an operator, Q, shown by Equations 1.50 and 1.51. Q obeys the commutation and anti-commutation relations shown in Equations 1.52 and 1.53. Here, P^{μ} is the space-time translation generator.

$$Q|boson\rangle = |fermion\rangle \tag{1.50}$$

$$Q|fermion\rangle = |boson\rangle \tag{1.51}$$

$$[P^{\mu}, Q] = [P^{\mu}, Q^{\dagger}] = 0 \tag{1.52}$$

$$\{Q_{\alpha}, Q^{\dagger}_{\beta}\} = 2\sigma^{\mu}_{\alpha\beta}P^{\mu} \tag{1.53}$$

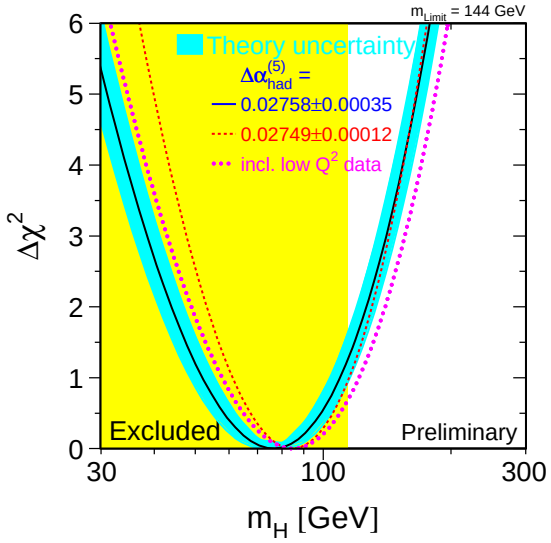

Figure 1.3: A $\Delta\chi^2$ fit of the SM to existing electroweak data as a function of Higgs boson mass [12]. The region in yellow has been excluded using data from LEP2. The blue band represents an estimate of the theoretical error due to missing higher order corrections.

Fermions and bosons paired off in this way are called *superpartners*. Since Q and $Q^\dagger$ commute with the mass operator (P^2) and the SM gauge generators, the superpartners must have equal mass, electric charge, weak isospin and colour.

Particles of this nature would be very simple to detect so a symmetry of this form must be broken. Though the breaking mechanism is not fully understood, if divergences in the Higgs boson mass are to be controlled the relationships between coupling constants must hold.

1.3.1 The Minimal Supersymmetric Model

The Minimal SuperSymmetric Model (MSSM) aims to minimise the number of fields and interactions within SUSY. Within this theory, every SM particle has a corresponding superpartner. Two Higgs doublets, $\Phi_u = (\phi_u^+, \phi_u^0)$ and $\Phi_d = \left(\phi_d^0, \phi_d^-\right)$, are now required to give mass to the up and down type fermions respectively. The quadratic sum of their vacuum expectation values can be related to the W-mass in the form shown in Equation 1.54. Although the ratio of these values, $tan\beta$ (Equation 1.55), is not directly predicted the requirement that fermion masses do not diverge

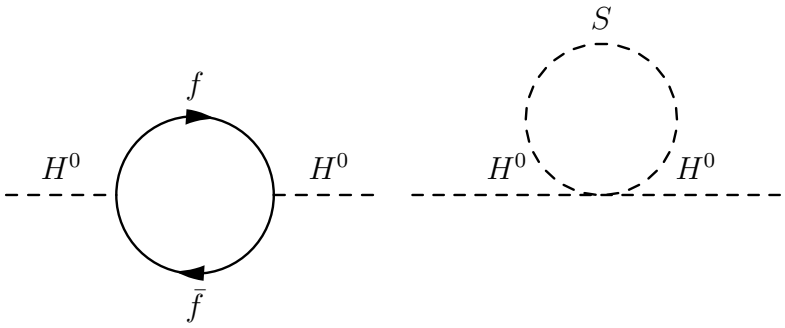

Figure 1.4: Feynman diagrams for second order corrections to the Higgs boson mass. Here, f is a fermion and S a boson.

leads to the constraint $1.2 \lesssim \tan\beta \lesssim 60$.

$$v_u^2 + v_d^2 = \frac{2m_W^2}{g_{weak}} \tag{1.54}$$

$$\frac{v_u}{v_d} \equiv \tan\beta \tag{1.55}$$

Within the MSSM, after electroweak symmetry breaking three degrees of freedom disappear leaving five Higgs bosons defined as linear combinations of the doublet fields: a CP odd neutral scalar, A (Equation 1.56); two CP even neutral scalars h and H (Equation 1.57) and charged scalars $H^{\pm}$ (Equation 1.58). Here, both α and β are mixing angles.

$$A = \sqrt{2}\left(\cos\beta Im[\phi_u^0] + \sin\beta Im[\phi_d^0]\right) \tag{1.56}$$

$$\begin{pmatrix} h \\ H \end{pmatrix} = \sqrt{2}\begin{pmatrix} \cos\alpha & -\sin\alpha \\ \sin\alpha & \cos\alpha \end{pmatrix}\begin{pmatrix} Re[\phi_u^0] - v_u \\ Re[\phi_d^0] - v_d \end{pmatrix} \tag{1.57}$$

$$H^{\pm} = \cos\beta\phi_u^{\pm} + \sin\beta\phi_d^{\mp} \tag{1.58}$$

For further details on supersymmetry, good overviews can be found in [18], [19] and [14].

Chapter 2

Introduction to the LHC and CMS

The Large Hadron Collider, a proton - proton collider based at CERN, Switzerland and the Compact Muon Solenoid, one of its general purpose detectors, are discussed. The Compact Muon Solenoid detector is described in detail, including an overview of each sub-detector, its purpose, technology and resolution. More emphasis is placed on the silicon strip tracker since it is the focus of much of this thesis. Finally, the data acquisition and triggering systems are reviewed.

2.1 The LHC

The Large Hadron Collider (LHC) [20], [21] is a proton - proton (p-p) and lead-ion (Pb-Pb) collider, based at CERN across the Swiss - French border. Particles will collide in bunches at four nominal interaction points, where two general purpose detectors: Compact Muon Solenoid (CMS) [22], [23] and ATLAS [24] and two dedicated detectors: LHCb [25] (for studying B physics) and ALICE [26] (for studying quark-gluon plasmas) will be positioned to record the product of each physics *event*.

The particle beam reaches its final energy and structure via four stages. Linear accelerators LINAC2 and LINAC3 are used for the initial acceleration of protons and lead ions, respectively. 26 GeV particle bunches, circulating in their final beam structure, are then formed in the Proton Synchrotron (PS). The Super Proton Synchrotron (SPS) pre-accelerates the beam to 450 GeV and injects it into the LHC where they will reach their nominal energy of 7 TeV (for the proton beam). Twelve injections are performed for each counter-rotating beam. Approximately 4 SPS cycles are necessary to fill the LHC taking $\sim$3 minutes.

beam parameter	p-p high	p-p low	Pb-Pb high	units
luminosity	10^{34}	2×10^{33}	10^{27}	$cm^{-2}s^{-1}$
energy per nucleon	7	7	2.76	TeV
bunch spacing	25	25	99.8	ns
particles per bunch	1.15×10^{11}	4×10^{10}	7.0×10^{7}	-
inelastic collisions per crossing	17.5	3.5	7	-

Table 2.1: Beam parameters for low and high luminosity p-p and high luminosity Pb-Pb collisions.

The LHC itself consists of two beam pipes passing through 1232 dipole magnets, with RF cavities to provide a *kick* and increase the proton energy by 0.5 MeV per revolution. To achieve collision conditions, each beam is focused by a complex array of magnets before they cross at every interaction point. On collision, bunches are highly elongated, with a diameter of $\sim$16 μm and length $\sim$8 cm.

The bunch structure of the LHC is fairly complicated. A bunch separation of 25 ns is maintained with trains of 72 occupied and 12 empty bunches. A longer gap of 950 ns is allocated every 3-4 trains for LHC injections. Other large gaps between trains exist for synchronisation, electronic resets and obtaining calibration data. Of the 3564 bunch spaces available during each cycle, 2808 are filled.

The experiment is scheduled to start taking p-p data in late 2008. After an initial physics run at a centre-of-mass energy of 10 TeV, it will reach its *low luminosity* running phase during 2009. Low luminosity running will provide 20 fb^{-1} of data per year. After the first year at least 100 pb^{-1} of 14 TeV data are expected. The LHC is designed to reach *high luminosity* in 2010. These conditions will provide an integrated luminosity of 100 fb^{-1} per year. Pb-Pb collisions are expected in 2009 reaching high luminosity after a year. Beam parameters for all three configurations are given in Table 2.1.

2.2 CMS physics motivation

2.2.1 Higgs physics searches

A key feature of the Standard Model of particle physics is the existence of the Higgs boson. The detection of this particle was a primary design consideration for CMS. According to the SM, Higgs boson production will proceed via four main processes at

the LHC: gluon fusion ($gg \rightarrow H$, via a top-quark loop), top-quark fusion ($gg \rightarrow Ht\bar{t}$), weak boson fusion ($qq \rightarrow Hqq$) and weak boson bremsstrahlung ($q\bar{q}' \rightarrow HW$). Figure 2.1 shows that gluon fusion will be dominant for Higgs boson masses up to about 700 GeV/c^2. Weak boson fusion becomes significant for a higher Higgs mass. This process involves central, high p_T hadronic jets in the final state, a distinct feature that can be used to suppress backgrounds.

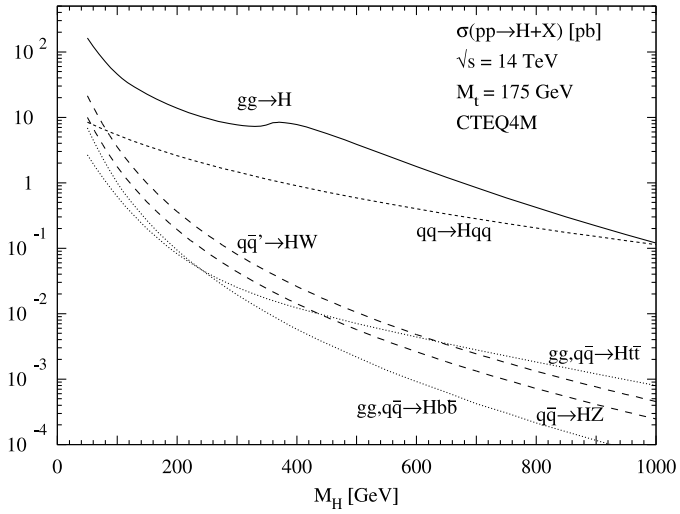

Figure 2.1: Higgs boson production channel cross sections as a function of mass [27].

The branching ratios of the most likely Higgs boson decay channels are shown in Figure 2.2. In the *low* mass range ($m_H < 114$ GeV/c^2 i.e. below the LEP mass lower-limit), hadronic decays dominate. These channels are difficult to use for discovery due to large QCD backgrounds and the poor energy resolution available for hadronic jets. Instead, decays into lepton or photons are preferable, despite the smaller branching ratios.

In the *intermediate* mass range (114 GeV/c^2 $< m_H < 2\ m_Z$), the signal resolution will be dominated by the detector resolution since the natural width of the Higgs boson is just a few MeV. This is clearly shown by Figure 2.3. Here, the two photon or ZZ* decay modes are the most promising and the detection of the decay products, muons and photons in particular, will become crucial. The *high* mass range ($m_H > 2\ m_Z$) yields the ZZ* decay as the mode of choice.

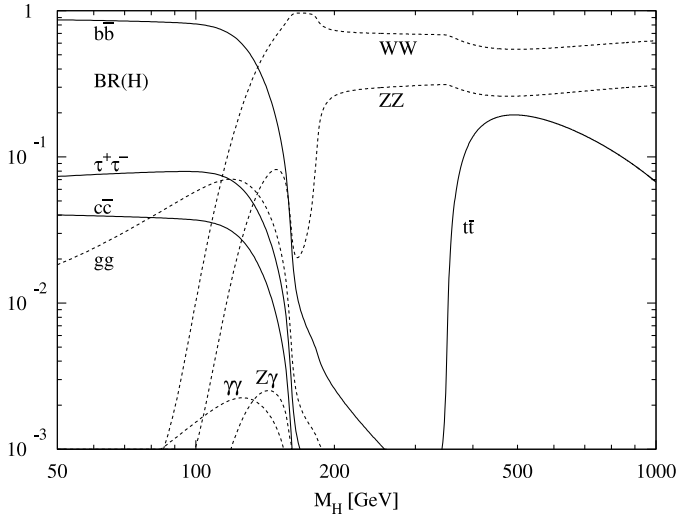

Figure 2.2: Higgs boson decay channel branching ratios as a function of mass [27].

2.2.2 Further physics studies

As well as Higgs searches, the LHC will allow studies of the SM through QCD, electroweak and flavour physics. Precision studies in these areas can also give indications of physics beyond the SM. For example, top quarks will be produced at a rate of several Hz allowing tests of SM couplings and spin. QCD can be thoroughly tested through the measurement of jet production with transverse energies of up to 4 TeV. These conditions are unprecedented in particle physics detectors. Other examples include flavour changing neutral currents and lepton flavour violation. Also, through TOTEM, an experiment designed to measure the p-p elastic cross section [28], diffractive physics can be studied.

The CMS detector also has the capability of dealing with heavy-ion collisions, which can be used as a probe for quark-gluon plasmas. Recent results from RHIC [29] indicate very strongly interacting nuclear matter is produced in high energy heavy-ion collisions. A striking feature of this is the phenomenon of jet-quenching, the suppression of high p_T particles. The increase in the nucleon collision energy from 200 GeV/c^2 at RHIC to 5500 GeV/c^2 at LHC will allow the extension of these studies to much higher p_T.

As well as providing a thorough and precise probe of the Standard Model, a core CMS objective is to allow for direct searches beyond. For example, the various supersymmetry theories can be tested via searches for the new particles they predict.

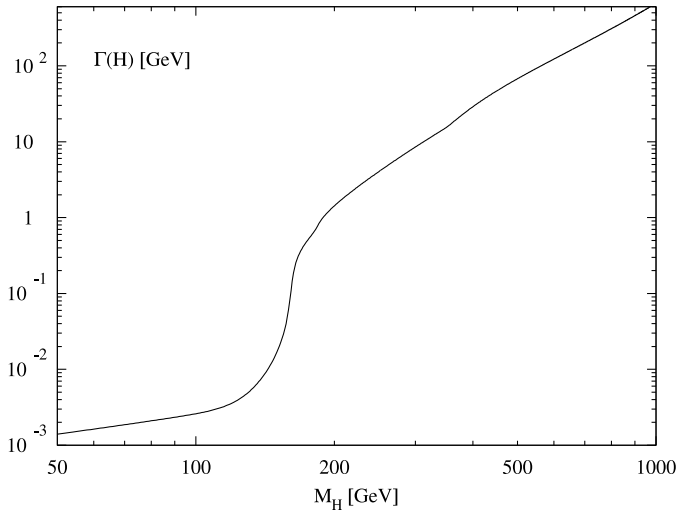

Figure 2.3: Higgs boson decay width as a function of mass [27].

Supersymmetric signatures are expected to include significant E_T^{miss}, if R-parity is conserved.

2.3 The CMS detector

CMS has been designed to cleanly detect the signatures of new physics by identifying and precisely measuring $\mu^{\pm}$, $e^{\pm}$ and γ over a large energy range. Robust tracking and precise vertex reconstruction and electromagnetic calorimetry are essential for the channels of interest. The high data rate (1 - 2 MB/event [30]) and intense radiation environment make the design of CMS very challenging.

The general structure of CMS is shown in Figures 2.4, 2.5 and 2.6. It is 15 m in diameter and 21.6 m in length with roughly cylindrical symmetry around the LHC beam-pipe. The detector consists of distinct *sub-detectors*, each with a well defined set of properties to measure within a given physics event. All measurements are performed as accurately as possible and with minimal disruption to the underlying physics processes. Each sub-detector labelled in these figures will be discussed in detail throughout this chapter.

The origin of the CMS co-ordinate system is the nominal LHC interaction point at *Point 5*. The z-axis points along the beam axis towards the Jura mountains and the y-axis vertically upwards. The x-axis points radially inward toward the centre

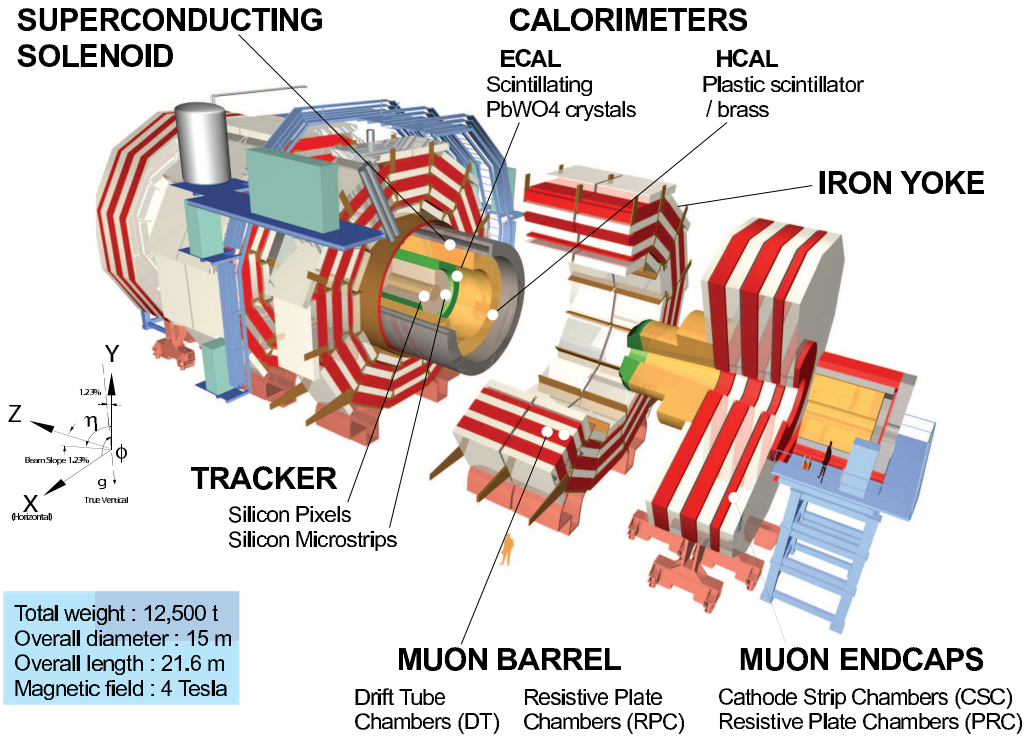

Figure 2.4: The CMS detector and its co-ordinate system.

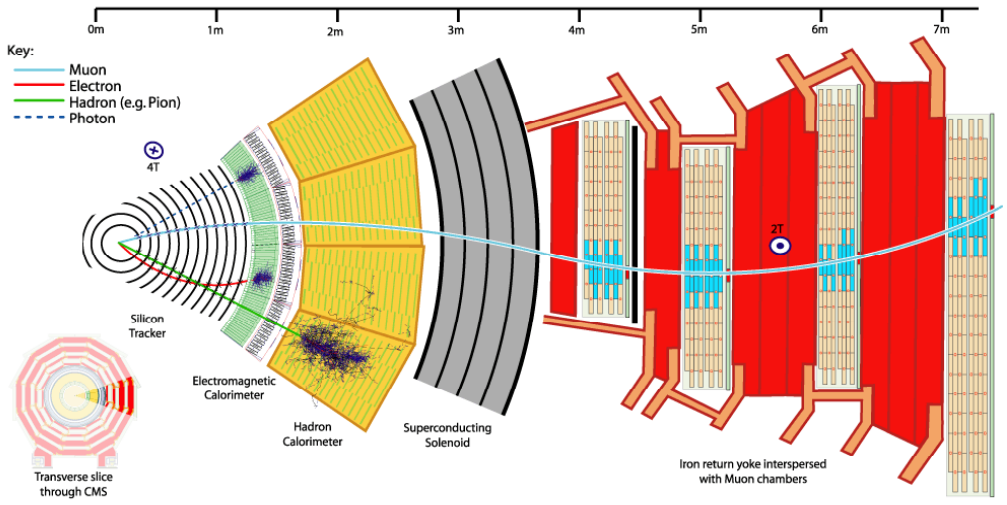

Figure 2.5: A slice of the CMS barrel in the x-y plane. Trajectories of a muon, electron, hadron and photon are illustrated.

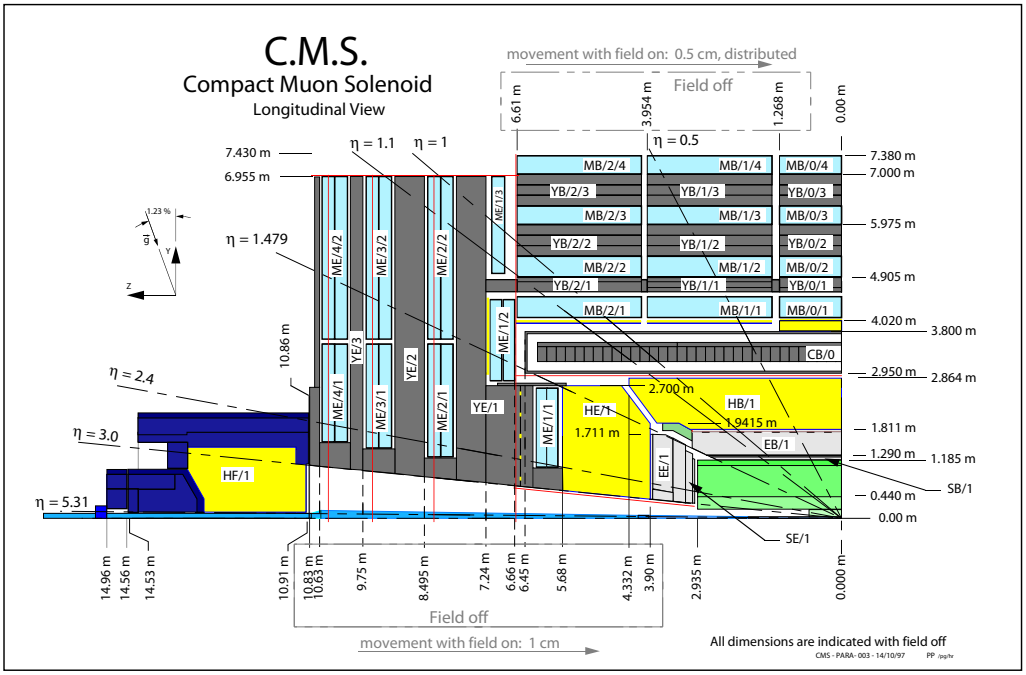

Figure 2.6: A quadrant of CMS in the x-z plane. Here, the tracker is green, electromagnetic calorimeter is light grey, hadronic calorimeter yellow and muon chambers blue. The iron return yoke (YB,YE) and magnet (CB) are both dark grey.

of the LHC, therefore. Cylindrical co-ordinates are defined with the polar angle, θ, measured from the z-axis and the azimuthal angle, ϕ, from the x-axis in the x-y plane. Pseudorapidity, η, is defined as $\eta = -\ln\tan(\theta/2)$. Transverse quantities such as transverse momentum, p_T, transverse energy, E_T and missing transverse energy, E_T^{miss}, are therefore computed from the x-y components.

2.3.1 The pixel tracker

On the inside of CMS, encompassing the beam-pipe, is the tracker which measures the trajectories and momenta of charged particles up to $\mid \eta \mid \simeq 2.4$ [31], [32]. A silicon detector is used, which provides a relatively low number of precisely measured points, rather than continuous tracks. Fine granularity pixels are placed closest to the interaction point, where the particle flux is highest to maintain a low channel occupancy and minimise track ambiguities [33].

The pixel system consists of 3 barrel layers: 4.4 cm, 7.3 cm and 10.2 cm from the beam-pipe with a length of 53 cm and 2 endcap discs extending from 6 cm to

15 cm in radius, at $\mid z \mid = 34$ cm and 46.5 cm. 66 million pixels of size $\sim$100 $\times$ 150 μm^2 are arranged across 768 and 672 modules in the barrel and endcaps respectively. To maximise vertex resolution an *almost square* pixel shape has been adopted. A Lorentz angle of 23° in the barrel improves the $r-\phi$ resolution through charge sharing. The endcap discs are assembled with a turbine-like geometry with blades rotated by 20° to also benefit from the Lorentz effect. The resultant spatial resolution is 10 μm in $r-\phi$ and 20 μm in z, allowing a primary vertex resolution of $\sim$40 μm in z. The layout is illustrated in Figure 2.7.

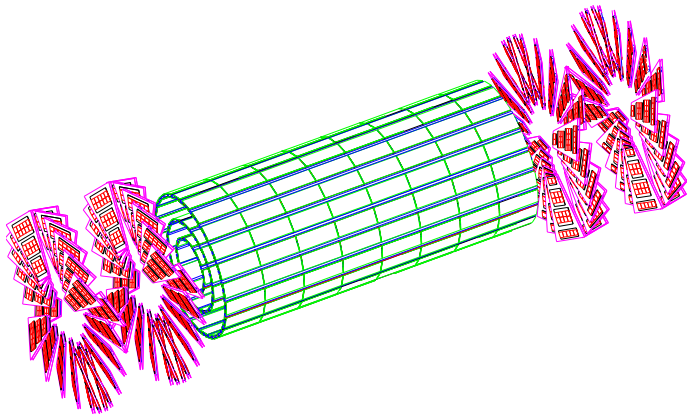

Figure 2.7: A schematic of the pixel tracker. The barrel is coloured green, the endcaps red.

2.3.2 The strip tracker

Outside the pixel detector the particle flux is low enough to justify the use of a Silicon Strip Tracker (SST). Here, 9.6 million microstrips are divided between six regions: an inner barrel (TIB), two inner discs (TID), an outer barrel (TOB) and two end-caps (TEC). One quadrant of this setup is shown in Figure 2.8. The TIB and TOB are sometimes referred to as the *barrel* of the SST, each TEC/TID pair constitutes an *endcap*.

The basic unit of this sub-detector is called a *module*, each housing 512 or 768 strips, depending on the location, and a Front-End Hybrid (FEH) [34]. On the latter a number of chips are mounted, namely four or six APV25 readout chips (discussed in detail in Section 2.4.2), two or three APVMUX chips, a Detector Control Unit (DCU) for environmental monitoring and a Phase-Locked Loop (PLL) for clock regeneration.

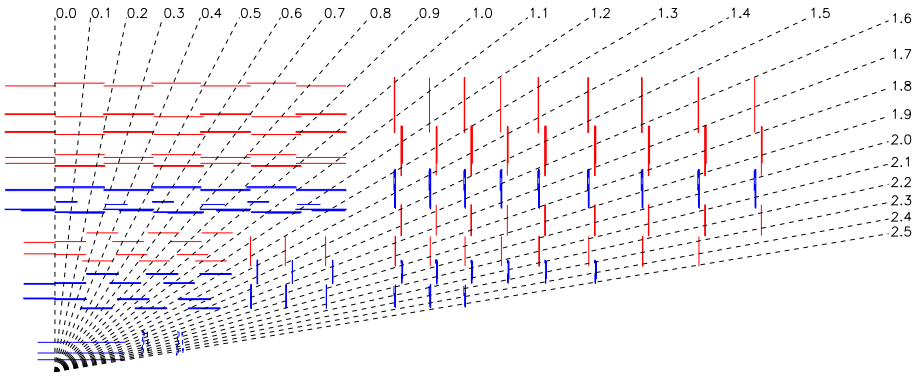

Figure 2.8: A schematic of one quadrant of the SST in the x-z plane. Modules in red are single-sided, blue are double-sided. The scale along the top and right is for pseudorapidity.

region	layers	radius m	modules	thickness μm	mean pitch μm	resolution $r-\phi\times$ z μm$\times\mu$m
TIB	4	0.2-0.55	2724	320	81/118	23-34×230
TOB	6	0.55-1.16	5208	500	81/183	35-52×530
TID	3	0.55-1.20	816	320	97/128/143	23-34×230
TEC	9	1.20-2.80	6400	320/500	96-183	35-52×530

Table 2.2: Breakdown of the SST.

The entire SST consists of almost 15 200 modules, mounted on carbon-fibre support structures. A breakdown of the module content of each partition is given in Table 2.2. In order to maintain a mean strip occupancy below 2 % under p-p collisions and hence minimise pile-up*, the silicon strip length is set to 11.9 cm for radii < 59 cm and 18.9 cm beyond. The longer detectors are thicker (500 μm) than the shorter ones (320 μm) to maintain a constant signal/noise.

Barrel and endcap strips lie along the z and radial directions, allowing precise measurement of the r and ϕ and the z and ϕ coordinates respectively. To improve the track resolution and resolve tracking ambiguities, some double sided modules are used, shown in blue in Figure 2.8. Each consists of a regular module and a *stereo* counterpart angled at 100 mrad back-to-back. They provide a three-dimensional hit measurement in r, ϕ and z. This layout allows the tracker to provide 8 to 14 *hit* measurements for high momentum tracks. On average about half of these are three-dimensional.

*Throughout this thesis the term *pile-up* is used to describe two independent effects: the influence of lower p_T p-p collisions within the same event (sometimes referred to as the *underlying event*) and the influence of neighbouring bunch-crossings.

Each module contains a sensor, which is an n-bulk device with p^+ doped microstrips on the surface [35]. A reverse-bias of up to 450 V is applied, such that the silicon is fully depleted and the charge collection time ($\sim$10 ns) is within the LHC bunch-crossing period. Incoming charged particles interact with the silicon, depositing energy and creating electron-hole pairs in the medium. This electrical signal is then channelled to the front-end electronics for readout.

CMS is a high radiation environment. After 10 years of operation, the inner layers of the microstrip tracker will have received a fluence of about 3.3×10^{13} cm^{-2} and 12.6×10^{13} cm^{-2} neutral and charged hadrons respectively. Surviving these conditions was a great challenge in the detector and readout design and there are several key features in the technologies used that reflect this. The silicon sensors suffer both surface damage, due to the trapping of charge released by ionising particles and bulk damage as a consequence of hadronic collisions with the lattice itself. The former increases strip capacitance and noise, the latter increases leakage current (and hence noise), decreases charge collection efficiency (and hence signal) and increases the depletion voltage [36], [37].

The APV25 readout chip is intrinsically radiation hard due to its small feature size. Since all active components of the integrated circuit are fabricated in a thin layer at the silicon surface, bulk damage to the silicon is of little concern. Charge trapping within transistor gate oxides can result in increased noise, voltage threshold shifts, decreased mobility and increased leakage currents. However, 5.5 μm thick gate oxides allow the neutralising of the trapped charge through electron tunnelling [38].

The leakage current for the bulk of the detector increases exponentially with temperature and linearly with radiation dose. In turn, the sensor temperature increases with the power dissipated within it. This cyclic dependency can result in a *thermal runaway* and requires the sensor temperature to be maintained at $\sim$ -15°C. This is achieved through a distributed cooling system.

Beyond the tracker sits the calorimetry. This consists of an electromagnetic calorimeter designed to measure the energies of electrons and photons and a hadron calorimeter designed to measure the energies of hadronic jets. These sub-detectors will be discussed in the following sections.

2.3.3 The electromagnetic calorimeter

The Electromagnetic CALorimeter (ECAL) is a homogeneous scintillating detector [39]. It has been designed for high precision energy measurements of photons and electrons and has been optimised to discover the Higgs boson through the $H^0 \rightarrow \gamma\gamma$ channel. A good energy resolution is required for this channel to distinguish it from the large $\pi^0 \rightarrow \gamma\gamma$ background.

The ECAL consists of almost 76, 000 crystals which are divided between two main sections, the barrel and the endcaps. Lead Tungstate, $PbWO_4$, crystals have been chosen as the sensitive material here due to its short radiation length, $X_0 \simeq 0.89$ cm, small Molière radius of 2.2 cm and radiation hardness. The barrel contains 61, 200 crystals over a pseudorapidity $\mid \eta \mid < 1.48$. Crystals are 23 cm long with a front-face of $\sim$2.2 cm $\times$ 2.2 cm corresponding to $\Delta\eta \times \Delta\phi \simeq 0.0175 \times 0.0175$ rad. Each has a pseudo-pointing geometry toward the mean interaction point, but with a 3° offset. This offset ensures particles cannot escape between the inter-crystal *dead* areas. Crystals are grouped into sets of 2 $\times$ 5 termed *submodules*. A set of 10 $\times$ 17 submodules forms a single *supermodule*.

The endcap section contains 14, 648 crystals covering a pseudorapidity $1.48 < \mid \eta \mid < 3.0$. Endcap crystals are 22 cm long with a front-face of 2.86 cm $\times$ 2.86 cm corresponding to $\Delta\eta \times \Delta\phi$ starting at $\simeq 0.0175 \times 0.0175$ rad, increasing to a maximum of 0.05×0.05 rad. Again, all are tapered in the same way to produce a pseudo-pointing geometry toward the main interaction point with a 3° offset. Crystals are grouped into 5 $\times$ 5 arrays called *supercrystals*, which in turn are grouped into 4 quadrants, each known as a *dee*.

Finally, a preshower region sits in front of each endcap covering a pseudorapidity of $1.48 < \mid \eta \mid < 2.6$. These are designed to improve rejection of the $\pi^0 \rightarrow \gamma\gamma$ background. They are necessary to resolve the low opening angle between photons in this region.

The energy resolution of the ECAL can be parameterised by Equation 2.1. Here E is in GeV and a, b, and c are constants where $a = 3.63$ %, $b = 0.124$ % and $c = 0.26$ % (taken from beam-test measurements [30]).

$$\frac{\sigma}{E} = \frac{a}{\sqrt{E}} \oplus \frac{b}{E} \oplus c \tag{2.1}$$

On the right-hand side, the first term represents stochastic uncertainties in the signal generating process. The second term represents digitisation, pre-amplifier

and pile-up noise. The third term is a constant and represents uncertainties in: longitudinal containment (energy leakage through the rear of the crystals); inter-crystal calibration and non-uniformity of light collection along the crystal. The first two terms here decrease at high energies making the third dominant. A constant term of no more than 0.5 % is required for precision physics studies, especially the resolution of the expected $H^0 \rightarrow \gamma\gamma$ peak.

2.3.4 The hadronic calorimeter

The Hadronic CALorimeter (HCAL) is a sampling calorimeter composed of plastic scintillator tiles interspersed with copper absorber plates [40]. Like the ECAL, the HCAL barrel region covers a pseudorapidity of $\mid \eta \mid < 1.48$ and the endcap $1.48 < \mid \eta \mid < 3.0$. The barrel section is ∼79 cm deep, with ∼6.5 λ of brass plates. This is rather thin, so to produce a larger sampling depth another layer is placed beyond the solenoid, extending the HCAL depth to 10 λ.

To cover the pseudorapidity region $3.0 < \mid \eta \mid < 5.0$ for an improved E_T^{miss} measurement, a forward calorimeter is located 6 m downstream of the HCAL endcaps. It allows jet measurements if the jet axis lies within $\mid \eta \mid < 4.5$. Quartz fibres are used as the scintillating material here, along with a copper absorber.

Throughout the barrel and most of the endcaps, the readout granularity is $\Delta\eta \times \Delta\phi \simeq 0.0870 \times 0.0873$ rad. The energy resolution of the HCAL is parameterised by Equation 2.2.

$$\frac{\sigma}{E} = \frac{a}{\sqrt{E}} \oplus b \tag{2.2}$$

Here, E is in GeV and a and b are constants where $a = 115$ % and $b = 5.5$ % (measured for 20 - 300 GeV pions using beam-test data [41]). On the right hand side of the equation, the first term represents statistical fluctuations in the signal generating process. The second term is for intrinsic fluctuations in the number of shower components and can vary significantly from shower to shower. The energy recorded by the HCAL, E_{vis} is related to the actual energy of the incident jet E by Equation 2.3.

$$E_{vis} = eE_{em} + \pi E_{ch} + nE_n + NE_{nucl} \tag{2.3}$$

Here E_{em}, E_{ch}, E_n and E_{nucl} are respectively the electromagnetic, charged hadronic and low energy neutron components of the showers and the energy lost in breaking up nuclei. Each component has its corresponding sampling fraction.

parameter	value
Field	4 T
Inner bore	5.9 m
Length	12.9 m
Number of turns	2168
Current	19.5 kA
Stored energy	2.7 GJ
Hoop stress	64 atm

Table 2.3: Parameters for the superconducting solenoid.

2.3.5 The magnet

At a radius of ∼3 m is a large superconducting solenoid, designed to contribute information on particle momenta from the bending of the tracks [42]. Its design parameters, given in Table 2.3, are defined by the unambiguous determination of the sign of a muon with momentum 1 TeV/c. The entire calorimetry system (with the exception of an extra layer in the HCAL barrel) is situated within this coil to minimise resolution loss as the particles pass through.

2.3.6 The muon tracker and return yoke

The muon tracker is the outer most sub-detector [43]. It is designed to record the trajectories of muons, the only particle flavour (except neutrinos) to penetrate the calorimeter, in the region $| \eta | < 2.4$. Due to the low particle flux compared with the silicon tracker and the large surface to be covered ($\sim$ 25 000 m^2), the granularity is far coarser consisting of nearly 1 million readout channels. The schematic of one quadrant of the muon system is shown in Figure 2.9.

Muon detectors are distributed over four layers in the barrel and endcap sections of the return yoke. The latter is made of iron and designed to prevent a strong magnetic field from extending to a very large distance. Three gaseous muon tracking technologies are employed, reflecting the varying radiation and magnetic environments. In the barrel, where $| \eta | < 1.2$, 250 Drift Tube (DT) chambers are operated in up to 12 planes. Each has a cross section of 42 × 13 mm^2, consists of a central anode wire surrounded by an aluminium cathode and is filled with Ar and CO_2. The induced charge has a maximum drift time of 400 ns (cf. the 25 ns LHC bunch

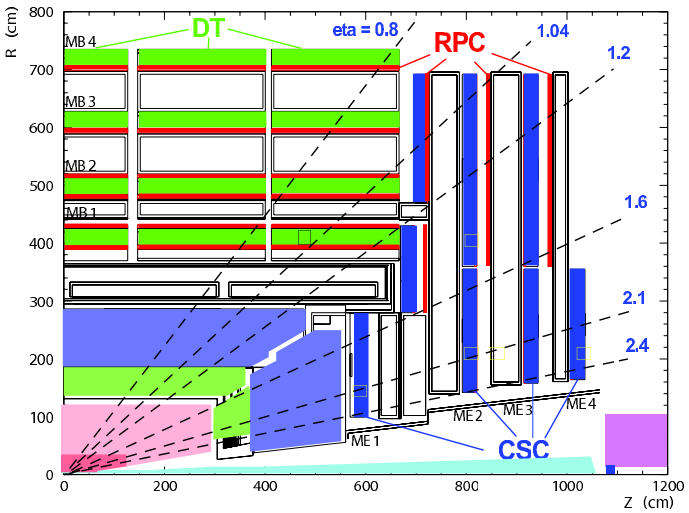

Figure 2.9: A schematic of one quadrant of the muon system. It shows three different detector technologies: Drift Tubes (DT), Resistive Plate Chambers (RPC) and Cathode Strip Chambers (CSC).

spacing). Each chamber provides a muon vector of resolution 100 μm in $r - \phi$ and 1 mrad in direction.

In the endcaps, where $1.2 < \mid \eta \mid < 2.4$, the high muon and neutron background environment means 468 Cathode Strip Chambers (CSCs) are used instead. CSCs provide precise space and time information in the presence of a high magnetic field and high particle rate. Each is trapezoidal and contains 6 gas gaps. Each gap has a plane of radial cathode strips with perpendicular anode wires. A single chamber provides a muon vector of resolution 200 μm in $r - \phi$ and 10 mrad in direction.

To achieve a fast time response and hence accurate bunch-crossing identification, Resistive Plate Chambers (RPCs) are also used in both barrel and endcap regions. An RPC consists of a gas gap enclosed by two graphite-coated bakelite plates forming cathodes, operated in avalanche mode. This gives them a time resolution of $\sim$1 ns.

2.4 SST control, readout and data acquisition

This section aims to describe the SST control system [44], which allows the monitoring of its environment and configuration of its electronics. The readout of strip data, front-end processing and subsequent construction of the raw data event will also be summarised. The control and readout systems for the SST are illustrated in Figure 2.10.

