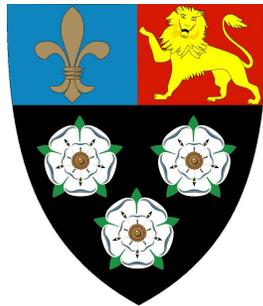


New Physics Searches in the ZZ Sector with the ATLAS Experiment

Thomas James Barber
of King's College



A dissertation submitted to the University of Cambridge
for the degree of Doctor of Philosophy
February 2011

Abstract

This thesis investigates the prospects of measuring anomalous triple gauge boson couplings in the ATLAS detector at the Large Hadron Collider (LHC). The most general ($VZZ, V = Z, \gamma$) vertex is parametrised by four couplings, $f_{i=4,5}^{V=Z,\gamma}$, all of which are zero in the Standard Model. Non-zero couplings would manifest themselves as an excess of events in ZZ diboson channels, and, if observed, would be a direct probe of new physics beyond the Standard Model.

A set of criteria are outlined to select events recorded by ATLAS in two such channels, $ZZ \rightarrow lll (l = e, \mu)$ and $ZZ \rightarrow ll\nu\bar{\nu}$. With 1 fb^{-1} of integrated luminosity at a centre of mass energy of $\sqrt{s} = 7 \text{ TeV}$, ATLAS can expect to observe 10 ± 1 events in the $ZZ \rightarrow lll$ channel, with $0.5_{-0.2}^{+0.9}$ background events. In the $ZZ \rightarrow ll\nu\bar{\nu}$ channel, 6.2 ± 0.7 signal events are expected, with a background of $1.9_{-0.2}^{+2.0}$ events.

The expected sensitivity of ATLAS to non-zero anomalous couplings is calculated by performing extended, unbinned maximum-likelihood fits to the Z boson transverse momentum spectrum. For 1 fb^{-1} of integrated luminosity at $\sqrt{s} = 7 \text{ TeV}$, ATLAS has the potential to place constraints on the coupling parameters of $|f_i^Z| < 0.06$ and $|f_i^\gamma| < 0.07$ at the 95% confidence level. These limits assume a form factor with a cutoff of $\Lambda_{FF} = 1.2 \text{ TeV}$.

As a prelude to ZZ observation, criteria are defined to select $Z \rightarrow ll (l = e, \mu)$ events in the first 315 nb^{-1} of ATLAS pp collision data at $\sqrt{s} = 7 \text{ TeV}$. In total 57 events are observed in the electron channel, with 109 in the muon channel, leading to cross section measurements of $\sigma(Z \rightarrow ee) = 0.70 \pm 0.09 \text{ (stat)} \pm 0.10 \text{ (syst)} \pm 0.08 \text{ (lumi)} \text{ pb}$ and $\sigma(Z \rightarrow \mu\mu) = 0.90 \pm 0.09 \text{ (stat)} \pm 0.07 \text{ (syst)} \pm 0.10 \text{ (lumi)} \text{ pb}$, both of which are consistent with the Standard Model predictions.

In addition, this thesis presents a summary of developments made to the Data Acquisition (DAQ) system of the ATLAS Semiconductor Tracker (SCT). These include the construction of a test system, involving a scaled-down version of the entire SCT read-out chain. The test system was subsequently used to develop a number of new DAQ features, including a hardware-based event simulator and monitoring framework.

Declaration

This dissertation is the result of my own work, except where explicit reference is made to the work of others, and has not been submitted for another qualification to this or any other university. This dissertation does not exceed the word limit for the respective Degree Committee.

Thomas Barber

Acknowledgements

I would first of all like to thank the Science and Technology Research Council (STFC) for their financial support under STFC student number (PPA/S/S/2006/04325), including a very enjoyable two years spent on Long Term Attachment (LTA) based in Geneva. I am also extremely grateful to King's College for travel grants which have allowed me to present my work and expand my cultural horizons at conferences in Israel and China.

I would also like to express my gratitude to my supervisor, Richard Batley, for his support over the last four years and speedy proof-reading in the last four weeks. Thanks also go to Pat Ward for productive and stimulating discussions on all things diboson. During the time I spent at CERN, I am greatly indebted to Dave Robinson for his patience explaining the SCT timing, and Bruce Gallop for his encyclopedic knowledge of the SCT ROD registers.

Finally, I would like to dedicate this thesis to my parents, Bridget and Graham, for their unending support over the last 27 years.

Preface

This thesis represents the culmination of work I have carried out at the Cavendish Laboratory, Cambridge between October 2006 and February 2011. My research is based on a number of aspects of the ATLAS experiment, one of the detectors designed to record proton-proton collisions in the Large Hadron Collider at The European Organisation for Nuclear Research (CERN) in Geneva.

At the start of my studies, the final components of the ATLAS detector were being installed in the Point-1 experimental cavern, 100 m beneath the Franco-Swiss border near Geneva. After extensive commissioning and calibration work, the Large Hadron Collider (LHC) is now accelerating and colliding protons at unprecedented energies of 7 TeV. The field of high energy particle physics has truly moved into a new era.

A vital aspect of the ATLAS detector, described in chapter 2, is the measurement of the position and momentum of charged particles produced in LHC collisions. During the two years I spend based at CERN, I was particularly involved with the ATLAS silicon strip detector, the Semiconductor Tracker (SCT). An introduction to the Data Acquisition (DAQ) system required to read out the 6 million channels of the SCT is given in chapter 3. Chapter 4 describes the construction of a scaled-down test system, which was used to develop many new features of the SCT DAQ. The features that I was particularly involved in are documented in chapter 5, and focus on developments made on the SCT Readout Driver (ROD). These include a hardware-based simulation of SCT events and the construction of a new monitoring framework.

One of the aims of the ATLAS experiment is to search for as-yet undiscovered physics. Our current understanding of the laws of the Universe is described by the Standard Model (SM), which states that each force is mediated by a set of particles known as gauge bosons. For example, as discussed in chapter 1, the photon, γ , is the force carrier for electromagnetism, while the W and Z bosons mediate the weak force. Interactions between three neutrally charged gauge bosons are forbidden in the Standard Model, and

are hence known as anomalous couplings. Discovery of non-zero anomalous couplings would be a direct indication of new physics.

Interactions in which a pair of Z bosons are produced, known as dibosons, are sensitive to such anomalous couplings. Chapter 6 outlines a set of criteria for selecting and reconstructing diboson events. Using detailed computer simulations of the ATLAS experiment, the expected number of events after 1 fb^{-1} of data (approximately one year) is calculated. Anomalous couplings have the effect of increasing the expected number of diboson events, and the momenta of the bosons perpendicular to the beam direction. In chapter 7, I show ATLAS should be able to rule out anomalous couplings with a sensitivity twice that of the current measurements.

The LHC began high energy collisions in March 2010, six months before the end of my doctorate. As the majority of this thesis uses simulated collisions, I couldn't resist the opportunity to get my hands dirty with some real data. Chapter 8 presents the first observation of single Z bosons in the ATLAS detector. I show that the measured cross section (production rate) is in good agreement with theoretical calculations.

Contents

1	Anomalous Triple Gauge Couplings	3
1.1	Introduction	3
1.2	The Electro-weak Interaction	3
1.2.1	The Weak Interaction	5
1.2.2	Electro-weak Unification	7
1.2.3	Spontaneous Symmetry Breaking	8
1.2.4	The Standard Model	9
1.3	Triple Gauge Couplings	9
1.3.1	Anomalous Neutral Boson Couplings	10
1.3.2	Origin of Anomalous Couplings	11
1.4	ZZ Diboson Production	13
1.5	Event Generation	14
1.5.1	Matrix Element Calculation	14
1.5.2	Parton Distribution Functions	14
1.5.3	Parton Showering and Hadronisation	15
1.5.4	Monte Carlo Generators	16
1.5.5	Anomalous Coupling Generators	17
1.6	ZZ Inclusive Cross Section	18
1.6.1	Gluon-Gluon Fusion	19
1.6.2	ZZ Decay Modes	20
1.6.3	Backgrounds to ZZ Production	20
1.6.4	Anomalous Triple Gauge Couplings at the LHC	21
1.7	Current Experimental Results	24
1.7.1	Results from LEP	24
1.7.2	Results from the Tevatron	25
1.7.3	ATLAS Sensitivity Studies	28
1.7.4	Diboson Background to New Physics	29
1.8	Conclusions	30

2	The ATLAS Experiment	31
2.1	The Large Hadron Collider	31
2.1.1	Context	32
2.2	The ATLAS Experiment	33
2.2.1	Co-ordinate Definition	34
2.2.2	Prospects for New Discovery	34
2.2.3	Experimental Requirements	35
2.3	The ATLAS Detector	35
2.3.1	The Inner Detector	36
2.3.2	Calorimeters	37
2.3.3	Muon Spectrometer	39
2.3.4	Forward Detectors	40
2.4	Trigger, Data Acquisition and Control Systems	40
3	The ATLAS Semiconductor Tracker	43
3.1	Introduction	43
3.2	Silicon Detectors	44
3.2.1	Physics of Semiconductors	44
3.2.2	p - n Junctions	44
3.2.3	Semiconductors as Particle Detectors	45
3.2.4	Radiation Damage	45
3.3	The ATLAS SCT	46
3.4	SCT Modules	48
3.4.1	Silicon Sensors	48
3.4.2	ABCD3T Chips	50
3.5	The SCT Data Acquisition System	51
3.5.1	Hardware	52
3.5.2	Software	57
3.6	Calibration of the SCT	59
3.6.1	Optical Tuning	59
3.6.2	Digital Tests	61
3.6.3	Analogue Calibration	61
3.6.4	Calibration Results	62
3.7	Data-taking	64
3.7.1	Milestone Runs	64
3.7.2	First Circulating Beam in the LHC	64

3.7.3	Cosmic Muons	64
3.7.4	First LHC Collisions	65
3.8	Conclusions	67
4	The SCT Test Sectors	69
4.1	Installation	70
4.2	Barrel Sector	70
4.2.1	Barrel Sector Layout	71
4.2.2	Power Supply	72
4.2.3	Optical Connections	72
4.2.4	Harness Testing	72
4.2.5	Final Assembly	73
4.3	End-cap Disk	73
4.3.1	Disk Installation	73
4.4	Detector Control System	74
4.5	Data Acquisition System	75
4.6	Calibrating the Sectors	75
4.7	Conclusions	77
5	SCT Data Acquisition Developments	79
5.1	Simultaneous Loading of Slave Firmware	79
5.2	Errors During Scans	80
5.3	Masking Noisy Links	80
5.3.1	Hot-Link Monitoring	80
5.3.2	Examples of Use	82
5.4	ABCD3T Simulation	82
5.4.1	Simulator Design	83
5.4.2	Event Counter Generation	84
5.4.3	Linear Feedback Shift Register	87
5.4.4	The Simulator Engine	88
5.4.5	Flagging Simulated Events Offline	91
5.4.6	Examples of Simulator Use	91
5.5	ROD-Level Monitoring	94
5.5.1	Digital Signal Processor (DSP) Histogramming	94
5.5.2	Crate-Level Modifications	97
5.5.3	Analysis Service	99
5.5.4	The SCT ROD Monitoring Application	100

5.6	Results from ROD-Monitoring	101
5.6.1	Noise Occupancy Measurements	101
5.6.2	Coincidence Rate Measurements with Cosmics	103
5.6.3	Beam Splash Measurements	108
5.6.4	First Collision Measurements	109
5.6.5	SCT Beam Monitoring	109
5.6.6	Timing Ratio Measurements	111
5.7	Conclusions	114
6	ZZ Diboson Selection in ATLAS	115
6.1	Introduction	115
6.2	Technical Overview	116
6.2.1	Distributed Analysis	116
6.2.2	ROOT Analysis	116
6.3	Simulated Data Samples	117
6.3.1	Simulating the ATLAS Detector	117
6.3.2	10 TeV Datasets	118
6.3.3	7 TeV Datasets	122
6.3.4	Cross Section Calculations	122
6.4	Common Preselection	126
6.4.1	Electron Selection	127
6.4.2	Muon Selection	129
6.4.3	Jet Selection	134
6.4.4	Missing Energy Selection	135
6.4.5	Z Reconstruction	136
6.4.6	Truth Selection	136
6.5	$ZZ \rightarrow lll$ Event Selection	137
6.5.1	Main Backgrounds	138
6.5.2	Cut Definition	138
6.5.3	Expected Yields	141
6.6	$ZZ \rightarrow ll\nu\bar{\nu}$ Event Selection	144
6.6.1	Main Backgrounds	144
6.6.2	Cut Definition	146
6.6.3	Expected Yields	148
6.7	Trigger Analysis	150
6.8	Systematic Errors	151

6.8.1	Monte Carlo Uncertainties	151
6.8.2	Luminosity Uncertainties	152
6.8.3	Detector Uncertainties	152
6.9	Summary of Expected Yields	153
6.10	Conclusions	156
7	Anomalous Triple Gauge Coupling Sensitivity	157
7.1	Introduction	157
7.2	Anomalous Coupling Monte Carlo	158
7.2.1	Form Factors	158
7.3	Extended Likelihood Function	160
7.4	Calculation of $P(f_i^V p_T^k)$	161
7.4.1	Detector Effects	163
7.5	Calculation of $\mu(f_i^V)$	166
7.6	Anomalous Coupling Sensitivity	167
7.6.1	Toy Monte Carlo Samples	167
7.6.2	Fitting Method	168
7.6.3	Results	169
7.7	Systematic Uncertainties	176
7.7.1	Uncertainties on $\mu(f_i^V)$	176
7.7.2	Uncertainties on $P(f_i^V p_T^k)$	176
7.7.3	Summary of Systematic Uncertainties	177
7.8	Gluon-Gluon Fusion	179
7.9	Conclusions	179
8	Observation of $Z \rightarrow ll$ in Early ATLAS Data	181
8.1	Introduction	181
8.2	Z Production	181
8.3	Event Selection	182
8.3.1	Pre-selection	182
8.3.2	Selection Cuts	184
8.4	Z Cross Section	188
8.5	A Diboson Prelude	190
8.5.1	$ZZ \rightarrow lll$ Channel	190
8.5.2	$ZZ \rightarrow ll\nu\bar{\nu}$ Channel	191
8.5.3	Anomalous Coupling Limits	192
8.6	Conclusions	192

9 Summary	195
A Thesis Statistics	199
A.1 Word Counting	199
A.2 Measuring the Length of the Day	202
Bibliography	207
Acronyms	213
List of Figures	217
List of Tables	221

*“You know, when I was a kid,
I always thought I was gonna grow up to be a hero.”*
— Butch Cassidy

Chapter 1

Anomalous Triple Gauge Couplings

*“Cause we’re all just
Protons, Neutrons, Electrons
That rest on a Sunday
Work on a Monday”
— The Cat Empire*

1.1 Introduction

Production of pairs of Electro-weak (EW) gauge bosons at the Large Hadron Collider (LHC) are predicted by the Standard Model (SM) of particle physics. This chapter gives a brief description of how EW interactions arise as a unified description of two of the four fundamental forces of nature: electromagnetism and the weak interaction. Self-couplings between three neutral gauge bosons (ZZZ , γZZ) are forbidden in the SM, but will be introduced as a possible consequence of new physics. The aim of this chapter is to provide sufficient theoretical and experimental background in order to motivate the studies presented in chapters 6 and 7.

1.2 The Electro-weak Interaction

In Quantum Field Theory (QFT), equations of motion can be obtained from the Lagrangian density of the system, \mathcal{L} , and the Euler-Lagrange equation. For example, the

Lagrangian density for a non-interacting Dirac fermion is

$$\mathcal{L}_D = \bar{\psi}i\gamma^\mu\partial_\mu\psi - m\bar{\psi}\psi, \quad (1.1)$$

for a field operator ψ , where γ^μ are the Dirac γ -matrices, and $\bar{\psi} = \psi^\dagger\gamma^0$ is the conjugate field. Applying the Euler-Lagrange equation

$$\frac{\partial\mathcal{L}}{\partial\psi} - \partial_\mu\left(\frac{\partial\mathcal{L}}{\partial(\partial_\mu\psi)}\right) = 0, \quad (1.2)$$

to equation 1.1 yields the Dirac equation

$$(i\gamma^\mu\partial_\mu - m)\psi = 0. \quad (1.3)$$

Considering a free electromagnetic field, the required Lagrangian density is

$$\mathcal{L}_{\text{EM}} = -\frac{1}{4}F_{\mu\nu}F^{\mu\nu}, \quad (1.4)$$

where $F^{\mu\nu} = \partial^\mu A^\nu - \partial^\nu A^\mu$ is the field strength tensor, and A^μ is the electromagnetic 4-potential. To introduce interactions between fermions and the electromagnetic field, an additional interaction term can be added to the Lagrangian density by replacing ∂_μ with the *covariant derivative*

$$D_\mu = \partial_\mu + ieA_\mu \quad (1.5)$$

where e is the electric charge of the electron. The overall Lagrangian for a fermion interacting with the electromagnetic field is

$$\mathcal{L} = \bar{\psi}i\gamma^\mu D_\mu\psi - m\bar{\psi}\psi - \frac{1}{4}F_{\mu\nu}F^{\mu\nu} \quad (1.6)$$

Imposing a gauge transformation of the electromagnetic field, so that for an arbitrary function $\chi(x)$,

$$A_\mu \rightarrow A'_\mu = A_\mu + \partial_\mu\chi, \quad (1.7)$$

also requires a simultaneous change of phase of the fermion field

$$\psi \rightarrow \psi' = e^{-ie\chi}\psi, \quad (1.8)$$

to preserve gauge invariance of the Lagrangian. Conversely, demanding that the Lagrangian is unchanged under local phase transformations of ψ , also requires the existence of a (massless) vector field to cancel terms involving $\partial_\mu\chi$. Noether's theorem states that there is a conserved current for every continuous symmetry of the Lagrangian, which in this case implies the conservation of electric current and charge.

Successive gauge transformations commute,

$$e^{-ie\chi_1}e^{-ie\chi_2} = e^{-ie\chi_2}e^{-ie\chi_1}, \quad (1.9)$$

and the infinite set of transformations forms the unitary Abelian gauge group $U(1)$.

1.2.1 The Weak Interaction

As an extension of the $U(1)$ gauge theory of electromagnetic interactions, transformations that mix together two fermion fields can also be considered. In the case of the weak interaction, mixing of leptons (or quarks) of a given generation is possible, for example the electron and electron neutrino:

$$\Psi_e = \begin{pmatrix} \psi_{\nu_e} \\ \psi_e \end{pmatrix}. \quad (1.10)$$

The mixing of the two states can be described by the transformation

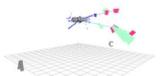
$$\Psi \rightarrow \Psi' = U\Psi, \quad (1.11)$$

where U is a 2×2 matrix, which must be unitary ($UU^\dagger = I$) to preserve normalisation and orthogonality. This matrix can be written in terms of the *non-commuting* Pauli spin matrices $\tau = \tau_1, \tau_2, \tau_3$, where

$$\tau_1 = \begin{pmatrix} 0 & 1 \\ 1 & 0 \end{pmatrix}, \tau_2 = \begin{pmatrix} 0 & -i \\ i & 0 \end{pmatrix}, \tau_3 = \begin{pmatrix} 1 & 0 \\ 0 & -1 \end{pmatrix} \quad (1.12)$$

so that

$$U = \exp\left(i\frac{g}{2}\boldsymbol{\omega} \cdot \boldsymbol{\tau}\right). \quad (1.13)$$



where $\boldsymbol{\omega} = \boldsymbol{\omega}(x)$ is an arbitrary vector about which rotations take place in the so-called weak isospin space. g is a constant, analogous to the electric charge e . Unlike the electromagnetic case, successive rotations do not commute, and so the infinite set of transformations are described by the non-Abelian group of traceless 2×2 matrices with unit determinant, $SU(2)$.

A gauge invariant $SU(2)$ theory is obtained by introducing a massless isovector field, $\mathbf{W}^\mu = W_1^\mu, W_2^\mu, W_3^\mu$ with both charged and neutral components. Invariance of the Lagrangian $\mathcal{L}_D = \bar{\psi} i \gamma^\mu D_\mu \psi$ requires a covariant derivative of the form

$$D_\mu = \partial_\mu + i \frac{g}{2} \boldsymbol{\omega} \cdot \boldsymbol{\tau} \quad (1.14)$$

so that these new fields are now coupled to the fermions. The corresponding gauge transformation of the field itself is more complex, and takes the form

$$\mathbf{W}^\mu \rightarrow \underbrace{\mathbf{W}^\mu - \partial^\mu \boldsymbol{\omega}}_{\text{EM-like piece}} - \underbrace{g \boldsymbol{\omega} \times \mathbf{W}^\mu}_{\text{extra term}}. \quad (1.15)$$

The fact that the Pauli matrices $\boldsymbol{\tau}$ do not commute has made it necessary to include an extra term (when compared to equation 1.7) in order to preserve overall gauge invariance.

To allow propagation of these gauge fields, a free-field term must be added to the Lagrangian

$$\mathcal{L}_W = -\frac{1}{4} \mathbf{G}^{\mu\nu} \mathbf{G}_{\mu\nu}, \quad (1.16)$$

where we require

$$\mathbf{G}^{\mu\nu} = \partial^\mu \mathbf{W}^\nu - \partial^\nu \mathbf{W}^\mu - \underbrace{g \mathbf{W}^\mu \times \mathbf{W}^\nu}_{\text{self-interactions}} \quad (1.17)$$

to satisfy gauge invariance. A striking result of this theory is that the Lagrangian now contains terms representing self-interactions of the gauge fields, which are discussed in section 1.3. As it stands, however, the above formalism does not fully describe the weak interactions and a complete description requires the unification of the weak and electromagnetic forces.

1.2.2 Electro-weak Unification

In the 1960s, Weinberg [1] and Salam [2] proposed that the weak and electromagnetic interactions derive from a single electro-weak force based on a combination of an $SU(2)_L$ weak isospin I group and a $U(1)$ group of hypercharge Y . In this theory, the fundamental bosons are a massless triplet $\mathbf{W}^\mu = W_1^\mu, W_2^\mu, W_3^\mu$ for $SU(2)$ and a massless singlet B^μ for $U(1)$. The covariant derivative takes the form

$$D_\mu = \partial_\mu + ig\mathbf{I} \cdot \mathbf{W}_\mu + ig'B_\mu Y \quad (1.18)$$

where $I = \frac{1}{2}\tau$ and Y are the isospin and hypercharge respectively, and g and g' are their couplings to the \mathbf{W}_μ and B_μ .

The physical bosons observed in nature consist of linear combinations of the fundamental fields \mathbf{W}_μ and B_μ . The charged bosons W^\pm can be constructed as

$$W_\mu^\pm = \frac{1}{\sqrt{2}}(W_\mu^1 \mp iW_\mu^2). \quad (1.19)$$

Combining the above expression with equation 1.10 and equation 1.18 gives an interaction term of the form $\frac{g}{\sqrt{2}}\bar{\psi}_{\nu_e}\gamma^\mu W_\mu^+\psi_e$, with corresponding Feynman diagram shown in figure 1.1.

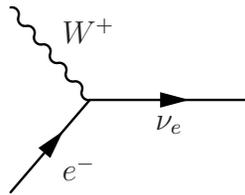
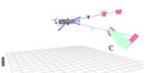


Figure 1.1: Feynman diagram of the weak charged current interaction.

Experimentally it is known that the W^\pm bosons actually only interact with the left-handed fermion states $\psi_L = \frac{1}{2}(1 - \gamma^5)\psi$ and right-handed anti-fermions, thus violating parity (P). This is incorporated into the theory by assigning a weak isospin $I_w = \frac{1}{2}$ to all left-handed fermions, while assigning $I_w = 0$ for the right-handed states.



The neutral gauge bosons, Z and γ interact with both left and right-handed fermions, and are a combination of W_μ^3 and B_μ :

$$Z_\mu = W_\mu^3 \cos \theta_w - B_\mu \sin \theta_w \quad (1.20)$$

$$A_\mu = W_\mu^3 \sin \theta_w + B_\mu \cos \theta_w \quad (1.21)$$

where θ_w is the weak-mixing or Weinberg angle. If these physical definitions are now added to equation 1.18, the covariant derivative takes the form:

$$D_\mu = \partial_\mu + \underbrace{i \frac{g}{\sqrt{2}} (W_\mu^+ + W_\mu^-)}_{\text{weak CC}} + \underbrace{i \frac{g}{\cos \theta_w} (I_3 - Q \sin^2 \theta_w) Z_\mu}_{\text{weak NC}} + \underbrace{igQ \sin \theta_w A_\mu}_{\text{EM NC}}, \quad (1.22)$$

where $Q = I_3 + Y$ is the electric charge and I_3 is the third component of the weak isospin. This shows us that the interaction contains three pieces, the weak Charged Current (CC), the weak Neutral Current (NC) and the electromagnetic neutral current. By comparison with equation 1.5, equating the couplings to the electromagnetic interaction gives

$$e = g \sin \theta_w. \quad (1.23)$$

1.2.3 Spontaneous Symmetry Breaking

One remaining problem with electro-weak theory is that of gauge boson masses. Experimentally, the W and Z gauge bosons are known to be massive, with $m_W \sim 80$ GeV and $m_Z \sim 90$ GeV. To generate a massive gauge boson field requires an additional term in the Lagrangian compared to equation 1.4:

$$\mathcal{L} = -\frac{1}{4} V_{\mu\nu}^\dagger V^{\mu\nu} + m_V^2 V_\mu^\dagger V^\mu, \quad (1.24)$$

where $V = Z, W$ and $V_{\mu\nu} = \partial_\mu V_\nu - \partial_\nu V_\mu$. The mass term $m_V^2 V_\mu^\dagger V^\mu$, however, is forbidden by gauge invariance.

One way in which to reconcile this problem is to consider the case where a symmetry of the Lagrangian is not a symmetry of the vacuum, known as spontaneous symmetry breaking. This is done by introducing a complex scalar field, ϕ , with Lagrangian,

$$\mathcal{L} = (\partial^\mu \phi^\dagger)(\partial_\mu \phi) - \mu^2 \phi^\dagger \phi - \lambda(\phi^\dagger \phi)^2, \quad (1.25)$$

which can interact with massless real vector fields. If this symmetry is now broken ($\mu^2 < 0$), a massive gauge field is recovered. A real scalar field is also predicted, with an associated massive boson known as the Higgs boson. The discovery of such a particle would justify the spontaneous symmetry breaking theory, and as such is a major goal for the LHC.

1.2.4 The Standard Model

The Standard Model is completed by the strong interaction, an $SU(3)$ gauge theory describing colour interactions between quarks and gluons. The combination of the strong and electro-weak gauge theories yields an overall gauge symmetry of $SU(3)_C \times SU(2)_L \times U(1)$.

1.3 Triple Gauge Couplings

Terms of the form $g\mathbf{W}^\mu \times \mathbf{W}^\nu$ in equation 1.17 represent self-interactions between the gauge bosons, with both triple and quartic couplings possible. The triple gauge coupling vertex takes the form

$$-\frac{1}{2} \epsilon_{jkl} (\partial_\mu W_j^\nu - \partial_\nu W_j^\mu) W_{k\mu} W_{l\nu}, \quad (1.26)$$

where the factor ϵ_{jkl} only allows interactions involving three different types of boson. For example, a Z or γ can interact with a W^\pm pair as shown in figure 1.2. This also

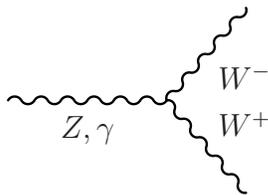
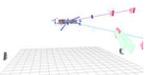


Figure 1.2: Feynman diagram of the WWZ triple gauge boson interaction.

means that production of a Z pair via a triple gauge vertex is forbidden in the SM at tree level. Anomalous interactions between three neutral gauge bosons, however, may arise as a result of new physics beyond the SM.



1.3.1 Anomalous Neutral Boson Couplings

The most general $Z_1 Z_2 V$ vertex is defined in figure 1.3, where Z_1 and Z_2 are both on-shell, while $V = Z, \gamma$ is in general off shell, but coupled to a conserved current. The vertex function, Γ_{ZZV} takes the form [3]:

$$\Gamma_{ZZV}^{\alpha\beta\mu} = e \frac{p^2 - m_V^2}{m_Z^2} [i f_4^V (p^\alpha g^{\mu\beta} + p^\beta g^{\mu\alpha}) + i f_5^V \epsilon^{\mu\alpha\beta\rho} (q_1 - q_2)_\rho], \quad (1.27)$$

where m_V is the boson mass, $p = \sqrt{s}$ is the centre-of-mass energy of the interaction and $q_{i=1,2}$ are the momenta of the outgoing bosons, Z_1 and Z_2 , as shown in figure 1.3. The effective Lagrangian generating the vertex function in equation 1.27 is [4]

$$\mathcal{L} = \frac{e}{m_Z^2} \left[f_4^V (\partial_\mu V^{\mu\beta}) Z_\alpha (\partial^\alpha Z_\beta) + f_5^V (\partial^\sigma V_{\sigma\mu} \tilde{Z}^{\mu\beta} Z_\beta) \right], \quad (1.28)$$

where $\tilde{Z}^{\mu\beta} = \frac{1}{2} \epsilon_{\mu\nu\rho\sigma} Z^{\rho\sigma}$. The Neutral Triple Gauge Couplings (NTGCs), f_4^V and f_5^V , are both complex functions of q_1^2 , q_2^2 and p^2 . All couplings violate charge conjugation, C . f_4^Z and f_4^γ are CP violating, and so contributions to the helicity amplitudes will not interfere with SM ZZ production. The f_5^V couplings are CP conserving, but violate parity, P .

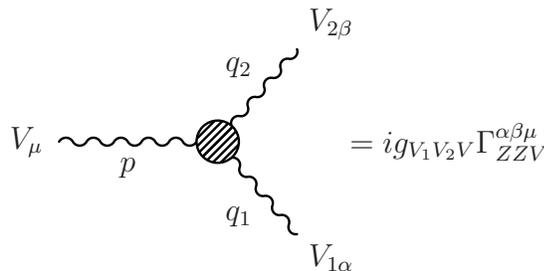


Figure 1.3: Feynman rule for the anomalous ZZV vertex.

If the Z_1 and Z_2 bosons are allowed to be off-shell, five additional anomalous ZZZ couplings and five additional anomalous $ZZ\gamma$ couplings contribute [5]. In this case, the factor $p^2 - m_V^2$ in equation 1.27 is replaced by $q_1^2 - q_2^2$, strongly suppressing the effect of these additional couplings, which will not be considered further.

It should be noted that there are also four $Z\gamma Z$ couplings ($h_i^Z, i = 1..4$) which appear in $Z\gamma$ production, where the Z is assumed on-shell. These are completely independent of the two $ZZ\gamma$ couplings. If all three bosons in the vertex are assumed to be off-shell, then a total of seven couplings contribute [5]. Four of these remain when considering $Z\gamma$ production, and two ($f_i^\gamma, i = 1, 2$) in ZZ production.

The contribution of the anomalous couplings to the cross section will grow with the centre-of-mass energy, \sqrt{s} . To avoid unphysical results which violate unitarity, an energy dependant form factor is typically introduced [6]:

$$f_i^V(s) = \frac{f_{i0}^V}{(1 + s/\Lambda_{FF}^2)^n} \quad (1.29)$$

where $i = 4, 5$, n is a constant and Λ_{FF} is a cutoff related to the scale of the new physics generating the anomalous couplings. The requirement that tree-level unitarity is satisfied across the entire range of s puts upper limits on the coupling parameters [7]:

$$|f_{4,5}^Z| \leq \frac{4}{\alpha} \sqrt{\frac{3}{10}} \sin \theta_w \cos \theta_w \times \left(\frac{m_Z}{\Lambda_{FF}} \right)^3 \frac{\left(\frac{2}{3}n\right)^n}{\left(\frac{2}{3}n - 1\right)^{(n-3/2)}}, \quad (1.30)$$

$$|f_{4,5}^\gamma| \leq \frac{1}{\alpha} \left[\frac{3}{5}(3 - 6 \sin^2 \theta_w + 8 \sin^4 \theta_w) \right]^{\frac{1}{2}} \times \left(\frac{m_Z}{\Lambda_{FF}} \right)^3 \frac{\left(\frac{2}{3}n\right)^n}{\left(\frac{2}{3}n - 1\right)^{(n-3/2)}}, \quad (1.31)$$

where α is the fine structure constant. Equations 1.30 and 1.31 require $n > 3/2$ to satisfy unitarity. A common choice [7] is $n = 3$, which will be used throughout this thesis. Figure 1.4 shows the unitarity limits as a function of Λ_{FF} , taking $\alpha = 1/128.93$, $\sin^2 \theta_w = 0.2310$ and $m_Z = 91.187$ GeV [8].

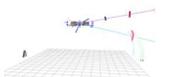
1.3.2 Origin of Anomalous Couplings

The simplest method for generating anomalous couplings via the vertex in equation 1.27 is via virtual effects of heavy fermions at the one-loop level [9], as shown in figure 1.5. Heavy fermions can generate the f_5^V couplings, whereas higher order terms are required to produce the CP -violating couplings, f_4^V .

Considering a single heavy fermion, F , interacting with Z and γ , leads to the relation of [4]:

$$f_5^V \propto \frac{\alpha}{4\pi} \frac{m_Z^2}{M_F^2}, \quad (1.32)$$

where M_F is the heavy fermion mass. In a new physics scenario with a completely degenerate family of heavy quarks and leptons, the couplings would vanish due to cancellations. This is known as the unbroken $SU(2) \times U(1)$ situation.



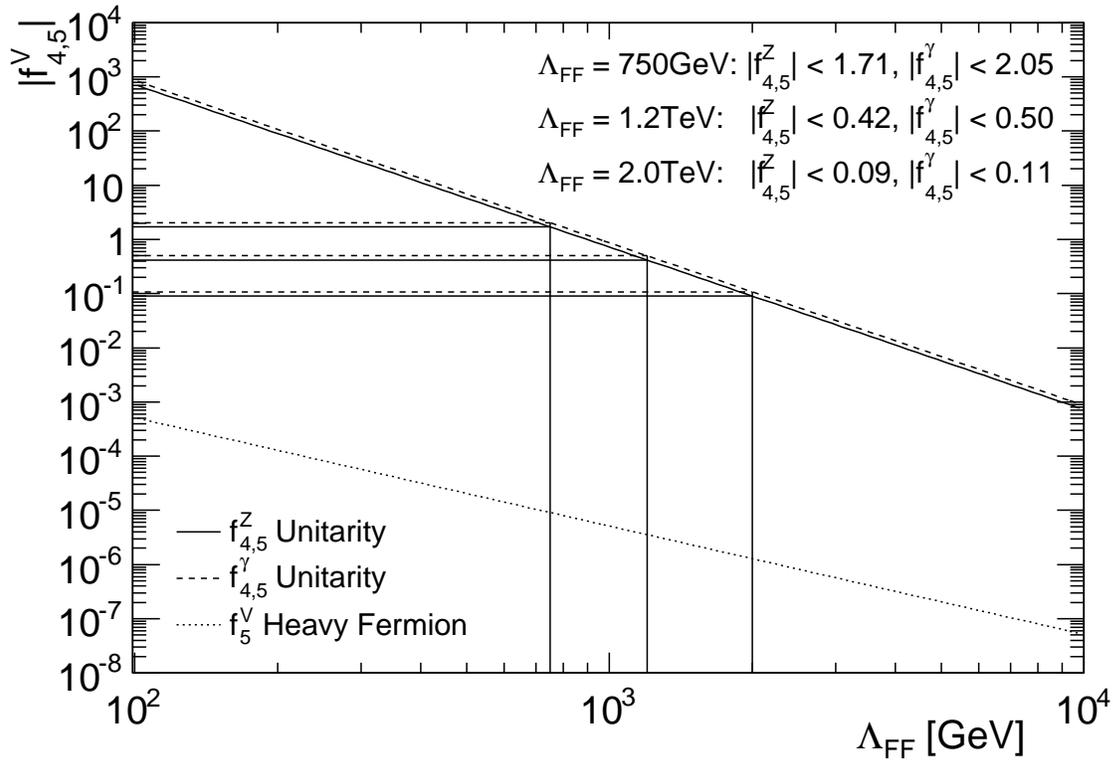


Figure 1.4: Unitarity limits on anomalous coupling magnitude against the form factor cutoff, Λ_{FF} . The limits for a number of popular cutoff choices are shown on the plot. The dotted line shows the order of the f_5^V couplings given the heavy fermion model in equation 1.32.

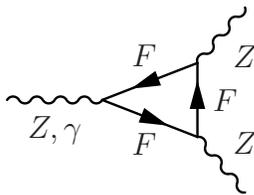


Figure 1.5: Production of an anomalous coupling vertex via heavy fermions.

The couplings can be restored by introducing a mass splitting between heavy lepton and quark doublets on the EW scale, m_Z . This will, however, produce couplings which are suppressed to m_Z^4/M_F^4 , known as the broken $SU(2) \times U(1)$ situation.

The final possibility arises where one fermion has a much lighter mass than the other fermions in the family, so that the couplings appear as m_Z^2/M_F^2 . This is the case where the $SU(2) \times U(1)$ symmetry is strongly broken. Even in this experimentally optimistic

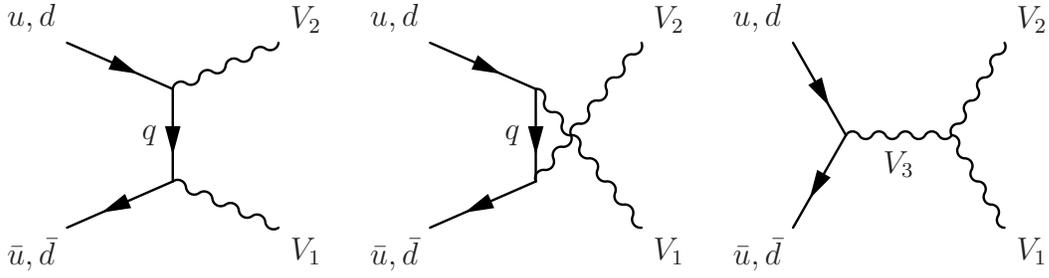


Figure 1.6: Leading order Feynman diagrams for ZZ production at the LHC.

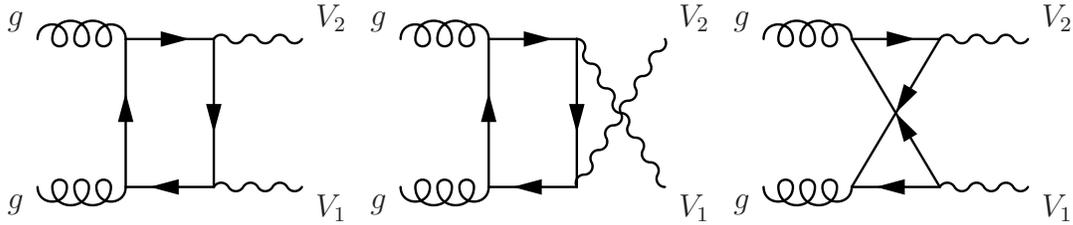


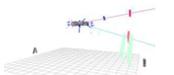
Figure 1.7: Feynman diagrams for ZZ production via gluon-gluon fusion.

case, equation 1.32 predicts coupling values of $\mathcal{O}(10^{-3})$ for heavy fermions in the 100 GeV mass range. The order of the coupling strength over a range of heavy fermion masses is shown in figure 1.4.

One New Physics (NP) theory which is able to generate the triangle diagram in figure 1.5 is the Minimally Supersymmetric Standard Model (MSSM) [9], where the heavy fermions are charginos and neutralinos. The charginos, $\tilde{\chi}_{1,2}^{\pm}$, contribute to the f_5^V couplings, while the neutralinos, $\tilde{\chi}_{1-4}^0$, contribute to f_5^Z only. Non-perturbative effects could enhance the f_4^V vertices by coupling the photon and Z to axial and vector resonances predicted by technicolour (TC) models.

1.4 ZZ Diboson Production

Experimentally, the anomalous NTGCs can be measured by studying ZZ diboson production. At hadron colliders, $q\bar{q} \rightarrow ZZ$ diboson production proceeds at tree level via the t-channel Feynman diagrams shown in the two left-hand diagrams of figure 1.6, with $V_{1,2} = Z$. The right-hand diagram involves the anomalous coupling ($V_3 = Z, \gamma$), which is zero at tree level in the SM. Gluon-gluon fusion, $gg \rightarrow ZZ$, is also expected to contribute to diboson production [10, 11], via the quark box diagrams shown in figure 1.7.



The expected number of events produced for a given process is given by

$$N = \int \mathcal{L} dt \sigma, \quad (1.33)$$

where $\mathcal{L} = \int L dt$ is the integrated luminosity, and σ is the cross section for that process. The cross section for ZZ production at the LHC is calculated using Monte Carlo generators, as discussed below.

1.5 Event Generation

Given the inherent probabilistic nature of QFT, predictions about the properties of any single event cannot be made. Instead, a large statistical sample of events for a given process is built up using random numbers to reproduce quantum mechanical properties, hence the name Monte Carlo generators. In brief, a Monte Carlo generator will typically calculate particle production in a number of stages, described below.

1.5.1 Matrix Element Calculation

In the first stage, the hard process is calculated from Feynman diagram matrix elements using perturbation theory. The hard process is usually the physics interaction of interest, for example, $q\bar{q} \rightarrow ZZ$. The hard process can be calculated at Leading Order (LO), which contains the minimal number of vertices possible for that process. Higher order calculations can be calculated, for example Next-to-Leading Order (NLO), which contain corrections due to additional interactions.

Two choices need to be made regarding the scale of the calculation. The renormalisation scale determines the energies at which higher-order, non-perturbative Quantum Chromodynamic (QCD) divergences are suppressed. The factorisation scale describes a second cutoff introduced to deal with divergences caused by collinear gluon radiation.

1.5.2 Parton Distribution Functions

In the case of hadron-hadron collisions, such as the proton-proton interactions at the LHC, an additional complication arises due to the fact that the colliding particles are not point like, but instead consist of quarks and gluons, collectively known as partons.

In the most simplistic picture, protons are composite particles consisting of two up quarks, and a single down quark, known as valance partons. This is not the whole story, as QCD interactions can spontaneously radiate gluons, which in turn produce short-lived $q\bar{q}$ pairs, together known as sea partons. The momentum carried by a given type of parton relative to the proton as a whole is known as the longitudinal momentum fraction, x .

The probability density for finding a parton f with a certain value of x , and at momentum transfer Q^2 are known as Parton Distribution Functions (PDFs), where

$$PDF = PDF(x, Q, f). \quad (1.34)$$

These distributions are typically extracted using data from a number of hard scattering experiments in the framework of perturbative QCD. The extracted PDF tables are then interfaced to event generators which calculate the interaction cross section between different combinations of partons.

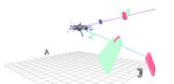
Typically both Leading and Next-to-Leading order PDFs are available, for example CTEQ6L [12] and CTEQ66 [13] respectively. In general the PDF order is matched with the order of the hard process being generated, although this is not a strict requirement. One exception are the MRST Modified Leading Order (LO*) PDFs [14], which aim to provide NLO-like PDF shapes with leading order generators. This is useful to produce events with more accurate kinematic distributions, but the cross sections require rescaling to the true NLO values.

1.5.3 Parton Showering and Hadronisation

An additional ingredient is the parton shower, which adds soft, co-linear corrections due to QCD processes, such as gluon radiation. If partons are also produced as part of the hard process, then an appropriate matching algorithm (such as CKKW [15]) is necessary to avoid double counting.

The final stage in the process is the hadronisation of bare partons into stable particles, and the addition of any particles from the underlying event, such as remnants from the incoming particles.

An example illustrating each element in event generation is shown in figure 1.8. The hard process of the event is the $pp \rightarrow ZZ \rightarrow ee\mu\mu$ interaction, created using



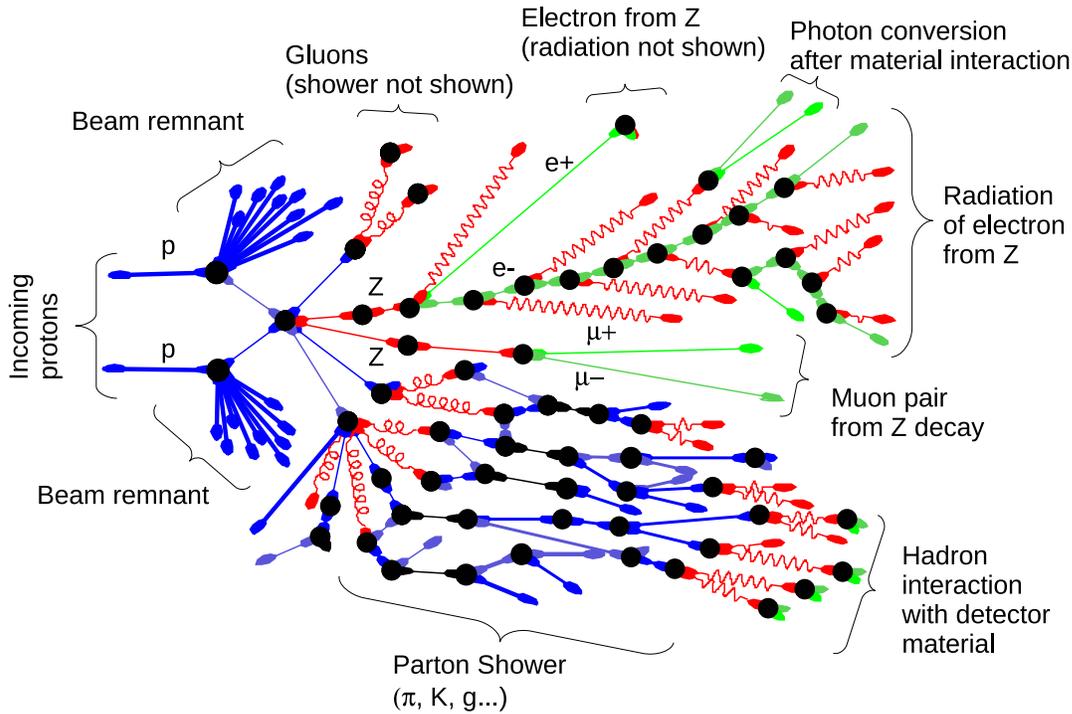


Figure 1.8: A cartoon of a ZZ event generated using MC@NLO. The event has been visualised from the generator event record, using HepMCVisual [16].

the MC@NLO generator. The beam remnants, parton showering and also material interaction effects are all clearly visible.

1.5.4 Monte Carlo Generators

Many such generators are available to the modern high energy physicist, each with its own benefits and drawbacks. The output will typically be a cross section calculation, together with a list of particles involved in each generated event, and their corresponding 4-vectors.

Pythia The PYTHIA [17] event generator produces events at LO from a built in library of almost all SM $2 \rightarrow 1$ and $2 \rightarrow 2$ processes. It includes parton shower models for both initial- and final-state radiation, with necessary matrix element matching. Hadronisation is performed using the string or Lund Model [18].

In chapter 6, Pythia is used to generate a large sample of $ZZ \rightarrow lll$ signal events, as it includes resonance decays of the Z , and interference with γ to give off-shell Z

production, labelled as Z^* . A minimum mass cutoff is required to preserve unitarity, and for Pythia this is internally set at $m_Z > 12$ GeV.

MC@NLO The MC@NLO generator [19,20] is an NLO event generator, which calculates the matrix element with additional single parton corrections to the hard process. These corrections give rise to weighted events $w = \pm 1$, which need to be taken into account when calculating expected yields. Parton shower matching is also performed with Herwig [21], with the JIMMY [22] package used to generate the underlying event.

In chapter 6 MC@NLO is used to generate $ZZ \rightarrow ll\nu\bar{\nu}$ signal sample, but unlike Pythia, does not include resonance decays, and so each Z is produced on-shell, with a fixed mass. MC@NLO is also used to generate a number of background samples, notably $t\bar{t}$, and other diboson processes.

AlpGen The AlpGen [23] generator is used to calculate leading order matrix elements for multiple particle final states ($2 \rightarrow n$). In chapter 6 it is used to generate events containing multiple partons, such as $Z + n$ Jets.

MCFM The MCFM generator [24,25], although not used to produce any large event samples, is useful in the calculation of cross sections. Both Leading and Next-to-Leading Order calculations are available, with on- and off-shell Z decays possible.

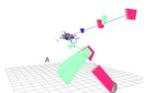
FEWZ The FEWZ generator [26] calculates cross sections for Z and W production at Next-to-Next-to-Leading Order (NNLO). It is used in chapter 8 to calculate the cross section for $Z/\gamma \rightarrow ll$ production.

GG2ZZ is a generator specialising in calculating the gluon-gluon fusion process $gg \rightarrow (Z/\gamma)(Z/\gamma)$ in hadron colliders [27]. It will be used in section 1.6.1 to estimate the contribution of this channel to ZZ production at the LHC.

1.5.5 Anomalous Coupling Generators

There are a number of event generators available with the capability of producing events containing anomalous coupling vertices, described briefly below.

Baur and Rainwater (BR) have made available a leading order generator, dedicated to the production of events containing anomalous coupling vertices [7]. The program generates the hard scattering and Z boson decays only; no underlying event or initial state radiation is included, hence the p_T of the two Z bosons is identical. The Z width and Z/γ interference are both included in the generator.



Sherpa The SHERPA [28] program also has the ability to produce events with anomalous couplings. SHERPA is a leading order generator, and unlike BR, includes parton showering, initial state radiation and the underlying event.

1.6 ZZ Inclusive Cross Section

The cross section for $q\bar{q} \rightarrow ZZ$ in proton-proton collisions at a centre-of-mass energy of $\sqrt{s} = 7$ TeV has been calculated with both the MC@NLO and MCFM generators as described in [29], using the MSTW2008 [30] PDFs. The NLO MC@NLO cross section for on-shell ZZ production was calculated as

$$\sigma(q\bar{q} \rightarrow ZZ) = 5.86_{-0.10}^{+0.14} \text{ (scale)} \pm 0.22 \text{ (PDF) pb.} \quad (1.35)$$

The NLO order MCFM cross section for on-shell ZZ production was calculated as

$$\sigma(q\bar{q} \rightarrow ZZ) = 6.04_{-0.14}^{+0.16} \text{ (scale)} \pm 0.20 \text{ (PDF) pb,} \quad (1.36)$$

with the off-shell ($m_{Z/\gamma} > 12$ GeV) cross section calculated as

$$\sigma(q\bar{q} \rightarrow ZZ) = 9.27_{-0.09}^{+0.19} \text{ (scale)} \pm 0.20 \text{ (PDF) pb.} \quad (1.37)$$

The cross sections have been calculated using central values of $\mu_F = \mu_R = m_Z$ for both the renormalisation and factorisation scales. This central value is then adjusted by a factor of two ($\mu_F = \mu_R = 2m_Z$ and $\mu_F = \mu_R = m_Z/2$) in order to estimate uncertainties due to the choice of scale. The quoted PDF uncertainties have been calculated using the prescription given in reference [31] in the case of MC@NLO and in reference [32] for MCFM.

The two generators are consistent with each other for the on-shell calculation at the level of 3%. Both show scale uncertainties of $\sim 2\%$, and PDF uncertainties at the 2-3% level. Re-calculation of the cross sections with the CTEQ66 [13] PDF set gave values which were consistently lower by 3.6%.

The theoretical centre-of-mass energy dependence of the inclusive ZZ production cross section, for proton-proton and proton-(anti-)proton collisions is shown in figure 1.9. The ratio of ZZ cross sections for pp collisions at $\sqrt{s} = 7$ TeV to $p\bar{p}$ collisions at

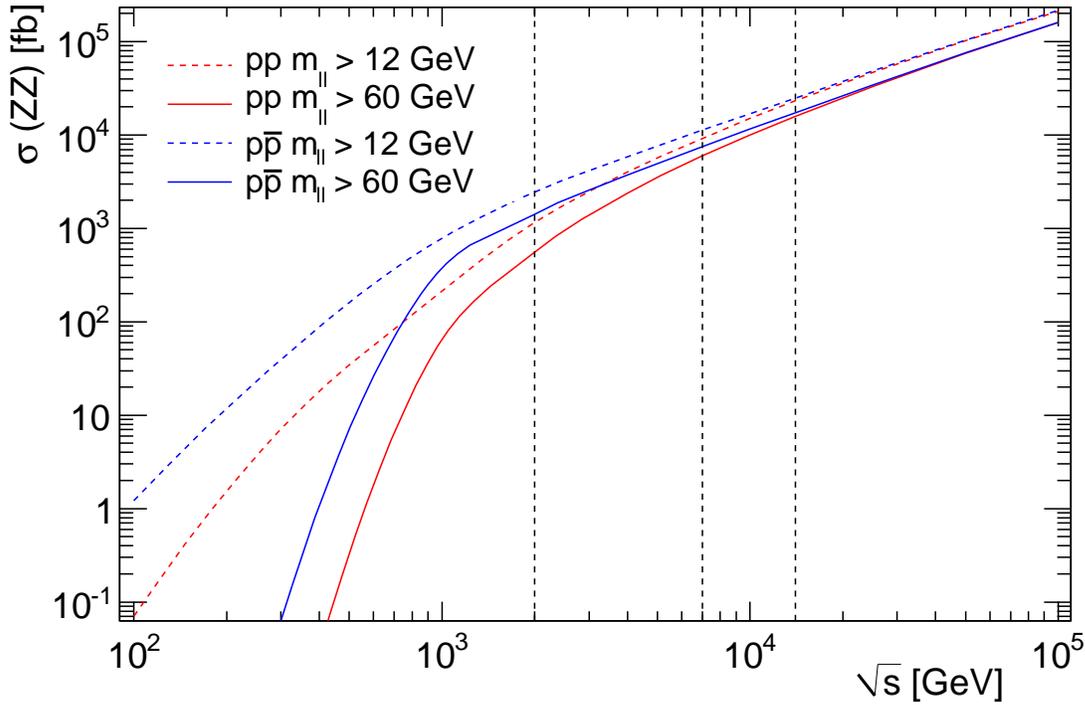


Figure 1.9: Next-to-leading order cross section for ZZ production in proton-(anti-)proton collisions as a function of centre-of-mass energy, using the MCFM generator [24, 25] with the MSTW2008 [30] PDF. The vertical lines are shown at 2, 7 and 14 TeV.

$\sqrt{s} = 2$ TeV is

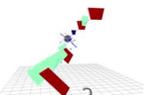
$$\frac{\sigma(pp \rightarrow ZZ, \sqrt{s} = 7 \text{ TeV})}{\sigma(p\bar{p} \rightarrow ZZ, \sqrt{s} = 2 \text{ TeV})} = 4.2. \quad (1.38)$$

1.6.1 Gluon-Gluon Fusion

It should be noted that the cross sections calculated above do not include the contribution from gluon-gluon fusion $gg \rightarrow ZZ$, as shown in figure 1.7. Although this process is suppressed by the square of the strong coupling, α_S^2 , the cross section is still considerable due to the high gluon content of the proton.

The LO cross section for gluon-gluon fusion was calculated at $\sqrt{s} = 7$ TeV using GG2ZZ, with a mass cut of $m_Z > 12$ GeV and the MSTW2008 PDF:

$$\sigma(gg \rightarrow ZZ) = 1.2_{-0.3}^{+0.4} (\text{scale}) \pm 0.3 (\text{PDF}) \text{ pb}, \quad (1.39)$$



corresponding to a contribution to the $pp \rightarrow ZZ$ cross section of

$$\frac{\sigma_{\text{LO}}(gg \rightarrow ZZ)}{\sigma_{\text{NLO}}(q\bar{q} \rightarrow ZZ)} = 13_{-4}^{+5}\%, \quad (1.40)$$

when comparing to equation 1.37. The error has been calculated from the combination of the scale and PDF errors in equation 1.39.

1.6.2 ZZ Decay Modes

Each Z boson can decay via a number of modes: quark-antiquark (70%), neutrino-antineutrino (20%) and oppositely charged leptons (10%). The combined branching fractions of the ZZ final state are shown in figure 1.10.

One set of decay modes studied in this thesis are those in which both Z bosons decay to a pair of same-flavour, oppositely-charged leptons, $ZZ \rightarrow ll\bar{l}$ ($l = e, \mu$). A second decay mode is considered, where one Z decays to oppositely charged leptons, and one decays to a neutrino-antineutrino pair. The branching fractions for these two channels are

$$\mathcal{B}(ZZ \rightarrow ll\bar{l}) = 0.452\%, \quad (1.41)$$

$$\mathcal{B}(ZZ \rightarrow ll\nu\bar{\nu}) = 2.69\%. \quad (1.42)$$

A method for identifying the $ZZ \rightarrow ll\bar{l}$ and $ZZ \rightarrow ll\nu\bar{\nu}$ channels in ATLAS is presented in chapter 6.

1.6.3 Backgrounds to ZZ Production

The most prominent backgrounds to ZZ production are those in which bosons are present in the final state. The impact of these backgrounds will be discussed further in chapter 6, but are included here for completeness. Figure 1.11 shows the LO Feynman diagrams for single Z and W production, with a charged lepton in the final state. Semi-leptonic top decay $t \rightarrow W(\rightarrow l\nu)b$ in the $t\bar{t}$ channel (figure 1.12) and $W^+W^- \rightarrow l\nu\bar{l}$ diboson production (figure 1.6, with $V_{1,2} = W^\pm$ and $V_3 = Z, \gamma$) both contain the same final state particles as the $ZZ \rightarrow ll\nu\bar{\nu}$ signal. WZ diboson production (figure 1.6, with $V_1 = W^\pm$, $V_2 = Z$ and $V_3 = W^\pm$) will also have the same signature as the $ZZ \rightarrow ll\nu\bar{\nu}$ channel if the lepton from the W decay is not measured by the detector. The $Zb\bar{b}$

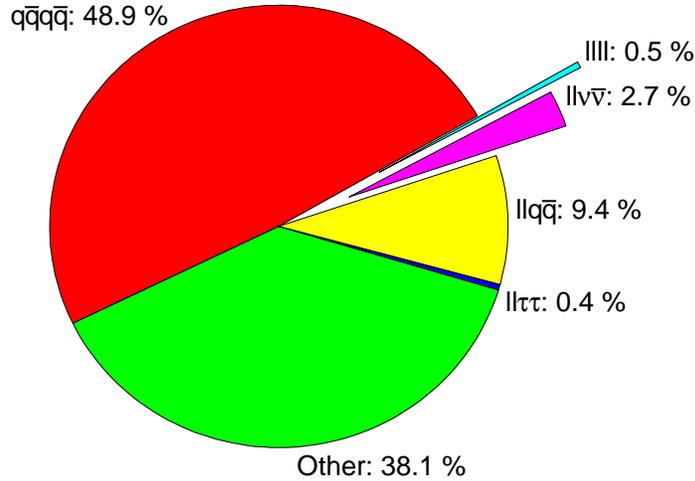


Figure 1.10: The ZZ final state branching fractions ($l = e, \mu$), calculated from [8].

channel, shown in figure 1.13, is a potential background to $ZZ \rightarrow llll$ if the b -quark decay products are reconstructed as leptons in the detector.

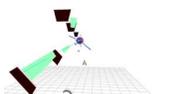
1.6.4 Anomalous Triple Gauge Couplings at the LHC

In addition to the two SM $q\bar{q} \rightarrow ZZ$ production diagrams in figure 1.6, anomalous interactions could also contribute to the cross section via the Feynman diagram in figure 1.6. From equations 1.27 and 1.29, the anomalous coupling contribution to the ZZ cross section is quadratic in the coupling,

$$\sigma(f_i^V) \propto (f_i^V)^2, \quad (1.43)$$

assuming a single non-zero coupling.

Figure 1.14 shows the ratio of the LO $q\bar{q} \rightarrow ZZ$ cross section with and without anomalous couplings included, as a function of centre-of-mass energy. Each coupling parameter is set in turn to $f_{i=4,5}^{V=Z,\gamma} = 0.1$, while the remaining parameters are fixed at zero. This values corresponds approximately to the current experimental limits, which



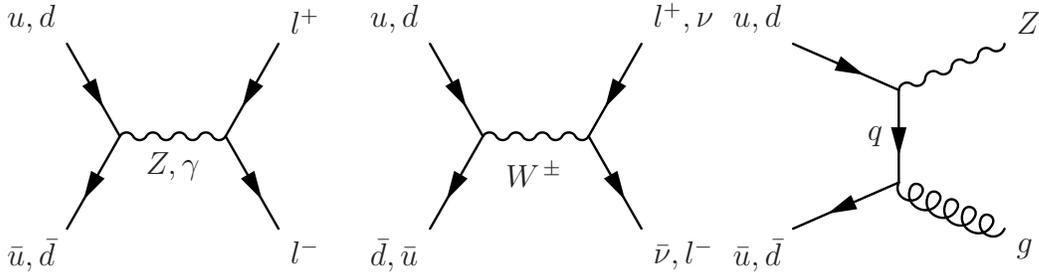


Figure 1.11: Leading order Feynman diagrams for leptonically decaying Z and W production. The right hand diagram shows an example of Z production in which a hard jet is also present.

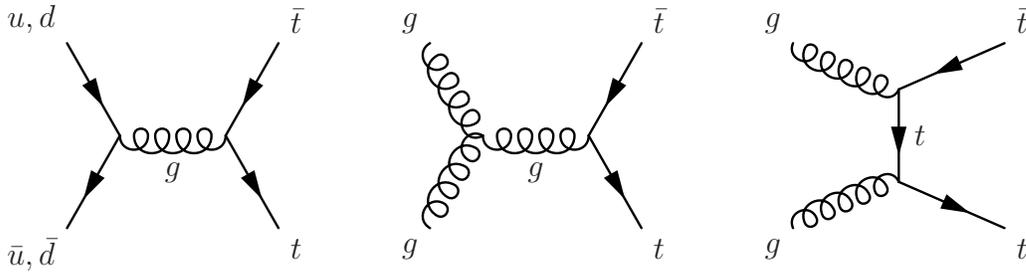


Figure 1.12: Leading order Feynman diagrams for $t\bar{t}$ production. An additional crossed process associated with the right-hand diagram has not been shown.

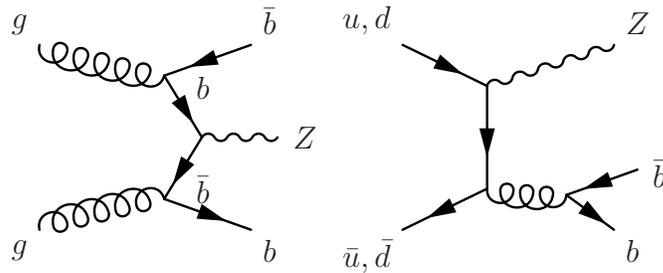


Figure 1.13: Leading order Feynman diagrams for $Zb\bar{b}$ production. The left-hand diagram provides the dominant contribution to the cross section.

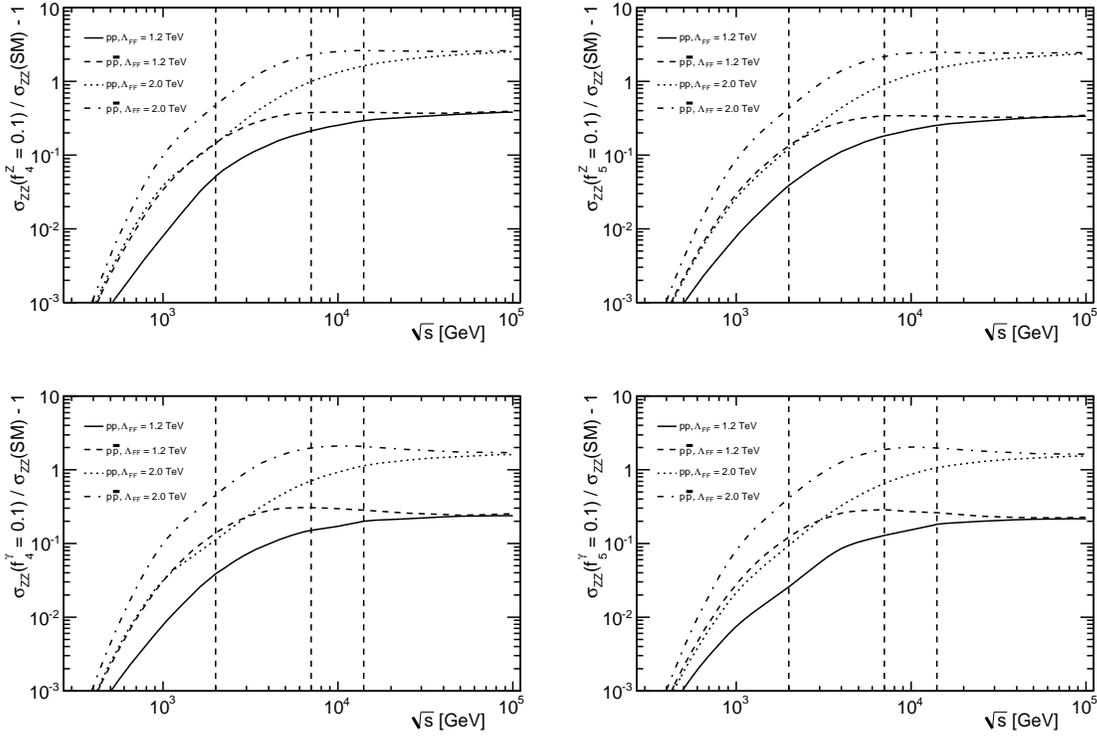
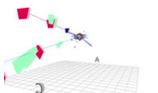


Figure 1.14: Ratio of inclusive ZZ cross section with anomalous couplings $f_{i=4,5}^{V=Z,\gamma} = 0.1$ to the SM, as calculated by the Baur and Rainwater Monte Carlo generator. Ratios are shown for $p\bar{p}$ and pp collisions, with form factor cutoffs of $\Lambda_{FF} = 1.2, 2.0$ TeV. All calculations use a form factor with $n = 3$.

will be discussed further in section 1.7. Cross sections are calculated with the BR Monte Carlo program, using the CTEQ6L PDFs. A form factor is applied as shown in equation 1.29, taking $\Lambda_{FF} = 1.2, 2$ TeV and $n = 3$.

At low values of $s \ll \Lambda_{FF}^2$, the cross section ratio increases with \sqrt{s} , as expected from equation 1.43. By design, at large $s \gg \Lambda_{FF}^2$, the form factor suppresses the anomalous coupling contribution to the cross section. For anomalous couplings $f_{i=4,5}^{V=Z,\gamma} = 0.1$ with $\Lambda_{FF}^2 = 1.2$ TeV, the increases in the ZZ cross section for pp collisions at $\sqrt{s} = 7$ TeV



are

$$\frac{\sigma_{ZZ}(f_4^Z = 0.1)}{\sigma_{ZZ}(\text{SM})} = 1.21 \quad (1.15), \quad (1.44)$$

$$\frac{\sigma_{ZZ}(f_5^Z = 0.1)}{\sigma_{ZZ}(\text{SM})} = 1.18 \quad (1.13), \quad (1.45)$$

$$\frac{\sigma_{ZZ}(f_4^\gamma = 0.1)}{\sigma_{ZZ}(\text{SM})} = 1.15 \quad (1.14), \quad (1.46)$$

$$\frac{\sigma_{ZZ}(f_5^\gamma = 0.1)}{\sigma_{ZZ}(\text{SM})} = 1.13 \quad (1.12), \quad (1.47)$$

where values in parentheses show the corresponding ratio for $p\bar{p}$ collisions with $\sqrt{s} = 2.0$ TeV.

The anomalous NTGCs also increase the ZZ cross section at high Z transverse momentum. This is exploited in chapter 7 to calculate the expected sensitivity of the ATLAS experiment to non-zero anomalous coupling parameters.

1.7 Current Experimental Results

The Z boson was first discovered in 1983 by the UA1 [33] and UA2 [34] experiments of the Super Proton Synchrotron (SPS) at The European Organisation for Nuclear Research (CERN). Precision measurements of Z properties have been made by experiments at the Large Electron-Positron (LEP) collider [35], an e^+e^- machine with centre of mass energies up to $\sqrt{s} = 200$ GeV.

Diboson searches and measurements of the relevant TGCs have been undertaken by experiments at two accelerators prior to the LHC. The first of these was at LEP, while the most recent results have come from the $D\bar{O}$ and Collider Detector at Fermilab (CDF) experiments at the Tevatron, a $p\bar{p}$ collider with $\sqrt{s} = 1.96$ TeV. The results of ZZ production studies from these experiments, together with results from previous ATLAS sensitivity studies, are described in detail in the following sections.

1.7.1 Results from LEP

Measurements of the ZZ production cross section from each of the four LEP experiments have been combined using the Best Linear Unbiased Estimate method, which takes into

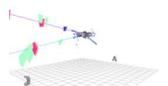
account systematic error correlations [36]. The cross section was measured over a range of energies across the threshold for ZZ production, as shown in figure 1.15a. At an energy of $\sqrt{s} = 200$ GeV, the combined cross section is measured to be $\sigma(e^+e^- \rightarrow ZZ) = 0.90 \pm 0.12$ pb.

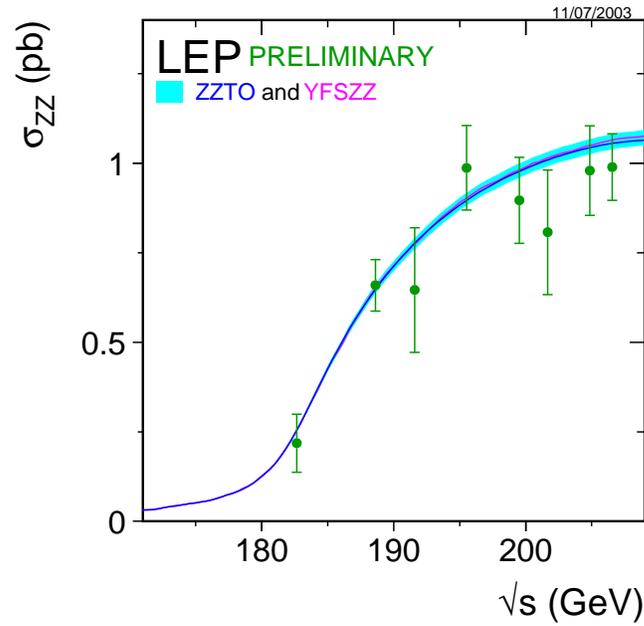
Confidence limits on the anomalous triple gauge coupling parameters have been obtained from each of the four LEP experiments [35], where all visible Z decay channels are used except for $\tau^+\tau^-q\bar{q}$, $\tau^+\tau^-\nu\bar{\nu}$ and $l^+l^-l^+l^-$. The neutral triple gauge boson couplings f_i^V are measured using an extended maximum likelihood fit of the Z production angle ($\cos\theta_Z$) distribution for the $q\bar{q}l^+l^-$ and $l^+l^-\nu\bar{\nu}$ channels, for which a cut-based method is used to select events. For the $q\bar{q}q\bar{q}$ channel, which represents 50% of the ZZ final states, the $\cos\theta_Z$ distribution is simultaneously fitted to the ZZ probability used to distinguish the signal from $WW \rightarrow q\bar{q}q\bar{q}$ background. A similar technique is applied to the $q\bar{q}\nu\bar{\nu}$ channel, where $\cos\theta_Z$ is fitted simultaneously to a combined discriminant variable.

Confidence limits were calculated for each coupling parameter individually, with the remaining couplings fixed to the SM value. The log-likelihood curves from each of the four LEP experiments are shown in figure 1.15b, together with a combined fit. The resulting combined 95% confidence limits are listed in table 1.1. No deviations from the expected SM values are observed. Small correlations were found between the f_4^V and f_5^V parameters when fitting pairs of couplings simultaneously.

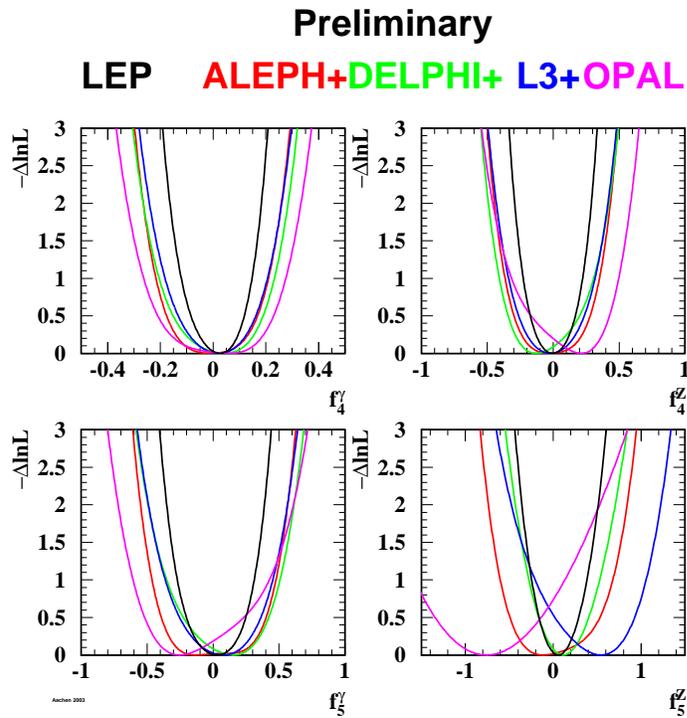
1.7.2 Results from the Tevatron

Evidence for ZZ production in a hadron collider experiment was first observed by two detectors at the Tevatron accelerator. The CDF Collaboration reported evidence for ZZ production in $p\bar{p}$ collisions at $\sqrt{s} = 1.96$ TeV, using a combination of $ZZ \rightarrow lll$ and $ZZ \rightarrow ll\nu\bar{\nu}$ channels ($l = e, \mu$) [37]. For an integrated luminosity of 1.9 fb^{-1} , three $ZZ \rightarrow lll$ events were observed, which in combination with a calculation of the relative ZZ and WW probabilities in the $ZZ \rightarrow ll\nu\bar{\nu}$ channel, yields an overall cross section of $\sigma = 1.4_{-0.6}^{+0.7}$ pb. The $ZZ \rightarrow lll$ analysis was recently updated to an integrated luminosity of 4.8 fb^{-1} [38], where a total of five events were observed, as shown in figure 1.16a. The updated cross section remains consistent with the earlier measurement, and is in good agreement with the SM prediction.





(a) ZZ production cross section as a function of \sqrt{s} , from [36].



(b) $-\Delta \log \mathcal{L}$ curves for the NTGC parameters, for each of the four LEP experiments (coloured) and a combined fit (black) [35].

Figure 1.15: Combined ZZ results from the LEP experiments.

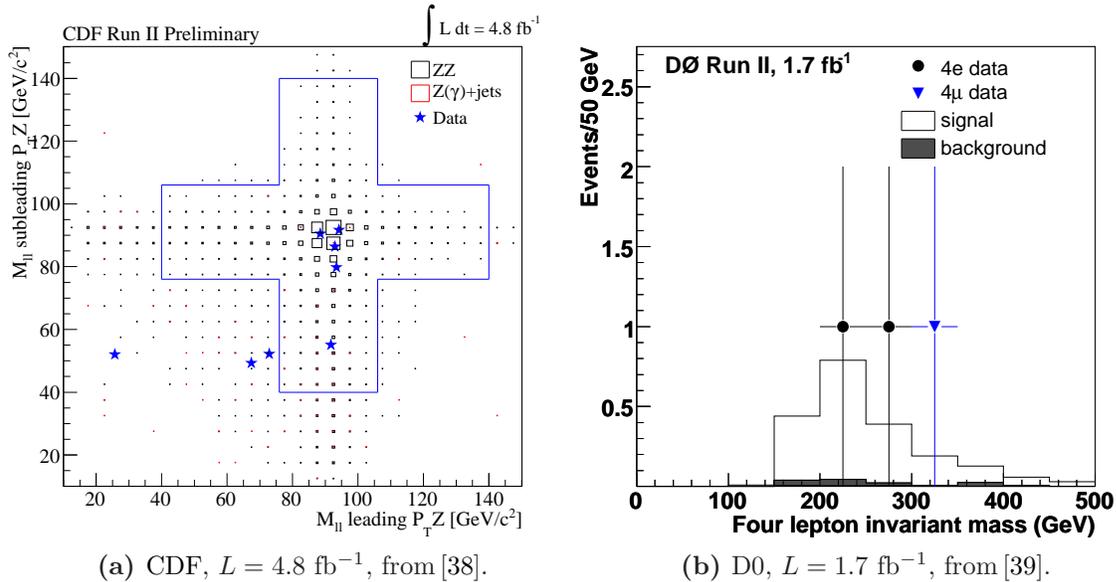
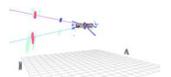


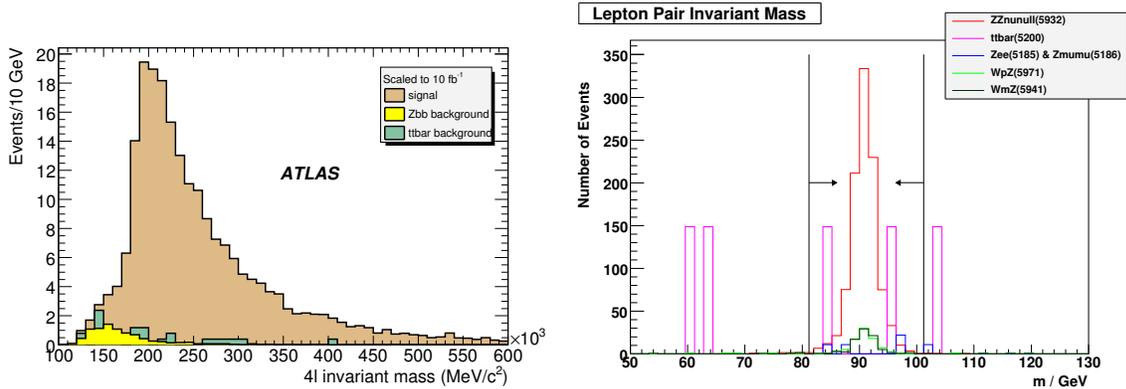
Figure 1.16: $ZZ \rightarrow lll$ events observed at the Tevatron.

The D0 experiment has also reported evidence for ZZ production. For 1.7 fb^{-1} of data, three events are observed in the $ZZ \rightarrow lll$ channel [39], reproduced in figure 1.16b. The combination of this result with an independent study of the $ZZ \rightarrow ll\nu\bar{\nu}$ channel using 2.7 fb^{-1} of integrated luminosity, yields a combined cross section of $\sigma = 1.60 \pm 0.63(\text{stat})_{-0.17}^{+0.16}(\text{sys}) \text{ pb}$.

Measurements have also been made at the Tevatron experiments of the anomalous triple gauge couplings. The CDF collaboration have used 1.9 fb^{-1} of integrated luminosity in the $ZZ \rightarrow llq\bar{q}$ channel [40]. 95% confidence limits on each of the four couplings are calculated by fitting the dijet invariant mass spectrum, using a form factor with $\Lambda_{FF} = 1.2 \text{ TeV}$ and $n = 3$. The resulting limits are reproduced in table 1.1, and show a good agreement with the expected sensitivity from Monte Carlo studies.

The D0 experiment has also derived limits on the anomalous couplings with 1 fb^{-1} of integrated luminosity in the $ZZ \rightarrow lll$ channel. Confidence limits are calculated by forming a likelihood from the expected number of events for models with non-zero coupling values. The resulting one-parameter 95% confidence limits are shown in table 1.1, again with form factor parameters $\Lambda_{FF} = 1.2 \text{ TeV}$ and $n = 3$.





(a) $ZZ \rightarrow llll$ four lepton invariant mass normalised to 10 fb^{-1} . (b) $ZZ \rightarrow ll\nu\bar{\nu}$ Di-lepton invariant mass normalised to 100 fb^{-1} .

Figure 1.17: Simulated ATLAS diboson events after cuts at $\sqrt{s} = 14 \text{ TeV}$, from [43, 44].

1.7.3 ATLAS Sensitivity Studies

A number of studies have been performed to predict the expected yields of ZZ events in the ATLAS detector and to estimate the sensitivity to non-zero NTGC parameters.

Early investigations using a fast simulation of the ATLAS detector [41] estimated that the LEP limits could be improved by a factor of $\mathcal{O}(10^3)$ after an integrated luminosity of 100 fb^{-1} at an energy of $\sqrt{s} = 14 \text{ TeV}$ [42].

A more detailed study at $\sqrt{s} = 14 \text{ TeV}$ used simulated events with a full description of the ATLAS detector to define a set of cuts to select the $ZZ \rightarrow llll$ and $ZZ \rightarrow ll\nu\bar{\nu}$ channels [43, 44]. After 1 fb^{-1} of integrated luminosity, ATLAS was expected to observe 17.0 ± 0.5 events in the $ZZ \rightarrow llll$ channel, with a total background of 2.0 ± 0.2 , corresponding to a significance of 6.8 standard deviations. Figure 1.17a shows the expected four-lepton invariant mass distribution for events after selection cuts. For the same integrated luminosity with the $ZZ \rightarrow ll\nu\bar{\nu}$ channel, 10 ± 0.2 signal events were expected, with a background of 5 ± 2 . The dilepton invariant mass distribution expected after cuts is shown in figure 1.17b.

The expected sensitivity of ATLAS to anomalous triple gauge couplings at $\sqrt{s} = 14 \text{ TeV}$ is also presented in [43, 44]. In this study the shape of the reconstructed Z transverse momentum distribution is used to perform a binned fit to both the $ZZ \rightarrow llll$ and $ZZ \rightarrow ll\nu\bar{\nu}$ channels. Single parameter sensitivities are predicted by fitting a large number of fake Standard Model distributions using the two channels separately and a

Measurement	f_4^Z	f_5^Z	f_4^γ	f_5^γ
LEP Combined [35]	[-0.31:0.29]	[-0.19:0.20]	[-0.17:0.16]	[-0.36:0.40]
CDF (1.9 fb ⁻¹) [40]	[-0.12:0.12]	[-0.13:0.12]	[-0.10:0.10]	[-0.11:0.11]
DØ (1.0 fb ⁻¹) [46]	[-0.28:0.28]	[-0.31:0.29]	[-0.26:0.26]	[-0.30:0.30]
ATLAS [44] 1.0 fb ⁻¹	[-0.023:0.023]	[-0.024:0.024]	[-0.028:0.028]	[-0.029:0.028]
10.0 fb ⁻¹	[-0.009:0.009]	[-0.009:0.009]	[-0.010:0.010]	[-0.011,0.010]

Table 1.1: Current 95% confidence limits on anomalous triple gauge couplings. The ATLAS limits represent the expected sensitivity at $\sqrt{s} = 14$ TeV.

combined fit, with $\Lambda_{FF} = 2$ TeV. Expected 95% confidence limits using combined fits after 1 fb⁻¹ are presented in table 1.1. With this integrated luminosity, it was expected that ATLAS could improve the LEP limits by an order of magnitude.

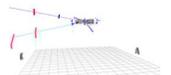
The main aim of chapter 6 is to define a set of cuts which can be used to select $ZZ \rightarrow lll$ and $ZZ \rightarrow ll\nu\bar{\nu}$ events in ATLAS, but with improved detector modelling, and at the 2010 LHC energy of $\sqrt{s} = 7$ TeV. This study draws extensively on experience gained from these early ATLAS studies, with many common cuts. Chapter 7 updates and refines the $\sqrt{s} = 14$ TeV ATLAS anomalous coupling sensitivity studies, using similar techniques.

Diboson channels containing τ leptons in the final state have been considered in ATLAS [45] at $\sqrt{s} = 14$ TeV. For an integrated luminosity of 10 fb⁻¹, 9.9 signal events are expected from the $ZZ \rightarrow \tau\tau\mu\mu$ channel. The level of background contamination was found to be below 12% for most types of τ decay. The $llq\bar{q}$ channel is another possibility for study, but a high background from $Z \rightarrow ll + \text{jets}$ excludes it from early ATLAS analysis.

1.7.4 Diboson Background to New Physics

SM diboson production is also the background to a number of other new physics searches. For example, the $ZZ \rightarrow lll$ channel is the main background to Higgs boson production via $H \rightarrow ZZ^{(*)} \rightarrow lll$ (see chapter 12 of [43]). In this regard, observation of SM diboson production will be an important prerequisite to any Higgs searches in this channel.

SM diboson production is also the background to a number of more exotic new physics models. For example, the $ZZ \rightarrow ll\nu\bar{\nu}$ channel forms the background to graviton



production. In [47], the expected yields from [44] are used to predict the sensitivity of ATLAS to events in which a graviton is produced in conjunction with a Z boson.

1.8 Conclusions

This chapter has given a brief introduction to electro-weak field theory and how self-couplings between electro-weak gauge bosons arise in the SM. Interactions between three neutrally charged bosons ($ZZZ, \gamma ZZ$) are forbidden in the SM at tree level, but can be introduced in the anomalous coupling formalism, parametrised by the couplings $f_{i=4,5}^{V=Z,\gamma}$. The simplest way of introducing such couplings is via heavy-fermion triangle diagrams, which are predicted in NP models such as the MSSM.

The next-to-leading order cross section for on-shell SM ZZ production for pp collisions at $\sqrt{s} = 7$ TeV was calculated in equation 1.36, and is ~ 4 times larger than at the Tevatron. Setting the anomalous coupling parameters to $f_i^V = 0.1$ in turn results in an enhancement in the SM cross section by a factor 13–21% at this energy.

An overview of ZZ cross section measurements and anomalous coupling limits from LEP and the Tevatron has been given, putting into context the event selection and sensitivity study of chapters 6 and 7.

Chapter 2

The ATLAS Experiment

*“Let’s test superstring theory,
Oh yoi yoi accelerate the protons,
stir it twice and then just add me!”*

— Gogol Bordello, Supertheory of supereverything.

2.1 The Large Hadron Collider

The Large Hadron Collider (LHC) [48] at CERN will expand the horizons of particle physics at unprecedented high energy and luminosity. It is designed to collide bunches of up to 10^{11} protons 40 million times every second at a centre of mass energy of $\sqrt{s} = 14$ TeV. It is anticipated that this energy will be reached after an initial year long run at 7 TeV, after which additional safety features will be installed to allow higher energy operation. The nominal LHC beam luminosity is 10^{34} $\text{cm}^{-2}\text{s}^{-1}$, which provides an annual integrated luminosity of ~ 100 fb^{-1} .

The proton beams used in the LHC are produced in the CERN Linear Accelerator (LINAC) at energies of 50 MeV, before being transferred first to the Proton Synchrotron (PS) and then the SPS where the beams reach 450 GeV in energy. They are then injected into the LHC itself, which is situated up to 80 m underground in the 27 km circular tunnel that was built to house the LEP collider. In total, 1300 superconducting dipole magnets are required to steer protons around the ring, in conjunction with 392 quadrupoles used for focusing. The peak magnetic field is 8.3 T, achieved by cooling the magnets to 1.9 K

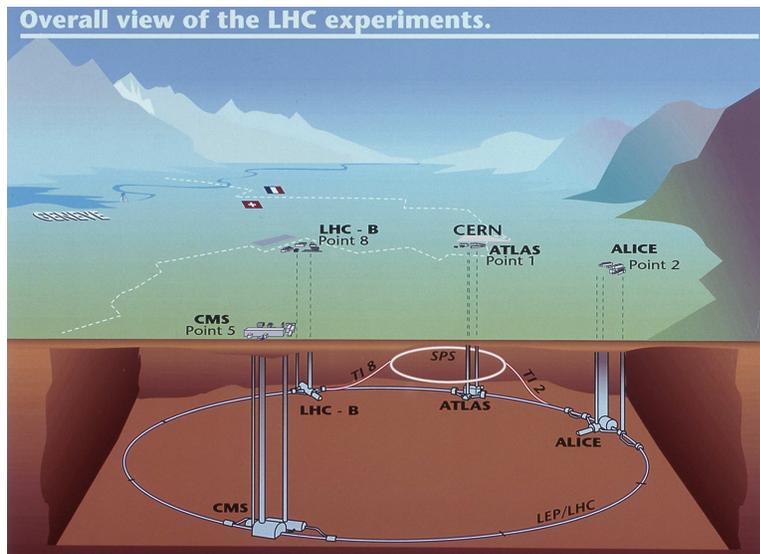


Figure 2.1: A cartoon showing the layout of the LHC underground areas.

with a total of 96 tonnes of super-fluid helium. The total energy stored in the beam and magnets is ~ 1 GJ, equivalent to a 400 tonne *Train à Grande Vitesse* travelling at 300 kmh^{-1} .

The beams intersect at four different points around the ring, where they provide collisions for the LHC detectors, as shown in figure 2.1. ATLAS [49] and the Compact Muon Solenoid (CMS) [50] are two general purpose, high-luminosity experiments designed to be sensitive to a broad range of new physics scenarios. LHCb [51] is a dedicated B -physics experiment designed to investigate CP violation, at reduced luminosity. The LHC also has a dedicated ion physics program, and is expected to provide one month of lead-lead collisions per year, for which the ALICE [52] experiment has been purposely built. In addition, the LHC is also home to two forward detectors, LHCf [53] and TOTEM [54].

2.1.1 Context

First LHC injection tests were performed in September 2008, with the intention of accelerating the proton beams to $\sqrt{s} = 14$ TeV. These plans were dramatically cut short two weeks after first injection, when a faulty connection between two magnets caused a catastrophic explosion inside the tunnels. Repair and replacement of damaged magnets meant a delay in LHC start-up until November 2009, when first collisions were achieved at injection energy. The first high-energy collisions at the LHC were delivered in March 2010, at a record energy of $\sqrt{s} = 7$ TeV.

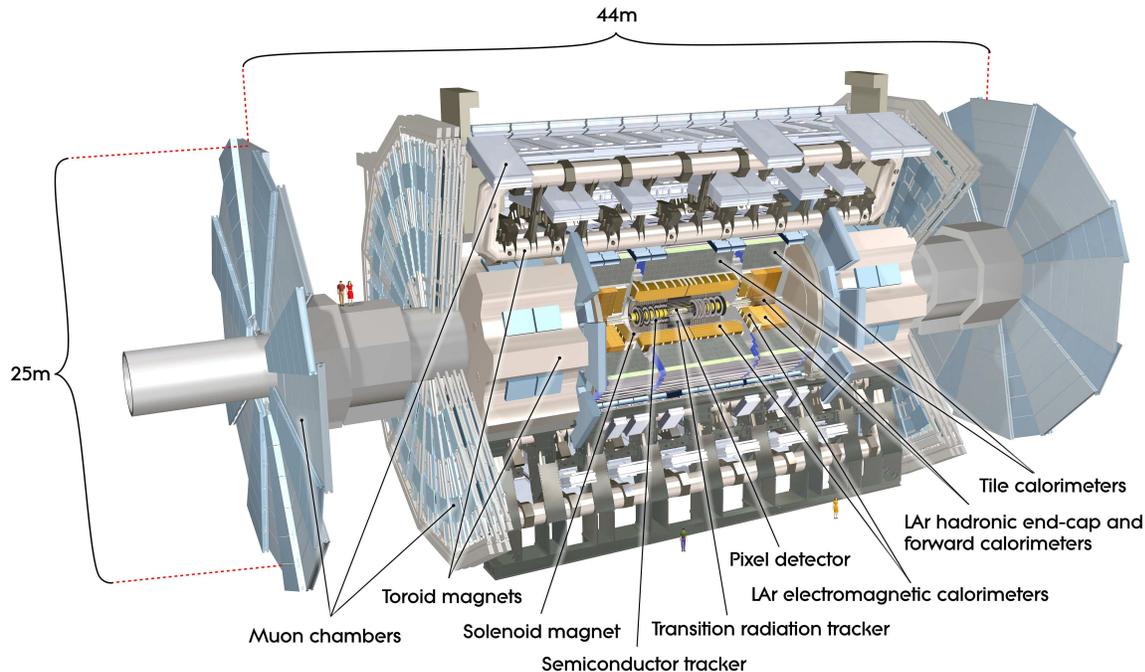
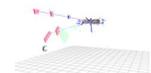


Figure 2.2: A diagram of the ATLAS Detector, from [49].

At the time of writing, it is foreseen that the LHC will continue $\sqrt{s} = 7$ TeV collisions until the end of 2011. During this time it is expected that the LHC experiments will record a total integrated luminosity of $\sim 1 \text{ fb}^{-1}$. This period will be preceded by a year-long shutdown to install the additional safety measures required for 14 TeV operation.

2.2 The ATLAS Experiment

ATLAS is an experiment of superlatives: it is the largest experiment ever built, probing matter at the highest energies, with the most dubious acronym, A Toroidal LHC Apparatus. The ATLAS detector is 44 m long and 25 m high (see figure 2.2) and weighs 7 kT. Construction was completed in 2008 when the final pieces were lowered into the underground experimental hall at Point-1, putting the finishing touches to the world's largest and most intricate ship in a bottle.



2.2.1 Co-ordinate Definition

The co-ordinate system used in ATLAS is defined by the beam direction, which lies along the z -axis. The $x-y$ plane lies transverse to the beam direction, with the positive x -axis pointing from the interaction point towards the centre of the LHC ring, and the y -axis pointing upwards. The azimuthal angle ϕ is measured around the beam axis, and the polar angle θ is the angle from the beam axis. The pseudorapidity is defined as

$$\eta = -\ln \left(\tan \left(\frac{\theta}{2} \right) \right). \quad (2.1)$$

The transverse momentum p_T , transverse energy E_T and missing transverse energy \cancel{E}_T are defined in the $x-y$ plane. The distance ΔR between two directions in pseudorapidity-azimuthal space is defined as

$$\Delta R = \sqrt{\Delta\phi^2 + \Delta\eta^2}, \quad (2.2)$$

which is independent of Lorentz boosts along the beam axis.

2.2.2 Prospects for New Discovery

The design of the ATLAS detector has been driven by the discovery potential of a number of new phenomena at the TeV energy scale. The most prominent of these is the search for the Higgs boson, which has been used as a benchmark in establishing the performance of the ATLAS subsystems. There are a number of promising channels for first Higgs discovery at low masses ($m_H < 2m_Z$), such as $t\bar{t}H$, WH and ZH , with $H \rightarrow b\bar{b}$. $H \rightarrow \gamma\gamma$ is also an important channel for low mass Higgs production. The higher mass range, $m_H > 130$ GeV is more relevant to this thesis, where $H \rightarrow ZZ^{(*)}$ with each Z decaying to charged leptons. SM ZZ observation in ATLAS, as discussed in chapter 6, will be an important prerequisite to Higgs discovery in this channel.

One extension of the SM by anomalous NTGCs was introduced in chapter 1. Other new physics models that could be accessible with ATLAS include searches for heavy gauge bosons W' and Z' , up to masses of ~ 6 TeV. Supersymmetry (SUSY) is another popular model to resolve the hierarchy problem in the SM. Experimentally, Supersymmetric cascade decays will always contain the Lightest Supersymmetric Particle (LSP). In R -parity conserving models, the LSP interacts weakly with the detector, measured

as significant missing transverse energy, accompanied by leptons and jets. Alternative theories to SUSY include the existence of extra dimensions and miniature black-hole production, both of which have the potential to be observed at the LHC.

2.2.3 Experimental Requirements

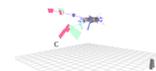
The formidable LHC luminosity is required to produce a statistically significant number of observed events, as most of the above processes have small cross-sections (see equation 1.33). However, the cross section for inelastic scattering of protons at the LHC is 80 mb, corresponding to a total rate of 10^9 inelastic events every second at design luminosity. At this rate, each candidate event for new physics will be accompanied by an average of 23 inelastic events per bunch crossing. This imposes a significant experimental challenge when designing an LHC detector.

The benchmark physics goals described above can be turned into a number of general requirements for the design of the ATLAS detector. These include fast, radiation-hard, electronics and sensors with high granularity. The detector should have large acceptance in pseudorapidity and good charged particle momentum resolution and reconstruction efficiency in the inner tracker. The electromagnetic calorimetry should have good electron and photon identification, complemented with full-coverage hadronic calorimeters for jet and missing transverse energy measurements. In addition, ATLAS should have good muon identification and momentum resolution, including the ability to determine the charge of high p_T muons. The trigger system should be able to achieve an acceptable rate for most physics processes of interest.

A summary of the performance requirements of the ATLAS detector is given in table 2.1.

2.3 The ATLAS Detector

The ATLAS detector is forward-backward symmetric with respect to the interaction point and consists of the inner detector for tracking and momentum measurements, electromagnetic and hadronic calorimeter systems, and a large muon spectrometer. The magnet system consists of a 2 T solenoidal field for the inner detector and a set of three



Detector Subsystem	Required Resolution	η Coverage
Tracking	$\sigma_{p_T}/p_T = 0.05\%p_T \oplus 1\%$	± 2.5
EM calorimetry	$\sigma_E/E = 10\%/\sqrt{E} \oplus 0.7\%$	± 3.2
Hadronic calorimetry barrel and end-cap forward	$\sigma_E/E = 50\%/\sqrt{E} \oplus 3\%$ $\sigma_E/E = 100\%/\sqrt{E} \oplus 10\%$	± 3.2 $3.1 < \eta < 4.9$
Muon spectrometer	$\sigma_{p_T}/p_T = 10\%$ at 1 TeV	± 2.7

Table 2.1: Performance goals of the ATLAS detector, from [49].

superconducting toroids providing bending for muon tracks. The individual subsystems will be briefly summarised in the following sections.

2.3.1 The Inner Detector

High precision spacepoint measurement is essential in ATLAS for track reconstruction and particle momentum measurements. In ATLAS this task is performed by the Inner Detector (ID): the Pixel subdetector, the Semiconductor Tracker (SCT) and the Transition Radiation Tracker (TRT). The Pixels lie closest to the interaction point and allow secondary vertex reconstruction. The TRT is furthest from the primary vertex and contributes around 36 track measurements based on straw tubes. The SCT is a silicon strip tracking detector housed between the Pixels and TRT. Table 2.2 shows the intrinsic accuracy of each of the ID subsystems, based on the requirements given in table 2.1

Figure 2.3 shows how the three subsystems are combined to make up the Inner Detector. The ID is contained within a 7 m long cylinder with radius 1.15 m, covering a pseudorapidity range of $|\eta| < 2.5$. A central solenoid magnet provides a 2 T magnetic field for the inner detector, and lies within the ID volume itself.

The commissioning and initial operation of the Semiconductor tracker will be described in more detail in chapter 3.

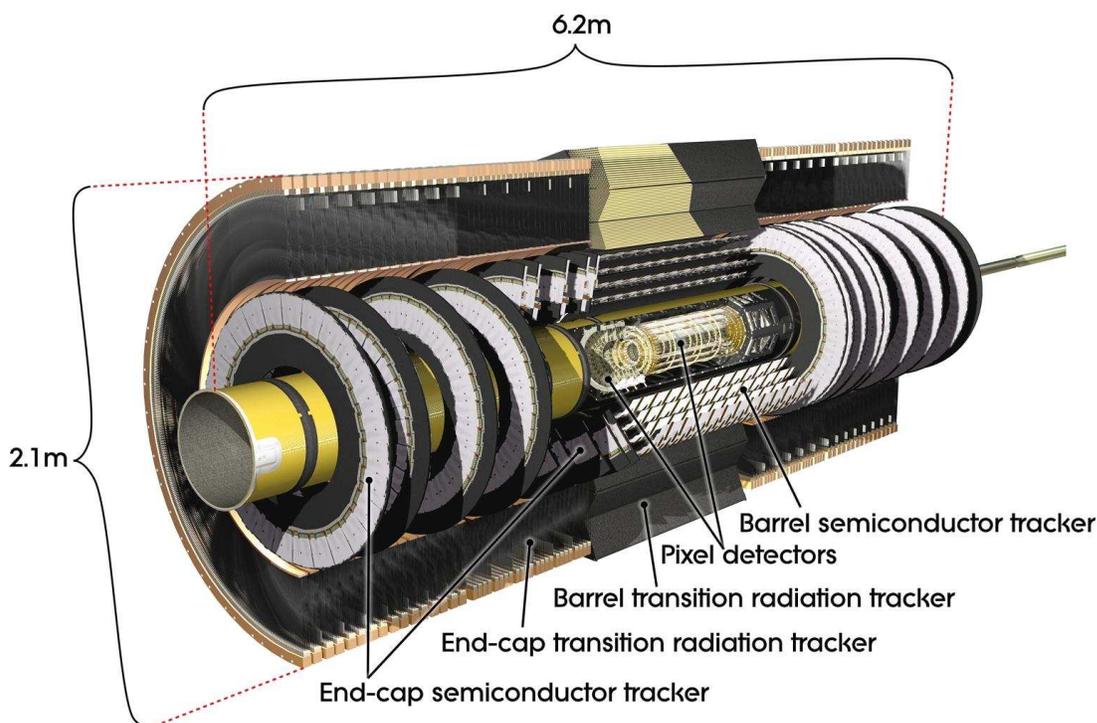


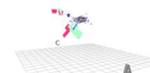
Figure 2.3: A diagram of the ATLAS Inner Detector, showing the structure of the pixel detector, SCT and TRT. Taken from [49].

Subsystem	Accuracy / μm
Pixel	
Barrel	10 ($R - \phi$) 115 (z)
Disks	10 ($R - \phi$) 115 (R)
SCT	
Barrel	17 ($R - \phi$) 580 (z)
end-caps	17 ($R - \phi$) 580 (R)
TRT	130

Table 2.2: Intrinsic measurement accuracies of the Inner Detector subsystems, from [49].

2.3.2 Calorimeters

The ATLAS calorimeters, shown in figure 2.4, cover the range $|\eta| < 4.9$, and are constructed using a number of different technologies to suit varying requirements in the large η -range. In general, the ATLAS calorimeters must provide good containment of electromagnetic and hadronic showers, and limit punch-through to the muon system.



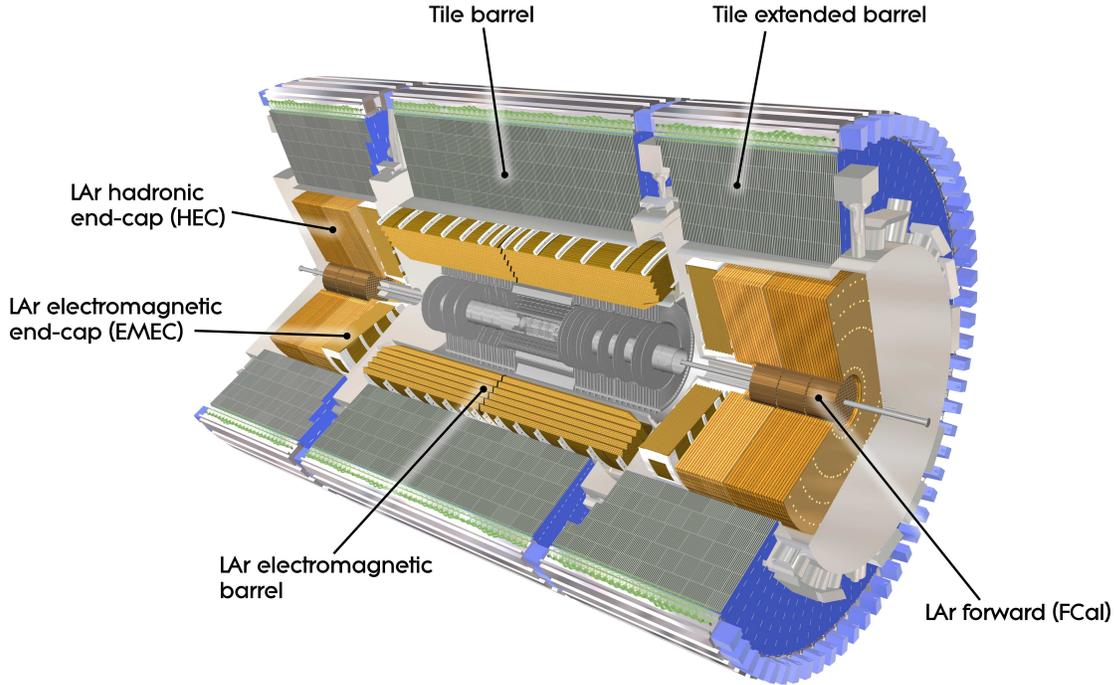


Figure 2.4: A diagram of the ATLAS calorimeter system, from [49].

The total thickness of the calorimeters at $\eta = 0$ is 11 interaction lengths, which simulations have shown to be sufficient to reduce punch-through well below the rate of muon production.

The electromagnetic (EM) calorimeter is designed to match the η coverage of the inner detector, with fine granularity for precision measurements of electrons and photons. It is constructed using a combination of lead material which acts as an absorber and Liquid Argon (LAr) as the active detector material. The lead plates are arranged in an accordion shaped geometry which provides complete ϕ symmetry with no azimuthal cracks. In the precision region of $|\eta| < 2.5$, the EM calorimeter is divided into a barrel region and two end-cap wheels, segmented into three sections in depth. In the region $2.5 < |\eta| < 3.2$, the inner end-cap wheel is segmented into just two sections in depth.

Hadronic calorimetry is provided by three separate sub-detectors. The tile calorimeter is situated directly outside of the EM calorimeter envelope, consisting of a barrel and two extended barrel regions, with an overall coverage out to $|\eta| < 1.7$. It is a sampling calorimeter which uses steel as the absorber and scintillating tiles as the active material. The Hadronic End-cap Calorimeter (HEC) uses LAr scintillator, consisting

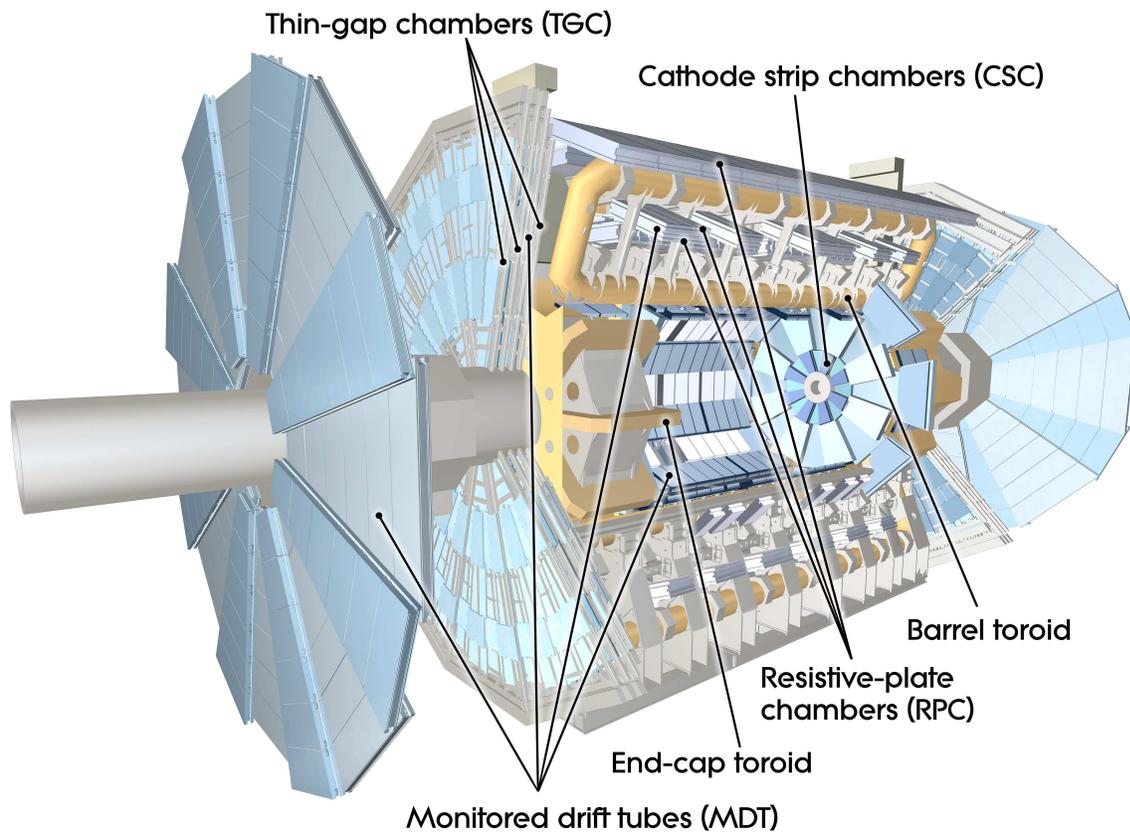


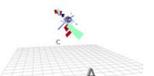
Figure 2.5: A diagram of the ATLAS muon spectrometer, from [49].

of two independent wheels per end-cap, and extends out to $|\eta| < 3.2$. Finally, the LAr Forward Calorimeter (FCal) comprises three sections in each end-cap, with one made of copper for electromagnetic measurements, and two of tungsten for hadronic energy determination. The Forward Calorimeter covers the range $3.1 < |\eta| < 4.9$.

2.3.3 Muon Spectrometer

The muon system is shown in figure 2.5, and consists of large superconducting air-core toroid magnets, which provide deflection for muon tracks, instrumented by separate high-precision and trigger chambers. In total there are three large magnets, a central barrel system comprising eight racetrack shaped coils, with two end-cap toroids at either end.

Precision muon track coordinates in the principle bending direction of the magnetic field are provided by Monitored Drift Tubes (MDTs) in the central η region, and Cathode



Strip Chambers (CSCs) from $2 < |\eta| < 2.7$. The trigger system covers the range $|\eta| < 2.4$ and uses Resistive Plate Chambers (RPCs) in the barrel and Thin Gap Chambers (TGCs) in the end-cap regions.

The accuracy of muon momentum measurements requires a $30 \mu\text{m}$ alignment precision, which is provided by 12000 sensors mounted on the MDT chambers. The magnetic field strength in the spectrometer volume is measured by a total of 1800 Hall sensors.

2.3.4 Forward Detectors

In addition to the main sub-detectors, ATLAS also includes three forward systems. The LUCID detector, situated at ± 17 m from the interaction point, provides online luminosity measurements using Cherenkov imaging. ALFA is located at ± 240 m, supplying luminosity information via Roman pot detectors designed to approach within 1 mm of the beam. The Zero-Degree Calorimeter (ZDC), at ± 140 m, is designed to determine the centrality of heavy-ion collisions.

2.4 Trigger, Data Acquisition and Control Systems

The ATLAS trigger system is designed to reduce the recorded event rate from a maximum of 40 MHz to 200 Hz, and is divided into three distinct levels. The first, L1, is a hardware based trigger which searches for high transverse momentum muons, electrons, photons, jets and large missing transverse energy. Trigger information is derived from both the muon trigger chambers and reduced granularity calorimeter cells, and can be pre-scaled to lower rates if luminosity and background conditions change. The L1 decision is made in less than $2.5 \mu\text{s}$, reducing the data rate to 75 kHz.

For each triggered event, a number of Regions of Interest (RoIs) are defined where interesting features have been identified. This information seeds the software-based L2 trigger decision, reducing the rate to ~ 3.5 kHz. The final stage, the event filter, is based on offline analysis techniques and reduces the final event rate to ~ 200 Hz.

After an event has been accepted by the Level-1 trigger (L1A), the Data Acquisition (DAQ) system receives and buffers data from on-detector pipelines. These are received by sub-detector Readout Drivers (RODs), which buffer and derandomise the event data before sending them on to Readout Subsystem (ROS). Events subsequently selected

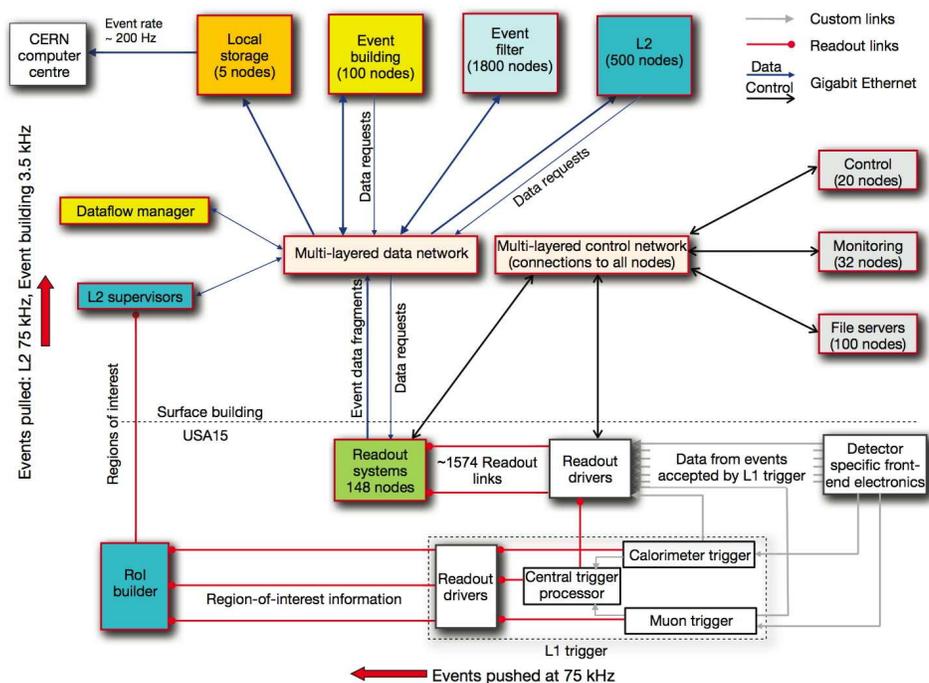
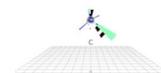


Figure 2.6: Block diagram of the ATLAS TDAQ subsystems, from [49].

by L2 triggers are transferred on to the event-building system. Finally, data passing the event filter are moved to the CERN computer centre for permanent storage and distribution via the LHC Computing Grid (LCG) [55]. An overview of the ATLAS Trigger and DAQ (TDAQ) systems is shown in figure 2.6.

The Detector Control System (DCS) permits operation of the ATLAS hardware with a homogeneous interface to all sub-detector components. These include high- and low-voltage power supplies, and environmental conditions such as temperature and humidity monitoring. Communication between the DAQ and DCS systems is also possible to coordinate data-taking depending on experimental conditions. The state of the DCS is also available for use offline, allowing physicists to require particular detector states for their analysis.



Chapter 3

The ATLAS Semiconductor Tracker

*“You just keep trying until you run out of cake,
and science gets done.”*

— Jonathan Coulton, Still Alive

3.1 Introduction

The ATLAS Semiconductor Tracker (SCT) is a silicon micro-strip detector designed to provide four high-resolution spacepoints for ATLAS track reconstruction. It is part of the ATLAS ID, as described in section 2.3.1.

The purpose of this chapter is to introduce the design and operation of the SCT in sufficient detail to understand the developments to the DAQ system described in chapters 4 and 5. The author was present at CERN during a large fraction of the commissioning and early operation periods of the SCT, and contributed via shifts both in the ATLAS Control Room (ACR) and as an on-call DAQ expert.

A more detailed description can be found in the SCT section of the ATLAS detector paper, chapter 4 of [49]. An excellent introduction and summary of the SCT DAQ can also be found in [56], upon which this chapter is based.

3.2 Silicon Detectors

An introduction to semiconductors and their properties can be found, for example, in [57]. A brief introduction will be given here in the context of their use as particle detectors.

3.2.1 Physics of Semiconductors

In the atomic crystal lattice structure of a semiconductor, the discrete energy levels of individual electrons merge to form continuous energy bands. In pure silicon, two distinct bands are formed, known as the valence band, which is completely filled, and the conduction band. The two are separated by an energy gap of around 1 eV. Thermal excitations across this gap are possible, forming two types of charge carrier: the promoted electron, and the “hole” of positive charge that it leaves behind in the valence band.

It is possible to control the electronic properties of intrinsic semiconductors by introducing impurities into the structure, a process known as doping. Addition of elements from group V of the periodic table, such as phosphorous, leave an unbonded electron which populates a state near the bottom of the conduction band. These atoms are therefore known as donors, and the resulting material is called n-type. On the other hand, acceptor atoms from group III, such as aluminium, will create a surplus of holes near the top of the valence band. This material is known as p-type.

3.2.2 *p-n* Junctions

When p-type and n-type materials are brought together, electrons will flow from n to p until electrochemical equilibrium is reached. This will leave a depletion region around the junction almost completely devoid of charge carriers. The width of the depletion region can be found by solving the one-dimensional Poisson equation:

$$\frac{\partial^2 V}{\partial x^2} = -\frac{\rho}{\epsilon_0 \epsilon_{\text{Si}}}, \quad (3.1)$$

where ρ is the charge density, ϵ_0 is the permittivity of free space and ϵ_{Si} is the relative permittivity of silicon. Assuming that the edges of the depletion region are sharp, and

integrating equation 3.1 twice leads to the relation:

$$d \approx \sqrt{\frac{2\epsilon_0\epsilon_{\text{Si}}V_{\text{bias}}}{eN_{\text{eff}}}}, \quad (3.2)$$

where $N_{\text{eff}} = \rho/e$ is the effective dopant concentration and V_{bias} is the applied bias voltage across the junction.

3.2.3 Semiconductors as Particle Detectors

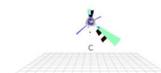
A p - n junction can be reverse biased by ensuring that the voltage applied at the cathode (n-type) is higher than that at the anode (p-type). In this way, holes in the p-type material and electrons in the n-type material are pulled in opposite directions away from the junction, and by equation 3.2, the depletion width will increase.

Such a junction forms the basis for semiconductor particle detectors, where it is common to increase the bias voltage such that the depletion region extends throughout the entire depth of material. In this fully depleted mode, current flow across the junction is minimised, although some leakage current is possible as thermally generated electron-hole pairs can migrate across.

When a charged particle traverses such a reverse-biased junction, the energy deposited by the particle creates a series of electron-hole pairs which migrate due to the electric field provided by the bias voltage. The induced charge is collected on the electrodes, and can be measured with supporting electronics. Silicon is a popular choice of semiconductor used in the construction of particle detectors, as it is radiation hard, and used technology that is highly developed by the microelectronics industry.

3.2.4 Radiation Damage

Detectors at the LHC have to be robust to damage caused by unprecedented levels of radiation. Over 10 years of operation, the maximum expected radiation dose is equivalent to 1.4×10^{14} 1 MeV neutrons cm^{-2} . Damage is primarily caused by interactions with particles which may leave a nucleus permanently displaced from its lattice position. This displacement will leave behind it a vacant site, which can migrate through the crystal lattice, forming stable defect complexes. These complexes can behave as



either donors or acceptors, which, over time, can cause the material to undergo a type inversion, changing from n- to p-type or vice versa.

These defects can also lead to increases in the leakage current, which increases the noise of the system, and loss of mobility, which reduces the charge collection efficiency. Both of these effects lead to a reduction in the signal to noise ratio.

In the short term (~ 2 days) after irradiation, thermal annealing actually produces beneficial effects by reducing the number of acceptor-like defects. On longer timescales, however, an effect known as *reverse-annealing* has a detrimental effect on the number of charge carriers [58]. Both of these effects are dependant on temperature, with time constant $\tau \propto \exp E_a/k_B T$, where E_a is the activation energy, T is the temperature and k_B is Boltzmann's constant. By maintaining silicon detectors at low temperatures ($< 0^\circ\text{C}$), reverse-annealing can be effectively suppressed.

3.3 The ATLAS SCT

The geometrical layout of the ATLAS ID, including components of the SCT, is shown in figure 3.1. In total the SCT comprises 4088 modules, with 63 m^2 of silicon sensors and 6.2 million read-out channels.

The central region, covering $|\eta| < 1.4$, is instrumented by the barrel subsystem, consisting of four concentric cylinders, with radii from 299 mm to 514 mm. The barrels are labelled 3,4,5,6 in increasing radius. Mounted on the barrel are a total of 2112 identical modules [59], with the module distribution per layer shown in table 3.1a.

The regions either side of the barrel, out to $|\eta| < 2.5$, are covered by two end-cap subsystems, labelled as A ($+z$) and C ($-z$). Each end-cap consists of nine discs of varying sizes, with a total of 988 modules [60]. Each disc contains modules mounted in up to three rings, known as inner, middle and outer. Due to geometrical constraints, three different module designs are required for each ring. The number of modules per disk and ring are shown in table 3.1b. For geometrical acceptance reasons, disk 9 only contains outer modules, while disks 1,7 and 8 have no inner modules. In addition, the middle modules of disk 8 are shorter as they contain only a single sensor.

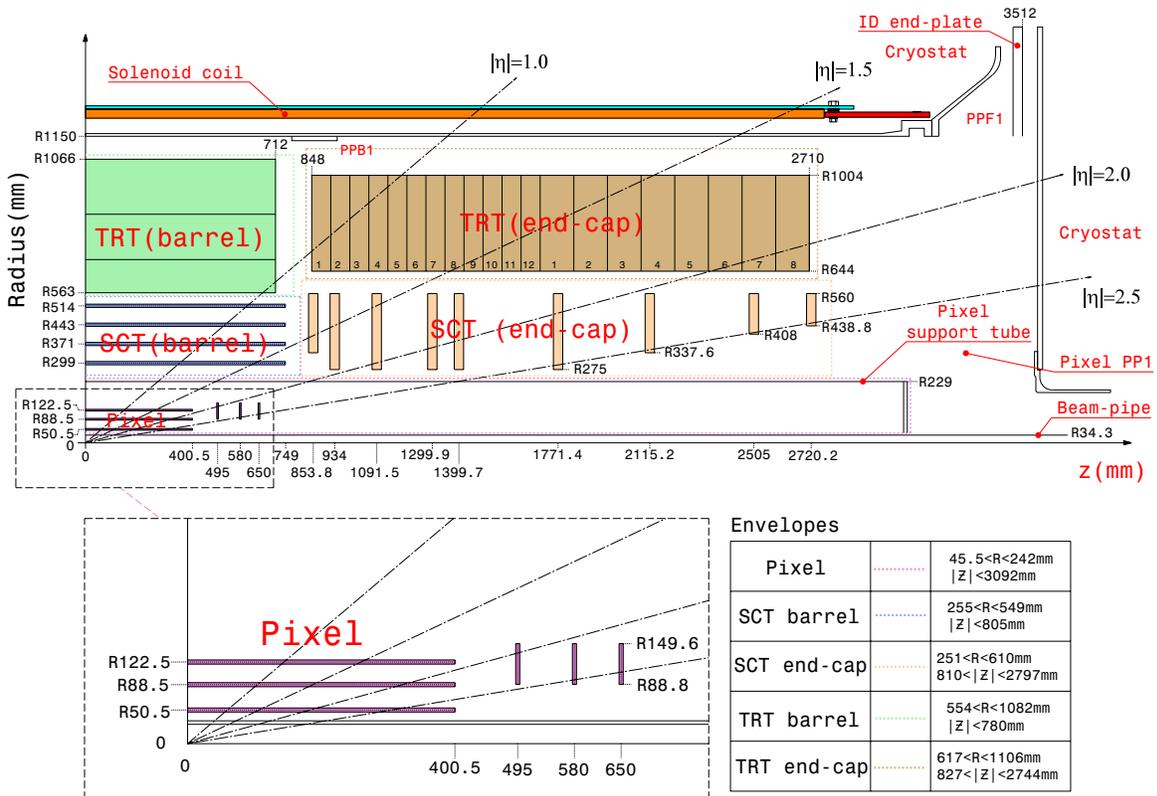


Figure 3.1: Plan view of a quarter-section of the ATLAS inner detector from [49], showing each of the major detector elements with its active dimensions and envelopes. The labels PP1, PPB1 and PPF1 indicate the patch-panels for the ID services.

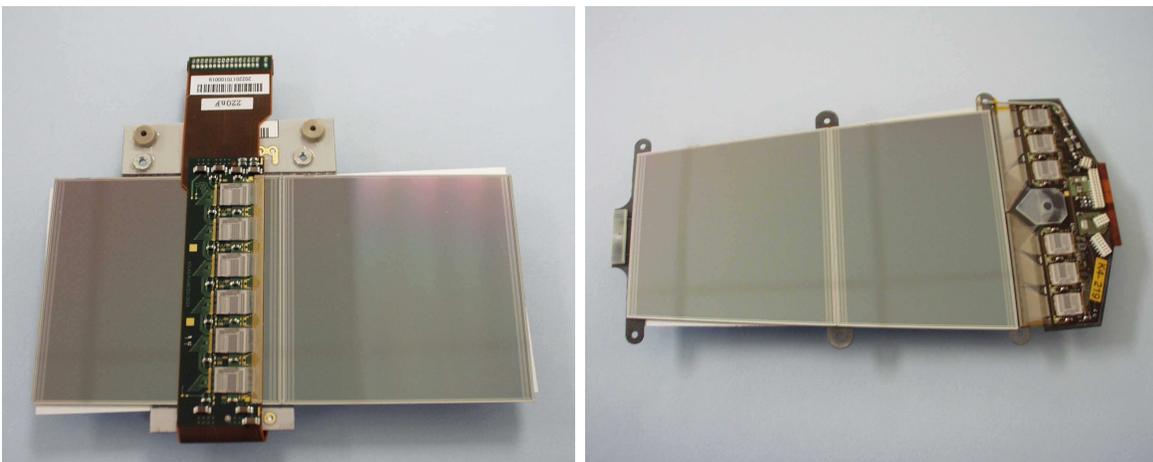
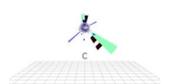


Figure 3.2: An SCT barrel (left) and end-cap module (right). The barrel modules are all identical, whereas the end-cap modules come in three different types: inner, middle (shown here) and outer. The six ABCD3T chips are clearly visible mounted on the hybrid.



Barrel Layer	Sensor Radius (mm)	Number of modules
3	299	384
4	371	480
5	443	576
6	514	672

(a) SCT barrel radii and number of modules.

Disk	1	2	3	4	5	6	7	8	9
$ z $ (mm)	853.8	934.0	1091.5	1299.9	1399.7	1771.4	2115.2	2505.2	2720.2
Outer	52								
Middle	40								None
Inner	None	40					None		

(b) SCT end-cap nominal z position and number of modules per disk.**Table 3.1:** SCT dimensions and distribution of modules.

3.4 SCT Modules

The SCT module is the basic unit of the detector and is comprised of the silicon wafers, front-end read-out chips and support structure. Photographs of a barrel and end-cap module are shown in figure 3.2.

3.4.1 Silicon Sensors

Each module has two planes of silicon, each with 770 $80 \mu\text{m}$ strips of heavily doped p-type (p+) implant on n-type bulk, with a heavily doped n-type (n+) base, denoted by (p+, n, n+). The cross section of an SCT module can be found in figure 3.3. After irradiation, the lightly doped n-type bulk will undergo type-inversion to behave more like p-type [61]. In this way the p - n junction will move from the top of the wafer to the bottom. This design allows the SCT to function as a particle detector even after radiation damage. At the start of LHC operations, a nominal bias voltage is applied across the junction of 150 V, which can be increased to compensate for irradiation effects.

Each plane of silicon consists of two sensors, with thickness of $285 \mu\text{m}$. The sensors are joined together with wire bonds, effectively doubling the strip length, but leaving a small inactive region across the centre of 2 mm. The total length of active silicon

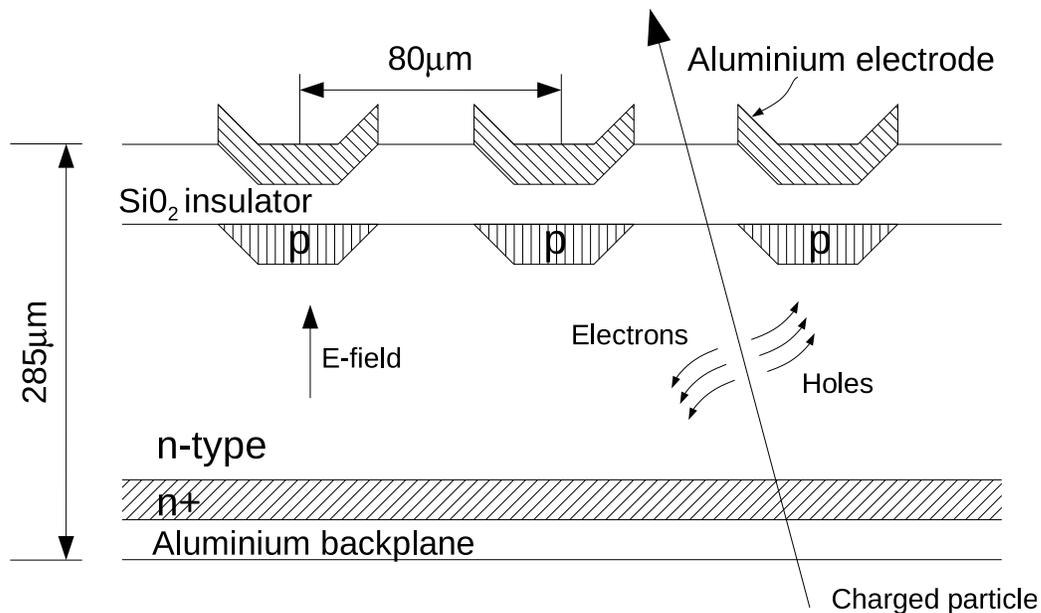


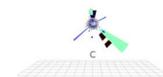
Figure 3.3: The silicon layers used in the construction of an SCT module.

Module Type	Active length (mm)
Barrel	126.09
Outer end-cap	119.14
Middle end-cap	115.61
Short-middle end-cap	52.48
Inner end-cap	59.1

Table 3.2: Active length of silicon sensors in the SCT.

is given in table 3.2, and differs between module types. Two pairs of such sensors are mounted back-to-back and offset from each other by a small stereo angle of 40 mrad to provide spacepoint reconstruction.

Barrel module sensors are mounted with silicon strips parallel to the beam direction, along the z -axis. In this orientation, the direction of the charge carrier drift is perpendicular to the solenoidal field. To compensate for Lorentz forces on charge carriers,



barrel modules are tilted by a small angle of $\sim 4^\circ$, known as the Lorentz angle [62]. In the end-cap, the strips are mounted radially in the $x - y$ plane, with the charge carriers drifting almost parallel to the magnetic field.

The module is supported by a baseboard which provides structural stability as well as dissipating heat. The modules are operated at a nominal temperature of -7°C , which primarily reduces damage from irradiation, in addition to minimising leakage current. A network of thin-walled copper pipes containing C_3F_8 in thermal contact with each module maintain this temperature by evaporative cooling.

3.4.2 ABCD3T Chips

The implant strips are read out by ABCD3T chips, which are radiation hard, Durci Mixte sur Isolant Logico-Lineaire (DMILL) technology [63] Application-Specific Integrated Circuits (ASICs), each with 128 channels. A total of 12 chips are mounted on a flex circuit known as the hybrid, which in turn is folded around the four-sensor baseboard assembly. On each side of the module, the central 768 strips are wire-bonded to the ABCD3T input channels. The remaining two outermost strips are unconnected, but are grounded to maintain field uniformity.

The charge generated in each strip on the sensor is first amplified, shaped and then forms the input to a discriminator, as shown in figure 3.4. The discriminator output is a binary signal that is passed through a mask register and stored in a pipeline buffer each time a clock signal is received by the chip. A high value of the binary output is known as a *hit*, as sufficient charge has been deposited above threshold.