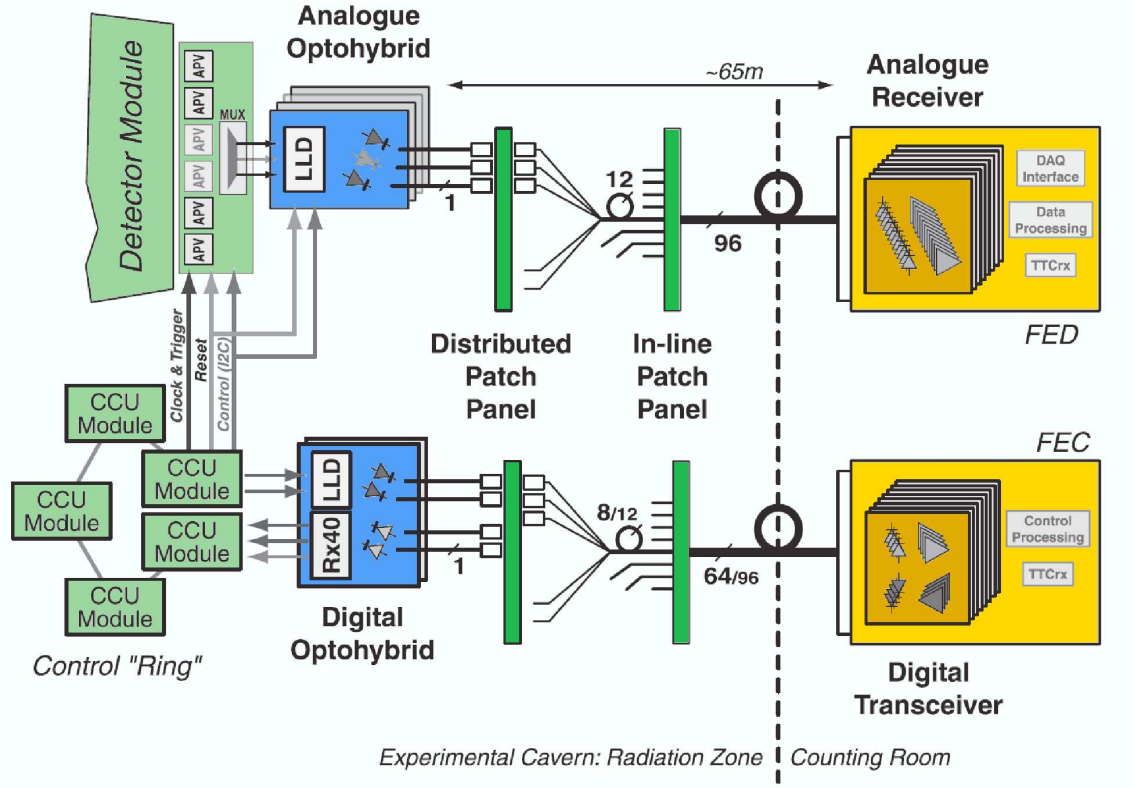

Figure 2.10: The SST control and readout systems. The control chain is shown at the bottom of the figure, the readout chain at the top.

2.4.1 The control system

The control system for the SST must be capable of monitoring the environmental conditions of almost 16k modules and enable the correct configuration of the readout electronics via the I^2C protocol [45]. It is also responsible for distributing clock and trigger signals to the detector front-end.

The hardware architecture is based on a single board: the Front-End Controller (FEC). This provides the interface with the global Timing, Trigger and Command (TTC) system [46], which distributes the LHC master clock and level 1 trigger to the front-end electronics. Encoded clock and trigger information is distributed by the FEC via optical links. Signals are then converted to electrical via a Digital OptoHybrid (DOH). Each FEC hosts several token *control rings*, with a series of Communications and Control Units (CCU) modules. A CCU module will communicate directly with several detector modules via a PLL Application Specific Integrated Circuit (ASIC). The PLL will perform clock regeneration and phase adjustment if necessary [47].

2.4.2 Front-end electronics readout

The SST data acquisition system is based on an analogue front-end readout chip known as the APV25 [48], [49]. The APV25 is an ASIC designed in 0.25 μm CMOS technology for low-noise and fast signal readout in a high radiation environment. The microstrip signals are read out in groups of 128. Each group is read by a single APV25 chip: tracker modules consisting of either 4 or 6 APV25s in total. The signal from each microstrip is pulse shaped in its corresponding APV25 channel to closely match the ideal CR-RC waveform which peaks after ∼50 ns. These signals are then stored inside buffers, awaiting a level 1 trigger. The buffers (128 pipelines) have a 192 location memory (with 32 locations reserved for storing events awaiting readout) into which samples are written at 40 MHz. This corresponds to a programmable level 1 trigger latency of up to 4 μs.

Following a trigger the relevant data are read out from their pipelines by an Analogue Pulse Shape Processor (APSP) in one of two ways. In peak mode, only the sample at the pulse peak is read. In deconvolution mode, three samples are read out from the memory and are processed with a deconvolution filter [50]. This outputs a weighted sum of the measured signal amplitudes to produce an effective one-sample reading from a signal with a shorter time constant of ∼25 ns. Strip pulse shapes in both readout modes are shown in Figure 2.11. With a 25 ns bunch crossing time, the deconvolution process should reduce pile-up significantly at high luminosities.

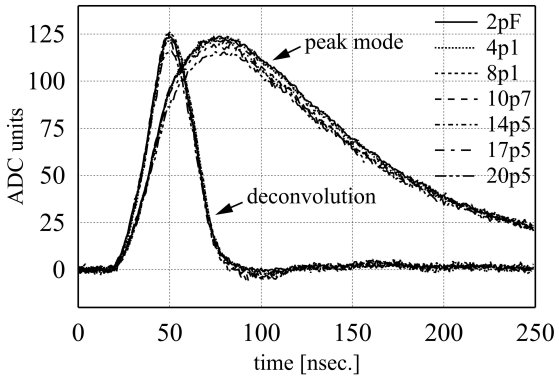

Figure 2.11: APV25 peak and deconvolution mode pulse shapes with varying input capacitance.

After the APSP, the data from pairs of APV25 chips are fed into a 256:1 multiplexer (APVMUX) which outputs at 40 MHz along a single stream. An example data frame from a multiplexed APV25 pair is shown in Figure 2.12 in both multiplexed and channel order. 12-bits of binary data are prefixed to the output data stream

including a 3-bit digital header and an 8-bit pipeline address. The entire data frame is ∼7 μs long, consisting of the 12-bit header and 256 25 ns analogue samples ending with a digital footer.

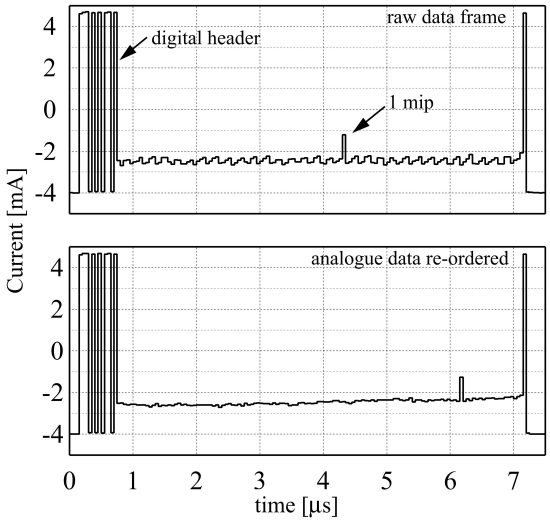

Figure 2.12: An example data frame from two multiplexed APV25s in both its raw form (after multiplexing) and re-ordered by channel number. It contains a 3-bit digital header, 8-bit pipeline address and 256 analogue samples. Signal from a minimally ionising particle (mip) is shown.

2.4.3 The optical links

The Analogue OptoHybrid (AOH), shown in Figure 2.10, converts electrical signals received from each APVMUX into optical for transmission along ∼65 m fibres to the off-detector electronics. Each module requires a single AOH which in turn consists of a Linear Laser Driver (LLD) ASIC and two or three edge-emitting laser diodes (one per multiplexed APV25 pair). The LLD provides currents to each laser to bias them at their quiescent operating points. The lasers generate the required optical power for transmission of the signals along the optical fibres. To compensate for component performance spread, temperature fluctuations and damage induced by radiation over time, the bias current of each laser and the gain of the LLD are configurable.

2.4.4 The Front End Driver

The Front-End Driver (FED) is designed to read out, digitise and process optical signals transmitted by the APV25 analogue pipeline. The entire SST requires ∼440

FEDs, each capable of receiving multiplexed, optical signals from 96 channels or 192 APV25s. These are converted to electrical levels, then sampled and digitised at 40 MHz. 8 Front-End (FE) Field Programmable Gate Array (FPGA) chips are then used to synchronise and re-order the data, before pedestal and common mode corrections are applied. Finally the data are zero-suppressed to reduce the event data size. This is essential when operating at high trigger rates. The zero-suppression algorithms include a cluster finding stage which accepts only clusters of neighbouring strips passing configurable amplitude thresholds. For more information on the FED firmware algorithms, please refer to [51]. The FED has been designed to operate within 4 different readout modes listed below.

- **Zero Suppressed** This is the normal FED operating mode and is designed for use with proton - proton and Pb-ion collisions. Data are 8-bit digitised in the 0 mV to 1023 mV range, re-ordered into physical channel order and zero-suppressed. Pedestals and common-mode noise are also removed.
- **Processed Raw** Data are 10-bit digitised in the 0 mV to 1023 mV range, re-ordered into physical channel order and pedestals are subtracted. No zero-suppression, pedestal or common-mode noise are subtracted, however.
- **Virgin Raw** Data are 10-bit digitised in the 0 mV to 1023 mV range. Only the *raw* data are produced i.e. no re-ordering, pedestal/common-mode noise subtraction or zero-suppression is performed. This mode will be used for testing, commissioning and calibration.
- **Scope** Data are 10-bit digitised in the 0 mV to 1023 mV range. The three modes above are all performed within the FEDs *frame-finding* data capture facility. Here, data capture starts as the FED recognises the header of the APV data frame, and continues until the frame ends. Another method of data capture involves commencing capture with the arrival of a trigger and continuing for a programmable number of samples. A full, multiplexed APV data frame consists of 281 samples. No re-ordering, pedestal/common-mode noise subtraction or zero-suppression is performed.

FED data are read out via a single back-end FPGA. This combines the processed information from all front-end channels within a single packet for each level 1 trigger and stores them within an external memory buffer to deal with the Poissonian fluctuations in the data rate. These buffered data are then transmitted via a fast S-Link64

interface [52] and event-builder to a computer farm for further data processing and event filtering.

Each back-end FED data packet contains, as well as the strip levels, various registers describing event, FED, APV25 and Trigger Throttling System (TTS) information. This *event fragment* is packed into groups of 64-bit words. For more details on the content and structure of a FED data packet, please refer to [53].

2.4.5 The event builder

Event fragments from each FED back-end are combined into a single event object within an event builder, illustrated by Figure 2.13. This consists of 64 FED builders and 8 Readout Unit (RU) Builders. A single FED builder combines event fragments from up to 8 FEDs, via a FED Builder switch into a *superfragment* on a RU. 64 superfragments distributed across 64 RUs are built per event. These superfragments are then combined into a full event on a Builder Unit (BU) via the 64 × 64 Readout Builder Network (RBN). 64 RUs connected via a RBN to 64 BUs is known as a RU builder. There are 8 RU builders in the full CMS system. The full event is delivered to a single filter unit node on the filter farm via the filter farm network. Here event selection algorithms known as the high-level trigger are run.

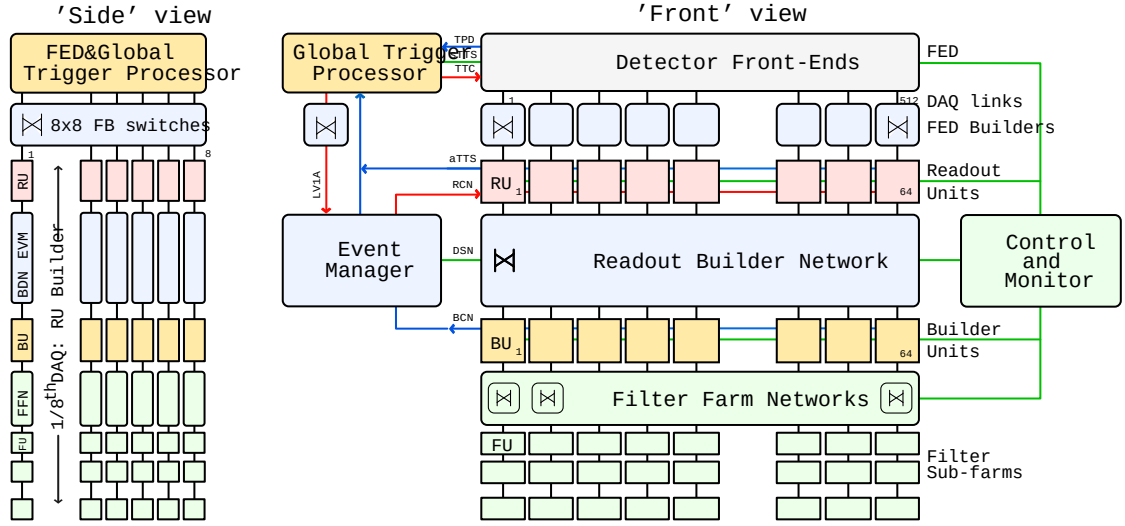

Figure 2.13: The CMS event builder. Raw data packets from 440 FED back-ends are built into a single event for processing on a Filter Unit (FU).

2.5 The trigger system

At the centre-of-mass energy and luminosity expected from the LHC, the interaction rate is orders of magnitude larger than what can be reasonably analysed in later offline analysis. The LHC bunch crossing frequency of 40 MHz and event data sizes of ∼1.5 MB would also result in data rates of ∼57 TBs^{-1}, which far exceeds the current maximum write-to-disc rate of ∼150 MBs^{-1}. The trigger system is designed to provide a high event selectivity of 10^{-4} - 10^{-5} with respect to the active bunch crossings at the LHC by accepting only those events deemed physically interesting.

In the CMS design, this selectivity is achieved in two physical steps: the level 1 trigger (L1) analyses each 25 ns crossing within a latency of 4 μs, using calorimetry and muon information to reduce the event rate from 40 MHz to 100 kHz. The High-Level Trigger (HLT) must operate at the L1 accept rate to combine information from the full detector and trigger at the write-to-disc rate of 100 Hz. Both trigger rates will incur Poissonian fluctuations.

2.5.1 Level 1

The level 1 trigger [54] is built of mostly custom-made hardware dedicated to analysing the detector information with a fairly coarse granularity. It operates at two levels. Firstly the calorimeter and muon triggers process information from their respective sub-detectors. The output is then fed into the global trigger which makes the final trigger decision. A schematic of the decision making process is shown in Figure 2.14.

The calorimeter trigger uses information from individual trigger towers. A trigger tower consists of a single HCAL cell and various ECAL cells within the same η, ϕ region (5×5 in the barrel and larger in the endcaps). For the ECAL this involves the energy sum and information on the transverse extent of the electromagnetic energy deposit, for the HCAL the energy sum and presence of minimum ionising energy. The Regional Calorimeter Trigger (RCT) combines this to find candidate electrons and photons (isolated and non-isolated) along with taus and jets. The candidate information, along with their E_T, are forwarded to the Global Calorimeter Trigger (GCT). An η,ϕ grid of quiet regions is also forwarded to the global muon trigger for isolation cuts. The GCT sorts the RCT candidates by energy and calculates

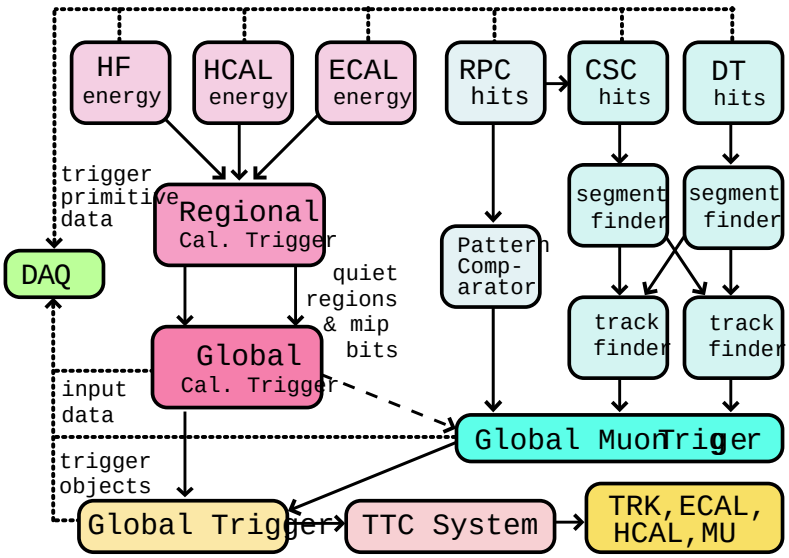

Figure 2.14: A schematic of the L1 trigger. Arrows represent the flow of trigger information. The final decision made by the *global trigger* is propagated via the TTC system to the various sub-detectors.

the total E_T and E_T^{miss} from them. The top 4 candidates, along with the global E_T information is forwarded to the global trigger.

The global muon trigger sorts the RPC, DT and CSC muon tracks, normalises them to the same p_T, η and ϕ scale and validates the muon sign. It then correlates the CSC and DT tracks with those of the RPC. Tracks are also deemed isolated if they fall within an area of quiet calorimeter towers. The final set of tracks are sorted by their quality, correlation and p_T. The top 4 muons are sent to the global trigger.

The global trigger synchronises trigger information arriving from the calorimeter and muon triggers and forwards the final L1 trigger decision to the TTC via the TTS. The latter allows the reduction of individual trigger rates (prescaling) to deal with large cross section triggers and the monitoring of trigger efficiencies. All trigger candidates are accompanied by their η,ϕ coordinates to allow the global trigger to vary thresholds based on location. It also permits the HLT trigger decision to focus on a region-of-interest.

2.5.2 High-level

The high-level trigger [55], [56], [57] processes all events accepted by L1 in a single processor farm. The fully programmable nature of the processors on the filter farm

enables the implementation of very complex algorithms utilising any and all information in the reconstructed event. The strategy follows that of traditional multi-level trigger systems, where the selection process is optimised by rejecting uninteresting events as quickly as possible. With this in mind, each trigger path consists of a sequence of software modules with increasing complexity and physics sophistication. Each module fulfils a well defined task such as reconstruction, intermediate trigger decisions or the final trigger decision for that path. If an intermediate decision on a trigger path is negative, the remainder will not be executed.

The HLT reconstruction and selection chain is implemented in steps which can be conceptually grouped into distinct trigger systems: Level 2 (L2) and Level 3 (L3). It should be noted that this terminology is used for convenience only and the division is entirely arbitrary. The convention used is that L2 algorithms refer to the first selection step in the HLT process and use information from the calorimeter and muon detectors only. In contrast, L3 trigger algorithms include fully reconstructed tracks. Since track reconstruction is a process that demands large amounts of CPU time this analogy can be carried even further. Level 2.5 (L2.5) triggers are those using information from only partially reconstructed tracks e.g. tracks reconstructed in the signal cone of a τ candidate only.

All HLT algorithms have been implemented using the CMSSW software suite, which will also be used for reconstruction and offline analysis. Chapter 3 describes some of the more relevant online and offline algorithms in detail.

Chapter 3

CMS simulation and reconstruction software

The CMS simulation, reconstruction and analysis software framework, CMSSW, is described with a focus on the reconstruction algorithms relevant to this thesis.

3.1 The software framework

"The high-level goals of the CMS software are to process and select events inside the High-Level Trigger Farm, to deliver the processed results to experimenters within the CMS Collaboration, and to provide tools for them to analyse the processed information in order to produce physics results." [30]

The CMS software infrastructure is broken down into 3 distinct categories: Monte-Carlo (MC) generation, simulation and reconstruction. MC generation involves the simulation of single p-p (or Pb-Pb) collisions within independent bunch-crossings and the decay of the collision products until they reach a stable state. The collision vertex is fixed at (0,0,0). The simulation phase propagates the MC particles through the detector material, applying the magnetic field and smearing the primary vertex. The response of the detector's sensitive material to the particle interactions, the front end electronics readout and the digitisation are also introduced. Finally, effects from previous and future bunch crossings as well as multiple inelastic collisions, collectively known as *pile-up*, can be modelled. Both the MC generation and simulation are designed to reproduce the readout state of the detector for a given

bunch-crossing. The final stage, reconstruction, converts the raw detector data into physics objects suitable for analysis. The reconstruction phase is transparent to real or simulated data.

Since the MC generation algorithms are not CMS specific, they have not been developed within the CMS software framework. Instead, the latter has been designed to interface with external generators such as Pythia [58] or CompHEP [59].

Until 2004 the simulation and reconstruction algorithms of CMS were written within the object-oriented COBRA framework. Within this scheme, the collection of algorithms themselves were known as ORCA [60]. Since then, a new single software suite, known as CMS SoftWare (CMSSW) has been developed. CMSSW is written in C++ and designed around the Event Data Model (EDM). The basic premise of the EDM is that all user-defined types are contained within a single object called the *event* [30], [61]. All such types relate to a single triggered readout of the CMS detector. The contents of an event can include objects representing the MC itself, various stages of the simulation, the raw detector output, reconstruction or analysis. They are processed by passing the event object through a series of user-defined *plugin modules* within a CMSSW *path*.

Each plugin module has a well defined and distinct functionality. No direct communication between plugin modules exists, instead they communicate only via the event. This allows them to be independently tested and used. User access to auxiliary conditions data e.g. cabling information is given via an independent object called the *event setup*.

Finally, the information provided by the reconstructed objects can then be stored in ntuples and analysed using ROOT [62]. This final stage can either be performed within a dedicated CMSSW module or externally to the framework. ROOT provides fast and convenient functions to both handle and present large amounts of data.

3.2 Track reconstruction

The raw data output from each SST FED encodes the basic hit information necessary for tracking, along with information on the bunch crossing number, the status of the tracker hardware and the status of the data acquisition system. Prior to reconstruction, these data are interpreted and the zero-suppressed hit information (pulse

heights and corresponding strip/pixel numbers) known as a *digis*, are extracted. This process is known as raw data *unpacking*. Track reconstruction itself can be artificially divided into two phases: hit reconstruction (local) and track formation (global).

3.2.1 Hit reconstruction

Hit reconstruction is itself a two-step process involving *clusterising*, that is grouping together neighbouring, gain-corrected digis and incorporation of the geometrical position of the hit.

Though the choice of clusterising algorithm is open, a robust option commonly used for the SST is the *three-threshold algorithm*. This selects hits based on thresholds applied to both the signal from each strip and the group of strips as a whole. Its clusters satisfy the following criteria, where t_1, t_2 and t_3 are constants [30].

- *signal/noise* of each strip $> t_1$.
- At least one *seed* strip with *signal/noise* $> t_2$.
- A combined *signal/noise* $> t_3$.
- A maximum number of internal *holes* (from dead strips) which is configurable.
- If the product is bounded by a dead strip on either side, they are appended to the cluster.

Each cluster is then converted into a tracking hit by including information on its position within 3D space. These objects are often termed *rechits* to highlight their status as fully reconstructed hits.

3.2.2 Track formation

Track formation can be subdivided into four steps: seed generation, fitting, ambiguity resolution and smoothing.

A track seed defines the initial trajectory parameters and their errors. The pixel detector is best suited for building track seeds due to its low occupancy, the precision of its hits and its proximity to the beam pipe. To obtain the five parameters

necessary for starting a trajectory, the resultant seed requires at least three hits or two hits and a beam constraint. The first hit falls within any predefined search area and usually on the innermost layer. For each selected first hit, compatible second (and third) hits are chosen.

To streamline this process for the HLT, tracking regions-of-interest can be identified within the pixel detector through seeds identified in external sub-detectors, such as the ECAL. Track seeds can also be built entirely from the strip tracker. This could be useful during detector start-up, for example, when the pixel detector may be poorly understood. It is also possible to combine the pixel layers with some strip layers. The hit finding algorithm for the strip tracker works along the same principle as the pixel hit finder.

Track fitting involves the grouping together of hits through a pattern recognition algorithm. One such algorithm commonly used by CMS is a combinatorial Kalman filter [63]. This proceeds from the coarse estimate of tracking parameters provided by the seed, determines which layers are compatible and extrapolates the trajectory to these layers according to the equations of motion of a charged particle in a magnetic field through material. As the trajectory encounters each hit, it is updated according to the Kalman filter formalism. At each layer, a new trajectory candidate is created for every compatible hit. An extra trajectory candidate is also generated in the case of a missing hit for that layer.

To limit the rapid growth in the number of track candidates several stopping conditions are introduced such as goodness-of-fit, number of invalid hits and total number of valid hits. The latter is especially useful for the HLT where, for example, the required performance is often reached after 5 or 6 hits. A minimum p_T threshold of 0.9 GeV eliminates *looping* tracks with insufficient p_T to reach the calorimeter.

Once the tracks are formed, ambiguities and instances of double counting are resolved. Since the full tracking information is only available at this point, any bias introduced with the imposed constraints is removed by a final fit. More details on the full tracking procedure can be found in [30], [64].

3.2.3 Vertex formation

Vertex reconstruction comes in two stages, finding: grouping tracks from a common vertex together, and fitting: determining the vertex parameters and their fit quality based on the associated tracks.

The offline vertex finding algorithm, which is most relevant to this thesis, uses a preselected set of tracks based on fit quality and p_T. They are then clusterised based on the z co-ordinate of their point of closest approach to the beam line, with tracks closer than 1 mm grouped within the same vertex. These candidates are then sorted into descending order of the sum of the ${p_T}^2$ of their associated tracks. In most cases the primary vertex of the signal is the first in the list.

The fitting process then determines the vertex parameters and their fit quality. The most commonly used fitting algorithm for CMS is the Kalman vertex filter [65]. Tracks incompatible with the candidate vertex are discarded here. Poor fits and vertices incompatible with the beam line are also excluded.

3.2.4 Track reconstruction performance

The track p_T resolution can be parameterised by Equation 3.1. p_T is in TeV and a and b are constants where a = 15 TeV^{-1} for $\mid \eta \mid < 1.6$ and 60 for $1.6 < \mid \eta \mid < 2.4$ and $b = 0.005$ [64].

$$\frac{\delta p_T}{p_T} = a p_T \oplus b \tag{3.1}$$

On the right hand side, the terms represent the dependence of p_T resolution on track curvature and uncertainties due to multiple scattering in the tracker material, respectively. Figure 3.1 shows the global track reconstruction efficiency for muons and pions as a function of η. The drops in efficiency at low and high η are mainly due to drops in coverage of the pixel discs. For $\eta \simeq 0$, for example, the efficiency decreases slightly due to gaps between the ladders of the pixel detector. For muons, the efficiency is about 99 % over most of the acceptance. In general, the efficiency is lower for pions due to elastic nuclear interactions with the tracker material inhibiting the Kalman filter. In fact, 20 % of 1 GeV pions do not reach the outer tracking layers [66].

Figure 3.2 shows the track p_T and transverse impact parameter resolutions as a function of the number of reconstructed hits. Although in each case the performance degrades with fewer hits, reconstruction becomes significantly faster as the number of channels and combinatorics are reduced. A reduced number of reconstructed hits also serves to increase tracking efficiency (until about 6 hits) and increase the fake track rate. This is shown by Figure 3.3. The former is mainly due to reduced nuclear interactions with the detector material, the latter due to increased ambiguities.

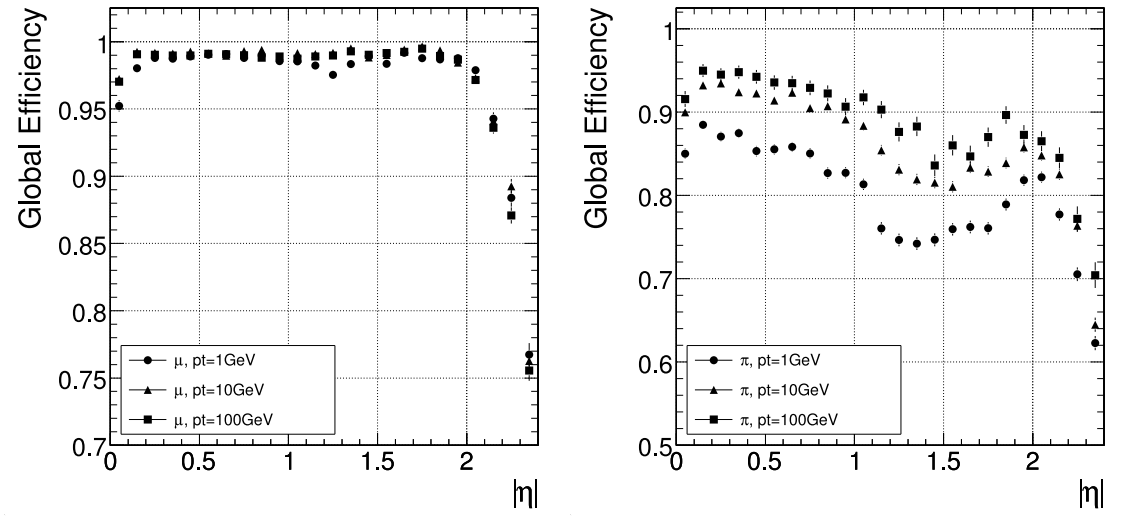

Figure 3.1: The muon (left) and pion (right) global track reconstruction efficiency for transverse momenta 1, 10 and 100 GeV/c. Global efficiency is defined as the reconstruction efficiency for all tracks including tracker acceptance, hit efficiency and pattern recognition efficiency. [30]

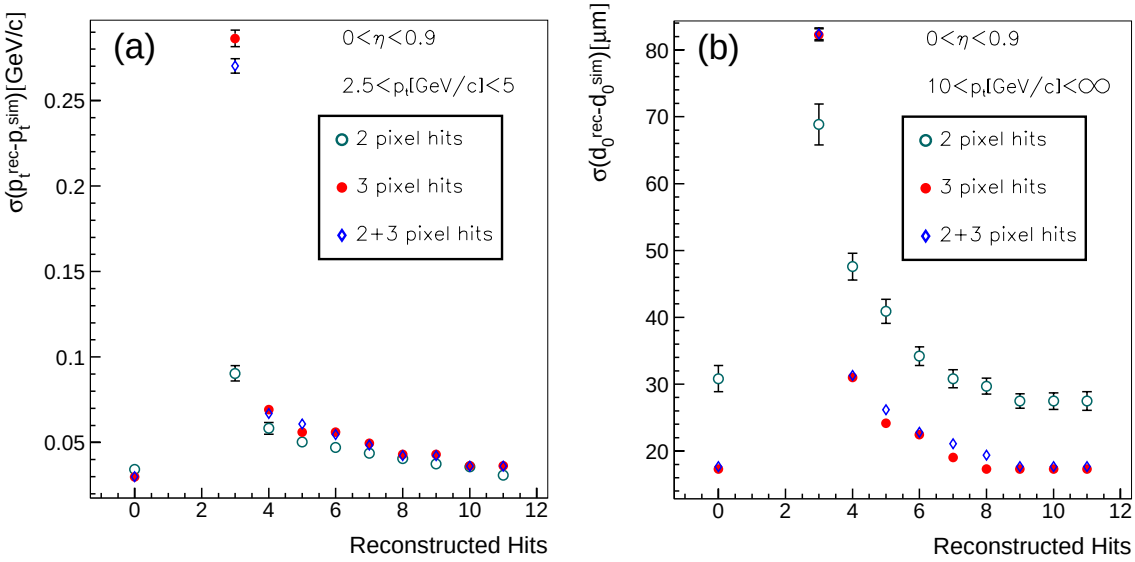

Figure 3.2: The p_T (a) and transverse impact parameter (d_0) (b), resolutions for muon tracks as a function of number of reconstructed hits. The left-most point, for 0 hits, represents the full tracker performance. [55]

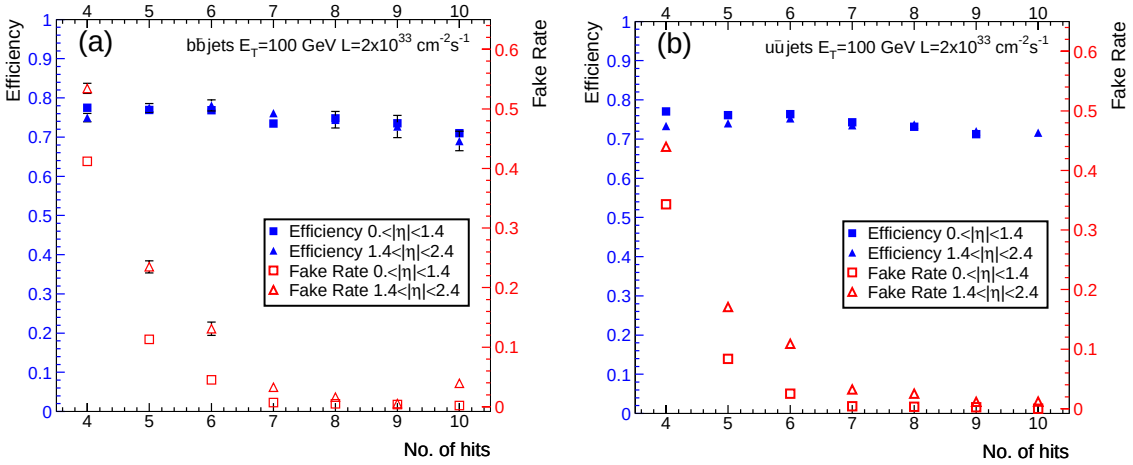

Figure 3.3: The track reconstruction efficiency and fraction of ghost tracks as a function of number of reconstructed hits for various η ranges. Tracks from b-jets and u-jets are shown in (a) and (b) respectively. [55]

3.3 Electron reconstruction

Electron showers deposit their energy in several crystals in the ECAL. The presence in CMS of material in front of the calorimeter results in bremsstrahlung and photon conversions. Combined with the electron's track curvature in the 4 T magnetic field, the energy reaches the ECAL spread in ϕ. These effects can hinder electron identification and energy resolution and have to be dealt with using sophisticated reconstruction algorithms.

A reconstructed primary electron is composed of a single track emerging from the interaction vertex, matched to an ECAL energy deposit or *supercluster*. A supercluster is a network of crystals designed to include a *seed* formed by the electron and peripheral clusters formed by radiated photons spread in ϕ. Two superclusterising algorithms, *hybrid* and *island*, are commonly used for the barrel and endcaps respectively [67]. The electron energy measurement is obtained by simply adding the contributions from the energy deposits in each crystal.

A strong correlation exists between the energy containment of the supercluster and the volume of the shower integration, measured approximately by the number of contributing crystals [68]. This relationship can be used to apply corrections to electron energies. For showering electrons, where a significant amount of energy is lost within the tracker material, a residual parametrised η-dependent correction must also be applied.

To include tracking information, dedicated track building techniques are required [68]. This is since the Kalman filter assumption of Gaussian fluctuations in track parameters at each layer is invalidated by the presence of bremsstrahlung radiation emitted along the electron trajectory. Instead, the Gaussian Sum Filter (GSF) is used, which is a nonlinear generalisation of the Kalman filter. To improve algorithmic efficiency on the HLT, the ECAL superclusters can be used to drive the track seed finding in the pixel detector.

From this basic electron object, features such as supercluster shape, ratio of energy deposited within the HCAL and ECAL and ratio of cluster energy and outer track momentum can be used to distinguish real electrons and fakes. The performance of this process depends highly on the nature of the considered background, however. More details on electron identification can be found in [30], [69].

3.4 Muon reconstruction

Muon reconstruction is performed in three stages: local, standalone and global [30], [70], [71]. For use with the HLT, muon candidates found by the level 1 trigger are used to seed local reconstruction, meaning only chambers compatible with the identified seed are reconstructed. Offline, local reconstruction is performed for the entire sub-detector. Track seeds are then constructed from patterns of segments reconstructed in the CSC and DT chambers.

Standalone reconstruction involves building tracks from the existing seeds. Tracks are built, from inside out, using the same Kalman filter technique described in Section 3.2.4. Despite the coarser spatial resolution, reconstructed hits from the RPC chambers are also included. In a similar fashion to SST track building, the tracks are then re-fitted through a backward Kalman filter: working outside in. Once the track has been built, it is extrapolated to the nominal interaction point and a final vertex constrained fit to the track parameters is performed.

Global reconstruction extends the muon tracks by including hits in the silicon tracker to the interaction point. The Kalman filter is used as before, which takes the magnetic field, multiple scattering and energy loss through material into account when projecting the track. This is especially important when finding the outermost hit in the SST - since the trajectory crosses the calorimeter and magnetic coil.

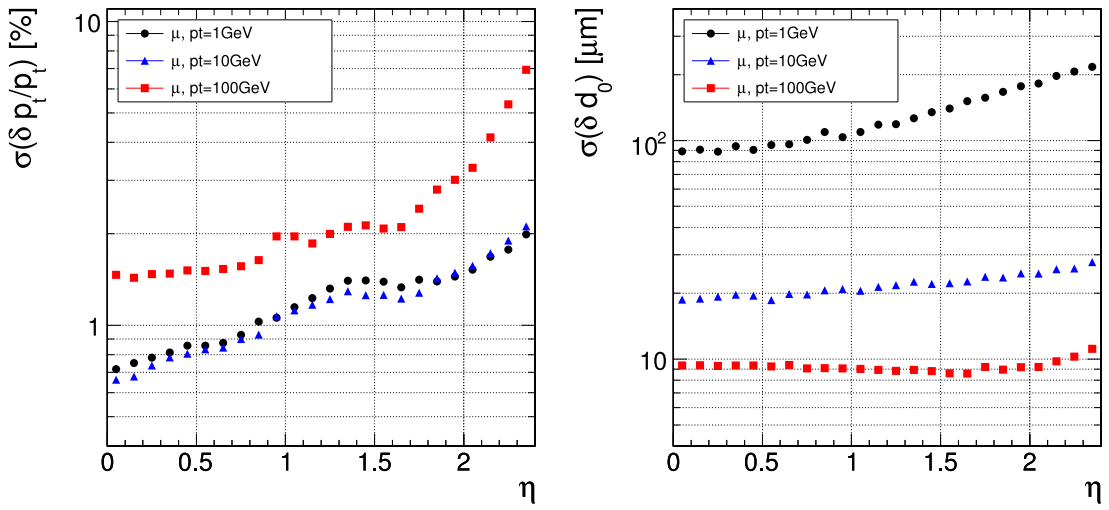

Figure 3.4: The p_T (left) and transverse impact parameter (d_0) (right) resolutions for muon tracks with transverse momenta 1, 10 and 100 GeV/c. [30]

Figure 3.4 shows the p_T and transverse impact parameter resolution for reconstructed muons as a function of η. For high momentum tracks (100 GeV/c) the p_T resolution is around 1 - 2 % up to $\eta \sim 1.6$. Beyond this it degrades due to the reduced lever arm. At 100 GeV/c, multiple scattering accounts for 20 - 30 % of the resolution, whereas at lower p_T it dominates. The resolution of the transverse impact parameter, d_0, is dominated by the first pixel hit at high p_T, resulting in a value of 10 μm.

3.5 Jet reconstruction

The basic element used to reconstruct hadronic jets is the calorimeter tower. Since the ECAL granularity is much finer than that of the HCAL, a tower combines ECAL and HCAL signals in an η,ϕ bin corresponding to a single HCAL cell. There are 4176 such towers in total.

Two schemes exist for combining energy measurements in calorimeter towers [72]. The energy scheme involves adding the constituent four-vectors to produce a massive jet. In the E_T scheme, the constituents are assumed massless and the measured E_T is equated to the sum of the constituents. Within this scheme, the jet direction is defined by $\eta = \Sigma E_{T_i}\eta_i/\Sigma E_T$ and $\phi = \Sigma E_{T_i}\phi_i/\Sigma E_T$, where i corresponds to a single tower.

Three principal jet reconstruction algorithms have been prepared for CMS: the *iterative cone*, the *midpoint cone* and the *inclusive* k_T. The latter two are widely used in offline analysis in current hadron collider experiments. The former, which is simpler and faster, is more suited to online reconstruction and the early phase of CMS data taking where a robust approach will be beneficial. This thesis will only use the iterative cone algorithm which is described below. A detailed description of the others can be found in [30], [73].