An optional edge-detect circuit can be applied which ensures that each detector hit only generates one output pulse. On receipt of a L1A, the result at the end of the pipeline is moved into an output buffer and compressed before being sent to the readout.

To aid synchronisation throughout the detectors, hits can be read-out in three consecutive 25 ns time bins centred on the L1A trigger. This information is recorded as a three-bit hit-pattern in the data-stream. For example, a hit pattern of 011 denotes no hit in the bunch crossing before the trigger, and hits coincident with, and following the trigger. Hits with time-bins of 01X = 010, 011 are considered to be timed-in to the trigger signal.

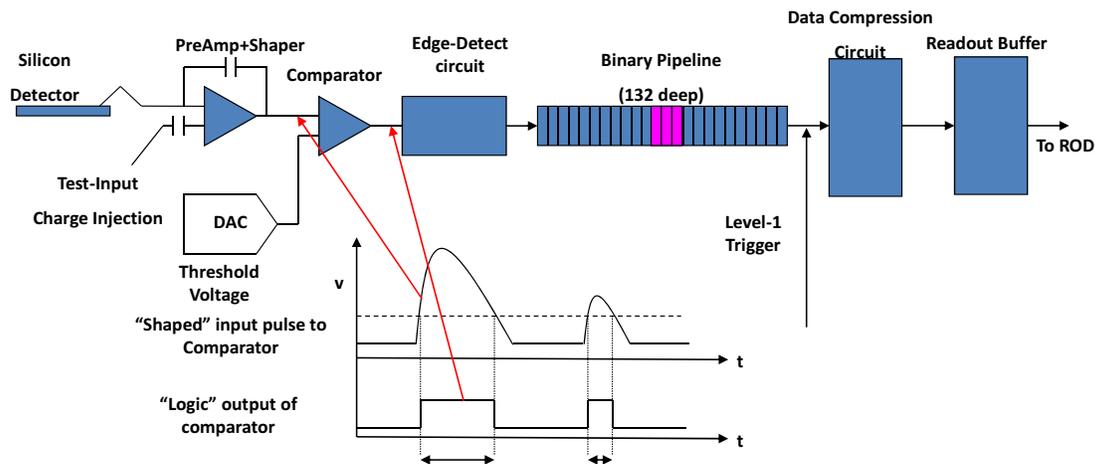


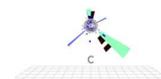
Figure 3.4: The components of the ABCD3T chip.

For calibration purposes, known charges can be injected into every fourth channel by applying voltages across a calibration capacitor.

There are two ABCD3T chips connected to the optical readout (one on each side), denoted as *master* chips, with the remaining 10 designated as *slaves*. During normal readout, data from each slave are passed in series to the master. Single failed chips can be bypassed to prevent loss of data from other chips on the module.

3.5 The SCT Data Acquisition System

The SCT DAQ system [64] is comprised of hardware components on the silicon detector modules, optical links and off-detector hardware and the DAQ software package which controls and configures these components for data-taking.



3.5.1 Hardware

Information from the SCT is read out via optical links which connect modules in the main detector hall with off-detector electronics in the adjacent counting rooms.

Optical Links

Each SCT module is connected to the off-detector read-out system by three infra-red optical links [65].

A module receives information over a single optical transmission (Tx) fibre, with the clock and command signals combined using a Bi-Phase Mark (BPM) encoding. A PiN¹ diode converts the incoming data from optical into electrical signals, which are then decoded by the Digital Optical Receiver Integrated Circuit (DORIC) and distributed to all chips on the module. A second set of clock and command signals is also generated which can be sent to the adjacent module on the detector, providing redundancy in the system should a link fail.

Each master chip on a module is connected to a Vertical Cavity Surface Emitting Laser (VCSEL), driven by the VCSEL Driver Chip (VDC), allowing transmission of hit data off-detector via two optical receiver (Rx) links. Under normal configuration, the hit data from six ABCD3T chips are transmitted per link, corresponding to one side of a module. Should one of the links fail, the chips on a module can be configured to read out information for all chips through a single link.

The entire optical package containing the PiN, DORIC, VCSEL and VDC and fibre connections are housed on the detector adjacent to the module.

Off-Detector Read-out Hardware

The off-detector DAQ system is installed in the ATLAS large service cavern, around 100m from the SCT itself, which resides in the experimental hall. The off-detector hardware consists of eight 9U Verse Module European (VME) crates and nine rack mounted servers. The crates contain various VME modules that communicate via a custom backplane. An overview of the most important components in the system is shown in figure 3.5.

¹PiN: a wide intrinsic semiconductor between p- and n-type regions.

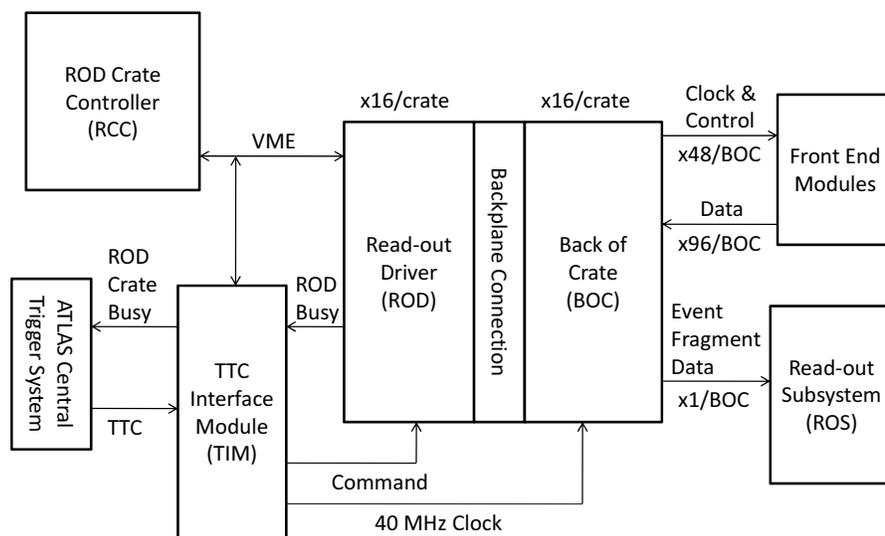


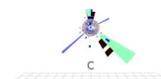
Figure 3.5: Off-detector electronics of the SCT.

An SCT crate will typically contain a single Timing Interface Module (TIM), a number of ROD and Back-of-Crate (BOC) pairs, and a single ROD Crate Controller (RCC).

ROD Crate Controller The RCC is a commercial 6U Single Board Computer (SBC) running Linux, which acts as the VME host, and as a result occupies the first slot in the crate. Dedicated software runs on the RCC providing an overall interface for control, configuration and communication of all the hardware present in the crate.

During data-taking physics mode runs, the RCC is used to configure modules, after which the data flow is handled solely by the Field Programmable Gate Arrays (FPGAs) on the ROD. After a run has started, the SCT DAQ software takes on the role of monitoring.

Timing Interface Module There is a single TIM in each crate, which is responsible for distributing the Trigger, Timing and Control (TTC) signals from the ATLAS Central Trigger Processor (CTP) system to every ROD over the backplane. For every L1A trigger



sent, the TIM also distributes Level-1 trigger ID (L1ID) and Bunch Crossing ID (BCID) counters to assist event synchronisation. If a ROD cannot keep up with the read-out rate, it will assert a BUSY signal that is sent to the TIM. This signal is propagated back to the ATLAS trigger system, halting triggers until the BUSY is cleared. The TIM can also generate a BUSY internally to veto fixed-frequency triggers, which may induce harmful wire-bond resonances [66].

Back-of-Crate Card The BOC is responsible for transmitting commands and data between the optical fibre connections and the ROD with which it is paired.

Each command designated for the front-end modules is routed via a Tx plug-in, which converts the clock and command into a single BPM signal. This electrical signal is converted to optical form by a 12-way VCSEL array before being transmitted to the modules along a 12-way fibre ribbon. The intensity of the optical signal can be configured at the level of individual fibre connections, using a Digital-to-Analogue Converter (DAC) on the BOC. Timing of the outgoing signals can be adjusted to ensure that the clock signal received by the modules has the correct phase relative to the particles from LHC collisions. This is set on a module-by-module basis to allow for differences in fibre lengths and time-of-flight variations from different module locations.

Incoming data from modules are received in optical form from 48 modules (96 input links) per BOC, and converted into electrical signals by eight 12-way PiN diode arrays. The signals are then discriminated and sampled at a variable phase to ensure reliable reconstruction of the binary data stream. The electrical data are finally forwarded to the ROD module that is paired with the BOC for event processing.

In addition, the BOC is also responsible for transmitting formatted event fragments from the ROD to the first level of the ATLAS high-level trigger system, the ROS. The ROD generates a single, formatted data stream which is forwarded to the ROS via an S-link connection [67].

The 40 MHz ATLAS clock is normally distributed directly to the front-end modules from the TIM via the BOC. In the absence of this clock, a phase-locked loop on the BOC can generate a local replacement. This ensures that the modules always receive a clock signal, without which they generate less heat which could effect detector alignment due to thermal distortions.

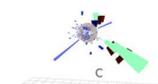
Readout Driver One basic function of the ROD [68] is to transmit the control commands and configuration data through the BOC to the front-end modules. These data can be Level 1 triggers, Event Count or soft resets (ECRs) or Bunch Clock Resets (BCRs), calibration commands or module register data. The second function of the ROD is to receive and format data streams from the front-end modules, which can either be event data or module register data. The ROD boards themselves are a hybrid architecture of FPGAs, which implement the data path, and Digital Signal Processors (DSPs) for control and calibration. Each ROD can control and process data from 48 SCT modules, with up to 16 RODs per crate.

The Master DSP (MDSP) oversees the operation of the entire ROD, although it does not explicitly take part in data taking during physics runs. It provides access to the FPGAs and can run tasks such as histogram control and link masking, communicating between the host via transfer of so-called primitives.

The ROD Controller FPGA (RCF) coordinates all of the control path operations required for data-taking, module calibration and on-board diagnostics, provides connections from the other FPGAs, slave DSPs and BOC to the MDSP and interfaces with the TIM for access to clock and trigger data. In normal data-taking one of the key functions of the RCF is to distribute clock and trigger signals to the modules via the BOC.

The remaining FPGA components are used to implement the data path on the ROD, as shown in figure 3.6. The formatters receive serial data from the module input links, convert them into a 32-bit data word format and provide derandomising buffering of the event fragments for each link in parallel. The formatters also detect module errors and can send a ROD BUSY to the TIM if one of the buffers is close to maximum occupancy. The formatter can output data in one of three modes. In expanded mode, the three bit time-bin information for each hit is retained, while in condensed mode it is discarded in order to produce a more compact data stream. While condensed mode is the default for SCT data-taking, the majority of physics runs up to summer 2010 have been performed in expanded mode to retain module timing information. In the third mode, raw data sent from modules is packaged and forwarded through the ROD, and can be triggered if certain synchronisation bits in the data stream are corrupted.

There are eight formatter chips on each ROD, each capable of reading out data from 12 links. The formatters are arranged into two banks of four devices, an architecture that allows the processing bandwidth to approach 80 MHz.



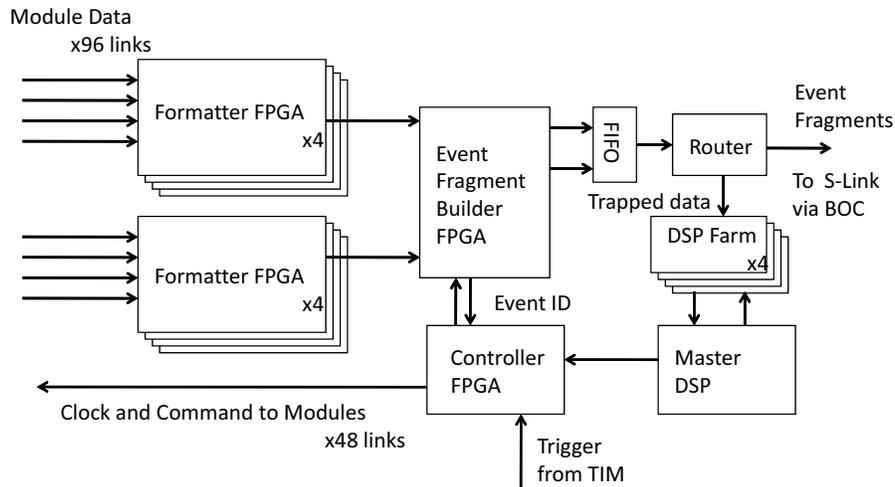


Figure 3.6: SCT ROD data path implemented by a hybrid architecture of FPGA and DSP components.

The next stage in the data path is the Event Fragment Builder (EFB) FPGA, which is responsible for creating ATLAS standard event fragments [69]. The EFB contains two processing engines that can each collect the output from a bank of four formatters. The data stream is also monitored and flagged for errors, including checks of L1ID and BCID synchronisation. As the event fragment is constructed, it is stored in a First-In, First-Out (FIFO) buffer until a complete fragment is ready to be sent on to the router.

The router FPGA is the final stage of data processing on the ROD, with the primary function of transmitting event fragments to the S-link, via the BOC. As data flow through the router, error flags are added to each link header, and link information replaces the L1ID and BCID counts. The router also contains four event traps which can be configured to syphon events at variable frequency and, if desired, filter them for specific trigger source (either ATLAS, TIM or ROD). Each trap is connected independently to a Slave DSP (SDSP), where event fragments can be transferred for counting and monitoring.

The SDSPs have a similar program structure to the MDSP, sharing the same primitive and task functionality. During normal operation, the MDSP is continuously com-

municating with the four slaves, and so direct access to them is not possible from the host. Instead, primitives intended for the slaves are sent using the master DSP as a middleman.

The main function of the slaves is to histogram event data from the router during calibration scans. The DSP retrieves data frames of 256 words using a Direct Memory Access (DMA) transfer with an Interrupt Service Routine (ISR) on completion. Once a complete event has arrived, the DSP ISR places it onto an event queue. On each pass through the DSP polling loop, the next event in the queue is transferred for processing by dedicated histogramming tasks, after which the event is removed and memory released. These tasks can either be in the C programming language, or in the faster assembly language, and can be configured to process hits from all three time bins together, or in separate histograms. At the end of a calibration scan the histogram data can then be read out to the host through the MDSP using data-packets known as primitives.

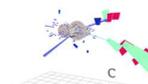
The ROD and TIM module designs are common to both the SCT and Pixels, with differing firmware tailored to suit the needs of the two sub-detectors.

3.5.2 Software

In addition to software running on the RCC, other distributed calibration and analysis tasks are run on the rack-mounted servers, written in a combination of C++, Java and Python. These machines are connected to each other and the user in the ATLAS control room via Ethernet.

The DAQ software is written using the ATLAS TDAQ framework, which allows it to be smoothly integrated into central ATLAS data-taking. In particular, the SCT DAQ makes extensive use of the Inter-process Communication (IPC) and Information Server (IS) components of TDAQ. IPC uses Common Object Resource Broker Architecture (CORBA) [70] to enable network communication between processes written in a number of languages which may be running on different machines. IS provides the framework for data sharing between applications, providing a repository in which data may be published and accessed.

To run the SCT detector at maximum efficiency, the optimum values must be found for various crate and module parameters, as described in section 3.6. The customary procedure to find the optimum value of some parameter is to perform a scan over all



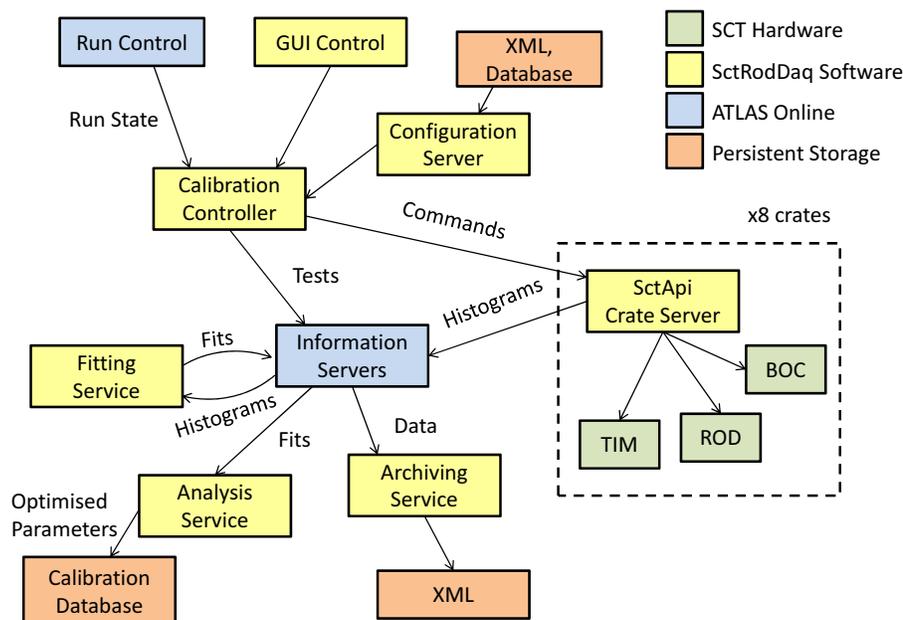


Figure 3.7: The structure of the SCT DAQ software.

possible settings of that parameter, using algorithms to fit and analyse the results. In some cases, a test is performed which involves running a series of scans in succession.

During an SCT calibration run, the SCT DAQ package is used to control and implement the starting of scans. An overview of the SCT DAQ software is shown in figure 3.7. The configuration of the SCT hardware consists of individual crate and module settings, and can be stored either in a Conditions Data Storage Model (COOL) database [71] or as a series of XML files. A single application, known as the SctApi Crate Server, runs on each of the eight SBCs, providing an interface between the SCT DAQ software and hardware. Before the start of a run, each crate server retrieves the relevant hardware configuration via the Configuration Service.

During calibration, the crate servers control setting up and varying scan parameters and collecting data into occupancy histograms. For each setting of a scan parameter, a sequence of internally generated triggers is sent to each module. The resulting hit data are trapped on the router, which are then histogrammed on the SDSP and finally read out by the host at the end of the scan.

The Calibration Controller is responsible for overseeing the process of setting up tests, starting scans on the crate servers, and choosing the relevant fitting and analysis algorithms.

Once a test is completed on the crate servers, the histogrammed results are published to an IS. The Fitting Service contains a listener thread which takes these histograms and adds them to a queue. Worker threads then undertake the processor-intensive job of performing fits, while the listener threads are free to respond to further data. This listener and worker implementation allows the computing to be done in close to real-time.

In a similar way, fit data published by the Fitting Service are picked up by the Analysis Service which then extracts the optimum operating parameters from the scan, and decides if the test has been successful. The Analysis Service may also run on raw histograms in the case where a fitting procedure is not needed.

Finally, the Calibration Controller monitors the test results when they are available and updates the configuration with the updated optimum values.

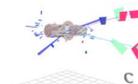
3.6 Calibration of the SCT

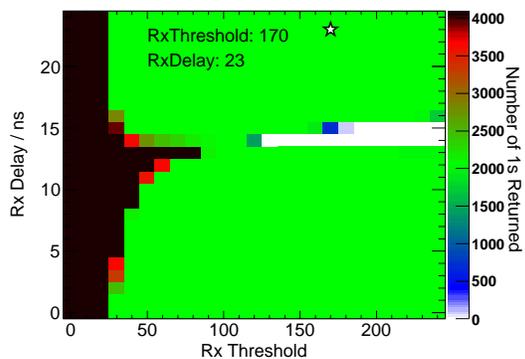
A number of different tests are necessary to fully calibrate the SCT detector. These tests fall into the three main categories of optical tuning, digital tests, and analogue characterisation. An excellent description of all the calibration tests can be found in [72]. A brief summary of some of the most important tests will be given here.

3.6.1 Optical Tuning

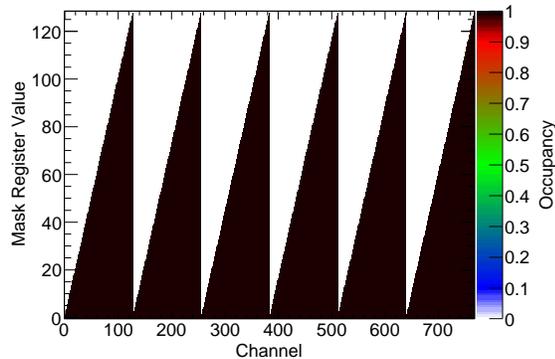
To ensure reliable communication between off-detector electronics and the front-end modules, it is essential to ensure that the optical links are well configured.

The most significant variables in the optical tuning are the current supplied to the VCSEL chips on the BOC (Tx current), the threshold at which a received signal is discriminated, and the phase that the received signal is sampled (Rx threshold and delay respectively). As the Rx threshold and delay are correlated, an important optical scan varies the two parameters simultaneously to find a region of parameter space where reliable communication occurs.

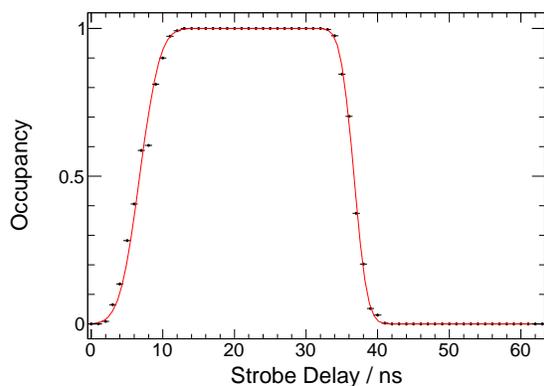




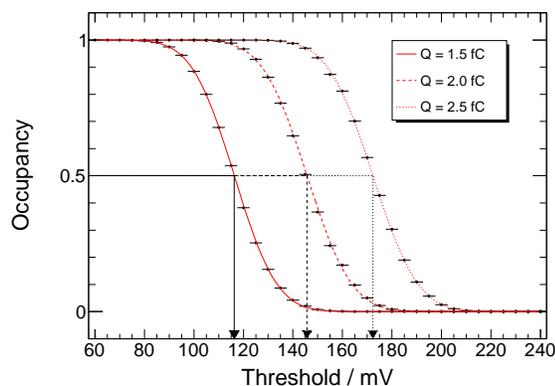
(a) 2D RxDelay and RxThreshold scan for link 1.



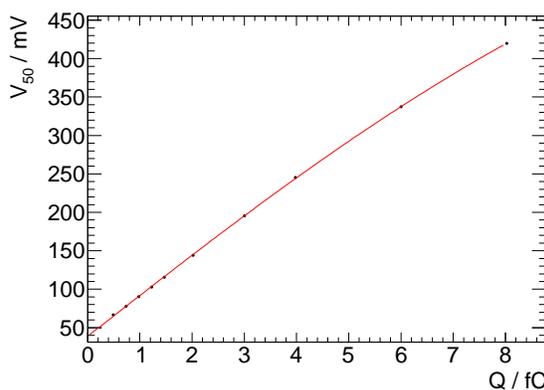
(b) NMask register scan for link 0.



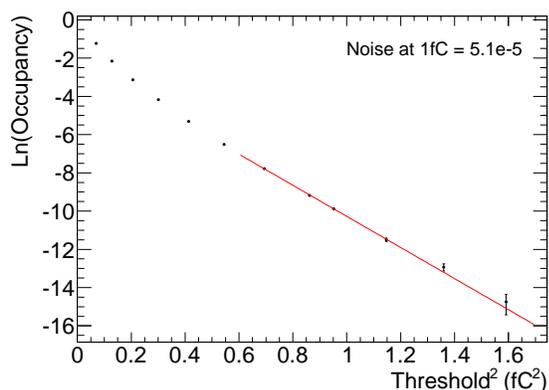
(c) Strobe delay scan for chip 0.



(d) Three point gain threshold scans for chip 0.



(e) Full response curve (10-point gain) for chip 0.



(f) Noise Occupancy scan for chip 0.

Figure 3.8: Example calibration plots for barrel sector module 20220170200123

For the optical tests, the modules are put into “clock/2” mode, in which they return a clock signal which is half the frequency of the 40 MHz input signal. For each setting of threshold and delay, counters on the MDSP record the number of 1’s within a set time window. As no triggers are sent, this is one of the fastest tests performed.

An example of a simultaneous Rx threshold and delay scan can be seen in figure 3.8a. Here, the MDSP records the number of 1’s in 4000 clock cycles, which should be 2000 for reliable communication. The figure clearly shows regions of optical noise at low threshold, and the region around the rising and falling edge of the clock signal, in this case at delay values of 15 ns. The optimum value found by the Analysis Service is also indicated on the plot by a star.

3.6.2 Digital Tests

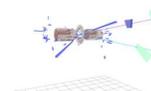
Once communication is established, the next stage is to confirm that the ABCD3T chips are functioning correctly. This is done with a variety of tests which exercise the channel mask registers (NMask), trigger and bunch crossing counters, pipeline cells and chip token-passing logic.

An example of an NMask test is shown in figure 3.8b, where module strips are sequentially masked off to form the characteristic pattern shown in the plot. This register is used to mask noisy strips out of the SCT configuration.

3.6.3 Analogue Calibration

An essential part of calibration is to characterise the detector’s response to injected charge. Each front-end ABCD3T chip has an 8-bit DAC which sets the threshold globally across that chip. Channel-by-channel variations can be compensated for by using a 4-bit DAC known as the TrimDAC.

One of the simplest tests, known as a Strobe Delay, shown in figure 3.8c, varies the phase of charge injection relative to the trigger command. A charge of 4 fC is injected at a threshold of 2 fC, and for the correct timing there should be 100% efficiency. A smeared top hat function is fitted to the data from each chip, with the optimum value set at 40% of the full width.



Disabled Readout Component	End-cap A	Barrel	End-cap C	SCT	Fraction (%)
Disabled modules	5	10	15	30	0.73
Disabled chips	5	24	4	33	0.07
Masked strips	3364	3681	3628	10673	0.17
Total					0.97

Table 3.3: The state of the SCT configuration as of May 2010, from [73].

By injecting various known charges and measuring the occupancy at different thresholds, the analogue properties of each channel can be determined. For each value of charge injected, the threshold is scanned over the entire range and, with the assumption of Gaussian noise, a complementary error function is fitted to the average occupancy for each chip. The threshold at which occupancy is 50% corresponds to the median of the injected charge, with the sigma giving the characteristic noise. By repeating this scan at different injected charges, the per-chip response can be built up. An example of a fast scan with three injected charges is shown in figure 3.8d. The longer full response curve contains data points from ten different charges, as shown in figure 3.8e.

The parameters used to fit the response curve are stored by the module configuration and can be used to set the threshold at an arbitrary charge. For normal operation, the threshold is set to a value of 1.0 fC. The overall chip response is characterised by the gain, which is the slope of the response curve, and the Equivalent Noise Charge (ENC), which is the width of the complementary error function, both quoted at 2 fC.

Similar scans are used to optimise the TrimDAC thresholds, and threshold scans without injected charge are used to find the noise occupancy of the detector. An example of the noise occupancy test is shown in figure 3.8f, where the noise is measured over a range of thresholds. The plot of $\ln(\text{occupancy})$ against the square of the threshold is reasonably linear, as would be expected for Gaussian noise. The noise occupancy at 1 fC is extracted from this plot, and the slope of the linear fit is used as an independent measurement of the ENC.

3.6.4 Calibration Results

Extensive calibration of the SCT modules was undertaken during 2008-09 in preparation for first LHC collisions. Table 3.3 shows the number of modules, chips and strips disabled

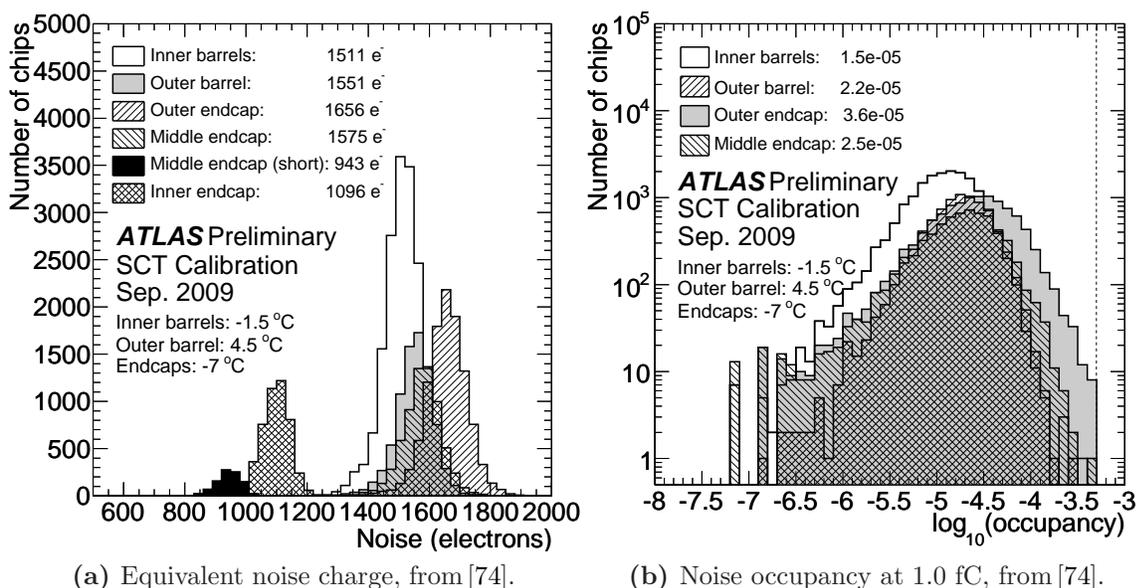


Figure 3.9: Results of SCT noise characteristics after extensive calibration during summer 2009.

or masked in the SCT configuration, as of May 2010. Of the disabled modules, 13 are excluded due to a defective cooling loop in end-cap C, 13 due to low- and high-voltage issues, and 4 are due to problems with optical communication. Of the 33 chips disabled, 31 of these are bypassed, and two chips are present for which all strips are masked. Strips are masked if they are particularly noisy, and are typically isolated and distributed across the entire SCT. The exception are six groups comprising between 10 and 20 adjacent strips. Overall fewer than 1% of the SCT detector channels are disabled.

Characterisation tests during 2008-09 resulted in a factor of two reduction in the noise occupancy of modules in the SCT, with smaller tails in the distributions. Figure 3.9 shows the ENC and average strip noise per module above a 1 fC threshold. Variations in noise related to strip length in the different types of modules are clearly seen. Temperature effects are also visible. While the end-caps were cooled to the nominal temperature of -7 °C, the barrels were maintained at -1.5 °C and 4.5 °C for barrels 3–5 and barrel 6 respectively. The higher outer barrel value is required to minimise the temperature difference between the SCT and neighbouring TRT. This temperature difference is reflected in the noise occupancy results, with the outer barrel 68% noisier than the inner layers. Overall, the majority of noise occupancy values lie within the range $2\text{--}5 \times 10^{-5}$, an order of magnitude better than the specification of 5×10^{-4} . This allows some headroom for noise occupancy increase after irradiation.



3.7 Data-taking

In parallel with the calibration runs, the SCT has also collected triggered events in data-taking mode, from simulated readout to high energy LHC collisions.

3.7.1 Milestone Runs

During 2007 and 2008 a number of week-long milestone runs took place, with the aim of integrating different ATLAS sub-detectors for combined read-out using either random or cosmic triggered events.

Participation of the SCT was limited to the latter of these runs, where the DAQ was exercised with the rest of ATLAS. In the first of these, a small number of modules in a test box were read-out, testing the performance of the trigger system and synchronisation issues. In the next run almost the whole of the SCT barrel took cosmic data, after synchronisation of the readout phase with the rest of ATLAS. For the final two runs the formatter simulator was used, enabling validation of software changes and the successful testing of 70 kHz high rate triggers. Further details of the SCT simulator are given in section 5.4.

3.7.2 First Circulating Beam in the LHC

On the 10th September 2008, the LHC began injecting protons into the accelerator, which were stopped on collimators close to ATLAS. Splashes of particles originating from the collimators reached the ATLAS cavern and were detected by a number of subsystems, including the SCT. Due to concerns about module safety, the bias voltage was reduced to 20 V with a threshold of 1.2 fC, and only the end-caps were powered. An example of a so-called splash event is shown in figure 3.10. Further splashes were provided at each subsequent start-up of the LHC, in late 2009 and early 2010, to enable coarse timing offsets to be determined.

3.7.3 Cosmic Muons

Between September 2008 and November 2009, the SCT was mainly concerned with data-taking using cosmic muons, with either the whole of ATLAS, or as part of dedicated

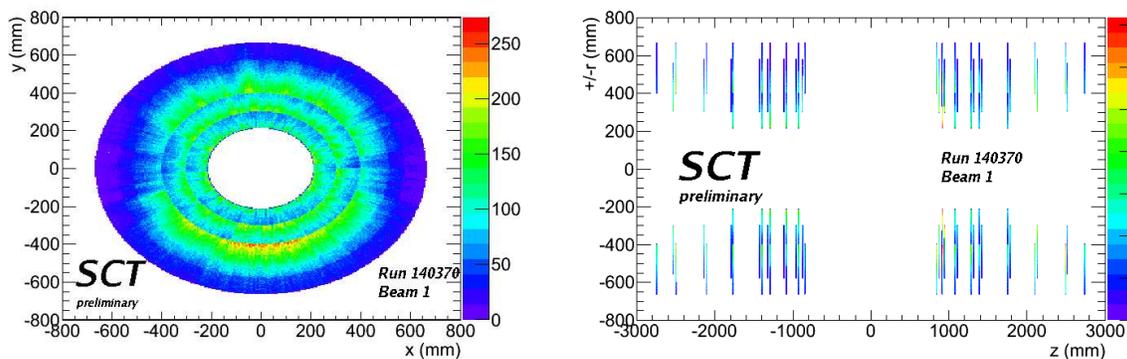


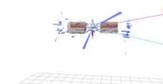
Figure 3.10: An example of an LHC beam splash event as seen by the SCT, showing the number of spacepoints, with the average noise from empty events subtracted. The left-hand plot shows the spacepoints in all 18 disks, projected into the plane perpendicular to the beam axis. The right-hand plot shows the same event from the side, with the beam originating from the negative z direction. Plots from [75].

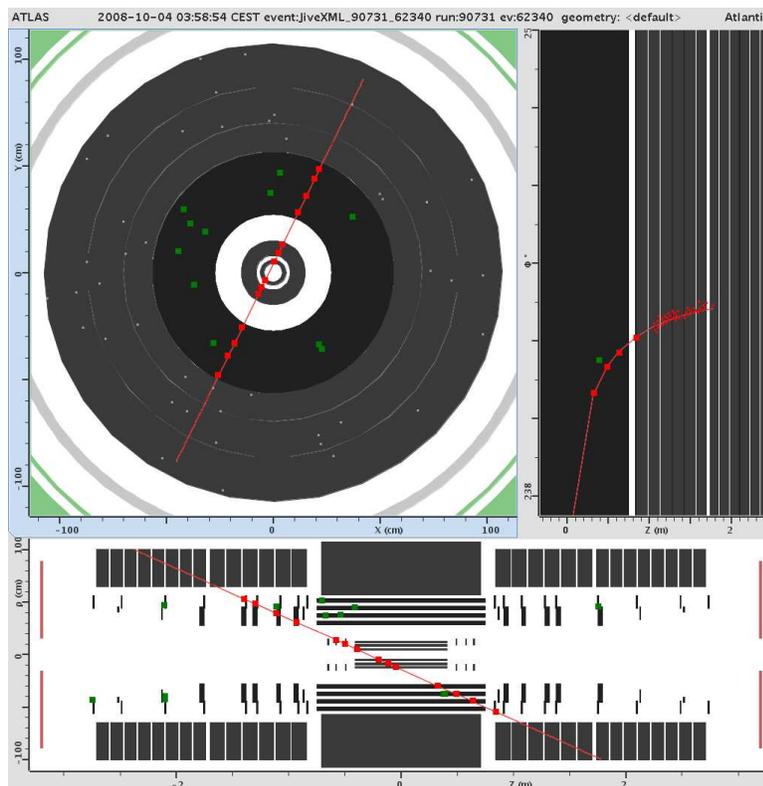
ID cosmic runs. An example of a cosmic-ray event traversing the SCT is shown in figure 3.11a. In 2008, the SCT DAQ collected around 1.15 million cosmic tracks without the magnetic field switched on, and 0.88 million with the magnetic field at 2 T. During this period, the efficiency of the barrel modules was measured to be well above the 99% nominal value. By examining the distance between a fitted track and the measured SCT spacepoint, detector alignment studies were also performed.

3.7.4 First LHC Collisions

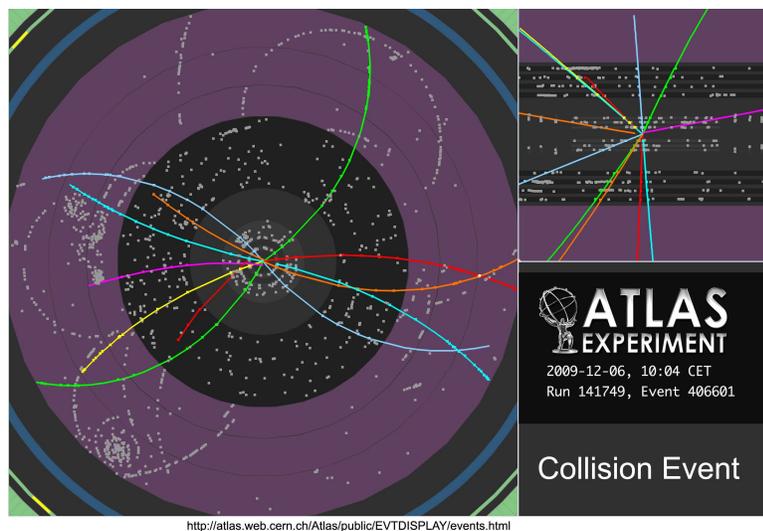
Colliding beams arrived in December 2009, giving the SCT the opportunity to finally test its mettle. Proton-proton bunches were synchronised in ATLAS at the injection energy of $\sqrt{s} = 900$ GeV, with one of the first collisions shown in figure 3.11b. This low energy run lasted until the end of 2009, during which the SCT performed to specifications. Figure 3.12a shows the intrinsic module efficiency for tracks in the SCT barrel. The efficiency is calculated as the ratio of measured hits on a track to the number of possible hits, taking into account excluded modules and chips. Two different types of track are displayed, both standalone SCT tracks, and combined tracks from all three inner detector sub-systems. The overall barrel hit efficiency is $99.8 \pm 0.1\%$. The end-caps also show efficiencies of $99.8 \pm 0.1\%$ each.

Figure 3.12b shows the number of strips with hits per module side during $\sqrt{s} = 900$ GeV collisions, with the solenoid magnet switched on. A minimum bias Monte



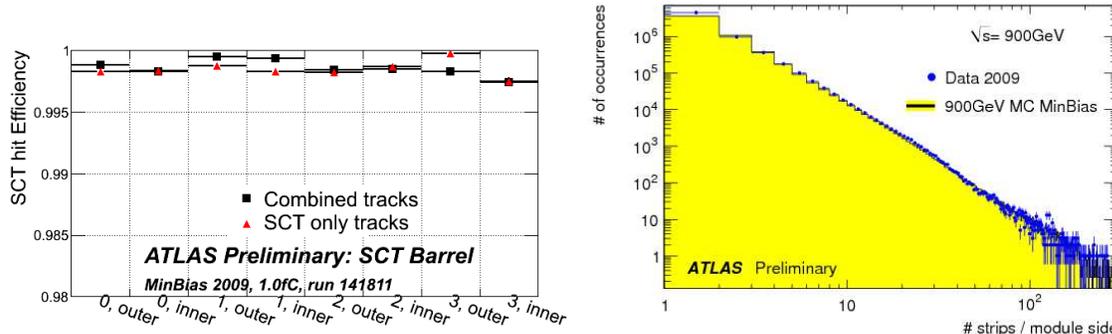


(a) A cosmic muon event, from [75]. The solenoid magnetic field was switched off, producing a straight track which traverses both the SCT and pixel barrels. This event is of particular interest as it also traverses one of the SCT end-caps.



(b) One of the first collision events in ATLAS at $\sqrt{s} = 900$ GeV, zoomed to the inner detector, from [76].

Figure 3.11: Atlantis event displays highlighting SCT performance.



(a) The SCT hit efficiency during run 142165. (b) The number of strips with hits per module side, from run 14291.

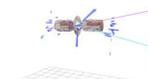
Figure 3.12: SCT performance plots during LHC collisions at $\sqrt{s} = 900$ GeV, from [73].

Carlo sample is shown for comparison, normalised to the number of data events. A good agreement is obtained across a wide range of strip numbers, with the discrepancy at low N due to the underestimation of noise in the simulation.

3.8 Conclusions

This chapter has introduced the SCT detector, with a focus on the DAQ system used to read-out and calibrate its 4088 modules. The tests required to optimise the operation of the SCT have been described, including results from recent calibrations. In total, over 99% of the SCT detector channels are operational as of May 2010. The noise occupancy of modules in the SCT ranges from $2\text{--}5 \times 10^{-5}$, within the specification of 5×10^{-4} .

During data-taking runs, the SCT DAQ has successfully triggered and read out millions of events in standalone and global ATLAS commissioning runs. In late 2008, the SCT recorded 2 million cosmic muon events with a barrel efficiency of over 99%. During the first LHC collisions in December 2009, the SCT measured a hit efficiency of $99.8 \pm 0.1\%$ at $\sqrt{s} = 900$ GeV.



Chapter 4

The SCT Test Sectors

“Watch your soul 'cos science is golden”

— The Grates

After 2007, the focus of the SCT collaboration shifted from installation and commissioning of the detector to its calibration and operation. Consequently, this reduced the availability of the SCT for DAQ developments, highlighting the need for a small scale test system.

In order to maintain a platform for developing and testing new features, two systems were constructed with the purpose of recreating the SCT on a small scale. The first to be commissioned was a sector of 48 barrel modules, followed by a quadrant of 33 end-cap modules mounted on a spare disk. The sectors were constructed with DAQ, DCS and cooling systems closely matching the real detector in order to gain operational experience in a more accessible environment.

Many of the SCT developments described in chapter 5 were prototyped using the test sectors described in this chapter. The author was heavily involved in the construction of the sectors, including routing and connection of cables, testing of the power supply and readout connections, and configuration of the DAQ system.

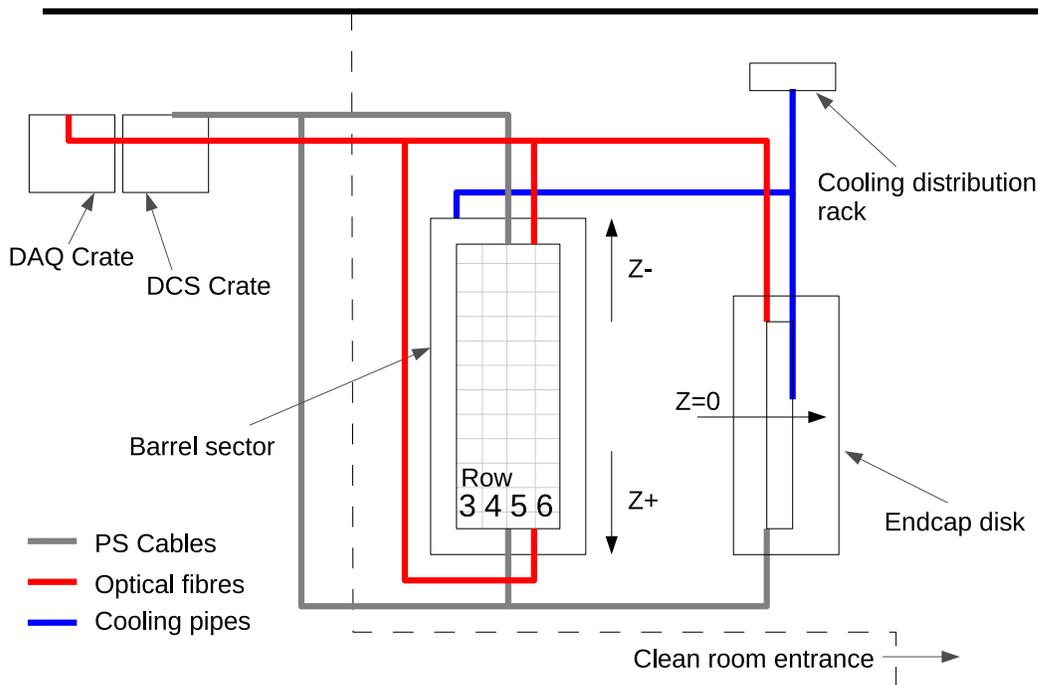


Figure 4.1: The location and orientation of the barrel sector and end-cap disk in the SR1 clean room.

4.1 Installation

The SCT test sectors were constructed and installed in the SR1 building at CERN, close to the main ATLAS control room on the Point-1 site. The building was used previously to house the SCT during reception tests before final installation underground in the ATLAS experimental hall. As a result, much of the infrastructure required for operating the detector was already in place. The building consists of a clean room, which houses the two sectors, and a rack area containing DAQ and DCS crates. Optical fibres and power supply cables are routed between the two, and shown in figure 4.1.

4.2 Barrel Sector

During the early phase of SCT detector development, a barrel sector was constructed to allow testing and prototyping of modules and connections. This was stripped down



(a) Two harnesses mounted on the sector. The orange low mass tapes and black readout fibres are clearly visible.



(b) The sector mounted with loopback cards, during harness testing.



(c) The completed sector mounted with 48 barrel modules.



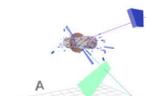
(d) A close up of modules on the barrel sector.

Figure 4.2: Photos of the barrel sector during construction.

during SCT construction, and subsequently resurrected during early 2008 to provide a platform for the new test sector.

4.2.1 Barrel Sector Layout

The barrel sector support structure consists of $\frac{1}{8}$ of a carbon fibre cylinder, with the same length and curvature as the innermost SCT barrel (layer 3). The sector can accommodate 48 modules, in four rows of 12, labelled 3–6. The orientation of the sector is shown in figure 4.1. The entire sector is housed inside a thermally insulated enclosure, which is kept at low humidity with a constant flow of dry air.



The modules on the sector are powered and controlled in groups of six by a single harness. Each harness consists of a series of low mass Kapton [77] tapes from each module, providing the power supply via embedded aluminium tracks. The harnesses also include optical fibre connections to the modules, and house the optical packages used for module communication. The barrel sector mounted with two such harnesses can be seen in figure 4.2a.

4.2.2 Power Supply

The power supply tapes from each harness are connected to a single patch-panel, which is mounted outside of the thermal enclosure. The patch-panels provide connector pins to power supply cables which are routed out of the clean room and are connected to the backplane of a power supply crate in the rack area. A single crate of low- and high-voltage cards is used to power all of the modules on the sector.

4.2.3 Optical Connections

As described in section 3.5.1, each module has a single Tx fibre and two Rx fibres. For each harness, these fibres are grouped into one ribbon from each of the six Tx fibres, and one for each of the twelve Rx connections. The ribbons are relatively short, and terminate in a connector which is housed outside the main enclosure, close to the power supply patch-panels.

4.2.4 Harness Testing

During the first stage of construction, module support brackets were mounted on the sector, followed by the harnesses which provide power supply connections and optical fibres. Before final module mounting, the harnesses underwent testing to ensure that all power supply and optical connections were operating correctly.

The tests were performed using loop-back cards attached to the harnesses, which apply fixed resistors across a number of sensor lines. A small number of broken connections were identified, which required some manual reworking before operation. The cards also return the incoming clock and command signals to and from the DAQ crates, allowing

the fibre mapping to be checked before module mounting. A photograph of the sector during the harness testing phase can be seen in figure 4.2b.

4.2.5 Final Assembly

After the harness tests were completed, 48 spare barrel modules were selected for the sector. These modules were mounted by hand using a purpose-built jig to allow positioning before being secured to the support brackets. Photographs of the barrel sector after module mounting can be seen in figure 4.2c, with detail visible in figure 4.2d.

After module mounting, the power supply cables were re-connected and re-tested. The final connections were made to the cooling system, which provides evaporative cooling using C_3F_8 , as with the real detector. The cooling plant consists of a single compressor and cooling fluid is distributed via a rack of valves located behind the sector.

4.3 End-cap Disk

A spare end-cap disk, used for initial module testing at The National Institute for Subatomic Physics (NIKHEF), was transported to CERN in autumn 2008 to allow construction of a second test setup of end-cap modules. The spare is a replica of a disc from position 2 in table 3.1b, and contains a single quadrant (top-right) populated with modules. A total of 33 modules are mounted, with 13, 10 and 10 in outer, middle and inner positions respectively. A photograph of the end-cap sector is shown in figure 4.3. As the end-cap modules have a different geometry to those from the barrel, the spare disk makes a useful addition to the SR1 test system.

4.3.1 Disk Installation

The spare disk is mounted inside an environmental enclosure containing many of the services necessary for operation. Power supply connections from the modules were routed to patch-panels mounted on the outside of the enclosure, and subsequently connected to cables from the crates. Some re-routing of power supply cables was required in the rack area which originally housed a much larger system. The end-cap disk requires a second power supply crate of high- and low-voltage cards.

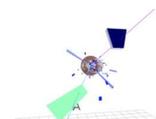




Figure 4.3: The spare end-cap disk showing outer and inner modules, viewed from $z = 0$ as indicated in figure 4.1. The middle modules are mounted on the reverse of the disk and are not visible here.

Cooling pipes were already mounted on the support structure for the disk, connected to valves on top of the enclosure. These were subsequently connected to a second outlet valve on the cooling distribution rack.

The optical connections required some rerouting inside the enclosure, with small patch fibres between the harness and the long fibres from the racks.

4.4 Detector Control System

A scaled down version of the DCS was installed in SR1 to control and monitor powering of the two sectors. This includes a version of the ATLAS DCS Finite State Machine (FSM), which is used to set and monitor the voltages and currents from the two crates. A graphical representation of the FSM from the barrel sector is shown in figure 4.4a.

The DCS system also monitors the temperature and humidity of the sector environment. Inside the two enclosures, sensors are connected to an interlock matrix card in the rack area, which automatically cuts power to the modules if they exceed safe operational limits. Hardware interlock values are set by physically changing resistor plug-ins mounted on the card itself. An additional software interlock is also present in the DCS

FSM, which also cuts power to modules if the module temperatures are dangerously high.

4.5 Data Acquisition System

The DAQ system in SR1 consists of two VME crates, one responsible for barrel sector readout, and one for the end-cap. Each crate contains a ROD-BOC pair and a single TIM, which is sufficient to control and read out all the modules in each sector. In fact, the end-cap ROD only uses 33 out of 48 readout channels, allowing the remaining 12 to readout a test box containing four water cooled modules originally from the SCT test beam [78]. The two-crate setup allows for simultaneous running of both the barrel and end-cap in separate DAQ partitions, enabling two developers to work on the system in parallel.

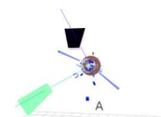
In addition, there is a single TTC crate in the rack area which can distribute clock signals and resets and L1A triggers to the two crates. A series of logic cards are also present, which veto triggers around clock resets. This is a requirement of the modules, and in the SCT detector itself is done during the long gap between LHC bunch orbits. The timing crate allows the system to be run in physics mode, and proved invaluable in developing and testing the ROD simulator and monitoring as described in chapter 5.

The SCT DAQ software used to control and calibrate the sectors is identical to that used in Point-1, under the ATLAS Trigger and DAQ framework. The software itself runs on a number of servers and terminals located in the SR1 control room. A screenshot of the end-cap sector as it appears in the DAQ Graphical User Interface (GUI) is shown in figure 4.4b.

4.6 Calibrating the Sectors

Many of the SCT calibration scans can be developed, tested and run using the sectors. As an example, the plots in figure 3.8 show the characteristic scans from one of the barrel sector modules.

This section presents the results of running standard calibration scans on sector modules. As the modules used for the test sectors did not pass quality assurance tests



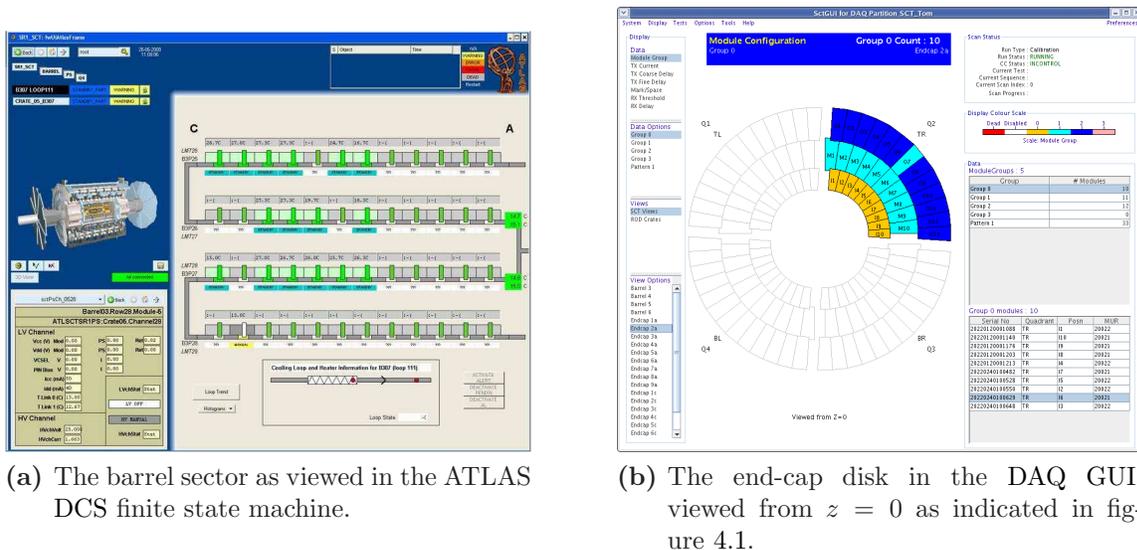


Figure 4.4: The SCT sectors as seen by the DAQ and DCS software.

for the real detector, many have problems that makes running the full calibration chain difficult. However, this also makes them excellent for ensuring the DAQ system is robust against such problematic modules. In total, 47 modules on the two sectors were available to run the entire chain of calibration scans. The remaining modules were excluded due to a variety of problems, such as readout issues, high voltage trips and chip errors.

Figure 4.5 shows optimised parameter distributions after the calibration tests described in section 3.6. The strobe delay test shows a mean value of 18.6 ns, a little lower than the value for the full SCT, which is 19.8 ns. After a three-point gain test, the mean sector module gain was 57.1 ns/fC, consistent with SCT values of 56 ns/fC. The three point gain test also yielded the ENC, with values of 1528, 1631 and 1116 electrons for the barrel, outer or middle end-cap and inner end-cap modules respectively. These values are in excellent agreement with those measured from the SCT, whose values are shown in figure 3.9.

In addition to the three-point gain test, a response curve was also run to fully calibrate the threshold of the modules. The optimal response curve parameters were applied to the modules, before performing a final noise occupancy test. This yields results of 7.2×10^{-5} , 1.3×10^{-4} and 2.5×10^{-5} , again for the barrel, outer/middle end-cap and inner end-cap modules respectively. These noise values are ~ 3 times higher for the barrel and outer end-cap modules, and $\mathcal{O}(10)$ times higher for the inner modules than the SCT itself. This difference can be attributed to the higher temperature of the

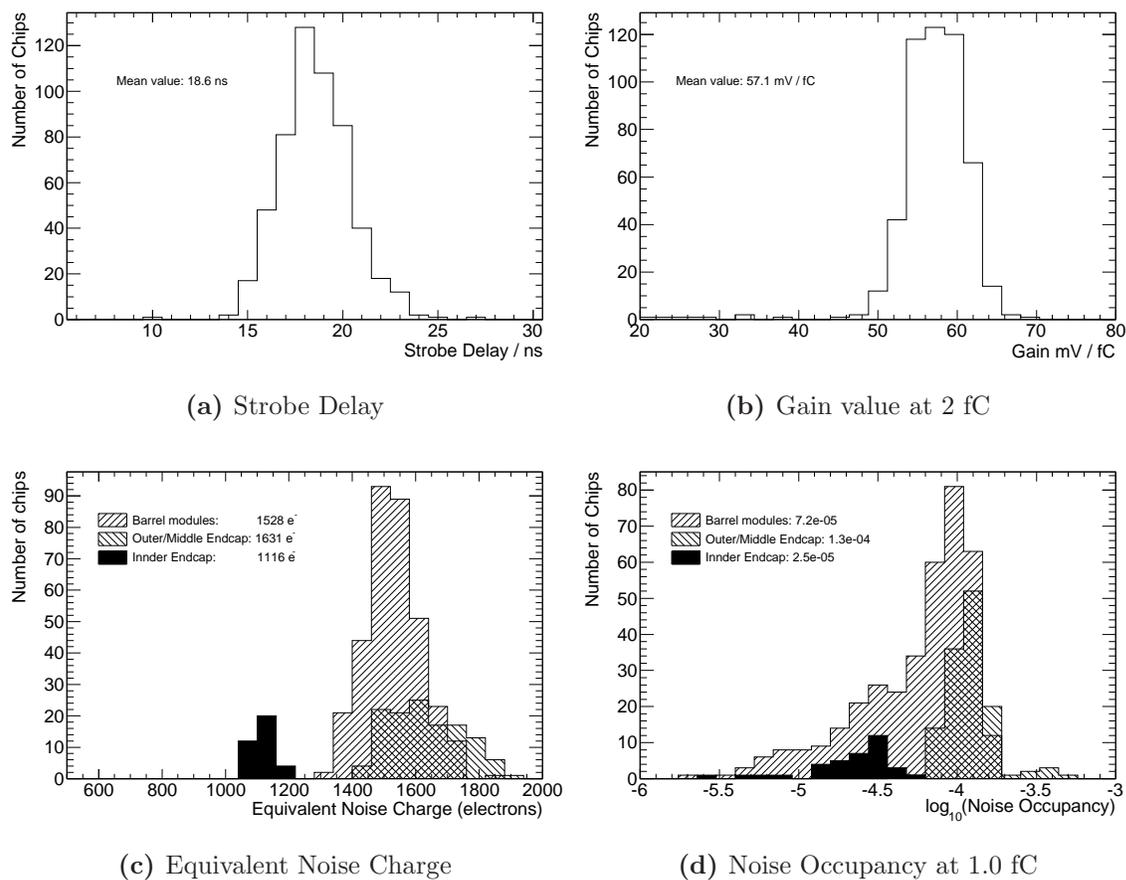
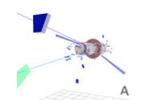


Figure 4.5: Characteristic calibration parameters for 47 modules from the SR1 sectors.

sector module. Due to constraints on the SR1 compressor, it is only possible to cool the modules to a temperature of $+15\text{ }^{\circ}\text{C}$ during operation.

4.7 Conclusions

The SCT barrel and end-cap test sectors in the surface buildings around Point-1 were constructed and commissioned during 2008-09. Their use as a DAQ development platform has proved to be extremely useful for the testing of new hardware and software features, some of which will be described in the following chapter.



Chapter 5

SCT Data Acquisition Developments

*“A stream of numbers hit a screen,
and you’re expected to know what they mean.”*

— Maximo Park, *Our Velocity*

In 2006, the SCT detector was installed in its final position in the ATLAS experimental hall at Point-1. From then until first LHC collisions, effort was focused on calibrating and preparing the detector for data taking.

During SCT commissioning, a number of developments were made to the DAQ allowing it to evolve into a faster and more reliable system. These include reduction in ROD configuration time and improving robustness to scan errors. Larger projects included development of an FPGA data simulator, and a DSP based monitoring framework.

This chapter represents a summary of the author’s main contributions to the SCT DAQ system.

5.1 Simultaneous Loading of Slave Firmware

The master DSP firmware takes the form of a binary file which must be loaded to the ROD before use, where it will remain until overwritten. The slave DSP firmware, however, must be loaded each time the ROD is configured. The firmware is stored as

a set of binary files which are copied over to each ROD in turn by the SctApi, via the MDSP. Previously, the same files had to be loaded to the ROD for each of the four SDSPs. To decrease configuration time, the MDSP firmware was modified to load the slave binaries to each SDSP in turn, thereby only sending one set of files to each ROD. This modification reduced the overall configuration time by ~ 1 s per ROD, an improvement of 10% compared to loading each SDSP in turn.

5.2 Errors During Scans

During a calibration scan, errors in the data stream may occur due to badly tuned optical links or noisy module strips. When calibrating large numbers of modules, it is important that scans are robust against these errors.

Events which contain errors are not immediately histogrammed by the SDSP, but sent to a separate routine that keeps track of which links might be causing the error. If an event is marked as containing errors it is picked up by the histogram control task running on the master, which attempts to reset erroneous modules before resending triggers. If there are still error events being sent after a fixed number of resets, the error handling routine will mask links until an error free event is received. This allowed calibration scans to run much more reliably in the presence of large numbers of modules. Information regarding the location and types of error are published to IS, where they can be viewed via the SCT DAQ GUI.

5.3 Masking Noisy Links

During physics mode data taking, the high data volume produced by noisy modules, chips and strips can cause the ROD to exert a BUSY signal which propagates to the ATLAS central DAQ via the TIM. During combined ATLAS runs, the only way to clear these BUSY signals is to mask off an entire ROD from the readout.

5.3.1 Hot-Link Monitoring

As an alternative to this, a more efficient masking system was developed whereby a task running on the MDSP monitors the status of all of the input links to the ROD. This

task can be started and stopped by sending primitives to the MDSP. It implements a simple FSM, polling the status of each formatter FIFO buffer. There is one such buffer per input link, and data from the modules arrive here before being derandomised and sent on to the EFB.

If a formatter buffer fills beyond a certain fraction of its maximum capacity, the link will enter a mode whereby only the header and trailer are stored (Header-Trailer Limit (HTL)), and the hit information is discarded. If the buffer continues to fill, then the formatter will assert a BUSY, which will halt triggers.

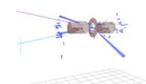
To prevent this happening, the MDSP will mask any links that enter the HTL so that no data are received from the noisy link. Instead, a warning flag is raised in the data stream, marking that link as having been masked. This flag is encoded as a high value in bit nine of the link header that is sent to the S-link, as shown in table 5.1.

Bits	Meaning
0-2	001 (fixed value)
3	Preamble error
4	Timeout error
5	L1ID Error
6	BCID Error
7	Condensed mode
8	Link Masked by DSP
9-15	Link Number

Table 5.1: The meaning of bits in the link header word, with the masked link bit highlighted.

In this way, the ATLAS offline software can access which links are active in the detector, in order to make accurate efficiency measurements and perform track reconstruction. In addition, a record of masked links is kept on the MDSP and transferred to the crate server, as a notification to the DAQ operator.

The task can also be configured when it is started, via the input primitive sent to the ROD. For example, the polling period can be set in milliseconds, with a limit on the maximum frequency obtainable at ~ 30 Hz.



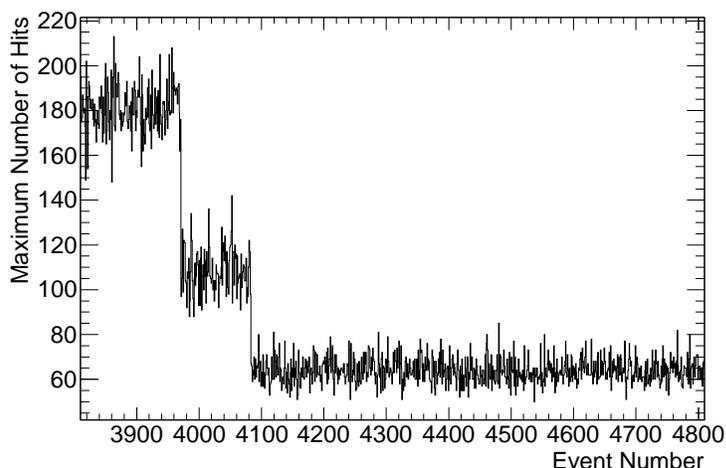


Figure 5.1: The number of hits for the noisiest module in the SCT against event number, for run 89335. The steps in the number of hits occur when the noisiest module is masked out of the run, which occurs here for two modules.

Control of the hot-link monitoring task was integrated into the SctApi, so that tasks are started on each ROD in a crate automatically at the beginning of a physics mode run.

5.3.2 Examples of Use

The hot-link monitoring proved particularly useful in the initial phase of global ATLAS data taking. During summer 2008, before the SCT had been fully calibrated, there were a number of noisy modules which would have consistently caused the SCT to enter a BUSY state. Figure 5.1 shows an example of such a run, during the course of which two noisy modules are masked out.

5.4 ABCD3T Simulation

During SCT commissioning in 2008, it became apparent that a hardware-based event simulator would be extremely useful in testing the SCT DAQ chain during periods when the detector itself was inoperative. With this incentive, an SCT ROD simulator was developed, based on a similar design by the Pixel group [79].