The iterative cone algorithm involves an E_T ordered list of calorimeter towers above an adjustable seed threshold. A cone of fixed radius, ΔR, defined by Equation 3.2 is formed around the highest E_T object and a *proto-jet* energy and direction is calculated within the E_T scheme. The cone is then adjusted to centre around the proto-jet direction. The procedure is repeated until the jet energy changes by less than 1 % between iterations and the proto-jet direction changes by $\Delta R < 0.01$. If a stable proto-jet is found, all calorimeter towers used to form it are removed from the list and the proto-jet is added to the list of jets in the event. The process comes to a finish when no more stable proto-jets can be produced.

$$\Delta R = \sqrt{(\Delta\eta)^2 + (\Delta\phi)^2} \tag{3.2}$$

Jet energy corrections can also be applied to restore equality between the reconstructed jet and its parton source. This calibration takes place in two steps. Firstly, the energy of the reconstructed jet is corrected to the energy of the particles produced by the hard scatter (which are subject to the same clusterising criteria as the reconstructed jet). This includes effects due to pile-up and residual electronics noise along with the intrinsic energy dependence of the detector response. An η and E_T dependent parametrisation is used by CMS to account for this correction [74]. Secondly, the energy of the particle level jet is corrected to that of the original parton. This accounts for particles that are not included in the particle level jet and is dependent on the hadronisation model.

3.5.1 b-jet identification

The identification of jets originating from a b-flavoured hadron relies upon the distinction of b-hadron properties from those of lighter flavours. Such distinctions include the large proper lifetime, large mass, high track multiplicity of the final

states, large semi-leptonic branching ratio and hard fragmentation function. The fastest and most robust *b-tagging* algorithm used by CMS and hence most useful for online reconstruction and the early phase of CMS data taking, is based on *track counting* [30], [75].

The track counting algorithm discriminates jets using the impact parameters of their constituent tracks. Firstly, tracks falling within a cone centred on a given jet are associated with it. A cone size of 0.5 is typically used for this purpose. The impact parameter for each track can then be computed in the transverse plane or in three dimensions, the latter providing a better overall b-tagging performance. To calculate the impact parameter of any track the following methodology is followed. Firstly, it is extrapolated back to the point of closest approach to the jet direction. This corresponds to the decay point of the B hadron. The track is then linearised (i.e. the tangent is taken) and the minimum distance of this new trajectory from the primary vertex is computed. This value provides a powerful means of determining whether a track originates from the primary interaction vertex or from the decay of a particle which has travelled a significant distance since then.

The impact parameters are also given a sign: a positive impact parameter means a positive displacement from the primary vertex, along the direction of the jet, a negative sign indicates the opposite. Cases of the latter arise most frequently from mismeasured track parameters, badly reconstructed primary vertices, pattern recognition failures and multiple scattering in the beam pipe. At this point, the tracks are ordered by decreasing impact parameter significance ($\frac{\text{impact parameter}}{\text{resolution}}$). If the Nth track in the list (where N is usually chosen to be 2 or 3) has a significance exceeding a given threshold then the jet is tagged as a b-jet.

3.5.2 τ-jet identification

The τ-lepton decays hadronically 65 % of the time, producing a τ-jet. If the p_T of the τ-jet is large compared with the τ mass, the hadrons will have a relatively small momentum transverse to the τ-jet axis. In 77 % of cases, the τ-jet consists of only one charged hadron, termed a *one-prong* decay. For these reasons, τ-jets are highly collimated and produce narrow energy deposits in the calorimeter.

Various discriminators are used both online and offline to identify τ-jets. A detailed description of these may be found in [30], [76]. The primary discriminator of interest

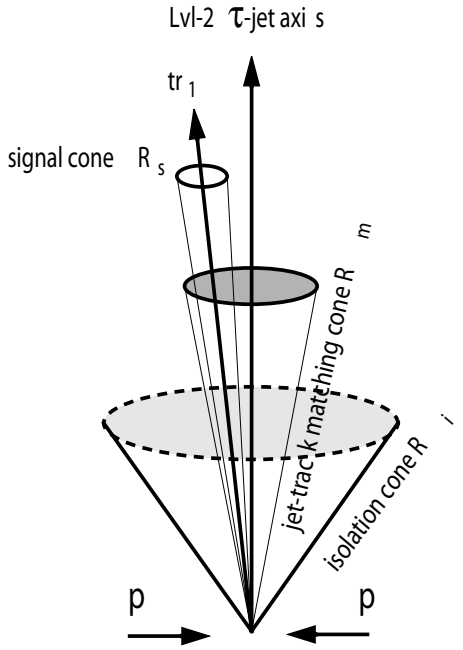

Figure 3.5: Illustration of a τ-jet identification technique using tracker isolation.

for this thesis, however, is tracker isolation. The principle of τ-jet identification using tracker isolation is shown in Figure 3.5. Tracks above a tunable p_T threshold, p_T^m, and falling within the jet matching cone, R_m, are considered in the search for signal tracks. Tracks within a narrow cone, R_s, of the *leading* (highest p_T) track and with an impact parameter $\mid z_{tr} - z_{tr}^{leading} \mid < \Delta z_{tr}$ are assumed to come from the τ decay. Tracks with a matching impact parameter and $p_T > p_T^i$ are then considered to fall within a larger cone R_i. If no tracks are found within this cone, except for those already within R_s, the isolation criteria is satisfied.

3.6 E_T^{miss} reconstruction

An imbalance in the E_T sum of the calorimeter towers, ${E_T}^{miss}$, indicates the presence of one or more energetic, weakly-interacting stable particles in the event. Measurement of ${E_T}^{miss}$ at the LHC will be complicated by the presence of pile-up collisions. In CMS, it is also expected to be hindered by the calorimeter resolution [77]. ${E_T}^{miss}$ corrections can be made by removing all reconstructed jet energies and replacing them with their corrected equivalents. Also, the combined E_T of reconstructed muons (which do not deposit their energy in the calorimeter) can be accounted for.

Chapter 4

Commissioning the CMS silicon strip tracker

Many physics analyses to be performed at CMS, such as $b\bar{b}Z^0$ production which is discussed in Chapter 6, are heavily reliant on the precise tracking of charged particles. With almost 10 million channels, the CMS strip tracker is the most complex detector of its kind ever built. Prior to any physics data taking, however, it will require thorough performance profiling and optimisation. This chapter describes the development of automated commissioning procedures for the strip tracker and their implementation during the sub-detector assembly at the Tracker Integration Facility, CERN.

4.1 Integration and commissioning at the TIF

Final assembly of the SST was performed in early 2007 at the Tracker Integration Facility (TIF), CERN. The facility provided an environmentally controlled clean room with infrastructure and services to power and operate as many as 1.6 M channels (cf. 9.6 M in the final system) using 82 FEDs and 11 FECs. It also provided a limited cooling capacity, sufficient to perform cold tests down to -15°C. The complete sub-detector (developed as TIB, TOB, TEC and TID sub-systems over a period of many years) was assembled, commissioned and operated over a period of several months. $\sim 5 \times 10^6$ cosmic ray events were accumulated during this time to validate detector performance. The commissioning software suite, developed over a timescale of several years, was also tested and with great success.

The tracker was one of the last sub-detectors to be installed within the underground experimental cavern at *Point 5* on the LHC ring. The TIF has allowed the tracker community to maximise its experience in operating the detector before the start-up phase and thoroughly validate and optimise the related software. There is far too much to detail here in full, so this chapter focuses on the author's own contributions to the commissioning effort.

4.2 Commissioning procedures

Many of the components used within the CMS tracker readout system are configurable by design. This is to manage the inherent spread in characteristics within the devices themselves; the sensitivity to temperature variations and the harsh radiation environment whose effects will vary with location and time. In addition, other non-configurable parameters require regular measurement such that the entire detector is characterised for a given *interval-of-validity*. The parameters of interest in this chapter are summarised in Table 4.1 *.

*Other configurables exist which are not commissioned within the CMSSW framework and so not discussed here.

parameter	units	beam requirement	number of channels CMS (TIF)	range	increment
control connectivity	-	no	37.5k (6135)	-	-
readout connectivity	-	no	37.5k (6135)	-	-
FED TTCrx delay	ns	no	3.5k (648)	0 : 25	1.04
PLL delay	ns	no	15.2k (2562)	0 : 25	1
AOH gain	V/V	no	37.5k (6135)	0.3 : 1.3	0.25
AOH bias	mA	no	37.5k (6135)	0 : 22	0.45
APVMUX gain	mA/MIP	no	37.5k (6135)	0.7 : 1.3	0.15
APV VPSP	mV	no	75k (12270)	-650 : +1250	7.5
strip pedestal	ADC	no	9.6×10^6 (1.6×10^6)	-	-
strip noise	ADC	no	9.6×10^6 (1.6×10^6)	-	-
HV	V	no	15.2k (2562)	0 : 500	1
L1 latency	ns	yes	1 (1)	0 : 4000	1
APV VFS	mV	yes	75k (12270)	-1250 : +650	7.5
APV ISHA	μA	yes	75k (12270)	0 : 255	1

Table 4.1: SST parameters to be configured, measured or validated by the CMSSW commissioning procedures. The list is in the order the parameters will be commissioned in (where order is relevant).

To keep the detector in an operational state for physics data taking, a large number of commissioning operations are required to configure, measure and validate the various parameters at regular intervals. To do this in a manageable time, and for consistency, automated procedures have been implemented. They can be coarsely divided into two categories: those that require LHC beam and those that do not. Those that do not are listed below in the order they should be performed.

During operation of the experiment, these procedures can be performed in between data taking periods to guarantee optimal detector performance in the following run. The frequency will depend on the type of procedure. For example, a connectivity run will only be performed once at the start, whereas pedestal and noise values will be monitored weekly, taking about 20 minutes. Complete detector commissioning without beam takes about 8 hours.

- **1. Control connectivity and detector partitioning.** All front-end devices are identified via the control network. It is also verified that no power cables have been swapped.
- **2. Readout connectivity.** Determination of the optical link connectivity between each multiplexed APV25 pair and the FED readout channels.
- **3. FED sampling.** Calibration of the FED TTCrx delay settings (which receive signals from the L1 trigger system) to find the optimum sampling time for each FED channel.
- **4. Internal front-end synchronisation.** Synchronisation of all APV25 clock and trigger signals through the PLL delay settings.
- **5. Optical readout links.** Calibration of each optical link through the LLD gain and bias settings.
- **6. APV25 multiplexer.** Calibration of the gain setting for each APVMUX.
- **7. APV25 data frame.** Adjustment of the data-stream baseline for each APV25 through its VPSP parameter. Measurement of the pedestal and noise of each strip.
- **8. High voltage.** Tuning of the depletion voltage (HV) for each sensor.

Procedures that do require LHC beam must be performed in dedicated calibration runs. Again, these are listed below in the order they should be performed. Most of this commissioning will remain valid for a long period of time, so these tasks will be run fairly infrequently. The global synchronisation, for example, will remain fixed as long as no modification is made to the logic of the trigger system.

- **9. Global front-end synchronisation.** Synchronisation of all APV25 clock and trigger signals to the LHC beam via the L1 trigger latency.
- **10. APV25 pulse shape.** Optimising the pulse-shape of each APV25 through its VFS and ISHA parameters.

The commissioning procedures listed above are those considered to be most crucial to the running of CMS. It is expected that more will be added in the future. The author has provided substantial input to the development of 2, 4, 5 and 7 under the guidance of R. Bainbridge (Imperial College London). These will be described in detail alongside results from the TIF in Sections 4.4 to 4.7. Although the effort for the remaining tasks came from others in the collaboration, they still rely on a core set of analysis and summary tools provided by myself. For this reason they may be discussed briefly.

Before specifics of these tasks are discussed, the central commissioning software framework will be reviewed.

4.3 Commissioning software architecture

The data acquisition software for the strip tracker is implemented within the XDAQ framework [78]. XDAQ is a platform for the development of distributed data acquisition systems. It also allows the commissioning of the detector through dedicated data acquisition *loops* over configuration parameters. Each loop is defined by a single run, interspaced with periods where the trigger is inhibited. During these gaps, the configuration of at least one device is incremented via the control system and the I^2C protocol. The configurations of several devices are usually tuned in parallel.

At the analysis level, various features of the APV25 data stream are used to reconstruct the appropriate *commissioning histograms*. At the end of the scan, commissioning *monitorables*, which reflect the optimised detector state, are extracted and

fed back into the online tracker configurations database for use in further commissioning or physics runs.

The first instance of the commissioning analysis software for the SST, called the *ROOT analyzer*, was developed for small beam-test and laboratory set-ups that yielded small event data sizes, typically < 1 MB. For this reason, a single process was sufficient for all events and the development proceeded on this premise. This architecture did not allow parallel event processing, however, and was therefore not scalable. The ROOT analyzer has been used to commission systems of up to 50 FEDs. Beyond this, a change in design was required.

In 2006, the commissioning software was migrated to the CMSSW framework. CMSSW provides many useful services and tools, such as the Data Quality Monitoring (DQM) framework which allows distributed histogram-based monitoring. In addition, CMSSW is installed on the filter farm, which runs the HLT algorithms during data taking. This provides hundreds of CPU units which are available for commissioning purposes in between runs. The CMSSW framework also provides access to the complete set of reconstruction algorithms developed for online and offline use. These have proven immensely useful in several cases. For example, when synchronising the detector to LHC collisions, the charge associated with reconstructed tracks is used.

The DQM framework allows remote *consumers* to subscribe to histograms that are defined and published by one or more *producers*. This idea is illustrated in Figure 4.1. The histograms from multiple producers can also be *collated* by a *collector*. For the purposes of commissioning, this permits real-time distributed analysis, with partial histogram building over multiple filter farm nodes, collation of information at regular intervals and monitorable extraction in a single dedicated consumer.

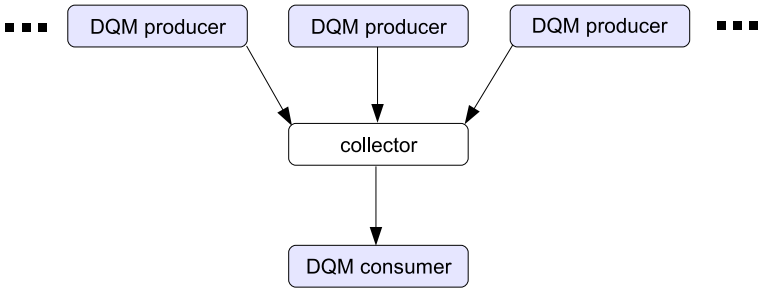

Figure 4.1: The DQM software model. Histograms are produced by multiple sources, collated at the collector and delivered to the consumer.

The commissioning software is illustrated in detail in Figures 4.2 and 4.3. Raw data are retrieved from the DAQ and added to the event by an input source. They are then unpacked into digi format (by the *raw-to-digis* module), giving digital hit information for each strip of interest. When operating the FEDs in raw or scope mode this corresponds to all strips; in zero suppressed mode, only those strips passing the FEDs zero suppression algorithm are included.

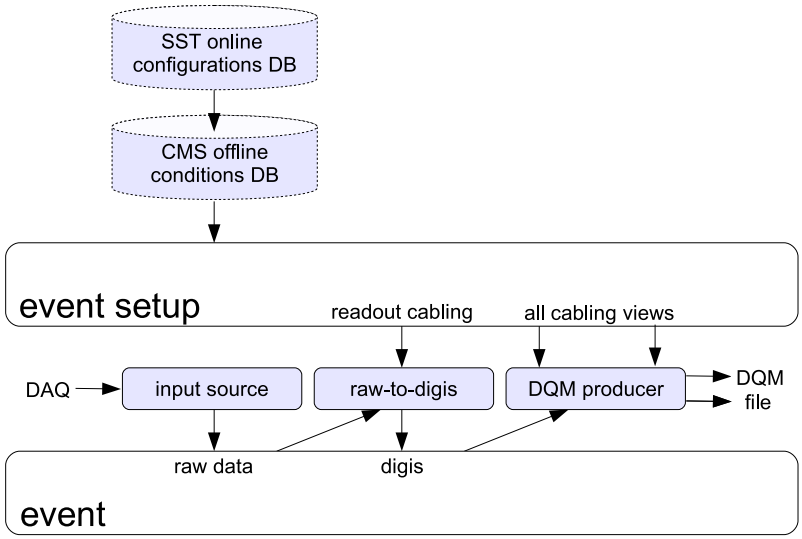

Figure 4.2: Commissioning software model for a DQM producer node.

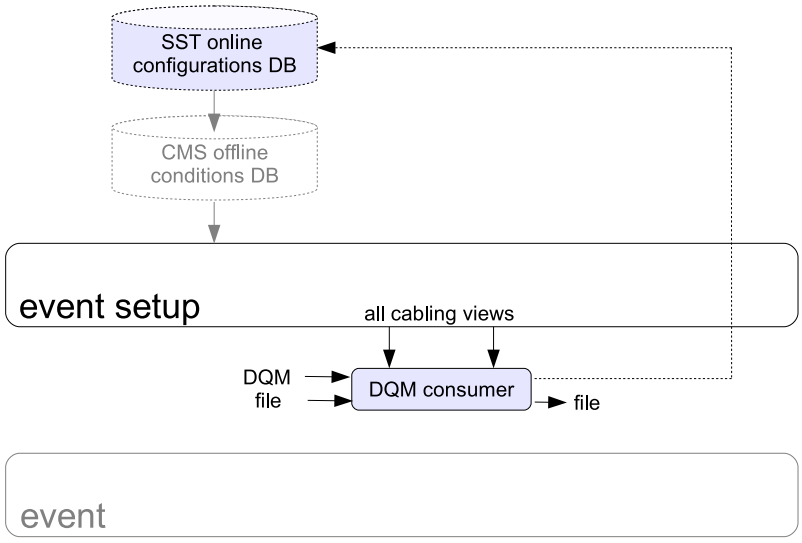

Figure 4.3: Commissioning software model for the DQM consumer (after collation).

Commissioning histograms, which are procedure dependent, are then defined, filled and published by a DQM producer. After DQM collation, a single consumer retrieves

and interprets collated commissioning histograms and extracts commissioning monitorables. It also optionally writes the collated histograms to file for later offline analysis. Monitorables can be presented in global summary plots for examination by the user or alternatively written to the online configurations database.

4.3.1 Extracting and summarising commissioning monitorables

This section describes the development of two software facilities, both implemented by the author, which are at the heart of the new commissioning software architecture.

The DQM consumer module, shown in Figure 4.3, is a CMSSW plugin that uses a separate analysis facility to interpret commissioning histograms and extract monitorables. The first instance of this software was designed for the ROOT analyzer and has therefore been developed and validated over a long period. To fully benefit from this experience, the new CMSSW algorithms were closely modelled on the old code.

One drawback of the old system was its lack of verbosity. Since it had been designed for small scale systems, unusual trends in the commissioning histograms could be identified by eye and dealt with accordingly. For full scale commissioning, however, this type of problem solving would be far too slow and unreliable. With this in mind, a detailed set of checks are now performed for each task to ensure the correct behaviour at each stage of the calculation.

To validate this new facility, it was run in parallel with the ROOT analyzer over small commissioning files from early beam-test setups. A one-to-one correspondence over all monitorables was achieved. Since then, the old ROOT analyzer has been phased out and the new analysis code used in its place. It has also been developed beyond its original form in response to experience gained by the tracker community at the TIF.

As with the ROOT analyzer, the new CMSSW analysis facility is dependent on ROOT only (i.e. no further software libraries are required). This was motivated by experience to ensure flexibility of use and light software management. It has also enabled offline commissioning analysis outside of CMSSW entirely, using histogram files produced by the DQM consumer. The ability to analyse offline has proved highly useful when facing DQM framework issues in the early phases of development. It also simplifies the debugging and validation procedures.

The ability to summarise hardware monitorables (optimised hardware parameters extracted by the analysis facility) is vital to the commissioning process. Establishing patterns of unusual behaviour can highlight defective hardware or problems with the procedures themselves. Though such software tools have been available in the past, they have been tailored for individual tracker partitions (TIB, TOB, TEC and TID) and so several have been developed in parallel. The new facility described here aims to incorporate the full functionality of each and therefore combine the experience of the full tracker community.

These summaries are presented in the form of user-defined histograms. Five parameters have been established to configure each histogram and can be coded in eXtensible Markup Language (XML), for simplicity of use. They are listed below.

- **Monitorable.** For each procedure more than one possible hardware parameter may be under scrutiny. This identifies which should be summarised.
- **Presentation.** The type of histogram to be used. Possibilities include sum, profile and scatter plots.
- **View.** This corresponds to the cabling view-point (discussed in detail in the text that follows) and identifies the ordering of modules in the x-axis of the summary.
- **System.** The region of the SST to be summarised. In general, the options depend on the cabling view. Examples include the entire system, the TIB, a single FEC-slot or a single FE-FPGA.
- **Granularity.** This defines the histogram binning. Commissioning monitorables can vary from the level of a single strip (e.g. strip pedestals and noise) up to the entire sub-detector (e.g. L1 trigger latency). Due to the vast number of strips, the summary facility allows users to set granularity as fine as a single APV25, or as coarse as the entire tracker. Levels in between depend on the view being considered.

When summarising the outcome of each commissioning analysis across a part of, or the entire detector, it is useful to define distinct cabling view-points. Each view defines a system of grouping modules together, usually according to some structure within the detector. Examining commissioning monitorables in the appropriate

view can make it much simpler to identify abnormal behaviour and effects which are common to modules of the same grouping.

Several such hierarchies exist. The two of most relevance to this thesis are itemised below.

- **Readout view.** Maps multiplexed APV25 pairs to their corresponding FED readout channel via the optical links. Problems common to a particular FED or FE-FPGA are easily diagnosed within this scheme. The readout hierarchy is as follows:

  FED ($\sim$440) $\rightarrow$ FE-FPGA (8 per FED) $\rightarrow$ FED channel (12 per FE-FPGA) $\rightarrow$ APV25-pair (1) $\rightarrow$ APV25 (2).

- **Control view.** Maps each APV25 chip within the control structure of the detector. This is especially useful for commissioning purposes as it can highlight effects due to the trigger and clock which are distributed through the control system. The control hierarchy is as follows, the numbers shown can vary across the detector:

  FEC-crate ($\sim$4) $\rightarrow$ FEC-slot ($\sim$11 per FEC-crate) $\rightarrow$ CCU-ring ($\sim$8 per FEC-slot) $\rightarrow$ CCU-module ($\sim$16 per CCU-ring) $\rightarrow$ CCU-channel ($\sim$10 per CCU-module) $\rightarrow$ LLD channel (3 or 2 per CCU-channel) $\rightarrow$ APV25 (2).

Other possibilities, which will not be discussed further, include the *detector view*, which maps each APV25 within the geometry of the detector. This can be useful when considering synchronisation effects due to particle time-of-flight, for example. It is the default cabling scheme used for event reconstruction and monitoring purposes. Also, the *cooling view* maps APV25s within the structure of the cooling system and the *HV view* maps them according to the distribution of power.

The sections that follow summarise various commissioning tasks of crucial importance for LHC startup. Where interesting, the algorithms used to extract monitorables will be reviewed. The summary histograms shown in these sections will also give an impression of how this facility is useful for *live commissioning*.

4.4 Readout connectivity

This procedure is used to map the connectivity of the optical links. These join each multiplexed APV25 pair with a single off-detector FED channel and are illustrated in Figure 2.10. In principle, each connection has been made in accordance with a readout scheme designed to minimise the variation in data rate across FEDs (for example, a single FED will not read out strips from only the highest occupancy layers near the beam pipe). In practice, with ∼38k links, mistakes may have been made.

By biasing the laser for each fibre high or low, unique patterns can be used to identify the corresponding FED channels. This procedure must be performed at the start of the commissioning process and is only required again if the cabling is altered. Figure 4.4 indicates that all 82 FEDs have been successfully detected. The variation in FED connectivity is also evident. FEDs are not always fully connected to optimise DAQ performance [55]. Usually this corresponds to 80 connections per FED or more. The low entries in the figure are due to the unusual configuration at the TIF.

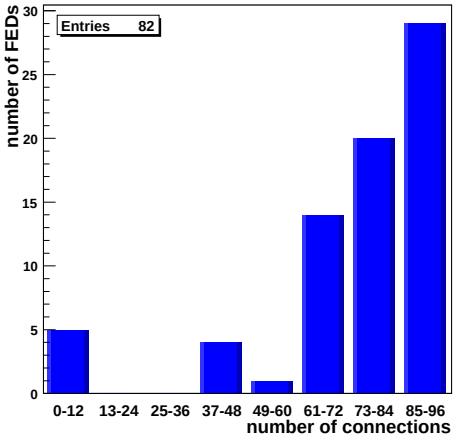

Figure 4.4: A summary of the number of connected channels for each FED at the TIF.

This procedure also serves as a diagnostic tool to identify abnormal light thresholds. When biased high, light levels should saturate the 10-bit ADCs. Those falling significantly below 1023 ADC counts can highlight dirty or broken fibres, badly responding front-ends or under-powered components.

It is important that the zero-light level is clearly visible within the FED ADC range. More specifically, it should never be set at zero ADC counts to avoid cutting off the

base of the signal. It should not be too high either, to maintain dynamic range and reduce noise effects. For this reason, each FED channel has a *TrimDAC* chip which raises the incoming analogue signal by 30 - 40 ADC counts[†]. When biased low, light levels should correspond to the TrimDAC threshold.

Figure 4.5 shows both high and low light level distributions for the entire system. 62% of connections recorded a high light level of 1023 ADC counts and 99% above 800. Any connections recording values below 800 could be studied. The mean low light level recorded, using a Gaussian fit, was 37.14 ADC counts. 0.15 % of connections lie above 100 counts and 1.1 % below 10. Likewise, these connections would require further study and possible re-commissioning of their TrimDAC settings.

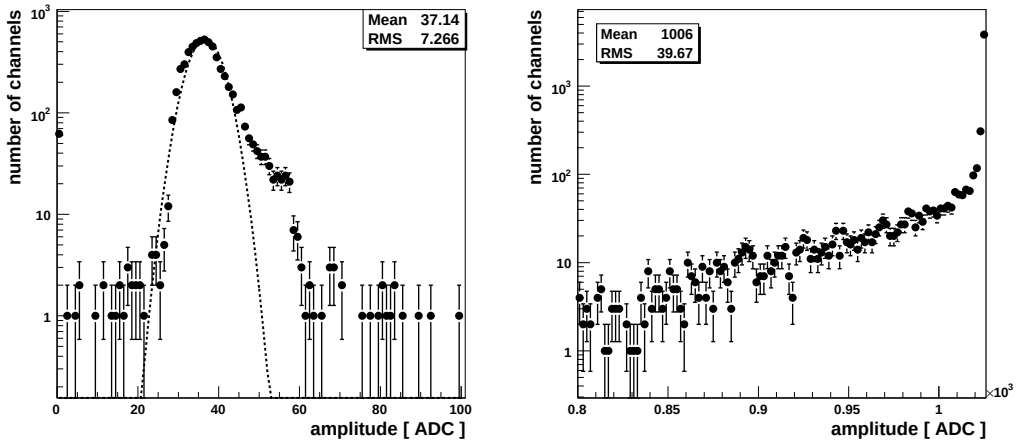

Figure 4.5: Digitised optical link readout with low (left) and high (right) biased light levels. The fit for the low bias distribution is Gaussian. The mean of this fit indicates the TrimDAC settings of the FEDs.

An example of poorly commissioned FED TrimDAC settings is shown in Figure 4.6. Two FEDs are shown for comparison. The low light levels for FED 067 are clearly less correlated than those of FED 068. A single FED channel exhibits low light levels as high as 300 ADC counts.

4.5 Internal synchronisation

A vital requirement of the SST is the synchronisation of each of its readout channels to LHC collisions. The end objective is to maximise tracking efficiency and minimise

[†]The TrimDAC settings are calibrated in an automatic procedure defined elsewhere.

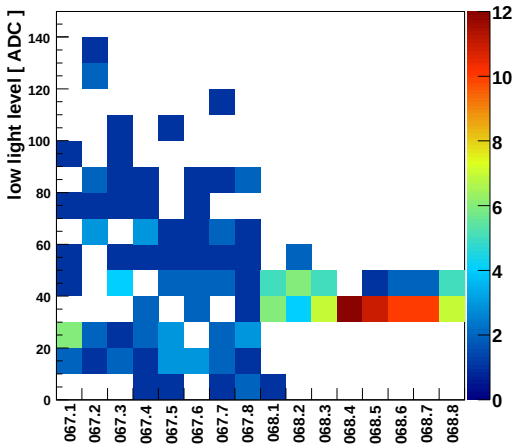

Figure 4.6: Digitised optical link readout with low biased light levels plotted against position in the readout cabling view (FED.FE-FPGA). FEDs 067 and 068 only are shown. The colour scale represents the number of entries in each bin.

fake hits from neighbouring bunch crossings. To ensure sampling the correct bunch crossing demands full synchronisation to much less than 25 ns. Within this interval, an asynchronous readout will result in sampling of the signal pulse shape off-centre and a loss in signal/noise of $\sim$4 %ns^{-1} in deconvolution mode. With this in mind a target precision of < 1 ns has been set.

For simplicity and efficiency, this procedure comes in three phases. This is a strategy originally proposed in [79]. Firstly, the entire sub-detector needs internal synchronisation of its readout to account for discrepancies in the propagation of clock and trigger signals to the front-end. This process is the subject of the current section. Secondly, a coarse synchronisation to the correct LHC bunch crossing is required to achieve the initial 25 ns target and finally a fine tuning to include differential effects such as particle time-of-flight. The last two steps are not possible without LHC beam and so were not performed at the TIF. A similar approach which synchronises readout to a cosmic trigger is used as a substitute in Section 4.8.

While receiving a clock signal, but in the absence of a trigger, the APV25 front-end readout chips generate dedicated analogue synchronisation pulses, or *tick-marks*, every 1.75 μs. Tick-mark data streams from pairs of APV25 chips are then multiplexed into a single readout channel and digitised by the FED. The ticks are output following the arrival of a reset signal at each APV25, making them a precise probe to determine the relative timing offsets.

The reconstruction of a multiplexed APV25 tick-mark is shown in Figure 4.7. This is an example of a commissioning histogram. Each filter node would own a subset of the statistics before it is collated into its final form at the DQM consumer. For the tick-mark data stream to be sampled by the FED, the latter needs to be operated in scope mode. To produce the curved shape of the tick-marks, however, the average amplitude of various positions along the tick must be measured by scanning through the programmable delay settings present on the TTCrx ASIC leading into each FE-FPGA [‡].

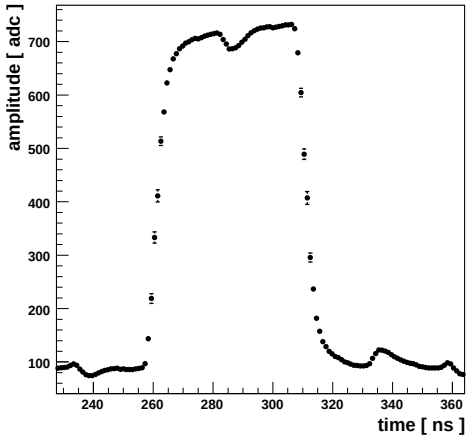

Figure 4.7: A reconstructed multiplexed tick-mark. The measurements were taken after the bias and gain commissioning of the optical links. In this example, the tick base sits at ∼100 ADC, the top at ∼700 ADC and the rising edge at 262 ns.

Once a multiplexed tick mark has been reconstructed for each channel the commissioning monitorable is extracted. For the case of synchronisation, this is the rising edge of the tick, defined as the position of the maximum derivative, measured relative to a predefined reference time. Discrepancies in rising edge times between readout channels are due to the differing analogue optical fibre lengths connecting each APV25 to its FED, combined with the staggered position of each module within the control structure. The former can be subtracted assuming the length of the fibre is known. The readout can then be synchronised by imposing a compensating delay through the programmable PLL chip positioned on the front-end hybrid.

The effect of clock propagation time through the control network is clearly seen in Figure 4.8. On the left, data from a single FEC-slot are shown, corresponding in

[‡]This is the same delay setting used in commissioning task 3 to tune the sampling time of each FED channel.

this case to 8 CCU-rings. The linear decrease in delay settings across the CCU-rings corresponds to the uniform fibre lengths connecting each neighbouring ring via their DOHs.

On the right of Figure 4.8, data from the full system at the TIF are shown. The PLL delay setting for each CCU-module within its ring is plotted. The downward sloping trends indicate that the highest index modules receive control signals later and are therefore closer to the chosen reference time. TEC modules receive the latest trigger signals and therefore require the greatest compensation. Three outlier modules are clearly visible in the TEC system (CCU index 2). These have been left underpowered for the purpose of illustration. Again, this highlights the usefulness of commissioning summary histograms, binned in a sensible order, in understanding a large scale detector.

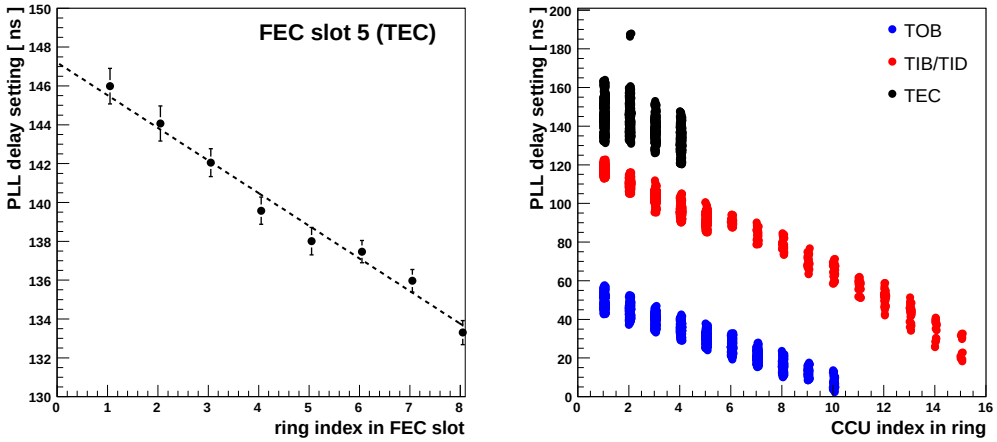

Figure 4.8: (left) The mean tick mark rising edge time (or required PLL delay setting) for each CCU-ring connected to one FEC-slot. The error bars show the corresponding standard error. The fit is linear. (right) Tick mark rising edge times for all modules at the TIF, as a function of the physical position of their local CCU-module within its ring. In both cases, measurements were made before internal synchronisation.

Figure 4.9 shows the internal synchronisation of the full TIF system after the commissioning process. Distributions from the TOB, TIB/TID and TEC partitions are very similar. Tick marks from only 2 modules were recognised by the analysis software, described in Section 4.3.1, as having an unusual structure. In each case, this was due to incomplete tick reconstruction resulting in no rising edge. These were flagged and are not shown in the final distributions. The combined Gaussian fit indicates a standard deviation of 0.68 $\pm$ 0.01 ns. This is within the desired 1 ns window and is encouraging for the final system.

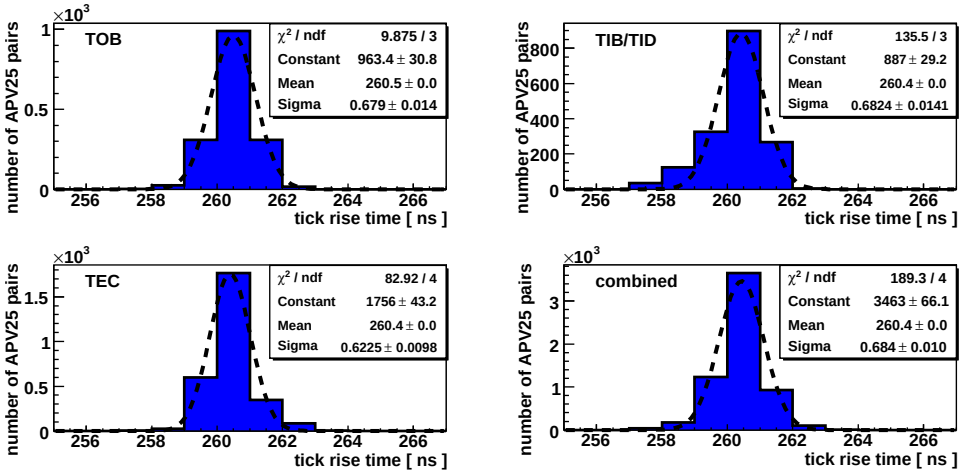

Figure 4.9: The distribution of tick mark rising-edge times after synchronisation. The fits are Gaussian.

There is a final complication to this process. During a synchronisation run, as well as the rising edge of each tick mark, the tick top and base levels are also monitored. These are used to tune the frame finding thresholds of each FED channel. Commissioning the optical links is discussed in detail in the next section, but the process of adjusting optical gain will affect the measurement of the height of each tick. Without a synchronised system, however, the process of event building will break down and no further commissioning tasks are possible. Experience has shown that the best solution to this conflict is to perform at least two iterations of the internal synchronisation procedure, one before optical link commissioning and one after.

4.6 Optical readout links

Each of the ∼38k optical links requires configuration of its AOH to achieve optimum transfer performance prior to physics data taking. There are two parameters that can be adjusted to alter the properties of each link. The bias current of the laser transmitter provided by the LLD and the LLD gain setting. The effect of scanning through various bias and gain settings on the link output is to *shift* and *stretch* respectively, the data frame within the FED ADC range. This is illustrated by the commissioning histograms for this task, shown in Figure 4.10.

There are two basic considerations that guide the analysis algorithms in the gain and bias selection for each link. Firstly, the overall gain of each readout link should

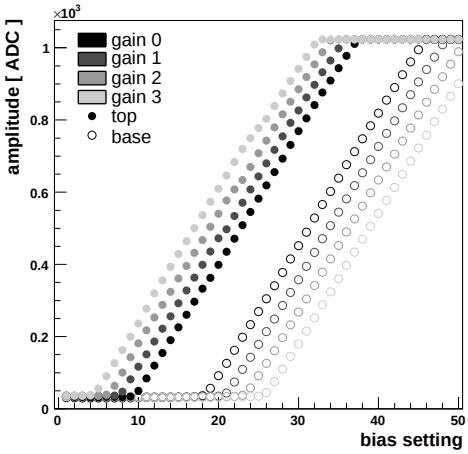

Figure 4.10: Multiplexed tick mark base and top amplitude as a function of LLD bias and gain settings. An example of a reconstructed tick mark is shown in Figure 4.7, with the top and base levels highlighted.

be as close as possible to 0.8 V/V [80], which corresponds to being able to view 3.2 MIP signals in an 8 bit dynamic range (1 MIP = 80 ADC) [31]. This is to optimise the use of the dynamic range available without compromising the integrity of i.e. *clipping* the signal. Also, maintaining an approximately constant gain throughout the tracker ensures uniform performance across all detectors. To equalise the links in this way, the LLD was designed with four gain settings uniformly distributed in the range 0.3 - 1.3 V/V[§]. Despite this, each LLD is unique and a small amount of gain spread is still anticipated.

The physical gain of each link can be measured from the ratio between the initial (800 ADC counts) and final tick heights. The final tick height can be inferred from Figure 4.10 as the shift in ADC counts from the tick-base to the tick-top. In the general case that the top plateaus before the base *lifts-off*, the top can be projected beyond its cut-off. This process is described in detail in [80].

The gain measurements made in this way are shown in Figure 4.11. On the left is a simulation of the full tracker at -10 °C. On the right are data taken from the TIF at +20 °C. The distributions are similar for both figures. In each case, the range in LLD gain measurements is between 0.64 and 0.96 V/V. Gain 1 is also the dominant setting, being optimal for more that 85 % of the channels. This is desirable as it allows for flexibility to compensate for both low and high gain links. For the data taken at the TIF, there is a slight shift towards the higher gain settings due to the higher environmental temperatures.

[§]The actual gain settings are temperature dependent [80].

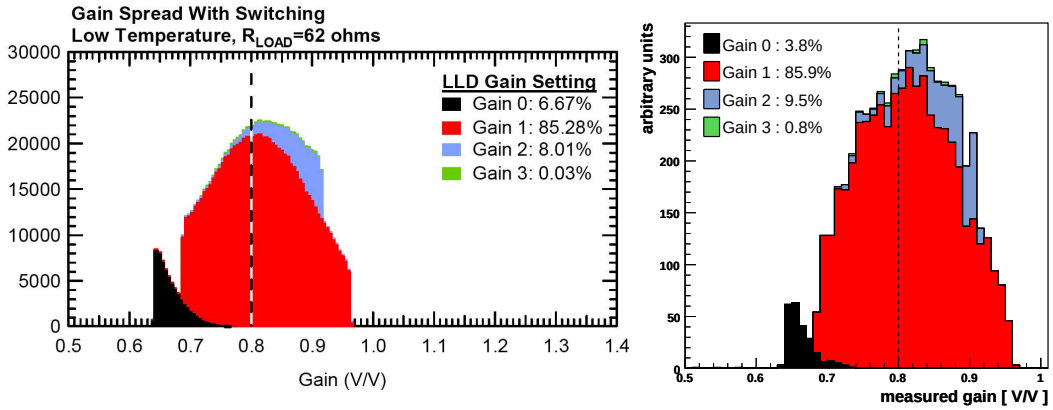

Figure 4.11: The distribution in measured LLD gain when commissioning to a target of 0.8 V/V. Contributions from each gain setting are shown. Both simulation at -10 ˚C (left) [80] and measurements taken at the TIF at +20 ˚C (right) are shown.

Secondly, the laser bias must be set just after the base lift-off: i.e. at the lowest setting where the base of the tick is clearly visible. This is to ensure it is high enough to avoid clipping signals from the bottom and to avoid falling below the lasing threshold during operation with temperature fluctuations. Furthermore, it must also be low enough to avoid compromising the top of the signal and to benefit from the better noise performances and the linearity of the receiving amplifier in this region.

Figure 4.12 shows the distribution of bias settings (left) and laser thresholds (right) at the TIF. The mean bias setting increases with each gain increment. Figure 4.10 shows the base will lift-off at higher bias settings as the gain increases, so this is to be expected. The laser threshold for each LLD is taken at halfway between the base and peak lift-off points ($\times$ 0.45 to convert to mA). Again, Figure 4.10 shows the base and peak lift-off points diverge fairly symmetrically as gain is increased, so the corresponding laser thresholds should be similar.

4.7 APV25 data frame

The output data-frame of the APV25, illustrated in Figure 2.12, contains a sub-structure which is unique to each chip. The mean levels of the constituent strips, known as *pedestals*, are defined in Equation 4.1. Here, p_i is the pedestal value for the ith strip and a_i is the corresponding sample recorded at time t.

$$p_i = \langle a_i(t) \rangle \tag{4.1}$$

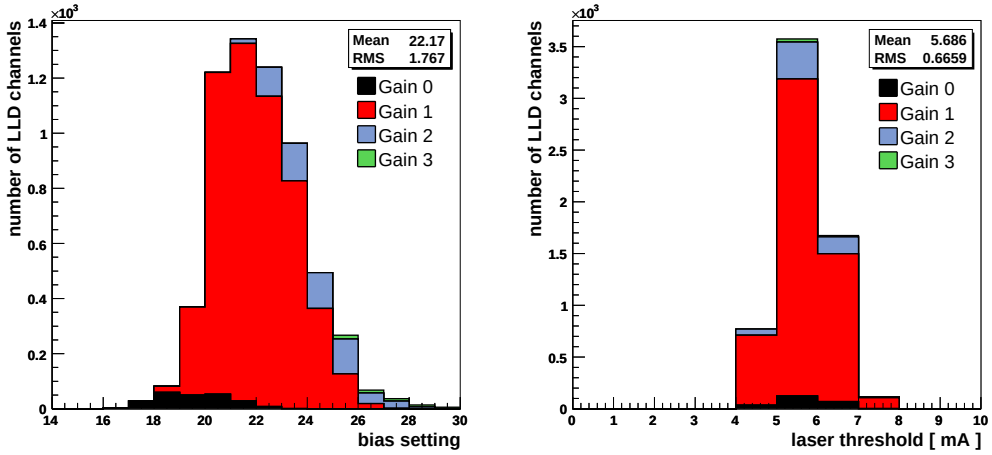

Figure 4.12: (left) The distribution of bias settings at the TIF after gain configuration. (right) The laser threshold distribution after gain and bias configuration. Both histograms are stacked according to gain setting.

Pedestals are the result of static DC levels observed on readout and are expected to remain constant for long periods of detector running. The chip *baseline*, b, is defined as the mean of the constituent pedestals, shown by Equation 4.2.

$$b = \frac{\sum_i p_i}{128} \tag{4.2}$$

These measurements can be combined to calculate a strip *residual*, r_i, through Equation 4.3. Residuals indicate systematic pedestal shifts from the baseline.

$$r_i = p_i - b \tag{4.3}$$

Although pedestals are expected to remain constant for long periods, the shape of each data frame will still vary from event to event. This is due to fluctuations at both the chip and strip levels. The former is known as *common mode*, $m(t)$, and is defined by Equation 4.4.

$$m(t) = \frac{\sum_i a_i(t)}{128} - b \tag{4.4}$$

Fluctuations in common mode are called *common mode noise*, $\sigma_{m(t)}$, and are usually attributed to a source external to the chip itself. The *strip noise*, $\sigma_{n_i(t)}$, comes from fluctuations in the DC strip levels themselves, $n_i(t)$. For a single strip, the combination of these two effects, called *raw strip noise*, is taken as the standard deviation of the sample distribution, σ_{a_i}. It is shown by Equation 4.5.

$$\sigma_{a_i(t)} = \sigma_{n_i(t)} \oplus \sigma_{m(t)} \tag{4.5}$$

Finally, the signal recorded by charged particles in the active material will be superimposed onto the pedestal and noise levels, in the manner shown in Equation 4.6. Here, s is the signal produced by the active material.

$$a_i(t) = m(t) + p_i + n_i(t) + s_i(t) \tag{4.6}$$

4.7.1 Analogue baseline calibration

The baseline, described by Equation 4.2, can be configured by adjusting the *VPSP* setting of the corresponding APV25. Figure 4.13 (top left) shows an example of a commissioning histogram for this task. The level is adjusted to maximise the dynamic range, minimise power consumption and ensure signal stability. If the level is too high, large signals may exceed the ADC limit of the FED. The chip power consumption is also dependent on baseline level, with a ∼30 % increase measured across the baseline range [81]. If the level is set too low, however, the risk of cutting the signal from below increases due to large negative common mode fluctuations. This can be especially problematic for highly ionising particles which can introduce large non-linear shifts [82].

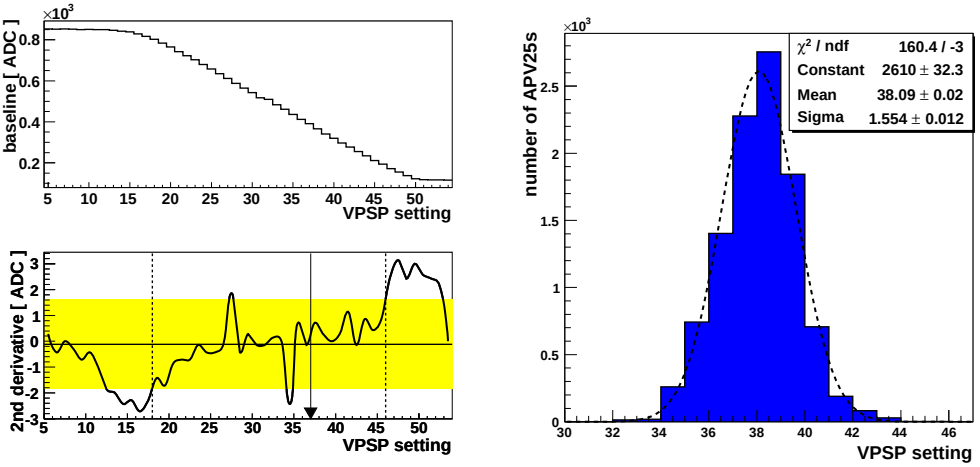

Figure 4.13: (top left) Measured baseline for a TOB APV25 as a function of the VPSP setting. (bottom left) Illustration of the optimum VPSP setting calculation for the same TOB APV. The solid line is the second derivative of a smoothed version of the baseline histogram. The horizontal line represents the mean second derivative and the yellow band ± 2σ from this. Second derivatives outside this region are identified as turning points. (right) The distribution in VPSP settings across all APV25s at the TIF. The fit is Gaussian.

The optimal baseline position has been fixed at an offset of 33 % of the baseline range, from the minimum allowed value. Figure 4.13 (bottom left) illustrates the

calculation of this optimal value for a single APV25. It is also an example of how the commissioning software searches for unusual hardware behaviour in an automated way. Firstly, the commissioning histogram is smoothed by taking a mean value of each bin and its two closest neighbours. The second derivative of this noise-reduced version is shown as the solid line in the Figure. The truncated mean (considering only the middle 10-90 % values in an ordered list) is shown as the horizontal line and $\pm$ 2σ as the yellow band. Values exceeding this band are considered to be *turning points* in the trend.

At this point some basic checks are performed. At least two turning points should exist. The maximum and minimum should be of opposite sign and in the correct order (the minimum first). The example shown is interesting as two other turning points seem to be present. These correspond to slight bumps in the diagonal trend of the source histogram. They are a common feature of this analysis and can be hazardous. Their amplitude and width are almost always smaller than those of the points of interest, however. The chosen VPSP setting for this example is highlighted with an arrow.

On the right of the same figure is the distribution of settings for all APV25s at the TIF. The Gaussian fit shows a standard deviation of less than 2 settings, which is very small for such a large number (over 12k) of devices.

4.7.2 Pedestals and noise measurement

To identify signal (plus strip noise), Equation 4.6 indicates that the pedestal and common mode noise contributions must first be subtracted. This operation is performed within the FED firmware algorithms. Furthermore, the FED zero-suppression algorithm requires the noise of each strip for use as a discriminator in the search for signal clusters. The measurement of these three sets of values (pedestals, noise and common mode noise) is the subject of the remainder of this section. Inaccurate measurement will result in poor signal extraction and ineffective zero-suppression. The former will mean a reduction in signal/noise and therefore a decrease in tracking efficiency. Fake rates would also be expected to increase. The latter would involve the loss of entire signal strips, reducing tracking efficiency, and inclusion of unnecessary strips meaning an increased data rate. Clearly increases in data rate must be compensated for by lowering L1 trigger rates. The measurement of these configuration parameters therefore has a direct impact on all physics studies.

Figure 4.14 shows pedestal, raw noise, noise and residual measurements for a single APV25 pair. The residual gradient is present due to a voltage droop across the chip. The pedestal distribution across all channels, shown in Figure 4.15 indicates a mean value of 380 ± 50 ADC counts. Pedestals which are fixed at abnormally high or low values can be used to identify faulty strips.

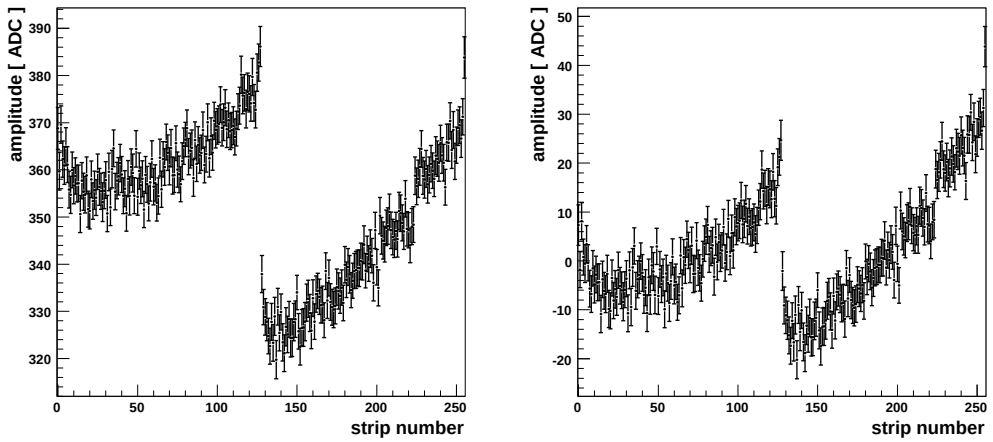

Figure 4.14: (left) Strip pedestals and raw noise shown by the y-values and their error bars respectively. (right) Strip residuals and noise shown by the y-values and their error bars respectively. Both are for a single multiplexed APV25 pair.

Figure 4.15 also shows the strip noise as a function of the tracker's control structure. Variations between TIB, TOB and TEC partitions are due to differing strip dimensions. The mean strip noise for the entire detector is 4.5 ± 0.8 ADC counts.

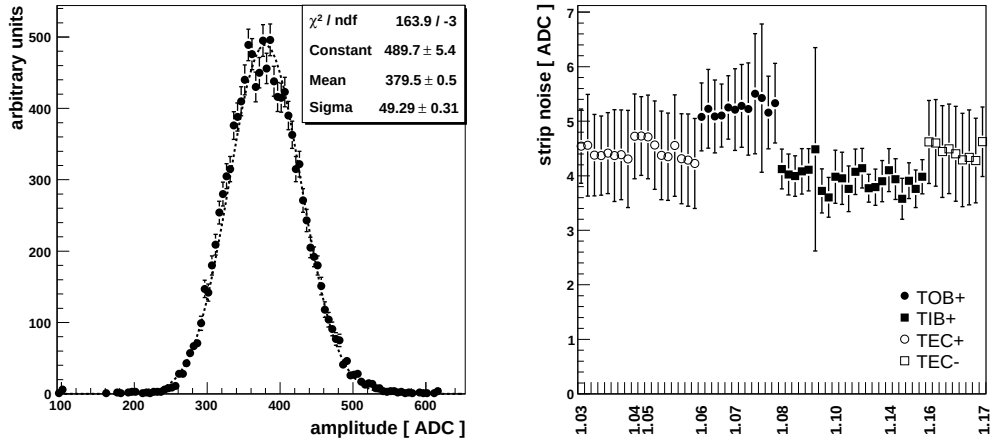

Figure 4.15: (left) The pedestal distribution for all strips. The fit is Gaussian. (right) The noise for all strips in the control cabling view (FEC-crate.FEC-slot).

Noise measurements can also be useful to diagnose *dead* or *noisy* strips: identified at less than and greater than 5σ from the mean noise respectively. These are also catalogued in the online configurations database and considered when building strip clusters. Within the TIF, 0.2 % of strips were identified as bad.

Figure 4.16 shows the m(t) + b (common mode + baseline) distribution for a single APV25. The standard deviation in this measurement is the common mode noise, measured to be 0.9 ADC counts. This is about 20 % of the mean strip noise. The points follow a Gaussian distribution with a high common mode tail.

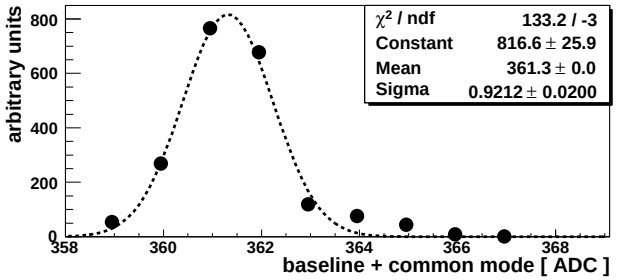

Figure 4.16: Baseline + common mode measurements for strips 0-127 of the multiplexed APV25 pair in Figure 4.14. The fit is Gaussian.

One key aim of TIF operation was to monitor the stability of these parameters over time. Any large variations between commissioning runs will have an impact on data quality. It has been shown that the stability is excellent with the only observed variations within 6 months of running corresponding to steps in temperature [83].

4.8 Synchronisation to the LHC

As described in Section 4.5, once the tracker has been synchronised internally, synchronisation to within 1 ns of LHC collisions is vital to optimise the signal/noise and minimise fake hits recorded from neighbouring bunch-crossings.

Firstly, a coarse global scan is performed in 25 ns steps to allow the correct bunch crossing to be located. The total charge deposited in the tracker is counted for each setting and the correct latency inferred from the peak. A fine delay scan is then performed in 1 ns steps with a similar track-based approach [84].

Since the TIF was recording cosmic data, no such bunch structure was available. It was still possible to test the fine delay approach, however, using a scintillating

cosmic trigger. The measurement was performed using both the TIB and TOB partitions, due to their favourable orientation for vertically propagating tracks, and is presented in Figure 4.17.

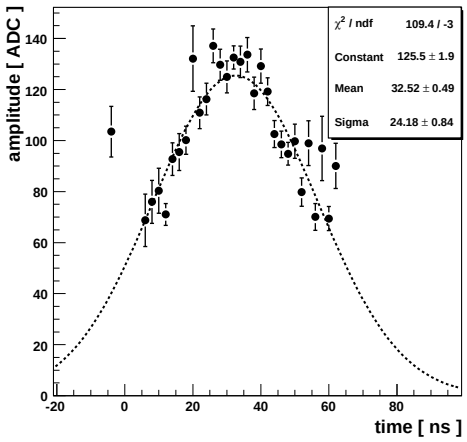

Figure 4.17: The mean ADC count in triggered cosmic muon events as a function of the trigger latency correction. Here, ADC counts are taken from the clusters used in reconstructed muon tracks.

The mean number of ADC counts from tracked clusters is shown as a function of trigger latency correction. Corrections for time-of-flight are also performed. A correction of +32.5 ns was added to the pre-existing L1 latency. The standard error in this measurement is ± 0.5 ns. Combined with the dominant internal synchronisation spread of ± 0.7 ns gives a total standard deviation of ± 0.9 ns. For the setup at the TIF, this is within the predefined target of 1 ns.

4.9 APV25 pulse shape tuning

The configurations of registers on all front-end APV25 readout chips are tuned to optimise the signal amplification and shaping. This is vital to maximise the signal/noise, enable clear bunch-crossing identification and to ensure homogeneity of the detector response. Also, by controlling the time behaviour of the signal, one ensures the response of out-of-time particles is the same as in simulations.

Since the TIF was equipped with a cosmic trigger, it was possible to take data containing actual signal clusters. This operation was performed in peak mode as the longer pulse is easier to compare with an optimal reference curve. Excellent

signal/noise values were recorded. Values are dependent on the partition under study and the corresponding strip dimensions, but were typically near 30. They were also, crucially, stable with time.

4.10 Summary

The CMS SST is the most complex sub-detector of its kind ever built and vital to the physics aspirations of CMS. It was assembled at the Tracker Integration Facility, CERN. Large-scale tests were also performed at the TIF to prepare for commissioning and data handling with the full system at Point 5. During this time it was demonstrated that the data acquisition software is stable and the architecture is scalable. Automated commissioning procedures, developed within the CMSSW framework and using generic data quality monitoring tools were also tested and used to calibrate 1.6 M channels and synchronise their readout of cosmic muon data to within 1 ns.

Chapter 5

Local reconstruction software for the CMS silicon strip tracker

One of the most CPU-intensive tasks within the CMS High-Level Trigger is the reconstruction of hits from charged particles propagating through the SST. These hits are then used in the formation of tracks. Even at low luminosities, the average processing time for the standard reconstruction chain is 5.5 s, which overwhelms the total HLT budget of 40 ms. A new scheme, optimised for speed and tracking performance, has been developed to reconstruct hits within tracking regions-of-interest only. This chapter describes the new approach in detail.

5.1 CMS event data

The CMS detector is unprecedented in terms of its size and complexity. For high luminosity running (10^{34} cm^{-2}s^{-1}) online event data sizes of 1-2 MB are expected [30]. Raw data will contribute a significant fraction of this. A summary of the contribution of the most dominant sub-detectors is shown in Table 5.1. Although the pixel detector has almost a factor 10 more channels than the strip tracker, larger occupancies in the latter result in almost 7 times the data rate.

5.1.1 The HLT time budget

During the first year of operating at low luminosity (2×10^{33} cm^{-2}s^{-1}), the L1 trigger is expected to provide an input rate of 50 kHz to the HLT farm. The HLT

Sub-detector	Channels	Event data size (kB)
Pixel tracker	6.6×10^7	70
Strip tracker	10^7	500
ECAL	7.6×10^4	100

Table 5.1: Raw event data summary for the dominant sub-detectors at high luminosity running.

farm will comprise ∼2000 CPUs, which results in an average time quota of 40 ms (2000 / 50 kHz) per event for each HLT processing node [55]. This must include raw data unpacking, event reconstruction and the trigger decision itself.

5.2 Local reconstruction software architecture

The original CMSSW software design for the local phase of track reconstruction (hit formulation) is shown in Figure 5.1. The raw data produced by the tracker FEDs and combined into a single data packet by the DAQ infrastructure were added to the CMSSW event by a dedicated module (the input source). The remaining CMSSW modules shown in the figure performed the unpacking (raw-to-digis), clusterising (digis-to-clusters) and hit production (clusters-to-hits). Each interacted directly with the event and event setup. The structure and functionality of these algorithms are discussed in more detail below.

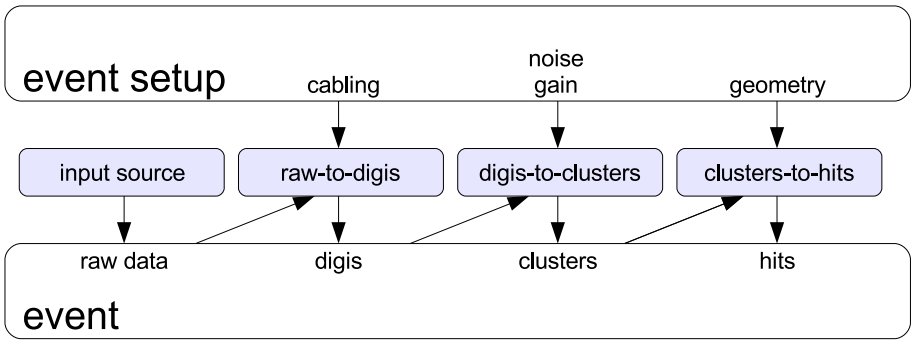

Figure 5.1: The software design for local reconstruction in the SST. Boxes represent CMSSW framework objects. Arrows represent the flow of data. The shaded boxes are plugin framework modules that perform a single event reconstruction task. Here *raw-to-digis*, *digis-to-clusters* and *clusters-to-hits* represent unpacking, clusterising and hit-production processes respectively.

The unpacking of raw data into digis was performed by the raw-to-digis module. Crucially, this module was capable of interpreting raw data from the FEDs operating

in all readout modes: virgin raw, processed raw, zero-suppressed and scope (the reader is referred back to Section 2.4.4 for a description of these). Although zero-suppressed mode will be universally used for physics runs, the other modes are essential for commissioning purposes and problem solving during early runs with low trigger rates.

The raw-to-digis module was also developed to handle data readout via two different FED output channels: either the 64-bit *S-LINK64* connections or the 32-bit *VME* [52],[85]. The latter, which allow data rates of only ∼10 Hz, were essential before the full DAQ infrastructure became available and are still vital for debugging purposes. Though the information contained in both types of data packet are identical, the 64-bit *words* that each raw event is built from are 32-bit swapped when using the VME path. This is for reasons of code legacy, but must be handled correctly during the unpacking process.

On a final note, since data are being extracted at the bit level, a small bug in the unpacking code can mean the completely incorrect identification of strip signals without any subsequent software crash. For example, if the incorrect readout path is assumed (VME instead of S-LINK64 perhaps), registers used to identify the status of the tracker could be identified as signal! For this reason, a large number of data quality checks were performed during the unpacking process.

The clusterising of unpacked digis into neighbouring groups was performed by the digis-to-clusters module. The definition of *cluster* depends on the algorithm in use, but the *three-threshold* description, defined in Section 3.2.1, is the most widely chosen and has been used throughout this thesis. The three thresholds are configurable by the user along with the number of allowed dead strips within a single cluster. Values chosen here (and also the default) were $t_1 = 2.0$, $t_2 = 3.0$ and $t_3 = 5.0$ with no dead strips allowed.

The production of hits which can be used by the tracking algorithms was performed by the clusters-to-hits module. This process involved inclusion of the geometrical position of each cluster and the matching of neighbouring clusters on stereo tracker modules.

The first version of the raw-to-digis algorithm was developed within the ORCA framework by R. Bainbridge (Imperial College London). The author of this thesis was then responsible for its migration to CMSSW and subsequent development. The digis-to-clusters and clusters-to-hits modules were developed by third parties within

the CMS collaboration*. Also, unless explicitly stated otherwise, all performance measurements presented in this chapter were made solely by the author.

5.2.1 Software performance

All studies in this chapter were performed in version 1.6.0 of CMSSW. Unless stated otherwise, times quoted are wall-clock times measured on a dual core AMD Opteron^{TM} processor 248, with a 2.14 GHz CPU and a 1 MB cache.

The software performance can be measured in two ways: through its timing cost to the HLT and through direct validation of the hits produced. For the purpose of validation, the author and R. Bainbridge developed a test facility that adds fake raw data to the event. This consists of randomly spaced clusters throughout the detector. This uniform spacing is unphysical since, in a typical event, a higher cluster density would be found in the innermost tracker layers. It should suffice for this study, however. Furthermore, the strip signals were fixed at 255 ADC to saturate the 8-bit FED registers. Again, although unphysical, this will have no effect on the software performance. Though allowing the raw data unpacker to be tested exhaustively, the clusteriser and hit producer cannot be validated in such a rigorous fashion. In general, the clusteriser algorithms reject digis below threshold and so the input and output data are not directly comparable.

The integrated reconstruction time for the original reconstruction scheme is shown in Figure 5.2. It shows the timing measurements up to the level of digis, clusters and hits for low luminosity minimum bias events as a function of strip occupancy. The time cost for unpacking the raw data (raw-to-digis) increased with strip occupancy, albeit with a relatively shallow gradient, due to a single loop over all occupied digis stored within the raw data. The intercept at 0 % occupancy is relatively large due to the time spend interpreting the raw data for each event. Since the FEDs are being operated in zero-suppressed mode, the raw data packets are inherently changeable in size and the physical position of registers of interest also vary. For this reason a scan of the full data packet is always required before unpacking can commence.

The time to clusterise the event data (raw-to-clusters - raw-to-digis) was strongly dependent on strip occupancy. This is to be expected since the algorithm performed

*Several people were involved in the development of the digis-to-clusters and clusters-to-hits algorithms, but my direct collaboration was with D. Giordano (INFN and Università di Bari) and C. Genta (INFN) respectively.

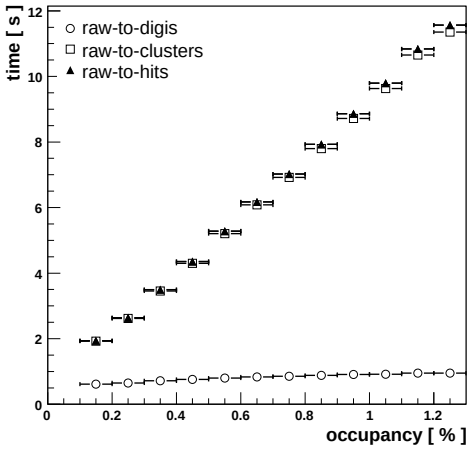

Figure 5.2: The variation in hit reconstruction time with strip occupancy for low luminosity minimum bias events. The mean cluster width for minimum bias events is 3.8 ± 0.3 strips.

two full loops over the zero-suppressed data: one to search for *seeds* and the second to find neighbouring strips above threshold. Multiple use of conditional statements also served to make the clusteriser module the dominant time cost. The overhead at low occupancies is large indicating a large amount of time spent initialising C++ container objects.

The hit production algorithm (raw-to-hits - raw-to-clusters) was clearly the least time consuming. This is for several reasons: its task was simple with little memory overhead or CPU requirement; it dealt with a smaller data-set since below threshold digis could be rejected by the clusteriser and it looped over clusters rather than digis reducing the number of loop iterations further still. For example, over a sample of 2000 minimum bias events, a mean of 5.2×10^4 digis were measured, 4.0×10^4 clusterised digis and only 1.0×10^4 clusters.

Since the expected mean tracker occupancy at low luminosity running is 0.56 % [86], the corresponding mean local reconstruction time within this scheme was 5.5 s. Although as few as 10 % of L1 accepted events are expected to require track reconstruction on the HLT [56], the resulting event-averaged time of 550 ms still exceeds the total HLT budget of 40 ms. To address these performance issues, the primary algorithmic inefficiencies and bottlenecks in the timing were identified by the author and are listed below:

- The templated container class $DetSetVector\langle T\rangle$ used by most CMSSW software modules within the reconstruction chain is effective, though highly inefficient when dealing with such large amounts of data. Every insertion into this

collection forces a sort and copy. Since the number of insertions was large for the SST (for a low luminosity minimum bias event, approximately 5.2×10^4 digis and 1.0×10^4 clusters), a new container class was needed.

- At the heart of the clusteriser's inefficiency was its use of an initial loop over strips to find cluster seeds. A more streamlined algorithm was required to clusterise within a single loop.

- The entire software scheme enforced a copy of a large fraction of the raw data at both the unpacking and clusterising stages. Since C++ copy times are linear with the data size, this became a highly significant effect for a 10^7 channel detector. Furthermore, each stage of the low level reconstruction chain introduced an extra loop over channels. Integrating the unpacking and clusterising algorithms within a single loop would serve to halve these effects.

- The majority of the unpacked raw data were not used within the tracking algorithms. A more sophisticated, *regional* approach to unpacking was necessary to reduce this overhead, though it could not be eliminated entirely, since the smallest unpacking granularity is a single FED.[†]

  It should be noted that, in the event data model, no interaction between software modules is allowed. Therefore, there is no direct communication between the local reconstruction modules listed in Figure 5.1 and the downstream track reconstruction modules. Additionally, no information from the muon and calorimetry systems on the event topology is communicated to the tracking algorithms. This means a regional reconstruction approach will require a completely new software design.

5.3 Regional reconstruction

To address the issues itemised in Section 5.2.1 and hence streamline the local reconstruction chain, a new software architecture was devised. The original concept was proposed by C. Jones [87] as a generic example of regional reconstruction within the CMSSW framework. It was then tailored for the SST and further optimised by the author. These optimisations (such as eliminating all surplus copies of raw

[†]In principle, unpacking a single channel is also possible, however since the raw data are zero-suppressed (and hence dynamic in size) the location of each channel within the packet will always require one full scan.

data) though barely necessary in the small-scale example of C. Jones, proved hugely important for the SST. In fact before the improvements were made, this new algorithm was *slower* than the original (when reconstructing the full tracker)! As will be shown, it is now much faster. Other features of this new software suite include a test facility for timing studies, trigger efficiency measurements and validation against the old chain. Finally, it should be noted that all the functionality available to the old reconstruction scheme (described in the previous section) has been maintained throughout these developments. This includes the capability of handling all FED readout modes and paths and various data quality checks.

Within the new reconstruction scheme, the SST modules are grouped into geometrical *regions*, defined by a grid on the η-ϕ plane with configurable dimensions. Each component of the grid extends from (0,0,0) to the outermost tracker radius. This is then further broken down into layers, each layer defining a single region. The new design allows raw data from *regions-of-interest* only to be considered. More specifically, any FED raw data packet with at least one channel connected to a region-of-interest is fully reconstructed. A schematic of the regional reconstruction concept is shown in Figure 5.3.

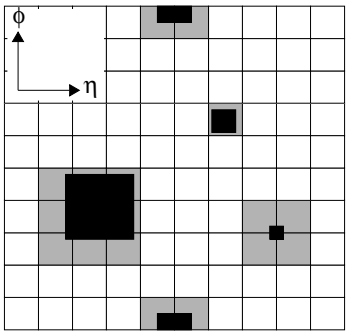

Figure 5.3: An illustration of the *regional* view of the SST in the η-ϕ plane. Each square corresponds to a section of the detector containing up to 12 layers, or *regions*. Each black shape represents a window of interest *seeded* by a physics object identified in an external sub-detector. The window size is configurable and will vary with object flavour. Grey squares are the corresponding *regions-of-interest*. The number of regions defined by a given window may vary with the position of the seed. Windows may also *wrap around* the ϕ dimension, as shown.

A schematic of the new software design is shown in Figure 5.4. At the heart of this scheme is a new, optimised cluster container which is added to the event by a *clusterizer facility*. Storing clusters by region (numbered sequentially from zero) means read and write operations can be performed by direct access to the relevant memory address rather than through a search, which is inherently slow. Also, by initialising the container at the start of every event, the sort and copy involved

with *DetSetVector*⟨*T*⟩ insertions is no longer necessary. A major concern with the concept of regional reconstruction has always been possibility of reconstructing the same region more than once. Having a single record of reconstructed regions makes overcoming this issue trivial. As well as providing fast access to clusters of a given region, this container also has access to the detector configurations (noise, strip gain and cabling) and the raw data themselves. This allows it to unpack and clusterise the raw data from any region on request[‡]. Furthermore, the process of unpacking and clusterising is now performed within a single loop.

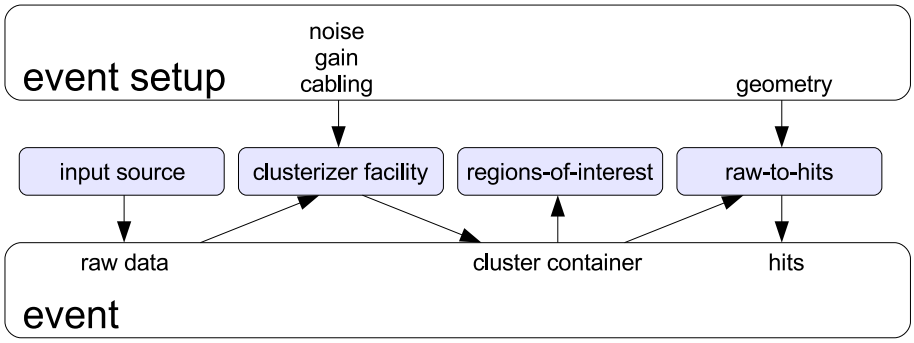

Figure 5.4: The new local reconstruction software design for the SST. Boxes represent CMSSW framework objects. Arrows represent the flow of data. The shaded boxes are plugin framework modules that perform a single event reconstruction task.

To inform the cluster container of the regions to be unpacked, a regions-of-interest module is required. The regions-of-interest themselves are usually identified by the positions of physics objects reconstructed in external sub-detectors. For example, superclusters identified in the calorimeter require association with a track before they can be identified as electrons. As such, all tracking regions neighbouring a supercluster in η-ϕ space would be added to the list of regions-of-interest in an electron trigger stream. To avoid the same region being declared more than once, the cluster container performs a *check-for-duplicates* on each insertion. Although duplicating regions-of-interest will not result in reconstructing the same region twice (this is avoided by a simple check in the cluster container) it would mean the same clusters are retrieved multiple times in a single loop. This would be confusing to the end user and propagate inefficiencies downstream.

Importantly, hit reconstruction is not performed at the regions-of-interest stage. Only when the clusters themselves are requested for a particular region is the re-

[‡]The inclusion of hit geometry has not been integrated within this loop. Instead, it must be performed independently. This avoids adding dependencies on the CMS geometry which greatly simplifies the software management.

construction process triggered. This final stage of the chain (looping over regions-of-interest and reconstructing clusters in the process) usually occurs during the hit reconstruction module so has been labelled *raw-to-hits* in the figure.

This design has been specifically chosen for use within the CMSSW HLT. Within each of the HLT's trigger paths, a unique reconstruction chain will be executed interspaced with trigger decisions. Before any of these streams are initiated, the clusterizer facility will be run. Then, each stream that contains a tracking module will define its own regions-of-interest just before the tracking is required. At the track building stage, any previously reconstructed clusters from earlier streams will be immediately available. Previously undeclared regions will be reconstructed from scratch.

The reconstruction algorithms have been developed in such a way as to allow various modes of operation, namely *global*, *regional*, *layers* and *on-demand*. The remainder of this section will describe each in detail and profile their performances. Comparisons will also be made with the old reconstruction chain, referred to from this point as *standard mode*.

5.3.1 Performance in *global mode*

When operating in *global mode*, all tracking regions are flagged for reconstruction by default. This therefore bypasses the regional functionality making it equivalent to the old scheme (standard mode), but optimised. Running in this mode can be especially useful for debugging the reconstruction chain and commissioning purposes.

Before comparing the performance of this mode with the standard scheme, the number of tracking regions must be configured. The total number of tracker regions has a significant effect on the overall timing, demonstrated by Figure 5.5. Hit reconstruction times for low luminosity minimum bias events were recorded in global mode. The quadratic increase in time with the number of regions is mainly due the overhead introduced by initialising the cluster container (which has a new entry for each region) at the start of every event. Although not shown, a slightly more rapid increase is recorded when operating in regional mode. This is a direct result of the check-for-duplicates performed as each new region-of-interest is defined.

In general, a balance clearly exists between increasing reconstruction times and improved detector granularity. With this in mind, a 20×20 configuration has been

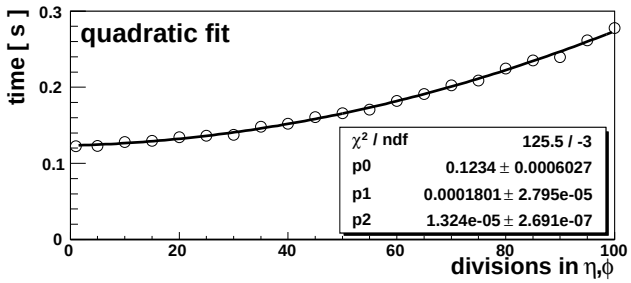

Figure 5.5: The variation in hit reconstruction time with regional granularity for minimum bias events at low luminiosity. The software is run in global mode. 20 divisions in η and ϕ corresponds to a regional granularity, $\Delta\eta \times \Delta\phi$ of 0.250×0.314.

chosen meaning a $\Delta\eta \times \Delta\phi$ regional granularity of 0.24×0.31. For reference, a calorimeter tower in the barrel has granularity 0.0870×0.0870. Since each of these divisions is further divided into between 9 and 12 layers, this corresponds to ~ 4000 regions in total.

Figure 5.6 shows the distribution in hit reconstruction times for low luminosity minimum bias events when unpacking globally. The fit is the sum of two Gaussians. This is motivated by the two dominant contributions to the CPU time: raw data unpacking and clusterising, each described by a Gaussian distribution. Over 2000 events, all times were measured at under 400 ms with a mean value of 140 ms. This is almost a factor 40 improvement on the standard reconstruction scheme. Since only 10 % of L1 accepted events are expected to require track reconstruction on the HLT, the average contribution per event is therefore ∼14 ms. This is ∼30 % of the full budget.

The variation of time with strip occupancy and number of clusters is shown in Figure 5.7. The figure on the left shows a linear increase in reconstruction times with occupancy, analogous to the trend of Figure 5.2. For example, it shows that high luminosity conditions, with a mean strip occupancy of 1.7 %, require 380 ms. Figure 5.7 (left) also serves to validate the fake cluster source described in Section 5.2.1. It shows the behaviour of the local reconstruction software over these clusters, with a fixed width and uniform distribution across the tracker, closely matches that for clusters from minimum bias events.

Figure 5.7 (right) uses the fake cluster source to show that the global reconstruction time is not just dependent on the number of digis (occupancy), but also how these digis are clusterised. More specifically, for events with an occupancy fixed at 0.56 %,

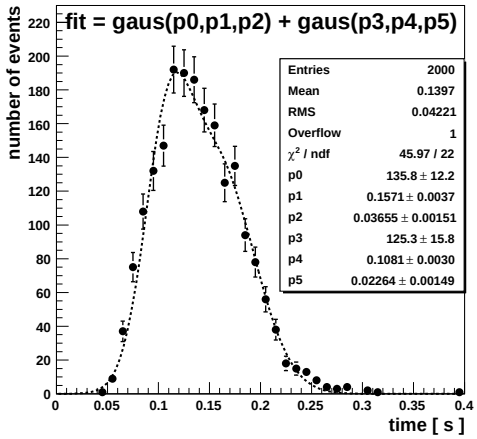

Figure 5.6: The spread in hit reconstruction time for minimum bias events at low luminosity. Here, the new regional software scheme is used, but in global mode. The fit is the sum of two Gaussians, reflecting the two dominant contributions to the overall CPU time: raw data unpacking and clusterising.

there is a linear relationship between the reconstruction time and the number of clusters. Also, since the mean cluster width and total number of clusters are inversely proportional, an increase in mean cluster width results in a decrease in reconstruction time.

Table 5.2 gives a breakdown of the dominant time costs measured over a characteristic HLT path. Due to the nature of the new SST reconstruction scheme, the hit reconstruction time is absorbed by the *L2.5 track reconstruction* module. Of the 60 ms measured here, ∼50 % goes toward reconstructing hits. In comparison, the full ECAL unpacking and hit production chain costs 22.6 ms with only $1/5^{th}$ the data rate (Table 5.1). The averaged times demonstrate the true timing cost of each module after considering the fraction of events passing L1 they are required for. Many of the minimum bias events are rejected early in the chain, therefore substantially reducing the average time taken in later steps. The table clearly shows the strip tracker is no longer the primary reconstruction bottleneck.