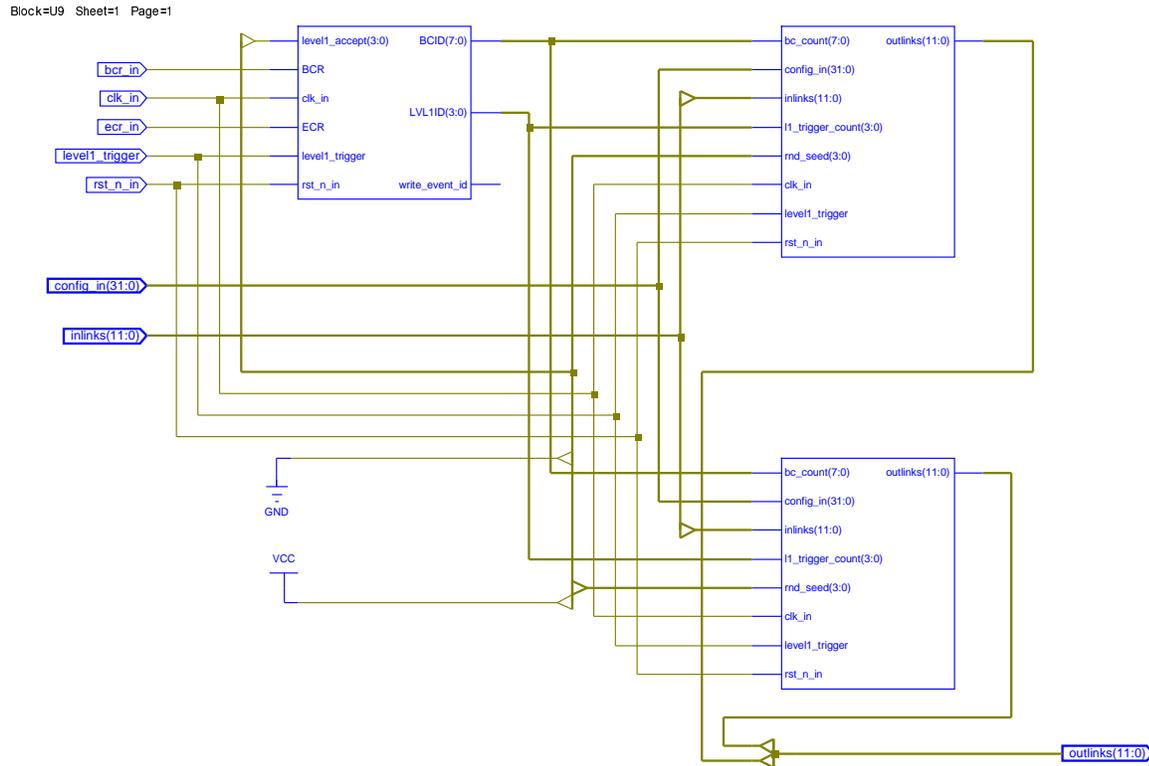
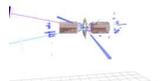


Figure 5.2: The SCT Formatter simulation block, here shown with two independent simulation engines on the right hand side. The left hand block shows the Event Counter handler, which generates L1ID and BCID numbers.

5.4.1 Simulator Design

The earliest stage at which simulated events can be injected into the SCT readout chain is in the Formatter FPGA, where each of the input links from the BOC can be replaced by internally generated data. The simulator itself has been written in Very High-Speed Integrated Circuit (VHSIC) Hardware Description Language (VHDL) as an additional block in the Formatter, and consists of a number of sub-components, as shown in figure 5.2. The main components are the simulator engine, which generates SCT format event data, and the Event Counter handler, which internally generates Bunch Crossing and Level 1 ID numbers to be inserted into the data-stream.

Each simulation engine generates two streams of data, corresponding to two sides of a module. When compiling the VHDL source, the number of engines can be specified.



If a single engine is specified, identical data are created for each of the six pairs of links associated with that formatter. If required, it is possible to include up to six engines at compile time, each producing separate data streams. The trade-off for increased variety is the space usage on the FPGA, as shown in table 5.2. It was decided that a design containing two formatter engines provided a good balance between variety and utilisation.

# Engines	Slice Usage
1	67%
2	71%
6	85%

Table 5.2: The dependence of occupied slices on number of simulation engines. The fraction of occupied slices gives an indication of how much space is used on the FPGA. A slice in this context refers to a collection of basic logic structures, such as lookup tables, flip-flops and memory.

The operation of the simulator is controlled by a single 32-bit simulator register, with table 5.3 showing how to configure the simulator for various different data-generation modes. Table 5.4 shows the location of the simulator register for each formatter chip on the ROD.

5.4.2 Event Counter Generation

One of the main challenges of implementing an SCT data simulator was matching the event counters to those from the incoming trigger. On the modules themselves, each ABCD3T chip contains a four-bit L1ID counter and an eight-bit BCID counter, which is written to the header of each event fragment sent to the ROD. Both of these counters can be reset by software or hardware resets. The L1ID can also be reset by an ECR, corresponding roughly to a luminosity block, and the BCID can be zeroed by a BCR which occurs every LHC orbit at a rate of 11.1kHz.

To ensure synchronisation is maintained, the ROD EFB checks that the counters sent match the ones received, and if not, the fragment is marked with a corresponding L1ID or BCID error. The formatter itself has no such internal counters, and so in order to produce an error-free data stream, one needs to be implemented within the simulator.

Bit	Value	Meaning
0	0	Simulation disabled
	1	Simulation enabled
1	0	“clock/2” mode disabled
	1	“clock/2” mode enabled
2	0	Debug mode disabled (Write BCID and L1ID to hit data)
	1	Debug mode enabled
7-3		Custom Header (field value translates to link header)
9-8		Hit Probability
	00	Always send a hit
	01	Half hit / half empty
	10	Mostly empty, some hits (1 per strip every 128 events)
	11	Always an empty event
11-10		Error Probability
	00	Always send hit/empty, so do nothing
	01	Mostly OK, but some errors
	10	Half OK / half error
	11	Always an error
15-12		Number of chips (0000 = 1 chip)
22-16		Number of Hits/Chip
25-23		Hit Map
	000	Fixed at 011
	001	Fixed at 010
	010	Detector alignment 1XX
	011	Detector alignment X1X
	100	Detector alignment XX1
	101	Level mode X1X
	110	Edge mode 01X
	111	Test mode XXX
27-26		Cluster Size
29-28		Bit Flip Probability in Link 0
31-30		Bit Flip Probability in Link 1
	00	No bit flips
	01	Flip 1 in 4 bits
	10	Flip 1 in 2^9 bits
	11	Flip 1 in 2^{19} bits

Table 5.3: The SCT formatter simulation configuration register.



Chip Number	Register Location
0	0x004000B0
1	0x004004B0
2	0x004008B0
3	0x00400CB0
4	0x004010B0
5	0x004014B0
6	0x004018B0
7	0x00401CB0

Table 5.4: The location of the simulator registers on the SCT ROD.

Command Control Protocol

The commands for L1A triggers, ECRs and BCRs are sent from the ATLAS CTP, propagating through the ROD via the control command serial stream, with the format as shown in table 5.5. As these commands are not normally required on the formatter, a

Type	Field 1	Field 2	Description
Level 1	110		Level 1 Trigger
Fast	101	0100	Soft Reset
		0010	BC Reset

Table 5.5: The trigger and fast commands of the SCT Module command protocol.

special decoder block was added, which converts them into three separate signals. These signals then form the input to the simulator event counter block.

Event Counter Block

The Event Counter Block contains a four-bit L1ID counter, and an eight-bit BCID counter. Both counters are reset if a soft-reset is received, while the BCID is also reset on a BCR signal. The L1ID is incremented on receipt of a L1A signal, and the BCID is incremented with each 40 MHz clock cycle. These values of the counters can then be written to the header of the simulated data stream, ensuring synchronised, error-free events are generated.

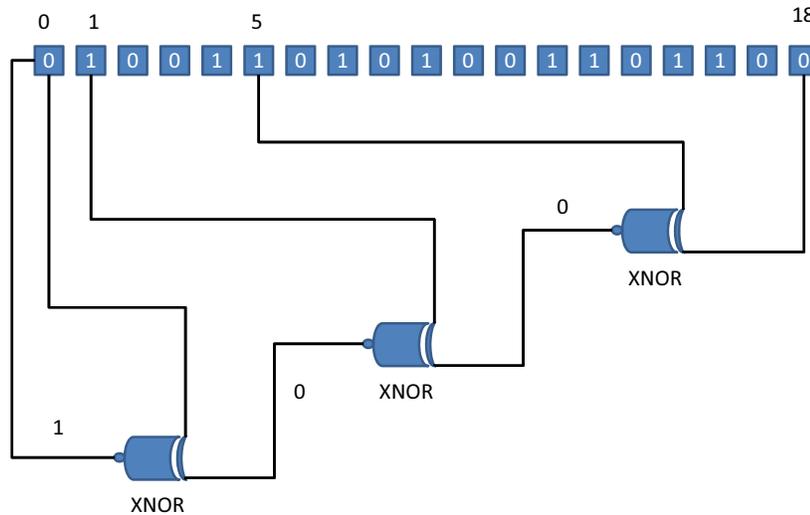


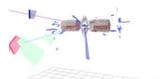
Figure 5.3: A 19-bit XNOR LFSR, with taps at bits 0, 1, 5 and 18.

5.4.3 Linear Feedback Shift Register

In order to produce a variety of different events and hit patterns in simulated data, a hardware based random number generator was implemented within each simulator module. One of the simplest methods for generating-pseudo random numbers is by using a Linear Feedback Shift Register (LFSR). In general, an LFSR takes the initial value and shifts the bits one bit to the right, with a new leftmost bit generated from a linear combination of previous bit values. The bit positions which effect the next state are known as the taps, and certain combinations of taps give maximal length generators, which cycle through all possible values without repetition.

In this case, a 19-bit maximal-length LFSR has been constructed using XNOR gates [80], as shown in figure 5.3. Such a generator has been implemented in software to demonstrate its functionality, with some sample output shown in table 5.6 and figure 5.4.

The LFSR used in the simulator is the same design as used for the Pixel simulator, with the additional feature that the initial random seed can be specified. This allows



Iteration	Binary Number	Decimal Number
0	0011011001010110010	111282
1	0110110010101100101	222565
2	1101100101011001011	445131
3	1011001010110010110	365974
4	0110010101100101101	207661

Table 5.6: The first 5 numbers generated with a 19-bit XNOR LFSR, with a seed of 111282.

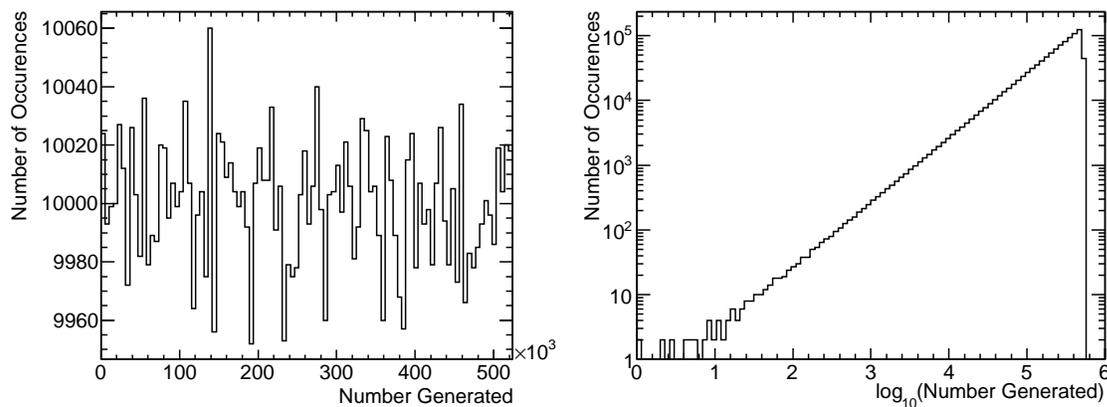


Figure 5.4: The output of a software-implemented 19-bit XNOR LFSR. The initial seed was set to 111282 and 10^6 random numbers were generated.

multiple simulation engines per ROD to generate different random sequences, and hence produce a unique data streams.

Various bits from the LFSR are used in the Simulator Engine to generate hits in random strips, and used to make random decisions regarding behaviour of the simulator FSM. This randomness is extremely useful, allowing the simulator to produce a wide range of different data patterns for testing the DAQ.

5.4.4 The Simulator Engine

The purpose of the simulator engine is to produce a bit stream of SCT-like data, which can be configured to generate events with hits of varying occupancies, time bins and error rates. This is implemented as a FSM, which remains in an idle state until a Level-1 trigger signal arrives.

After a trigger signal is received, the FSM will move to sending an event header, which although nominally 11101, can be set to a custom value to test the robustness of the DAQ system to noisy optical data. Following the header, the event identifier is sent, which consists of the L1ID and BCID numbers strobed from the event counter block, as described in section 5.4.2.

Once the header and counter information has been sent, the simulator then generates a number of data blocks depending on the settings in the configuration register. Events can be generated either with hits every trigger, entirely empty events, or a mixture of hits and empty events. If an event occurs with a hit, then a fraction of these events can be set to generate error data instead. Finally, a trailer word is sent, which completes the FSM cycle, returning it to an idle state. The overall data pattern has the following format:

$$\underbrace{11101}_0 \underbrace{lll}_{\text{L1ID}} \underbrace{bbbbbbb}_1 \underbrace{ddddddd}_{\text{Data Blocks}} \underbrace{1000000000000000}_{\text{Trailer}} \quad (5.1)$$

where the most significant bit is to the left. A brief discussion of the possible data blocks will be discussed in the following sections. For a more complete description of SCT module data formats, see the ABCD3T manual [81].

Empty Data

For empty events, the simulator simply generates a no hit data packet

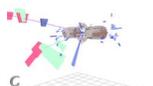
$$001 \quad (5.2)$$

before moving on to sending a trailer.

Hit Data Packets

A data packet can also be generated containing hits, with the following format

$$01 \underbrace{aaaa}_{\text{Chip Number}} \underbrace{ccccc}_{\text{Strip Number}} \underbrace{1 \underbrace{ddd}_{\text{Hit Pattern}} \quad 1 \quad \underbrace{ddd}_{\text{Hit Pattern}}}_{\text{Cluster}} \quad (5.3)$$



A number of different data packets can be produced per link for each event. The first loop is over each chip associated with a particular link, up to a total of six. In fact, two independent data streams are generated for each engine, simulating data from the two sides of an SCT module, denoted as link 0 and link 1. The first stream will produce data for chips from 0–5, whereas the second will produce events with chip numbers 13–8. This configuration produces events with hits in chips on opposite sides of a module. Spacepoints are therefore more likely to be generated, which is useful for offline reconstruction and monitoring algorithms.

In order to produce some variation in the data stream, the strip number is generated by the random number module, as discussed in section 5.4.3. Again, the strip number for link 1 is generated such that it will produce data corresponding to a spacepoint. An additional debug mode is available whereby the L1ID and BCID numbers are reported as the chip and channel, to check for synchronisation errors. Once the FSM has generated the chip and channel number, it moves on to generating a cluster of hits, consisting of a synchronisation bit (set high) and the three bit time-bin pattern described in section 3.4.2. These bits are also configurable, as it is an important variable to monitor accurately when trying to calibrate the overall timing of the SCT. An example of the simulator working to produce different time bin information is shown in figure 5.5.

The number of hits in the cluster can be set by the user, as can the number of clusters per channel. It should be noted, however, that real SCT modules read out hits from sequential strips, and so a series of clusters generated in non-ascending order will be flagged with a non-sequential chip error.

Error Data Packets

In addition to generating hit data, the simulator can also be set up to generate a set fraction of error events with the following format

$$\underbrace{000}_{\text{Leader}} \underbrace{aaaa}_{\text{Chip Address}} \underbrace{eee}_{\text{Error Code}} \underbrace{1}_{\text{Separator}} . \quad (5.4)$$

On a real module, this will occur if the chip has not received a L1A trigger (code 001), if there is a buffer overflow (code 010) or a buffer error (code 100). Being able to simulate error events allows testing of the DAQ system’s robustness to such errors, and allows the monitoring to be tested in a controlled way.

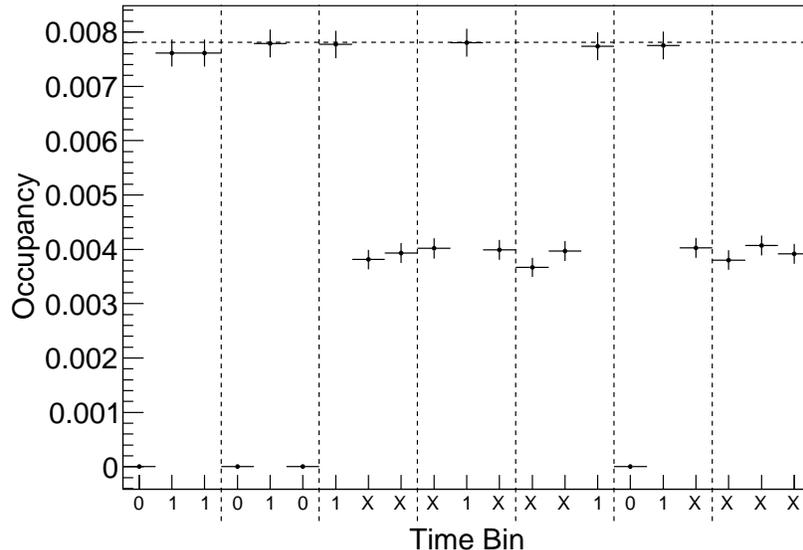


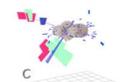
Figure 5.5: Measured occupancy (hit rate) of the simulator for a number of different time-bin settings. The horizontal line indicates the expected noise from the simulator, which was set to generate one hit per chip per trigger. For each 'X', the simulator will generate equal numbers of 1 and 0.

5.4.5 Flagging Simulated Events Offline

As the simulator can, in principle, be enabled at any time during an SCT physics mode run, it was also necessary to mark events which contain simulated events, so that they are not confused with real data-taking runs. In the ATLAS event format header, which is generated by the ROD before being sent via the S-link to the ROS, there is a data word reserved for the detector event type. Specific event types can be set in the EFB FPGA, using the register denoted `DFLT_ROD_EVT_TYPE` [82]. It was decided that bit 5 of this register would be used to mark simulated events, which is done when the user enables the simulator in the SCT GUI.

5.4.6 Examples of Simulator Use

Allowing the SCT ROD to produce events with specific size and error properties has been very useful in testing the robustness of the system to large, noisy and erroneous events. In addition to testing the ROD-Level Monitoring, as described in section 5.5, and many other DAQ developments, there are two particular scenarios that highlight the simulator's use.



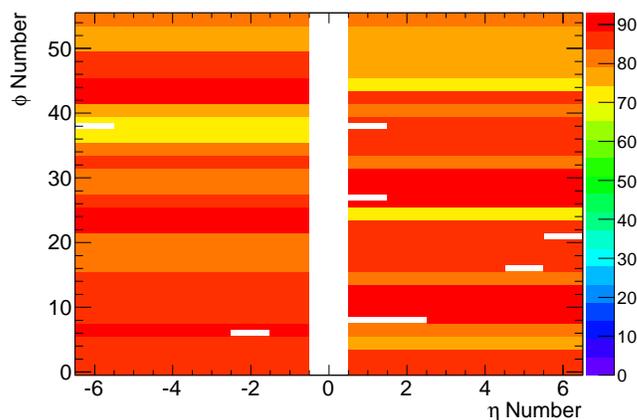


Figure 5.6: An SCT Monitoring plot of the hit-map for barrel 6, side 0, using simulated data to generate 1 hit per chip every 128 triggers. The striped pattern is due to the simulator producing the same data stream for each ROD. White modules are excluded from data-taking.

Global ATLAS Running

During times when the SCT itself was not available, the simulator was used to test the SCT DAQ chain. After a serious cooling plant incident in early 2008, the simulator was used to enable the SCT to take part in a number of global ATLAS milestone runs, allowing validation of software changes and the successful testing of 70 kHz triggers. Figure 5.6 shows an example of a monitoring plot generated during a global ATLAS run whilst using the SCT simulator.

Debugging Lost Fragments

During the first LHC injection tests in September 2008, it was noticed that there was occasional loss of ROS data in parts of the SCT. This was most pronounced during beam splash events, and in some cases data for an entire end-cap were missing.

The problem was traced to insufficient memory allocated to processes running on the SCT ROS which build events from a number of separate ROD fragments. If an event is received that is larger than this memory buffer, it is simply discarded. This phenomenon was successfully reproduced in a controlled manner using the SCT simulator to generate events with unusually high occupancy.

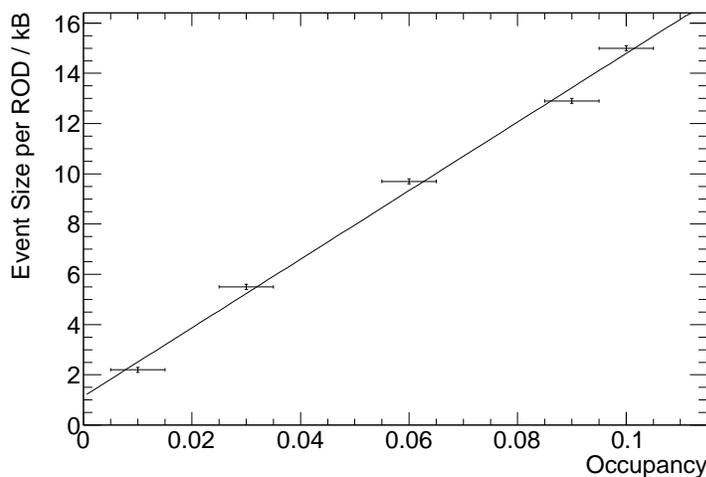


Figure 5.7: The ROD event fragment size against the occupancy as set with the simulator. The linear fit has parameters: $\text{size/kB} = 140 \times \text{Occupancy} + 1.1$.

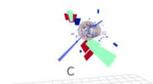
In order to set a safe buffer limit for SCT operation, events of different occupancies were generated using the simulator. The corresponding ROD event size was measured for a given occupancy, as shown in figure 5.7. The linear fit to these data has a gradient of

$$\frac{\text{Event Size}}{\text{Occupancy}} = 140 \text{ kB.} \quad (5.5)$$

The non-zero offset of 1.1 kB is due to the fact that headers and trailers are required by the event format specifications even if there are no hits present.

The new memory size was set to 15 kB after September 2008, a factor eight increase compared to the previous value. This allows running of the SCT up to an occupancy of 10 % before fragments are lost, and is a safe limit considering the occupancy is expected to be < 0.1 % during collision data-taking. This was then further increased to 64 kB before November 2009 to ensure no fragments were lost during the first beam splashes after the LHC restart.

As a follow-up, the SCT end-cap also participated in data-taking for the first LHC operation of 2010. During a beam loss event on the morning of 28th February, five RODs out of 45 saw lost fragments, corresponding to an occupancy exceeding 47%.



5.5 ROD-Level Monitoring

During calibration scans, the SCT ROD DSPs are used extensively to histogram module data in order to optimise performance. During data-taking, the DSPs are mainly idle, only being used to configure the modules at the start of the run.

The aims of the ROD-Level monitoring were to utilise this computing power to facilitate a monitoring framework for the SCT, providing measurements of occupancy, spacepoint rate, timing and errors. The advantage of such a system is that every L1A triggered event can be monitored, and so statistics will be high and can be collected quickly. A second advantage is that by using existing SCT DAQ analysis infrastructure, derived monitoring variables can be quickly calculated and fed back to the user within minutes of a scan completing.

A number of modifications and additions to the SCT DAQ software were required to enable the ROD to histogram events during physics mode.

5.5.1 DSP Histogramming

As described in section 3.5.1, event data fragments passing through the router FPGA on the ROD can be copied into one of four event trap buffers, each connected to a slave DSP. The DSPs can be configured to run a histogramming task which copies event fragments from the router to a larger section of memory. Each event fragment contains a trailer with a marker that is set by the ROD if the event contains an error. The SDSP histogramming routine checks this bit, sending the event to an error decoding routine if errors are detected, or an occupancy histogramming function if the fragment is error-free. In the case of the occupancy routine, the slave DSPs have the ability to histogram the following quantities.

Strip-Level Occupancy

A histogram is stored in the DSP memory for each module on the ROD, with a bin for each module strip. Every hit over threshold is read-out from the SCT modules and is marked with the relevant chip and strip number. The histogramming routine decodes this information from the event fragment and keeps count of the number of hits recorded

by each strip. The number of triggers is also recorded, with the occupancy calculated as

$$\text{occupancy} = \frac{n_{\text{hits}}}{n_{\text{triggers}}} \quad (5.6)$$

If the formatter is running in expanded mode, the three-bit timing information is used to fill three separate histograms for each module, corresponding to hits in bunch crossings before, coincident with, and after the trigger, as described in section 3.4.2. Histogramming of occupancy in different time bins is a useful measurement of how accurately each module is synchronised with the central ATLAS trigger.

Pseudo-spacepoint Reconstruction

The SDSP also has the ability to perform rudimentary spacepoint reconstruction. Spacepoints occur when a particle passes through a module and induces hits in strips on both sides. For a given strip, the total number of overlapping strips on the opposite side of the module is given by:

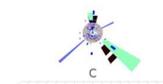
$$n_{\text{strips}} = \frac{l \tan \theta}{w} \quad (5.7)$$

$$= 64 \quad (5.8)$$

where $l = 128$ mm is the length of a strip, $\theta = 40$ mrad is the stereo angle between the two sides, and $w = 80\mu\text{m}$ is the width of a single strip.

A full spacepoint reconstruction would therefore require checking of 64 strips adjacent to each hit. As this is computationally expensive for the DSP, the number of coincident chips is computed as an approximation. In this algorithm, the total number of hits per chip on both sides is counted for each event. For each chip with a hit strip on side 0 (chips 0–5), a coincidence is recorded if there is also a hit in the corresponding chip on the opposite side. As a hit may occur in a strip on the edge of a chip, checks for hits are also made on the two chips adjacent to the one directly opposite. This is illustrated in figure 5.8, and the chips used in the coincidence calculation listed in table 5.7.

Each chip is responsible for reading out 128 strips, which compared to the 64 strip overlap from equation 5.8, means that the coincident chip approximation will successfully find all real coincidences. There will also be some additional noise due to fake spacepoints, as discussed below.



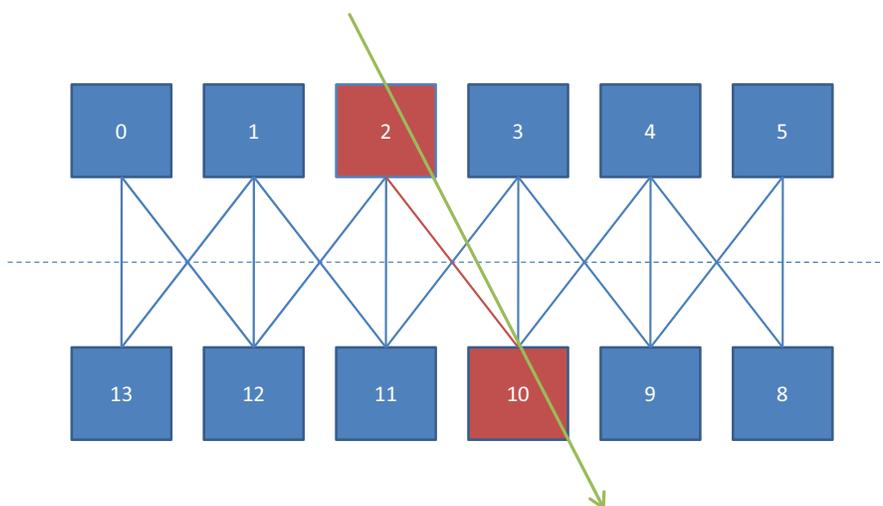


Figure 5.8: Diagram showing the numbering of ABCD3T chips on side 0 (top) and side 1 (bottom) of an SCT module.

Chip (side 0)	Coincident Chips (side 1)	Non-coincident Chips (side 1)
0	12, 13	8, 11
1	11, 12, 13	8, 9, 10
2	10, 11, 12	8, 9, 13
3	9, 10, 11	8, 12, 13
4	8, 9, 10	11, 12, 13
5	8, 9	10, 13

Table 5.7: Chips used to determine coincidences for the pseudo-spacepoint reconstruction, and non-coincident chips used for noise-subtraction. The chip numbering scheme is the same as in figure 5.8

Using this algorithm, the number of spacepoints per event is histogrammed, up to a maximum of 16 in a single event, due to the possible chip combinations if every chip contains at least one hit strip.

In expanded mode, the number of spacepoints is calculated for hits in each of the three individual time bins. In addition, coincidences are also recorded for hits with the 01X (= 010, 011) time structure and coincidences with hits in any time bin.

Noise Subtraction

During offline SCT analysis, noise hits and real hits can be distinguished by fitting a track to a series of hits, and focusing solely on the hits associated with a track. The SDSP has no knowledge of the geometrical layout of the modules it is reading out, and module events are distributed across several RODs, making tracking impossible at this level, and noise separation more of a challenge.

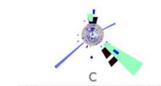
In order to estimate the coincidences that are due to noise, a calculation can be made using only chips that have no overlapping strips. The chips used in this non-coincident calculation are listed in table 5.7. Any coincidences from these combinations of chips will always be fake, and can be attributed to noise. This noise rate can then be used as a pedestal to subtract from the rate obtained from coincident chips.

It is possible to have an event where two particles cause hits in two different chips of one module. In this case, the coincidence algorithm would register two real spacepoints, but also two noise-spacepoints, which would cancel each other out. To prevent this from happening, the non-coincident calculation is only counted for events with no real coincidences in a particular module.

5.5.2 Crate-Level Modifications

A number of modifications were required to both the DSP firmware and the SctApi software to adapt the histogramming for physics mode running, as described below.

A key modification was to prevent the MDSP from starting the histogram control task. In calibration mode, this task is responsible for sending triggers to the modules, which is unnecessary in physics mode, as triggers are generated externally. The histogram control task also masks off links with too many errors in calibration mode, (section 5.3), which is also undesirable during physics runs. Instead, the histogramming tasks are set up solely on the slaves, which will count the number of hits and errors until read out is prompted by the SctApi.



During calibration scans, the modules are separated into four groups of approximately equal numbers, with triggers sent to each group successively. Events from each group are distinguished by an internal trigger type, which is detected by the router and allows events from different modules to be distributed across the four slave DSPs. In physics mode, such a grouping is not possible as triggers are sent to the modules simultaneously, and pass through the router as a single event fragment. For this reason, only a single event trap and corresponding slave is used to histogram the data on each ROD. This could be modified in the future to distribute the events across all four slaves, with perhaps each slave collecting a different trigger type, but would require an extensive rewrite of the raw histogram data handling on the SctApi side.

An additional modification to the physics mode histogramming was to prevent the readout of a histogram until after the router traps had been stopped. In calibration mode the sending of triggers is controlled by the ROD, and no triggers are sent during histogram readout, so this issue was not relevant in the past. In physics mode, event fragments still pass through the router after histogramming on the slave has been stopped. As the slave memory where the histogram data are stored only has one access path, it cannot be written to and read out at the same time. This can cause problems if the histogramming task is still running during readout. The simple solution to this was to stop the event traps in the SctApi before readout.

Two distinct configurations are available, automatic and manual. For automatic monitoring, the SctApi regularly polls the state of the DSP event counters to check how many triggers have been collected. After a set number of triggers has been collected, a primitive is sent to the RODs requesting the histogramming moves to the next bin. To prevent overly long monitoring periods in the case of an unexpectedly low trigger rate, the bin will also be changed if the number of triggers has not been reached after a specified timeout. Once the required number of bins has been filled, primitives are sent to read the histograms out, and the monitoring terminates. In the manual configuration, however, it is up to the user to change bins and initiate histogram readout, via a custom GUI. This option is useful if one-off monitoring is required or if bins need to be changed at irregular intervals.

ATLAS Event Type

The router event traps can be configured to capture either all events passing through the S-link, or only events with a particular ATLAS event type. This event type is an 8-bit

number assigned to each L1A trigger by the CTP, denoting the trigger's origin. The meanings of the bits in this word are shown in table 5.8. All of the physics triggers have bit 7 set high. For example, a random trigger would be selected with 0x81, cosmics with 0xC0 and minimum bias with 0xA0. This functionality was added to the SCT DAQ histogram options to allow trapping of all triggers together, or just a single event type. Most of the cosmic runs described in the following sections were taken simultaneously with a 30 kHz random trigger. Without the ability to filter out the random events, the cosmic signal would have been lost in noise.

Bit	Trigger Source
0	Random
1	Beam Pickup (BPTX)
2	Level-1 Calorimeter
3	Thin Gap Chamber (TGC)
4	Resistive Plate Chamber (RPC)
5	Minimum Bias (MBTS)
6	Cosmics
7	Physics

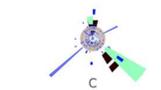
Table 5.8: Meaning of bits in the ATLAS L1A event type, correct as of December 2009.

Some modifications were required in the SctApi to allow configuration of these trap parameters at the start-up of the histogramming task.

5.5.3 Analysis Service

As mentioned in section 3.5.2, results from raw histogramming are published to IS, where they are picked up by the Analysis Service, which extracts parameters of interest. A dedicated analysis algorithm was written to calculate a number of quantities to be monitored during physics mode running.

The average occupancy at the chip level is calculated by taking the number of hits and dividing by the number of total events histogrammed. As the occupancy changes with threshold and silicon bias voltage, these parameters are also saved to the final test result.



Noise corrections are applied to the number of coincidences per event and the ratio of 01X spacepoints to the number with hits in any time-bin is calculated:

$$f_{01X} = \frac{N_{01X}}{N_{\text{Any Hit}}}, \quad (5.9)$$

which gives an estimate of how well synchronised a module is with the incoming trigger. If it is well timed-in, then this ratio should be close to unity for hits from particles. A second timing quantity is also derived from the number of coincidences in each time bin. For the bin occurring before the trigger, the coincidences are given a weight of -25 ns, and those after, a weight of $+25$ ns. Coincidences in the central bin are in-time, and so are given a weight of zero. Weighted coincidences are summed and divided by the total number of coincidences from all bins to obtain the average timing offset for each module. Modules that are well timed-in should have a timing offset that is close to zero. The two quantities are complementary, as the first gives the magnitude of the timing offset, while the second shows whether the module is reading out before or after the trigger.

The final results of the analysis are available to the SCT shifter as with any other calibration mode scan, via the SCT GUI. The final results are also uploaded to the SCT Calibration database, a mySQL database within the point 1 firewall, and mirrored on the General Public Network (GPN). The results are then available for browsing via a web interface¹, and for offline analysis. This also allows long term monitoring of occupancies and spacepoint rates over the lifetime of the SCT.

5.5.4 The SCT ROD Monitoring Application

During the development of the ROD-Level monitoring, histogramming could be started via a simple GUI. This is a useful tool if monitoring is only required over a relatively short timescale, where interesting events only occur within a narrow window. For long term monitoring, this method is inefficient as it requires human intervention, and no matter how much chocolate you give shifters, they still forget to start it.

With this in mind, an application was developed to automatically start histogramming during physics runs, based on the ATLAS TDAQ Run Control Application, allowing it to respond to changes in the run state. When a physics run is started, a separate

¹<https://pc-sct-www01.cern.ch/cgi-bin/CalibrationDB/RunList.pl?test=11>

thread is invoked which periodically starts the histogramming task via a call to the SctApi Server over IPC.

The thread then waits for the task to complete and read out before starting the next monitoring run. The thread is terminated at the end of a run, and if the histogramming is running at the time, it is immediately aborted so that the SCT can be promptly reconfigured if required.

To avoid conflicts, the application also checks whether the histogramming is currently running before starting a new monitoring period, and in this way can be restarted or killed during a run if necessary. It also provides a complement to user-invoked histogramming, as it offers semi-continuous monitoring of the SCT performance over the course of an entire run. The application can be configured using a number of command line options, which can be specified using the ATLAS TDAQ configuration database, OKS.

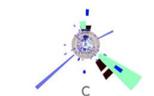
In order to provide additional flexibility, the ROD monitoring can also be configured on-the-fly by the shifter, during a physics run. The monitoring application connects to an IS server into which the user can publish control objects via a GUI. The monitoring will loop over all objects on the server in sequence, and will continue waiting if no objects are present. In this way a number of different monitoring configurations can be specified during a shift if, for example, the beam configuration changes. More detailed information on configuration of the ROD monitoring can be found in [83].

5.6 Results from ROD-Monitoring

During the SCT commissioning and operation during the late 2009 and early 2010, the ROD-Level monitoring was extensively tested and executed on a number of different occasions.

5.6.1 Noise Occupancy Measurements

The first test of the monitoring was to evaluate its performance in measuring the noise occupancy of the SCT. This could be done in both standalone and combined physics runs.



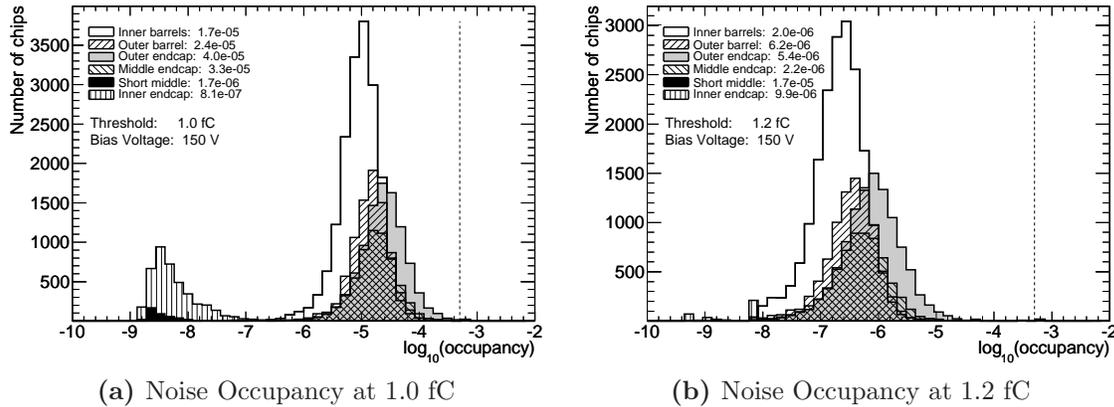


Figure 5.9: The noise occupancy of the SCT measured by the ROD Monitoring in physics mode. The dotted line shows the specification noise level at 5×10^{-4} .

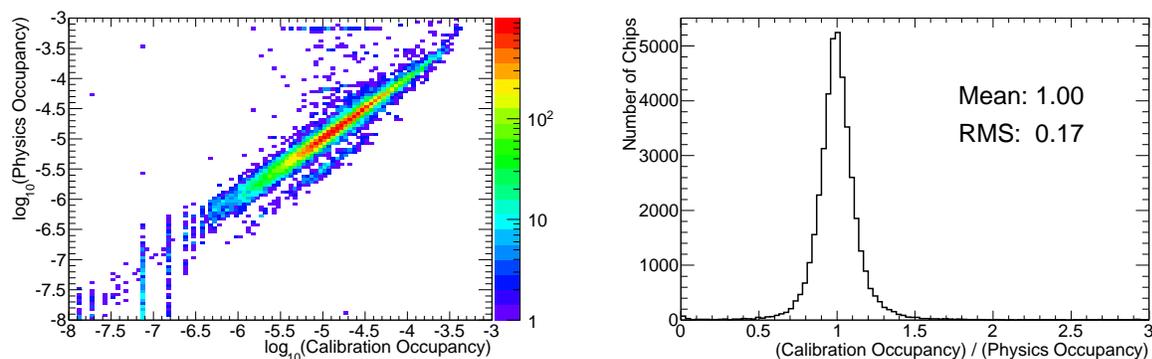
Figure 5.9 shows the noise occupancy measured by the ROD monitoring at the nominal threshold of 1.0 fC, and at the higher threshold of 1.2 fC, after extensive calibration work over summer 2009 in order to prepare the SCT for first LHC collisions. The data were taken using random triggers at rates of between 1–10 kHz.

Figure 5.9a shows the same expected noise characteristics as described in section 3.6.4, and can be compared to figure 3.9. A notable difference is that due to the high trigger rate, the short-middle and inner-end-cap modules are now visible. The noise occupancy of these modules peaks between 3–4 orders of magnitude lower than the other modules due to their shorter strip length. At 1.2 fC, the noise occupancy is reduced by around an order of magnitude, with the short-middle and inner end-cap distributions no longer visible.

Comparison with Calibration Scans

To check the consistency of the noise occupancy measurements obtained during physics mode, the results can be compared to those obtained from the calibration mode noise occupancy test introduced in section 3.6. This value should be a direct comparison to the occupancy obtained in physics mode, as shown in figure 5.10.

The ratio of measurements from the two methods gives a mean value close to 1. The relatively large Root Mean Square (RMS) spread of 0.17 is due to statistical fluctuations between the two methods. The additional structure seen in the left-hand plot of



(a) Noise occupancy measured in calibration against the same physics measurement. (b) Ratio of noise occupancy measured during calibration over the physics measurements.

Figure 5.10: Noise occupancy comparison at nominal SCT conditions.

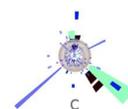
figure 5.10 can be attributed to differences in the way the noise occupancy is extracted at 1.0 fC between the two methods.

5.6.2 Coincidence Rate Measurements with Cosmics

During autumn 2009, the SCT took part in extended ATLAS cosmic data-taking runs. A number of different cosmic triggers were available, including calorimeter and muon triggers. These, however, have a relatively low inner detector acceptance. The most effective trigger for SCT cosmic studies is the TRT fast-OR, which triggers when the sum of TRT hits is above a certain threshold. Unfortunately, this trigger was not always available, and on the occasions where it was available, it is indistinguishable from other cosmic triggers at L1A. Thus most of the time the SCT cosmic track rate of ~ 0.5 Hz was swamped beneath an overall ATLAS rate of up to 100 Hz. In this case, the noise subtraction was essential to extract a real cosmic signal using the ROD monitoring.

Expected Noise Coincidence

Before attempting to measure the real spacepoint rate using the ROD monitoring, it was first necessary to estimate the rate of coincidences from noise, before the noise subtraction is performed. Referring to figure 5.8, it is possible to calculate the expected coincidence rate per event assuming random noise triggers. The probability that a chip



has at least one noise hit is

$$P(N_{\text{hits}} > 0)_{\text{chip}} = 1 - P(N_{\text{hits}} = 0)_{\text{chip}} \quad (5.10)$$

$$= 1 - (1 - p)^{3s}, \quad (5.11)$$

where p is the strip-level noise occupancy, and s is the number of strips per chip. The additional factor of three is included as we are considering noise hits in all three SCT time-bins, which it is assumed are uncorrelated for noise. The probability that a coincident chip also measures at least one hit is

$$P(N_{\text{Hits on Adjacent Chip}} > 0)_{\text{chip}=1-4} = (1 - (1 - p)^{9s}), \quad (5.12)$$

$$P(N_{\text{Hits on Adjacent Chip}} > 0)_{\text{chip}=0,5} = (1 - (1 - p)^{6s}). \quad (5.13)$$

By multiplying these probabilities, and assuming that p is small, it can be shown that

$$P(N_{\text{Coinc}} > 0)_{\text{chip}=1-4} = 27p^2s^2, \quad (5.14)$$

$$P(N_{\text{Coinc}} > 0)_{\text{chip}=0,5} = 18p^2s^2. \quad (5.15)$$

By combining the two above results for the entire module, the probability of at least one coincident noise hit on a module to first order is given as

$$P(N_{\text{Coinc}} > 0)_{\text{module}} = 144p^2s^2 \quad (5.16)$$

$$= 6.8 \times 10^{-5}, \quad (5.17)$$

taking $p = 1.7 \times 10^{-5}$ from the value for inner barrel modules. Figure 5.11 shows the ratio of the uncorrected coincidence rate measured with the SDSP compared to the predicted value as given in equation 5.16, using the module average occupancy in each case. There is a good agreement at both 1.0 and 1.2 fC, indicating the origin of noise is well modelled.

Coincidence Rates with noise-subtraction

To test the noise subtraction routine, the ROD monitoring was run in two configurations, trapping random and cosmic triggers respectively. During these runs, the SCT was at nominal voltage and threshold. The TRT trigger was not active, and so the SCT cosmic acceptance was $\sim 1\%$. The coincidence rate distributions for barrel modules are shown

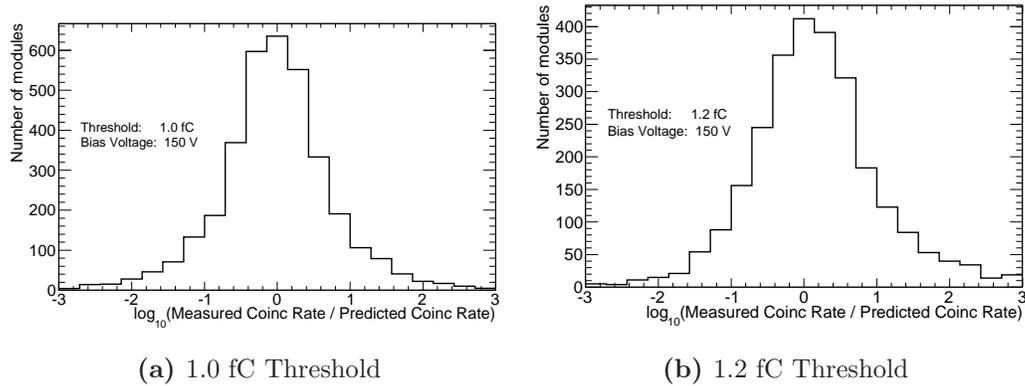


Figure 5.11: Comparison of the coincidence rate obtained using noise triggers with the predicted value from the noise occupancy.

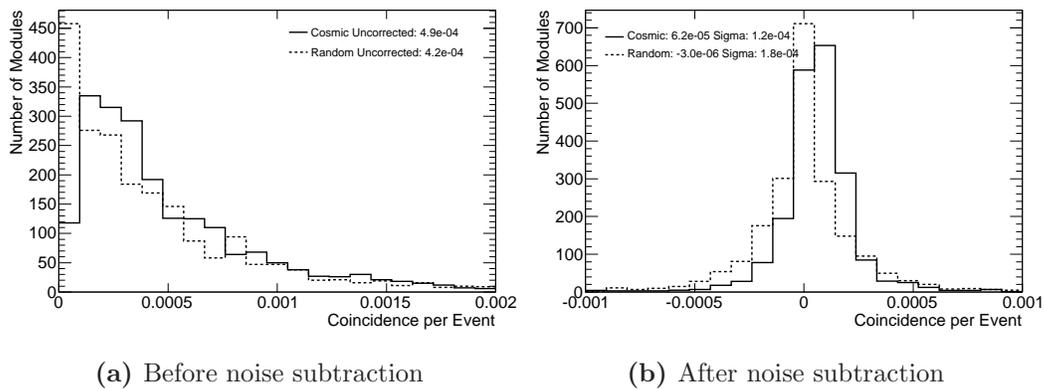
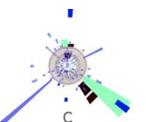


Figure 5.12: The coincidence rate per module in the SCT with cosmic and random triggers (run 138790, scans 3 and 6 respectively).

in figure 5.12, before and after noise subtraction. The corrected distributions were fitted with Gaussian curves, the parameters of which are shown in table 5.9. The expected width of the distribution due to noise variations can be calculated by applying the subtraction method described in section 5.5.1. The noise-subtracted coincidence rate is given as:

$$c = \frac{n_{\text{coinc}} - n_{\text{non-coinc}}}{n_{\text{events}}} \quad (5.18)$$



Denoting $d = n_{\text{coinc}} - n_{\text{non-coinc}}$ and assuming binomial errors, the uncertainty on c is given as:

$$\frac{\sigma(c)}{c} = \sqrt{\left(\frac{\sigma(d)}{d}\right)^2 + \left(\frac{\sigma(n_{\text{events}})}{n_{\text{events}}}\right)^2} \quad (5.19)$$

$$\approx \frac{\sigma(d)}{d} \quad (5.20)$$

$$= \frac{\sqrt{n_{\text{coinc}} + n_{\text{non-coinc}}}}{n_{\text{coinc}} - n_{\text{non-coinc}}} \quad (5.21)$$

$$\sim \frac{\sqrt{2n_{\text{non-coinc}}}}{n_{\text{coinc}} - n_{\text{non-coinc}}}, \quad (5.22)$$

where the assumption is made that the number of coincidences is small compared to the number of events, and that the rate of coincident and non-coincident chips is equal to first order. Multiplying through by c gives

$$\sigma(c) \sim \frac{\sqrt{2n_{\text{non-coinc}}}}{n_{\text{events}}}. \quad (5.23)$$

Substituting the noise coincidence rate given in equation 5.16 for $n_{\text{non-coinc}}/n_{\text{events}}$, gives the following expression

$$\sigma(c) = \sqrt{\frac{288p^2s^2}{n_{\text{events}}}}. \quad (5.24)$$

The expected width for the two monitoring runs is shown in table 5.9, and can be compared to the measured widths from the Gaussian fit. In both cases the correct order of magnitude is predicted.

The noise-corrected coincidence rate taken with random triggers is consistent with zero, as expected, and justifies use of the non-coincident chip calculation. The measured coincidence rate for cosmic triggered events is low, with a mean rate of $(6.2 \pm 0.3) \times 10^{-5}$ per trigger. A similar monitoring run was also undertaken at a threshold of 1.2 fC, with a factor 10 reduction in the noise rate compared to 1.0 fC. The cosmic coincidence rate was measured in this configuration as $(2.34 \pm 0.02) \times 10^{-5}$ per trigger, which is the same order of magnitude as the noise-corrected 1.0 fC measurement.

Module-by-module variations in the coincidence rate can be measured by the ROD monitoring. The left-hand plot of figure 5.13 shows the spacepoint map of the SCT

Parameter	Noise Triggers	Cosmics Triggers
Events	22178	90502
Mean	$(-3.0 \pm 4.2) \times 10^{-6}$	$(6.2 \pm 0.3) \times 10^{-5}$
Sigma	$(1.75 \pm 0.06) \times 10^{-4}$	$(1.15 \pm 0.03) \times 10^{-4}$
Expected sigma	2.5×10^{-4}	1.2×10^{-4}

Table 5.9: Comparison of noise-corrected coincidence rates for random and cosmic triggers.

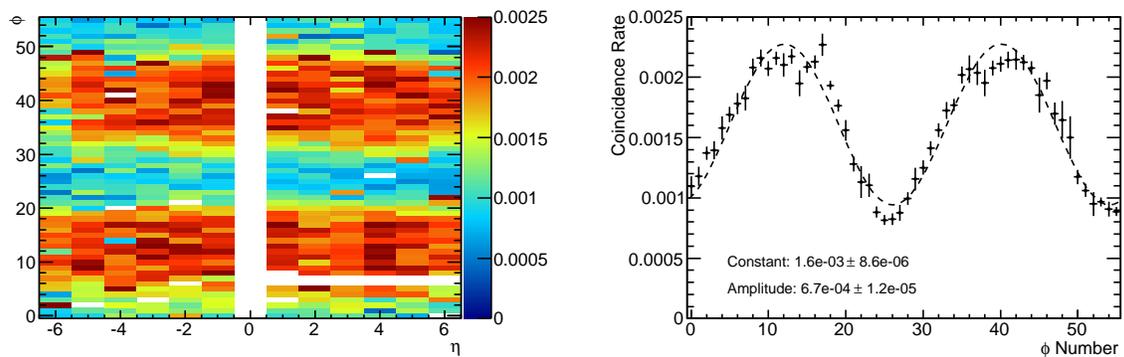


Figure 5.13: Coincidence rate map with cosmic triggers for barrel 6 (left). The right-hand plot shows the mean over all η values, with errors displaying the RMS. The other barrel layers show similar distributions. (From run 145720, scan 3)



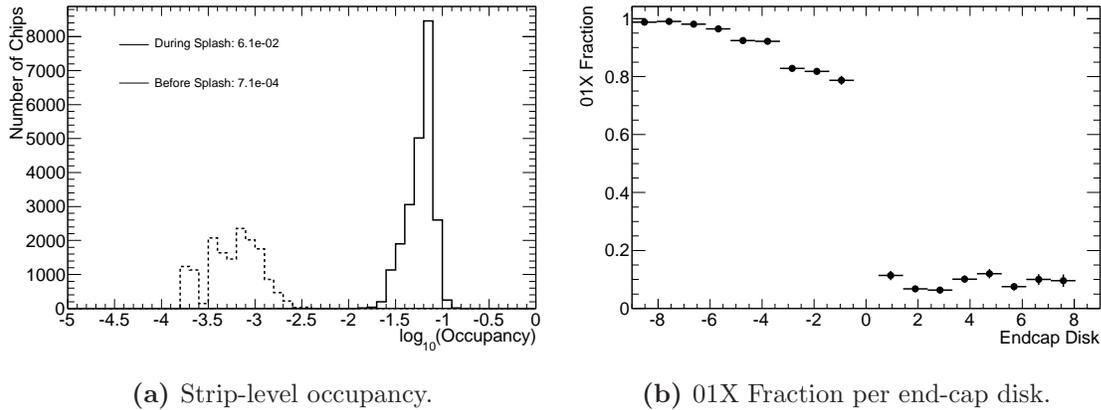


Figure 5.14: A 2009 ATLAS beam splash event as recorded by the ROD monitoring, run 140370.

barrels during a long monitoring run of around four hours. The right-hand plot shows the coincidence rate as a function of the SCT ϕ co-ordinate. A clear variation in ϕ can be seen, with maxima at the top and bottom of the barrel, as expected assuming cosmic travel vertically (in the $-y$ direction) through the SCT. A sinusoidal fit to the data gives a constant rate of $(1.6 \pm 0.01) \times 10^{-3}$ with an amplitude variation of $(6.7 \pm 0.1) \times 10^{-4}$. The mean rate is higher than the previous measurement thanks to the presence in this run of the TRT cosmic trigger, which gives a higher SCT acceptance of $\sim 10\%$. The rate remains non-zero even at the barrel edges as cosmic rays do not fall exactly vertically.

5.6.3 Beam Splash Measurements

The ROD monitoring was active during the LHC beam splash events of 2009 (section 3.7.2), and provided fast feedback of the SCT timing. The left-hand plot of figure 5.14 shows the occupancy of the end-caps before and during such a beam splash. The average value is $\sim 6\%$, far higher than the pre-splash noise occupancy of 0.1%.

The right-hand plot of figure 5.14 shows the average 01X timing ratio as a function of end-cap disk, for a beam splash originating from side C ($-z$). Hits from disks on side C are almost coincident with the incoming particles, with an 01X fraction close to one. The ratio falls for modules closer to $z = 0$, as the timing delays are corrected for time-of-flight, but in the opposite direction to the beam splash particles. Modules on side A show a lower timing ratio of 0.1, as the whole end-cap is read out before the majority of particles arrive. The distribution is flat, as the time-of-flight corrections are now in

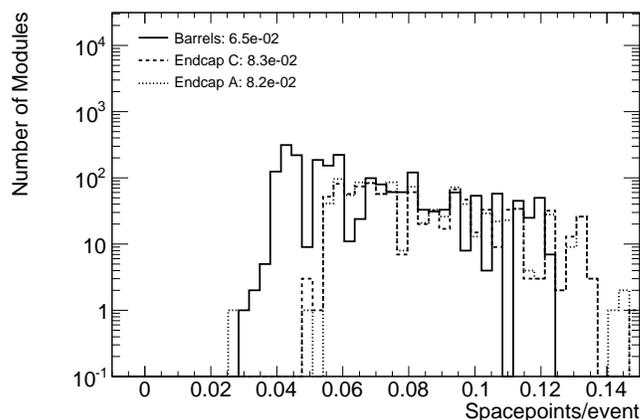


Figure 5.15: Spacepoint rates in the SCT for 7 TeV collisions (Run 154822, Scan 18).

the same direction as the beam splash. This led to the conclusion that the SCT had been trigger early with respect to the incoming splashes. This suspicion was confirmed by other sub-detectors, resulting in a change of the overall ATLAS trigger delay.

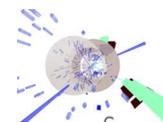
5.6.4 First Collision Measurements

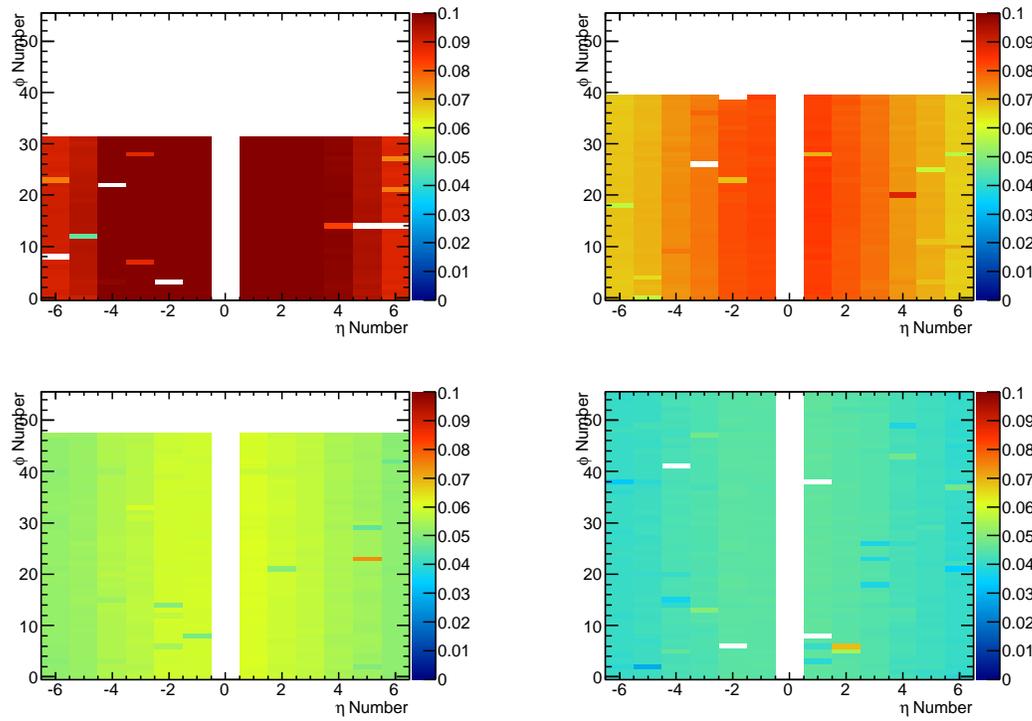
The ROD monitoring reached its full potential when it was active during the first high-energy LHC collisions in early 2010. Figures 5.15 and 5.16 are examples of the monitoring plots produced during 7 TeV pp collisions.

The spacepoint rate shown in figure 5.15 is now on average 0.07 for barrel modules, and 0.08 for the end-caps, corresponding to a spacepoint every 150 and 120 triggers in a single module respectively. Assuming that each track leaves four hits in the SCT, this equates to a multiplicity of ~ 80 charged particles per event. The spacepoint map for the barrel layers shows a higher intensity closer to the interaction point, as expected. The outer layers also show a sequential reduction in rate, as each module covers a smaller solid angle with increasing radius. The end-cap maps show a good symmetry between the two sides of the detector, as expected with well-centred beams.

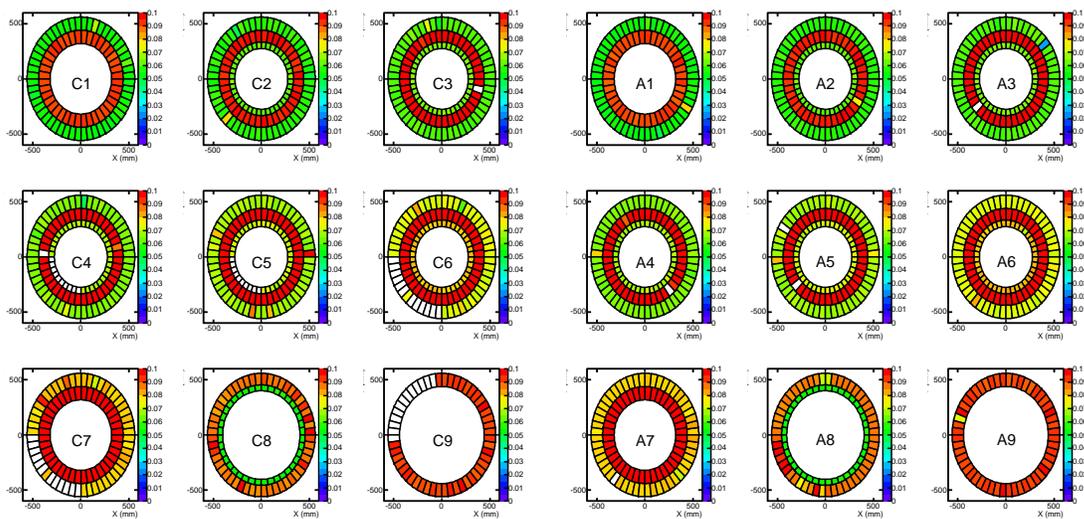
5.6.5 SCT Beam Monitoring

The spacepoint measurement can be used as a rudimentary beam monitoring system. The particle flux per event can be calculated by dividing the spacepoint rate by the





(a) Spacepoint rate map for barrel 3 (top-left), 4 (top-right), 5 (bottom-left) and 6 (bottom-right).



(b) Spacepoint rate map for the end-cap modules.

Figure 5.16: Spacepoint rate maps in the SCT for 7 TeV collisions (Run 154822, Scan 18). The z -axis units are the number of spacepoints per triggered event. White modules are excluded from the readout due to the reasons described in section 3.6.4.

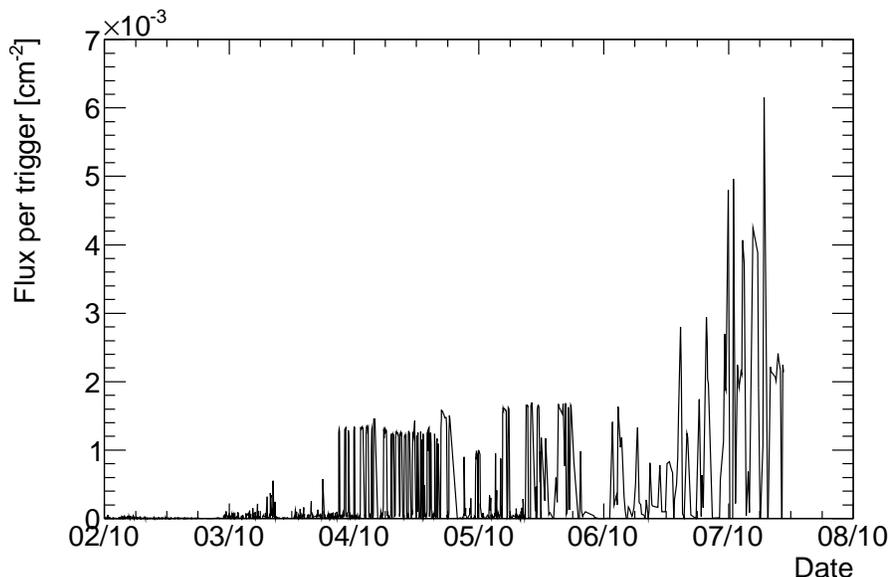


Figure 5.17: Derived particle flux against time for module 20220330200427, a barrel module from the innermost layer 3, η number 1.

module area. Figure 5.17 shows how the flux has changed over time for a single module in the innermost barrel (layer 3) and η index of 1. The measured flux can be directly correlated to beam activity. During the first months of 2010 only cosmic and noise runs were undertaken, which show a very low recorded flux. During March, there are some small peaks from single beam tests. The first collisions at $\sqrt{s} = 7$ TeV can be seen at the start of April, after which the flux remains at $\sim 1.5 \times 10^{-3} \text{cm}^{-2}$ until mid-May. From mid-May onwards, the beams were squeezed to low β^* , increasing the instantaneous luminosity. During June, the first high intensity bunches were collided with 10^{11} protons, which are clearly seen as the higher intensity peaks, the highest of which reaches over $6 \times 10^{-3} \text{cm}^{-2}$.

5.6.6 Timing Ratio Measurements

The ROD monitoring has also been used as a cross check of the SCT timing. As with other measurements, cosmic runs were used to test and validate the monitoring before collisions.

The goodness of timing can be quantified by the 01X timing ratio defined in equation 5.9. The expected ratio from purely random triggers can be estimated using a similar method as used to derive equation 5.16. The probability of a noise hit producing



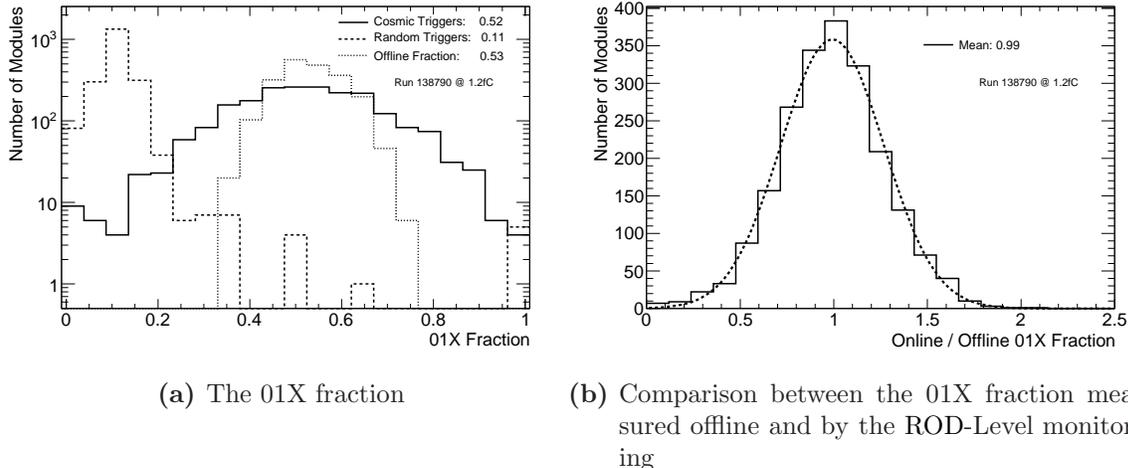


Figure 5.18: The fraction of hits with 01X time structure for noise and cosmic triggers.

a coincidence with a 01X noise pattern is:

$$P(N_{01X} > 0)_{\text{module}} \approx P(N_{010} > 0)_{\text{module}} = 16p^2s^2, \quad (5.25)$$

and so the expected ratio from random noise is $f_{01X} = \frac{1}{9}$.