Table 5.2 also shows that for minimum bias events without pile-up, the HLT quota of 40 ms has been largely achieved. This is indicated by the sum of all averaged times, corresponding to the mean single τ plus missing E_T processing time over minimum bias events. 41.5 ms exceeds the 40 ms budget by < 4 %. This was a very important milestone for the development of the HLT in 2007 and was highlighted in the collaboration's HLT review [56]. This work was at the very heart of the achievement. Tracker occupancies were very low for this measurement (<1 %), however,

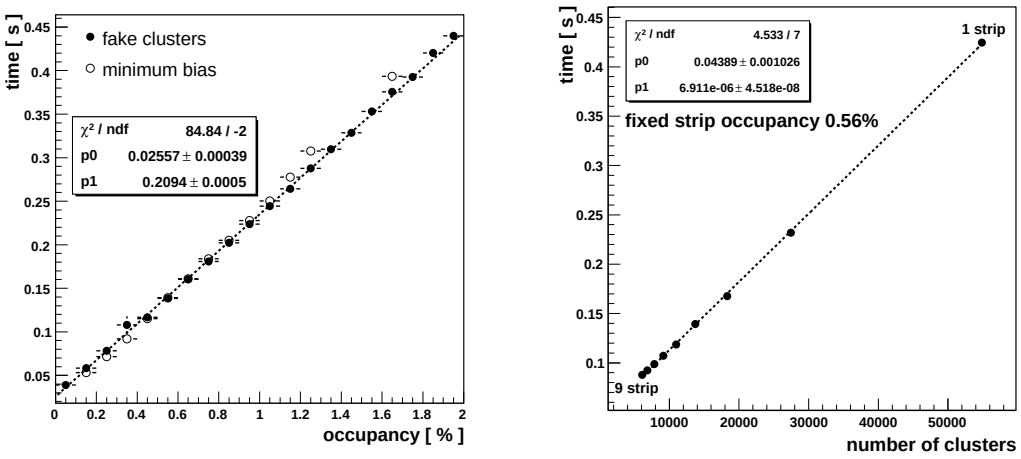

Figure 5.7: (left) The variation in hit reconstruction time with strip occupancy. The minimum bias events were generated at low luminosity giving a mean occupancy of 0.56 ± 0.18 % and a mean cluster width of 3.8 ± 0.3 strips. The fake clusters, with linear fit, are all 4 strips in width and evenly distributed throughout the detector. (right) The variation in hit reconstruction time with the number of fake clusters in the event. Here, occupancy is fixed at 0.56 % hence the cluster width varies with the number of clusters as labeled. For both figures the new regional software scheme is used, but in *global* mode.

so to stay within the 40 ms budget for pile-up conditions would require using the software's regional functionality. With this in mind, the regional functionality of the local reconstruction software is profiled in the following section.

5.3.2 Performance in *regional mode*

Figure 5.8 demonstrates the benefits seen with the regional approach. It shows the measured hit reconstruction time as a function of the fraction of regions flagged as interesting. The fit is for a quadratic increase in time with regions-of-interest over evenly distributed fake clusters. Naively, one would expect a linear relationship, however the check-for-duplicates described in Section 5.3.1 adds an, albeit small, second-order term. An overhead of 1.7 ± 0.2 ms exists in the case of no reconstruction being performed. This is due to the initialisation of the cluster container for each event. This overhead is much smaller than the corresponding value using the standard scheme, since the FED raw data packets do not require interpreting (scanning to understand the zero suppressed structure) unless a region is actually requested for reconstruction.

For real physics events, regions-of-interest must be defined around a set of η-ϕ point *seeds*. These points are chosen to represent the objects at the L2 stage of the trigger

HLT process	running time (ms)	averaged time (ms)	fraction of all events passing L1
ECAL unpacking	13.6	13.6	0.11
ECAL hit production	9	9	0.11
HCAL unpacking	1	1	0.11
HCAL hit production	3	3	0.11
Calorimeter tower maker	4.5	4.5	0.11
Jet reconstruction	3	3	0.11
Pixel unpacking	2	0.1	0.006
Pixel clusterising	6	0.3	0.006
Pixel hit production	2	0.1	0.006
Pixel tracks	10	0.5	0.006
L2.5 regional seeding	11	0.5	0.006
L2.5 track reconstruction	60	2.7	0.006
L3 regional seeding	26	0.2	0.0008
L3 track reconstruction	280	2.0	0.0008
Total time	421	41.5	-

Table 5.2: The time cost of the dominant processes within the single tau plus missing E_T HLT path. Results were measured over minimum bias events without pile-up. Here, *running time* represents the mean processing time for each step over events passing the previous step only and *averaged time* is the mean processing time for each step over all events passing L1 [56].

which require tracking. These candidates, their centre points and the physics objects they represent are listed in Table 5.3.

Centred on each seed, a window-of-interest is defined by a $\Delta\eta \times \Delta\phi$ *rectangle*. The coverage of this window will then define the regions to reconstruct in each case. For simplicity, the window dimensions are only dependent on the L2 candidate type. Also for this study we fix $\Delta\eta = \Delta\phi$, thus making the rectangle a *square*. This reduces the optimisation scan to 1-dimension. However, it is foreseen that for low p_T tracks with a large curvature, the relation $\Delta\eta < \Delta\phi$ will be more efficient. The tuning of these dimensions is the subject of the remaining section.

For the purpose of configuration, four trigger paths have been studied, each representative of a physics object type. These, along with the appropriate physics channel, number of events and corresponding L1, L2.5 (after global hit and track reconstruction) and full HLT trigger efficiencies are listed in Table 5.4.

For each L2 candidate, the goal is to minimise the number of declared regions-of-interest whilst maintaining an HLT trigger efficiency close to that achieved with global reconstruction. L1 $\times$ HLT trigger efficiencies relative to that for full tracker

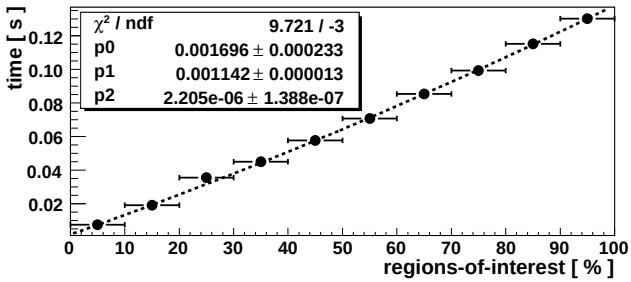

Figure 5.8: The variation in hit reconstruction time as a function of the fraction of the tracking regions that are reconstructed. The fit is quadratic. Fake, 4-strip clusters are evenly distributed throughout the tracker at low luminosity (0.56 % strip occupancy).

Object	L2 candidate	Seed
$e^{\pm}$	*supercluster*	$\frac{\sum_i x_i W_i}{\sum_i W_i}$
$\mu^{\pm}$	*standalone muon*	*outermost hit position*
b-jet	*iterative-cone jet*	*cone centre*
τ-jet	*iterative-cone jet*	*cone centre*

Table 5.3: Summary of the seeds used to define regions-of-interest for a given trigger path. More details on the reconstruction chain for each physics object can be found in Chapter 3. Here i represent component clusters of a supercluster and x are the cluster centres. Each cluster has corresponding weight, W, representing its energy. The cluster centres themselves are calculated using a similar approach. Full details can be found in [67].

hit reconstruction have been recorded whilst scanning through various $\Delta\eta \times \Delta\phi$ windows. These scans are shown for leptons, b-jets and τ-jets in Figures 5.9 and 5.10 and the chosen window sizes have been highlighted. In each case, a clear step in trigger efficiency can be seen as the window size passes a particular threshold. In principle, these steps will continue as the window size grows until the entire tracker is covered (reflecting the granular nature of the reconstruction). Only at this point would the measured trigger efficiency match the global efficiency for a large number of events. However, a trade-off exists with the fraction of the tracker to be reconstructed, so for this scheme to be successful the first plateau in trigger efficiency should be used as the optimum window size.

Figure 5.9 (top) shows that the relative efficiency for the isolated dielectron path is greater than the relaxed dimuon path for a given $\Delta\eta \times \Delta\phi$ configuration. This is thought to be due to the choice of seed for each object. Slight differences in the alignment of the reconstruction window and the particle trajectory in the SST will

Trigger	Channel	Events	L1 (%)	L1 × L2.5 (%)	L1 × HLT (%)
isolated dielectron	$Z^0 \rightarrow e^+e^-$	2000	77.3	-	62.6
relaxed dimuon	$Z^0 \rightarrow \mu^+\mu^-$	2000	99	-	98.1
b-jet plus 4 jets	$t\bar{t}$	5025	87	-	20
single τ-jet	$H^{\pm}(200) \rightarrow \tau^{\pm}\nu_{\tau}$	4975	68	22	16

Table 5.4: Summary of the data sets used for the trigger studies performed in this chapter. Here the term *relaxed* indicates the absence of isolation criteria. The corresponding trigger efficiencies for L1, L1 and L2.5 and for the entire trigger system are also shown.

be manifested as a variation in relative trigger efficiency. It is possible that other seed choices would prove more effective (e.g. the projection of the muon trajectory toward the outermost SST layer) but this investigation will be left to a later study. The shape of the *b-jet plus 4 jets* curve in Figure 5.9 (bottom) is the same as for the electron and muon case, only the final plateau occurs with a larger window. This is to be expected with a multiple-track jet.

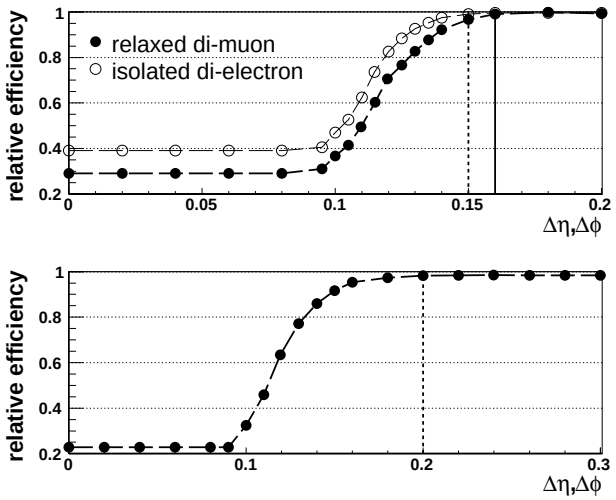

Figure 5.9: The variation in L1 × HLT isolated dielectron and relaxed dimuon (top) and b-jet plus 4 jets (bottom) trigger efficiencies with reconstruction window size, $\Delta\eta \times \Delta\phi$. Here, efficiency is relative to that measured for global reconstruction over the events listed in Table 5.4. The optimised window sizes are marked with the vertical lines.

Figure 5.10 shows the same scan performed using L2 τ-jets as the seeds. As described in Section 3.5.2, the tracking aspect of τ-jet identification occurs in two stages. At L2.5, tracking uses seeds from pixel hits found in a small rectangle $\Delta\eta \times \Delta\phi = 0.1 \times 0.1$ around the jet centre only. For high p_T tracks this corresponds to less than the size of a single tracking region, though low p_T tracks, with a greater curvature,

may extend beyond. Tracks falling within the matching cone of the jet, of size R_m = 0.5, and within a signal cone centred on the highest p_T track of size $R_s = 0.1$, are then considered to come from the τ decay.

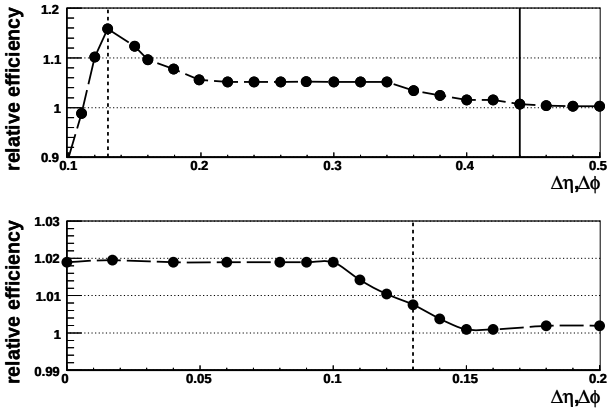

Figure 5.10: The variation in L1 × HLT single τ-jet trigger efficiency with reconstruction window size, $\Delta\eta \times \Delta\phi$. (top) Efficiency is relative to that measured for global reconstruction over the events listed in Table 5.4. The inclusion of signal tracks, required for L2.5 tracking, is indicated by the dotted limit. The threshold required to satisfy the L3 isolation criteria is indicated by the solid limit. (bottom) The same $\Delta\eta \times \Delta\phi$ scan, but considering the HLT efficiency up to L2.5 only. The dotted limit indicates the L2.5 threshold.

The peak in the L1 × HLT efficiency curve (top of the figure) with a 0.13 × 0.13 reconstruction window indicates the inclusion of all signal tracks at this point. This is therefore the required window size for L2.5. It is far smaller than the equivalent b-jet requirement (0.2 × 0.2) since since τ-jets are highly collimated.

L3 tracking uses seeds from pixel hits found in a rectangle $\Delta\eta \times \Delta\phi = 0.5 \times 0.5$ around the jet centre. Tracks already reconstructed at L2.5 are re-used. Any tracks outside the signal cone that fall within a larger cone of radius 0.4 (centred on the highest p_T track reconstructed at L2.5) are considered to violate the isolation criteria.

A relative efficiency above 1 in Figure 5.10 (top) (meaning the absolute trigger efficiency after regional reconstruction is larger than the equivalent using global track reconstruction) can be attributed to the absence of reconstructed tracks in the isolation cone. This means all events pass L3 by default. Interestingly, the same is true for Figure 5.10 (bottom) which describes the relative efficiency at L2.5. The regional efficiency is larger than the global measurement (by a factor of 1.008 at the chosen threshold). The cause of this slight increase is the incomplete reconstruction of tracks falling within less collimated jets. These will miss their outermost tracks and may therefore mimic a τ-jet. To correct for both of these effects, the window

size used at L3 must be larger than 0.13 × 0.13. It can be seen from Figure 5.10 that the relative trigger efficiency approaches its global value as more of the isolation cone is reconstructed. A threshold of 0.44 × 0.44 has been chosen.

The optimised window sizes for each trigger path studied are shown in Table 5.5. An extension to the regional mode of operation makes use of the layered structure of the tracker regions. This *layers mode* is discussed in the next section.

5.3.3 Performance in *layers mode*

To see further improvements in local reconstruction times, the layered structure of the tracking regions can be used. To examine the possibility of ignoring several outer tracking layers for HLT purposes, a scan was performed for each trigger path, with the $\Delta\eta \times \Delta\phi$ window sizes fixed at their optimal values. Beginning with a full tracker, at each stage of the scan the outermost layer was omitted until only a single layer at the centre remained. Figure 5.11 shows the measured effect on the final trigger efficiency. All trigger paths plateau in their efficiency using between 5 and 8 SST layers.

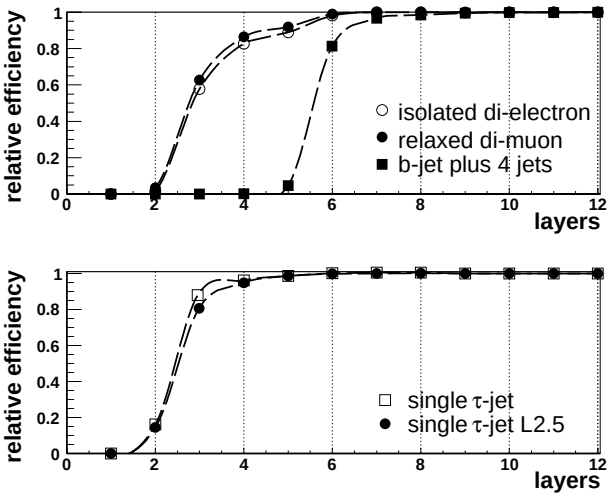

Figure 5.11: The variation in HLT dielectron trigger efficiency with the number of tracker layers for electrons, muons and b-jets (top) and τ-jets (bottom). Here, efficiency is relative to that measured for global reconstruction over the events listed in Table 5.4.

Figure 3.2 shows the variation in p_T and impact parameter resolution for muons as a function of the number of reconstructed hits (SST and pixel combined). It clearly shows a plateau in both resolutions beyond 5 or 6 hits. Also Figure 3.3 shows the

tracking fake rate drops below 1 % for tracks with at least 8 hits. These results indicate that no dramatic change in the track parameters is expected as a result of this procedure.

A summary of the chosen $\Delta\eta \times \Delta\phi$ window sizes and layer numbers for each trigger path studied, along with the corresponding mean fraction of tracking channels included, the channel occupancy of these regions, the unpacking time per event and measured relative HLT trigger efficiency is shown in Table 5.5 for *regional* and *layers* modes. The case of *global* and *standard* (before any optimisations) *modes* are also shown for comparison. The unpacking times increase with both the fraction of tracker reconstructed and the mean occupancy. Importantly, reducing the former by introducing regions-of-interest has dramatically improved hit reconstruction times whilst maintaining relative trigger efficiencies within 2 % of their global values in all cases. Although not ideal, this is considered acceptable if it allows tracking information to be used on the HLT.

As an example, the single τ-jet trigger path costs only 4.73 ± 0.06 ms as a result of hit reconstruction in layers mode. This is a factor 19 improvement on the hit reconstruction time in global mode, and 177 over the old standard. The cost is a 0.1 % drop in trigger efficiency.

A further reconstruction mode: *on-demand*, is also shown in this table[§]. This will prove to be the optimal case, with reconstruction times similar to layers mode but zero loss in the trigger efficiency with respect to global or standard modes. It is discussed in detail in the next section.

5.4 *on-demand* reconstruction

To further optimise the local hit reconstruction process, the regional reconstruction algorithms described in Section 5.3 have been introduced into the tracking code itself. In this way, detector regions-of-interest can be identified within the tracking algorithms using the projected track trajectory and its uncertainty, both defined by a Kalman filter. Since regions-of-interest are established through the needs of the tracking algorithms directly, this scheme is known as reconstruction *on-demand*.

[§]Thanks to Nicholas Cripps (Imperial College London) for assisting with these measurements

Trigger	Mode	Layers	$\Delta\eta \times \Delta\phi$	Fraction of tracker (%)	Occupancy (%)	Hit reconstruction (ms)	Full track reconstruction (ms)	Relative efficiency
isolated dielectron	standard	all	-	100	0.126±0.001	729±3	922±4	1
	global	all	-	100	0.126±0.001	65.9±0.4	255±1	1
	regional	all	0.15×0.15	2.81±0.02	0.241±0.003	5.95±0.04	205±1	0.99±0.01
	layers	6	0.15×0.15	1.21±0.01	0.275±0.003	3.41±0.03	183.4±0.8	0.98±0.01
	on-demand	-	-	3.02±0.02	0.237±0.003	7.99±0.04	224±1	1
relaxed dimuon	standard	all	-	100	0.122±0.002	705±3	902±4	1
	global	all	-	100	0.122±0.002	64.6±0.5	255±1	1
	regional	all	0.16×0.16	3.22±0.01	0.126±0.002	5.26±0.04	213±1	0.991±0.004
	layers	6	0.16×0.16	1.441±0.008	0.169±0.003	3.27±0.03	195±1	0.989±0.004
	on-demand	-	-	1.175±0.009	0.208±0.003	5.77±0.03	215±1	1
b-jet plus 4 jets	standard	all	-	100	0.241±0.001	883±3	1155±5	1
	global	all	-	100	0.241±0.001	96.7±0.4	363±2	1
	regional	all	0.2×0.2	28.0±0.1	0.389±0.002	46.7±0.3	340±2	0.982±0.006
	layers	8	0.2×0.2	17.23±0.09	0.452±0.002	34.1±0.2	328±2	0.984±0.006
	on-demand	-	-	0.995±0.009	1.025±0.007	7.55±0.05	310±2	1
single τ-jet	standard	all	-	100	0.218±0.002	838±5	1045±5	1
	global	all	-	100	0.218±0.002	90.0±0.6	293±2	1
	regional	all	0.44×0.44	4.99±0.06	0.302±0.005	8.3±0.1	240±2	1.007±0.005
	layers	6	0.44×0.44	2.19±0.03	0.350±0.006	4.73±0.06	237±2	0.999±0.005
	on-demand	-	-	0.51±0.01	0.68±0.01	5.13±0.05	239±2	1
single τ-jet L2.5	regional	all	0.13×0.13	0.95±0.01	0.57±0.01	3.74±0.05	218±1	1.008±0.006
	layers	5	0.13×0.13	0.307±0.005	0.64±0.02	2.06±0.03	214±1	0.987±0.006
	on-demand	-	-	0.274±0.005	0.75±0.01	4.52±0.04	218±1	1

Table 5.5: Performance summary of the local and full track reconstruction software in all modes over a selection of trigger paths. For this table, measurements were made within CMSSW 1.7.5 using a dual core Pentium D^{TM} with 3.4 GHz CPU. Times quoted for the single τ-jet trigger include tracks reconstructed at L2.5. Relative efficiency represents the full trigger efficiency (L1 and full HLT, or L1, L2 and L2.5 for single τ-jet L2.5) relative to that recorded in *standard* mode. Here, occupancy corresponds to the strip occupancy over reconstructed regions only.

As information external to the SST is no longer necessary the software management is simplified considerably. Also, since all hits from detector modules intersecting the track's path are reconstructed by definition, track reconstruction efficiencies are guaranteed to remain unaffected. In this sense it is the best possible scenario.

As for Figure 5.4, the scheme uses a single clusterizer facility at the start of each path, with access to the raw data and cabling information. Regions-of-interest are defined during the tracking process, however, instead of via a dedicated CMSSW module. This idea is illustrated in Figure 5.12. At each layer, a central position is defined by the track projection through the magnetic field along with a 1σ uncertainty represented by the red bar. Regions that lie within this coverage are declared as *interesting*. Corresponding hits are then reconstructed, a combinatorial fit is applied and the best candidate is added to the track. The track trajectory is finally updated and projected to the next layer. As the track propagates outwards more hit information is included in the trajectory and hence the uncertainties lessen.

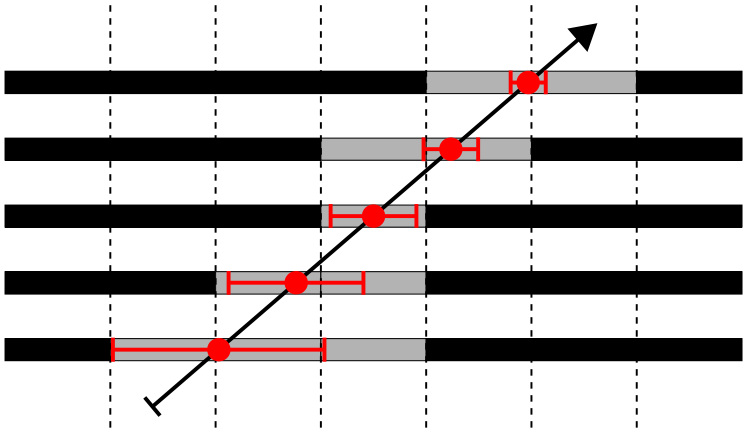

Figure 5.12: An illustration of the on-demand hit reconstruction algorithms. Regions-of-interest are defined by the track trajectory. This example is for several layers within the barrel geometry, shown in the x-z plane. Tracking regions are defined by the boundaries (dotted-lines) and each layer. The projected track trajectory intersection with each layer is represented by a red dot, along with the uncertainty bar. Regions-of-interest are light grey.

Table 5.5 compares both the hit and full track reconstruction times for the trigger paths discussed in Section 5.3.2. For clarity, the path times are compared over all reconstruction modes in Figure 5.13. It is clear that the software developments made for this chapter are a great improvement on the old scheme. *on-demand* performs as well as, or slightly worse than, the *regional* and *layers* modes for $e^{\pm}$, $\mu^{\pm}$ and $\tau^{\pm}$ events, but better for the $t\bar{t}$ events, where the number of tracks is greater. For example, the single τ-jet trigger path costs 5.13 ± 0.05 ms compared with 4.73 ± 0.06 in layers mode.

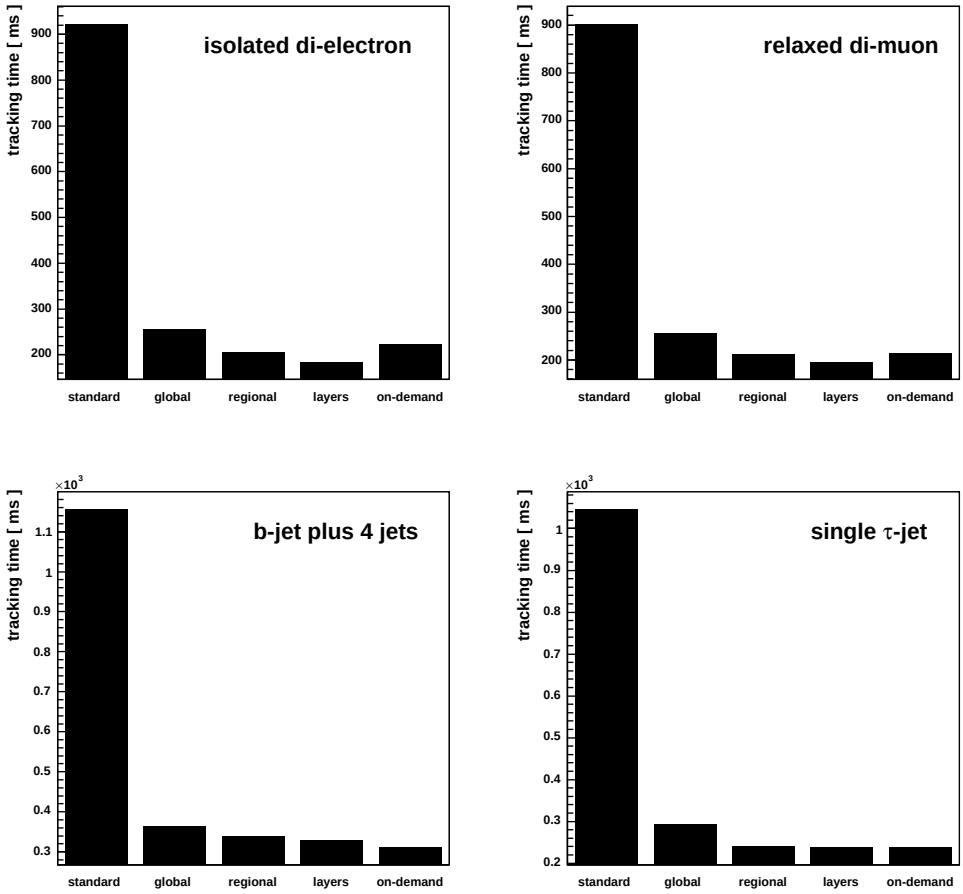

Figure 5.13: The variation in tracking times for various HLT paths.

For reasons of software management, CPU performance and optimal trigger efficiency, the *on-demand* mode has been adopted by the tracker community as its default online local reconstruction solution. It is expected that it will also be used as the offline solution in the near future.

5.5 Extension to other sub-detectors

This new reconstruction scheme developed for the SST has proven highly successful. It is efficient through both its minimal use of loops and copies and the processing of tracking regions-of-interest only. The improvements seen in the SST which have been discussed in detail in this chapter, are a direct reflection of the number of occupied strips, or digis, expected per event. As such, other sub-systems (such as

the ECAL or pixel tracker) with fewer digis would be expected to see reduced levels of improvement.

Nevertheless, the ECAL community undertook an investigation into an identical approach and recently adopted it as their solution. This uses a common set of software tools originally developed for the SST. The Pixel and HCAL communities are considering doing likewise.

5.6 Summary

Raw data unpacking and local reconstruction of the CMS silicon strip tracker is subject to the stringent requirements of the HLT computer farm, with only 40 ms processing time per event. To minimise their contribution to the former, the local reconstruction algorithms have been redesigned, reducing the time for minimum bias events at low luminosity from 5.5 s to 140 ms. Since approximately 10 % of level 1 accepted events are expected to require track reconstruction, the average contribution per event is ∼14 ms. This is ∼30 % of the full budget.

Additionally, in order to reduce the time cost further, the new scheme is also capable of reconstructing only physics regions-of-interest. For example, when running the single τ trigger path over $H^{\pm} \rightarrow \tau^{\pm}\nu_{\tau}$ events, 99.9 % of the global trigger efficiency is maintained whilst unpacking only 2.19 ± 0.03 % of the tracker. In this case, the mean hit reconstruction time per event is reduced from 90.0 ± 0.6 ms when considering the entire detector to only 4.73 ± 0.06 ms. This time is also expected to reduce by a factor of 10 when considering the fraction of events requiring track reconstruction, meaning ∼0.4 ms hit reconstruction times for this path.

Finally, the regional hit reconstruction code was integrated into the tracking algorithms themselves. With this approach, the detector regions-of-interest can be identified internally *on-demand* via the projected track trajectory at each layer. Reconstruction times in this mode are comparable to the regional case (5.13 ± 0.05 ms compared with 4.73 ± 0.06 for the single τ-jet path), but the tracking efficiency is maintained at the level of standard or global modes. In fact, it is believed that the use of the regional algorithms within the tracking code can be optimised further still. If this is true then the reconstruction times may drop below that of the other regional modes. For these reasons, and due to simpler software management, the *on-demand* approach has been permanently adopted by the tracker community as its online solution.

Chapter 6

Toward a 100 pb^{-1} b$\bar{\text{b}}$Z^0 cross section measurement at the LHC

$b\bar{b}Z^0$ production at the LHC is of great interest as a benchmark process and background to the neutral MSSM Higgs production $b\bar{b}h/H/A$. The cross section of the latter is enhanced at large values of tanβ making it a potential avenue for discovery. This chapter focuses on preparing for a measurement of the $b\bar{b}Z^0$ cross section with the first 100 pb^{-1} of CMS data at $\sqrt{s}$ = 14 TeV. This integrated luminosity is expected by the end of 2009. If the next-to-leading-order theoretical predictions are validated with real data then the predictions for the associated Higgs boson production should also apply.

6.1 $b\bar{b}H^0$ and $b\bar{b}Z^0$ production

Many signals searched for in hadron colliders involve electroweak gauge bosons in association with one or more jets containing heavy quarks (c or b). A good example of this is the discovery of the top [88],[89] through W + 4 jets with one or more containing a b tag. Indeed, the discovery of new physics beyond the Standard Model often relies on a detailed understanding of these processes as backgrounds.

In the Standard Model, the coupling of the Higgs boson to b-quarks is proportional to the bottom Yukawa coupling $\frac{m_b}{v}$, where m_b is the b-quark mass and v the vacuum expectation value of the Higgs field. It is therefore very weak due to the small b mass ($4.79^{+0.19}_{-0.08}$ GeV/c^2). In a two Higgs-doublet model however, some or all of the Higgs

particle couplings can be greatly enhanced [90]. In the MSSM, for example, as b-quarks only couple to the neutral Higgs ϕ_u^0, the Yukawa term to first order is $h_b\phi_u^0 b\bar{b}$, where h_b is the coupling constant. Upon SUSY breaking however, a correction can be added to the Lagrangian density of the form $\Delta h_b \phi_d^0 b\bar{b}$. Upon electroweak symmetry breaking, the Higgs doublets acquire their vacuum expectation values and this term scales as $\Delta h_b v_2$. Though Δh_b is still small, in the region of large $\tan\beta$ this correction can become significant.

If the coupling is sufficiently enhanced, the production of a Higgs boson in association with bottom quarks can become an important process at the LHC. This process, shown by Figure 6.1 and its subsequent decay into a pair of τ-leptons provides a clear signal for neutral Higgs bosons (h, H or A) and is therefore a potential discovery channel. As such, a great deal of attention has been paid towards this channel both at the Tevatron $p-\bar{p}$ collider in Fermilab, Chicago [91] and the LHC [92].

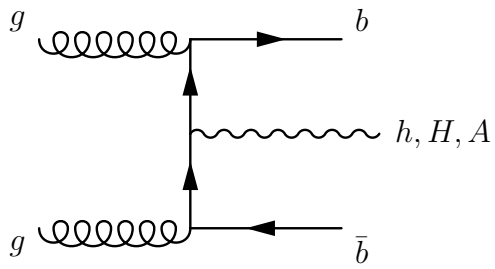

Figure 6.1: Example LO Feynman diagram contributing to h,H,A production in association with two b-quarks.

The process of Z boson production in association with two b-quarks is topologically very similar to the Higgs production described above. Feynman diagrams for the dominant Leading-Order (LO) production mechanisms are shown in Figure 6.2. Of the two processes, it transpires that the $q\bar{q} \to b\bar{b}Z$ cross section is much smaller at the LHC. This is illustrated by Table 6.1 where the NLO contributions to Z boson production with a b-jet at the LHC and Tevatron are shown. At the LHC, the $gg \to b\bar{b}Z$ cross section exceeds that of $q\bar{q} \to b\bar{b}Z$ by a factor of 21.

The similarity of the dominant $b\bar{b}H^0$ and $b\bar{b}Z^0$ production mechanisms can be exploited when trying to validate the theoretical predictions for the former process, such as its kinematics and cross section, using real data. This chapter, for example, will focus on a $b\bar{b}Z^0$ production cross section measurement to verify the Next-to-Leading-Order (NLO) calculations. This, in turn, will give confidence in

the predictions made for Higgs production. In this sense it is a testing-ground for a Higgs discovery. Furthermore, the measurement of this process is interesting in its own right as the primary $b\bar{b}H^0$ production background.

Despite these similarities, the case of Z^0 production is more complicated for two reasons [93],[94]. Firstly, although the Higgs production in association with charm and light quarks is negligible, this is not true for the Z boson. Since both charm and light quarks can fake a b-quark, this is a major additional source of background for $b\bar{b}Z^0$ and will be discussed in more detail in Section 6.2. Secondly, there are a variety of contributions to the inclusive production of a Z^0 with associated b-quarks which have no high cross section analogue in the Higgs case. For example, $q\bar{q} \rightarrow b\bar{b}Z$, where q are light quarks. This complication can be somewhat overlooked, however, due to the low $\sigma(q\bar{q} \rightarrow b\bar{b}Z)/\sigma(gg \rightarrow b\bar{b}Z)$ ratio.

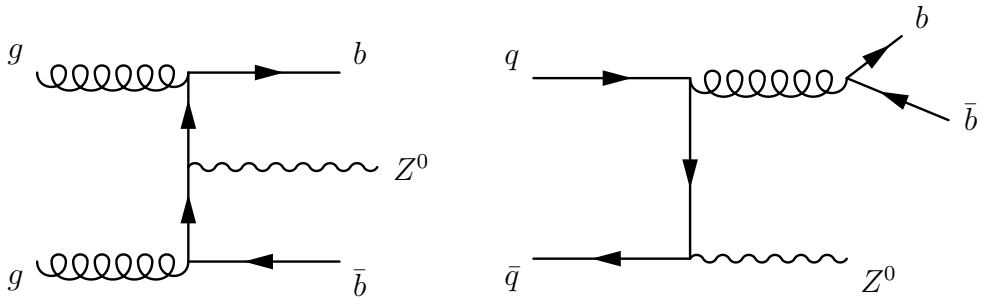

Figure 6.2: Example LO Feynman diagrams contributing to Z^0 production in association with two b-quarks.

6.1.1 $gg \rightarrow b\bar{b}Z^0$ matrix element calculation

Two different methodologies exist for calculating the gg $\rightarrow b\bar{b}Z^0$ matrix element. The *four-flavour* scheme is based on the LO process $gg \rightarrow b\bar{b}Z^0$. It is so labelled after the proton's Parton Distribution Function (PDF) which is assumed to have only four types of quark: u, d, c and s. Here, the gluon splitting introduces logarithmic terms of the form $ln(m_{Z^0}/m_b)$, which diverge the perturbation series. They can be regulated by keeping the b-quark mass non-zero, however this does complicate the calculation. At present, the four-flavour scheme has been used (with $m_b = 0$) to NLO [96].

Channel	Collider	NLO cross section (pb)				
		QZ	$(Q\bar{Q})Z$	$Qj Z$	$Q\bar{Q}Z$	inclusive
$gg \to b\bar{b}Z$	LHC	649	11.3	304	78.1	$1040\ ^{+70}_{-60}\ ^{+70}_{-100}\ ^{+30}_{-50}$
	Tevatron	10.4	0.169	2.19	0.631	$13.4\ \pm 0.9 \pm 0.8 \pm 0.8$
$q\bar{q} \to bbZ$	LHC	24.3	13.5	-	11.4	49.2
	Tevatron	3.32	1.92	-	1.59	6.83

Table 6.1: NLO cross section calculations for Z boson production in association with a b-jet at the LHC ($\sqrt{s}$ = 14 TeV pp) and Tevatron ($\sqrt{s}$ = 1.96 TeV $p\bar{p}$) [95]. The $gg \to b\bar{b}Z$ process is calculated in the five flavour scheme, described in Section 6.1.1. Two final state partons are merged into a single jet if $\Delta R < 0.7$. No branching ratios or tagging efficiencies have been factored in. Here Q is exactly one heavy jet; $(Q\bar{Q})$ is one jet containing two heavy quarks; Qj is two jets, one containing a heavy quark and $Q\bar{Q}$ is exactly two jets both containing one heavy quark. A jet lies in the range $p_T > 15$ GeV and $\mid \eta \mid < 2.5$. The uncertainties are due to renormalisation scale, factorisation scale and parton distribution functions respectively.

The *five-flavour* scheme is based on the LO process $gb \to Z^0 b$. By absorbing one of the gluon splitting terms into the b-quark PDF of the proton (this demands an extra quark flavour, hence "five-flavour"), the logarithmic divergences associated with the four-flavour scheme no longer exist. For this reason, including the b-quark mass is no longer necessary. In this sense the calculation is far simpler. NLO corrections are also easier to obtain as there is one less particle in the final state. As for the four-flavour scheme, the matrix element for this process has been calculated to NLO with $m_b = 0$. The calculation was considered sufficiently straightforward, however, that the Next-to-Next-to-Leading-Order (NNLO) calculation may be performed within the foreseeable future [95]. It should be noted that, since each scheme is simply a rearrangement of the others terms, the two calculations should approach one another at higher orders.

The four-flavour scheme is based upon the assumption that both b-quarks have a p_T significantly larger than their mass. Theorists usually place this cut at 15 GeV/c. However, the five-flavour scheme only requires a single b-quark (which is not included in the proton sea) to be above this threshold. This thesis aims to probe the five-flavour calculation and so the event selection strategy must require at least one high p_T b-jet.

6.1.2 Z^0 decay modes

The choice of Z^0 decay mode is based on the ability of CMS to reconstruct the decay products. A high purity event selection is vital to avoid systematic uncertainties

which will be large, especially during early data taking. CMS has been designed with excellent muon and electron reconstruction capability, making the $Z^0 \to e^+e^-$ and $Z^0 \to \mu^+\mu^-$ modes an obvious choice. Since events with other leptons in the final state, including $\tau^\pm$ and neutrinos of all flavours and indeed hadronic final states, are expected to have a significantly larger background, these decay channels will not be considered in this thesis.

6.1.3 Previous studies

Both the CDF and D0 experiments have performed $b\bar{b}Z^0$ production cross section measurements at the Tevatron [97], [98]. Both studies rely on the $Z^0 \to e^+e^-$ and $Z^0 \to \mu^+\mu^-$ modes and probe the five-flavour scheme by requiring at least one jet with $E_T > 20$ GeV (after energy corrections to the hadron level).

To minimise the effect of both theoretical and experimental uncertainties, they calculated the $\sigma(Z^0 + b\text{-jet})/\sigma(Z^0 + \text{jet})$ ratio yielding values of 0.021 ± 0.004 (stat) $^{+0.002}_{-0.003}$ (sys) and 0.0236 ± 0.0074 (stat) ± 0.0053 (sys) respectively. In both cases, the uncertainty is dominated by limited statistical precision. Both results are in good agreement with the NLO Standard Model prediction of 0.018 ± 0.004 [95].

Although successful studies have already been performed at the Tevatron, further scrutiny at the LHC is vital to probe these interactions at the highest possible energies. One study has already been performed for CMS [99]. This too relies on the $Z^0 \to e^+e^-$ and $Z^0 \to \mu^+\mu^-$ modes. Its signal selection strategy differs from that of this analysis, however, making it complementary. The key difference is the requirement for at least 2 reconstructed b-jets ($E_T > 20$ GeV) in a selected event (probing the four-flavour calculation), whereas this study will look for only one.

6.2 Background processes

The purpose of this chapter is to prepare for a $b\bar{b}Z^0$ cross section measurement with early LHC data. The cross section measurement of any process (discussed in detail in Section 6.4) relies on the measurement of the process rate at a given LHC luminosity. This, in turn, depends on the identification of signal events and rejection of background.

If the emphasis is placed on signal selection, statistical errors in the final result are expected to be minimised. Conversely, if the emphasis is instead placed on background rejection the effect will be to minimise the systematic uncertainties. After 100 pb^{-1} integrated luminosity, ∼5000 signal events are expected. If we assume a signal selection efficiency of order 5 %, this corresponds to a statistical uncertainty of 6 %. Systematic uncertainties are expected to outweigh this, especially during early LHC running. A detailed study of their impact on a $b\bar{b}Z^0$ cross section measurement is given in Section 6.7.3, but it is clear from the outset that the emphasis of this analysis should be on background rejection.