This is confirmed by the ROD monitoring as shown in the figure 5.18a, where the timing ratio for randomly triggered events peaks at the expected value of 0.11. The same plot also shows the 01X fraction for cosmic triggers as measured by the ROD monitoring, and by the offline measurement. The offline measurement is expected to be more precise, as the timing ratio is calculated only for those hits associated to a reconstructed track, hence reducing the contribution from noise. Indeed, a much narrower distribution is observed for the offline measurement compared to the ROD monitoring. Figure 5.18b shows the ratio of f_{01X} as measured by the ROD monitoring and offline. The two measurements are in good agreement, with the mean of a Gaussian fit being 0.992 ± 0.006 .

Synchronisation of SCT Timing

The timing of the SCT can be changed using a combination of different hardware registers which control the delay between a trigger being received by the TIM or ROD, and its arrival on the modules. The coarse delay was set during cosmic runs, by adjusting the overall trigger delay in units of 25 ns, until the track rate reached a maximum. Module-

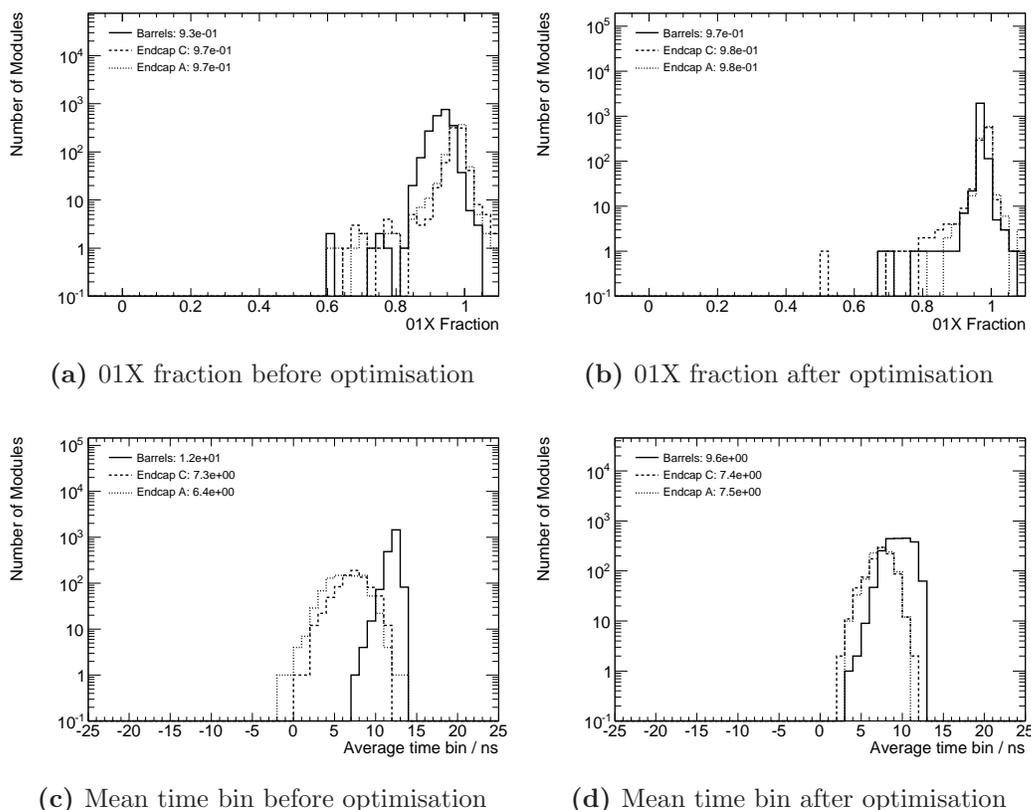
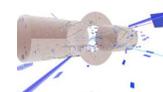


Figure 5.19: Timing variables for the SCT during collisions with Minimum Bias triggers, before (left) and after (right) an initial timing optimisation.

by-module variations in this delay were initially calculated using estimates derived from fibre lengths. This initial estimate proved to be remarkably accurate, as shown in the left-hand plots of figure 5.19, which show the timing fraction and the timing offset during collisions. The 01X fraction was already 93% in the barrels and 97% in the end-caps, putting the timing within 12 ns and 7 ns of the trigger respectively.

In order to improve the timing of the SCT, a timing scan was performed, altering the module delay in steps of 5ns, between -20 ns and 20ns from the initial value. The bulk of the analysis was performed offline, choosing the delay which gives the highest 01X fraction for each module. Although the ROD monitoring was not involved in the final analysis, it was used during the scan to verify that the timing had been changed. After the analysis, the optimal offsets were applied to each module. The timing shift can also be observed with the ROD monitoring in the right-hand plots of figure 5.19. The barrels now show a 97% 01X fraction, with the endcaps at 98%, with smaller tails. As the 01X fraction was used to find the optimum offset, the expected range of timing offsets is 0ns



(010) to 12.5ns (011). This is confirmed by the plot, which shows no modules outside of this range.

5.7 Conclusions

This chapter has described some of the developments made to the SCT DAQ system in the run up to first collisions at the LHC. A number of smaller developments, such as simultaneous loading of slave DSP firmware and masking of noisy links during calibration scans, helped to improve the overall speed and reliability of the system. Masking of noisy links increased the stability of SCT data-taking during early combined ATLAS physics runs, before the detector was fully calibrated. The addition of an FPGA simulator facilitated injection of data at the earliest possible stage in the DAQ chain, allowing complete testing of the system even when the detector itself was unavailable. The simulator was also used to generate events of fixed occupancy in order to determine a safe buffer size during physics running. A ROD-based monitoring framework was developed using the existing histogramming routines used during calibration mode. The monitoring was tested using random triggers to make high statistics noise occupancy measurements, in which even the short end-cap modules were visible. A coincident chip calculation allowed spacepoint measurements to be made using both cosmic and minimum bias trigger types. Finally, the monitoring was used to confirm changes during timing optimisation scans.

Chapter 6

ZZ Diboson Selection in ATLAS

“Am I just making up numbers?”

— The Pigeon Detectives

6.1 Introduction

The production of Z pairs in proton-proton collisions at the LHC has been introduced in section 1.4. Pairs of Z bosons cannot be directly observed in the ATLAS detector; instead their presence must be inferred by searching for combinations of their decay products.

In this study, two complimentary channels will be considered. In the first, both Z bosons decay into pairs of oppositely-charged, same-flavour leptons, where a lepton can either be an electron or a muon. This channel has the advantage of being experimentally very clean, but has a relatively low branching ratio. At 7 TeV, the product of production cross section from equation 1.36, with the branching ratio in figure 1.10, gives $\sigma\mathcal{B}(ZZ \rightarrow lll) \sim 27$ fb.

The second channel contains one $Z \rightarrow ll$ ($l = e, \mu$) decay, with the second Z decaying to neutrinos, which produce significant missing transverse energy (\cancel{E}_T). All three neutrino generations contribute, giving a cross section of $\sigma\mathcal{B}(ZZ \rightarrow ll\nu\bar{\nu}) \sim 160$ fb at 7 TeV, ~ 6 times higher than the $ZZ \rightarrow lll$ channel.

This chapter proposes two sets of cuts designed to select ZZ events, in the $ZZ \rightarrow ll\bar{l}l$ and $ZZ \rightarrow ll\nu\bar{\nu}$ channels, using simulated ATLAS events at both 7 and 10 TeV centre-of-mass energy, in order to separate the signals from background. The cuts are motivated by the need to separate the signal from background channels, which in some cases have cross sections up to 10^5 times larger than the signal.

The expected yields obtained at the end of this chapter will be used in chapter 7 to estimate the sensitivity of ATLAS to the anomalous triple gauge couplings introduced in section 1.3.1.

6.2 Technical Overview

The analysis of large data sets presents a considerable logistical challenge, and is performed in two stages, described below.

6.2.1 Distributed Analysis

In the distributed analysis computing model, simulated datasets are generated and subsequently stored at a number of different computing sites around the world. Analysis jobs are then sent to the site where the relevant dataset is stored and the results are returned, without having to copy large datasets locally. A job can also be split into a number of subjobs, allowing parallel data processing and reducing the overall job time.

For this analysis, a custom set of tools and algorithms were developed in C++, based on the ATLAS analysis framework, Athena [84]. These tools extract information relevant to the analysis from the grid datasets to produce an output n -tuple. Using the Ganga [85] job management frontend, jobs were sent to run these algorithms on signal and background samples situated at computing sites around the world, via both the LCG and the Production and Distributed Analysis Framework (PanDA) [86].

6.2.2 ROOT Analysis

The smaller, more manageable n -tuple files are then copied locally and analysed using a standalone custom C++ package [87], based on the ROOT [88] framework. This final

stage in the analysis chain makes cuts on particular variables, produces histograms and calculates expected event yields.

6.3 Simulated Data Samples

The simulated data sets used in this study have been produced centrally by the ATLAS collaboration, using the Athena software package to run a chain of computation steps. This involves production of events for a given process, using generators described in section 1.5, and simulating the response of the ATLAS detector, as explained in section 6.3.1.

Datasets in this study fall into two groups, those generated at 10 TeV centre-of-mass energy, and those at 7 TeV, described in detail in sections 6.3.2 and 6.3.3 respectively.

6.3.1 Simulating the ATLAS Detector

Simulated events are first generated using one of the Monte Carlo programs described in section 1.5. For a given hard process, a generator will produce a set of events containing lists of final-state particles and their four-momenta with respect to the origin.

In the simulation stage, the generated four-vectors are fed into a GEANT [89] model of the ATLAS detector, where interactions between particles and the detector material are modelled. This includes simulating charge deposits in the tracking detectors and showering of particles in the calorimeter material. Interactions between particles and inactive material such as support structures and cabling are also modelled. For example, photons frequently convert into electron-positron pairs on passage through the detector material.

In the next stage, charge deposits in active detector regions are digitised [90] to mimic the experiment's read-out systems. For example, the SCT models charge drift in the silicon sensors and the response of detector electronics. Random noise hits are also added using results from the noise occupancy test. Events from the digitisation stage are output with the same format as events from ATLAS collisions.

The reconstruction stage can be run both on simulated events and real data. Patterns of hits in the Inner Detector are used to recreate charged particle tracks and measure their momenta, and energy deposits in the calorimeters are grouped together into clus-



ters. Reconstructed electrons, muons, jets and missing energy are also formed using combinations of tracking and calorimetry information. In the case of simulated events, a record of the original generated particle kinematics is also retained, referred to as the “Truth”.

6.3.2 10 TeV Datasets

At the start of this study, it was foreseen that the LHC would begin to collide protons at a centre of mass energy of 10 TeV, and as a result a set of simulated data samples were produced by the ATLAS collaboration at this energy. These samples are part of the mc08 production, and were reconstructed using Athena version 15.3.1.6. This particular round of event generation used the CTEQ6LL [12] PDFs as an input to the matrix element calculation for LO generators, and CTEQ6M [91] at NLO. A total of 12 million simulated events were available at 10 TeV for use in this analysis.

A summary of the simulated 10 TeV signal and background processes relevant to this analysis is shown in table 6.1. It also describes any filters applied to the dataset during event generation and the corresponding efficiency. The preselection efficiency in this table is described further in section 6.4.5.

Re-weighting 10 TeV Monte Carlo

It became apparent in early 2010 that 10 TeV would be an unrealistic prospect for early LHC running, with 7 TeV decided upon as the highest safe energy. In addition to using the dedicated 7 TeV data sets described in section 6.3.3, the 10 TeV data sets can be effectively scaled down in energy using a re-weighting technique. As the two samples are statistically independent, they can be combined to decrease uncertainties. In some cases 10 TeV re-weighted backgrounds will be used where the 7 TeV equivalent sample was unavailable.

Each 10 TeV data set is rescaled by assigning an event weight which depends on the underlying PDFs, which were introduced in section 1.5.2. The re-weighting is performed with the Athena `PDFReweight` tool [92], applying a weight w to each event

$$w = \frac{\text{PDF}(x'_1, Q, f_1) \times \text{PDF}(x'_2, Q, f_2)}{\text{PDF}(x_1, Q, f_1) \times \text{PDF}(x_2, Q, f_2)}, \quad (6.1)$$

Name	Generator	Dataset	σ / fb @ 10 TeV	$\langle w \rangle$	σ / fb @ 7 TeV	Generator filter	Filter ϵ	k-Factor	N Events	Preselection ϵ	N Preselected
$ZZ^* \rightarrow ll\bar{l}l$ ($l = e, \mu, \tau$)	Pythia	109291	103.3	0.649	67.0852		1	1.53 (1.55)	187723	0.602	113053
$ZZ \rightarrow ll\nu\bar{\nu}$ ($l = e, \mu$)	MC@NLO	105932	247.3	0.618	152.772		1	1	19872	0.588	11692
$Z \rightarrow ee$	Pythia	106050	1.14396×10^6	0.710	812753	$m_{ll} > 60$ GeV $n_l > 0$	0.96	1.48 (1.50)	5187647	0.405	2099637
$Z \rightarrow \mu\mu$	Pythia	106051	1.14396×10^6	0.710	811802	$m_{ll} > 60$ GeV $n_l > 0$	0.96	1.48 (1.50)	4907023	0.490	2406743
$Z \rightarrow ee$ (0 Jets)	Alpgen	107650	898180	0.721	647262		1	1	269280	0.367	98835
$Z \rightarrow ee$ (1 Jet)	Alpgen	107651	206570	0.636	131428		1	1	61767	0.427	26369
$Z \rightarrow ee$ (2 Jets)	Alpgen	107652	72500	0.573	41557.7		1	1	216945	0.440	95375
$Z \rightarrow ee$ (3 Jets)	Alpgen	107653	21080	0.515	10847.5		1	1	63412	0.453	28696
$Z \rightarrow ee$ (4 Jets)	Alpgen	107654	6000	0.454	2726.78		1	1	18470	0.448	8269
$Z \rightarrow \mu\mu$ (0 Jets)	Alpgen	107660	900210	0.720	648286		1	1	270098	0.448	121000
$Z \rightarrow \mu\mu$ (1 Jet)	Alpgen	107661	205210	0.635	130214		1	1	61936	0.502	31095
$Z \rightarrow \mu\mu$ (2 Jets)	Alpgen	107662	69350	0.571	39617.8		1	1	207173	0.513	106305
$Z \rightarrow \mu\mu$ (3 Jets)	Alpgen	107663	21630	0.513	11090.4		1	1	64956	0.522	33884
$Z \rightarrow \mu\mu$ (4 Jets)	Alpgen	107664	6080	0.456	2771.69		1	1	18470	0.513	9484
$Z \rightarrow \tau\tau$	Pythia	105062	1.14396×10^6	0.708	810359	$m_{ll} > 60$ GeV	0.96	1.48 (1.50)	978438	0.016	15664
$t\bar{t}$	MC@NLO	105200	373600	0.395	147459	No all-hadronic decays	0.543	1	2225238	0.047	103816
$b\bar{b} Z \rightarrow ee$ (0 Jets)	Alpgen	109300	12220	0.537	6564.83	$m_{ll} > 30$ GeV	1	1	299757	0.539	161601
$b\bar{b} Z \rightarrow ee$ (1 Jet)	Alpgen	109301	4947	0.497	2456.3	$m_{ll} > 30$ GeV	1	1	147838	0.520	76831
$b\bar{b} Z \rightarrow ee$ (2 Jets)	Alpgen	109302	1960.2	0.462	905.862	$m_{ll} > 30$ GeV	1	1	39985	0.501	20013
$b\bar{b} Z \rightarrow ee$ (3 Jets)	Alpgen	109303	950	0.408	388.002	$m_{ll} > 30$ GeV	1	1	10000	0.485	4849
$b\bar{b} Z \rightarrow ee$ (0 Jets)	Alpgen	109400	12220	0.488	5964.02	$m_{ll} > 30$ GeV, $N_l = 4, p_T(e, \mu) > 5$ GeV, $p_T(\tau) > 10$ GeV	0.0044	1	99801	0.654	65316
$b\bar{b} Z \rightarrow ee$ (1 Jet)	Alpgen	109401	4947	0.448	2215.07	$m_{ll} > 30$ GeV, $N_l = 4, p_T(e, \mu) > 5$ GeV, $p_T(\tau) > 10$ GeV	0.0104	1	49892	0.608	30287
$Z \rightarrow \mu\mu b\bar{b}$ (0 Jets)	Alpgen	109305	12280	0.536	6580.01	$m_{ll} > 30$ GeV	1	1	299714	0.629	188571
$Z \rightarrow \mu\mu b\bar{b}$ (1 Jet)	Alpgen	109306	4924	0.495	2438.2	$m_{ll} > 30$ GeV	1	1	144742	0.597	86399
$Z \rightarrow \mu\mu b\bar{b}$ (2 Jets)	Alpgen	109307	1917.2	0.460	882.365	$m_{ll} > 30$ GeV	1	1	39952	0.569	22722
$Z \rightarrow \mu\mu b\bar{b}$ (3 Jets)	Alpgen	109308	936.9	0.411	385.461	$m_{ll} > 30$ GeV	1	1	10000	0.543	5432
$Z \rightarrow \mu\mu b\bar{b}$ (0 Jets)	Alpgen	109405	12280	0.485	5952.83	$m_{ll} > 30$ GeV, $N_l = 4, p_T(e, \mu) > 5$ GeV, $p_T(\tau) > 10$ GeV	0.0049	1	99671	0.709	70667
$Z \rightarrow \mu\mu b\bar{b}$ (1 Jet)	Alpgen	109406	4924	0.445	2190.54	$m_{ll} > 30$ GeV, $N_l = 4, p_T(e, \mu) > 5$ GeV, $p_T(\tau) > 10$ GeV	0.0123	1	49855	0.651	32466
$ZZ \rightarrow eebb$	Pythia	109288	81.55	0.649	52.9589		1	1	89717	0.487	43682
$ZZ \rightarrow \mu\mu b\bar{b}$	Pythia	109289	81.55	0.646	52.6407		1	1	99919	0.559	55814
$ZZ \rightarrow \tau\tau b\bar{b}$	Pythia	109290	81.36	0.642	52.1961	$p_T(e, \mu) > 5$ GeV, $p_T(\tau) > 10$ GeV, $ \eta < 6$	0.592	1	99861	0.039	3900
$W^+ Z \rightarrow l^+ \nu ll$	MC@NLO	105941	264.73	0.612	162.123		1	1	24697	0.670	17037
$W^- Z \rightarrow l^- \bar{\nu} ll$	MC@NLO	105971	155.96	0.560	87.3162		1	1	19725	0.707	13943
$W^+ W^- \rightarrow e^+ \nu e^- \bar{\nu}$	MC@NLO	105921	828.38	0.639	529.729		1	1	24833	0.507	12593
$W^+ W^- \rightarrow e^+ \nu \mu^- \bar{\nu}$	MC@NLO	105922	828.38	0.602	498.889		1	1	23642	0.007	158
$W^+ W^- \rightarrow e^+ \nu \tau^- \bar{\nu}$	MC@NLO	105923	828.38	0.627	519.735		1	1	24616	0.072	1781
$W^+ W^- \rightarrow \mu^+ \nu \mu^- \bar{\nu}$	MC@NLO	105924	828.38	0.636	527.136		1	1	24796	0.601	14906
$W^+ W^- \rightarrow \mu^+ \nu e^- \bar{\nu}$	MC@NLO	105925	828.38	0.627	519.324		1	1	23905	0.006	151
$W^+ W^- \rightarrow \mu^+ \nu \tau^- \bar{\nu}$	MC@NLO	105926	828.38	0.633	524.111		1	1	24515	0.0805	1974
$W^+ W^- \rightarrow \tau^+ \nu \tau^- \bar{\nu}$	MC@NLO	105927	828.38	0.639	529.563		1	1	24467	0.0218	534
$W^+ W^- \rightarrow \tau^+ \nu e^- \bar{\nu}$	MC@NLO	105928	828.38	0.636	527.014		1	1	22958	0.073	1682
$W^+ W^- \rightarrow \tau^+ \nu \mu^- \bar{\nu}$	MC@NLO	105929	828.38	0.632	523.228		1	1	24052	0.080	1932

Table 6.1: Summary of signal and background simulated data samples at $\sqrt{s} = 10$ TeV. The reweighted cross section at $\sqrt{s} = 7$ TeV is also shown.



where f_1 and f_2 are the flavours of the two partons involved in the hard process of the original event, with corresponding longitudinal momenta x_1 and x_2 . The rescaled parton momenta are given by

$$x'_1 = x_1 \times (E/E'), \quad (6.2)$$

$$x'_2 = x_2 \times (E/E'), \quad (6.3)$$

where E and E' are the original and new beam energies respectively. Here we take $E = 5$ TeV and $E' = 3.5$ TeV. The weight w is then multiplied by any internal generator weight (for example MC@NLO) to give an overall event weight.

Re-weighting Validation

This re-weighting procedure has been used for a number of analyses in ATLAS, originally to scale simulated events from $14 \rightarrow 10$ TeV, and more recently from $10 \rightarrow 7$ TeV.

To validate this procedure, two LO PYTHIA samples were generated for the $ZZ \rightarrow ll\bar{l}l$ ($l = e, \mu$) process at 7 TeV and 10 TeV, both using the CTEQ6L PDF, with the 10 TeV sample also re-weighted to 7 TeV for comparison.

As shown in the upper plots of figure 6.1, the re-weighted $Z p_T$ spectrum and invariant dilepton mass show some deviation from the 7 TeV kinematics at the generator level. As demonstrated by the lower right-hand plot, the event weight shows a strong dependence on the incoming quark flavour.

The lower left-hand plot shows the ratio of re-weighted cross sections as a function of the $Z p_T$. From this plot it appears that the re-weighting method underestimates the cross section at low $p_T < 50$ GeV by up to $\sim 2\%$, while at higher p_T , the cross section is overestimated by the same magnitude.

The technique described in [93] was used to compare the 7 TeV and re-weighted $Z p_T$ histograms in the range $p_T < 300$ GeV. The comparison yielded a $\chi^2 = 54.3$ for 49 degrees of freedom, with a corresponding p -value of 0.28.

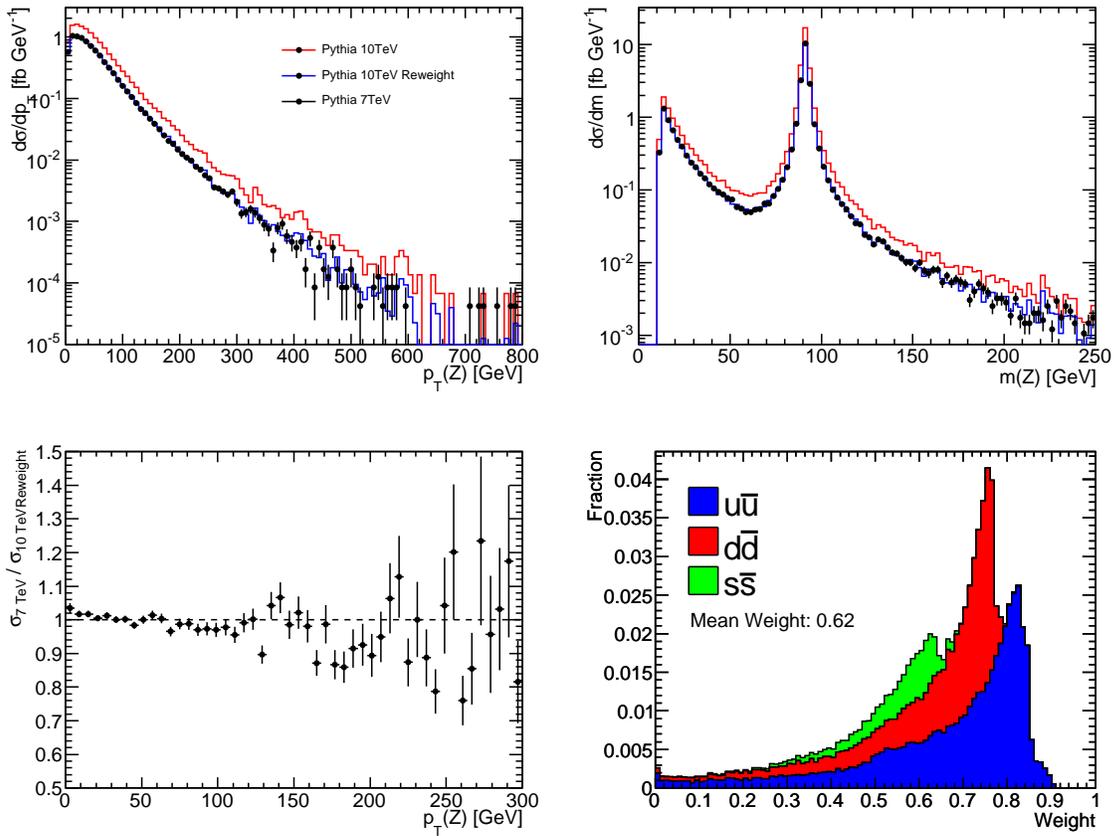


Figure 6.1: Example of re-weighting simulated samples from 10 to 7 TeV, using $ZZ \rightarrow lll$ events generated with PYTHIA. The top-left and top-right plots show the $p_T(Z)$ and m_Z differential cross sections respectively. The lower left plots shows the ratio of 7 TeV to re-weighted spectra, with the dashed line at 1 for guidance. A stacked histogram of event weights is shown in the lower right-hand plot, separated by the flavour of the incoming quarks.

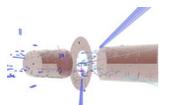
The re-weighting method gives a good agreement when considering the overall cross section. The cross sections were calculated as follows:

$$\sigma(10 \text{ TeV}) = 53.4 \pm 0.2 \text{ (stat) fb}, \quad (6.4)$$

$$\sigma(7 \text{ TeV}) = 33.2 \pm 0.1 \text{ (stat) fb}, \quad (6.5)$$

$$\sigma(7 \text{ TeV (re-weighted)}) = 33.2 \pm 0.3 \text{ (stat) fb}, \quad (6.6)$$

which is consistent with the mean event weight of 0.62 from figure 6.1. The re-weighted cross section agrees with the 7 TeV calculation with statistical errors. The mean weight and re-weighted cross section for each 10 TeV data set are listed in table 6.1.



6.3.3 7 TeV Datasets

In addition to simulated events at 10 TeV, dedicated 7 TeV simulated datasets, produced with Athena version 15.6.6.5, were available at the time of writing, including signal samples and many of the backgrounds. These datasets were generated as part of the ATLAS mc09 production, and unlike the 10 TeV samples, were produced using the LO* PDFs from MRST2007 [14] in the case of PYTHIA, CTEQ66 [13] for MC@NLO and CTEQ6LL for Alpgen. In this case, a total of 50 million simulated events were available at 7 TeV for use in this study. A summary of the relevant datasets and their cross sections is shown in table 6.2.

6.3.4 Cross Section Calculations

This section discusses the choice of cross sections used for the signal and background channels as shown in tables 6.2 and 6.1.

***ZZ* → *llll* Cross Section**

For the four lepton channel, a large simulated PYTHIA sample is used, containing decays to all three lepton families ($Z \rightarrow ee, \mu\mu, \tau\tau$). The PYTHIA generator has the advantage that, unlike MC@NLO, it contains off-shell Z bosons (Z^*) and Z/γ interference, which can be used to increase the acceptance, as described in section 6.5.2.

The disadvantage of using PYTHIA is that it only generates events at leading order, and so an overall correction is required to scale the cross section to NLO, known as the k -factor.

Each row of table 6.3 shows cross section calculations at LO and NLO, using the MCFM generator with the indicated PDFs. Z/γ interference has been included in the cross section. The LO cross section using the PYTHIA generator is also included for comparison, producing a cross section which is lower than MCFM by 3–4%. The k -factor is defined as

$$k = \frac{\sigma_{\text{NLO}}}{\sigma_{\text{LO}}}, \quad (6.7)$$

Name	Generator	Dataset	$\sigma / \text{fb} @ 7 \text{ TeV}$	Generator filter	Filter ϵ	k-Factor	N Events	Preselection ϵ	N Preselected
$ZZ^* \rightarrow ll\bar{l}l$ ($l = e, \mu, \tau$)	Pythia	106050	76.2955		1	1.25	99976	–	99976
$ZZ \rightarrow ll\nu\bar{\nu}$ ($l = e, \mu$)	MC@NLO	105932	151.582		1	1	99840	–	99840
$Z \rightarrow ee$	Pythia	106050	850525	$m_{ll} > 60 \text{ GeV } n_l > 0$	0.97	1.17	299870	0.424	127106
$Z \rightarrow ee$	Pythia	106046	851019	$m_{ll} > 60 \text{ GeV}$	1	1.17	4758621	0.413	1963918
$Z \rightarrow ee$	MC@NLO	106087	952394		1	1	4959337	0.411374	2040141
$Z \rightarrow \mu\mu$	Pythia	106051	850603	$m_{ll} > 60 \text{ GeV } n_l > 0$	0.96	1.17	298952	0.537	160390
$Z \rightarrow \mu\mu$	Pythia	106047	851011	$m_{ll} > 60 \text{ GeV}$	1	1.17	4998410	0.521	2606449
$Z \rightarrow \mu\mu$	MC@NLO	106088	952526		1	1	4957942	0.519	2575600
$Z \rightarrow ee$ (0 Jets)	Alpgen	107650	659583		1	1	304216	0.340	112487
$Z \rightarrow ee$ (1 Jet)	Alpgen	107651	132462		1	1	63440	0.431	27350
$Z \rightarrow ee$ (2 Jets)	Alpgen	107652	41355.5		1	1	19497	0.458	8934
$Z \rightarrow ee$ (3 Jets)	Alpgen	107653	10789.9		1	1	10998	0.472	5196
$Z \rightarrow ee$ (4 Jets)	Alpgen	107654	3128.18		1	1	1499	0.446	669
$Z \rightarrow ee$ (5 Jets)	Alpgen	107655	753.425		1	1	500	0.478	239
$Z \rightarrow \mu\mu$ (0 Jets)	Alpgen	107660	652731		1	1	303947	0.479	145471
$Z \rightarrow \mu\mu$ (1 Jet)	Alpgen	107661	133855		1	1	62996	0.536	33755
$Z \rightarrow \mu\mu$ (2 Jets)	Alpgen	107662	40756.8		1	1	18993	0.544	10329
$Z \rightarrow \mu\mu$ (4 Jets)	Alpgen	107664	2832.34		1	1	1499	0.550	825
$Z \rightarrow \mu\mu$ (5 Jets)	Alpgen	107665	756.621		1	1	998	0.523	522
$W \rightarrow e\nu$	Pythia	106043	8.89409×10^6		1	1.15	7694356	0.00322	24739
$W \rightarrow \mu\nu$	Pythia	106044	8.871×10^6		1	1.15	6993798	0.000619	4326
$W^+ \rightarrow e^+\nu$	MC@NLO	106081	5.86829×10^6		1	1	3955495	0.00294	11637
$W^- \rightarrow e^-\nu$	MC@NLO	106082	3.9973×10^6		1	1	2960648	0.00253	7479
$W^+ \rightarrow \mu^+\nu$	MC@NLO	106083	5.86829×10^6		1	1	4016691	0.000580	2331
$W^- \rightarrow \mu^-\nu$	MC@NLO	106084	3.9973×10^6		1	1	2929863	0.000537	1572
$t\bar{t}$	MC@NLO	105200	144069	No all-hadronic decays	0.562	1	201634	0.0475	9585
W^+Z	MC@NLO	105941	161.155		1	1	24556	0.718	17628
W^-Z	MC@NLO	105971	86.3766		1	1	128878	0.723	93162
$W^+W^- \rightarrow e^+\nu e^-\bar{\nu}$	MC@NLO	105921	507.852		1	1	49549	0.524	25944
$W^+W^- \rightarrow e^+\nu\mu^-\bar{\nu}$	MC@NLO	105922	507.852		1	1	45348	0.00831	377
$W^+W^- \rightarrow e^+\nu\tau^-\bar{\nu}$	MC@NLO	105923	507.852		1	1	58925	0.0751	4423
$W^+W^- \rightarrow \mu^+\nu\mu^-\bar{\nu}$	MC@NLO	105924	507.852		1	1	49725	0.647	32152
$W^+W^- \rightarrow \mu^+\nu e^-\bar{\nu}$	MC@NLO	105925	507.852		1	1	46197	0.00723	334
$W^+W^- \rightarrow \mu^+\nu\tau^-\bar{\nu}$	MC@NLO	105926	507.852		1	1	49314	0.0869	4283
$W^+W^- \rightarrow \tau^+\nu\tau^-\bar{\nu}$	MC@NLO	105927	507.852		1	1	48434	0.0209	1010
$W^+W^- \rightarrow \tau^+\nu e^-\bar{\nu}$	MC@NLO	105928	507.852		1	1	50091	0.0722	3616
$W^+W^- \rightarrow \tau^+\nu\mu^-\bar{\nu}$	MC@NLO	105929	507.852		1	1	50731	0.0832	4219

Table 6.2: Summary of signal and background simulated data samples at $\sqrt{s} = 7 \text{ TeV}$.

\sqrt{s} / TeV	Generator	Cuts	LO PDF	NLO PDF	σ_{LO} / fb	σ_{NLO} / fb	k -factor
7	PYTHIA	✗	MRST LO*	-	76.3	-	-
	MCFM	✗	CTEQ6LL	CTEQ66	67.4	91.9	1.36
	MCFM	✓	CTEQ6LL	CTEQ66	24.1	36.7	1.53
	MCFM	✗	MRST LO*	CTEQ66	78.9	91.9	1.16
	MCFM	✓	MRST LO*	CTEQ66	29.5	36.7	1.25
10	PYTHIA	✗	CTEQ6LL	-	103.3	-	-
	MCFM	✗	CTEQ6LL	CTEQ66	109.3	150.7	1.38
	MCFM	✓	CTEQ6LL	CTEQ66	33.8	52.4	1.55

Table 6.3: Cross sections for $ZZ \rightarrow lll$ ($l = e, \mu, \tau$) production at Leading and Next-to-Leading Order.

in order to scale the cross section to a common order and PDF. A cross in the cuts column indicates that the only requirement made was that $m_{ll} > 12$ GeV is satisfied for both Z bosons, in order to match the PYTHIA cuts. A tick indicates that additional cuts have been applied, requiring the highest p_T lepton in the event to have $p_T > 20$ GeV, with a cut of $p_T > 10$ GeV on the remaining leptons. Each lepton is required to have $|\eta| < 2.5$. The mass of one Z is required to be $70 \text{ GeV} < m_Z < 110 \text{ GeV}$, with the second needing $m_Z > 20 \text{ GeV}$. The intention of the cuts is to derive a k -factor which matches the region of phase-space of selected events described in section 6.5.2.

The MRST LO* produce cross sections which are $\sim 17\%$ higher than CTEQ6LL, and consequently predict lower k -factors. This is because the LO* PDFs are constructed to produce events which give a closer approximation to NLO kinematics using LO generators, and hence have larger-than-LO cross sections. The predicted k -factors at 10 TeV are $\sim 3\%$ higher than the corresponding 7 TeV calculation. It also appears that NLO effects are more significant when cuts are applied. With CTEQ6LL as the LO PDF, the k -factor increases by 13% with cuts compared to without. This difference is 8% in the case where MRST LO* PDFs are used.

In conclusion, k -factors calculated with generator cuts applied will be used to scale the PYTHIA $ZZ \rightarrow lll$ cross sections to NLO. k -factors of 1.53, 1.55 and 1.25 will be used for the 10 TeV, re-weighted and 7 TeV signal samples respectively. The error on the cross sections from PDF uncertainties is typically 4%, which corresponds to a k -factor uncertainty of 6%.

$ZZ \rightarrow ll\nu\bar{\nu}$ Cross Section

In the case of the $ZZ \rightarrow ll\nu\bar{\nu}$ channel, tight cuts need to be made on the invariant mass of the observed lepton pair to reduce background; thus little is gained from using a Monte Carlo sample containing off-shell Z decays. A sample of $ZZ \rightarrow ll\nu\bar{\nu}$ ($l = e, \mu$) events was available from the MC@NLO program, which generates Z bosons with zero width and no Z/γ interference. As MC@NLO already generates events at next-to-leading order, no k -factor corrections are necessary.

\sqrt{s} / TeV	Generator	Order	PDF	σ / fb
7	MC@NLO	NLO	CTEQ66	151.6
7	MCFM	NLO	CTEQ66	157.9
10	MC@NLO	NLO	CTEQ6M	247.3
10	MCFM	NLO	CTEQ6M	259.6

Table 6.4: Cross sections for $ZZ \rightarrow ll\nu\bar{\nu}$ ($l = e, \mu$) for proton-proton collisions.

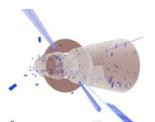
The cross-section calculations for $ZZ \rightarrow ll\nu\bar{\nu}$ production are shown in table 6.4. As previously, the MCFM result is slightly higher than MC@NLO by 4–5%. The MC@NLO values for the cross section will be used to estimate yields after cuts.

Background Cross Sections

The background channels relevant to this analysis are described in section 1.6.3. As some of the background samples have also been generated at LO, additional k -factors are required.

For the MC@NLO and Alpgen background channels, the cross section used in each case will be taken from the generator producing the events. In the case of the single W and Z PYTHIA backgrounds, k -factors are also required to scale the given cross section to NLO.

In the case of single $W \rightarrow l\nu$, the k -factor is the ratio of the cross section at NLO with CTEQ66, to that at LO with the MRST2007 LO* PDF, for consistency with the MC@NLO samples. Cuts are applied on the leptons in the event, requiring $p_T(l) >$



20 GeV and $|\eta| < 2.5$. This gives values at 7 TeV of

$$k(W^+) = 1.17 \text{ (1.13)} \quad (6.8)$$

$$k(W^-) = 1.11 \text{ (1.09)} \quad (6.9)$$

$$k(W) = 1.15 \text{ (1.12)}, \quad (6.10)$$

with the final value calculated as the cross-section weighted mean W^+ and W^- for use with the PYTHIA sample. The values in parentheses show the k -factors without cuts.

A similar strategy is used for the $Z \rightarrow ll$ channel. Cross sections were calculated with cuts of $p_T(l_1) > 20$ GeV and $p_T(l_2) > 10$ GeV for the leading and trailing leptons, $|\eta| < 2.5$ and $70 \text{ GeV} < m_Z < 110 \text{ GeV}$.

For the 10 TeV and re-weighted samples, the LO cross section is calculated with CTEQ6L, scaled to NLO with CTEQ66. In the 7 TeV case, MRST2007 LO* is scaled to CTEQ66 at NLO. This gives results of

$$k(Z @ 10 \text{ TeV}) = 1.48 \text{ (1.25)}, \quad (6.11)$$

$$k(Z @ 7 \text{ TeV Re-weight}) = 1.50 \text{ (1.35)}, \quad (6.12)$$

$$k(Z @ 7 \text{ TeV}) = 1.17 \text{ (1.10)}, \quad (6.13)$$

where the values in parentheses are the k -factors calculated without the above cuts. Once again, the 7 TeV k -factor is lower as the LO* calculations tend to produce results that are closer in value to NLO. These values will be used to scale the PYTHIA cross section to NLO. For the 7 TeV background samples, an MC@NLO single Z sample is available, which will be used in preference to the PYTHIA sample as it is true NLO. The PYTHIA sample remains useful for cross checks as it includes the Z width.

The k -factors calculated with cuts will be used as they give a better approximation to the region of phase-space selected by the kinematic cuts described in sections 6.5.2 and 6.6.2.

6.4 Common Preselection

As there are many common features between the selection of the $ZZ \rightarrow lll$ and $ZZ \rightarrow ll\nu\bar{\nu}$ channels, a single n -tuple is created containing variables common to both. At this

stage, custom Athena algorithms loop over reconstructed objects of different types to see if they pass the pre-selection criteria defined in the following sections.

Objects are only added to the n -tuple if they do not overlap with existing objects. Two objects are defined as overlapping if one lies within a cone closer than $\Delta R < 0.1$ to the other. This requirement means that the insertion order of physics objects is important. The order used for insertion matches the order in which objects are described here, namely electrons first, followed by muons and finally jets.

6.4.1 Electron Selection

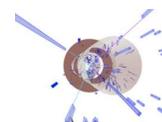
A typical electron signature in ATLAS comprises a charged track in the ID, matched to an electromagnetic shower in the calorimeter. A number of preselection cuts are required to match these two components and successfully reconstruct electrons.

The first requirement is that electrons must be reconstructed with the *egamma* algorithm (chapter 3 of [43]) which selects electron and photon candidates based on a number of different identification techniques including cuts, likelihood, H-matrix and a neural network. In this analysis electrons are selected using cuts based on the shower shape properties in the calorimeters as well as variables combining inner detector tracks with calorimeter deposits. A number of pre-defined electron selection cuts are provided with the aim of standardising ATLAS analyses, known as tight, medium and loose. In addition, the requirement can be imposed that the electron is isolated, using cuts on the energy deposited in a ΔR cone (equation 2.2) of fixed half-radius. This helps to reduce background where fake electrons are reconstructed from jets.

Two samples of electrons are defined, labelled as *physics* electrons and *veto* electrons. As will be discussed in section 6.6.2, only the *physics* electrons are used for Z formation, while the *veto* electrons are used to reject background in the $ZZ \rightarrow ll\nu\bar{\nu}$ channel.

Physics Electrons

The electrons used for Z formation are required to pass the standard criteria for medium electrons with isolation. Electrons must also have high transverse momentum ($p_T > 5$ GeV) and lie within the acceptance of the ID, with $|\eta| < 2.5$.



Physics electrons are also required to fulfil isolation requirements defined by the ATLAS *egamma* group [94]. The isolation ratio is defined as

$$I_R^{0.3} = \frac{E_T(\Delta R < 0.3)}{E_T}, \quad (6.14)$$

where the denominator is the total transverse energy of the electron, and the numerator is the transverse energy in a cone of $\Delta R < 0.3$ around the electron. In reference [94], cuts on the value of $I_R^{0.3}$ have been optimised in bins of p_T and η to maximise discrimination between electrons and jets.

Veto Electrons

Veto electrons are only required to pass the loose selection criteria, with no isolation requirement. The transverse momentum is again required to be $p_T > 5$ GeV. Note that the overlap requirement implies that veto electrons are only selected if they lie outside a cone of $\Delta R < 0.1$ around each physics electron.

Electron Selection Efficiency

The efficiency of electron pre-selection can be calculated by matching reconstructed electrons to truth electrons. An electron has a truth match if a true electron is found within a cone of $\Delta R < 0.01$ around it. For this calculation, only true electrons coming from the decay of the Z are considered, with an acceptance cut requiring $|\eta_e| < 2.5$. The efficiency is then defined as

$$\epsilon = \frac{N(\text{true matched to reconstructed})}{N(\text{true})}. \quad (6.15)$$

Similarly, the purity of the sample is defined as

$$p = \frac{N(\text{reconstructed matched to true})}{N(\text{reconstructed})}, \quad (6.16)$$

where the truth sample now consists of all electrons in the event, not just those from the true Z . This is done to avoid an artificially low purity arising from the fact that there are electrons and muons present from τ decays (the signal sample also contains $Z \rightarrow \tau\tau$).

Plots of the selection efficiency can be found in figure 6.2 as a function of true p_T , η and ϕ for physics and veto electrons. The purity of the electron sample is also shown as a function of p_T . The plots were made by selecting electrons from the $ZZ \rightarrow lll$ simulated sample at 7 TeV centre-of-mass energy.

Reconstruction Resolution

It is also possible to calculate the resolution of the electron p_T measurement by comparison with that of the matching true electron. The fractional difference between the reconstructed and true electrons is defined as

$$\frac{p_T^{\text{reco}} - p_T^{\text{true}}}{p_T^{\text{reco}}}, \quad (6.17)$$

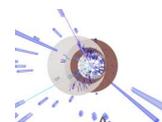
where p_T^{true} is the p_T of the true electron matched to the reconstructed electron. This distribution of this variable for electrons in the $ZZ \rightarrow lll$ sample is shown in the lower left plot of figure 6.2, with a Gaussian fit to physics electrons in the central region. The width of this Gaussian defines the reconstruction resolution, which is 2.8%, with an offset of -1.4%. The lower right-hand plot of figure 6.2 shows the offset binned as a function of p_T , with error bars representing the resolution. The effect of finite electron reconstruction resolution will be discussed in section 7.4.1, where it is included in the anomalous coupling limit calculations.

Preselection Summary

A summary of cuts used to select the two types of electron and the corresponding efficiency are shown in table 6.5. The overall selection efficiencies for electrons with $|\eta| < 2.5$ originating from a Z decay are also shown for the $ZZ \rightarrow lll$ signal, at $\sqrt{s} = 7$ TeV. The errors are statistical uncertainties from the number of Monte Carlo events in the sample.

6.4.2 Muon Selection

Muons have a very distinct signature in the ATLAS detector, as they traverse the calorimeters and pass through the dedicated muon tracking stations. A bewildering



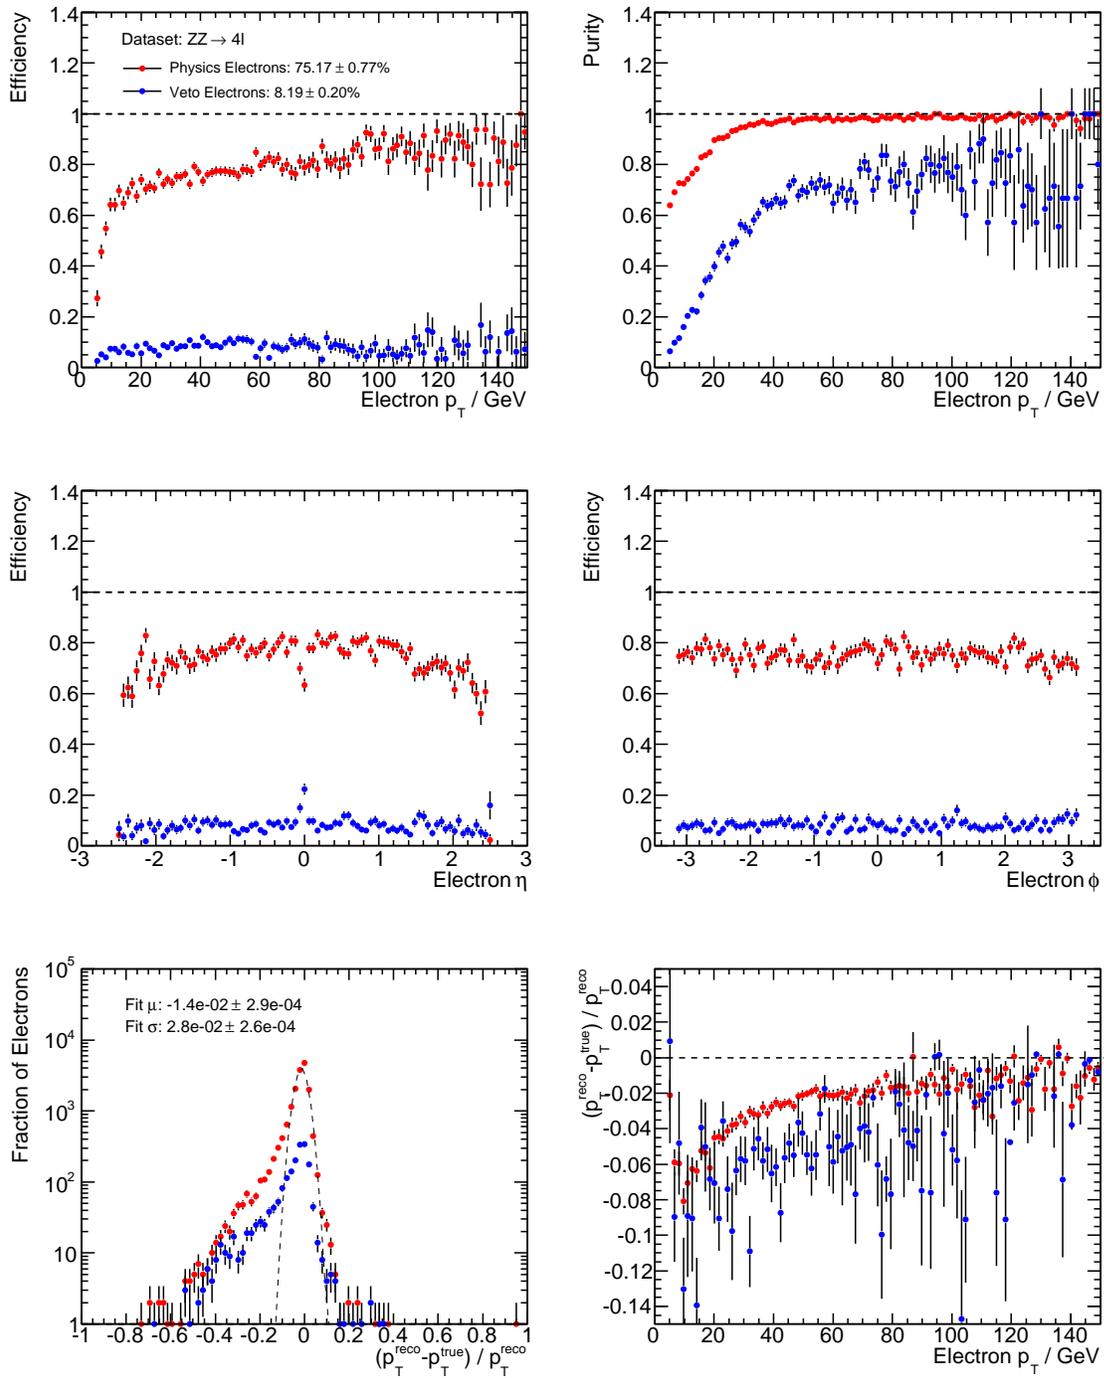


Figure 6.2: Efficiency, purity and resolution plots for electrons in the analysis.

Cut	Physics	Veto
Algorithm	<i>egamma</i>	<i>egamma</i>
Cuts	“MediumIso”	“Loose”
p_T	$> 5 \text{ GeV}$	$> 5 \text{ GeV}$
$ \eta $	< 2.5	< 2.5
Efficiency	$75.2 \pm 0.8\%$	$8.2 \pm 0.2\%$

Table 6.5: Summary of electron pre-selection cuts.

array of complimentary muon reconstruction algorithms are available to the ATLAS physicist. The main ingredients of a typical muon reconstruction algorithm involve combining track segments from the ID with tracklets constructed using hits in the muon chambers.

In this analysis, ID tracks are formed using the ATLAS New Tracking (NEWT) algorithm [95], while muon segments come from the Muonboy [96] algorithm¹. A χ^2 variable is computed reflecting the goodness-of-fit of the standalone muon tracklet. Here the requirement is made that $\chi_{\text{fit}}^2/n_{\text{dof}} < 15$ to ensure good quality tracks are chosen. A more detailed description of muon reconstruction is available in chapter 4 of [43]. For this analysis, two categories of muons, *combined* muons and *standalone* muons are defined as follows.

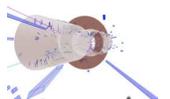
The two independent track measurements from the inner detector and muon systems described above are merged using the STATistical COmbination (STACO) routine [96], employing the following method: For two tracks at a reference location defined by their parameter vectors, P_1 and P_2 , and their covariance matrices, C_1 and C_2 , the parameter vector of the combined track, P , is the solution to the equation

$$(C_1^{-1} + C_2^{-1}) \times P = C_1^{-1} \times P_1 + C_2^{-1} \times P_2. \quad (6.18)$$

The combined covariance matrix, C is given by,

$$C = (C_1^{-1} + C_2^{-1})^{-1}, \quad (6.19)$$

¹No relation to the MuGirl algorithm!



with associated χ^2 of

$$\chi_{\text{match}}^2 = (P - P_1)^T \times C_1^{-1} \times (P - P_1) + (P - P_2)^T \times C_2^{-1} \times (P - P_2). \quad (6.20)$$

To ensure a good quality combined track, it is required that $\chi_{\text{match}}^2/n_{\text{dof}} < 15$. As information is used from two separate subsystems, these are known as *combined* muons.

As the muon chambers extend out to $|\eta| < 2.7$ compared to $|\eta| < 2.5$ of the inner detector, it is possible that muon segments are reconstructed in this region without a corresponding inner detector track. These are labelled as *standalone* muons.

Muon Selection Efficiency

In addition to the track quality cuts, kinematic cuts are also imposed, requiring $p_T > 5$ GeV, with acceptance cuts of $|\eta| < 2.5$ and $|\eta| < 2.7$ for the combined and standalone candidates respectively. Muons are also required to be isolated, using the ratio

$$I_R^{0.4} = \frac{p_T(\Delta R < 0.4)}{p_T}, \quad (6.21)$$

imposing $I_R^{0.4} < 0.2$ for both types of muon.

Similar truth matching criteria as used in the electron case can be used to define muon efficiency and purity, using expressions analogous to equations 6.15 and 6.16. Cuts are made on the true muons requiring $p_T > 1$ GeV and $|\eta| < 2.7$.

Figure 6.3 shows the muon efficiency as a function of true p_T , η and ϕ . Sharp drops in efficiency are observed at $|\eta| \sim 0, 1.5$, due to gaps in detector coverage due to services such as power cables and cooling systems. The standalone muons are also seen to extend the acceptance of the muons in regions $|\eta| > 2.5$. The slight inefficiencies seen at $\phi = -1.0, -2.2$ are due to the support feet of the ATLAS experiment.

The resolution of the muon reconstruction is shown in the lower plots of figure 6.3. A Gaussian fit to the combined muon p_T resolution in the central region gives a width of 2.8%, comparable to that of electrons. A tail is seen where the reconstructed p_T is underestimated, due to uncorrected energy losses from effects such as ionisation. Unlike the electron case, bremsstrahlung effects are suppressed due to the higher muon mass.

A summary of the cuts used to select the two types of muon and the corresponding efficiencies are shown in table 6.6.

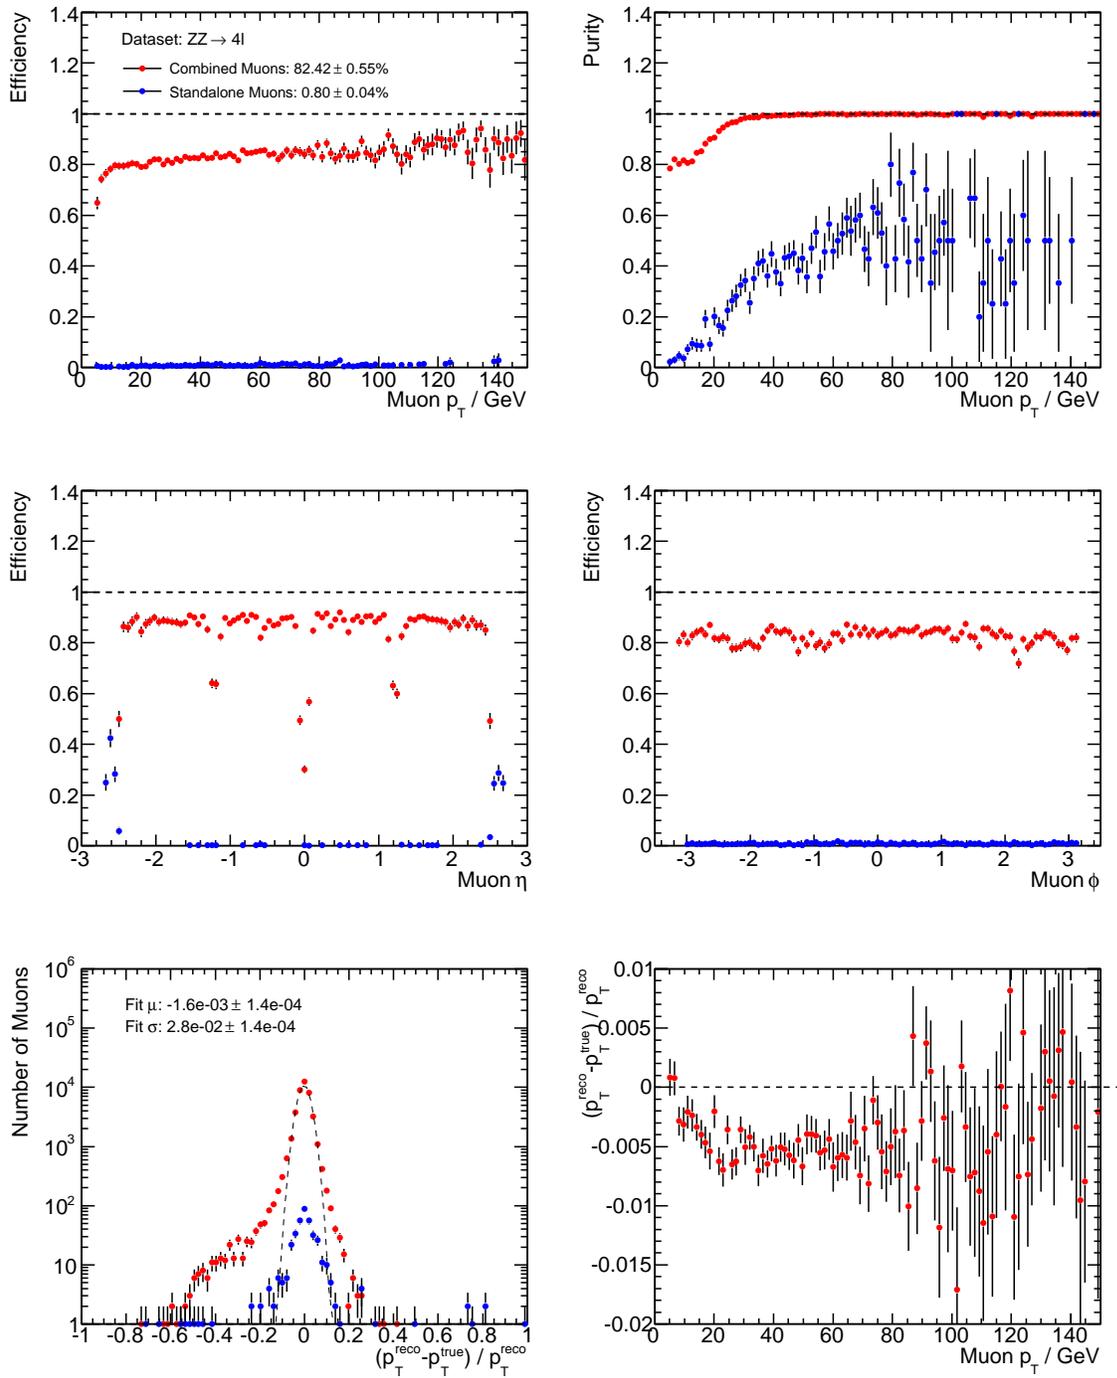


Figure 6.3: Efficiency, purity and resolution plots for muons in the analysis.



Cut	Combined	Standalone
ID track	yes	no
p_T	> 5 GeV	> 5 GeV
$ \eta $	< 2.5	< 2.7
Muon track χ^2/n_{dof}	< 15	< 15
Combined fit χ^2/n_{dof}	< 15	N/A
Isolation $I_R^{0.4}$	< 0.4	< 0.4
Efficiency	$82.4 \pm 0.6\%$	$0.80 \pm 0.04\%$

Table 6.6: Summary of muon pre-selection cuts.

6.4.3 Jet Selection

Quarks and gluons (collectively known as partons) are produced in high energy collisions, but unlike electrons and muons, are never observed directly. Partons carry the strong colour charge, and as such are confined to short distances corresponding to $\Lambda_{\text{QCD}} \sim 200$ MeV. At larger distances, high energy partons will hadronise into a jet of particles. In ATLAS, jets are reconstructed using specialised algorithms, designed such that the jet properties, such as energy and momentum, reflect as closely as possible the short distance dynamics of the underlying partons.

Jet Algorithms

A number of different algorithms are available for jet formation, as discussed in reference [97]. In this analysis, jets are chosen using the k_{\perp} algorithm [98], where calorimeter deposits are grouped by defining the following quantities between pairs of calorimeter energy towers:

$$d_{ij} = \min(k_{Ti}^{2p}, k_{Tj}^{2p}) \frac{(\Delta R_{ij})^2}{R^2} \quad (6.22)$$

$$d_{iB} = k_{Ti}^{2p} \quad (6.23)$$

where k_{Ti} is the transverse momentum of particle i and ΔR_{ij} is the angular distance between particles i and j . All possible combinations are computed, and the particles

in the pair with the smallest d_{ij} are combined. The computation is repeated until a combination has $d_{ij} < d_{iB}$, when it is labelled as a jet and removed from the list. In the specific case of the anti- k_{\perp} algorithm, $p = -1$, so that soft objects are merged with a central hard object, and the jet boundary is unaffected by soft radiation and hence is infra-red safe. The parameter R sets the resolution at which jets are resolved, where $R = 0.4$ for jets in this analysis.

As discussed in [99], the anti- k_{\perp} algorithm was found to be the best performer in a number of areas, such as reconstruction efficiency and computation time, when compared to other jet algorithms in ATLAS.

Jet Preselection Cuts

Jets reconstructed with the anti- k_{\perp} algorithm are pre-selected for this analysis if they are within $|\eta| < 3.0$ and satisfy $p_T > 1$ GeV.

6.4.4 Missing Energy Selection

The missing transverse energy (\cancel{E}_T) is used as a signature for particles which leave the detector without interacting, and is a good discriminant for the $ZZ \rightarrow ll\nu\bar{\nu}$ channel. The \cancel{E}_T is calculated by first constructing topological clusters from electromagnetic-scale calorimeter deposits [100]. The x - and y - components of the missing energy are calculated by summing over the transverse energy measured in the topological cells i :

$$\cancel{E}_{x,y} = \cancel{E}_{x,y}^{\text{Calo}} = - \sum_i E_{x,y}. \quad (6.24)$$

The missing transverse energy is then calculated as:

$$\cancel{E}_T = \sqrt{\cancel{E}_x^2 + \cancel{E}_y^2}. \quad (6.25)$$

In this analysis, an object-based \cancel{E}_T quantity is used (METRefFinal [101]), the recommended default for ATLAS analyses. To calculate METRefFinal, calorimeter energy deposits are associated with high p_T objects such as electrons, photons, muons and jets. The transverse energy of the objects is then used in equation 6.24 as a replacement for the original cell energy. As the objects have a more accurate calibration than the bare calorimeter cell calibration, the accuracy of \cancel{E}_T reconstruction is improved.



6.4.5 Z Reconstruction

Z bosons decaying into charged leptons can be reconstructed by summing the 4-momenta of their daughter particles. Pairs of same flavour, oppositely charged leptons, which can be electrons or muons as described in sections 6.4.1 and 6.4.2 respectively, are selected to form Z candidates. The ΔR between the two daughter leptons is required to be $\Delta R > 0.1$.

In multi-lepton events, a number of combinations of leptons are possible. For example, if two electrons and two positrons ($2e^-2e^+$) are present in an event, four Z candidates can be constructed. In the case where two candidates share the same daughter lepton, a choice must be made about which to keep. To resolve such ambiguities, the list of all possible lepton pairs is constructed, and ordered by $\Delta m = |m_{ll} - 91.2 \text{ GeV}|$. The candidate with the lowest Δm is retained and the second candidate examined. If the second candidate has any leptons in common with the first, it is rejected and the third candidate is examined, and so on. This process is repeated until up to two Z bosons have been selected, each time checking for overlap with the decay products of currently selected candidates.

Event Filtering

In order to reduce processing time of the n -tuples, a filter was applied to background datasets requiring at least one Z boson candidate to have been selected in the event. The filter efficiency is recorded on the n -tuple so that the overall event yield can be correctly calculated.

The preselection efficiency and number of events after preselection for the 10 TeV and 7 TeV simulated samples is shown in the last two columns of tables 6.1 and tables 6.2 respectively.

6.4.6 Truth Selection

True Z bosons in the event are also retained in the n -tuple for comparison with their reconstructed counterparts. For some generators, such as PYTHIA, this is a straightforward exercise as the event record contains the Z particles themselves. In this case the decay chain can be followed to find the daughter particles.