For this reason, a precise knowledge of the background processes is crucial. Backgrounds fall into two distinct categories: irreducible and reducible. Irreducible physics background processes involve two b-jets and two $e(\mu)$ leptons in the final state which have similar kinematics to that of the signal. Reducible backgrounds however, can be events from any process with misidentified leptons or light or c-jets that are misidentified as b-jets. The most prominent background processes of each type are included in this analysis. They are described in detail below.

The dominant irreducible background for this study is $t\bar{t}$ production. Example LO Feynman diagrams for this process are shown in Figures 6.3 and 6.4. The former is expected to be the primary contributor at the LHC at $\sim$ 90 % the total cross section with the latter at only $\sim$ 10 %. It should be noted that, strictly speaking, the irreducible $t\bar{t}$ background is constrained by the subsequent decay $t\bar{t} \rightarrow b\bar{b}l^+l^-\nu_l\bar{\nu}_l$, where l = e (μ). However, the case of single electron or muon decay should also be considered to account for fakes introduced during reconstruction or secondary leptons emitted from jets. For simplicity, *inclusive* $t\bar{t}$ (with all decay modes) is used.

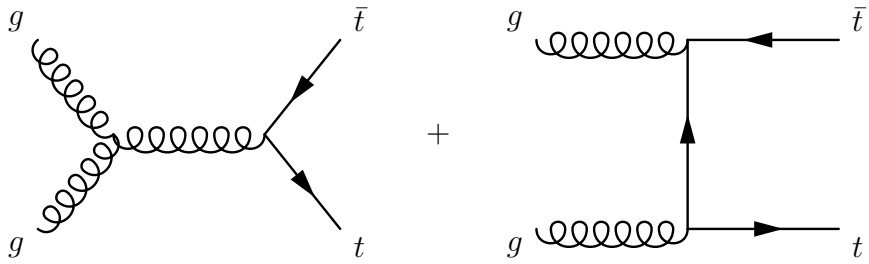

Figure 6.3: Example LO Feynman diagrams contributing to $t\bar{t}$ production from gluon fusion.

The reducible backgrounds studied are Z^0+1-jet, Z^0+2-jets and Z^0+3-jets production, with the subsequent decay $Z^0 \rightarrow l^+l^-$, where l = e (μ). Here, *jet* corresponds

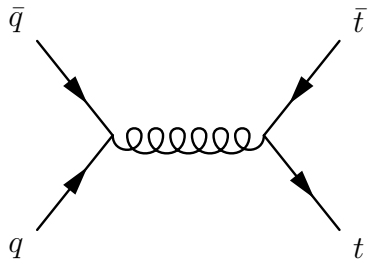

Figure 6.4: Example LO Feynman diagrams contributing to $t\bar{t}$ production from quark fusion.

to a light jet of flavour u, d, s or g (c-jets are considered separately). Example LO Feynman diagrams for these processes are shown in Figures 6.5, 6.6 and 6.7.

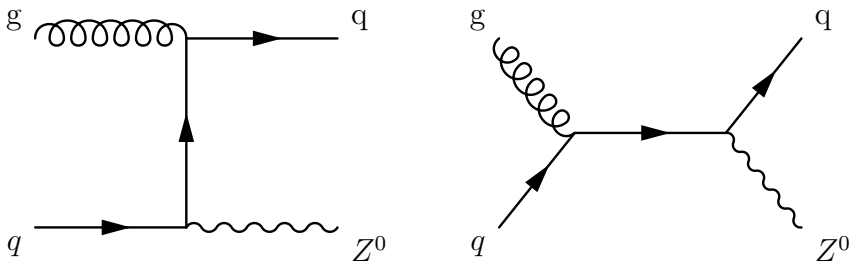

Figure 6.5: Example LO Feynman diagrams contributing to Z^0+1-jet production.

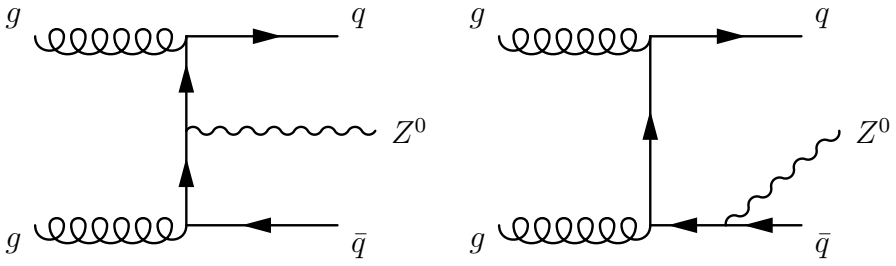

Figure 6.6: Example LO Feynman diagrams contributing to Z^0+2-jets production.

Another reducible background which should be considered in this analysis is $c\bar{c}Z^0$ production. Its production mechanism is analogous to that for $b\bar{b}Z^0$ shown in Figure 6.2, but with c-quarks substituted for b's. With a ratio of NLO cross sections $\frac{\sigma(c\bar{c}Z)}{\sigma(b\bar{b}Z)}$ of 0.26 (after event preselection) [100],[101], suppression of this relies on a small misidentification efficiency of c-jets as b-jets. Using CMS b-tagging algorithms, for a b-jet identification efficiency of 50 % a 10 % misidentified c-jet efficiency is expected

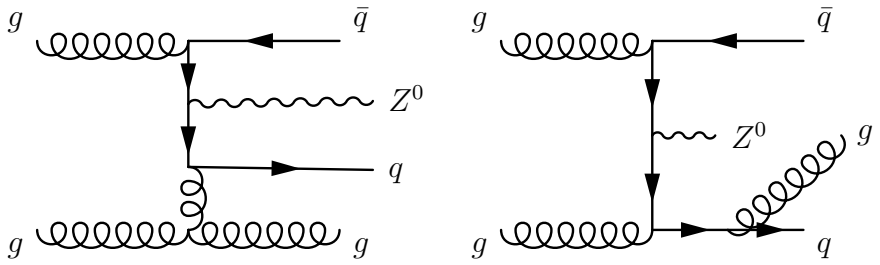

Figure 6.7: Example LO Feynman diagrams contributing to Z^0+3-jets production.

[30]. If both b-jets are identified this reduces the $c\bar{c}Z^0$ contribution to $\frac{10\%\times10\%}{50\%\times50\%} \times \frac{\sigma(c\bar{c}Z)}{\sigma(b\bar{b}Z)}$ = 1 %, making it negligible in comparison with other systematics. For a single b-tagging strategy, however, a contribution of order $\frac{10\%\times90\%+90\%\times10\%}{50\%\times50\%+50\%\times50\%} \times \frac{\sigma(c\bar{c}Z)}{\sigma(b\bar{b}Z)} = 9$ % is expected. This clearly is significant. However, since just ∼5000 $c\bar{c}Z^0$ events were made available for this analysis and only with the decay mode $Z^0 \rightarrow e^+e^-$, it is not discussed in great detail here. Instead, Section 6.9 will make an estimate of the actual contribution from Monte Carlo and discuss how an estimate of this background can be made from data.

Other sources of reducible background come from channels with neither leptons nor b-jets in their final state. To pass the signal selection, a single b-jet and two leptons must be misidentified. As the probability of such an occurrence is very low, this will only happen through channels with a very large cross section. The best example of this is QCD 3-jet events which have a LO cross section of $\sim 10^7$ pb [101]. If we consider the electron branch, the CMS electron identification algorithms incur a misidentified jet efficiency of 6×10^{-4}[30]. For two fake electrons, the combined efficiency of 3.6×10^{-7} renders this background negligible. Thorough studies on the rejection of QCD backgrounds to $Z^0 \rightarrow e^+e^-$ and $Z^0 \rightarrow \mu^+\mu^-$ events can be found in [102] and [103].

Finally, the reader should be aware that processes containing virtual Z-bosons always interfere with the analogous γ^* process. For example, $\gamma^* b\bar{b}$ production interferes with the signal channel. This will not always be explicitly stated however, as these contributions are easily suppressed by applying an acceptance window on the reconstructed boson mass.

6.3 Sample production and preselection

A summary of signal and background generation is listed in Table 6.2, including the generator used, preselection criteria, preselected cross section and total number of preselected events. It should be noted that the samples produced for this study are for LHC collisions with $\sqrt{s} = 14$ TeV. Although, in principle, a 50 pb^{-1} measurement at the initial physics run conditions of $\sqrt{s} = 10$ TeV is possible [104], this scenario is not considered here.

The event generators (CompHEP [59] and Alpgen [105]) simulate hard processes based on LO matrix elements. They are interfaced with PYTHIA [58] which simulates the parton showers and hadronisation. Jet matching (jet formation from four-vectors) can be either be performed at the ME or parton shower levels, depending on the separation of the jets within each event.

The preselection is based on four-vector information provided at the level of event generation. The purpose of preselection is to identify events of interest early and so avoid spending unnecessary time in analysis downstream. For signal this means avoiding $b\bar{b}\gamma^*$ interference by imposing a dilepton mass cut and only choosing events containing leptons within the geometrical limits of the CMS tracker. The latter implies they are reconstructible and hence have a chance of passing any future selection criteria. The flavour of the Z^0 decay products is also fixed as e or μ.

For the Z^0+*1,2,3*-jets backgrounds, as well as a dilepton mass cut and dilepton flavour preselection, a ${Z^0}_{p_T}$ window has been chosen. This loosely corresponds to the ${Z^0}_{p_T}$ of the signal and hence focuses on background events where the lepton kinematics are likely to be similar. To avoid overlaps between the signal and Z^0+*1,2,3*-jets background processes, events of the latter which contain b-quarks in the initial or final state are also removed. This process has been termed *b-veto* in the table.

The final, preselected samples have been analysed within CMSSW version 1.6.7.

Channel	ME generator	I	σ (pb)	II	σ (pb)	theoretical uncertainty (%)	Events generated	$\mathcal{L}_{int}$ (fb^{-1})
$b\bar{b}Z^0$	CompHEP	$p_T^b > 10$ $\mid \eta^b \mid < 10$ $40 < M_{Z^0}$ $Z^0 \rightarrow e^{\pm}(\mu^{\pm})$ $\mid \eta_{e^{\pm},\mu^{\pm}} \mid < 2.5$	50	-	50	11	645,252	12.91
$t\bar{t}$	Alpgen	-	840	-	840	6	1,473,797	1.75
Z^0+1-jet	Alpgen	$p_T^j > 20$ $\mid \eta^j \mid < 5$ $0 < {Z^0}_{p_T} <100$ $40 < M_{Z^0} < 200$ $Z^0 \rightarrow e^{\pm}(\mu^{\pm})$	319.6	*b-veto*	316	3	270,530	0.86
Z^0+2-jets	Alpgen	$p_T^j > 20$ $\mid \eta^j \mid < 5$ $0 < {Z^0}_{p_T} <100$ $40 < M_{Z^0} < 200$ $Z^0 \rightarrow e^{\pm}(\mu^{\pm})$	91.8	*b-veto*	89	5	91,150	1.02
Z^0+3-jets	Alpgen	$p_T^j > 20$ $\mid \eta^j \mid < 5$ $0 < {Z^0}_{p_T} <100$ $40 < M_{Z^0} < 200$ $Z^0 \rightarrow e^{\pm}(\mu^{\pm})$	23.1	*b-veto*	22	-	23,035	1.00

Table 6.2: Summary of signal and background production and preselection with the corresponding theoretical NLO cross section at each stage. Cross sections were calculated with MCFM using $\mu_F = \mu_R = m_Z$. For $b\bar{b}Z^0$ and Z^0*+1,2,3*-jets, PDF CTEQ6M and CTEQL1 were used, respectively. For $b\bar{b}Z^0$ a K factor of 1.51 was measured. This does not vary significantly with the chosen p_T^b or $\mid \eta^b \mid$ cuts [100]. Cross section uncertainties include contributions from the energy scale ($b\bar{b}Z^0$ [95], $t\bar{t}$, Z^0+1-jet [104] and Z^0+2-jets [106]) and parton distribution function ($b\bar{b}Z^0$ and $t\bar{t}$ only). No uncertainty is available for the Z^0+3-jets calculation. Here, *b*-veto corresponds to the veto of events containing *b*-quarks in the initial or final state. Preselection thresholds are quoted in GeV.

6.4 Cross section measurement

The differential cross section, $\frac{d\sigma_{fi}}{d\Omega}$, of a scattering process is defined as the probability to observe a scattered state, f, per unit solid angle, Ω, if a target state, i, is irradiated with unit luminosity, $\mathcal{L}$. The usual units are millibarns per steradian (mbsr^{-1}) where $1\ \text{mb} = 10^{-27}\ \text{cm}^2$.

Mathematically, the differential cross section is expressed in Equation 6.1. Here, W_{fi} is the transition rate from state i into state f and $\mathcal{L}$ is the collision luminosity.

$$\frac{d\sigma_{fi}}{d\Omega} = \frac{dW_{fi}}{\mathcal{L}d\Omega} \tag{6.1}$$

According to time-dependent perturbation theory, W_{fi} can be described by the matrix element and the Lorentz invariant phase space, *LIPS*. This is shown in Equation 6.2.

$$W_{fi} = |\mathcal{M}_{fi}|^2 \times LIPS \tag{6.2}$$

For a directionless measurement, integrating Equation 6.1 over solid angle gives Equation 6.3. Here σ_X is the cross section for a given channel, X and W_X is the corresponding production rate.

$$\sigma_X = \frac{W_X}{\mathcal{L}} \tag{6.3}$$

The rate itself can be rewritten through the number of produced events, N_X, during an infinitesimal change in time dt, shown by Equation 6.4.

$$W_X = \frac{dN_X}{dt} \tag{6.4}$$

In reality, a perfect selection of signal events will not exist and hence N_X cannot be measured directly. Instead, the selection will occur with some efficiency ϵ_X, between 0 and 1, yielding a number of signal events n_X after time, t. This is shown by Equation 6.5.

$$n_X(t) = \epsilon_X N_X(t) \tag{6.5}$$

Furthermore, the number of selected signal events n_X will be indistinguishable from the background events passing the same selection process, n_B. The latter is an impurity which must be either calculated or measured independently and removed. The total number of selected events, n_T, is therefore described by Equation 6.6.

$$n_T(t) = n_X(t) + n_B(t) \tag{6.6}$$

Finally, to simplify our expression the notion of *integrated luminosity* is introduced, $\mathcal{L}_{int}$, in Equation 6.7.

$$d\mathcal{L}_{int}(t) = \mathcal{L}(t)dt \tag{6.7}$$

Combining Equations 6.3, 6.4, 6.5, 6.6 and 6.7 for $b\bar{b}Z^0$ production and assuming $n(0)_{b\bar{b}Z^0} = 0$ yields an expression for the signal cross section, Equation 6.8.

$$\sigma_{b\bar{b}Z^0} = \frac{n_T(t) - \sum_B n_B(t)}{\epsilon_{b\bar{b}Z^0}\mathcal{L}_{int}(t)} \tag{6.8}$$

$n_T(t)$ should be measured directly from data. $\mathcal{L}_{int}(t)$ will be provided as an independent measurement based on the evolution of the LHC luminosity [30]. $\sum_B n_B(t)$ can either be calculated for each background channel using the prescription shown in Equation 6.9, obtained by combining Equations 6.3, 6.4 and 6.5, or measured independently.

$$n_B = \epsilon_B \sigma_B \mathcal{L}_{int} \tag{6.9}$$

Finally, the measurement of $\epsilon_{b\bar{b}Z^0}$, that is the selection efficiency of the signal, is to be performed on Monte Carlo.

The remainder of this chapter will describe how $\epsilon_{b\bar{b}Z^0}$ and $\sum_B n_B(t)$ are to be measured using Monte Carlo and data from CMS. All of the material presented was produced by the author alone, under the supervision of A. Nikitenko (Imperial College London).

6.5 A generator level signal selection strategy

To select signal events and suppress background, a series of selection rules or *discriminators* need to be developed. This thesis will rely on cut based decision making. Though often less efficient than more advanced tools such as neural networks [107] or decision trees, this approach does provide a robust, reliable and transparent methodology that will be especially useful during early data taking.

This section describes the development of a signal selection strategy based on information provided at the generator level. At this stage only the MC four-vectors are available and so it is independent of the CMS detector resolution and object identification efficiencies. It does, however, provide a good initial basis for understanding signal and background event kinematics and topology.

Before any offline selection strategy is devised, the trigger streams of interest should be identified. The primary purpose of these should be to provide a high signal efficiency. Figure 6.8 shows the Monte Carlo electron and muon kinematics for the signal. It is clear from the p_T distributions that the online *isolated dielectron* and *relaxed dimuon* triggers are suitable choices. Here, *relaxed* indicates the lack of isolation criteria (the isolated version was not available for this study). Each has p_T thresholds of 12 GeV and 3 GeV respectively, allowing a high signal selection efficiency. Many other analyses will also rely on these triggers, hence they will be well maintained making them a good choice for an early measurement.

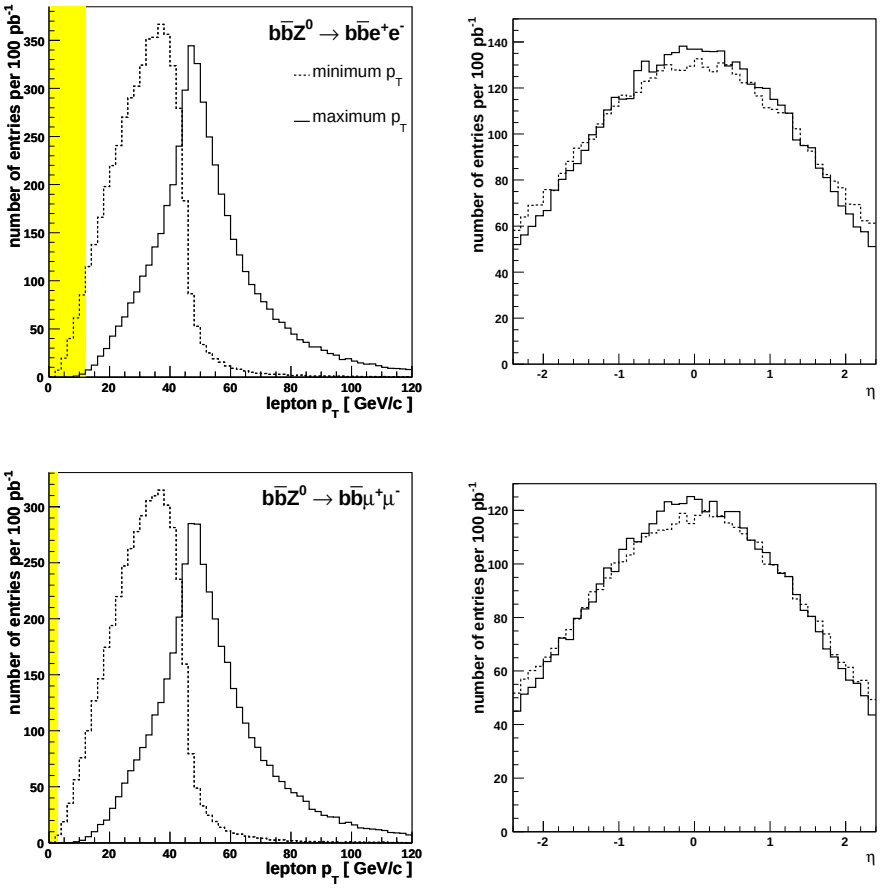

Figure 6.8: Monte Carlo lepton distributions in p_T (left) and η (right) for $b\bar{b}Z^0$ events with both $b\bar{b}e^+e^-$ (top) and $b\bar{b}\mu^+\mu^-$ (bottom) final states. Here, *minimum* p_T and *maximum* p_T correspond to the minimum and maximum p_T lepton in each event. Regions below the isolated double electron and relaxed double muon trigger thresholds are highlighted in yellow as appropriate.

Any further discrimination between signal and Z^0*+1,2,3*-jets or $t\bar{t}$ backgrounds is difficult with maximum and minimum p_T lepton information as the distributions are

similar. Instead the lepton information can be exploited by reconstructing the invariant dilepton mass. This can be calculated using the relations shown in Equation 6.10 and 6.11.

$$E_l^2 = p_l^2 + m_l^2 \tag{6.10}$$

$$(E_{l^+} + E_{l^-})^2 = (p_{l^+} + p_{l^-})^2 + m_{ll}^2 \tag{6.11}$$

Here E is energy, p is momentum and m is mass. l corresponds to either the electron or muon flavour. m_{ll} is therefore the dilepton invariant mass which corresponds to the Z^0 mass in signal and Z^0+*1,2,3*-jets background events. If we make the approximation $m_l \sim 0$, we arrive at the relation shown in Equation 6.12. Here, θ is the angle between the leptons.

$$m_{ll} \approx \sqrt{2 \mid p_{l^+} \mid\mid p_{l^-} \mid (1 - \cos\theta)} \tag{6.12}$$

Figure 6.9 shows the calculated dielectron invariant mass distributions at the generator level. It is clear that the $t\bar{t} \rightarrow b\bar{b}e^+e^-\nu_e\bar{\nu}_e$ contribution is relatively small compared with that of the other backgrounds. However, the possibility for fake electrons to be reconstructed in other $t\bar{t}$ decay modes still makes this background potentially significant. From the m_{ee} mass shape, it is clear that a mass window centred on the Z^0 peak would reject a significant fraction of $t\bar{t}$ events.

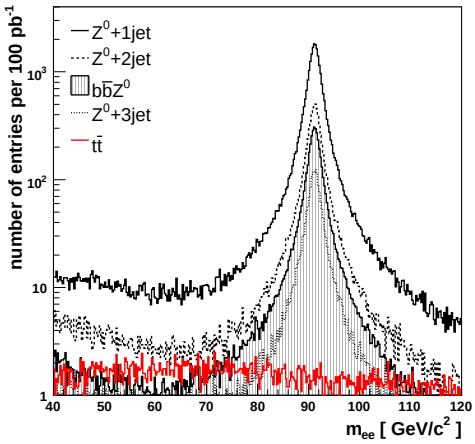

Figure 6.9: Monte Carlo dielectron invariant mass distribution. Only events with a Monte Carlo $e^{\pm}$ pair in the final state are considered.

Figure 6.10 shows the b-quark kinematics for both $b\bar{b}Z^0$ and $t\bar{t}$ events. It is clear that a common feature of the signal is the presence of a single central b-quark, albeit at low p_T which will make it hard to identify on reconstruction. Conversely, the $t\bar{t}$ background has a more symmetrical b-quark topology and far fewer escape

the tracker acceptance of $\eta = \pm\, 2.4$. Combined with the higher p_T of both b-quarks, and thus more efficient b-tagging, we expect most $t\bar{t}$ events to reconstruct at least 2 b-jets.

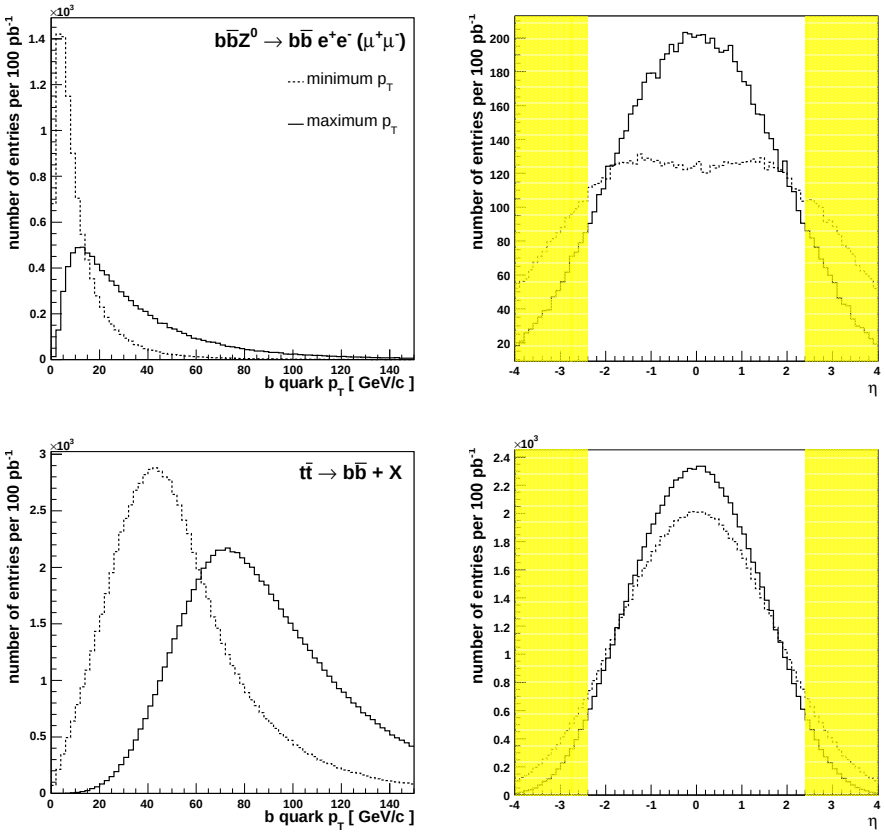

Figure 6.10: Monte Carlo b-quark distributions in p_T (left) and η (right) for $b\bar{b}Z^0$ (top) and $t\bar{t}$ (bottom) events. Here, *minimum* p_T and *maximum* p_T correspond to the minimum and maximum p_T b-quark in each event. Only signal events with the subsequent decay $Z^0 \rightarrow e^+e^-$ or $Z^0 \rightarrow \mu^+\mu^-$ and $t\bar{t}$ events containing 2 b-quarks are included. Regions in η which are beyond the tracker acceptance are highlighted in yellow.

The signal selection strategy will therefore include the identification of a single high p_T b-jet. One should recall from Section 6.1.1 that the five-flavour scheme for calculating the signal cross section demands at least one high-p_T b-quark. This scheme is *not* a suitable probe of the four-flavour calculation, however, which stipulates that both b-quarks should be above a 15 GeV/c cut.

Finally, the generator level E_T^{miss} distributions for $b\bar{b}Z^0$ and $t\bar{t}$ are shown in Figure 6.11. This is calculated using the four-vectors of stable particles in the final state as input. Neutrinos are then excluded and the imbalance in visible E_T is calculated.

The figure shows a larger tail in the $t\bar{t}$ distribution indicating that E_T^{miss} could also provide a powerful discriminant against this background.

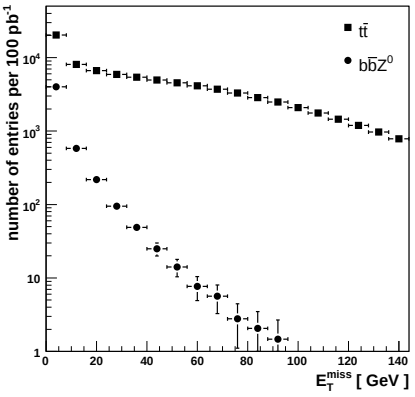

Figure 6.11: Monte Carlo E_T^{miss} distributions.

With these selection discriminators in mind, the full signal selection strategy adopted for this chapter is itemised below.

- The isolated isolated dielectron and relaxed dimuon triggers.
- 2 electrons ($e^{\pm}$) or muons ($\mu^{\pm}$) with p_T above a configurable threshold (offline).
- A dielectron or dimuon mass window with configurable boundaries (offline).
- 1 b-jet with p_T above a configurable threshold (offline).
- A veto of all events containing more than one jet (tagged or otherwise) above the b-jet threshold and within the tracker acceptance (offline).
- Missing E_T below a configurable threshold (offline).

Before using these variables as signal selection discriminators, they must first be tested for correlations. Two or more parameters which are strongly correlated will lose their discrimination power. For this purpose the Pearson product-moment correlation coefficient, defined by Equation 6.13, has been used. Here, x and y are the variables under scrutiny. Correlated, anti-correlated and uncorrelated variables will

return a value of $r_{xy} = 1, -1, 0$ respectively.

$$r_{xy} = \frac{n\sum_{i=1}^{n} x_i y_i - \sum_{i=1}^{n} x_i \sum_{i=1}^{n} y_i}{\sqrt{n\sum_{i=1}^{n} x_i^2 - \left(\sum_{i=1}^{n} x_i\right)^2 - n\sum_{i=1}^{n} y_i^2 - \left(\sum_{i=1}^{n} y_i\right)^2}} \tag{6.13}$$

The correlation of each combination is shown in Table 6.3 for signal and background events with an $e^{\pm}$ pair in the final state only. Except for the $p_{Te^{\pm}}$, m_{ee} and $p_{Tb,\bar{b}}$, E_T^{miss} pairs, all values are almost uncorrelated, with $\mid r_{xy} \mid$ values of 0.12 or less. The stronger correlation seen for the former is to be expected as they contain common information (the p_T of the highest p_T electron). Nevertheless, a lepton p_T cut is inevitable due to the use of the dielectron trigger stream and the $t\bar{t}$ rejection power provided by the dielectron mass constraint may prove powerful, especially after reconstruction when fake electrons are expected. The $p_{Tb,\bar{b}}$, E_T^{miss} pair also exhibits a non-negligible correlation. This is also unsurprising as high p_T jets are likely to leave a larger imbalance which will manifest itself as E_T^{miss}. This effect is larger for signal events than it is for the $t\bar{t}$ background, which the E_T^{miss} variable aims to suppress. For these reasons, all selection discriminants listed in this section will be imposed in the final scheme.

		r_{xy}				
x	y	$b\bar{b}Z^0$	Z^0+1-jet	Z^0+2-jets	Z^0+3-jets	$t\bar{t}$
$p_{Te^{\pm}}$	$p_{Tb,\bar{b}}$	0.12	-	-	-	0.10
$p_{Te^{\pm}}$	E_T^{miss}	0.04	0.00	-0.01	-0.01	-0.12
$p_{Te^{\pm}}$	m_{ee}	0.41	0.33	0.32	0.29	0.69
$p_{Tb,\bar{b}}$	E_T^{miss}	0.37	-	-	-	0.20
$p_{Tb,\bar{b}}$	m_{ee}	0.06	-	-	-	0.13
E_T^{miss}	m_{ee}	0.03	-0.01	-0.01	-0.01	-0.05

Table 6.3: Correlation of Monte Carlo event selection discriminators, calculated from Equation 6.13. Here, $p_{Te^{\pm}}$ corresponds to the minimum p_T electron and $p_{Tb,\bar{b}}$ is for the maximum p_T b-quark. Only events with a $e^{\pm}$ pair in the final state are considered.

Though a four-vector study is useful in understanding event kinematics, any accurate prediction of the discrimination power of a selection scheme requires the use of detector simulation and event reconstruction. The impact of these two processes is the subject of the following section.

6.6 Detector simulation and event reconstruction

Before any measurement of $\epsilon_{b\bar{b}Z^0}$ can be made, the signal and backgrounds must first be understood at the reconstruction level. The lepton and b-jet identification (or tagging) algorithms also require optimisation.

Detector simulation of each event has been performed within CMSSW 1.4.9. The signal reconstruction was run using version 1.6.7 and the backgrounds with version 1.5.2. No detector pile-up or misalignment scenarios have been included. These effects are expected to include degraded particle identification efficiencies and increased fake rates. Although misalignment will be a key factor for early data, pileup is expected to be less important. Both effects should be considered in the next iteration of this analysis.

6.6.1 Particle identification

To quantify the performance of the particle identification algorithms, identified objects must be qualified as *real* or *fake* against the generator level four-vectors. This is done by matching the position of the reconstructed objects in η-ϕ space with the direction of the Monte Carlo particle under study at its vertex.

Analogous to jet reconstruction, the MC matching cone has radius ΔR, defined by Equation 3.2. The size of the matching cone to be used is flavour specific, and has been tuned by examining the generator level - reconstructed object separation. In Figures 6.12 and 6.13, the Monte Carlo level particles of interest, $e^{\pm}$, $\mu^{\pm}$, b and $\bar{b}$ are paired with their closest reconstructed candidate of the same flavour and the ΔR separation is recorded. Through this process, ΔR cuts of 0.1, 0.05 and 0.4 have been fixed for electrons, muons and b-jets respectively. They have been chosen to maintain a high matching efficiency (of $\gtrsim$99 % for leptons and $\sim$80 % for bjets) whilst removing the tail in the distribution to avoid incorrect associations.

Prior to any signal selection study, the particle identification algorithms, run both during the reconstruction and analysis stages, require tuning. To assess the performance of the identification procedure in each case and to aid the tuning process, the identification efficiency, ϵ_{obj}, and collection purity, p_{obj}, have been defined in Equations 6.14 and 6.15. Here, a *candidate* is defined as a reconstructed object of a given flavour (e.g. electron, muon or b-jet) prior to its positive or negative identification by the algorithm under study.

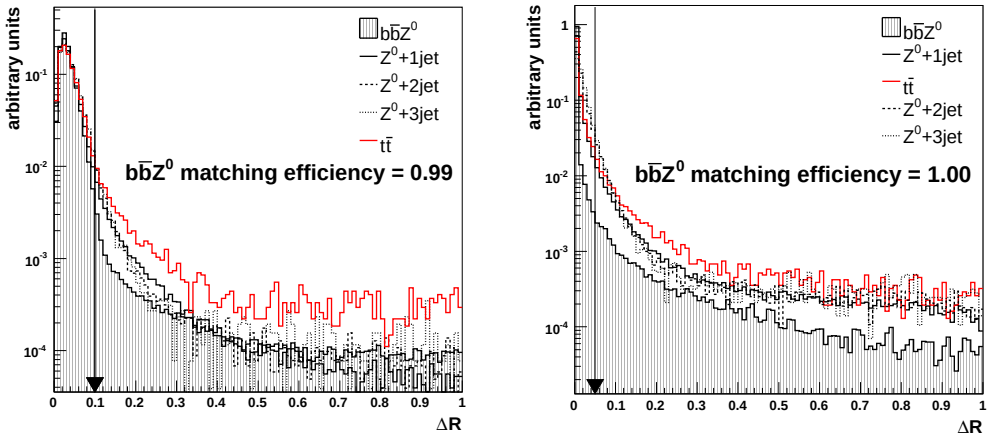

Figure 6.12: ΔR separation of generator level $e^{\pm}$ (left) and $\mu^{\pm}$ (right) from their closest reconstructed equivalent. Consequently, MC matching of reconstructed electrons and muons uses cones with radii 0.1 and 0.05 respectively.

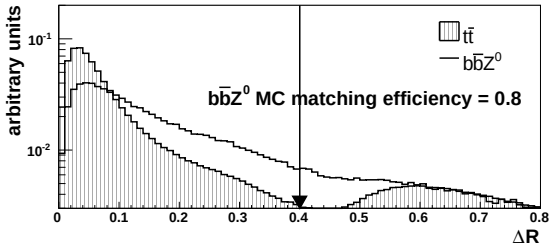

Figure 6.13: ΔR separation of b-quarks from their closest reconstructed jets. Consequently, MC matching of reconstructed jets uses a cone of radius 0.4.

$$\epsilon_{obj} = \frac{n\,(\text{MC matched candidates passing selection})}{n\,(\text{MC matched candidates})} \tag{6.14}$$

$$p_{obj} = \frac{n\,(\text{MC matched candidates passing selection})}{n\,(\text{candidates passing selection})} \tag{6.15}$$

At this point it should be emphasised that the variable p_{obj} will be used to represent purity rather than momentum. In general, as each identification discriminator is varied, a gain in efficiency will result in a loss in purity, or vice-versa. As this study is preparing for a high purity signal measurement, the tendency will be towards high purity particle identification. The tuning of the identification algorithms is discussed in Sections 6.6.2, 6.6.3 and 6.6.4.

6.6.2 Electron identification

For this analysis, the default set of cuts provided by the reconstruction software is used. The identification discriminants used by the CMSSW 1.6.7 software are listed below. The maximum and minimum thresholds for each discriminant depend on the classification of the electron *type* (e.g. showering or large bremsstrahlung). The *loose* set of default cuts have been chosen for this analysis. More information on the classification procedure can be found in [68].

- $\frac{E_{SC}}{p_{out}}$ The ratio of supercluster energy to the momentum of the electron track at the ECAL surface.
- $\frac{1}{E_{SC}} - \frac{1}{p_{out}}$ Similar to above, but this discriminant is less sensitive to the electron energy.
- $\eta_{SC} - \eta_{vertex}$, $\phi_{SC} - \phi_{vertex}$ The separation in η and ϕ of the supercluster centre from the track angle at its vertex (calculated using a helical extrapolation from the ECAL surface).
- $\phi_{SC} - \phi_{out}$ The separation in ϕ of the supercluster centre from the track position at the ECAL surface.
- $\frac{H}{E}$ The ratio of supercluster energy deposited in the HCAL to that in the ECAL.
- $\frac{E_{3x3}}{E_{5x5}}$ The ratio in energy of the central 3×3 and 5×5 crystal arrays.
- $\frac{p_{in} - p_{out}}{p_{in}}$ The loss in track momentum between the vertex and the ECAL. Also known as the *bremsstrahlung fraction*, this indicates the loss of electron energy through the tracker material.

Since this analysis uses the isolated dielectron trigger, to some extent the electrons have already been deemed isolated. However, the isolation information has been calculated from the calorimeter alone, so any tracking information is yet to be considered. As well as the discriminants listed above, the isolation of each electron according to the tracker should be included. This *tracker isolation* variable is described by the final term in the list, added below.

- $\frac{\sum_i p_i}{E_{SC}}$ Here i are tracks from the primary vertex which have an innermost direction falling within a cone of radius $0.01 < \Delta R < 0.2$, centred on the electron track (the value of 0.01 is chosen to avoid the electron track itself).

Figure 6.14 (top) shows the track isolation values for both real (MC matched) and fake (MC unmatched) electrons. It can be seen that real electrons dominate at low tracker isolation values and hence tend to be more isolated. By scanning through various isolation cuts (and accepting only electrons below this), values for the electron reconstruction efficiency, $\epsilon_{electron}$, and the collection purity, $p_{electron}$, can be recorded in each case. Figure 6.14 (bottom) shows the result of this procedure. Choosing an isolation discriminator of 0.3 optimises the trade-off between efficiency and purity and is shown on the figure. More specifically, with respect to a discriminator of ∞ (no isolation) $\epsilon_{electron}$ drops by 1.6 % to 0.984, but $p_{electron}$ increases by 10.0 % from 0.808 to 0.908.

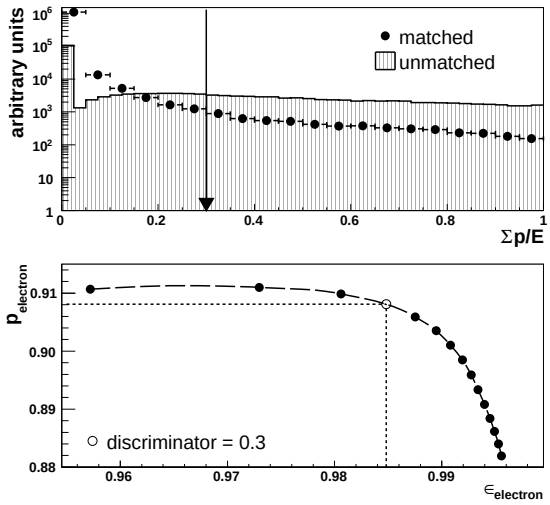

Figure 6.14: (top) Isolation discriminators for real (MC matched) and fake (MC unmatched) reconstructed electrons in signal events passing the dielectron trigger. (bottom) $p_{electron}$ vs. $\epsilon_{electron}$ for isolation discriminators 0.025 to 0.375 in increments of 0.025. The optimal discriminator of 0.3 is identified.

Figure 6.15 shows $\epsilon_{electron}$ and $p_{electron}$ as a function of E_T for electrons passing all discriminator thresholds, including track isolation. Only events passing the isolated dielectron trigger are considered. $\epsilon_{electron}$ plateaus to 0.98 for all channels above 12 GeV, which is the trigger threshold. Also, $p_{electron}$ is greater than 0.8 for signal electrons above ∼16 GeV. $p_{electron}$ for $t\bar{t}$ events plateaus with a 10 % lower value.

This is due to the large fraction of events without any MC electrons, which provide a plentiful source of potential fakes. It is worth noting that at the trigger threshold, $p_{electron}$ is still on a sharp rise, so it may be necessary to apply an increased offline threshold during event selection.