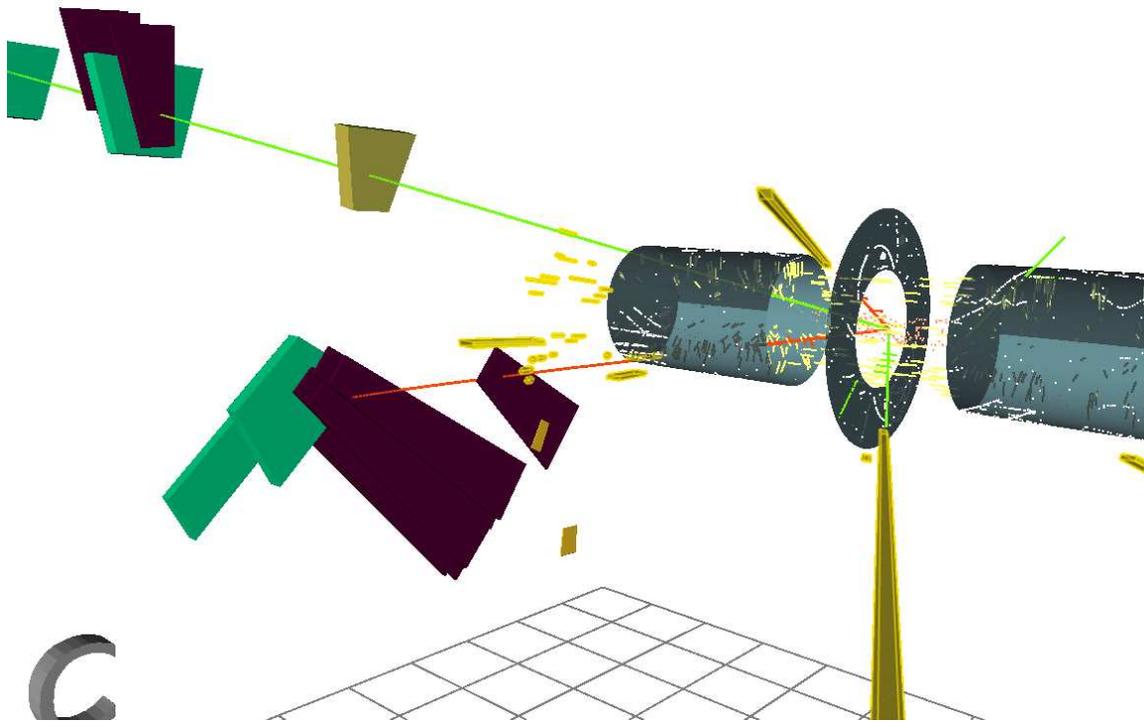


Figure 6.4: A simulated $ZZ \rightarrow e^+e^-\mu^+\mu^-$ event, displayed in the ATLAS event display Virtual Point 1. The event is from dataset 109291, event number 132740, at a centre-of-mass energy of 10 TeV. An animated version of this event can be found in the lower right corner of this thesis.

Some generators, for example SHERPA, do not contain Z particles, only vertices with four outgoing leptons. In this case a similar technique is applied as used with the reconstructed Z candidates, making pairs of leptons and removing any repeats.

6.5 $ZZ \rightarrow lll$ Event Selection

The $ZZ \rightarrow lll$ channel has a very distinct signature, with a final state containing four high transverse momentum leptons. An example of a simulated $ZZ \rightarrow e^+e^-\mu^+\mu^-$ event is shown in figure 6.4. In this section, a number of analysis cuts are defined with the aim of separating signal and background processes. For this channel, many of the cuts used are defined in common with the ATLAS Higgs group, for whom SM ZZ production is an important background to $H \rightarrow ZZ^{(*)} \rightarrow lll$.



6.5.1 Main Backgrounds

The main source of backgrounds in the $ZZ \rightarrow lll$ channel are events in which jets or other particles are mis-reconstructed as leptons. Fake electrons can arise if charged particles, such as pions, leave electron-like showers in the calorimeter. The fake rate is generally lower for muons, as few particles have sufficient energy to punch-through the calorimeters to the muon spectrometer.

The largest backgrounds are channels in which a single Z boson is produced, together with two fake electrons reconstructed from jets. In general, the mis-identification rate is very low, but the single Z cross section is $\sim 10^5$ times higher than for $ZZ \rightarrow lll$. Background channels of this type relevant to this analysis include the single Z channel (figure 1.11), and a number of processes which explicitly specify partons in the final state. These include the Alpgen $Z + n$ Jets samples, when n partons are included in the matrix element calculation. Figure 1.11 shows an example of Z production in which a single gluon is also present in the final state. $Zb\bar{b}$ production, as shown in figure 1.13, contains a single true Z and a pair of b -quark jets in the final state.

6.5.2 Cut Definition

The first stage in reducing the single Z background is the requirement that there are at least two pairs of same-flavour oppositely-charged leptons in the event to form a pair of Z bosons as described in section 6.4.5.

In the case of Z bosons decaying to an electron pair, an additional quality cut is applied requiring that at least one electron in each pair has a track with a hit in the innermost layer of the pixel detector, the b -layer. This cut reduces backgrounds from electrons which do not originate from the primary pp interaction. A particular example are high p_T photons, which can convert into two electrons when passing through detector material.

A cut is also applied for Z bosons decaying to muons, requiring at least one combined muon in each pair. This is applied to reduce combinatorial background arising from two poorly reconstructed standalone muon tracks.

The two Z candidates are ordered by the absolute difference between the true and reconstructed Z mass, and labelled Z_1 and Z_2 respectively. The daughter leptons from each Z are subsequently ordered by descending transverse momentum, and labelled l_1

and l_2 (also referred to as leading and trailing respectively). Overall, the four leptons in the event are labelled as:

$$\underbrace{l_1(Z_1), l_2(Z_1)}_{p_T(l_1) > p_T(l_2)} \text{ and } \underbrace{l_1(Z_2), l_2(Z_2)}_{p_T(l_1) > p_T(l_2)} \quad (6.26)$$

$$\underbrace{\hspace{10em}}_{|m_{Z_1} - 91.2 \text{ GeV}| < |m_{Z_2} - 91.2 \text{ GeV}|}$$

The first kinematic cuts are on the lepton transverse momenta, requiring that

$$p_T(l_1) > 20 \text{ GeV} \text{ and } p_T(l_2) > 10 \text{ GeV} \quad (6.27)$$

for both Z candidates. The distributions in the upper plots of figure 6.5 show the lepton p_T distributions for the signal and single Z background, after preselection. In the case of the single Z channel, the reconstructed Z_1 is likely to be from a true Z , while Z_2 is falsely reconstructed from other leptons in the event, which are not true Z decay products. These cuts reduce the fake lepton background, which is found predominantly at low p_T .

The next set of cuts are made to the dilepton invariant mass. The first requires that the mass of the lepton pair closest to the true Z mass lies within the window

$$|m_{Z_1} - 91.2 \text{ GeV}| < 20 \text{ GeV}. \quad (6.28)$$

This reduces background from channels in which a non-resonant lepton pair is produced, such as from $t\bar{t} \rightarrow WWb\bar{b} \rightarrow ll + X$, as shown in the lower-left plot of figure 6.5. A looser mass cut is imposed on the second Z , requiring that

$$m_{Z_2} > 20 \text{ GeV}, \quad (6.29)$$

which allows the off-shell Z^* to be retained, whilst also reducing combinatorial backgrounds. The effect of the Z^* can be seen in the lower-right plot of figure 6.5, which shows an enhanced signal cross section at low m_Z . The cut reduces background from channels where the second lepton pair does not come from a true Z , such as $Zb\bar{b}$.



Cut	$ZZ \rightarrow 4e$	$ZZ \rightarrow 4\mu$	$ZZ \rightarrow 2e2\mu$	$Z \rightarrow ll$	Z+Jets	WZ	Zbb
Before Cuts	17.78	17.78	35.56	$2.93 \times 10^{+6}$	$2.82 \times 10^{+5}$	879	$6.90 \times 10^{+4}$
$n_Z = 2$	4.63	6.22	10.3	36.1	8.01	0.713	6.79
Opposite Sign	4.38	6.20	9.98	19.6	4.67	0.281	4.34
B-layer electrons	4.35	6.20	9.84	8.77	3.67	0.252	2.72
Combined Muons > 0	4.35	6.18	9.82	8.77	3.67	0.252	2.67
$p_T(l_{Z_1}^1) > 20$ GeV	4.31	6.08	9.70	8.77	3.67	0.252	2.49
$p_T(l_{Z_1}^2) > 10$ GeV	4.25	5.94	9.48	8.15	3.34	0.252	2.40
$p_T(l_{Z_2}^1) > 20$ GeV	3.92	5.43	8.79	0.309	0.335	0.241	0.417
$p_T(l_{Z_2}^2) > 10$ GeV	3.65	5.08	8.18	0.0	0.0	0.0938	0.172
$ m_{Z_1} - 91.2$ GeV < 20 GeV	3.64	5.05	8.06	0.0	0.0	0.0794	0.171
$m_{Z_2} > 20$ GeV	3.55	4.93	7.86	0.0	0.0	0.0794	0.00862
Number in 1 fb^{-1}	3.55 ± 0.06	4.93 ± 0.06	7.86 ± 0.08	< 1	< 2	$0.08^{+0.06}_{-0.03}$	$0.009^{+0.4}_{-0.002}$
Events passing	4167	5776	9210	0	0	6	13
Overall Efficiency	20.0%	27.7%	22.1%	0%	0%	0.0045%	0.00034%

Table 6.7: Expected yields in 1 fb^{-1} at 10 TeV, for the $ZZ \rightarrow lll$ channel.

Cut	$ZZ \rightarrow 4e$	$ZZ \rightarrow 4\mu$	$ZZ \rightarrow 2e2\mu$	$Z \rightarrow ll$	Z+Jets	WZ	Zbb
Before Cuts	10.87	10.87	21.73	$2.10 \times 10^{+6}$	$1.56 \times 10^{+5}$	521	$3.55 \times 10^{+4}$
$n_Z = 2$	3.11	4.16	6.91	23.7	4.97	0.420	3.52
Opposite Sign	2.95	4.15	6.69	12.6	2.69	0.167	2.27
B-layer electrons	2.93	4.15	6.59	5.65	2.07	0.155	1.42
Combined Muons > 0	2.93	4.14	6.59	5.65	2.07	0.155	1.39
$p_T(l_{Z_1}^1) > 20$ GeV	2.90	4.07	6.49	5.65	2.07	0.155	1.29
$p_T(l_{Z_1}^2) > 10$ GeV	2.86	3.97	6.34	5.25	1.84	0.155	1.24
$p_T(l_{Z_2}^1) > 20$ GeV	2.63	3.61	5.85	0.187	0.163	0.147	0.207
$p_T(l_{Z_2}^2) > 10$ GeV	2.44	3.36	5.44	0.0	0.0	0.0587	0.0917
$ m_{Z_1} - 91.2$ GeV < 20 GeV	2.43	3.34	5.35	0.0	0.0	0.0500	0.0909
$m_{Z_2} > 20$ GeV	2.37	3.25	5.21	0.0	0.0	0.0500	0.00414
Number in 1 fb^{-1}	2.37 ± 0.04	3.25 ± 0.04	5.21 ± 0.06	< 0.7	< 2	$0.05^{+0.04}_{-0.02}$	$0.004^{+0.4}_{-0.001}$
Events passing	4167	5776	9210	0	0	6	13
Overall Efficiency	21.8%	29.9%	24.0%	0%	0%	0.0045%	0.00034%

Table 6.8: Expected yields in 1 fb^{-1} at 7 TeV, for the $ZZ \rightarrow lll$ channel, calculated using 10 TeV reweighted data.

Cut	$ZZ \rightarrow 4e$	$ZZ \rightarrow 4\mu$	$ZZ \rightarrow 2e2\mu$	$Z \rightarrow ll$ PYTHIA	$Z \rightarrow ll$ MC@NLO	WZ	$t\bar{t}$
Before Cuts	10.60	10.60	21.20	$2.79 \times 10^{+6}$	$2.66 \times 10^{+6}$	534	$1.15 \times 10^{+4}$
$n_Z = 2$	2.74	4.05	6.48	36.7	24.6	0.483	2.61
Opposite Sign	2.52	4.01	6.24	19.5	11.4	0.273	2.09
B-layer electrons	2.51	4.01	6.23	17.9	10.7	0.273	1.57
Combined Muons > 0	2.51	3.99	6.22	14.3	8.05	0.270	1.57
$p_T(l_{Z_1}^1) > 20$ GeV	2.48	3.93	6.14	13.9	7.82	0.270	1.04
$p_T(l_{Z_1}^2) > 10$ GeV	2.44	3.84	6.02	13.1	6.93	0.262	0.522
$p_T(l_{Z_2}^1) > 20$ GeV	2.24	3.49	5.51	1.42	0.447	0.240	0.522
$p_T(l_{Z_2}^2) > 10$ GeV	2.07	3.25	5.14	0.618	0.224	0.0705	0.0
$ m_{Z_1} - 91.2$ GeV < 20 GeV	2.06	3.24	5.07	0.618	0.224	0.0582	0.0
$m_{Z_2} > 20$ GeV	2.01	3.17	4.92	0.408	0.224	0.0564	0.0
Number in 1 fb^{-1}	2.01 ± 0.04	3.17 ± 0.06	4.92 ± 0.07	$0.4^{+0.7}_{-0.2}$	$0.2^{+0.8}_{-0.1}$	$0.06^{+0.03}_{-0.02}$	< 1
Events passing	2108	3328	5157	2	1	36	0
Overall Efficiency	19.0%	30.0%	23.2%	$6.8 \times 10^{-6}\%$	$3.4 \times 10^{-6}\%$	0.0078%	0%

Table 6.9: Expected yields in 1 fb^{-1} at 7 TeV, for the $ZZ \rightarrow lll$ channel.

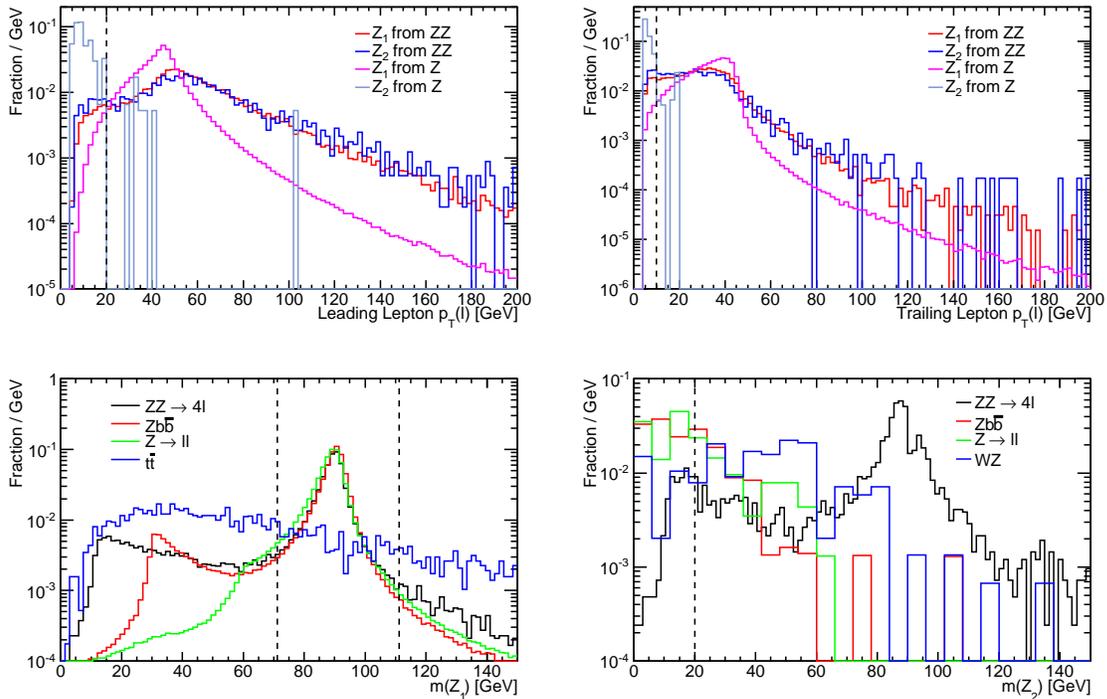


Figure 6.5: Shapes of kinematic distributions for variables used in the selection of $ZZ \rightarrow llll$ events. The upper-left and right plots show the leading and trailing lepton p_T respectively. The lower plots show the invariant mass distributions for lepton pairs closest to the true Z mass (left) and further from it (right). The plots are made from 7 TeV simulation and normalised to unit area, with the exception $Zb\bar{b}$, which is from the re-weighted sample.

6.5.3 Expected Yields

The expected event yields after 1 fb^{-1} of ATLAS data-taking at 10 TeV are shown in table 6.7. The expected results using 10 TeV samples re-weighted to 7 TeV are shown in table 6.8, while those for samples generated directly at 7 TeV can be found in table 6.9. The first three columns of each table show the signal channel, divided up into the three final state lepton combinations ($4e, 4\mu, 2e2\mu$). The remaining columns show the expected background yields, summed over the three lepton combinations.

For the tables constructed with 10 TeV datasets (6.7 and 6.8), the $Z \rightarrow ll$ channel comes from the sum of the yields from the $Z \rightarrow ee$ and $Z \rightarrow \mu\mu$ channels (datasets 106050 and 106051 in table 6.1). The $Z + \text{Jets}$ column is the sum of the Alpgen samples 107650-107654 and 107660-107664, while the WZ column presents the sum of the W^+Z and W^-Z MC@NLO samples 105971 and 105972. Finally, the $Zb\bar{b}$ column lists the



expected yields summed over the Alpgen samples 109300-109308, 109400-109401 and 109405-109406.

In the case of the 7 TeV datasets (table 6.9), slightly different samples were available for the background study. For the single Z background, large statistics datasets were available with PYTHIA (datasets 106078 and 106079) and MC@NLO (datasets 106087 and 106088). The other columns use the equivalent 7 TeV datasets as described in the 10 TeV case, except for the $Zb\bar{b}$ samples which were not available at the time of writing.

Each row of the tables shows the expected number of events after a specific cut is applied, calculated as follows

$$n_{\text{expected}} = \mathcal{L} \times \sigma \times k \times \epsilon_{\text{gen}} \times \epsilon_{\text{cut}}, \quad (6.30)$$

where \mathcal{L} is the integrated luminosity, σ is the cross section, k is the k-factor and ϵ_{gen} is the generator filter efficiency, all taken from tables 6.1 and 6.2. The remaining value, ϵ_{cut} is the efficiency of the selection after each cut, given by

$$\epsilon_{\text{cut}} = \frac{\sum_{i=0}^{n_{\text{pass}}} w_i^{\text{gen}} w_i^{\text{re-weight}}}{\sum_{i=0}^N w_i^{\text{gen}}}. \quad (6.31)$$

The numerator represents the sum over events passing all cuts, weighted by both the generator weight and the re-weighting weight. The generator weight is normally $w_i^{\text{gen}} = 1$, except in the case of MC@NLO, where it can take values ± 1 , as discussed in section 1.5. The second term, $w_i^{\text{re-weight}}$ is unity except when considering the 10 TeV \rightarrow 7 TeV re-weighted events. The denominator is the sum over all events before the pre-selection filter, weighted by the generator weight.

The ante-penultimate row of each table shows the expected number of events after all cuts, with errors representing the statistical uncertainty on the number of events selected. If more than ten events pass cuts, binomial errors are computed. If fewer than ten events remain in a particular channel, asymmetric errors are calculated from the 68.27% confidence limit intervals given in table II of [102]. In the case where no events remain after all cuts, only the upper 68.27% confidence limit is quoted. Errors from different channels in a single column are combined by addition in quadrature. The penultimate row of each table lists the number of Monte Carlo events remaining after all

cuts. The final row shows the selection efficiency, calculated as the ratio of events after cuts to the number before. The statistical uncertainty on the efficiency calculations for signal events is around 2%.

Signal Expectation

The total numbers of signal events expected for 1 fb^{-1} of data are shown in table 6.14 for the three different energy schemes. Overall it is expected that 10.1 ± 0.1 (stat) signal events will be seen at 7 TeV with a selection efficiency of 23.8%. It should be noted that the original PYTHIA Monte Carlo sample also contains $Z \rightarrow \tau\tau$ decays, which are not included in the yields or efficiency calculations. There are some differences between the selection efficiencies between the 7 TeV and re-weighted samples, which can be attributed to differences in reconstruction algorithms and re-weighting errors.

Background Estimation

This section aims to estimate the total background that is expected to pass cuts, using a combination of all the channels available. The total backgrounds for each of the three energy schemes is shown in table 6.14.

The only channels which pass selection cuts with more than two events are from the WZ and $Zb\bar{b}$ channels. The cross sections of these processes are such that the expected number of events is small compared to the signal.

The only other background passing all the cuts is the single Z channel. In fact, no events pass the cuts from the 10 TeV samples, whereas at 7 TeV a single event from the MC@NLO $Z \rightarrow \mu\mu$ sample passes the $ZZ \rightarrow 4\mu$ cuts, and two PYTHIA events pass the $ZZ \rightarrow 2e2\mu$ cuts. Such low statistics make a background estimation challenging, especially as a single simulated event scales to a significant fraction of the expected signal. Employing a Toy Monte Carlo method would be inappropriate here, as the background from the single Z comes from mis-identified leptons and not the tail of a true kinematic distribution. In some cases it is also possible to extrapolate the background distribution into the cut region, but this is also difficult with the single Z background as the number of events passing the earlier cuts is also small.

It is possible, however, to put upper confidence limits on the expected number of events passing cuts. In [102], the 68.27% confidence limit on the mean of a Poisson



distribution for no observed events is 1.29. When calculating the limit on the single Z yield, there are two background channels ($Z \rightarrow ee$ and $Z \rightarrow \mu\mu$) that could potentially pass three sets of cuts ($4e$, 4μ and $2e2\mu$). It will be assumed that of these, only four combinations are likely to be important, namely $Z \rightarrow ee$ passing the $4e$ selection, $Z \rightarrow \mu\mu$ passing the 4μ selection, and both backgrounds passing $2e2\mu$.

For the 10 TeV study, the estimate of the upper limit for the Z background is ~ 1.0 event, given that no events pass the selection. In the $10 \rightarrow 7$ TeV re-weighted case, this number is scaled down to ~ 0.7 events. In the case of the 7 TeV single Z background, confidence limits are calculated given one and two events passing cuts from the MC@NLO and PYTHIA samples respectively. The two generators give consistent predictions within the computed errors, as shown in table 6.9.

It should be noted that these limits are expected to be conservative, and the actual background levels may be much lower.

6.6 $ZZ \rightarrow ll\nu\bar{\nu}$ Event Selection

The experimental signature for the $ZZ \rightarrow ll\nu\bar{\nu}$ channel is a single pair of same-flavour, oppositely-charged leptons forming a Z , and missing transverse energy from the two neutrinos. The $ZZ \rightarrow ll\nu\bar{\nu}$ channel is a more challenging channel than $ZZ \rightarrow llll$, as the missing energy from the $Z \rightarrow \nu\bar{\nu}$ is more easily faked than two leptons. The increased cross section compared to the four lepton channel, however, makes it a viable prospect for analysis with early early ATLAS data. An example of a $Z \rightarrow e^+e^-\nu\bar{\nu}$ can be found in figure 6.6.

6.6.1 Main Backgrounds

There are a number of different backgrounds to the $ZZ \rightarrow ll\nu\bar{\nu}$ channel. The first are the single Z channels, where a true Z is successfully reconstructed, and the high cross section leads to the tail of the \cancel{E}_T distribution being accepted by cuts. In particular a spuriously high \cancel{E}_T measurement can be made if jets are present in the event that lie in poorly instrumented regions of the detector.

Backgrounds from the single W channel are also important. A true high p_T lepton is present in the event, combined with significant \cancel{E}_T . If a single additional lepton is

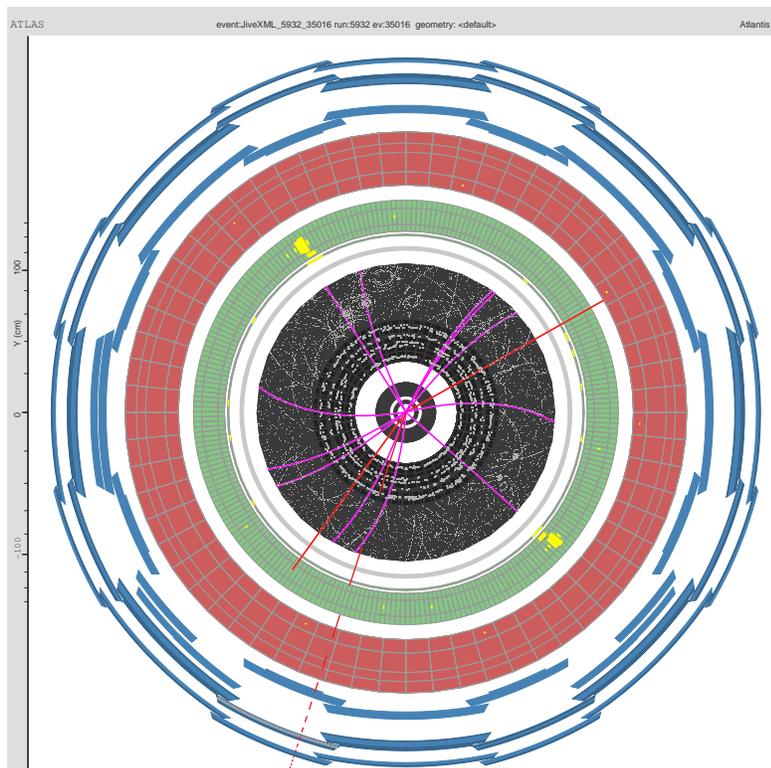


Figure 6.6: A simulated $ZZ \rightarrow e^+e^-\nu\bar{\nu}$ event, viewed along the z -axis in the Atlantis event display. The event is from dataset 105932, event number 35016. Reconstructed tracks with $p_T > 1$ GeV are shown in purple. The true neutrinos are displayed as solid red lines, with the reconstructed \cancel{E}_T direction shown as the dashed line.

reconstructed, this can also give a similar signature to the signal. The cross section for single W events is even higher than single Z , being $\sim 10^5$ times above that of the signal.

Another category of backgrounds are those in which a pair of leptonically decaying W bosons appear in the final state, either directly via $W^+W^- \rightarrow l\nu l\bar{\nu}$, or via the $t\bar{t} \rightarrow WWb\bar{b} \rightarrow l^+\nu l^-\bar{\nu}$ channel. The presence of two neutrinos in this channel also leads to a high \cancel{E}_T measurement. The $t\bar{t}$ background is particularly dangerous, as the cross section for this process is $\sim 10^3$ times higher than the signal.

One final type of background considered here is where a real Z is produced, with additional leptons which do not fall in the detector acceptance. These include the $ZZ \rightarrow ll\bar{l}l$ channel itself, where the lepton pair from the second Z is not reconstructed. More important is diboson production in the WZ channel, where the lepton from the W decay is not reconstructed. This leaves an event signature with two high p_T leptons from a true Z , and high \cancel{E}_T from the W . The cuts in the next section aim to reduce these backgrounds.



6.6.2 Cut Definition

The first selection requirement is that a pair of same-flavour, oppositely-charged leptons is present in the event. As with the $ZZ \rightarrow ll\bar{l}$ channel, an additional cut is applied to electron daughters, requiring that at least one has a hit in the pixel b -layer. For the $ZZ \rightarrow ll\nu\bar{\nu}$ channel, two combined muons are explicitly required to form a Z boson. The contribution of standalone muons was found to distort the missing transverse energy spectrum at values $\cancel{E}_T \sim 50$ GeV, in the same region where the signal becomes significant.

The first cuts are applied to the lepton kinematics, where it is required that both reconstructed leptons from the Z satisfy

$$p_T(l) > 20 \text{ GeV}, \quad (6.32)$$

and that the dilepton invariant mass is close to the true Z mass,

$$|m_Z - 91.2 \text{ GeV}| < 20 \text{ GeV}. \quad (6.33)$$

This mass window cut is effective at removing any non-resonant backgrounds, in particular WW and $t\bar{t}$ backgrounds, as illustrated by the plots in figure 6.7.

The two neutrinos present in the signal channel give rise to signal events with high \cancel{E}_T , which is exploited by applying the cut

$$\cancel{E}_T > 50 \text{ GeV}, \quad (6.34)$$

which serves to remove most of the single Z background, as shown in figure 6.7. The single Z \cancel{E}_T spectrum is shown to fall more sharply than the signal, as no high momentum neutrinos are present in the final state. As shown in the same plot, this cut is not effective at reducing the $t\bar{t}$ background, and so a second \cancel{E}_T related cut is introduced, requiring that

$$0.65 < \frac{\cancel{E}_T}{p_T(ll)} < 1.35, \quad (6.35)$$

the motivation being that, for the signal, the two Z bosons tend to be produced back-to-back, with similar transverse momenta. The distribution of this variable for the signal can be seen in figure 6.7, and does indeed show a peak around one. The same plot also

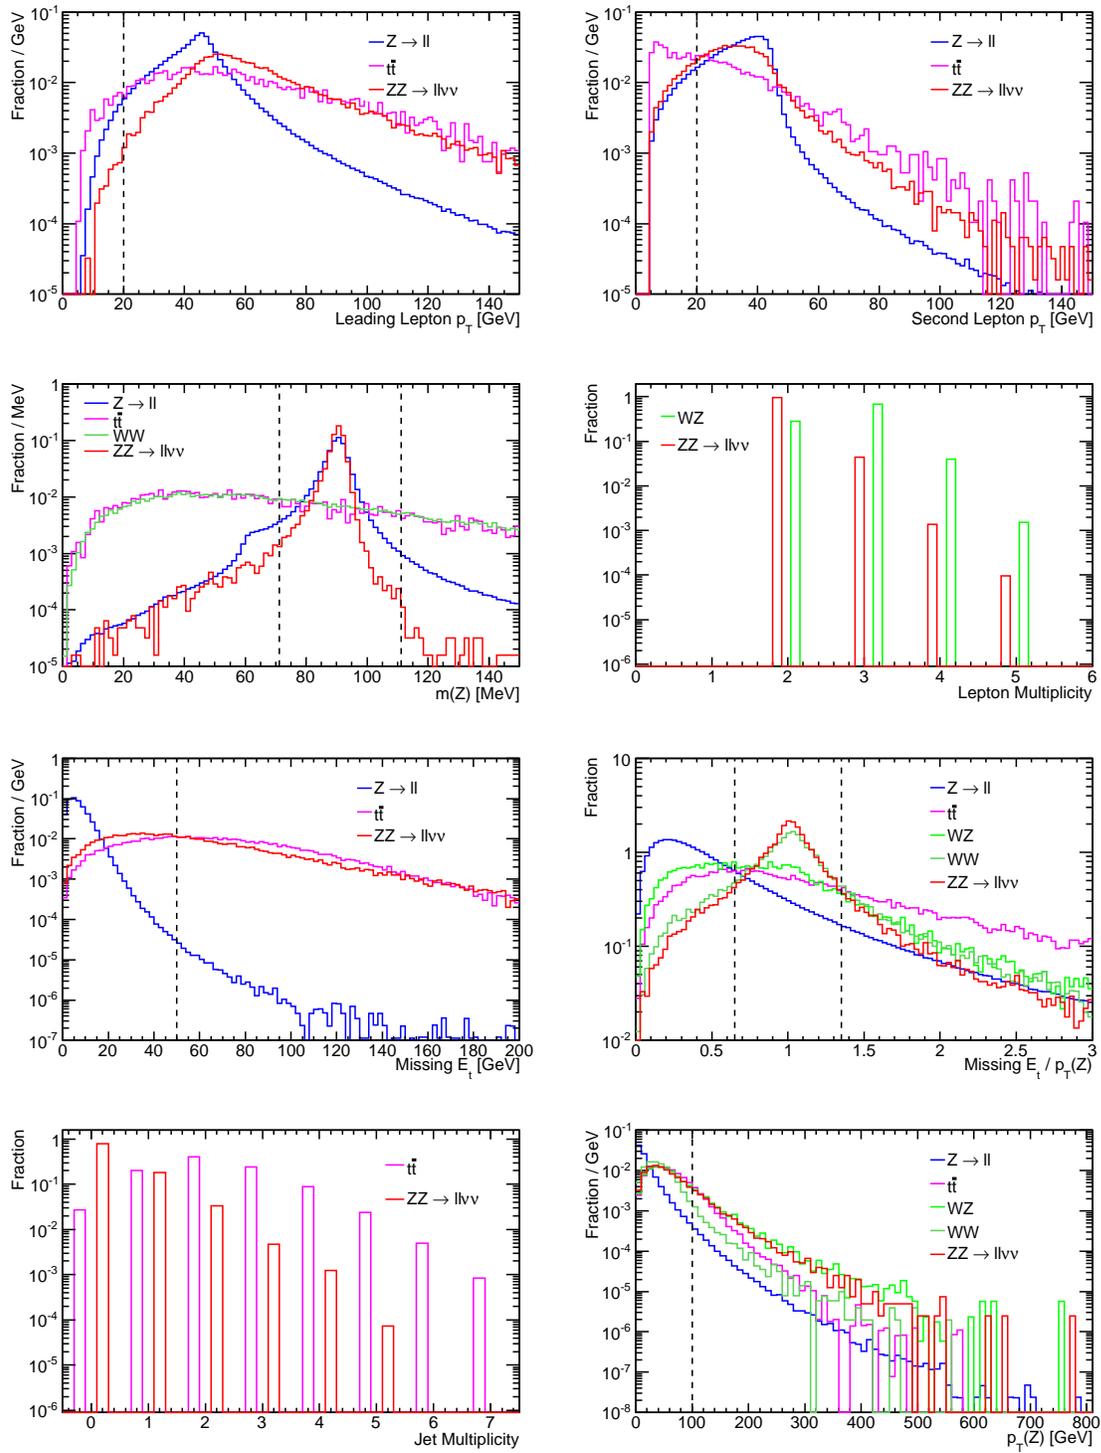
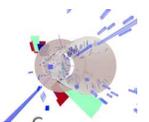


Figure 6.7: $ZZ \rightarrow ll\nu\bar{\nu}$ kinematic variables for signal and background channels.



shows the effectiveness of this cut at reducing the WW backgrounds, where no such peak is visible.

Two additional veto cuts are also required to reduce specific backgrounds, the first being WZ , which is reduced by requiring that the total number of reconstructed leptons is exactly two,

$$\begin{aligned} n(\text{physics electrons}) + n(\text{veto electrons}) + \\ n(\text{combined muons}) + n(\text{standalone muons}) = 2. \end{aligned} \quad (6.36)$$

The distribution of total lepton multiplicity is shown in figure 6.7. As the plots are made after the requirement that each event has $n_Z > 0$, no events are seen with $n(\text{leptons}) < 2$.

The second veto is placed on the number of jets in the event, to reduce events from the $t\bar{t}$ channel. In this channel, the top quarks can decay via $t \rightarrow bW (\rightarrow l\nu)$, producing two oppositely-charged leptons in the final state and two b -quark jets. The high multiplicity of jets in the $t\bar{t}$ channel can be exploited by vetoing on events which contain any jets satisfying

$$p_T > 30 \text{ GeV and } |\eta| < 3.0. \quad (6.37)$$

The multiplicity of these high p_T jets is shown in figure 6.7, and illustrates the effectiveness of such a veto in reducing the $t\bar{t}$ background.

The final cut to be applied relies on the fact that the two Z bosons are highly boosted as they are produced approximately back-to-back. This cut requires that

$$p_T(Z) > 100 \text{ GeV}, \quad (6.38)$$

as shown in figure 6.7, and serves to reduce the background from the single Z and $t\bar{t}$ channels.

6.6.3 Expected Yields

Tables 6.10, 6.11 and 6.12 show the expected number of events from the $ZZ \rightarrow ll\nu\bar{\nu}$ channel after each cut, calculated in an analogous way to those of section 6.5.3. The first two columns show the expected yields from the two signal channels ($ee\nu\nu$ and

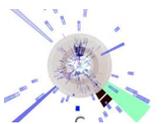
Cut	$ZZ \rightarrow 2e2\nu$	$ZZ \rightarrow 2\mu 2\nu$	$Z \rightarrow ll$	$Zb\bar{b}$	$t\bar{t}$	$ZZ \rightarrow lll$	WZ	WW
Before Cuts	123.7	123.7	$1.51 \times 10^{+6}$	$4.57 \times 10^{+4}$	$1.89 \times 10^{+4}$	187	586	$2.33 \times 10^{+3}$
$n(l^+l^-) > 0$	67.2	76.1	$1.46 \times 10^{+6}$	$2.25 \times 10^{+4}$	$8.33 \times 10^{+3}$	94.7	261	$1.17 \times 10^{+3}$
B-layer	67.0	76.1	$1.46 \times 10^{+6}$	$2.24 \times 10^{+4}$	$8.23 \times 10^{+3}$	94.1	260	$1.16 \times 10^{+3}$
Combined Muons > 0	67.0	70.8	$1.38 \times 10^{+6}$	$2.16 \times 10^{+4}$	$8.06 \times 10^{+3}$	90.2	250	$1.12 \times 10^{+3}$
$p_T(l_{\frac{1}{2}}^1) > 20$ GeV	66.8	70.6	$1.34 \times 10^{+6}$	$2.09 \times 10^{+4}$	$7.70 \times 10^{+3}$	85.4	249	$1.07 \times 10^{+3}$
$p_T(l_{\frac{1}{2}}^2) > 20$ GeV	55.6	59.9	$1.20 \times 10^{+6}$	$1.71 \times 10^{+4}$	$4.83 \times 10^{+3}$	68.1	210	736
$n_l = 2$	53.3	56.3	$1.16 \times 10^{+6}$	$1.54 \times 10^{+4}$	$4.07 \times 10^{+3}$	18.0	61.3	699
$ m_Z - 91.2$ GeV < 20 GeV	52.8	56.0	$1.11 \times 10^{+6}$	$1.46 \times 10^{+4}$	$1.19 \times 10^{+3}$	15.9	50.9	205
$\cancel{E}_T > 50$ GeV	27.5	29.5	283	48.2	781	2.24	20.5	73.9
$ p_T(Z)/\cancel{E}_T - 1 < 0.35$	19.9	21.3	54.6	11.3	270	1.22	11.8	45.4
$n_{\text{jets}} = 0$	16.0	18.1	10.9	1.50	8.73	0.655	7.57	35.8
$p_T(Z) > 100$ GeV	5.17	4.95	0.645	0.0819	0.860	0.112	1.41	0.527
Number in 1 fb^{-1}	5.2 ± 0.3	5.0 ± 0.4	$0.6^{+0.8}_{-0.3}$	$0.08^{+0.3}_{-0.05}$	$0.9^{+0.4}_{-0.3}$	0.11 ± 0.01	1.41 ± 0.1	$0.5^{+0.3}_{-0.2}$
Events passing	358	382	2	2	7	136	126	11
Overall Efficiency	4.2%	4.0%	$1.7 \times 10^{-05}\%$	0.0001%	0.00016%	0.036%	0.14%	0.0025%

Table 6.10: Expected yields in 1 fb^{-1} at 10 TeV, for the $ZZ \rightarrow ll\nu\bar{\nu}$ channel.

Cut	$ZZ \rightarrow 2e2\nu$	$ZZ \rightarrow 2\mu 2\nu$	$Z \rightarrow ll$	$Zb\bar{b}$	$t\bar{t}$	$ZZ \rightarrow lll$	WZ	WW
Before Cuts	76.4	76.4	$1.08 \times 10^{+6}$	$2.35 \times 10^{+4}$	$7.47 \times 10^{+3}$	121	347	$1.49 \times 10^{+3}$
$n(l^+l^-) > 0$	41.8	46.8	$1.05 \times 10^{+6}$	$1.16 \times 10^{+4}$	$3.29 \times 10^{+3}$	61.9	155	745
B-layer	41.7	46.8	$1.05 \times 10^{+6}$	$1.16 \times 10^{+4}$	$3.25 \times 10^{+3}$	61.5	155	742
Combined Muons > 0	41.7	43.8	$9.94 \times 10^{+5}$	$1.12 \times 10^{+4}$	$3.19 \times 10^{+3}$	59.2	149	714
$p_T(l_{\frac{1}{2}}^1) > 20$ GeV	41.6	43.6	$9.64 \times 10^{+5}$	$1.07 \times 10^{+4}$	$3.05 \times 10^{+3}$	55.7	149	685
$p_T(l_{\frac{1}{2}}^2) > 20$ GeV	34.6	36.9	$8.67 \times 10^{+5}$	$8.78 \times 10^{+3}$	$1.90 \times 10^{+3}$	44.3	125	471
$n_l = 2$	33.1	34.8	$8.35 \times 10^{+5}$	$7.89 \times 10^{+3}$	$1.60 \times 10^{+3}$	11.1	33.1	448
$ m_Z - 91.2$ GeV < 20 GeV	32.8	34.6	$7.97 \times 10^{+5}$	$7.49 \times 10^{+3}$	479	9.82	27.1	134
$\cancel{E}_T > 50$ GeV	16.6	17.6	179	17.6	313	1.26	10.5	46.6
$ p_T(Z)/\cancel{E}_T - 1 < 0.35$	12.1	12.7	35.3	4.35	110	0.676	6.09	29.3
$n_{\text{jets}} = 0$	9.90	10.8	6.66	0.633	3.81	0.348	3.91	23.5
$p_T(Z) > 100$ GeV	3.01	2.82	0.379	0.0314	0.331	0.0544	0.638	0.347
Number in 1 fb^{-1}	3.0 ± 0.2	2.8 ± 0.2	$0.4^{+0.5}_{-0.2}$	$0.031^{+0.3}_{-0.02}$	$0.3^{+0.2}_{-0.1}$	0.054 ± 0.005	0.64 ± 0.07	$0.3^{+0.3}_{-0.1}$
Events passing	358	382	2	2	7	136	126	11
Overall Efficiency	3.9%	3.7%	$1.7 \times 10^{-5}\%$	0.0001%	0.00016%	0.036%	0.14%	0.0025%

Table 6.11: Expected yields in 1 fb^{-1} at 7 TeV, for the $ZZ \rightarrow ll\nu\bar{\nu}$ channel, using datasets reweighted from 10 TeV.

Cut	$ZZ \rightarrow 2e2\nu$	$ZZ \rightarrow 2\mu 2\nu$	$Z \rightarrow ll$	$W \rightarrow l\nu$	$t\bar{t}$	$ZZ \rightarrow lll$	WZ	WW
Before Cuts	75.8	75.8	$1.86 \times 10^{+6}$	$7.84 \times 10^{+4}$	$7.70 \times 10^{+3}$	179	356	$1.51 \times 10^{+3}$
$n(l^+l^-) > 0$	41.0	50.4	$9.12 \times 10^{+5}$	$2.13 \times 10^{+4}$	$3.32 \times 10^{+3}$	56.9	157	751
B-layer	40.9	50.4	$9.11 \times 10^{+5}$	$2.11 \times 10^{+4}$	$3.32 \times 10^{+3}$	56.9	157	751
Combined Muons > 0	40.9	47.4	$8.67 \times 10^{+5}$	$2.00 \times 10^{+4}$	$3.26 \times 10^{+3}$	54.9	152	722
$p_T(l_{\frac{1}{2}}^1) > 20$ GeV	40.8	47.2	$8.44 \times 10^{+5}$	$1.63 \times 10^{+4}$	$3.07 \times 10^{+3}$	51.5	152	691
$p_T(l_{\frac{1}{2}}^2) > 20$ GeV	34.2	40.0	$7.63 \times 10^{+5}$	676	$1.90 \times 10^{+3}$	40.8	127	475
$n_l = 2$	32.6	38.4	$7.36 \times 10^{+5}$	656	$1.62 \times 10^{+3}$	10.3	35.3	454
$ m_Z - 91.2$ GeV < 20 GeV	32.4	38.2	$7.04 \times 10^{+5}$	193	472	9.10	28.0	138
$\cancel{E}_T > 50$ GeV	16.1	18.3	177	19.1	305	1.09	10.3	45.6
$ p_T(Z)/\cancel{E}_T - 1 < 0.35$	12.6	14.6	30.7	10.8	104	0.611	6.49	31.5
$n_{\text{jets}} = 0$	10.8	12.3	6.83	8.11	2.09	0.309	4.26	26.0
$p_T(Z) > 100$ GeV	2.95	3.26	0.209	0.0	0.0	0.0473	0.918	0.371
Number in 1 fb^{-1}	2.95 ± 0.09	3.26 ± 0.09	$0.2^{+0.6}_{-0.1}$	< 2	< 1	0.047 ± 0.007	0.92 ± 0.08	$0.37^{+0.1}_{-0.08}$
Events passing	1699	1875	1	0	0	53	384	34
Overall Efficiency	3.9%	4.3%	$5.1 \times 10^{-06}\%$	0%	0%	0.017%	0.12%	0.0038%

Table 6.12: Expected yields in 1 fb^{-1} at 7 TeV, for the $ZZ \rightarrow ll\nu\bar{\nu}$ channel.

$\mu\mu\nu\nu$), with the remaining columns giving the predicated number of background events, summing over the relevant datasets from table 6.2.

Signal Expectation

After all cuts have been applied, the expected number of reconstructed signal events is 6.2 ± 0.1 at 7 TeV energy with 1 fb^{-1} of data, with a selection efficiency of 4.1%. The total number of events is also consistent with the predictions from the re-weighted simulation.

Background Estimation

Considering first the 10 TeV and re-weighted samples, a combination of all background channels gives a total of $3.6_{-0.5}^{+0.9}$ (stat) events after cuts. The contribution from the $ZZ \rightarrow llll$, WZ and WW channels all have small statistical errors, and a total of seven Monte Carlo events pass the $t\bar{t}$ selection. Only one event passes cuts from each of the $Z \rightarrow ee$ and $Z \rightarrow \mu\mu$ channels, corresponding to limits of 0.8 and 0.5 at the 68.27% confidence level.

In the 7 TeV case, no events pass cuts from the single W or $t\bar{t}$ channels, with only a single event from $Z \rightarrow ll$. This corresponds to an upper 68.27% confidence limit of 0.3 events for the single Z channel, and 2.0 events for single W , where both the PYTHIA and MC@NLO statistics have been combined. In the case of $t\bar{t}$, the re-weighted limit will be used, as the 10 TeV sample contained ~ 10 times more events than at 7 TeV. As no 7 TeV $Zb\bar{b}$ sample was available, the re-weighted yields have been added to the 7 TeV total background.

6.7 Trigger Analysis

Before events from ATLAS are recorded on disk, they must pass a number of trigger levels, as explained in section 2.4. For this analysis, common triggers are used for the two channels, and are based on the requirement of a high p_T lepton in the final state.

For channels containing only electrons, events must pass the EM5 trigger at L1, followed by e5_medium at L2 and the event filter, referring to electrons with $E_T > 5 \text{ GeV}$.

In the case of muons, the triggers are MU6 at L1 and mu6 at L2 and event filter for muons with $p_T > 6$ GeV. For the $ZZ \rightarrow ee\mu\mu$ channel, events can pass either the electron or muon triggers. The trigger levels are cumulative, so that all three levels must pass for the event to be recorded.

Channel	L1	L2	EF	Total
$ZZ \rightarrow eeee$	100.00%	100.00%	100.00%	100.00%
$ZZ \rightarrow \mu\mu\mu\mu$	99.88%	99.97%	100.00%	99.85%
$ZZ \rightarrow ee\mu\mu$	100.00%	99.96%	100.00%	99.96%
$ZZ \rightarrow ee\nu\nu$	100.0%	99.92%	100.0%	99.92%
$ZZ \rightarrow \mu\mu\nu\nu$	99.15%	99.54%	99.80%	98.50%

Table 6.13: Trigger efficiency after selection cuts, relative to the previous trigger level. The total is the product of the efficiency at all three levels.

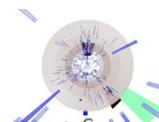
Table 6.13 shows the expected trigger efficiency after offline selection cuts for each trigger level. The efficiencies are calculated with respect to the previous level, so for example, the L2 efficiency is the number of events passing L1 and L2, divided by the number passing L1. In general, the trigger efficiency is very high for the $ZZ \rightarrow lll$ channels with typical values of $\sim 99.9\%$. In the case of $ZZ \rightarrow ll\nu\bar{\nu}$, the electron trigger efficiency is also high, while for the muons is slightly lower, but still $\sim 99\%$ overall. The overall trigger efficiencies are included in the final yields quoted in section 6.9.

6.8 Systematic Errors

The sources of systematic errors on the predicted number of events can be broadly divided into three categories, namely uncertainties in the theoretical cross section, luminosity measurements and detector reconstruction, as described below.

6.8.1 Monte Carlo Uncertainties

As mentioned in section 1.6, the uncertainty due to generator variations is expected to be 4%, from PDF and scale uncertainties of 3% and 2% respectively. An additional uncertainty of 6% is assigned to the $ZZ \rightarrow lll$ predictions from k -factor uncertainties as given in section 6.9.



6.8.2 Luminosity Uncertainties

At the time of writing, luminosity measurements were calibrated with 7 TeV collisions using a series of Van Der Meer scans [103]. The total systematic error on the luminosity was estimated to be 11%, of which the majority comes from uncertainty on the beam current measurements [104]. This value is expected to improve with operational experience, but will be used here as a conservative estimate.

6.8.3 Detector Uncertainties

Detector uncertainties will be derived from those used in references [105] and [106], where sources of uncertainty are assessed for cross section measurements in the $W \rightarrow l\nu$ and $Z \rightarrow ll$ channels respectively. The principle contributions are reproduced below:

- Reconstruction uncertainties were evaluated in the electron channel by shifting calorimeter shape variables, resulting in a contribution of $\pm 10\%$. In the muon channel, efficiency dependence on p_T and contamination from K/π yields an error of $\pm 7\%$.
- The uncertainty in trigger efficiency is expected to be negligible in the electron channel. In the muon channel, an uncertainty of $\pm 2\%$ is derived by varying the matching tolerance between tracks and trigger signals.
- Extra material in front of the electromagnetic calorimeters can have an effect on both electron reconstruction efficiency and the medium electron identification. An uncertainty of $\pm 8\%$ has been calculated in the electron channel using dedicated Monte Carlo samples.
- Problematic regions in the liquid argon calorimeter also affect the electron reconstruction efficiency, with an uncertainty of $\pm 4\%$.
- Pile-up occurs when more than one proton-proton interaction is present in a single bunch-crossing. This is evaluated to have an impact of $\pm 2\%$ on the medium electron identification efficiency.
- Energy scale and resolution is dominated in the electron channel by uncertainty in the electromagnetic calorimeter energy scale, with an uncertainty contribution of $\pm 2\%$. The corresponding muon uncertainty arises from the momentum scale, evaluated as $\pm 1\%$.

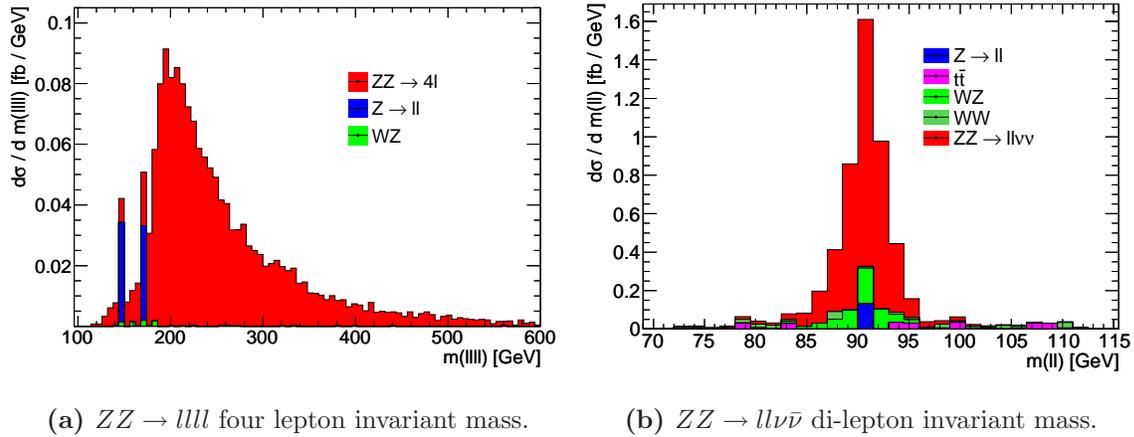


Figure 6.8: Distributions showing the predicted number of signal and background events passing all cuts at 7 TeV.

- The expected uncertainty of the missing transverse energy scale and resolution has been evaluated considering factors such as the energy scale of the clustering algorithms, problematic calorimeter regions and pile-up effects. The contribution to the uncertainty is estimated to be $\pm 2\%$.

These calculations yield a total systematic error of $\pm 14\%$ and $\pm 7\%$ for the $Z \rightarrow ee$ and $Z \rightarrow \mu\mu$ channels respectively (not including \cancel{E}_T). The total systematic error for each diboson channel is estimated by summing in quadrature the errors for each type of boson present in the event. This results in uncertainties of $\pm 19\%$, $\pm 10\%$ and $\pm 16\%$ for the $ZZ \rightarrow 4e$, $ZZ \rightarrow 4\mu$ and $ZZ \rightarrow 2e2\mu$ channels respectively. In the channels containing neutrinos, the single boson uncertainty is combined with a \cancel{E}_T contribution of $\pm 2\%$, resulting in overall uncertainties of $\pm 14\%$ and $\pm 8\%$ in the $ZZ \rightarrow 2e2\nu$ and $ZZ \rightarrow 2\mu2\nu$ channels respectively. A summary of the expected statistical and systematic uncertainties is given in table 6.15.

6.9 Summary of Expected Yields

This section summarises the event selections from the $ZZ \rightarrow llll$ and $ZZ \rightarrow ll\nu\bar{\nu}$ channels, presenting the final expected yields of signal and background events.

Figure 6.8a shows the predicted differential cross section against the invariant mass of the four reconstructed leptons, for signal and background events passing cuts in the



$ZZ \rightarrow lll$ channel. Figure 6.8b shows the equivalent plot for the $ZZ \rightarrow ll\nu\bar{\nu}$ channel, against the di-lepton invariant mass.

Expected Yields

Table 6.14 gives a summary of the expected signal and background yields in the combined $ZZ \rightarrow lll$ and $ZZ \rightarrow ll\nu\bar{\nu}$ channels after 1 fb^{-1} of ATLAS data-taking. The numbers expected for $\sqrt{s} = 14 \text{ TeV}$ are taken from [44], which used a similar cut based analysis with an earlier version of Athena. The 14 TeV study was performed with different version of the reconstruction software, and so the yields are not directly comparable, but are included in the table for completeness. For the remaining energies, numbers are taken from this study. The errors shown are estimates of the statistical uncertainty from Monte Carlo events passing cuts. In the case of the backgrounds, errors are the combination of one σ binomial errors for channels with sufficient statistics passing cuts ($n > 10$), with the 68.27% Poisson confidence limits for low statistics channels.

Channel	\sqrt{s} / TeV	7	7 (Re-weighted)	10	14
$ZZ \rightarrow lll$	Signal	10.1 ± 0.1	10.8 ± 0.08	16.3 ± 0.1	17.0 ± 0.5
	Background	$0.5^{+0.9}_{-0.2}$	$0.05^{+0.9}_{-0.02}$	$0.09^{+1}_{-0.03}$	2.0 ± 0.2
	$S/\sqrt{S+B}$	$3.1^{+0.05}_{-0.1}$	$3.3^{+0.02}_{-0.1}$	$4.0^{+0.02}_{-0.2}$	3.9 ± 0.1
$ZZ \rightarrow ll\nu\bar{\nu}$	Signal	6.2 ± 0.1	5.8 ± 0.3	10.1 ± 0.5	10.2 ± 0.2
	Background	$1.9^{+2.0}_{-0.2}$	$1.5^{+2.0}_{-0.2}$	$3.6^{+0.9}_{-0.5}$	5.2 ± 2.6
	$S/\sqrt{S+B}$	$2.2^{+0.05}_{-0.3}$	$2.2^{+0.09}_{-0.3}$	$2.7^{+0.1}_{-0.3}$	2.6 ± 0.2
Total	Signal	16.3 ± 0.1	16.6 ± 0.3	26.4 ± 0.5	27.2 ± 0.5
	Background	$2.4^{+2.3}_{-0.3}$	$1.5^{+2.3}_{-0.2}$	$3.7^{+3.2}_{-0.5}$	7.2 ± 2.6
	$S/\sqrt{S+B}$	$3.8^{+0.05}_{-0.2}$	$3.9^{+0.06}_{-0.3}$	$4.8^{+0.09}_{-0.3}$	4.6 ± 0.2

Table 6.14: Expected yields of signal and background after 1 fb^{-1} . The uncertainty on the signal yields are calculated assuming binomial errors on Monte Carlo statistical uncertainties. The errors quoted for the background yields and significance calculations represent the 68.27% confidence limits as described in sections 6.5.3 and 6.6.3.

In general, the predictions obtained using the 7 TeV and re-weighted samples are consistent, as most predictions lie within the stated errors. The $ZZ \rightarrow lll$ background is found to be larger for the 7 TeV samples, but is dominated by just three single Z events, and is still consistent within the stated uncertainties.

Channel	Events	Stat	MC	Syst	Lumi	Total Error
$ZZ \rightarrow 4e$	2.0	0.04	0.1	0.4	0.2	0.5
$ZZ \rightarrow 4\mu$	3.2	0.06	0.2	0.3	0.3	0.5
$ZZ \rightarrow 2e2\mu$	4.9	0.07	0.4	0.8	0.5	1.0
$ZZ \rightarrow lll$	10.1	0.1	0.5	0.9	0.6	1.2
$ZZ \rightarrow 2e2\nu$	3.0	0.09	0.1	0.4	0.3	0.5
$ZZ \rightarrow 2\mu2\nu$	3.3	0.09	0.1	0.2	0.4	0.5
$ZZ \rightarrow ll\nu\bar{\nu}$	6.3	0.1	0.1	0.5	0.5	0.7
Total	16.3	0.16	0.5	1.0	0.8	1.4

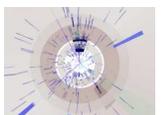
Table 6.15: The expected number of signal events with corresponding statistical and systematic errors for 7 TeV collisions and an integrated luminosity of 1 fb^{-1} .

For all three energies, the $ZZ \rightarrow ll\nu\bar{\nu}$ channel yields $\sim 60\%$ fewer events when compared to the $ZZ \rightarrow lll$ case. This is because $ZZ \rightarrow lll$ is a much cleaner channel, and so requires looser cuts to reduce background, with higher selection efficiency (24% compared to 4%).

The quantity $\frac{S}{\sqrt{S+B}}$ will be used here as an estimate of the significance of the signal, with values shown in the right-hand column of table 6.14. In this case, the significance is 3.0 for $ZZ \rightarrow lll$ and 2.2 for $ZZ \rightarrow ll\nu\bar{\nu}$ at 7 TeV, rising to 3.7 when considering the two channels combined. In general, a significance of at least five is considered necessary for a discovery, which would require a total integrated luminosity of 1.9 fb^{-1} of data.

Table 6.15 summarises the expected number of signal events and errors at $\sqrt{s} = 7 \text{ TeV}$ for 1 fb^{-1} of integrated luminosity. The errors are broken down into the uncertainty from Monte Carlo statistics, and the three categories of systematic errors as discussed in section 6.8.

The predicted yields given in this chapter do not include the contribution from gluon-gluon fusion. As mentioned in section 1.6.1, this process is not included in either the PYTHIA or MC@NLO generators, and so no ATLAS simulated events were available. From equation 1.40, the expected yields presented here could be between 9–19% higher due to contributions from gluon-gluon fusion. The effect of gluon-gluon fusion on the expected anomalous coupling sensitivity will be discussed in section 7.8.



6.10 Conclusions

Applying the cut-based event selection outlined in this chapter, it is expected that ATLAS will be able to observe 10 ± 1 events in the $ZZ \rightarrow lll$ channel for 1 fb^{-1} at 7 TeV centre of mass energy. The number of expected background events is $0.5_{-0.2}^{+0.9}$, with a large contribution from $Z \rightarrow ll$. The corresponding number of events in the $ZZ \rightarrow ll\nu\bar{\nu}$ channel is expected to be 6.2 ± 0.7 , with a background of $1.9_{-0.2}^{+2.0}$ events. With these yields, the significance of the signal will reach 5σ after an integrated luminosity of 1.9 fb^{-1} . The predictions at 10 TeV were found to be higher than those at 7 TeV by a factor of 1.6.

The results of this chapter, in particular the expected yields and selection efficiencies, will be used to calculate the expected sensitivity of the ATLAS experiment to Anomalous Triple Gauge Couplings in chapter 7.

Chapter 7

Anomalous Triple Gauge Coupling Sensitivity

“I’m sciencing as fast as I can!”

— Professor Farnsworth

7.1 Introduction

Measurement of the $pp \rightarrow ZZ$ differential cross section can be used to set limits on the anomalous NTGCs introduced in section 1.3. As these couplings are zero at tree level in the SM, their measurement provides a sensitive test of electro-weak theory. Non-zero measured values would be direct evidence for new physics beyond the SM.

This chapter presents the expected anomalous coupling sensitivity in ATLAS at an energy of $\sqrt{s} = 7$ TeV, for integrated luminosities of 0.01–10 fb⁻¹. Confidence limits on the anomalous coupling parameters, $f_{i=4,5}^{V=Z,\gamma}$, are obtained using extended unbinned maximum likelihood fits to the $Z p_T$ spectrum. Fits are made with both the $ZZ \rightarrow lll$ and $ZZ \rightarrow ll\nu\bar{\nu}$ channels, using results obtained in chapter 6.

7.2 Anomalous Coupling Monte Carlo

Non-zero values of the anomalous coupling parameters tend to increase the ZZ cross section at high Z boson p_T . Figure 7.1a shows the differential cross section $\frac{d\sigma}{dp_T}$ at $\sqrt{s} = 7$ TeV, for the Z boson transverse momentum, where the Z is chosen at random from the two in each event. Cuts are imposed requiring that $m_{ll} > 60$ GeV and $m_{llll} > 110$ GeV.

The predicted SM cross sections show a good agreement between the PYTHIA, SHERPA and BR generators for SM events. In addition, the $p_T(Z)$ spectrum is also shown for SHERPA and BR samples with $f_4^Z = 0.02$. A form factor is applied as in equation 1.29, with $n = 3$ and $\Lambda_{FF} = 2$ TeV.

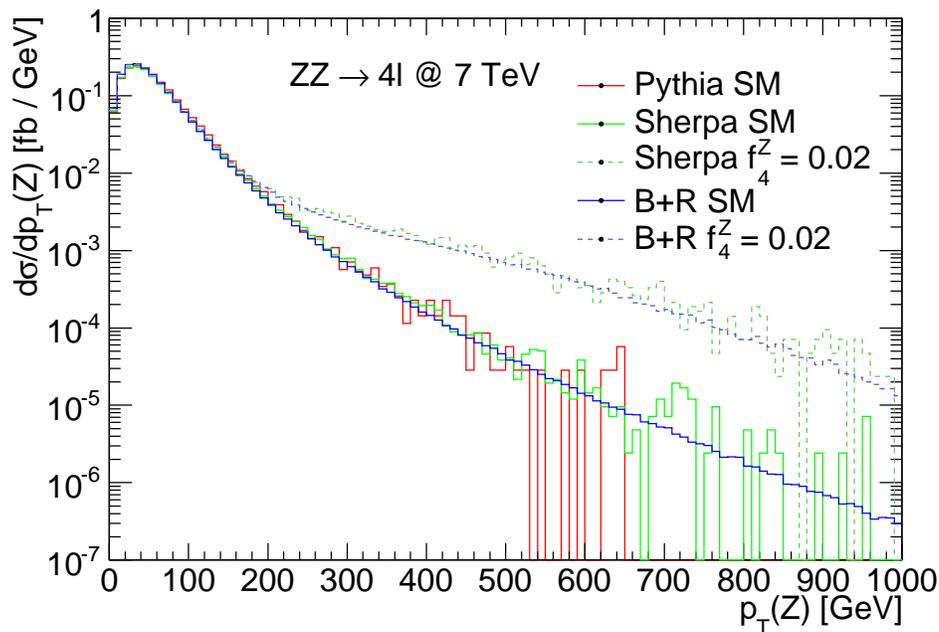
The program by BR, introduced in section 1.5.5, was chosen to generate the majority of Monte Carlo samples with anomalous coupling vertices. BR was chosen over SHERPA for ease of implementation. Unless otherwise stated, Monte Carlo events have been generated using the CTEQ6LL [12] parton density functions. The uncertainties due to the choice of generator and PDF on the anomalous coupling sensitivity is discussed in section 7.7.

7.2.1 Form Factors

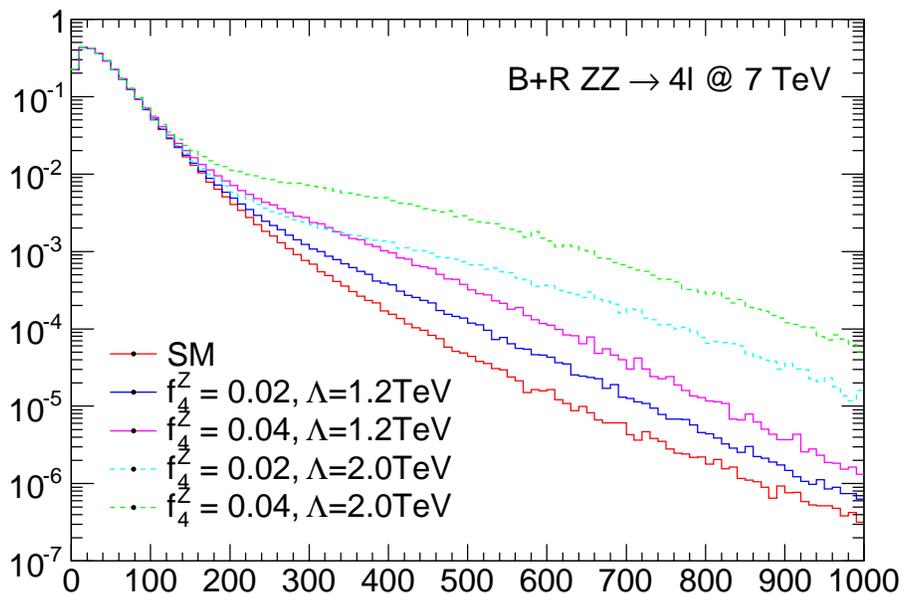
A form factor is applied for all events generated with anomalous coupling vertices, using the form shown in equation 1.29. As discussed in section 1.3.1, a form factor is necessary to avoid violation of unitarity at high energies.