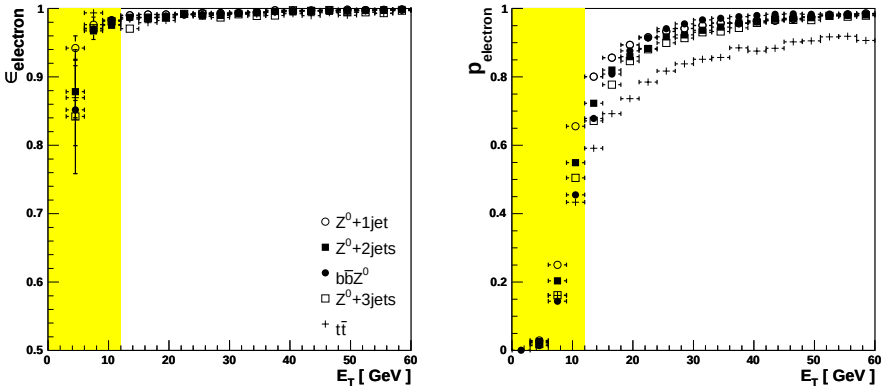

Figure 6.15: $\epsilon_{electron}$ (left) and $p_{electron}$ (right) as a function of reconstructed E_T. Only events passing the dielectron trigger are considered. Regions below the trigger threshold are highlighted.

6.6.3 Muon identification

Unlike the isolated dielectron stream, the relaxed dimuon trigger does not require the leptons to be isolated. The collection of reconstructed muons can therefore include secondary muons emitted from jets, or indeed fake muons reconstructed from the jets themselves. To minimise these contributions, a similar approach to electron identification can be taken.

Figure 6.16 is analogous to Figure 6.14 for muons. Figure 6.16 (top) shows a significant excess of Monte Carlo matched muons being isolated. By imposing track isolation with a discriminator of 0.3 (Figure 6.16 (bottom)), ϵ_{muon} drops by 1.3 % to 0.987, but the purity increases by 3.1 % from 0.902 to 0.933.

Figure 6.17 shows ϵ_{muon} and p_{muon} as a function of muon E_T. Only events passing the trigger are considered. ϵ_{muon} for signal events is consistently > 0.95 and rises with increasing p_T. p_{muon} plateaus above $\sim$16 GeV indicating, as for the electron branch, that an increased offline threshold may be necessary during event selection.

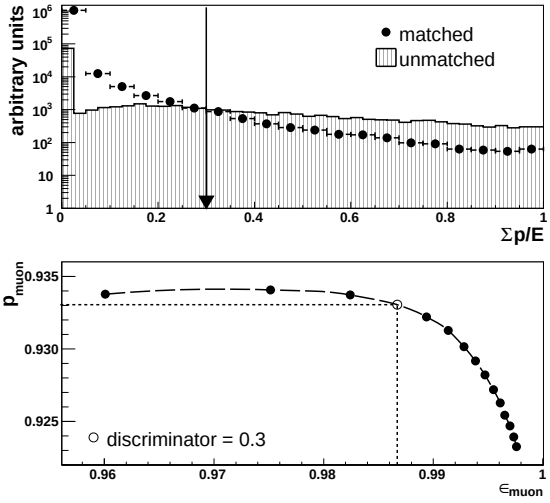

Figure 6.16: (top) Isolation discriminator for real (MC matched) and fake (MC unmatched) reconstructed muons over signal events passing the dimuon trigger. (bottom) p_{muon} vs. ϵ_{muon} for discriminators 0.025 to 0.375 in increments of 0.025. The optimal discriminator of 0.3 is identified.

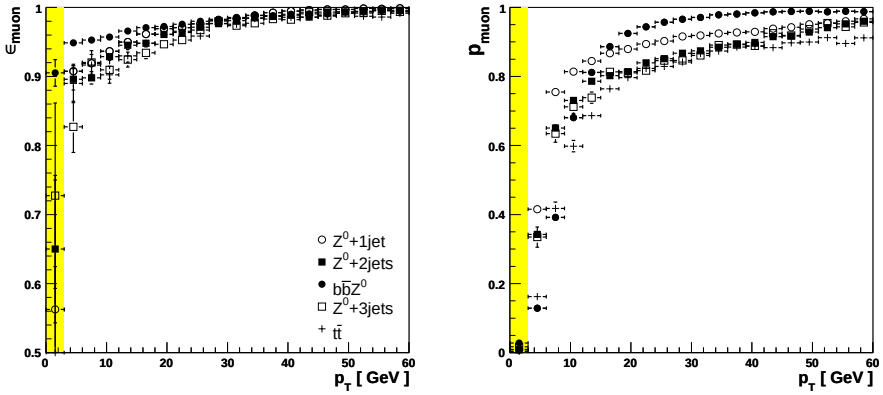

Figure 6.17: ϵ_{muon} (left) and p_{muon} (right) as a function of reconstructed E_T for all events passing the dimuon trigger. Regions below the trigger threshold are highlighted in yellow.

6.6.4 *b*-jet tagging

For b-tagging in CMS, various algorithms are available. For its robustness, which will be useful for a measurement with early data, the track counting approach has been taken (Section 3.5.1). This lists the impact parameter significances for all tracks associated with a given jet. They are put in decreasing order and the Nth value serves as the algorithm's discriminator. All jets have been reconstructed using the iterative cone algorithm with a cone size $\Delta R = 0.5$. N is usually chosen to be

2 or 3. $N = 3$ will be used here to achieve a high purity b-jet collection.

Figure 6.18 (left) shows the track counting discriminator distributions for both MC matched and unmatched jets over signal events. Only events containing a single reconstructed jet with $p_T > 20$ GeV and within the tracker acceptance have been considered, in line with the event selection strategy. 27.3 % of MC matched jets with 3 associated tracks or more were recorded with negative impact parameters. These are used to signify that the impact parameter of the discriminating track is situated on the opposite side of the primary vertex from the jet direction. They arise most frequently as a result of mismeasured track parameters, but also badly reconstructed primary vertex positions, fake tracks, or multiple scattering at low radii.

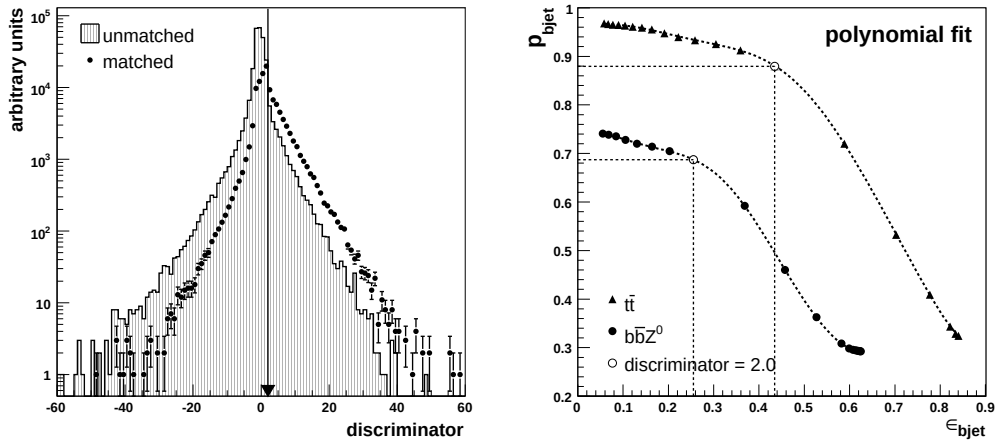

Figure 6.18: (left) b-jet discriminator for real (MC matched) and fake (MC unmatched) reconstructed jets over signal events containing a single reconstructed jet with $p_T > 20$ GeV and within the tracker acceptance. (right) p_{bjet} vs. ϵ_{bjet} for discriminators -10.0 to 9.0 in increments of 1.0. The optimal discriminator of 2.0 is identified.

Figure 6.18 (right) highlights the tuned discriminator threshold of 2.0, meaning jets with a value greater than this are tagged. With this configuration $p_{bjet} = 0.687$ and $\epsilon_{bjet} = 0.255$ for the signal. In comparison, over $t\bar{t}$ events $p_{bjet} = 0.879$ and $\epsilon_{bjet} = 0.435$. This large reduction in b-tagging performance over signal events is due to the lower b-jet E_T. This was alluded to in Section 6.5 and Figure 6.10.

Figure 6.19 shows ϵ_{bjet} and p_{bjet} as a function of jet E_T. The behaviour of ϵ_{bjet} is similar for both $b\bar{b}Z^0$ and $t\bar{t}$ events. This confirms the explanation of a lower signal ϵ_{bjet} due to softer jets. The behaviour of p_{bjet} with jet E_T does vary between channels, however. The loss of purity between ∼15 and 140 GeV must be explained by a greater source of fakes from initial and final state gluon emission.

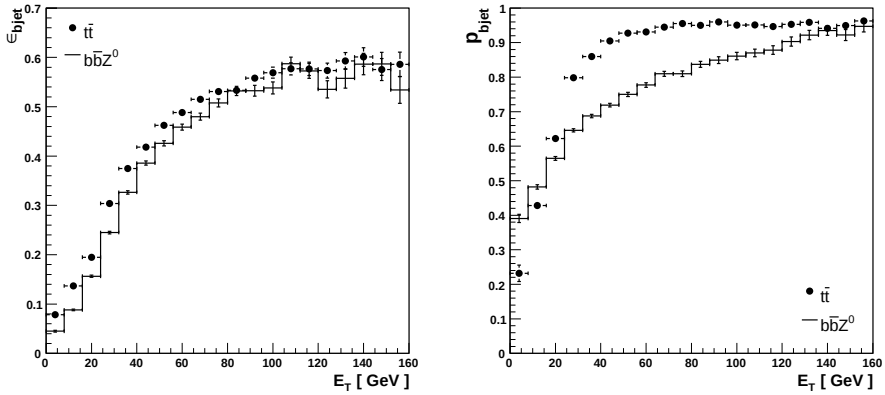

Figure 6.19: ϵ_{bjet} (left) and p_{bjet} (right) as a function of E_T for signal events containing a single reconstructed jet with $E_T > 20$ GeV and within the tracker acceptance.

6.6.5 Energy scaling

Sections 3.5 and 3.6 described how reconstructed jet and E_T^{miss} energies require an E_T and η parameterised scale to restore equality between the reconstructed jets and their original parton source. This process is of vital importance to the choice of b-jet E_T and E_T^{miss} event selection cuts.

Figure 6.20 (top) shows the ratio of b-jet E_T to the p_T of its closest b-quark as a function of b-quark p_T. Above ∼50 GeV, the corrected jets stabilise to 0.7 of the original quark p_T. This is for two main reasons. Firstly, the corrections applied here are just to the level of the component hadrons. The final stage (down to the parton at the source of the jet) has not been considered. Hadron level jets will tend to underestimate the energy of the original parton. Secondly, the correction algorithm used has been calibrated using jets from gluons. Since b-jets have a higher track multiplicity, they tend to lose more energy in the tracker material. As such, the corrections applied are not expected to compensate for the full energy loss. Once a set of corrections tuned specifically for b-jets become available, they should be used instead. Between 40 GeV and 70 GeV, the raw jets provide a closer description of the original b-quark p_T, than those after corrections. At ∼20 GeV, the corrected E_T exceeds that of its corresponding b-quark and the difference grows rapidly below. This is the most crucial region for $b\bar{b}Z^0$ production.

Figure 6.20 (middle) shows the deviation from the matched b-quark p_T increases with $\mid \eta \mid$. Again, this can be explained by an increase in tracker material in the forward regions. Although the corrected E_T can be more than twice as large as the

b-quark p_T at large η, the corrections do still offer a closer approximation than the raw jet energies.

The ratio of reconstructed to generator level E_T^{miss} as a function of generator E_T^{miss} is shown in Figure 6.20 (bottom) for $t\bar{t}$ events. Above ∼40 GeV, the corrections perform well, bringing the reconstructed value closer to that of the generator level and keeping within 0.9 and 1.2 of this target value. Above ∼50 GeV the value is underestimated as a direct result of the underestimated jet energies. Below 40 GeV, both corrected and raw E_T^{miss} tend to be significantly overestimated. Figure 6.11 showed that the majority of $b\bar{b}Z^0$ and $t\bar{t}$ events consist of < 20 GeV E_T^{miss}. Within this region we can expect an overestimation by more than a factor 3. However, since both $t\bar{t}$ and $b\bar{b}Z^0$ events exhibit similar E_T^{miss} performance on reconstruction, the discrimination power of this variable may still be preserved.

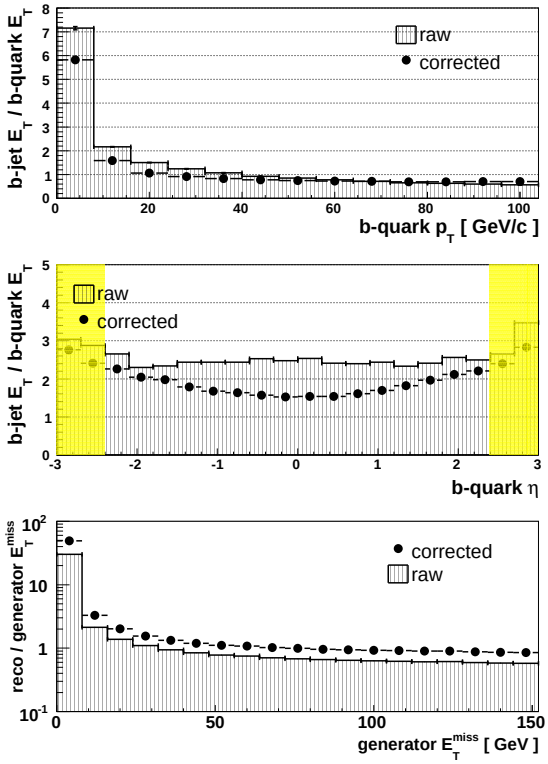

Figure 6.20: Performance of the b-jet E_T and E_T^{miss} correction algorithms. (top and middle) Ratio of MC matched b-jet E_T to its matched b-quark p_T. Shown vs. b-quark p_T and η for $b\bar{b}Z^0$ events. (bottom) Ratio of E_T^{miss} calculated at the reconstructed and generator level. Shown vs. generator E_T^{miss} for $t\bar{t}$ events.

6.6.6 E_T^{miss} and dilepton invariant mass reconstruction

This section profiles the performance of the E_T^{miss} and Z^0 mass reconstruction which are both integral to the chosen event selection strategy. The reconstructed dilepton invariant mass, calculated from Equation 6.11 using the two highest E_T electrons or muons from each event is shown in Figure 6.21. Only events where these two leptons have opposite charge are included.

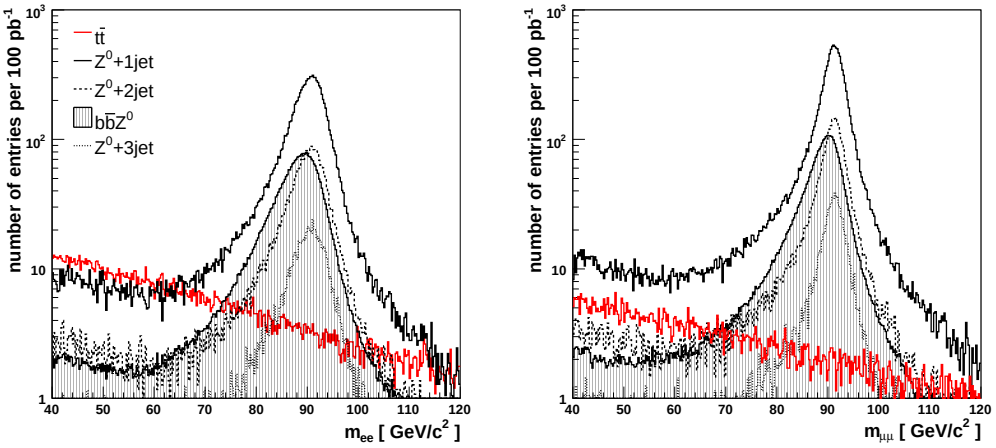

Figure 6.21: Reconstructed dielectron (left) and dimuon (right) mass distributions. For both $b\bar{b}Z^0$ and Z^0+*1,2,3*-jets channels we consider only the $Z^0 \rightarrow e^+e^-$ and $Z^0 \rightarrow \mu^+\mu^-$ branches as appropriate. For all channels we consider only events containing two or more reconstructed leptons of the correct flavour. The two with the highest E_T (p_T) are used in the calculation. If the charges of these two leptons are the same then the event is not included.

The Z^0 peaks reconstructed from the $Z^0 \rightarrow \mu^+\mu^-$ branch are narrower due to the better muon p_T resolution. From the figure it can be seen that more $t\bar{t}$ events are present for the dielectron mass distribution and that this background is more significant than first indicated by the Monte Carlo in Figure 6.21. This is due to the inclusive nature of the $t\bar{t}$ background. Events with fewer than two MC electrons or muons can still reconstruct a valid Z^0 mass if fakes are introduced. Since the reconstructed electron purity is slightly lower than the corresponding purity for muons at low E_T this reducible background will be larger. The figure also indicates that a dilepton mass window should be a powerful $t\bar{t}$ suppressant.

The reconstructed E_T^{miss} distributions for $b\bar{b}Z^0$ and $t\bar{t}$ events are shown in Figure 6.22. As well as applying jet energy corrections to the E_T^{miss}, reconstructed muons are deducted from the imbalance when they are present. The slightly larger mean E_T^{miss} in the $Z^0 \rightarrow \mu^+\mu^-$ branch of the signal is solely due to the finer p_T resolution

of the muons. Despite the various discrepancies with E_T^{miss} at the generator level, this event parameter still seems to offer a powerful selection discriminant for signal over $t\bar{t}$.

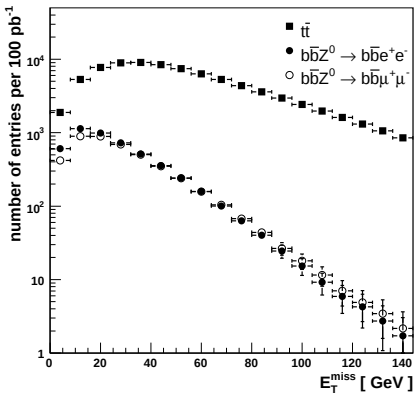

Figure 6.22: Reconstructed E_T^{miss} distributions.

6.7 Signal selection efficiency measurement

At this stage, the most distinctive features of the signal events compared with background have been understood at the reconstruction level. The b-jet tagging and lepton identification algorithms have also been tuned for our purposes. In order to finalise the selection strategy and make an $\epsilon_{b\bar{b}Z^0}$ measurement from Monte Carlo, the various event selection thresholds must first be optimised. This is the subject of the next section.

6.7.1 Tuning the selection parameters

The tuning process starts with a fixed set of cuts which are deemed *reasonable*, based on the reconstructed event distributions shown in the previous section. The selection efficiency of the signal and each background channel is then recorded. At this point, a single selection threshold is adjusted by a small increment and the updated efficiencies are measured. Once a full scan has been performed for the first parameter, the threshold is optimised and fixed. The process is then repeated for each of the remaining selection discriminators. Since some non-negligible correlations exist between discriminators, the process described above was repeated several

times, using the tuned values from the previous run as the new starting point in each case.

The m_{ll} window is slightly more complicated as it involves tuning two thresholds, an upper and a lower bound. Figure 6.21, however, shows that the high mass edge is steeper than the lower mass edge of the reconstructed Z^0 peak, making it possible to perform a more reliable estimate of the former. 105 GeV/c^2 has been chosen for both the dielectron and dimuon mass distributions.

Throughout this process, the goal is to maintain a high signal event purity without incurring a dramatic loss in signal selection efficiency. Selection efficiency for any channel, ϵ_X was defined in Equation 6.5. The corresponding channel purity, p_X is defined in Equation 6.16 below. Again, n_X is the number of events from channel X that pass the full selection process and the sum is over all considered channels X'.

$$p_X = \frac{n_X}{\sum_{X'} n_{X'}} \tag{6.16}$$

The selection efficiency for each channel (during the final tuning run), as a function of the lepton p_T, b-jet p_T, E_T^{miss} and low mass edge cuts are shown in Figures 6.23 to 6.26. The corresponding signal purity, $p_{b\bar{b}Z^0}$ values, calculated using Equation 6.16, are shown in Figure 6.27. The chosen optimal thresholds are indicated by the arrows. The choices are discussed in detail in the remaining section, but in general where the optimum threshold does not correspond to the maximum purity, a sharp decline in signal efficiency is present. It is also clear from the figures that the signal selection efficiency exceeds that of the closest background by more than an order of magnitude. The efficiencies measured in the muon branch are also a factor of $\sim$2 higher than for the electron branch. As will be seen, this is primarily due to a higher trigger efficiency. The hardest background to reject also seems to be Z^0+3-jets, but this has the smallest cross section so is less crucial to suppress.

It should be noted that during this process the $Z^0 \rightarrow e^+e^-$ and $Z^0 \rightarrow \mu^+\mu^-$ branches of the signal and backgrounds were considered separately. This was because, in principle, the optimum thresholds for each branch could have been entirely different. However, since the behaviour between the two branches was similar, identical thresholds were chosen for simplicity.

Figure 6.23 (left) and (right) shows the scan through offline electron E_T and muon p_T cuts. All efficiencies remain flat in the highlighted regions since the triggers

already impose a higher threshold. Both figures indicate a steady drop in efficiency above this whilst Figure 6.27 (top) shows a flat signal purity. An offline threshold of 15 GeV was chosen to ensure the electron and muon collection purity clears the sharp rises indicated in Figures 6.15 (right) and 6.17 (right).

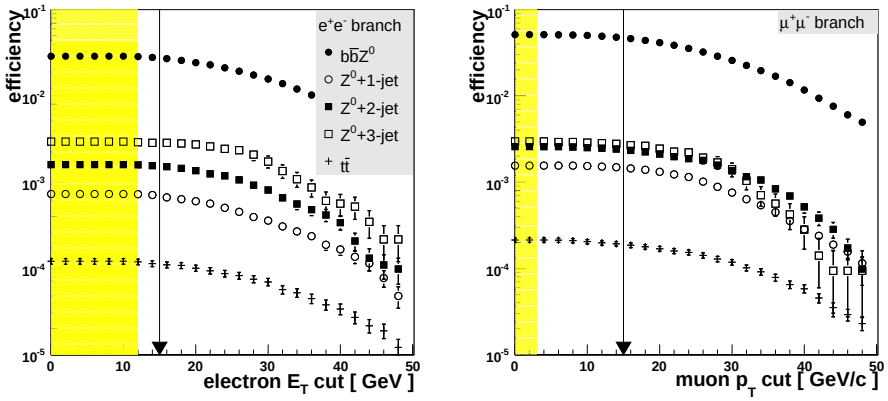

Figure 6.23: Event selection efficiency for all channels as a function of the offline lepton E_T cut. $Z^0 \rightarrow e^+e^-$ (left) and $Z^0 \rightarrow \mu^+\mu^-$ (right) selection streams are considered separately. The regions in yellow represent the cuts enforced by the 12 GeV and 3 GeV dielectron and dimuon triggers.

Figure 6.24 (left) and (right) shows the scan through single b-jet E_T cuts. In both cases a peak signal efficiency is visible at 35 GeV. The b-jet selection criterion requires one and only one jet above the chosen threshold (which is subsequently b-tagged). Since low E_T cuts will include the soft jets produced by initial and final state gluon radiation, a sharp decrease in the selection efficiency is measured in this region. A steady decline is measured beyond the 35 GeV peak due to a reduced number of b-jets with the required E_T. The signal event purity, shown by Figure 6.27 (second), reaches a minimum at ∼20 GeV and then begins to rise. The threshold of 50 GeV was chosen as an optimal trade-off between falling efficiency and increasing purity.

The scan through E_T^{miss} cuts is shown in Figure 6.25 (left) and (right). As the maximum E_T^{miss} threshold is increased the event selection efficiency for both signal and background increases, as expected. Importantly, the $t\bar{t}$ background increases more rapidly indicating that reconstructed E_T^{miss} will be a powerful suppressant. In comparison, Figure 6.27 (third) shows the signal purity rises to a maximum at 60 GeV before dropping slightly. With this in mind, an optimised threshold of 60 GeV has been chosen.

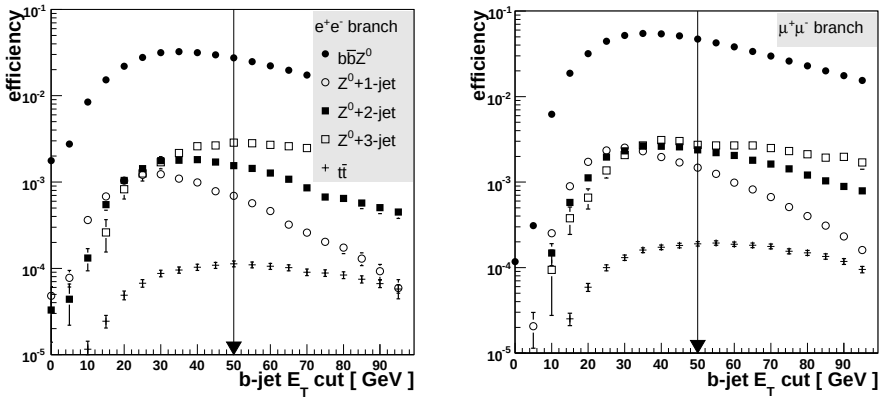

Figure 6.24: Event selection efficiency for all channels as a function of the b-jet E_T cut. $Z^0 \rightarrow e^+e^-$ (left) and $Z^0 \rightarrow \mu^+\mu^-$ (right) selection streams are considered separately.

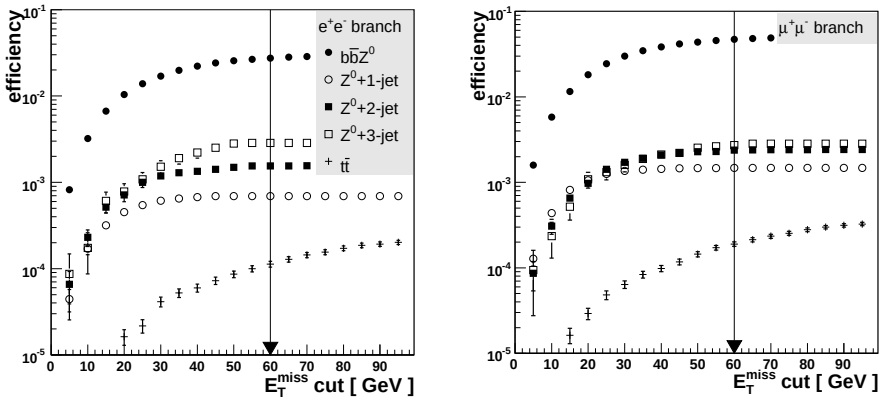

Figure 6.25: Event selection efficiency for all channels as a function of the E_T^{miss} cut. $Z^0 \rightarrow e^+e^-$ (left) and $Z^0 \rightarrow \mu^+\mu^-$ (right) selection streams are considered separately.

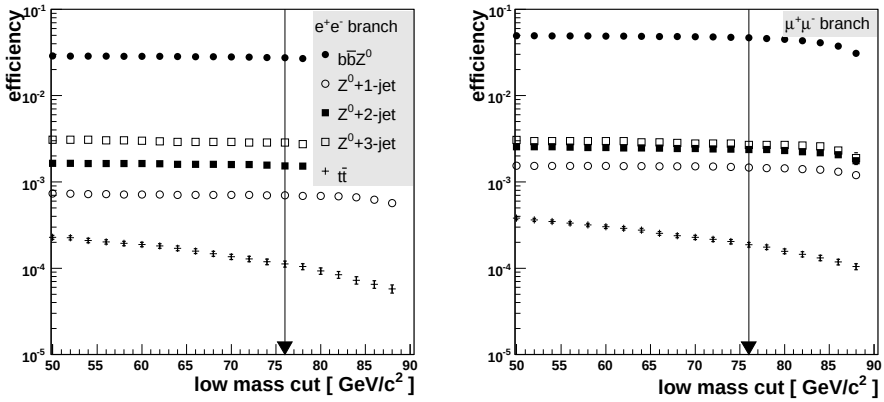

Figure 6.26: Event selection efficiency for all channels as a function of the low m_{ll} cut. $Z^0 \rightarrow e^+e^-$ (left) and $Z^0 \rightarrow \mu^+\mu^-$ (right) selection streams are considered separately.

Finally, the scan through the low edge of the dilepton mass window is shown in Figure 6.26. For both muon and electron branches a flat efficiency is recorded until ∼85 GeV. This corresponds to the rising edge of the reconstructed Z^0 peaks shown in Figure 6.21. The $t\bar{t}$ background follows a Landau distribution and so falls off exponentially with the increasing low mass cut. Figure 6.27 (bottom) shows the signal purity reaches its maximum at 76 GeV so this is the chosen cut.

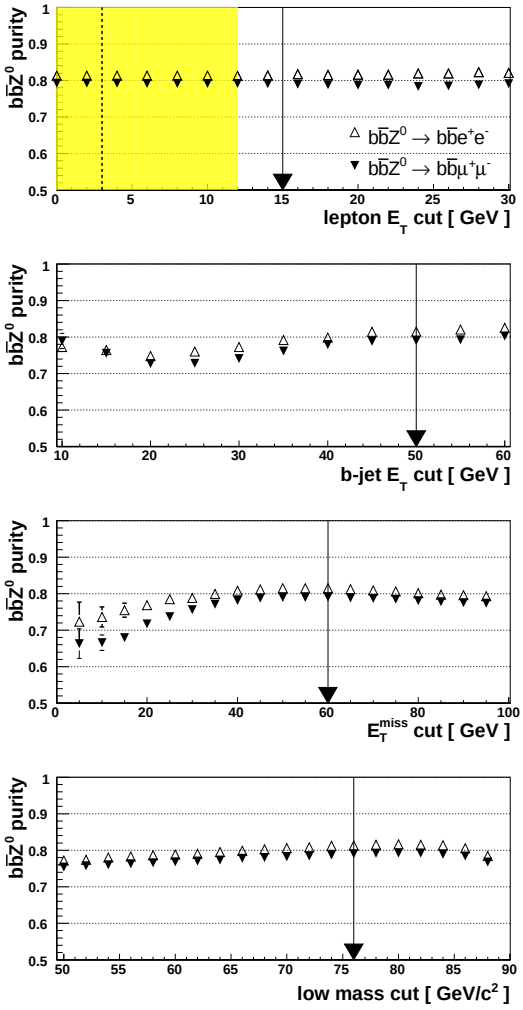

Figure 6.27: $b\bar{b}Z^0$ selection purity as a function of the (top) offline lepton E_T cut, (second) b-jet E_T cut, (third) E_T^{miss} cut, (bottom) low m_{ll} cut. $Z^0 \rightarrow e^+e^-$ and $Z^0 \rightarrow \mu^+\mu^-$ selection streams are considered separately. The region in yellow (top) represents the 12 GeV dielectron trigger. To the left of the dotted line marks the dimuon trigger.

Table 6.4 summarises the initial and tuned values alongside the tuning range and increment for each variable.

discriminator	initial value	tuning range	increment	tuned value
2 × lepton p_T (GeV/c)	12	0 - 50	2	15
m_{ll} lower cut (GeV/c^2)	60	50 - 90	2	76
E_T^{miss} (GeV)	80	0 - 100	5	60
single b-jet E_T (GeV)	80	0 - 100	5	50

Table 6.4: Summary of the tuning process for each event selection discriminator.

6.7.2 Monte Carlo signal and background selection efficiencies

Now the particle identification algorithms have been tuned and event selection thresholds have been optimised, the full selection chain can be run over both signal and background events to measure the performance at each stage. This is the subject of the following section.

Tables 6.5 and 6.6 break down the signal and background selection efficiencies for the dielectron and dimuon streams respectively. Both the cumulative efficiency and the efficiency relative to the previous stage are shown. Three separate streams are evident within each table, highlighting three different selection strategies. The first, designed to probe the five flavour cross section calculation scheme, looks for one reconstructed jet only, which is b-tagged, and has been the subject of this chapter so far. The second stream, which looks for multiple reconstructed jets of which at least one must be b-tagged, also probes the five-flavour scheme and has been included for comparison. It can be seen that, although the $b\bar{b}Z^0$ efficiency improves considerably: we expect 32 % and 35 % more events from the electron and muon branches respectively, the background efficiencies also increase. A 3.2 % and 5.9 % loss in signal purity is measured for the two branches. Finally, the third stream is designed to probe both the four-flavour and five-flavour calculation schemes since it passes events with two or more reconstructed b-jets above a 50 GeV E_T cut. A dramatic loss in signal efficiency can be seen from both tables, which is compensated for by a relatively smaller increase in purity.

If we examine each step of the selection process in more detail, we can measure the effectiveness of the selection criteria. The fields in white, which are implemented before the b-tagging and are common to all three selection streams, are most effective in suppressing the $t\bar{t}$ background. This is primarily due to the low trigger efficiencies and is to be expected since ∼99 % of these events do not contain two

final state electrons or muons. The trigger efficiencies are also much lower for the Z^0+*1,2,3*-jets bacgrounds than the signal (by almost a factor of 2 in each case). This is due to the kinematics of the Z^0 boson and its decay products. Figure 6.8 shows the p_T distribution for minimum p_T leptons increases rapidly near the trigger thresholds. The slight downwards shift present for the backgrounds (not shown in this thesis) serves to dramatically decrease the dilepton trigger efficiencies. The offline 2 $\times$ leptons stage is far more effective for the muon branch since the difference in offline and online p_T thresholds is greater. For both tables, the invariant dilepton mass window is the most effective offline stage. Over 90 % of signal events reaching this level pass through, compared with only $\sim$20 % of $t\bar{t}$ events. The E_T^{miss} limit is also effective as both a signal filter and background suppressant. The fields in grey focus on the jet content of each event. Their main purpose is to remove the Z^0+*1,2,3*-jets backgrounds. This is clearly achieved for all three selection streams.

For the purpose of a cross section measurement, the single b-jet selection strategy (the first stream) seems to be the preferred choice. It maintains a high $p_{b\bar{b}Z^0}$ of 0.847 and 0.874 for the electron and muon branches, which will keep effects from background systematics to a minimum. It also has a high signal efficiency, $\epsilon_{b\bar{b}Z^0}$, of 0.0275 $\pm$ 0.0002 and 0.0479 $\pm$ 0.0003 respectively. Undoubtedly, when data arrive, all three will be run in parallel. Before the signal selection efficiency quoted in this section is used to calculate the $b\bar{b}Z^0$ production cross section, however the statistical and systematic uncertainties must be estimated. Those quoted above consider binomial fluctuations only and are fairly small due to the large number of events generated. The effect of theoretical and experimental systematic uncertainties, however, are by far the most dominant and are discussed in the following section.

	selection efficiency (fraction from previous step)				
selection	bbZ^0	Z^0+1-jet	Z^0+2-jets	Z^0+3-jets	$t\bar{t}$
isolated dielectron trigger	0.58358 (1.000)	0.28921 (1.000)	0.28709 (1.000)	0.29156 (1.000)	0.00949 (1.000)
2 × electrons $p_T > 15$, $\mid \eta \mid < 2.4$	0.52475 (0.899)	0.24888 (0.861)	0.24434 (0.851)	0.24858 (0.853)	0.00659 (0.694)
opposite sign electrons	0.50570 (0.964)	0.23793 (0.956)	0.23186 (0.949)	0.23560 (0.948)	0.00555 (0.842)
$75 < m_{ee} < 105$	0.46062 (0.911)	0.21609 (0.908)	0.20927 (0.903)	0.21242 (0.902)	0.00117 (0.210)
$E_T^{miss} < 60$	0.43109 (0.936)	0.21535 (0.997)	0.20723 (0.990)	0.20482 (0.964)	0.00051 (0.438)
1 jet $p_T > 50$, $\mid \eta \mid < 2.4$	0.08933 (0.207)	0.04100 (0.190)	0.06793 (0.328)	0.07792 (0.380)	0.00024 (0.465)
1 b-jet $p_T > 50$, $\mid \eta \mid < 2.4$	0.02752 (0.308)	0.00069 (0.017)	0.00155 (0.023)	0.00287 (0.037)	0.00011 (0.474)
$\geq$ 1 jet p_T >50, $\mid \eta \mid < 2.4$	0.10710 (0.248)	0.04117 (0.191)	0.08041 (0.388)	0.11131 (0.543)	0.00035 (0.681)
$\geq$ 1 b-jet $p_T > 50$, $\mid \eta \mid < 2.4$	0.03646 (0.340)	0.00069 (0.017)	0.00233 (0.029)	0.00512 (0.046)	0.00019 (0.561)
$\geq$ 2 jet p_T >50, $\mid \eta \mid < 2.4$	0.01777 (0.041)	0.00017 (0.001)	0.01247 (0.060)	0.03338 (0.163)	0.00011 (0.215)
$\geq$ 2 b-jets $p_T > 50$, $\mid \eta \mid < 2.4$	0.00131 (0.074)	0.00000 (0.000)	0.00000 (0.000)	0.00004 (0.001)	0.00002 (0.210)
	events after 100 pb^{-1}				
	140 ± 10	22 ± 5	14 ± 4	6 ± 3	9 ± 3
	180 ± 10	22 ± 5	21 ± 5	11 ± 3	16 ± 4
	7 ± 3	0 ± 0	0 ± 0	0.1 ± 0.3	2 ± 1

Table 6.5: Cumulative selection efficiencies for the reconstructed signal and backgrounds after successive selection cuts. Only the $Z^0 \rightarrow e^+e^-$ decay branch is considered. The number of events expected after 100 pb^{-1} integrated luminosity with corresponding binomial errors are also shown. The various event selection criteria are introduced from top to bottom. Three selection streams exist shown in grey. Fields in white are common to both streams. The $p_{b\bar{b}Z^0}$ values calculated for each stream are, from top to bottom : 0.841, 0.814 and 0.956.

	selection efficiency (fraction from previous step)				
selection	bbZ^0	Z^0+1-jet	Z^0+2-jets	Z^0+3-jets	$t\bar{t}$
relaxed dimuon trigger	0.83612 (1.000)	0.45960 (1.000)	0.45530 (1.000)	0.44580 (1.000)	0.03923 (1.000)
2 × muons $p_T > 15$, $\mid \eta \mid < 2.4$	0.73565 (0.880)	0.39660 (0.863)	0.38736 (0.851)	0.37556 (0.842)	0.01019 (0.260)
opposite sign muons	0.73377 (0.997)	0.39651 (1.000)	0.38723 (1.000)	0.37542 (1.000)	0.00902 (0.885)
$75 < m_{\mu\mu} < 105$	0.66456 (0.906)	0.35654 (0.899)	0.34598 (0.893)	0.33636 (0.896)	0.00188 (0.209)
$E_T^{miss} < 60$	0.61715 (0.929)	0.35556 (0.997)	0.34293 (0.991)	0.32821 (0.976)	0.00082 (0.433)
1 jet $p_T > 50$, $\mid \eta \mid < 2.4$	0.15003 (0.243)	0.07780 (0.219)	0.11480 (0.335)	0.12823 (0.391)	0.00036 (0.443)
1 b-jet $p_T > 50$, $\mid \eta \mid < 2.4$	0.04703 (0.313)	0.00147 (0.019)	0.00238 (0.021)	0.00273 (0.021)	0.00019 (0.525)
$\geq$ 1 jet p_T >50, $\mid \eta \mid < 2.4$	0.18219 (0.295)	0.07833 (0.220)	0.13947 (0.407)	0.19301 (0.588)	0.00058 (0.708)
$\geq$ 1 b-jet $p_T > 50$, $\mid \eta \mid < 2.4$	0.06345 (0.348)	0.00148 (0.019)	0.00373 (0.027)	0.00815 (0.042)	0.00034 (0.591)
$\geq$ 2 jet p_T >50, $\mid \eta \mid < 2.4$	0.03216 (0.052)	0.00054 (0.002)	0.02468 (0.072)	0.06478 (0.197)	0.00022 (0.265)
$\geq$ 2 b-jets $p_T > 50$, $\mid \eta \mid < 2.4$	0.00257 (0.080)	0.00000 (0.000)	0.00001 (0.001)	0.00024 (0.004)	0.00004 (0.204)
	events after 100 pb^{-1}				
	230 ± 20	46 ± 7	21 ± 5	6 ± 2	16 ± 4
	320 ± 20	47 ± 7	33 ± 6	18 ± 4	29 ± 5
	13 ± 4	0 ± 0	0.0 ± 0.3	0.5 ± 0.7	3 ± 2

Table 6.6: Cumulative selection efficiencies for the reconstructed signal and backgrounds after successive selection cuts. Only the $Z^0 \to \mu^+\mu^-$ decay branch is considered. The number of events expected after 100 pb^{-1} integrated luminosity with corresponding binomial errors are also shown. The various event selection criteria are introduced from top to bottom. Three selection streams exist shown in grey. Fields in white are common to both streams. The $p_{b\bar{b}Z^0}$ values calculated for each stream are, from top to bottom : 0.874, 0.822 and 0.899.