Two separate cutoff schemes are investigated, both taking $n = 3$. The first applies a cutoff of $\Lambda_{FF} = 1.2$ TeV to allow direct comparison with the current experimental limits measured by the Tevatron experiments [40, 46]. A second value of $\Lambda_{FF} = 2$ TeV has been chosen to match the cutoff used in the high-luminosity ATLAS study at $\sqrt{s} = 14$ TeV [44]. The effect of this choice on $\frac{d\sigma}{dp_T}$ is highlighted in figure 7.1b, which shows different shapes for the two cutoff values.

It should be noted that the coupling limits measured in this chapter are the bare couplings, f_{i0}^V in equation 1.29. For clarity, the additional subscript will be dropped in subsequent notation.

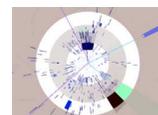


(a) The Leading Order $Z p_T$ spectrum for $ZZ \rightarrow llll$ events ($l = e, \mu$) for the SM and in the presence of anomalous couplings for the SHERPA and Baur and Rainwater event generators.



(b) Comparison of the $Z p_T$ spectrum with two choices of form factor cutoff, using the Baur and Rainwater generator.

Figure 7.1: Examples of $Z p_T$ spectra including anomalous couplings.



7.3 Extended Likelihood Function

Confidence limits on the values of the anomalous coupling parameters, $f_{i=4,5}^{V=Z,\gamma}$, are measured by constructing an extended, unbinned likelihood function,

$$\mathcal{L}(f_i^V | \{p_T^1, \dots, p_T^n\}) = \underbrace{\frac{e^{-\mu(f_i^V)} \mu(f_i^V)^n}{n!}}_{\text{Extended term}} \times \underbrace{\prod_{k=1}^{k=n} P(f_i^V | p_T^k)}_{\text{Unbinned likelihood}}, \quad (7.1)$$

with the terms described as follows:

$\{p_T^1, \dots, p_T^n\}$ is the set of measurements of the transverse momentum of the Z boson, where n is the number of observed events passing the cuts described in chapter 6. The events are divided into two sets: those satisfying the $ZZ \rightarrow lll$ selection criteria, and those satisfying $ZZ \rightarrow ll\nu\bar{\nu}$, with a separate likelihood function for each. For each $ZZ \rightarrow ll\nu\bar{\nu}$ candidate event, the p_T of the reconstructed $Z \rightarrow ll$ boson is used. In the case of $ZZ \rightarrow lll$ events, the p_T of a randomly selected Z is used. This is done to allow a more direct comparison between limits obtained from the two channels, and to match the method used in the 14 TeV ATLAS study [44].

$\frac{e^{-\mu(f_i^V)} \mu(f_i^V)^n}{n!}$ is the extended term, and represents the probability of measuring n events from a Poisson distribution with a mean of $\mu(f_i^V)$.

$\prod_{k=1}^{k=n} P(f_i^V | p_T^k)$ is the unbinned likelihood term, where $P(f_i^V | p_T^k)$ is the probability that a transverse momentum measurement, p_T^k , has been drawn from a probability distribution function (PDF¹) corresponding to an anomalous coupling of f_i^V .

The quantities $\mu(f_i^V)$ and $P(f_i^V | p_T^k)$ are calculated using Monte Carlo methods described in section 7.2. Likelihood functions are constructed using p_T measurements from the $ZZ \rightarrow lll$ and $ZZ \rightarrow ll\nu\bar{\nu}$ channel individually. A combined likelihood function is also constructed as the product of individual likelihoods from the two channels:

$$\mathcal{L}_{\text{Combined}} = \mathcal{L}_{ZZ \rightarrow lll} \times \mathcal{L}_{ZZ \rightarrow ll\nu\bar{\nu}}. \quad (7.2)$$

¹Not to be confused with Parton Density Function!

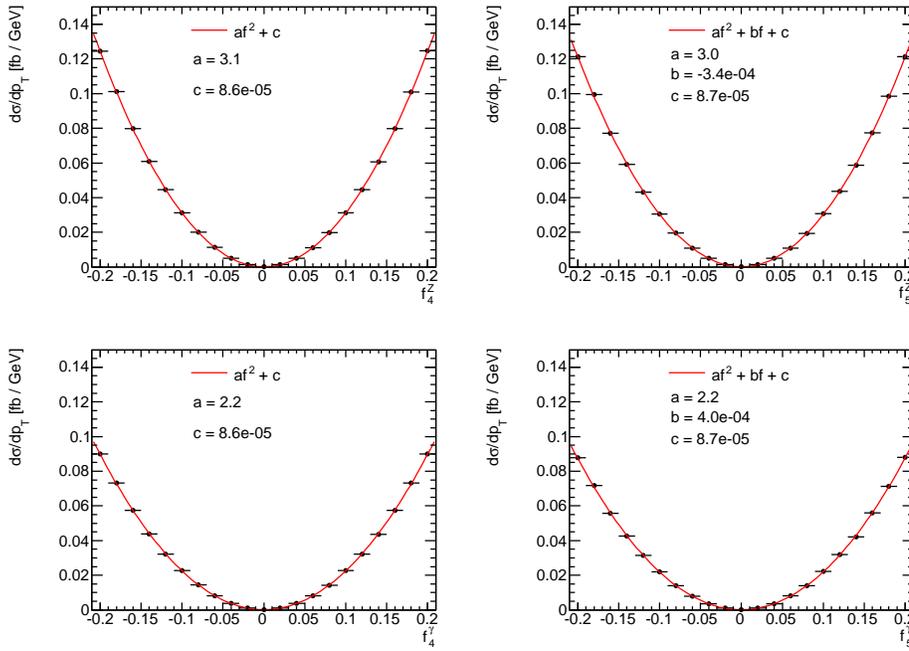


Figure 7.2: The cross section against anomalous coupling for events with $500 \text{ GeV} < p_T^Z < 510 \text{ GeV}$. A form factor with $n = 3$ and $\Lambda_{FF} = 2 \text{ TeV}$ has been used.

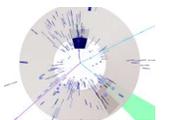
7.4 Calculation of $P(f_i^V | p_T^k)$

To construct $P(f_i^V | p_T^k)$, sets of $ZZ \rightarrow lll$ and $ZZ \rightarrow ll\nu\bar{\nu}$ Monte Carlo events were generated using BR across a range of couplings, $-0.2 < f_i^V < 0.2$, for pp collisions at $\sqrt{s} = 7 \text{ TeV}$. All of the couplings are assumed to be real, and only one coupling is non-zero at any time. Cuts are applied at generator level, requiring $m_Z > 12 \text{ GeV}$ for both channels, and $m_{lll} > 110 \text{ GeV}$ for the $ZZ \rightarrow lll$ channel.

Generated events are binned by the transverse momentum of a Z in the event, using the binning in figure 7.1a. In the case of $ZZ \rightarrow ll\nu\bar{\nu}$, the p_T of the Z decaying to charged leptons is used. In the case of $ZZ \rightarrow lll$, the p_T of a randomly chosen Z is taken. Figure 7.2 shows the cross section as a function of the anomalous coupling value for events with $500 \text{ GeV} < p_T < 510 \text{ GeV}$.

Equation 1.43 shows $\sigma \propto (f_i^V)^2$, and hence a quadratic fit is used to parametrise the cross section in each p_T bin, of the form:

$$\frac{d\sigma_{\text{Total}}(f)}{dp_T} = c(p_T) + b(p_T)f + a(p_T)f^2, \quad (7.3)$$



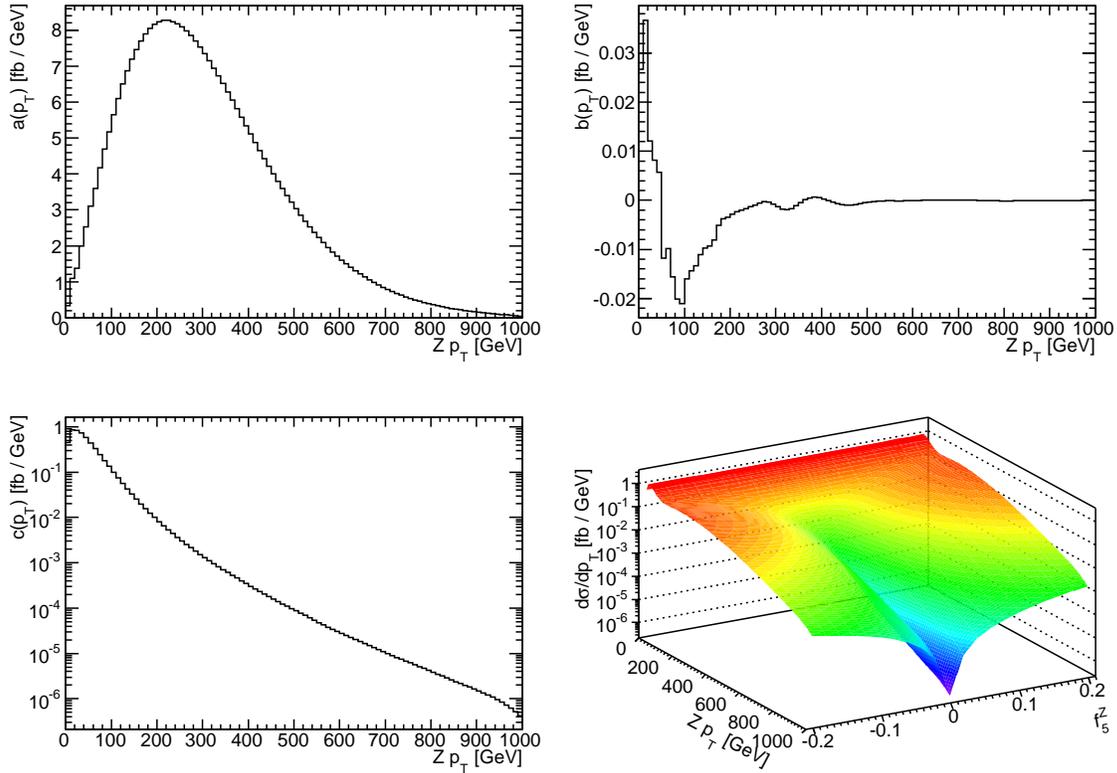


Figure 7.3: Fitted parameter values from equation 7.3 against $p_T(Z)$, using a form factor with $n = 3$ and $\Lambda_{FF} = 2$ TeV. Plots are shown for the f_5^Z couplings in the $ZZ \rightarrow lll$ channel, using the Baur and Rainwater event generator. Distributions are similar for the remaining coupling parameters as are those from the $ZZ \rightarrow ll\nu\bar{\nu}$ channel. The lower-right plot shows the final PDF, $P(f_5^Z|p_T)$.

where $f = f_i^V$, and $c = d\sigma_{\text{SM}}/dp_T$ is the SM cross section. The linear parameter, b , represents the parity violating term, and is only non-zero for the f_5^V couplings. Examples of such fitted functions for a single bin are shown in figure 7.2.

By interpolating fitted parameter values as a function of $p_T(Z)$, it is possible to evaluate the differential cross section at any value of $p_T(Z)$ and f_i^V . An example of the fitted parameters as functions of $p_T(Z)$ are shown in figure 7.3.

The generator-level probability density function, $P(f_i^V|p_T^k)$, is then constructed from equation 7.3,

$$P(f_i^V|p_T^k) = A(f_i^V) \frac{d\sigma_{\text{Total}}(f_i^V)}{dp_T}, \quad (7.4)$$

where $A(f_i^V)$ is a normalisation factor to ensure the total probability is unity. The normalisation is performed within the ROOFIT [107] framework, which is also used to generate toy Monte Carlo samples and perform the fitting described in section 7.6.1.

7.4.1 Detector Effects

The shape of $P(f_i^V|p_T^k)$ will be altered in a number of ways by detector effects, so that the PDF becomes

$$P(f_i^V|p_T^k, \epsilon, R) = [P(f_i^V|p_T^k) + B(p_T^k)] \epsilon(p_T^k) \otimes R(p_T^k), \quad (7.5)$$

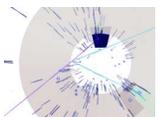
where $\epsilon(p_T^k)$, $R(p_T^k)$ and $B(p_T^k)$ are the detector selection efficiency, resolution and background contribution respectively, described below.

Selection Efficiency

The shape of $P(f_i^V|p_T^k)$ is changed by p_T dependant selection efficiencies due to detector reconstruction and background suppression cuts. The selection efficiency, $\epsilon(p_T)$, is calculated as the number of events passing all selection cuts divided by the total number of generated events, as a function of true Z p_T . The efficiencies corresponding to the event selection criteria of chapter 6 are shown in figure 7.4. The $ZZ \rightarrow lll$ efficiency is seen to increase with p_T , while the $ZZ \rightarrow ll\nu\bar{\nu}$ efficiency tends to decrease. Both efficiencies show a similar shape to the previous ATLAS study at 14 TeV [44].

Z p_T Resolution

As discussed in section 6.4.1, momentum measurements will differ from the true momentum due to finite detector resolution. The left-hand plot of figure 7.5 shows the fractional difference between the true $p_T(Z)$ and that of the closest reconstructed Z within a cone of $\Delta R < 0.1$, for $ZZ \rightarrow lll$ events passing cuts described in chapter 6. Events were divided into bins in true $p_T(Z)$, and fitted with a Gaussian. The right-hand plot of figure 7.5 shows the mean and sigma of each Gaussian as a function of true $p_T(Z)$.



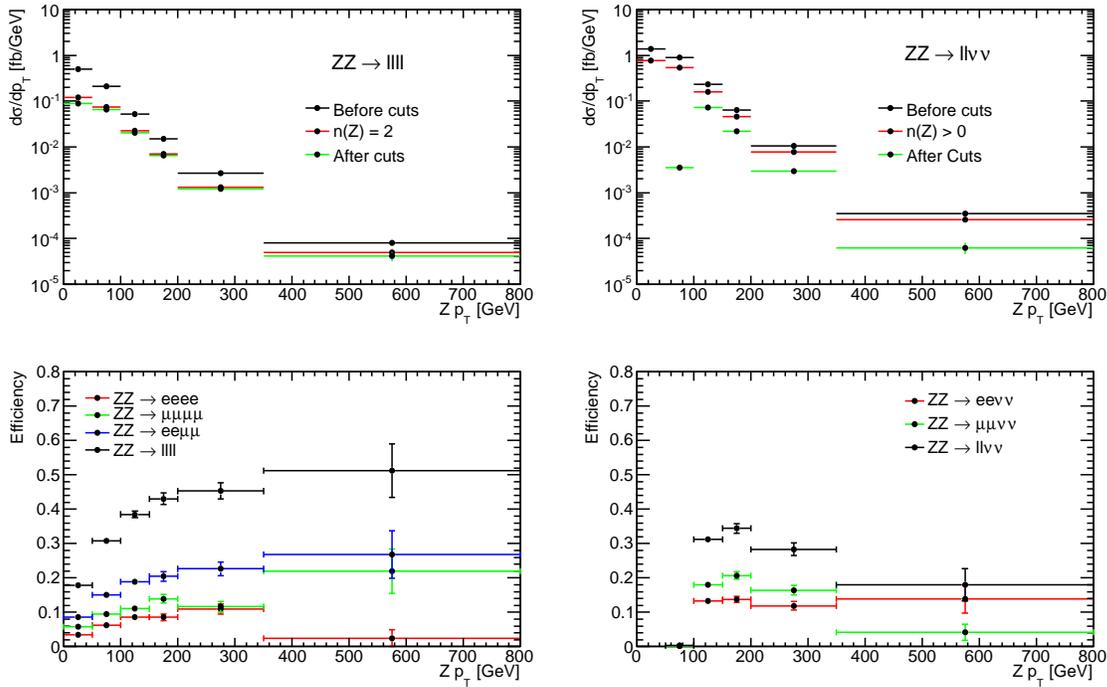


Figure 7.4: The upper plots show the differential cross section of $ZZ \rightarrow llll$ (left) and $ZZ \rightarrow ll\nu\bar{\nu}$ (right) as a function of $p_T(Z)$, before cuts, after preselection and after cuts described in chapter 6. The lower plots show the corresponding selection efficiency as a function of true $p_T(Z)$. The efficiency is always made with respect to the total number of true events (either $ZZ \rightarrow llll$ or $ZZ \rightarrow ll\nu\bar{\nu}$ ($l = e, \mu$)), and so the total efficiency can be calculated as the sum of efficiencies for each sub-channel.

To include the resolution in $P(f_i^V | p_T^k)$, each generated p_T distribution was convolved with a Gaussian, taking

$$\mu = -1\%, \sigma = 5\%. \quad (7.6)$$

The fitted parameters in the right-hand plot of figure 7.5 show no significant variation with transverse momentum, and so μ and σ will be assumed constant with $p_T(Z)$. The effect of varying these parameters will be discussed in section 7.7.

Background Contribution

As estimation of the background contribution, $B(p_T^k)$, is also included in the final value, using the predictions obtained from fully simulated Monte Carlo samples in section 6.9.

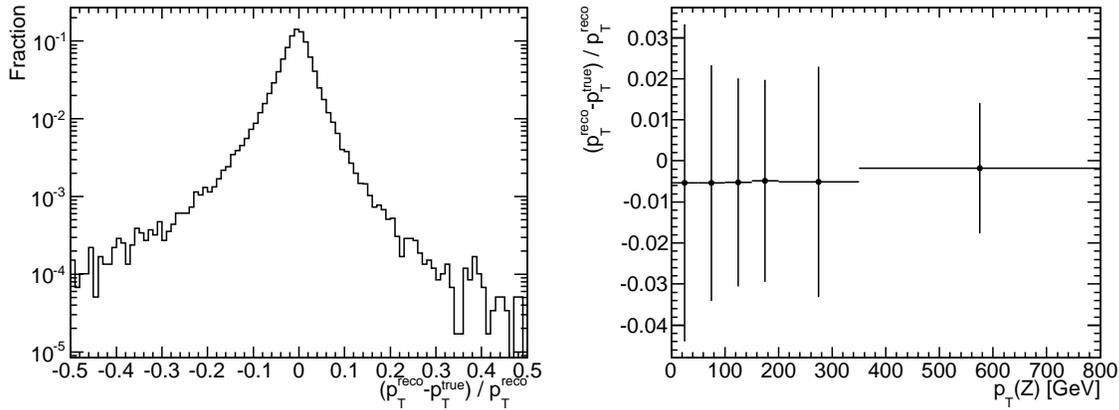
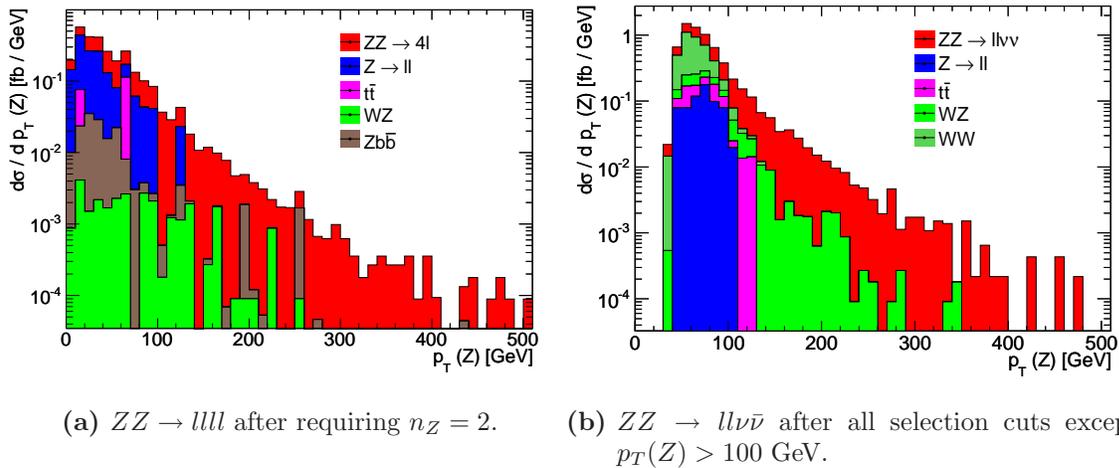


Figure 7.5: The left-hand plot shows the fractional difference between true and reconstructed Z bosons in the $ZZ \rightarrow lll$ simulated sample. The right-hand plot shows the mean (position) and width (error) of a Gaussian fit to the fractional difference as a function of $p_T(Z)$.



(a) $ZZ \rightarrow lll$ after requiring $n_Z = 2$.

(b) $ZZ \rightarrow ll\nu\bar{\nu}$ after all selection cuts except $p_T(Z) > 100$ GeV.

Figure 7.6: Stacked histograms of the reconstructed Z p_T spectrum for signal and background, weighted by the number of events expected in 1 fb^{-1} of integrated luminosity.

The signal:background ratio is 20^{+17}_{-14} for the $ZZ \rightarrow lll$ channel, compared to $3.6^{+0.8}_{-1.0}$ for $ZZ \rightarrow ll\nu\bar{\nu}$.

Figure 7.6 shows the Z p_T spectrum of fully simulated signal and background events from chapter 6. Figure 7.6a shows the p_T for the highest p_T Z in $ZZ \rightarrow lll$ events, after the cut requiring four fully reconstructed leptons. After additional cuts in this channel, the statistics become limited and the background shape is difficult to distinguish.

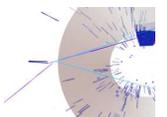


Figure 7.6b shows the differential cross section for signal and background channels as a function of $p_T(Z \rightarrow ll)$, after all cuts expect for $p_T(Z) > 100$ GeV. Many of the backgrounds display a steeply falling p_T spectrum below 100 GeV, in particular single Z and $t\bar{t}$ events. The WZ channel extends to higher p_T with a similar gradient to the signal. Importantly, none of the background distributions increase with p_T , where the shape is most sensitive to the anomalous coupling contribution.

The background contribution to the PDF will be taken as a uniform fraction of the SM signal expectation, so that

$$B(p_T^k) = f_{\text{bg}} \times P(\text{SM}|p_T^k). \quad (7.7)$$

f_{bg} is a constant, calculated for each channel using the central signal:background ratios quoted above. The effect of varying the background shape will be discussed further in section 7.7.

7.5 Calculation of $\mu(f_i^V)$

The mean expected number of events, $\mu(f_i^V)$, for each channel and for a given integrated luminosity, \mathcal{L} , is given by:

$$\mu(f_i^V) = \mathcal{L} \times \sigma_{\text{SM}}^{ZZ}(1 + f_{\text{ac}} + f_{\text{bg}}), \quad (7.8)$$

with the following definitions:

σ_{SM}^{ZZ} is the expected SM cross section after selection cuts in chapter 6, with values taken from table 6.15. The NLO cross section is used in the case of $ZZ \rightarrow ll\nu\bar{\nu}$. In the case of $ZZ \rightarrow llll$, k -factors have been used to scale the predictions to NLO.

$f_{\text{ac}} = f_{\text{ac}}(f_i^V)$ is the fraction of additional events due to anomalous couplings. It is calculated by integrating the differential cross section, including detector efficiency and resolution, over the entire p_T range as follows:

$$1 + f_{\text{ac}}(f_i^V) = \frac{\int_0^\infty \frac{d\sigma_{\text{Total}}(f_i^V)}{dp_T} \epsilon(p_T) \otimes R(p_T) dp_T}{\int_0^\infty \frac{d\sigma_{\text{Total}}(\text{SM})}{dp_T} \epsilon(p_T) \otimes R(p_T) dp_T}, \quad (7.9)$$

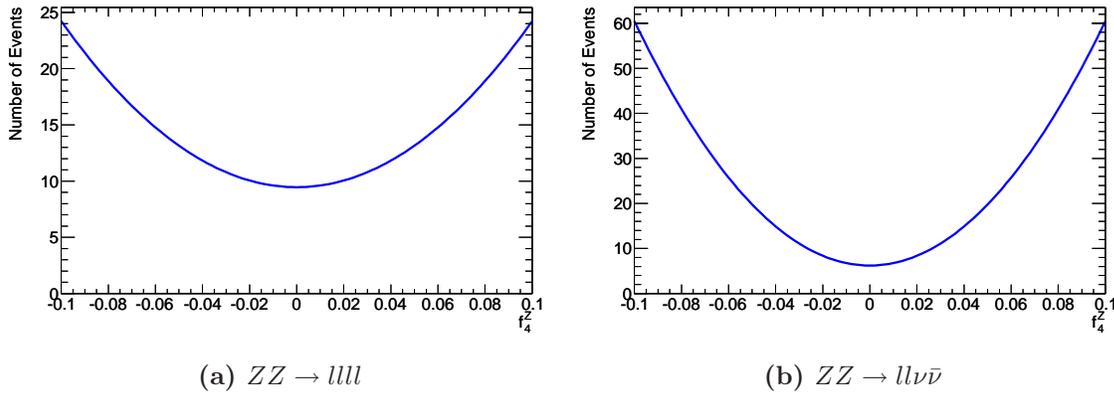


Figure 7.7: Expected total number of events after selection cuts for 1 fb^{-1} of data, as a function of the anomalous coupling parameter f_4^Z , using a form factor with $n = 3$ and $\Lambda_{FF} = 2 \text{ TeV}$.

where $\frac{d\sigma_{\text{Total}}(\text{SM})}{dp_T}$ is the SM p_T spectrum. $\epsilon(p_T^k)$ and $R(p_T^k)$ are the detector selection efficiency and resolution introduced in section 7.4.1.

f_{bg} is the fraction of events passing cuts from background channels, as defined in section 7.4.1.

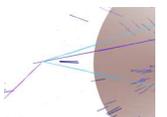
Figure 7.7 shows $\mu(f_i^V)$ as a function of f_4^Z , for $\mathcal{L} = 1 \text{ fb}^{-1}$, in the $ZZ \rightarrow lll$ and $ZZ \rightarrow ll\nu\bar{\nu}$ channels. Although fewer SM events are expected in the $ZZ \rightarrow ll\nu\bar{\nu}$ channel, it is more sensitive to non-zero coupling values due to the high Z p_T cut.

7.6 Anomalous Coupling Sensitivity

This section aims to calculate the expected anomalous coupling limits which could be obtained from early ATLAS data ($\mathcal{L} \sim 1 \text{ fb}^{-1}$) at $\sqrt{s} = 7 \text{ TeV}$. This is done by generating toy Monte Carlo samples and fitting them using the likelihood function described in section 7.3.

7.6.1 Toy Monte Carlo Samples

Toy Monte Carlo datasets are created by drawing p_T values from the probability density function, $P(f_i^V | p_T^k)$, with SM couplings, such that $f_i^V = 0$. The number of events in each dataset is generated from a Poisson distribution with mean given by equation 7.8,



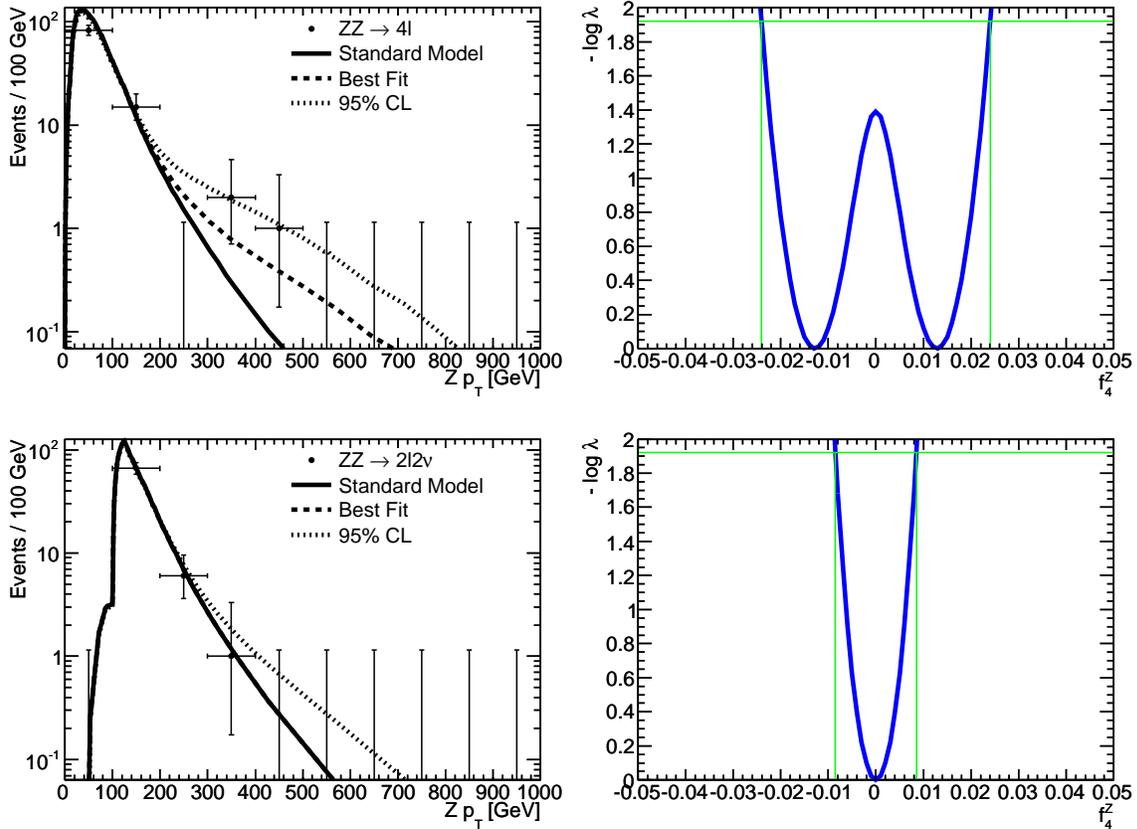


Figure 7.8: Toy Monte Carlo datasets in the $ZZ \rightarrow lll$ (top-left) and $ZZ \rightarrow ll\nu\bar{\nu}$ (bottom-left) channels with 10 fb^{-1} of data. The expected SM, best fit and 95% confidence limit shapes are also shown, using $n = 3$ and $\Lambda_{FF} = 2 \text{ TeV}$. The right-hand plots show the corresponding log-likelihood function over a range of f_4^Z values.

taking $f_{ac} = 0$. This procedure is used to generate fake datasets for both channels, with an example of a single dataset generated for both the $ZZ \rightarrow lll$ and $ZZ \rightarrow ll\nu\bar{\nu}$ channels is shown in figure 7.8.

7.6.2 Fitting Method

The best fitting anomalous coupling value is calculated using the extended maximum log-likelihood method, using numerical routines to minimise $-\log \mathcal{L}(f_i^V | p_T^k)$. The 95% Confidence Limits (CLs) on f_i^V are calculated by finding the values at which the negative log-likelihood is 1.92 above the minimum. Some fits contain two minima separated by a central maximum which is more than 1.92 above the minimum. In these cases, the two most extreme values are taken at which the log-likelihood rises to 1.92 above the

minimum. Fits are made assuming all couplings are real, and only one coupling is non-zero at any one time.

An example of a single toy Monte Carlo dataset are shown in figure 7.8, where fits have been made to the f_4^Z coupling parameter after an integrated luminosity of 10 fb^{-1} . This integrated luminosity has been chosen for illustrative purposes, as more events are observed at high p_T . The shapes of $P(f_i^V|p_T^k)$ corresponding to the best fitting coupling, the 95% CL, and the SM, are also included in the plots. The negative log-likelihood distributions as a function of f_4^Z are shown on right-hand side of the figure 7.8. Vertical lines indicate the couplings corresponding to the 95% confidence limits.

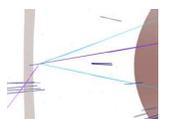
In the majority of cases, the negative log-likelihood function has a single minimum as exemplified in the lower right-hand plot of figure 7.8. On some occasions, statistical fluctuations lead to an excess of high p_T events. In this case the best fit corresponds to a non-zero coupling, with two minima in the negative log-likelihood function. An example of such a function can be seen in the upper right-hand plot of figure 7.8. In the case of the f_4^V couplings, the minima will be symmetric about $f_4^V = 0$, whereas for f_5^V there will be a small asymmetry due to the linear term in equation 7.3.

7.6.3 Results

The upper (lower) expected confidence limits are calculated by taking the mean 95% upper (lower) CL for 1000 toy Monte Carlo datasets. Limits are obtained for each of the four coupling parameters, using fits from each channel separately and an overall combined fit obtained using the likelihood function shown in equation 7.2.

The distributions of the best fitting f_4^Z coupling and its confidence limits are shown in figure 7.9, for 1 fb^{-1} of integrated luminosity. A small number of toy datasets fail to give good fit values, but these are normally $< 1\%$ depending on the luminosity. The best fit distributions show a large peak at the SM value of $f_4^Z = 0$, indicating the fits are performing as expected.

The best fit distributions also display two symmetric “shoulders” around zero, which correspond to models in which high- p_T Z bosons are generated from the SM PDF, such as the $ZZ \rightarrow lll$ example from figure 7.8. This is confirmed in figure 7.10, which shows the best fitting f_4^Z value against the largest measured p_T in the $ZZ \rightarrow lll$ sample. The plot shows only SM fits for $p_T < 150 \text{ GeV}$, above which the distribution bifurcates as



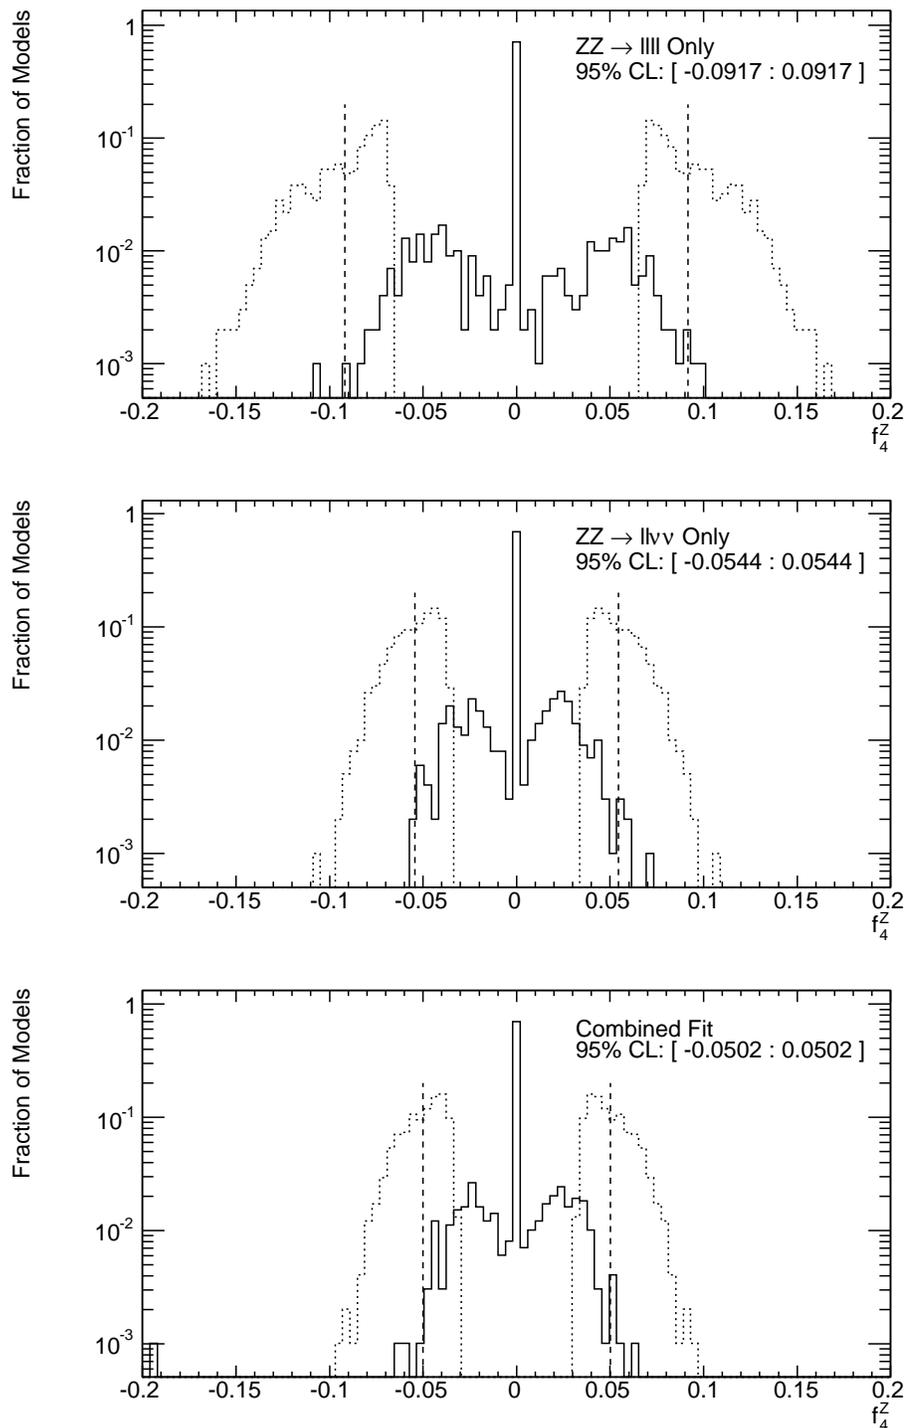


Figure 7.9: The best fit (solid) and 95% confidence limit (dotted) distributions for $ZZ \rightarrow llll$ (top), $ZZ \rightarrow ll\nu\bar{\nu}$ (middle) and combined (bottom) fits to f_4^Z after 1 fb^{-1} of integrated luminosity, using 1000 toy Monte Carlo datasets. Dashed vertical lines indicate the mean limit values. A form factor with $n = 3$ and $\Lambda_{FF} = 1.2 \text{ TeV}$ is used.

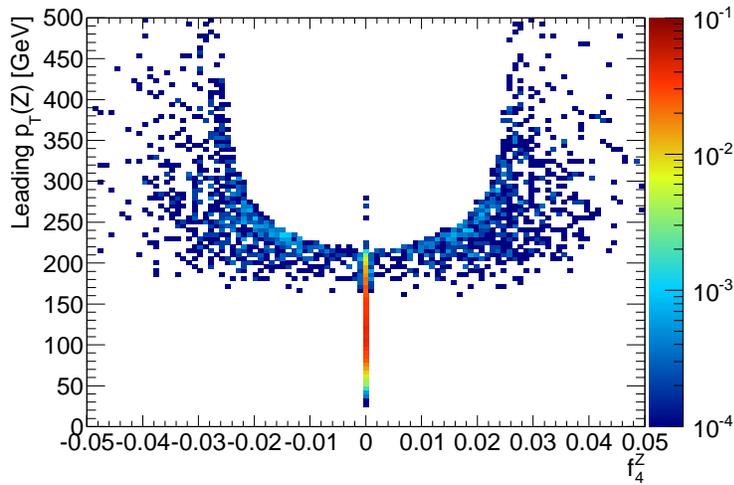


Figure 7.10: A histogram of the highest p_T Z boson against the best fitting f_4^Z coupling in the $ZZ \rightarrow lll$ channel, for 1000 toy Monte Carlo models for 1.0 fb^{-1} .

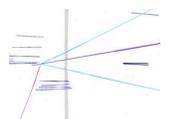
the log-likelihood function splits into two minima. If the highest Z has $p_T > 250$, then a SM fit becomes unobtainable.

The remaining three coupling parameters show similarly shaped distributions to the f_4^Z case. One exception is that the f_5^V distributions display a slight asymmetry in the best fit distribution due to interference with the SM ZZ production diagrams.

Distributions of measured coupling limits for other integrated luminosity values show similar shapes to those in figure 7.9. At low luminosity, where no events are expected to be observed, the limits show narrower peaks, as the only information available is the number of events. For example, with 0.01 fb^{-1} of data, the standard deviation of the measured limits is 15% compared to the mean limit, compared to 25% at 1 fb^{-1} .

Tests were also performed with non-zero anomalous couplings as the input parameter. Figure 7.11 shows the mean and standard deviation of the best fitting f_4^Z coupling across a range of input values. The cross shape arises due to the symmetry of the likelihood function, which will have two minima at $\pm f_4^V$. (The f_5^V likelihood functions will be slightly asymmetric due to the linear term in equation 7.3). As the starting value for the input is chosen randomly, there is an equal probability of obtaining a positive and negative best fit.

Tables 7.1 and 7.2 show the expected 95% confidence limits on the anomalous coupling parameters for integrated luminosities of $0.01\text{-}10 \text{ fb}^{-1}$, using $\Lambda_{FF} = 1.2, 2.0 \text{ TeV}$ respectively. Both sets of results use $n = 3$. Limits are quoted using fits to each channel



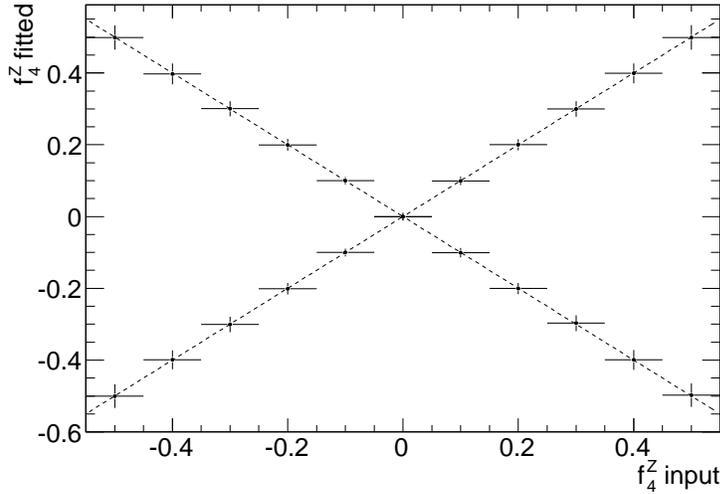


Figure 7.11: The best fitting f_4^Z coupling against the input value. Each point represents the mean of 1000 fits, with the errors bars equal to the standard deviation of the distribution. The dashed lines represent $y = \pm x$ for guidance.

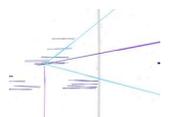
separately and from a combined fit using both channels. The combined limits as a function of integrated luminosity are also presented in figure 7.12, which displays the 68% and 90% confidence limits in addition to 95%, in the case where $\Lambda = 1.2$ TeV.

In general, the $ZZ \rightarrow ll\nu\bar{\nu}$ channel appears to be more sensitive to anomalous couplings than $ZZ \rightarrow llll$. At first this may seem counter-intuitive, as fewer events are predicted to pass the selection cuts and the background fraction is much higher. However, the $ZZ \rightarrow ll\nu\bar{\nu}$ channel contains more events in the kinematic region which is most sensitive to anomalous couplings, thanks to the $p_T(Z) > 100$ GeV cut used to select events in section 6.6. Retaining the low p_T events in the $ZZ \rightarrow llll$ channel is still important as a way of constraining the SM cross section. In the $ZZ \rightarrow ll\nu\bar{\nu}$ channel this would be impractical due to higher background rates below $p_T(Z) < 100$ GeV from WW production.

Overall, tighter limits are predicted for the f_i^Z couplings compared to f_i^γ by 20%, and f_4^V show slightly tighter limits than f_5^V . For 0.01 fb^{-1} of integrated luminosity, only 0.2 events are expected in total, with the majority of toy models observing no events. At these low luminosities, the limits are obtained using only the Poisson term in equation 7.1. Between $0.1\text{--}1.0 \text{ fb}^{-1}$, the expected number of events rises above one, and so p_T information is also typically available for the fit.

$\mathcal{L}/\text{fb}^{-1}$	Coupling	$ZZ \rightarrow lll$	$ZZ \rightarrow ll\nu\bar{\nu}$	Combined	Systematic
0.01	f_4^Z	[-0.688 : 0.688]	[-0.353 : 0.353]	[-0.319 : 0.319]	[-0.328 : 0.329]
	f_5^Z	[-0.724 : 0.726]	[-0.375 : 0.376]	[-0.338 : 0.339]	[-0.347 : 0.348]
	f_4^γ	[-0.807 : 0.807]	[-0.420 : 0.421]	[-0.380 : 0.380]	[-0.396 : 0.396]
	f_5^γ	[-0.860 : 0.856]	[-0.446 : 0.442]	[-0.404 : 0.400]	[-0.421 : 0.417]
0.1	f_4^Z	[-0.237 : 0.237]	[-0.129 : 0.129]	[-0.118 : 0.118]	[-0.126 : 0.126]
	f_5^Z	[-0.251 : 0.256]	[-0.135 : 0.136]	[-0.123 : 0.124]	[-0.133 : 0.134]
	f_4^γ	[-0.284 : 0.284]	[-0.155 : 0.155]	[-0.142 : 0.142]	[-0.149 : 0.149]
	f_5^γ	[-0.303 : 0.299]	[-0.165 : 0.162]	[-0.151 : 0.148]	[-0.159 : 0.156]
1.0	f_4^Z	[-0.092 : 0.092]	[-0.054 : 0.054]	[-0.050 : 0.050]	[-0.059 : 0.059]
	f_5^Z	[-0.095 : 0.095]	[-0.057 : 0.058]	[-0.053 : 0.053]	[-0.060 : 0.060]
	f_4^γ	[-0.108 : 0.108]	[-0.065 : 0.065]	[-0.059 : 0.059]	[-0.070 : 0.070]
	f_5^γ	[-0.114 : 0.112]	[-0.069 : 0.066]	[-0.063 : 0.061]	[-0.074 : 0.071]
10.0	f_4^Z	[-0.039 : 0.039]	[-0.026 : 0.026]	[-0.024 : 0.024]	[-0.034 : 0.034]
	f_5^Z	[-0.040 : 0.041]	[-0.027 : 0.027]	[-0.025 : 0.025]	[-0.034 : 0.035]
	f_4^γ	[-0.047 : 0.047]	[-0.031 : 0.031]	[-0.029 : 0.029]	[-0.040 : 0.040]
	f_5^γ	[-0.049 : 0.048]	[-0.033 : 0.031]	[-0.031 : 0.029]	[-0.043 : 0.041]

Table 7.1: The expected 95% confidence limits, with $\Lambda_{FF} = 1.2$ TeV. Results highlighted in red violate the unitarity requirements in equations 1.30 and 1.31, while those in blue are tighter than current measured limits by the CDF collaboration [40].



$\mathcal{L}/\text{fb}^{-1}$	Coupling	$ZZ \rightarrow lll$	$ZZ \rightarrow ll\nu\bar{\nu}$	Combined	Systematic
0.01	f_4^Z	[-0.379 : 0.379]	[-0.197 : 0.197]	[-0.177 : 0.177]	[-0.182 : 0.182]
	f_5^Z	[-0.396 : 0.396]	[-0.202 : 0.202]	[-0.182 : 0.183]	[-0.187 : 0.188]
	f_4^γ	[-0.451 : 0.451]	[-0.233 : 0.233]	[-0.210 : 0.210]	[-0.218 : 0.218]
	f_5^γ	[-0.468 : 0.467]	[-0.243 : 0.242]	[-0.218 : 0.217]	[-0.227 : 0.226]
0.1	f_4^Z	[-0.128 : 0.128]	[-0.070 : 0.070]	[-0.064 : 0.064]	[-0.068 : 0.068]
	f_5^Z	[-0.132 : 0.132]	[-0.072 : 0.073]	[-0.065 : 0.066]	[-0.070 : 0.071]
	f_4^γ	[-0.151 : 0.151]	[-0.083 : 0.083]	[-0.075 : 0.075]	[-0.079 : 0.079]
	f_5^γ	[-0.157 : 0.156]	[-0.086 : 0.085]	[-0.078 : 0.077]	[-0.082 : 0.081]
1.0	f_4^Z	[-0.046 : 0.046]	[-0.027 : 0.027]	[-0.025 : 0.025]	[-0.029 : 0.029]
	f_5^Z	[-0.047 : 0.047]	[-0.028 : 0.028]	[-0.025 : 0.025]	[-0.029 : 0.029]
	f_4^γ	[-0.055 : 0.055]	[-0.032 : 0.032]	[-0.029 : 0.029]	[-0.035 : 0.035]
	f_5^γ	[-0.056 : 0.055]	[-0.034 : 0.033]	[-0.031 : 0.030]	[-0.036 : 0.035]
10.0	f_4^Z	[-0.018 : 0.018]	[-0.012 : 0.012]	[-0.011 : 0.011]	[-0.015 : 0.015]
	f_5^Z	[-0.019 : 0.019]	[-0.012 : 0.012]	[-0.011 : 0.011]	[-0.015 : 0.016]
	f_4^γ	[-0.022 : 0.022]	[-0.014 : 0.014]	[-0.013 : 0.013]	[-0.018 : 0.018]
	f_5^γ	[-0.023 : 0.022]	[-0.015 : 0.014]	[-0.013 : 0.013]	[-0.019 : 0.018]

Table 7.2: The expected 95% confidence limits, with $\Lambda_{FF} = 2.0$ TeV. Results highlighted in red violate the unitarity requirements in equations 1.30 and 1.31, while those in blue are tighter than current measured limits by the CDF collaboration [40].

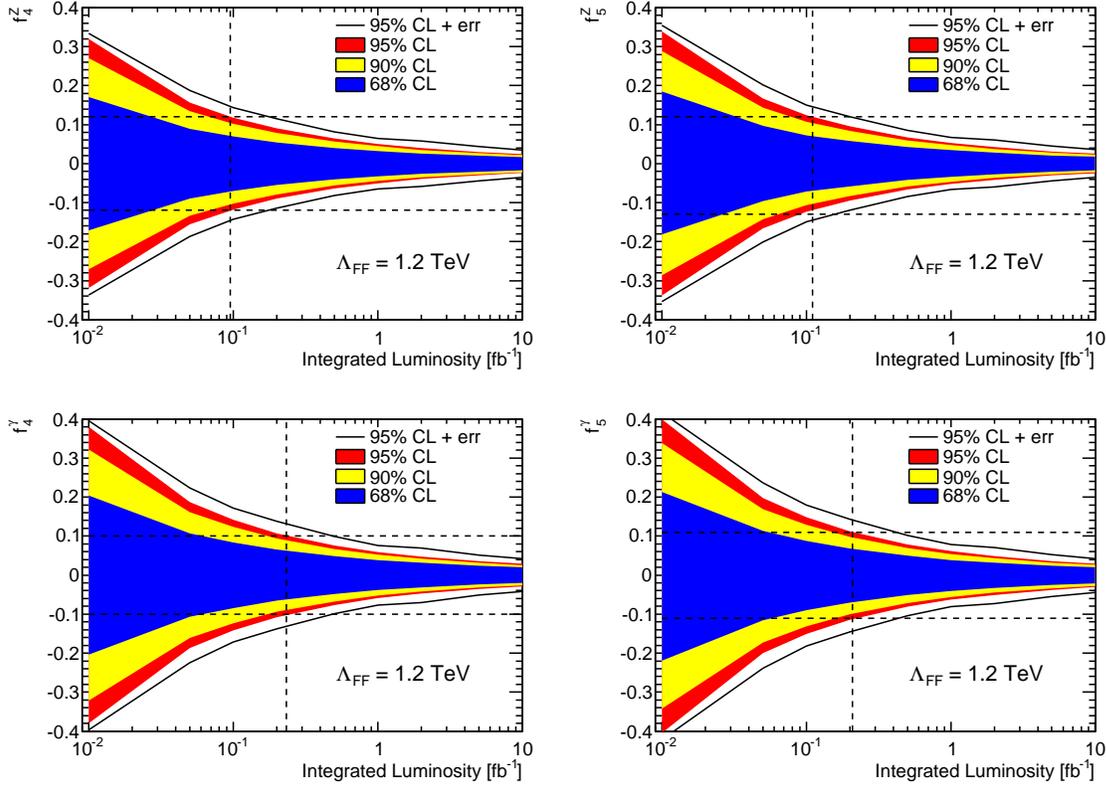


Figure 7.12: Expected anomalous coupling limits as a function of integrated luminosity using combined fits to both the $ZZ \rightarrow lll$ and $ZZ \rightarrow ll\nu\bar{\nu}$ channels. The horizontal dashed lines indicate the current limits obtained from the CDF experiment [40].

A comparison of tables 7.1 and 7.2 show that limits obtained with a cutoff of $\Lambda = 2.0$ TeV are approximately double those obtained with $\Lambda = 1.2$ TeV. At low luminosities, it is possible to obtain limits which violate the unitarity requirements given in equations 1.30 and 1.31. For $\Lambda = 1.2$ TeV, this occurs with the $ZZ \rightarrow lll$ only fits at the lowest luminosity of $\mathcal{L} = 0.01$ fb⁻¹. For $\Lambda = 2.0$ TeV, all of the results obtained for $\mathcal{L} = 0.01$ fb⁻¹ would violate unitarity, as would the $ZZ \rightarrow lll$ limits for $\mathcal{L} = 0.1$ fb⁻¹.

Comparing the expected limits with $\Lambda = 1.2$ TeV in figure 7.12, to the current measurements listed in table 1.1, ATLAS has the potential to set competitive limits at relatively low integrated luminosities. Similar limits as those set by CDF [40] could be reached with as little as 150 pb⁻¹ for the f_i^Z couplings, and 250 pb⁻¹ for f_i^γ . After 1 fb⁻¹ of integrated luminosity, ATLAS could improve on the current CDF measurement by a factor of ~ 2 .



7.7 Systematic Uncertainties

Systematic uncertainties on both $\mu(f_i^V)$ and $P(f_i^V|p_T^k)$ can effect the sensitivity of the anomalous coupling limits. All of the systematic studies have been performed with a form factor cutoff of $\Lambda = 1.2$ TeV.

7.7.1 Uncertainties on $\mu(f_i^V)$

The largest systematic uncertainties on the predicted anomalous coupling limits are from uncertainties on the expected event yields.

Uncertainties in the expected number of ZZ SM events, σ_{SM}^{ZZ} in equation 7.8, are given in table 6.15. Uncertainties arise from Monte Carlo statistics, PDF errors, detector systematics and luminosity. In total, the expected uncertainty on σ_{SM}^{ZZ} is 12% in the $ZZ \rightarrow lll$ channel, and 11% for $ZZ \rightarrow ll\nu\bar{\nu}$. To estimate the effect of these uncertainties, toy Monte Carlo datasets were generated with σ_{SM}^{ZZ} equal to $\pm 12\%$ ($\pm 11\%$) of the central values of 10.1 fb (6.2 fb) in the $ZZ \rightarrow lll$ ($ZZ \rightarrow ll\nu\bar{\nu}$) channel. Fits were made to the datasets, but with the central value of σ_{SM}^{ZZ} . The contribution to the systematic error on the combined limits is ${}_{-8\%}^{+10\%}$ for 1 fb^{-1} of integrated luminosity.

Uncertainties in the number of expected background events in table 6.14 are large due to a lack of sufficient Monte Carlo statistics. Their effect on the anomalous coupling limits was estimated by varying the background fraction, f_{bg} in equation 7.8, by ${}_{-43\%}^{+190\%}$ (${}_{-11\%}^{+104\%}$) in the $ZZ \rightarrow lll$ ($ZZ \rightarrow ll\nu\bar{\nu}$) channel. Fitting to the central value of f_{bg} gives a systematic uncertainty on the combined limits of ${}_{-1\%}^{+11\%}$ for 1 fb^{-1} of integrated luminosity.

7.7.2 Uncertainties on $P(f_i^V|p_T^k)$

Uncertainties on the shape of the p_T spectrum also contribute to the uncertainty on the predicted limits.

The shape of the background p_T spectrum is challenging to predict due to lack of Monte Carlo statistics for $p_T(Z) > 100$ GeV. To estimate the uncertainty due to different background shapes, an additional p_T dependant factor is introduced, so that equation 7.7

becomes

$$B(p_T^k) \propto P(\text{SM}|p_T^k) \times e^{-\lambda p_T^k}. \quad (7.10)$$

The expected limits are recalculated with background shapes for $\lambda = \pm 0.01 \text{ GeV}^{-1}$. These values are taken from the errors on exponential fits to the Monte Carlo p_T shape in the region $p_T(Z) > 100 \text{ GeV}$. The contribution to the systematic uncertainty is expected to be ${}^{+5\%}_{-1\%}$ for 1 fb^{-1} of integrated luminosity.

The detector resolution also gives a small contribution to the systematic uncertainty. To estimate this effect, the width of the Gaussian in equation 7.6 was changed from 5% to $5\%{}^{+5\%}_{-2.5\%}$. The contribution to the limits was ${}^{+0.5\%}_{-0.2\%}$ for 1 fb^{-1} of integrated luminosity.

Limits on f_4^Z were calculated using the SHERPA generator in place of BR. The SHERPA generator includes effects of parton showering, initial state radiation and the underlying event. For 1 fb^{-1} of integrated luminosity, the SHERPA limits were 2% larger than those obtained with BR.

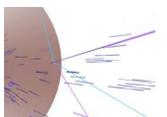
The effect of changing the Parton Density Function was also studied. The expected limits were recalculated using the MRST2004 [108] PDFs. The coupling limits obtained using MRST2004 are typically 2% larger than those using CTEQ6LL.

7.7.3 Summary of Systematic Uncertainties

The systematic uncertainties discussed above are summarised in table 7.3. A total systematic error has been calculated by adding the individual contributions in quadrature. Statistical errors are calculated from the standard deviation of the toy Monte Carlo confidence limit distributions. The mean systematic error on the combined coupling limits for 1 fb^{-1} of integrated luminosity is ${}^{+16\%}_{-9\%}$, compared to an expected statistical variation of $\pm 23\%$.

Systematic errors were found to be dependant on the integrated luminosity, ranging from $\pm 3\%$ for 0.01 fb^{-1} to ${}^{+40\%}_{-20\%}$ for 10 fb^{-1} . Expected combined limits including pessimistic systematic errors are shown in the right-hand column of tables 7.1 and 7.2.

The solid line in figure 7.12 shows the expected combined coupling limits including the worst-case combination of statistical and systematic errors. Even in this scenario, the limits measured by CDF should be reached with 200 pb^{-1} for f_i^Z and 500 pb^{-1} for



Systematic	f_4^Z	f_5^Z	f_4^γ	f_5^γ	Mean
σ_{SM}^{ZZ}	+11%	+9%	+10%	+10%	+10%
	-8%	-9%	-7%	-8%	-8%
f_{bg}	+11%	+9%	+13%	+11%	+11%
	-1%	-3%	-1%	-1%	-1%
$B(p_T^k)$	+6%	+4%	+6%	+6%	+5%
	-1%	-3%	-1%	-2%	-1%
$R(p_T^k)$	+0.4%	+1%	+2%	+0.4%	+0.5%
	-0.2%	-1%	-1%	-0.1%	-0.2%
SHERPA	$\pm 2\%$	$\pm 2\%$	$\pm 2\%$	$\pm 2\%$	$\pm 2\%$
PDF	$\pm 2\%$	$\pm 0.2\%$	$\pm 3\%$	$\pm 2\%$	$\pm 2\%$
Total	+17%	+14%	+18%	+17%	+16%
	-8%	-10%	-8%	-8%	-9%
Statistical	$\pm 23\%$	$\pm 22\%$	$\pm 22\%$	$\pm 23\%$	$\pm 23\%$

Table 7.3: Effects of systematic errors on predicted coupling limits for 1 fb^{-1} of integrated luminosity, with $\Lambda = 1.2 \text{ TeV}$. The statistical error is included for comparison, and is calculated as the RMS of the confidence limit distributions, for example those shown in figure 7.9.

f_i^γ . Including systematic errors results in coupling limits which are ~ 1.5 times tighter than CDF with 1 fb^{-1} of integrated luminosity.

7.8 Gluon-Gluon Fusion

As mentioned in section 1.6.1, gluon-gluon fusion could contribute to the ZZ cross section, but is not included in either the BR or SHERPA generators. The effect of the $gg \rightarrow ZZ$ process on the anomalous coupling sensitivity can be estimated by adding the expected contribution by hand.

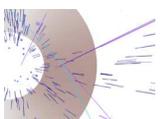
The mean number of events, $\mu(f_i^V)$, is increased by 13% as given in equation 1.40. The inclusion of $gg \rightarrow ZZ$ also changes the shape of the $p_T(Z)$ spectrum, as shown in figure 7.13. The $gg \rightarrow ZZ$ contribution was calculated using GG2ZZ, with $m_Z > 12 \text{ GeV}$, and then added to the BR histograms for $q\bar{q} \rightarrow ZZ$, before the quadratic fits described in section 7.4.

Repeating the expected limit calculation including the $gg \rightarrow ZZ$ contribution yields combined confidence limits which show a difference of $< 1\%$ when compared to the values quoted in tables 7.1 and 7.2.

7.9 Conclusions

In this chapter, unbinned maximum likelihood fits to the Z transverse momentum spectrum in the $ZZ \rightarrow lll$ and $ZZ \rightarrow ll\nu\bar{\nu}$ channels have been used to calculate the expected sensitivity of ATLAS to anomalous triple gauge couplings at $\sqrt{s} = 7 \text{ TeV}$ over a range of integrated luminosities and using two form factor cutoff schemes. ATLAS has good prospects to improve on the current experimental limits with as little as 500 pb^{-1} . For 1 fb^{-1} , ATLAS can expect to place limits of $|f_i^Z| < 0.06$ and $|f_i^\gamma| < 0.07$ at the 95% confidence level, using a form factor cutoff of $\Lambda = 1.2 \text{ TeV}$. This represents an improvement on the current experimental measurements by a factor of ~ 2 .

Statistical uncertainties from variations between toy Monte Carlo datasets are $\pm 23\%$. The systematic uncertainty on the expected limits are $^{+16\%}_{-9\%}$ for 1 fb^{-1} , of which the largest contributions are from uncertainties on the expected number of signal and background events.



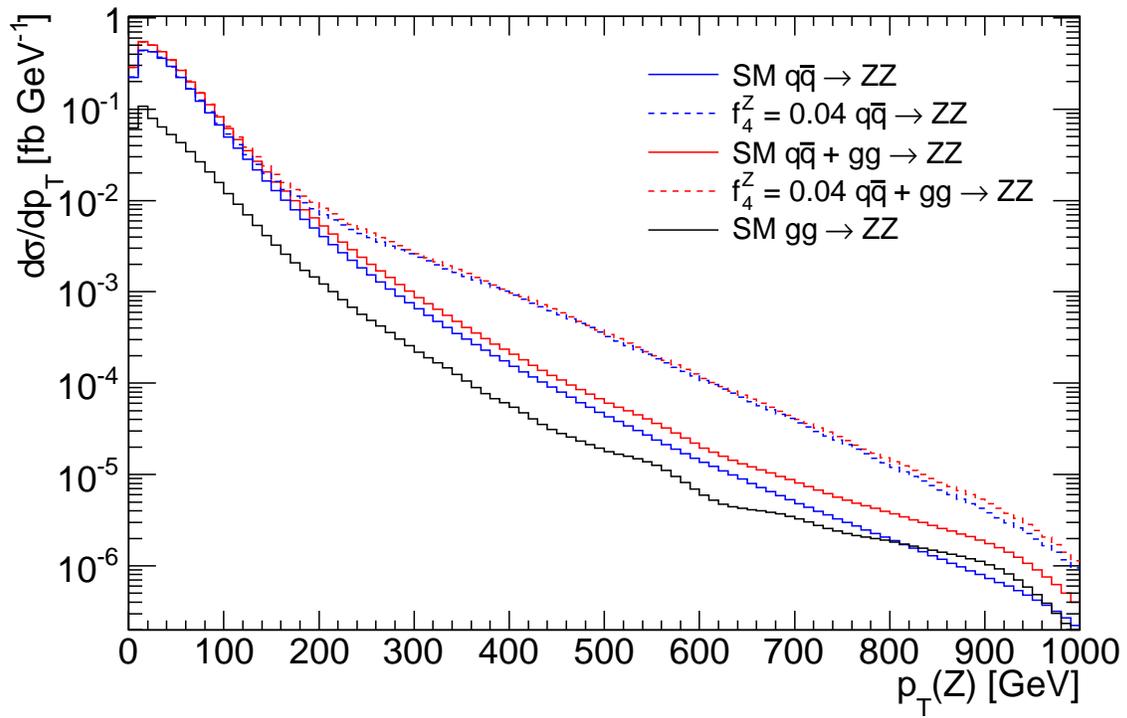


Figure 7.13: The $p_T(Z)$ spectrum for $gg \rightarrow ZZ$ and $q\bar{q} \rightarrow ZZ$ processes at $\sqrt{s} = 7$ TeV, after applying detector resolution smearing.

Chapter 8

Observation of $Z \rightarrow ll$ in Early ATLAS Data

“Guns don’t kill people, physics kills people!”

— Dick Solomon, Third Rock from the Sun

8.1 Introduction

The LHC began producing proton-proton collisions at $\sqrt{s} = 7$ TeV in March 2010, and by July 2010 ATLAS had recorded $328 \pm 36 \text{ nb}^{-1}$ of integrated luminosity [104].

This chapter describes criteria for selection of $Z \rightarrow ee$ and $Z \rightarrow \mu\mu$ events in the first ATLAS data, based on the analyses presented in [105] and [106]. The number of events passing selection cuts will be used to measure the cross section for $Z \rightarrow ll$ production. Finally, the lepton multiplicity and \cancel{E}_T distributions will be examined for evidence of ZZ production.