6.7.3 Systematic uncertainty in event selection efficiencies

To measure the effect of theoretical and experimental systematic uncertainties on the value of $\epsilon_{b\bar{b}Z^0}$, Table 6.7 has been constructed. This shows the primary sources of uncertainty, their magnitude and the percentage effect on the resultant selection efficiency for each channel. For the energy measurements (lepton and jet) this effect was measured by multiplying the reconstructed (corrected) energy by a constant factor defined by the predetermined uncertainty (labelled scale in the table). This was done in both the positive and negative directions and the full event selection was re-run. For E_T^{miss}, the uncertainty was studied by varying the energy of the unclustered calorimeter towers by ± 10 % and the (corrected) jet energies by $\pm$ 7 %.

For the systematics measurement, only the single b-jet selection scheme has been considered. For the signal, the b-tagging uncertainty is dominant. A combined value for $\sigma_{\epsilon_{b\bar{b}Z^0}}$ of 15 % means the final $\epsilon_{b\bar{b}Z^0}$ with the single b-jet strategy is 0.0275 ± 0.0002 (stat) $\pm$ 0.004 (sys) and 0.0470 ± 0.0003 (stat) $\pm$ 0.007 (sys) for the $Z^0 \rightarrow e^+e^-$ and $Z^0 \rightarrow \mu^+\mu^-$ branches respectively.

		selection uncertainty (%)				
source	scale (%)	bbZ^0	Z^0+1-jet	Z^0+2-jets	Z^0+3-jets	$t\bar{t}$
lepton energy	2(0)	0(0)	0(0)	0(0)	0(0)	0(0)
jet energy	7	8(8)	15(13)	7(7)	2(2)	0(0)
E_T^{miss}	10	5(4)	0(0)	3(3)	5(3)	18(16)
lepton id	2	2⊕2	2⊕2	2⊕2	2⊕2	2⊕2
b-jet tag	8	8⊕8	-	-	-	8⊕8
b-jet mis-tag	8	-	8	8⊕8	8⊕8⊕8	-
σ	Table 6.2	-	0(0)	0(0)	-	0(0)
	combined	15(15)	17(16)	14(14)	15(15)	21(20)

Table 6.7: Summary of the systematic uncertainties introduced into the final signal and background selection efficiencies of the electron (muon) streams. Uncertainties in the energy scale and b-tagging/lepton id were found in [30]. The components are added in quadrature to produce the final combination.

6.8 Background estimation

The final term in Equation 6.8 required for a cross section measurement is $\sum_B n_B(t)$. This describes the number of background events passing event selection. The purpose of the selection strategy developed was to minimise this number, and thus

reduce the impact of this term on the overall cross section systematics. As discussed in Section 6.4, this number can be calculated using the background selection efficiencies measured in Section 6.7.2, the theoretical cross sections and the integrated luminosity through Equation 6.9. However, relying on calculated background cross sections is not ideal since it relies on a complete understanding of the underlying physical processes, so a method to measure the background contributions directly from data is necessary.

Figure 6.28 shows an example of how to do this for the $t\bar{t}$ background. The invariant dilepton mass peak is for 100 pb^{-1} of signal and background events. The events included pass all selection criteria *excluding* the mass window itself. Both muon and electron branches are considered. By fitting the data points with a Breit-Wigner + Landau curve and extracting the Landau component, the $t\bar{t}$ contribution can be estimated. This is highlighted by the hashed area in the figure. This follows a methodology originally proposed for 30 fb^{-1} of data by S. Lehti in [99]. In the example shown, a $t\bar{t}$ contribution of 25 $\pm$ 8 events is projected by the fit, compared with the 29 events that are actually present. This strategy therefore also seems to apply for relatively small amounts of data. To estimate the Z^0*+1,2,3*-jets contribution, a similar approach can be taken.

The region highlighted in blue on Figure 6.28 corresponds to the Z^0*+1,2,3*-jets background *plus* the signal. The Breit-Wigner component of the fit can be used to measure this area as 470 $\pm$ 30 events. For comparison, 489 events were actually counted from the Monte Carlo. More information is required, however, to decouple the contributions from signal and background.

Figure 6.29 shows the same invariant mass peak, but with the b-tagging requirement withdrawn. The number of Z^0*+1,2,3*-jets and $b\bar{b}Z^0$ events here are estimated to be 4310 $\pm$ 80. This can be compared with 4501 events actually present. The difference is mainly due to the low mass edge of the peak which contains an excess of events. The event content of the regions in blue for Figures 6.28 and 6.29, N_{final} and N_{nobtag}, are described by Equations 6.17 and 6.18 respectively.

$$N_{final} = N_{b\bar{b}Z^0(final)} + N_{Z^0+1,2,3jet(final)} \tag{6.17}$$

$$N_{nobtag} = \frac{N_{b\bar{b}Z^0(final)}}{\epsilon_{bjet}} + \frac{N_{Z^0+1,2,3jet(final)}}{\epsilon_{nonbjet}} \tag{6.18}$$

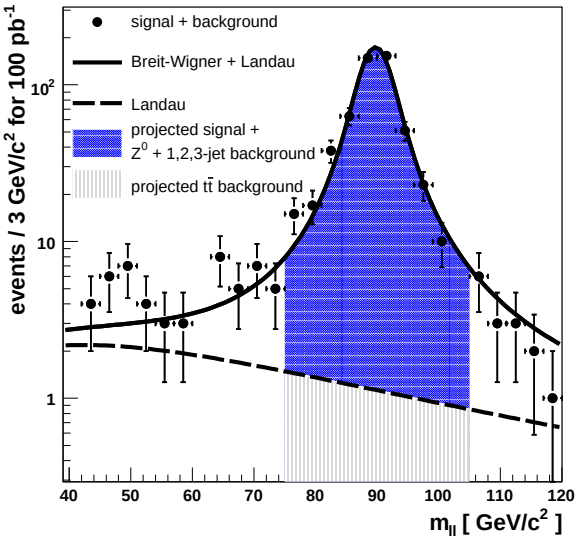

Figure 6.28: Invariant dilepton mass distribution for 100 pb^{-1} of signal and background events. Both $Z^0 \rightarrow e^+e^-$ and $Z^0 \rightarrow \mu^+\mu^-$ branches are considered. Only events passing the full selection procedure, with the exception of the invariant dilepton mass window, are chosen. Signal + background points are fitted with the Breit-Wigner + Landau function. The Landau component, corresponding to the $t\bar{t}$ background, is also drawn separately. The projected $t\bar{t}$ and $b\bar{b}Z^0$ + Z^0+*1,2,3*-jets contributions that sit within the selection mass window of 76 to 105 GeV/c^2 are highlighted. For this example, the Landau fit predicts 25 $\pm$ 5 events in the hashed region. 29 $t\bar{t}$ events are actually present. The Breit-Wigner fit predicts 470 $\pm$ 30 events in the blue region. 373 $b\bar{b}Z^0$ and 116 Z^0+N-jets events are actually present, corresponding to 489 in total. The errors in these values are estimated from uncertainties in the fit parameters.

Here, ϵ_{bjet} and $\epsilon_{nonbjet}$ are the measured b-jet and non-b-jet tagging efficiencies for jets with $E_T > 50$ GeV and $\mid \eta \mid < 2.4$. They can be measured from data using $t\bar{t}$ events [108]. These equations can be combined into an expression for the Z^0+*1,2,3*-jets background, shown by Equation 6.19.

$$N_{Z^0+1,2,3jet(final)} = \frac{\epsilon_{nonbjet}\epsilon_{bjet}N_{nobtag} - \epsilon_{nonbjet}N_{final}}{\epsilon_{bjet} - \epsilon_{nonbjet}} \tag{6.19}$$

The uncertainty in the $N_{Z^0+1,2,3jet(final)}$ estimate will come from the uncertainty in contributions to Equation 6.19, namely, $\sigma_{\epsilon_{nonbjet}}$, $\sigma_{\epsilon_{bjet}}$, $\sigma_{N_{nobtag}}$ and $\sigma_{N_{final}}$.

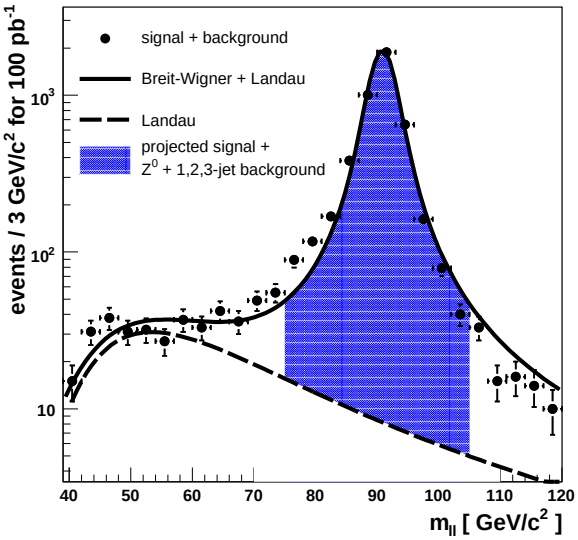

Figure 6.29: Invariant dilepton mass distribution for 100 pb^{-1} of signal and background events. Both $Z^0 \rightarrow e^+e^-$ and $Z^0 \rightarrow \mu^+\mu^-$ branches are considered. Only events passing the full selection procedure, with the exception of the invariant dilepton mass window and b-tagging, are chosen. Signal + background points are fitted with the Breit-Wigner + Landau function. The Landau component, corresponding to the $t\bar{t}$ background, is also drawn separately. The projected $b\bar{b}Z^0$ +Z^0+*1,2,3*-jets contribution that sits within the selection mass window of 76 to 105 GeV/c^2 is highlighted. For this example, the Breit-Wigner fit predicts 4310 ± 80 events. 4501 Z^0+*1,2,3*-jets and signal events are actually present. The errors in these values are estimated from uncertainties in the fit parameters.

6.9 Contributions from the $c\bar{c}Z^0$ background

Until this stage, the $c\bar{c}Z^0$ background discussed in Section 6.2 has not been considered. This is primarily due to the small number of available events. This section aims to establish the impact this background will have on a $b\bar{b}Z^0$ cross section measurement.

A 5280 event $c\bar{c}Z^0$ + 0jets sample was produced with ALPGEN using the thresholds $p_T^c > 20$ GeV, $\mid \eta^c \mid < 5$ and $m_Z > 40$ GeV/c^2. Only the $Z^0 \rightarrow e^+e^-$ branch was generated. The same PDF and scale settings used for $b\bar{b}Z^0$ production (Table 6.2) were applied. A NLO cross section of 11 pb was calculated using MCFM with PDF CTEQ6L1.

Firstly, the effect of the b-tagging algorithm on c-jets should be established. This measurement was made by counting the number of jets matched with a c-quark (as for b-jets, a matching cone of 0.4 was used) that were positively identified as b-jets.

Figure 6.30 shows how this efficiency, ϵ_{cquark}, varies with the chosen discriminator. The value of 2.0 used in previous event selection measurements is also highlighted. With this threshold a c-tagging efficiency of 0.046 ± 0.006 was measured. It is also clearly near its lowest value, so it can be concluded that the chosen discriminator is suitable for the $c\bar{c}Z^0$ background.

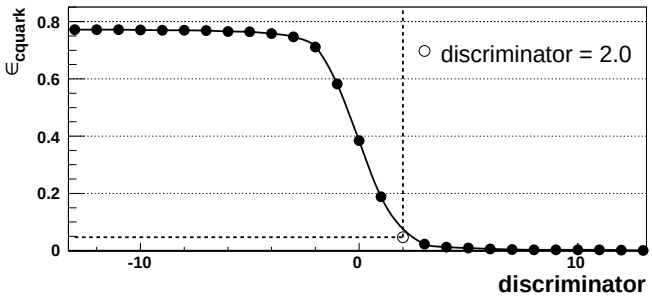

Figure 6.30: c-jet identification efficiency as a function of the chosen b-tagging discriminator (using the track counting approach). The value chosen in Section 6.6.4 for this analysis (2.0) is shown by the arrow. The corresponding c-tagging efficiency is measured to be 0.046 ± 0.006.

The full $c\bar{c}Z^0$ Monte Carlo sample was then run through the event selection filter. The results are shown in Table 6.8. As with the other Z^0+N-jets backgrounds, the most effective suppressant is the b-tagging stage. It is also clear, through comparison with Table 6.5 that the contribution is comparable with the Z^0+2-jets background.

Although not explicitly shown here, it is anticipated that the same background estimation technique demonstrated in Section 6.8 can be applied to $c\bar{c}Z^0$. However, the added terms in Equations 6.17 and 6.18 mean these two equations now have three unknowns: $N_{b\bar{b}Z^0(final)}$, $N_{Z^0+1,2,3jet(final)}$ and $N_{c\bar{c}Z^0(final)}$. Following the D0 approach [97], the theoretical ratio shown by Equation 6.20 can be used to reduce the number of unknown parameters to two.

$$R = \frac{\sigma(c\bar{c}Z(final))}{\sigma(Z^0 + 1,2,3jet(final))} \times \frac{\epsilon_{c\bar{c}Z}}{\epsilon_{Z^0+1,2,3jet(final)}} \tag{6.20}$$

The uncertainty in the $\frac{\sigma(c\bar{c}Z(final))}{\sigma(Z^0+1,2,3jet(final))}$ ratio is ± 3 % due to variations with the chosen factorisation and renormalisation scale [100]. The PDF uncertainty is not yet taken into account. The systematic uncertainties in the $c\bar{c}Z^0$ selection efficiency measurement, $\epsilon_{c\bar{c}Z}$ can be estimated using a similar approach to that adopted in Section 6.7.3. Finally, since the charm tagging efficiency cannot easily be measured from data the measurement from Monte Carlo should be used. In fact the ratio of b-jet to c-jet tagging efficiencies will be recorded and then scaled to reproduce the b-tagging efficiency measured with data.

selection	$c\bar{c}Z^0$ selection efficiency (fraction from previous step)
isolated dielectron trigger	0.467 (1.000)
2 × electrons $p_T > 15$, $\mid \eta \mid < 2.4$	0.422 (0.904)
opposite sign electrons	0.410 (0.972)
$75 < m_{ee} < 105$	0.392 (0.956)
$E_T^{miss} < 60$	0.388 (0.988)
1 jet $p_T > 50$, $\mid \eta \mid < 2.4$	0.131 (0.338)
1 b-jet $p_T > 50$, $\mid \eta \mid < 2.4$	0.012 (0.091)
$\geq$ 1 jet p_T >50, $\mid \eta \mid < 2.4$	0.164 (0.423)
$\geq$ 1 b-jet $p_T > 50$, $\mid \eta \mid < 2.4$	0.019 (0.113)
$\geq$ 2 jet p_T >50, $\mid \eta \mid < 2.4$	0.033 (0.086)
$\geq$ 2 b-jets $p_T > 50$, $\mid \eta \mid < 2.4$	0.001 (0.023)
	events after 100 pb^{-1}
	13 ± 4
	20 ± 5
	0.8 ± 0.9

Table 6.8: Cumulative selection efficiencies for the $c\bar{c}Z^0$ background after successive selection cuts. Only the $Z^0 \rightarrow e^+e^-$ decay branch is considered. The number of events expected after 100 pb^{-1} integrated luminosity with corresponding binomial errors are also shown. The various event selection criteria are introduced from top to bottom. Three selection streams exist shown in grey. Fields in white are common to both streams.

6.10 Summary

$b\bar{b}Z^0$ production at the LHC is of great interest, primarily as it is a benchmark process and background to the MSSM Higgs production $b\bar{b}h/H/A$. The measurement of the $b\bar{b}Z^0$ cross section with 100 pb^{-1} of data has been the focus of this chapter. Two competing calculation schemes have been developed for this channel, both being performed to NLO. The *five flavour* approach, which is simplest, and most readily extendable to NNLO, has been targeted here. This particular calculation assumes at least one high p_T b-quark (> 15 GeV) which is satisfied by the proposed signal selection model.

In this chapter, the $Z^0 \rightarrow e^+e^-$ and $Z^0 \rightarrow \mu^+\mu^-$ decay modes have been studied. Other decay modes involving $\tau^{\pm}$, neutrinos or hadronic decays are harder to identify and so have a larger background as a result. With this in mind, the isolated dielectron (2 × 12 GeV) and relaxed dimuon (2 × 3 GeV) trigger streams will be used.

The primary backgrounds to this process have been identified as $t\bar{t}$, Z^0+1-jet, Z^0+2-jets and Z^0+3-jets. The offline selection strategy chosen to suppress these requires two leptons (electrons or muons) of opposite charge with p_T> 15 GeV; a corresponding invariant dilepton mass between 76 and 105 GeV/c^2; the b-tagging of a single, central jet with p_T> 50 GeV and an E_T^{miss} below 60 GeV. Thresholds were chosen to maximise background suppression. A method of electron and muon identification through tracker isolation was discussed and the discriminators tuned for signal events. The track counting discriminator used to identify b-jets was also optimised to achieve a high purity collection in signal samples.

A final signal selection efficiency of 0.0275 ± 0.0002 (stat) ± 0.004 (sys) and 0.0470 ± 0.0003 (stat) ± 0.007 (sys) was measured from Monte Carlo for the $Z^0 \rightarrow e^+e^-$ and $Z^0 \rightarrow \mu^+\mu^-$ branches respectively. The systematic errors were estimated from lepton, jet and E_T^{miss} energy scales along with uncertainties in lepton and b-jet identification and non-b-jet misidentification efficiencies.

Finally, to avoid dependencies on theoretical cross section calculations, a method to estimate the size of the $t\bar{t}$ and *Z^0+1,2,3*-jets backgrounds from data has been discussed. It was shown to apply even for 100 pb^{-1} of data.

The remaining background yet to be studied was $c\bar{c}Z^0$ production. Just over 5000 events were generated in the $Z^0 \rightarrow e^+e^-$ branch. The charm mistagging efficiency was measured to be 0.046 ± 0.006. To estimate the contribution with real data a similar approach to the *Z^0+1,2,3*-jets backgrounds can be taken, but using a theoretical $\frac{\sigma(c\bar{c}Z(final))}{\sigma(Z^0+1,2,3jet(final))}$ ratio and a charm mistagging efficiency measured from Monte Carlo rather than data.

Chapter 7

Conclusions

The CMS experiment is a general purpose LHC detector that has been over 15 years in the making. It has been designed and optimised to discover the Higgs boson and signatures of new physics beyond the Standard Model. An early discovery of the Higgs boson is the collaboration's top priority and will require a good understanding of both the detector and the physics of the background processes, with a small integrated luminosity. This principle has been the driving force behind the work presented in this thesis.

Chapter 4 described the development of commissioning procedures for the silicon strip tracker, which sits at the heart of CMS. Their purpose is to configure, synchronise and calibrate its control and readout systems in preparation for physics data taking. The CMS SST consists of almost 10 M channels so the commissioning process must be fully automated for speed and consistency. The software structure was briefly discussed with a focus on the core algorithms that calculate the optimal calibration values. The process of *live* commissioning at the Tracker Integration Facility was then described for the remainder of the chapter, illustrated with the commissioning *summary* distributions which are useful for diagnosing hardware problems. During this time, 1.6 M channels were calibrated and synchronised to a cosmic muon trigger to within 1 ns.

This thesis has demonstrated the success of the commissioning software in preparing a relatively small system at the TIF for physics data taking. Since the author's involvement in this analysis, the full tracker has been installed and commissioned at Point 5 on the LHC ring with great success.

The SST is expected to produce five times more zero-suppressed data than any other CMS subdetector, necessitating its efficient handling within the HLT algorithms. Chapter 5 profiled the performance of the online hit reconstruction algorithms, characterised the inefficiencies and proposed a new schema to focus on physics regions-of-interest only. Four possible implementations of this schema were discussed. The *on-demand* approach was the final choice adopted by the tracker community as its online solution. As an example of its success, when running the single τ trigger path over $H^{\pm} \rightarrow \tau^{\pm}\nu_{\tau}$ events, hit reconstruction times were reduced from 838 $\pm$ 5 ms to only 5.13 $\pm$ 0.05 ms without any loss in tracking efficiency. It is expected that the new software design will also become the offline solution in the near future.

The dramatic improvement in hit reconstruction performance demonstrated in this chapter will have a direct impact on the physics aspirations of CMS. Without the streamlined handling of such a large amount of raw data in the HLT algorithms, the tracking information could not have been considered online. This would affect any channel relying on a track-based trigger, including the isolated dielectron and relaxed dimuon streams used when studying $b\bar{b}Z^0$ production. The times quoted in this thesis were for low luminosity conditions or (as in the example quoted above) without pileup entirely. It has been demonstrated that the hit reconstruction times increase linearly with strip occupancy so further developments may be necessary for high luminosity running.

Finally, Chapter 6 prepared for a 100 pb^{-1} $b\bar{b}Z^0$ production cross section measurement. This integrated luminosity is expected within the first year of LHC running so the analysis is heavily reliant on a well commissioned detector. The channel is interesting primarily due to its status as a background to a supersymmetric Higgs boson production process. Two competing matrix element calculation schemes have been developed. The four-flavour scheme, which assumes two high p_T b-quarks in the final state, has been studied elsewhere. This thesis has focused on probing the five-flavour scheme. With this approach only a single high p_T b-quark is required. The primary backgrounds to this process were identified as $t\bar{t}$ and Z^0+*1,2,3*-jets. The signal selection strategy was developed with background suppression in mind and after optimisation achieved an efficiency of $\epsilon_{b\bar{b}Z^0} = 0.0275 \pm 0.0002$ (stat) $\pm$ 0.004 (sys) and $\epsilon_{b\bar{b}Z^0} = 0.0470 \pm 0.0003$ (stat) $\pm$ 0.007 (sys) for the $Z^0 \rightarrow e^+e^-$ and $Z^0 \rightarrow \mu^+\mu^-$ modes respectively.

A background estimation procedure, originally proposed by S. Lehti for 30 pb^{-1} of data, was then tested over a 100 pb^{-1} sample. It performed well, with the estimations matching the actual background within errors with the exception of an underestimation of the number of Z^0+*1,2,3*-jets events when not using b-tagging. Finally, a $c\bar{c}Z^0$ background was considered using a small sample of generated events. A similar background estimation procedure to that for Z^0+*1,2,3*-jets is envisaged.

Before a cross section measurement can be made with real data, the $c\bar{c}Z^0$ background should be studied more thoroughly. The detector misalignment, a significant factor for an early measurement, should also be simulated. Despite these simplifications, the results presented in this thesis and the algorithms developed in the process will allow the community to make an early and competitive measurement.

Glossary

ADC Analogue to Digital Converter
ALICE A Large Ion Collider Experiment
ALPGEN A Monte Carlo event generator for hadronic collisions
AOH Analogue OptoHybrid
APSP Analogue Pulse Shape Processor
APV25 Analogue Pipeline Voltage 25
APVMUX Analogue Pipeline Voltage MUltipleXer
ASIC Application Specific Integrated Circuit
ATLAS A Toroidal Lhc ApparatuS
BU Builder Unit
CCU Communications and Control Unit
CDF Collider Detector at Fermilab
CMOS Complimentary Metal Oxide Semiconductor
CMS Compact Muon Solenoid
CMSSW CMS SoftWare
COBRA Coherent Object-oriented Base for Reconstruction and Analysis
CompHEP For the calculation of multi-particle final states in tree level processes
CPU Central Processor Unit
CSC Cathode Strip Chambers
D0 Detector at Fermilab (at point D0)
DAQ Data AcQuisition
DCU Detector Control Unit
DOH Digital OptoHybrid
DQM Data Quality Monitoring
DT Drift Tubes
ECAL Electromagnetic CALorimeter
EDM Event Data Model

FE Front End
FEC Front End Controller
FED Front End Driver
FEH Front End Hybrid
FPGA Field Programmable Gate Array
GCT Global Calorimeter Trigger
GSF Gaussian Sum Filter
HCAL Hadronic CALorimeter
HLT High Level Trigger
HV High Voltage
I^{2}C Inter-IC
ISHA APV25 shaper input current bias
L(1,2,2.5,3) Level (1,2,2.5,3)
LEP(2) Large Electron Positron collider (run 2)
LHC Large Hadron Collider
LHCb LHC beauty experiment
LINAC LINeAr Collider
LLD Linear Laser Driver
LO Leading Order
MC Monte Carlo
MCFM Monte Carlo for FeMtobarn processes
ME Matrix Element
MIP Minimally Ionising Particle
MSSM Minimal Supersymmetric Model
NLO Next-to-Leading Order
NNLO Next-to-Next-to-Leading Order
ORCA Object-oriented Reconstruction for Cms Analysis
PDF Parton Distribution Function
PLL Phase-Locked Loop
PS Proton Synchrotron
RBN Readout Builder Network
RCT Regional Calorimeter Trigger
RHIC Relativistic Heavy Ion Collider
RPC Resistive Plate Chambers
RU Readout Unit
S-LINK64 A data link used to connect a CMS FED to the event builder
SM Standard Model

SPS Super Proton Synchrotron
SST Silicon Strip Tracker
SUSY SUperSymmetrY
TEC Tracker End Caps
TIB Tracker Inner Barrel
TID Tracker Inner Discs
TIF Tracker Integration Facility
TOB Tracker Outer Barrel
TOTEM TOTal Elastic and diffractive cross section Measurement
TrimDAC An adjustable offset available on each FED ADC
TTC Timing Trigger and Command
TTCrx A TTC receiver ASIC for LHC detectors
TTS Trigger Throttling System
QCD Quantum ChromoDynamics
QED Quantum ElectroDynamics
QFT Quantum Field Theory
XDAQ Cross-platform DAQ
XML eXtensible Markup Language
VFS APV25 shaper feedback voltage bias
VME Versa Module Eurocard
VPSP APSP voltage level adjust

References

[1] H. Lehmann *et al.*, "On the formulation of quantized field theories," *Nuovo Cimento (Serie 10)* **1** (1955) 205.

[2] L. H. Ryder, "Quantum field theory," *Cambridge University Press* (1996).

[3] M. E. Peskin and D. V. Schroeder, "An introduction to quantum field theory," *Addison-Wesley* (1995).

[4] R. P. Feynman, "Quantum electrodynamics," *Westview Press* (1998).

[5] S. L. Glashow, "Partial-symmetries of weak interactions," *Nuclear Physics* **22** (1961) 579.

[6] S. Weinberg, "A model of leptons," *Physical Review Letters* **19** (1967) 1264.

[7] A. Salam, "In elementary particle theory," *Almqvist and Wiksells, Stockholm* (1968) 367.

[8] P. W. Higgs, "Spontaneous symmetry breakdown without massless bosons," *Physical Review* **145** (1966) 1165.

[9] C. P. Burgess, "A Goldstone boson primer," *arXiv:hep-ph/9812468* (1998).

[10] B. W. Lee *et al.*, "Weak interaction at very high energies: The role of the Higgs boson mass," *Physical Review* **D16** (1977) 1519.

[11] The LEP working group for Higgs boson searches, "Search for the standard model Higgs boson at LEP," *Physics Letters* **B565** (2003) 61.

[12] The LEP electroweak working group, "Precision electroweak measurements and constraints on the standard model," *arXiv:hep-ex/07120929* (2007).

[13] Tevatron electroweak working group, CDF and D0, "Combination of CDF and D0 results on the top-quark mass," *arXiv:hep-ex/0703034* (2007).

[14] S. P. Martin, "A supersymmetry primer," *arXiv:hep-ph/9709356* (1999).

[15] I. J. R. Aitchison and A. J. G. Hey, "Gauge theories in particle physics," *Institute of Physics Publishing* (1989).

[16] D. Griffiths, "Introduction to elementary particles," *John Wiley and Sons* (1987).

[17] F. Halzen and A. D. Martin, "Quarks and leptons," *John Wiley and Sons* (1984).

[18] I. J. R. Aitchison, "Supersymmetry and the MSSM: An elementary introduction," *arXiv:hep-ph/0505105* (2005).

[19] M. Drees, "Introduction to supersymmetry," *arXiv:hep-ph/9611409* (1996).

[20] LHC collaboration, "LHC Design Report," *CERN/LHCC* **03** (2004).

[21] L. Evans and P. Bryant, "LHC machine," *Journal of Instrumentation* **3** (2008).

[22] CMS collaboration, "The Compact Muon Solenoid Technical Proposal," *CERN/LHCC* **38** (1994).

[23] CMS collaboration, "The CMS experiment at the CERN LHC," *Journal of Instrumentation* **3** (2008).

[24] ATLAS collaboration, "The ATLAS experiment at the CERN Large Hadron Collider," *Journal of Instrumentation* **3** (2008).

[25] LHCb collaboration, "The LHCb detector at LHC," *Journal of Instrumentation* **3** (2008).

[26] ALICE collaboration, "The ALICE experiment at the CERN LHC," *Journal of Instrumentation* **3** (2008).

[27] M. Spira *Fortsch. Phys.* **46** (1998) 203–284.

[28] TOTEM collaboration, "TOTEM Technical Design Report," *CERN/LHCC* **002** (2004).

[29] E. V. Shuryak, "Strongly coupled quark-gluon plasma: The status report," *arXiv:hep-ph/0608177* (2006).

[30] CMS collaboration, "The CMS physics Technical Design Report, volume 1," *CERN/LHCC* **001** (2006).

[31] CMS collaboration, "The tracker project Technical Design Report," *CERN/LHCC* **06** (1998).

[32] CMS collaboration, "Addendum to the CMS tracker TDR," *CERN/LHCC* **016** (2000).

[33] K. Arndt *et al.*, "Silicon sensors development for the CMS pixel system," *Nuclear Instruments and Methods* **A511** (2003) 106.

[34] U. Goerlach, "Industrial production of front-end hybrids for the CMS silicon tracker," *Proceedings of the 9th workshop on electronics for LHC experiments* (2003).

[35] L. Borrello *et al.*, "Sensor design for the CMS silicon strip tracker," *CERN/LHCC* **020** (2003).

[36] C. Goßling, "Pion-induced damage in silicon detectors," *Nuclear Instruments and Methods* **A379** (1996) 116–123.

[37] S. Braibant *et al.*, "Investigation of design parameters for radiation hard silicon microstrip detectors," *Nuclear Instruments and Methods* **A485** (2002) 343.

[38] E. Noah *et al.*, "Total dose irradiation of a 0.25 μm process," *CERN/LHCC* **041** (2000).

[39] CMS collaboration, "The electromagnetic calorimeter project, Technical Design Report," *CERN/LHCC* **33** (1997).

[40] CMS collaboration, "The hadron calorimeter project, Technical Design Report," *CERN/LHCC* **31** (1997).

[41] V. D. Elvira, "Measurement of the pion energy response and resolution in the hadronic barrel calorimeter using CMS HCAL test beam 2002 data," *CERN/LHCC* **020** (2004).

[42] CMS collaboration, "The magnet project: Technical Design Report," *CERN/LHCC* **10** (1997).

[43] CMS collaboration, "The muon project, Technical Design Report," *CERN/LHCC* **32** (1997).

[44] K. Gill *et al.*, "Progress on the CMS tracker control system," *Proceedings of the 11th workshop on electronics for LHC experiments* (2005).

[45] Philips semiconductors, "The I^2C-bus specification 2.1." Document order number 9398 393 40011 (2000).

[46] S. Baron, "Timing, trigger and control for the LHC." http://ttc.web.cern.ch/TTC/intro.html.

[47] A. Marchioro *et al.*, "A PLL-delay ASIC for clock recovery and trigger distribution in the CMS tracker," *CERN/LHCC* **60** (1997).

[48] M. French *et al.*, "Design and results from the APV25, a deep sub-micron CMOS front-end chip for the CMS tracker," *Nuclear Instruments and Methods* **A** (2001) 359–365.

[49] M. Raymond *et al.*, "The CMS tracker APV25 0.25 μm CMOS readout chip," *Proceedings of the 6th workshop on electronics for the LHC experiments* (2000).

[50] N. Bingefors *et al.*, "Novel technique for fast pulse-shaping using a slow amplifier at LHC," *Nuclear Instruments and Methods* **A** (1993) 112.

[51] W. Gannon, "Compact Muon Solenoid Front End Driver front end FPGA Technical Description." Private note (2003).

[52] H. van der Bij, "The S-LINK interface specification." http://hsi.web.cern.ch/HSI/s-link/spec/spec/s-link.pdf.

[53] S. Taghavi, "Compact Muon Solenoid Front End Driver back end FPGA Technical Description." Private note (2005).

[54] CMS collaboration, "The TriDAS project - the level 1 trigger Technical Design Report," *CERN/LHCC* **38** (2000).

[55] CMS collaboration, "The TriDAS project Technical Design Report, volume 2: data acquisition and the high-level trigger," *CERN/LHCC* **26** (2002).

[56] CMS collaboration, "CMS High Level Trigger," *CERN/LHCC* **021** (2007).

[57] W. Adam *et al.*, "The CMS High Level Trigger," *European Physical Journal* **C46** (2005) 605.

[58] T. Sjostrand *et al.*, "PYTHIA 6.4. physics and manual," *arXiv:hep-ph/0603175* (2006).

[59] A. Pukhov *et al.*, "CompHEP - a package for evaluation of Feynman diagrams and integration over multi-particle phase space. users manual for version 33," *arXiv:hep-ph/9908288* (1999).

[60] W. Adam *et al.*, "ORCA: CMS reconstruction package." http://cmsdoc.cern.ch/orca/.

[61] K. Lassila-Perini *et al.*, "CMSSW application framework." https://twiki.cern.ch/twiki/bin/view/CMS/WorkBookCMSSWFramework.

[62] R. Brun and F. Rademakers, "ROOT: An object oriented data analysis framework." http://root.cern.ch.

[63] R. Fruhwirth, "Application of Kalman filtering to track and vertex fitting," *Nuclear Instruments and Methods* **A** (1987) 444–450.

[64] W. Adam *et al.*, "Track reconstruction in the CMS tracker," *CMS NOTE* **041** (2006).

[65] R. Fruhwirth *et al.*, "Vertex reconstruction and track bundling at the LEP collider using robust algorithms," *Nuclear Instruments and Methods* **A** (2005) 463–482.

[66] T. Speer *et al.*, "Track reconstruction in the CMS tracker," *CMS CR* **041** (2005).

[67] E. Meschi *et al.*, "Electron reconstruction in the CMS electromagnetic calorimeter," *CMS NOTE* **034** (2001).

[68] S. Baffioni *et al.*, "Electron reconstruction in CMS," *CMS NOTE* **40** (2006).

[69] M. Dittmar and A. S. Giolo-Nicollerat, "Identifying central electrons in the CMS detector," *CMS NOTE* **036** (2004).

[70] G. Bruno *et al.*, "Local reconstruction in the muon detectors," *CMS NOTE* **043** (2002).

[71] E. James *et al.*, "Muon identification in CMS," *CMS NOTE* **010** (2006).

[72] O. Kodolova, "Jet energy measurements in CMS," *CMS CR* **019** (2005).

[73] S. Chekanov *et al.*, "Jet algorithms: a mini review," *arXiv:hep-ph/0211298* (2002).

[74] A. Heister *et al.*, "Measurement of jets with the CMS detector at the LHC," *CMS NOTE* **036**.

[75] A. Rizzi *et al.*, "Track impact parameter based b-tagging with CMS," *CMS NOTE* **019** (2006).

[76] G. Bagliesi *et al.*, "Tau jet reconstruction and tagging at high level trigger and off-line," *CMS NOTE* **028** (2006).

[77] H. Haifeng *et al.*, "Measurement of missing transverse energy with the CMS detector at the LHC," *"CMS NOTE"* **035** (2006).

[78] L. Orsini and J. Gutleber, "XDAQ." http://xdaqwiki.cern.ch/index.php/Main_Page.

[79] K. Gill and N. Marinelli, "Start-up synchronization of the CMS tracker." http://www.te.rl.ac.uk/esdg/cms-fed/urd/synchronization_procedure.pdf.

[80] S. Dris, "Performance of the CMS tracker optical links and future upgrade using bandwidth efficient digital modulation." Ph.D. thesis of Imperial College London (2006).

[81] M. Raymond and M. Wingham, "I^2C settings for APV operation at cold temperatures - measurements on a TOB module." http://www.hep.ph.ic.ac.uk/ dmray/Tracker.html (2005).

[82] R. Bainbridge, "Influence of highly ionising events on the CMS APV25 readout chip." Ph.D. thesis of Imperial College London (2004).

[83] C. Delaere *et al.*, "CMS tracker commissioning and first operation experience," *16th International Workshop on Vertex Detectors* (2007).

[84] R. Bainbridge *et al.*, "Data acquisition software for the CMS strip tracker," *Proceedings from CHEP, British Columbia* (2007).

[85] J. Coughlan, "The Front-End Driver card for the CMS silicon strip tracker readout," *CERN/LHCC* **034** (2002).

[86] O. Kodolova *et al.*, "Expected data rates from the silicon strip tracker," *CMS NOTE* **047** (2002).

[87] C. Jones, "regional tracking proposal." https://twiki.cern.ch/twiki/bin/view/CMS/RegionalTrackingProposal.

[88] CDF collaboration, "Observation of top quark production in $\bar{p} - p$ collisions," *Physical Review Letters* **74** (1995) 2626.

[89] D0 collaboration, "Observation of the top quark," *Physical Review Letters* **74** (1995) 2632.

[90] M. Carena *et al.*, "MSSM Higgs boson phenomenology at the Tevatron collider," *arXiv:hep-ph/9808312* (1998).

[91] D0 collaboration, "Search for neutral supersymmetric Higgs bosons in multijet events at $\sqrt{s} = 1.96$ TeV," *arXiv:hep-ex/0504018* (2005).

[92] S. Lehti, "Study of H/A $\rightarrow \tau\tau \rightarrow e\mu$ + X in CMS," *CMS NOTE* **101** (2006).

[93] J. Campbell *et al.*, "Production of a Z boson and two jets with one heavy-quark tag," *arXiv:hep-ph/0510362* (2008).

[94] F. Maltoni *et al.*, "Inclusive production of a Higgs or Z boson in association with heavy quarks," *Physical Review* **D72** (2005) 074024.

[95] J. Campbell *et al.*, "Associated production of a Z boson and a single heavy-quark jet," *Physical Review* **D69** (2004) 074021.

[96] J. Campbell and R. Ellis, "Next-to-leading order corrections to W+2 jet and Z+2 jet production at hadron colliders," *Physical Review* **D65** (2002) 113007.

[97] D0 collaboration, "Measurement of the ratio of inclusive cross sections $\sigma(p\bar{p} \rightarrow Z + b\ jet)/\sigma(p\bar{p} \rightarrow Z + jet)$ at $\sqrt{s} = 1.96$ TeV," *Physical Review Letters* **94** (2005) 161801.

[98] CDF collaboration, "Measurement of the b jet cross section in events with a Z boson in $p\bar{p}$ collisions at at $\sqrt{s} = 1.96$ TeV," *Physical Review* **D74** (2006) 032008.

[99] S. Lehti, "Observability of bbZ events at CMS as a benchmark for MSSM bbH search," *CMS NOTE* **099** (2006).

[100] A. Nayak, "Measurement of Zb$\bar{\text{b}}$, Z→ll cross section with 100 pb^{-1} of early LHC data using CMS detector," *CMS AN* **020** (2008).

[101] M. Pierini *et al.*, "ALPGEN production in summer 07." https://twiki.cern.ch/twiki/bin/view/Main/AlpgenSummer07.

[102] N. Adam *et al.*, "Towards a measurement of the inclusive W → eν and Z → ee cross section in pp collisions at sqrt(s) = 14 TeV," *CMS AN* **026** (2007).

[103] N. Adam *et al.*, "Towards a measurement of the inclusive W → $\mu\nu$ and Z → $\mu^+\mu^-$ cross sections in pp collisions at $\sqrt{s} = 14$ TeV," *CMS AN* **031** (2007).

[104] A. Nikitenko. Private communication.

[105] M. L. Mangano *et al.*, "ALPGEN, a generator for hard multiparton processes in hadronic collisions," *arXiv:hepph/0206293* (2003).

[106] J. Campbell *et al.*, "Next-to-leading order QCD predictions for W + 2 jet and Z + 2 jet production at the CERN LHC," *arXiv:hep-ph/0308195* (2003).

[107] C. Bishop, "Neural networks for pattern recognition," *Oxford University Press* (1995).

[108] S. Lowette *et al.*, "Offline calibration of b-jet identification efficiencies," *CMS AN* **030** (2005).