8.2 Z Production

Single Z bosons are produced in pp collisions at LO via the Feynman diagrams shown in figure 1.11. The NNLO cross section for $Z \rightarrow ll (l = e, \mu)$ production at $\sqrt{s} = 7$ TeV

was calculated by the FEWZ program [26] as

$$\sigma \times \mathcal{B}(Z \rightarrow ll) = 0.94 \pm 0.01 \text{ (scale)} \pm 0.04 \text{ (PDF) nb}, \quad (8.1)$$

where the invariant mass of the lepton pair was required to satisfy $|m_{ll} - 90.1\text{GeV}| < 20 \text{ GeV}$. The calculation was made using renormalisation and factorisation scales of $\mu_F = \mu_R = m_Z$ with the MRST2001 PDF [108]. The PDF uncertainty has been estimated by recalculating the cross section using the CTEQ6M PDF [91]. The scale uncertainty has been obtained by repeating the cross section calculation with $\mu_F = \mu_R = 2m_Z$ and $\mu_F = \mu_R = m_Z/2$, as in section 1.6.

Simulated $Z \rightarrow ll$ signal events have been generated using PYTHIA, as described in section 6.3.3. The ATLAS detector response is fully simulated as described in section 6.3.1. Additional signal samples with more than one pp interaction per event (pile-up) are also produced and will be discussed further in section 8.5. The main backgrounds are from $Z \rightarrow \tau\tau$, $W \rightarrow l\nu$ and $t\bar{t}$ production, the LO Feynman diagrams of which are shown in figures 1.11 and 1.12.

8.3 Event Selection

The data used in this chapter were collected between March and July 2010, during the first months of high-energy LHC operation, as presented in figure 8.1. Events must pass beam and detector data-quality requirements which ensure the relevant sub-detectors are operational. This results in a total integrated luminosity of 316 nb^{-1} for the $Z \rightarrow ee$ channel and 314 nb^{-1} for the $Z \rightarrow \mu\mu$ channel, corresponding to $\sim 96\%$ of the total luminosity delivered to ATLAS. The uncertainty on the luminosity measurements is 11%, as discussed in section 6.8.2.

8.3.1 Pre-selection

Data from the ATLAS detector are written to disk and processed in a number of different streams. Events are pre-selected from calorimeter triggered events (L1Calo stream) if two or more electrons have been reconstructed with the *egamma* algorithm, both with $p_T > 5 \text{ GeV}$. Muon triggered events (MuonswBeam stream) are selected if two or more STACO muons are present, again with $p_T > 5 \text{ GeV}$. Details of the *egamma* and STACO

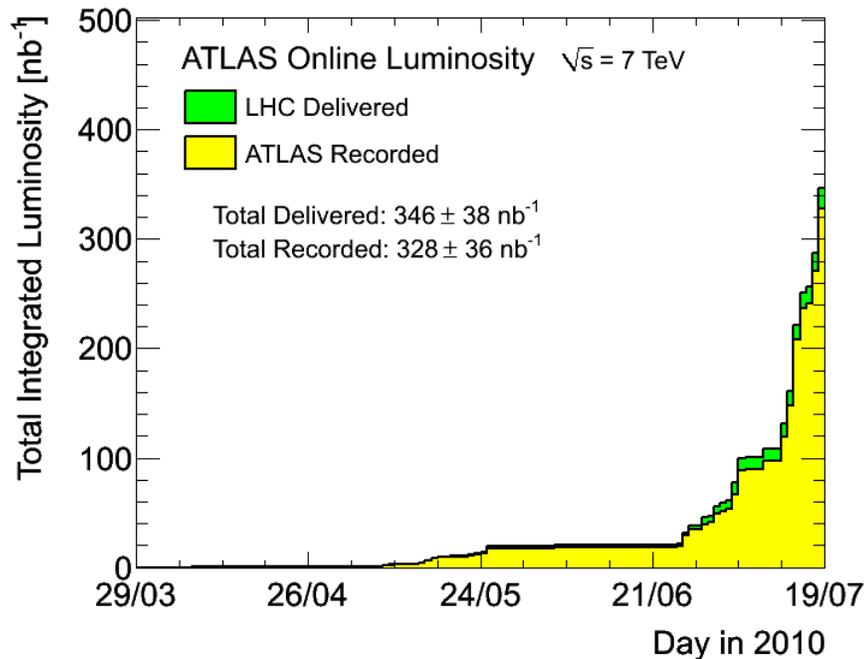
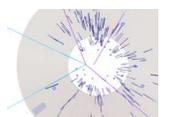


Figure 8.1: Cumulative integrated luminosity versus day delivered to (green), and recorded by ATLAS (yellow) during stable beams at $\sqrt{s} = 7$ TeV centre-of-mass energy, from [104].

muon algorithms can be found in sections 6.4.1 and 6.4.2 respectively. Pre-selection is performed on AOD files using the LHC computing grid, after which 2.5×10^6 $Z \rightarrow ee$ and 3.5×10^5 $Z \rightarrow \mu\mu$ events remain. The large difference between these numbers can be attributed to the fact that electron candidates can be more easily faked by jets compared to muon triggered events.

Events are required to have been recorded when the LHC and relevant ATLAS sub-detectors were operational, with separate periods defined for both electron and muon events. The periods take the form of lists of good runs and luminosity blocks, as defined in [109]. Collisions are selected by requiring that at least one reconstructed vertex is present with more than ten tracks.

Events can occasionally contain localised high-energy calorimeter deposits, which although rare, can have a significant impact on the \cancel{E}_T measurement. Such events are removed from the analysis using cleaning requirements described in [110]. Briefly, a jet with $E_T > 10$ GeV is defined as bad if it passes one of the following criteria:



- Noisy cells in the hadronic end-cap: The fraction of energy deposited in the hadronic end-cap, $f_{HEC} > 0.8$ and 90% of the energy deposit is contained in five or fewer cells ($n_{90} \leq 5$);
- Out-of-time: The jet timing, $|t_{\text{jet}}| > 50$ ns;
- Noisy cells in the Liquid Argon: The quality of a reconstructed jet > 0.8 and the fraction of energy in the electromagnetic calorimeters, $\text{emf} > 0.95$.

An event is rejected if it contains one or more bad jets reconstructed with the anti- k_{\perp} algorithm in a cone of $\Delta R < 0.4$. The cleaning criteria remove 0.02% and 0.4% of electron- and muon-triggered events respectively.

Problems in the liquid argon calorimeter can have a significant impact on the electron reconstruction. Events are rejected if a candidate electromagnetic cluster is located in a problematic region of the detector, including regions with high-voltage problems, isolated noisy or dead cells, and electronic front-end boards not providing output. The loss in electron acceptance due to this requirement is approximately 13% [105].

Finally, events are selected with an L1A hardware trigger as described in section 2.4. The instantaneous LHC luminosity up until July 2010 peaked at $\sim 10^{30} \text{ cm}^{-2}\text{s}^{-1}$, $\mathcal{O}(10^4)$ lower than the expected nominal value. This allowed ATLAS to record data with only the first level trigger active. The calorimeter trigger was used to select electron candidates with a threshold of ~ 10 GeV (L1EM10). The corresponding L1 muon trigger (L1MU6) selects patterns of hits consistent with $p_T > 6$ GeV muons.

8.3.2 Selection Cuts

In the electron channel, candidate events are selected by requiring an oppositely charged electron-positron pair. These leptons must be reconstructed with the *egamma* algorithm; electrons with energy deposits only in the forward calorimeters are not considered. In addition, both leptons must pass the medium criteria as described in section 6.4.1, and have a cluster energy $E_T > 20$ GeV within the pseudorapidity range $|\eta| < 2.47$, excluding the transition region between the barrel and end-cap calorimeters ($1.37 < |\eta| < 1.52$).

In the muon channel, oppositely charged muon candidates are selected with $p_T > 20$ GeV and $|\eta| < 2.4$. Both must be combined muons, reconstructed from standalone muon tracks associated to inner-detector tracks, as described in section 6.4.2. All muon

candidates must also satisfy isolation requirements $\sum p_T^{\text{ID}}/p_T < 0.2$ in a cone of $\Delta R < 0.4$.

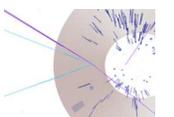
Tables 8.1 (8.2) show the number of electron (muon) events passing each cut. Monte Carlo simulated datasets have been normalised to the number of events expected for 316 (314) nb^{-1} in the electron (muon) channel. In total 57 events pass all selection criteria in the electron channel and 109 in the muon channel, within a dilepton invariant mass window $|m(ll) - m_Z| < 20$ GeV. Figure 8.2 shows examples of candidate events from in electron and muon channels.

Figure 8.3 shows the electron cluster E_T and muon p_T of the selected lepton candidates in data, with statistical errors shown. The signal and background distributions expected from Monte Carlo samples is also presented in the plot, normalised to the number of data events. As the backgrounds are small, they will not be presented in subsequent plots.

Figure 8.4 shows the dilepton invariant mass and transverse momentum distributions of the selected events. A Voigtian (Breit-Wigner convolved with a Gaussian) fit is made to the data with the following form:

$$f(m) = \frac{1}{(m - \mu)^2 + \frac{1}{4}\Gamma^2} \otimes \exp\left(-\frac{1}{2}\frac{m^2}{\sigma^2}\right), \quad (8.2)$$

where m is the invariant mass, μ is the resonant mass and σ is the width due to detector resolution. Γ is the natural Z boson width, which is fixed during the fits to the measured value of 2.4952 GeV [8]. Table 8.3 lists the resulting fitted values for both Monte Carlo and data distributions. The electron fits show a good agreement with Monte Carlo for both μ and σ . In the muon case, while μ is in good agreement, the width σ is larger in data than Monte Carlo due to mis-alignments in the forward region of the muon spectrometer and the inner-detector for high- p_T tracks.

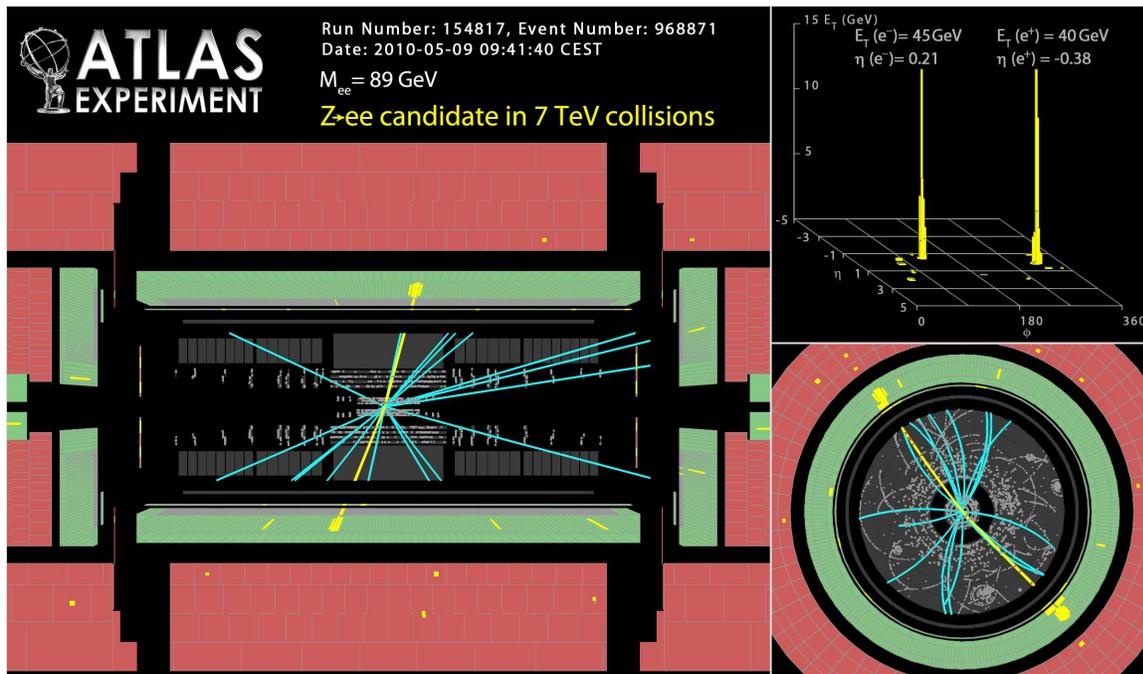


Cut	Data	$Z \rightarrow ee$	$Z \rightarrow ee$ (Pile-up)	$Z \rightarrow \tau\tau$	$W \rightarrow l\nu$	$t\bar{t}$
Preselection	2512052	293	293	60.1	326	23.5
Detector	2203582	293	293	60.1	326	23.5
Vertex	2195233	273	289	57.5	315	23.5
Jet Cleaning	2194774	273	289	57.5	315	23.4
LAr Cleaning	1994754	196	208	48.9	279	19.8
Trigger	1521584	193	204	43.7	266	19.5
$n_l \geq 2$	1521584	193	204	43.7	266	19.5
$q(ll) = 0$	475807	114	120	19.0	98.5	8.00
$egamma$	234148	112	117	14.8	62.3	5.55
Medium	157	94.3	97.4	0.872	1.17	0.155
$ \eta_{l_1} < 2.47$	157	94.1	97.2	0.871	1.16	0.155
$ \eta_{l_2} < 2.47$	156	94.0	97.1	0.869	1.16	0.154
$ \eta_{l_1} $ Crack	140	88.7	91.7	0.813	1.04	0.147
$ \eta_{l_2} $ Crack	130	83.9	86.5	0.765	0.926	0.138
$E_T(l_1) > 20$ GeV	82	82.6	85.2	0.566	0.785	0.135
$E_T(l_2) > 20$ GeV	63	75.9	78.5	0.188	0.167	0.0949
$ m(ll) - m_Z < 20$ GeV	57	72.8	75.3	0.0219	0.0493	0.0502
Number (316nb^{-1})	57 ± 7.5	72.8 ± 0.1	75.3 ± 0.2	0.022 ± 0.002	0.049 ± 0.005	0.050 ± 0.002
Events passing	57	1179913	256557	149	110	1931
Overall Efficiency	$4.5 \times 10^{-5}\%$	25%	24%	0.0075%	0.0016%	0.19%

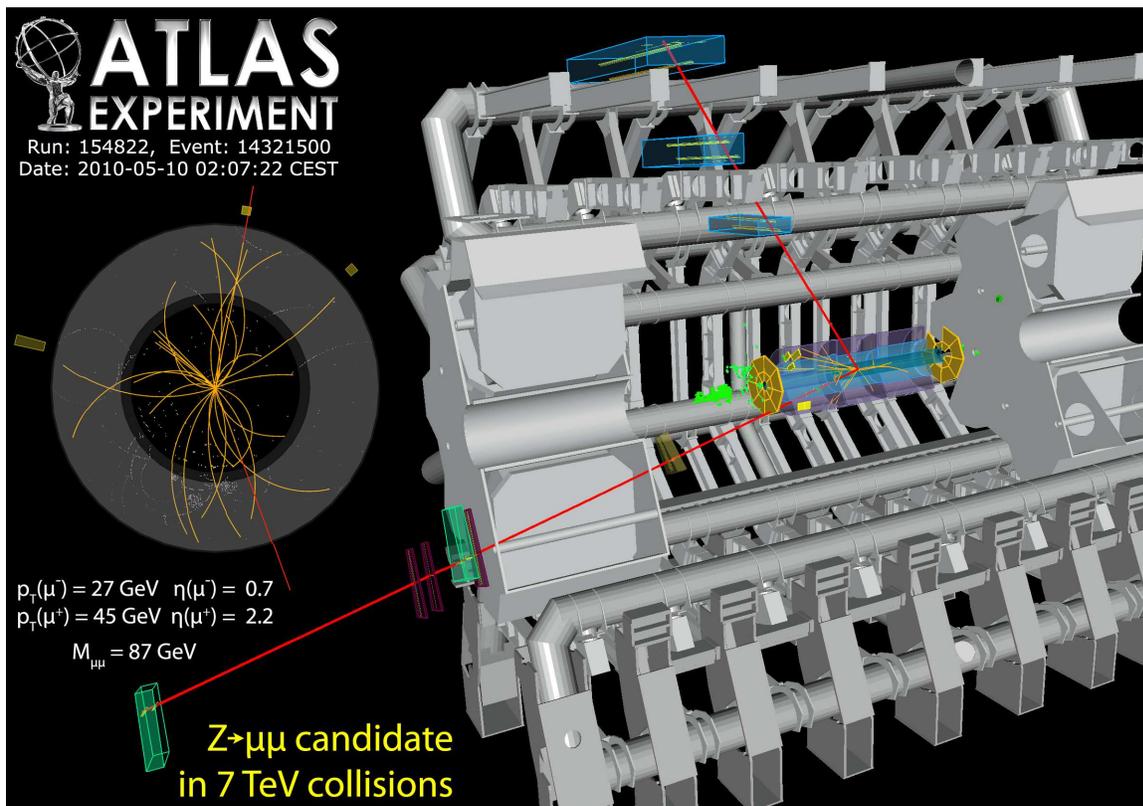
Table 8.1: Number of events passing $Z \rightarrow ee$ selection cuts for data and Monte Carlo samples. The Monte Carlo samples have been scaled to the integrated luminosity.

Cut	Data	$Z \rightarrow \mu\mu$	$Z \rightarrow \mu\mu$ (Pile-up)	$Z \rightarrow \tau\tau$	$W \rightarrow l\nu$	$t\bar{t}$
Preselection	352607	311	311	63.7	91.9	24.9
Detector	274980	311	311	63.7	91.9	24.9
Vertex	100175	289	306	61.0	90.4	24.9
Jet Cleaning	99787	289	306	60.9	90.3	24.8
Trigger	38701	249	263	6.63	65.4	11.3
$n_l \geq 2$	38701	249	263	6.63	65.4	11.3
$q(ll) = 0$	6925	169	178	1.87	3.69	1.31
Combined	563	146	153	1.60	1.14	1.15
Isolation μ_1	230	145	152	1.57	1.02	0.928
Isolation μ_2	197	143	149	1.50	0.144	0.467
$ \eta_{l_1} < 2.4$	195	138	144	1.46	0.141	0.460
$ \eta_{l_2} < 2.4$	184	133	139	1.41	0.135	0.452
$p_T(l_1) > 20$ GeV	142	130	136	1.00	0.123	0.442
$p_T(l_2) > 20$ GeV	122	119	125	0.326	0.0306	0.324
$ m(ll) - m_Z < 20$ GeV	109	113	119	0.0447	0.0132	0.101
Number (314nb^{-1})	109 ± 10	113.3 ± 0.1	118.9 ± 0.2	0.045 ± 0.003	0.013 ± 0.003	0.101 ± 0.002
Events passing	109	1813705	381749	287	28	3609
Overall Efficiency	0.0002%	36%	38%	0.014%	0.0004%	0.36%

Table 8.2: Number of events passing $Z \rightarrow \mu\mu$ selection cuts for data and Monte Carlo samples. The Monte Carlo samples have been scaled to the integrated luminosity.

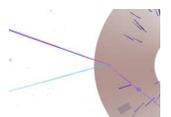


(a) Atlantis event display of a $Z \rightarrow ee$ candidate.



(b) Virtual Point 1 event display of a $Z \rightarrow \mu\mu$ candidate.

Figure 8.2: The first $Z \rightarrow ll$ candidate events in ATLAS, from [76].



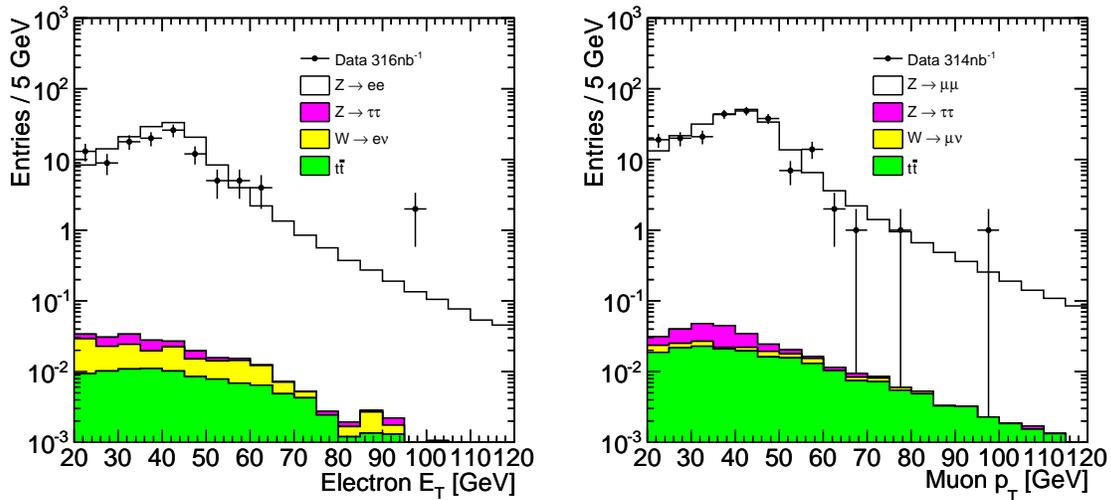


Figure 8.3: Electron transverse energy (left) and muon track transverse momentum (right) distributions for leptons in Z candidate events.

8.4 Z Cross Section

A simple measurement of the Z cross section can be obtained using the following formula:

$$\sigma_Z \times \mathcal{B}(Z \rightarrow ll) = \frac{N_{\text{sig}} - N_{\text{bg}}}{A \cdot \epsilon \times \mathcal{L}_{\text{int}}}, \quad (8.3)$$

where N_{sig} is the measured number of signal events, N_{bg} the number of background events, \mathcal{L}_{int} is the integrated luminosity, and $A \cdot \epsilon$ is the product of the geometrical acceptance and selection efficiency.

The number of background events is calculated from the number of Monte Carlo events passing cuts, from tables 8.1 and 8.2. Of the background samples, only events from the $t\bar{t}$, $Z \rightarrow \tau\tau$ and $W \rightarrow l\nu$ channels were found to pass the selection criteria. This yields total background expectations of 0.12 ± 0.01 events in the electron channel and 0.168 ± 0.004 events in the muon channel. Errors quoted are the from Monte Carlo statistics and have been summed in quadrature from each of the background channels.

The factor $A \cdot \epsilon$ represents a combination of the geometrical acceptance and the trigger, reconstruction and selection efficiencies for signal events. It is calculated as the number of Monte Carlo signal events passing all cuts divided by the total number of events generated with the true Z mass $|m_Z - 91.2 \text{ GeV}| < 20 \text{ GeV}$. In the electron (muon) channel this factor is 0.258(0.379). The systematic errors arising from detector uncertainties are

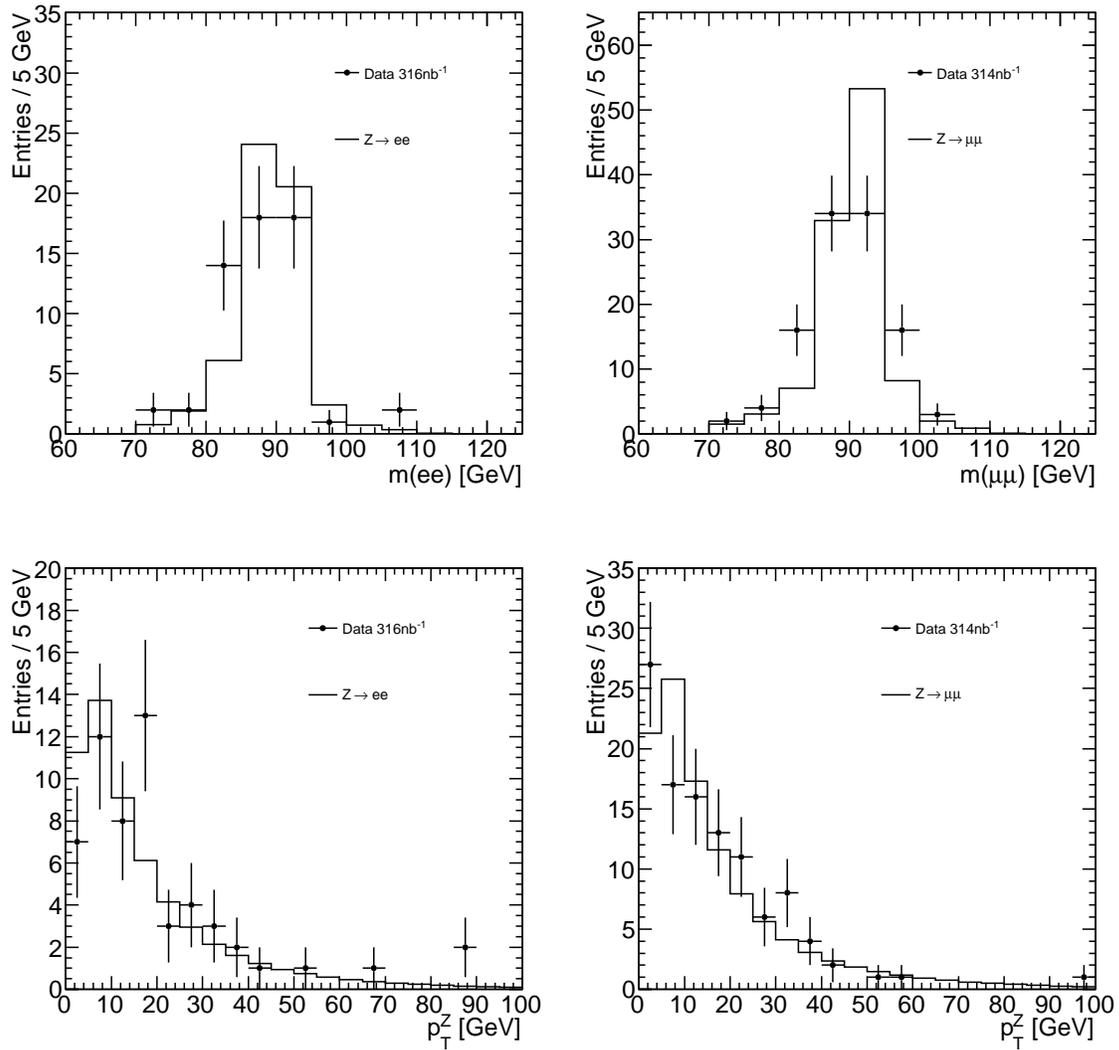
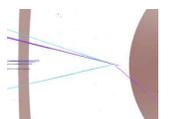


Figure 8.4: Dilepton invariant mass (top) and Z transverse momentum (bottom) for electrons (left) and muons (right).

Parameter	$Z \rightarrow ee$		$Z \rightarrow \mu\mu$	
	Data	Monte Carlo	Data	Monte Carlo
μ/GeV	87.5 ± 0.7	89.21 ± 0.01	89.6 ± 0.6	90.72 ± 0.01
σ/GeV	3.6 ± 0.6	3.07 ± 0.01	4.5 ± 0.5	2.77 ± 0.01

Table 8.3: Parameters obtained from fitting equation 8.2 to data and Monte Carlo dilepton invariant mass distributions in the $Z \rightarrow ee$ and $Z \rightarrow \mu\mu$ channels. Statistical errors are shown on each parameter.



described in section 6.8, and are expected to be $\pm 14\%$ and $\pm 7\%$ in the electron and muon channels respectively.

The resulting cross section for $Z \rightarrow ee$ production, where $|m_Z - 91.2\text{GeV}| < 20\text{ GeV}$, is given by equation 8.3 as

$$\sigma(Z \rightarrow ee) = 0.70 \pm 0.09 \text{ (stat)} \pm 0.10 \text{ (syst)} \pm 0.08 \text{ (lumi) pb.} \quad (8.4)$$

The corresponding cross section measurement for $Z \rightarrow \mu\mu$ is

$$\sigma(Z \rightarrow \mu\mu) = 0.90 \pm 0.09 \text{ (stat)} \pm 0.07 \text{ (syst)} \pm 0.10 \text{ (lumi) pb.} \quad (8.5)$$

Both cross section measurements show a good agreement with the values quoted in the ATLAS study [105].

The measurements are consistent with the NNLO prediction given in equation 8.1. The measured value for $\sigma(Z \rightarrow ee)$ is lower than both the predicted and the muon measurement. This is understood to be due to larger electron lateral shower shapes in data than in Monte Carlo [111].

8.5 A Diboson Prelude

The analysis of single Z candidates in the first 300 nb^{-1} of ATLAS data is a necessary prerequisite to diboson searches. The predicted yields of ZZ events outlined in section 6 rely heavily on Monte Carlo simulations. Figures 8.3 and 8.4 lend some support to these predictions, showing that ATLAS can reconstruct Z candidates in good agreement with Monte Carlo predictions. As a prelude to future studies, events passing single Z cuts were tested using the diboson selection criteria as follows.

8.5.1 $ZZ \rightarrow lll$ Channel

Of the single Z boson candidates, none were found to contain additional leptons which pass the cuts described in section 6.5. This is consistent with the prediction of 3×10^{-3} events from Monte Carlo.

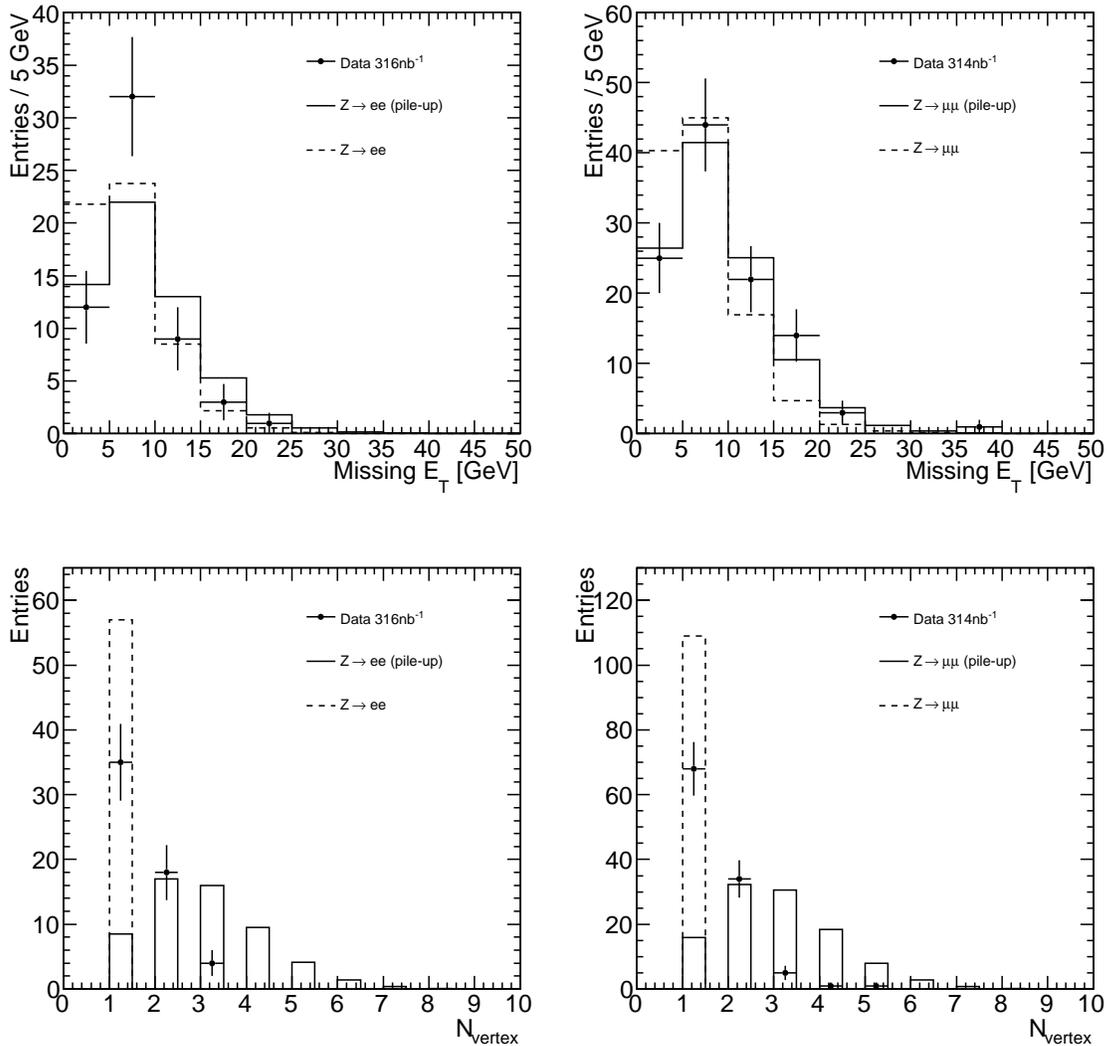
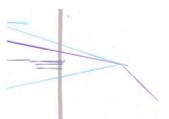


Figure 8.5: Missing transverse energy (top) and number of vertices (bottom) for electrons (left) and muons (right).

8.5.2 $ZZ \rightarrow ll\nu\bar{\nu}$ Channel

Events with a single reconstructed Z boson and high missing transverse momentum are a signature for $ZZ \rightarrow ll\nu\bar{\nu}$ events, as discussed in section 6.6. The \cancel{E}_T quantity described in section 6.4.4 is not appropriate for early data studies, as it is highly reliant on Monte Carlo optimised object identification. Instead, the \cancel{E}_T is calculated following the method used in [106]. For electron candidates, the \cancel{E}_T components are calculated as in equation 6.24, using energy measured directly from topological cells. For muon candidates, \cancel{E}_T is calculated from the reconstructed momenta of muons measured in the



range $|\eta| < 2.7$ and the calorimeter term in equation 6.24:

$$\cancel{E}_{x,y}^{\text{muon}} = - \sum_{\text{muon}} p_{x,y} + \cancel{E}_{x,y}^{\text{Calo}}. \quad (8.6)$$

The overall missing transverse energy is calculated as in equation 6.25. As described in section 8.3.1, events are vetoed if they contain localised high-energy calorimeter deposits, which can lead to anomalously high \cancel{E}_T measurements. The missing transverse energy spectrum of candidate Z events is shown in the upper plots of figure 8.5.

Multiple pp interactions within a single bunch-crossing, known as pile-up, can affect the \cancel{E}_T measurement. The lower plots of figure 8.5 show the number of reconstructed vertices with more than three associated tracks. The plots show a comparison between data and two types of $Z \rightarrow ll$ Monte Carlo samples. The first contains only a single pp interaction per event. The second is a pile-up sample, containing a Poisson distributed number of minimum bias events, with a mean of three, superimposed on the $Z \rightarrow ll$ interaction. The mean number of vertices per event from the data is 1.1 ± 0.5 . The effect on the \cancel{E}_T measurement can be seen in the upper plots of figure 8.5, where the data show a much better agreement to the Monte Carlo sample which includes pile-up.

None of the Z candidate events satisfy the high missing transverse energy cut, $\cancel{E}_T > 50$ GeV, and so no events pass $ZZ \rightarrow ll\nu\bar{\nu}$ cuts described in section 6.6, consistent with the expected number of 2.4×10^{-3} .

8.5.3 Anomalous Coupling Limits

Fits were made to the anomalous coupling parameters as outlined in chapter 7, given no observed events and for an integrated luminosity of 315 nb^{-1} . The resulting coupling limits with a form factor cutoff of $\Lambda_{FF} = 1.2$ TeV produce confidence limits which are four times larger than the unitarity requirements given in equations 1.30 and 1.31.

8.6 Conclusions

This chapter has applied a set of criteria to select $Z \rightarrow ll (l = e, \mu)$ candidate events from the first 315 nb^{-1} of collision data from ATLAS at $\sqrt{s} = 7$ TeV. In total, 57 and 109 events pass the electron and muon selection cuts respectively. The Z cross section

for $|m_Z - 91.2\text{GeV}| < 20\text{ GeV}$ was measured as

$$\sigma(Z \rightarrow ee) = 0.70 \pm 0.09 \text{ (stat)} \pm 0.10 \text{ (syst)} \pm 0.08 \text{ (lumi)} \text{ pb},$$

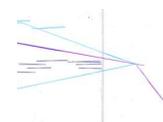
in the electron channel and

$$\sigma(Z \rightarrow \mu\mu) = 0.90 \pm 0.09 \text{ (stat)} \pm 0.07 \text{ (syst)} \pm 0.10 \text{ (lumi)} \text{ pb},$$

in the muon channel. Both results are consistent with the theoretical NNLO cross section prediction of

$$\sigma \times \mathcal{B}(Z \rightarrow ll) = 0.94 \pm 0.01 \text{ (scale)} \pm 0.04 \text{ (PDF)} \text{ nb}.$$

The search for single Z candidates is an important prerequisite for the ZZ diboson search described in chapter 6. None of the single Z events pass diboson selection requirements, consistent with the expected number of 5×10^{-3} . In general, there is good agreement between the data and Monte Carlo predictions, indicating that ATLAS has excellent prospects for ZZ observation with higher integrated luminosities.



Chapter 9

Summary

The LHC recently began colliding bunches of protons at a centre-of-mass energy of $\sqrt{s} = 7$ TeV, heralding a new era in particle physics. The ATLAS detector has successfully recorded collisions at the LHC, and the search for new physics beyond the Standard Model is underway.

One important component of the ATLAS detector is the Inner Detector, which tracks charged particles from the collision point to the calorimeters. The SCT, one of the three sub-detectors of the Inner Detector, was introduced in chapter 3. In order to read out the SCT's 4088 channels at trigger rates up to 100 kHz requires a robust and reliable DAQ. Thorough development and testing of new DAQ features was performed using the barrel and end-cap test systems described in chapter 4. A selection of such features was described in chapter 5, and included an event simulator, used to test the robustness of the DAQ at high occupancy. Another prominent development was a hardware-based monitoring framework, which is used to measure SCT module occupancy, spacepoint rates and timing on-the-fly during LHC collisions.

The production of events containing pairs of Z bosons provides sensitivity to new physics. Chapter 6 outlined a set of criteria to select such diboson events in the $ZZ \rightarrow lll(l = e, \mu)$ and $ZZ \rightarrow ll\nu\bar{\nu}$ channels. For 1 fb^{-1} of integrated luminosity at a centre of mass energy of $\sqrt{s} = 7$ TeV, ATLAS can be expected to observe 10 ± 1 events in the $ZZ \rightarrow lll$ channel, with $0.5_{-0.2}^{+0.9}$ background events. In the $ZZ \rightarrow ll\nu\bar{\nu}$ channel, 6.2 ± 0.7 signal events are expected, with a background of $1.9_{-0.2}^{+2.0}$ events.

Anomalous couplings between three neutral gauge bosons were introduced in chapter 1, and can be described by four parameters, $f_{i=4,5}^{V=Z,\gamma}$. All of the coupling parameters are zero at tree-level in the Standard Model, and hence their measurement provides a di-

rect probe of New Physics. Chapter 7 described a technique for measuring the strength of anomalous couplings by fitting the number of selected ZZ events and the $Z p_T$ spectrum of such events. It was shown that with 1 fb^{-1} of integrated luminosity at $\sqrt{s} = 7 \text{ TeV}$, ATLAS has the potential to place constraints on the coupling parameters of $|f_i^Z| < 0.06$ and $|f_i^\gamma| < 0.07$ at the 95% confidence level. These limits assume a form factor with a cutoff of $\Lambda = 1.2 \text{ TeV}$, and represent an improvement on current measurements by a factor of two.

By summer 2010, the first single Z candidate events were observed in ATLAS. In chapter 8, criteria were defined to select $Z \rightarrow ll$ events in the first 315 nb^{-1} of ATLAS collision data at $\sqrt{s} = 7 \text{ TeV}$. In total 57 events are observed in the electron channel, with 109 in the muon channel, leading to cross section measurements of

$$\sigma(Z \rightarrow ee) = 0.70 \pm 0.09 \text{ (stat)} \pm 0.10 \text{ (syst)} \pm 0.08 \text{ (lumi) pb}$$

and

$$\sigma(Z \rightarrow \mu\mu) = 0.90 \pm 0.09 \text{ (stat)} \pm 0.07 \text{ (syst)} \pm 0.10 \text{ (lumi) pb},$$

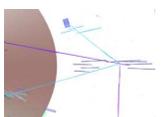
both of which are consistent with the NNLO calculation of

$$\sigma \times \mathcal{B}(Z \rightarrow ll) = 0.94 \pm 0.01 \text{ (scale)} \pm 0.04 \text{ (PDF) nb}.$$

In general, the measured lepton kinematics of selection $Z \rightarrow ll$ events show good agreement with Monte Carlo, an encouraging sign for future ATLAS diboson studies.

Looking further to the future, the LHC is expected to deliver pp collisions at $\sqrt{s} = 7 \text{ TeV}$ until the end of 2011, after which additional safety features will be installed during a year long shutdown. It is hoped that these safety measures will finally allow collisions at the LHC design energy of 14 TeV . At this energy, ATLAS has the potential to measure anomalous coupling limits of $|f_4^V| < 0.007$ and $|f_5^V| < 0.008$ for 30 fb^{-1} [44], improving on the current measurements by a factor of 10.

This thesis has spanned an exciting period in particle physics, from the final preparations and commissioning of the ATLAS SCT, to the analysis of data from the world's highest energy collisions. It has been shown that ATLAS has excellent prospects for measuring new physics in the ZZ sector, with the potential to push our understanding of the Universe further than ever before.



Appendix A

Thesis Statistics

“Well it’s too late now.”

— Sundance

A.1 Word Counting

This thesis was written in L^AT_EX over a period from December 2009 until February 2011. It comprises 236 pages, 8249 lines, 39112 words and 126 figures. It was compiled a total of 1151 times. The statistics are summarised in table A.1.

Quantity	Count
Pages	236
Lines	8249
Words	39112
Unique Words	3479
Figures	126
Iterations	1151
Days	98
Cups of tea	196

Table A.1: Summary of statistics for this thesis.

The evolution of the page, line, word and figure counts against time is shown in the upper-left plot of figure A.1. The lower-left plot shows the same quantities on a logarithmic scale to highlight recent activity. The upper-right plot shows the page and word rate as a function of time.

Word	Count	Fraction %
the	3607	9.296
of	1463	3.770
in	1064	2.742
to	995	2.564
a	924	2.381
physics	63	.162
diboson	28	.072
module	83	.213
bottom	3	.007
violation	2	.005
chocolate	1	.002

Table A.2: Frequency of selected words.

Table A.2 shows the frequency of selected words used in this thesis. Zipf's Law [112, 113], states that while only a few words are used very often, most are used rarely. This leads to the empirical law

$$g(r) = an^{-b}, \quad (\text{A.1})$$

where $g(r)$ is the frequency of a word ranked n^{th} in a body of text. a is a constant, while b is the exponent which is typically close to one. The distribution of $g(r)$ for this thesis is shown in the lower-right hand plot of figure A.1. Fitting equation A.1 yields an exponent of 0.912 ± 0.002 .

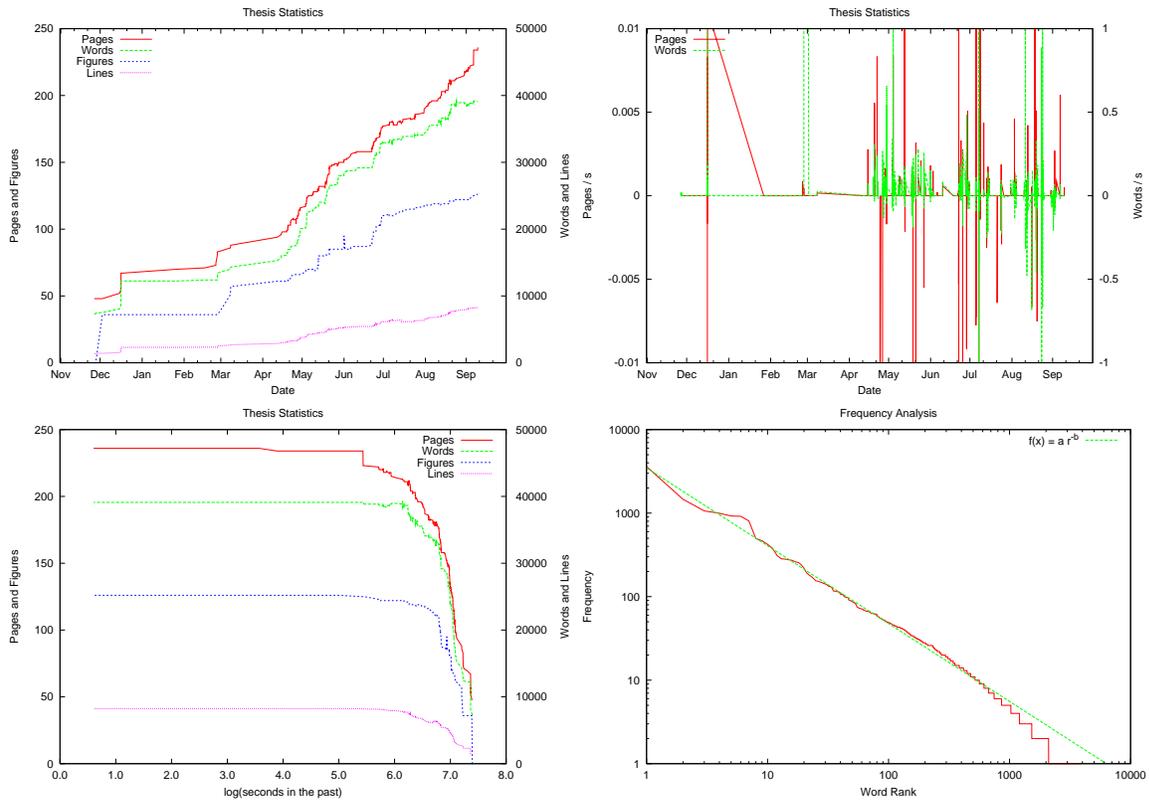
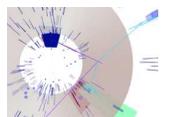


Figure A.1: Evolution the page, line, word and figure counts in this thesis.



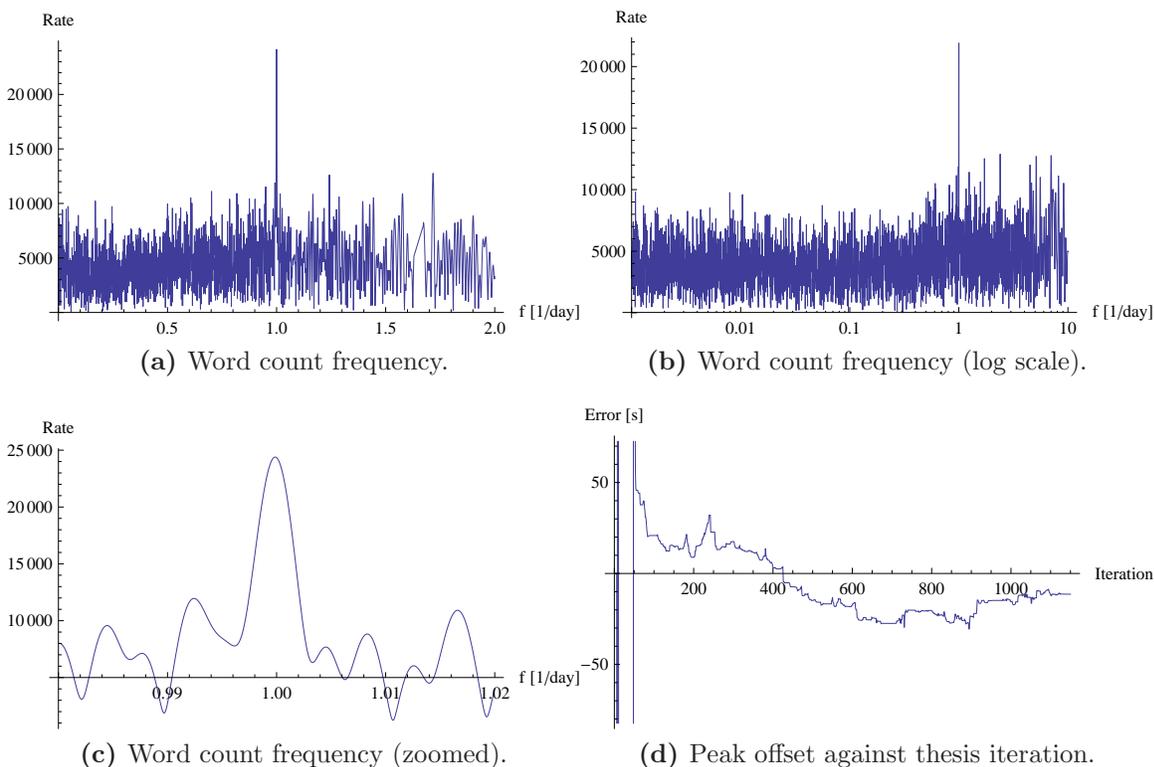
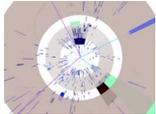


Figure A.2: Fourier transforms of the word rate.

A.2 Measuring the Length of the Day

Inspired by pioneering work in the field of procrastination [114], the thesis-writing habits of the author can be used to estimate the length of the day. By taking a Fourier transform of the word rate as shown in figure A.1, figure A.2a is obtained. There is a prominent peak at around one day, indicating strong daily writing habits.

Numerical routines can be used to find the exact position of the maximum in figure A.2c. The position of the peak is measured at 1.00013 days, which is accurate to within 11 s. This corresponds to a fractional discrepancy of 0.01%, making it the most accurate measurement in this thesis. The offset of the peak with respect to the true day as a function of time is shown in figure A.2d. It should be noted that the time has not been corrected for British Summer Time, or changes in time zone incurred during visits to CERN. Figure A.2b shows no evidence for the week.



Colophon

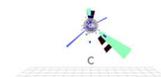
This thesis was made in L^AT_EX 2_ε using the “hepthesis” class [115].

Bibliography

- [1] S. Weinberg, Phys. Rev. Lett. **19**, 1264 (1967).
- [2] A. Salam and J. Ward, Physics Letters **13**, 168 (1964).
- [3] K. Hagiwara, R. D. Peccei, D. Zeppenfeld, and K. Hikasa, Nucl. Phys. **B282**, 253 (1987).
- [4] G. J. Gounaris, J. Layssac, and F. M. Renard, Phys. Rev. D **61**, 073013 (2000).
- [5] G. J. Gounaris, J. Layssac, and F. M. Renard, Phys. Rev. D **62**, 073012 (2000).
- [6] U. Baur and E. L. Berger, Phys. Rev. D **47**, 4889 (1993).
- [7] U. Baur and D. L. Rainwater, Phys. Rev. **D62**, 113011 (2000).
- [8] W. M. Yao *et al.*, Journal of Physics G: Nuclear and Particle Physics **33**, 1+ (2006).
- [9] G. J. Gounaris, J. Layssac, and F. M. Renard, Phys. Rev. D **62**, 073013 (2000).
- [10] E. W. N. Glover and J. J. V. der Bij, Nuclear Physics B **321**, 561 (1989).
- [11] C. Zecher, T. Matsuura, and J. J. van der Bij, Z. Phys. **C64**, 219 (1994).
- [12] J. Pumplin *et al.*, JHEP **07**, 012 (2002).
- [13] P. M. Nadolsky *et al.*, Phys. Rev. D **78**, 013004 (2008).
- [14] A. Sherstnev and R. S. Thorne, Eur. Phys. J. **C55**, 553 (2008).
- [15] S. Catani, F. Krauss, B. R. Webber, and R. Kuhn, Journal of High Energy Physics **2001**, 063 (2001).
- [16] S. Böser, <http://projects.hepforge.org/hepmcvisual>, 2009.
- [17] T. Sjostrand *et al.*, Comp. Phys. Comm. **135**, 238 (2001).
- [18] B. Andersson, G. Gustafson, G. Ingelman, and T. Sjostrand, Physics Reports **97**, 31 (1983).
- [19] S. Frixione and B. R. Webber, JHEP **06**, 029 (2002).

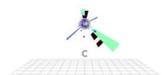
- [20] S. Frixione and B. R. Webber, (2006), hep-ph/0612272.
- [21] G. Corcella *et al.*, JHEP **01**, 010 (2001).
- [22] J. M. Butterworth, J. R. Forshaw, and M. H. Seymour, Z. Phys. **C72**, 637 (1996).
- [23] M. L. Mangano, M. Moretti, F. Piccinini, R. Pittau, and A. D. Polosa, JHEP **07**, 001 (2003).
- [24] J. Campbell and R. K. Ellis, Phys. Rev. D **62**, 114012 (2000).
- [25] J. M. Campbell and R. K. Ellis, Phys. Rev. D **60**, 113006 (1999).
- [26] C. Anastasiou, L. Dixon, K. Melnikov, and F. Petriello, Phys. Rev. D **69**, 094008 (2004).
- [27] T. Binoth, N. Kauer, and P. Mertsch, (2008), 0807.0024.
- [28] T. Gleisberg *et al.*, JHEP **02**, 007 (2009).
- [29] B. Mellado, J. Qian, D. Rebutzi, and R. Tanaka, Diboson Inclusive Cross Sections at 7 TeV, 2010.
- [30] A. D. Martin, W. J. Stirling, R. S. Thorne, and G. Watt, European Physical Journal C **63**, 189 (2009).
- [31] P. M. Nadolsky *et al.*, Phys. Rev. **D78**, 013004 (2008).
- [32] J. M. Campbell, J. W. Huston, and W. J. Stirling, Reports on Progress in Physics **70**, 89 (2007).
- [33] G. Arnison *et al.*, Physics Letters B **126**, 398 (1983).
- [34] P. Bagnaia *et al.*, Physics Letters B **129**, 130 (1983).
- [35] The LEP Collaborations, (2006), hep-ex/0612034.
- [36] The LEP Collaborations, CERN-PH-EP/2005-051 (2005).
- [37] The CDF Collaboration, Phys. Rev. Lett. **100**, 201801 (2008).
- [38] The CDF Collaboration, CDF/PHYS/ELECTROWEAK/PUB/9910 (2009).
- [39] The D0 Collaboration, Phys. Rev. Lett. **101**, 171803 (2008).
- [40] The CDF Collaboration,
<http://www-cdf.fnal.gov/physics/ewk/2008/ZZZatGC/ZZZaTGC.html> (2008).
- [41] D. Cavalli, ATL-PHYS-INT-2007-005 (2007).
- [42] S. Hassani, ATL-COM-PHYS-2002-012 (2003).
- [43] The ATLAS Collaboration, CERN Report No. CERN-OPEN-2008-020, 2009 (unpublished).

- [44] K. Bachas *et al.*, CERN Report No. ATL-PHYS-INT-2009-007, ATL-COM-PHYS-2008-041, 2008 (unpublished).
- [45] J. D. Chapman, CERN Report No. ATL-COM-PHYS-2007-105, 2007 (unpublished).
- [46] The D0 Collaboration, Phys. Rev. Lett. **100**, 131801 (2008).
- [47] S. Ask, Eur. Phys. J. **C60**, 509 (2009).
- [48] L. Evans and P. Bryant, Journal of Instrumentation **3**, S08001 (2008).
- [49] The ATLAS Collaboration, Journal of Instrumentation **3**, S08003 (2008).
- [50] The CMS Collaboration, Journal of Instrumentation **3**, S08004 (2008).
- [51] The LHCb Collaboration, Journal of Instrumentation **3**, S08005 (2008).
- [52] The ALICE Collaboration, Journal of Instrumentation **3**, S08002 (2008).
- [53] The LHCf Collaboration, Journal of Instrumentation **3**, S08006 (2008).
- [54] The TOTEM Collaboration, Journal of Instrumentation **3**, S08007 (2008).
- [55] J. Elmsheuser *et al.*, Journal of Physics: Conference Series **119**, 072014 (2008).
- [56] T. Barber, The data acquisition and calibration system for the ATLAS Semiconductor Tracker, in *16th IEEE-NPSS Real Time Conference*, pp. 465–471, 2009.
- [57] S. M. Sze and K. K. Ng, *Physics of Semiconductor Devices* (Wiley-Interscience, 2006).
- [58] M. Bruzzi, Nuclear Science, IEEE Transactions on **48**, 960 (2001).
- [59] A. Abdesselam *et al.*, Nucl. Instrum. Meth. **A568**, 642 (2006).
- [60] A. Abdesselam *et al.*, Nucl. Instrum. Meth. **A575**, 353 (2007).
- [61] M. Moll, DESY-THESIS-1999-040.
- [62] E. Coniavitis, A. Okamoto, and R. Tanaka, CERN Report No. ATL-COM-INDET-2009-039, 2009 (unpublished).
- [63] M. Dentan *et al.*, Nuclear Science, IEEE Transactions on **43**, 1763 (1996).
- [64] A. Abdesselam *et al.*, JINST **3**, P01003 (2008).
- [65] A. Abdesselam *et al.*, JINST **2**, P09003 (2007).
- [66] T. J. Barber *et al.*, Nucl. Instrum. Meth. **A538**, 442 (2005).
- [67] H. C. Van der Bij *et al.*, CERN Report No. ATL-DAQ-2002-007, 2001 (unpublished).



- [68] J. Joseph *et al.*, *ATLAS Silicon ReadOut Driver (ROD) Users Manual*.
- [69] C. Bee *et al.*, ATLAS Communication **ATL-DAQ-98-129** (2004).
- [70] The Object Management Group, <http://www.omg.org/spec/CORBA>.
- [71] Team-COOL,
<https://twiki.cern.ch/twiki/bin/view/Persistency/Cool>.
- [72] B. Demirkoz, CERN-THESIS-2008-001.
- [73] The ATLAS SCT Collaboration,
<https://twiki.cern.ch/twiki/bin/view/Atlas/SCTPublicResults>.
- [74] K. Hanagaki, P. S. Miyagawa, H. Pernegger, I. Torchiani, and M. White, CERN Report No. ATL-COM-INDET-2009-038, 2009 (unpublished).
- [75] The ATLAS SCT Collaboration,
<https://twiki.cern.ch/twiki/bin/view/Atlas/ApprovedPlotsSCT>.
- [76] The ATLAS Collaboration,
<https://twiki.cern.ch/twiki/bin/view/Atlas/EventDisplayPublicResults>.
- [77] DuPont, http://www2.dupont.com/Kapton/en_US/tech_info/index.html.
- [78] F. Campabadal *et al.*, Nucl. Instrum. Meth. **A538**, 384 (2005).
- [79] D. Dobos, <https://twiki.cern.ch/twiki/bin/view/Atlas/PixelSimulator>.
- [80] P. Alfke, Xilinx Report No. XAPP 52, 1996 (unpublished).
- [81] W. Dabrowski *et al.*, *ABCD3T Chip Specification Version V1.2*.
- [82] J. Joseph *et al.*, *ATLAS Silicon ReadOut Driver (ROD) Users Manual, Appendix A, Register Definitions*.
- [83] T. Barber, <https://twiki.cern.ch/twiki/bin/view/Atlas/SctRodMonitor>.
- [84] P. Calafiura, W. Lavrijsen, C. Leggett, M. Marino, and D. Quarrie, The athena control framework in production, new developments and lessons learned, in *Interlaken 2004, Computing in high energy physics and nuclear physics*, pp. 456–458, 2004.
- [85] J. Moscicki *et al.*, Computer Physics Communications **180**, 2303 (2009).
- [86] D. C. Vanderster *et al.*, Journal of Physics: Conference Series **219**, 072056 (2010).
- [87] T. J. Barber,
<https://svnweb.cern.ch/trac/atlasusr/browser/tbarber/DibosonAnalysis>.
- [88] R. Brun and F. Rademakers, Nucl. Inst. and Meth. in Phys. Res. A **389**, 81 (1997).

- [89] M. Asai, *Trans. Amer. Nucl. Soc.* **95**, 757 (2006).
- [90] The ATLAS Collaboration, CERN Report No. ATLAS-SOFT-2010-01-003, 2010 (unpublished).
- [91] J. Pumplin *et al.*, *Journal of High Energy Physics* **2002**, 012 (2002).
- [92] Y. Kataoka, <https://twiki.cern.ch/twiki/bin/view/AtlasProtected/PDFReweight>, 2009.
- [93] N. D. Gagunashvili, (2006), hep-ex/0605123.
- [94] J. Hartert and I. Ludwig, CERN Report No. ATL-COM-PHYS-2010-070, 2010 (unpublished).
- [95] T. Cornelissen *et al.*, CERN Report No. ATL-SOFT-PUB-2007-007, ATL-COM-SOFT-2007-002, 2007 (unpublished).
- [96] D. Levin, <https://twiki.cern.ch/twiki/bin/view/AtlasProtected/MuonRecoPedia>.
- [97] J. E. Huth *et al.*, Presented at Summer Study on High Energy Physics, Reearch Directions for the Decade, Snowmass, CO, Jun 25 - Jul 13, 1990.
- [98] M. Cacciari, G. P. Salam, and G. Soyez, *Journal of High Energy Physics* **2008**, 063 (2008).
- [99] L. Asquith *et al.*, CERN Report No. ATL-COM-PHYS-2009-630, 2010 (unpublished).
- [100] W. Lampl *et al.*, CERN Report No. ATL-LARG-PUB-2008-002, ATL-COM-LARG-2008-003, 2008 (unpublished).
- [101] D. Cavalli *et al.*, CERN Report No. ATLAS-COM-CONF-2010-059, 2010 (unpublished).
- [102] G. J. Feldman and R. D. Cousins, *Phys. Rev. D* **57**, 3873 (1998).
- [103] A. Bertin and other, CERN Report No. ATL-COM-LUM-2010-021, 2010 (unpublished).
- [104] G. Aad, CERN Report No. ATLAS-CONF-2010-060, 2010 (unpublished).
- [105] G. Aad *et al.*, CERN Report No. ATLAS-CONF-2010-076, 2010 (unpublished).
- [106] G. Aad *et al.*, CERN Report No. ATLAS-CONF-2010-051, 2010 (unpublished).
- [107] W. Verkerke and D. P. Kirkby, (2003), physics/0306116.
- [108] A. Martin, R. Roberts, W. Stirling, and R. Thorne, *Physics Letters B* **604**, 61 (2004).



- [109] M. Baak *et al.*, CERN Report No. ATL-COM-GEN-2009-015, 2009 (unpublished).
- [110] The ATLAS Collaboration, CERN Report No. ATLAS-CONF-2010-038, 2010 (unpublished).
- [111] M. Aharrouche *et al.*, CERN Report No. ATL-COM-PHYS-2010-701, 2010 (unpublished).
- [112] G. K. Zipf, *Human behavior and the principle of least effort* (Addison-Wesley, 1949).
- [113] Y.-S. Chen and F. F. Leimkuhler, *Information Processing and Management* **23**, 171 (1987).
- [114] C. G. Lester and M. A. Parker, *Model independent sparticle mass measurements at ATLAS*, PhD thesis, Cambridge Univ., Geneva, 2001, Presented on 12 Dec 2001.
- [115] A. Buckley, The hepthesis L^AT_EX class.

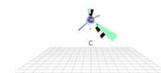
Acronyms

The field of High-Energy Physics (HEP) is rich with acronyms, and the use of Three Letter Acronyms (TLAs) can sometimes be overwhelming. This section lists the main acronyms used in this thesis, to help the reader tell their RCC from their NLO.

ACR	ATLAS Control Room
ASIC	Application-Specific Integrated Circuit
BCID	Bunch Crossing ID
BCR	Bunch Clock Reset
BOC	Back-of-Crate
BPM	Bi-Phase Mark
BR	Baur and Rainwater
CC	Charged Current
CDF	Collider Detector at Fermilab
CERN	The European Organisation for Nuclear Research
CL	Confidence Limit
CMS	Compact Muon Solenoid
COOL	Conditions Data Storage Model
CORBA	Common Object Resource Broker Architecture
CSC	Cathode Strip Chamber
CTP	Central Trigger Processor
DAC	Digital-to-Analogue Converter

DAQ	Data Acquisition
DCS	Detector Control System
DMA	Direct Memory Access
DMILL	Durci Mixte sur Isolant Logico-Lineaire
DORIC	Digital Optical Receiver Integrated Circuit
DSP	Digital Signal Processor
ECR	Event Count or soft reset
EFB	Event Fragment Builder
EM	electromagnetic
ENC	Equivalent Noise Charge
EW	Electro-weak
FCal	Forward Calorimeter
FIFO	First-In, First-Out
FPGA	Field Programmable Gate Array
FSM	Finite State Machine
GPN	General Public Network
GUI	Graphical User Interface
HEC	Hadronic End-cap Calorimeter
HEP	High-Energy Physics
HTL	Header-Trailer Limit
ID	Inner Detector
IPC	Inter-process Communication
IS	Information Server
ISR	Interrupt Service Routine
L1A	Level-1 trigger
L1ID	Level-1 trigger ID
LAr	Liquid Argon

LCG	.LHC Computing Grid
LEP	.Large Electron-Positron
LFSR	.Linear Feedback Shift Register
LHC	.Large Hadron Collider
LINAC	.Linear Accelerator
LO	.Leading Order
LO*	.Modified Leading Order
LSP	.Lightest Supersymmetric Particle
LTA	.Long Term Attachment
MDSP	.Master DSP
MDT	.Monitored Drift Tube
MSSM	.Minimally Supersymmetric Standard Model
NC	.Neutral Current
NEWT	.New Tracking
NIKHEF	.The National Institute for Subatomic Physics
NLO	.Next-to-Leading Order
NNLO	.Next-to-Next-to-Leading Order
NP	.New Physics
NTGC	.Neutral Triple Gauge Coupling
PanDA	.Production and Distributed Analysis Framework
PDF	.Parton Distribution Function
PS	.Proton Synchrotron
QCD	.Quantum Chromodynamic
QFT	.Quantum Field Theory
RCC	.ROD Crate Controller
RCF	.ROD Controller FPGA



RMS	Root Mean Square
ROD	Readout Driver
RoI	Region of Interest
ROS	Readout Subsystem
RPC	Resistive Plate Chamber
Rx	optical receiver
SBC	Single Board Computer
SCT	Semiconductor Tracker
SDSP	Slave DSP
SM	Standard Model
SPS	Super Proton Synchrotron
STACO	STATistical COmbination
STFC	Science and Technology Research Council
SUSY	Supersymmetry
TC	technicolour
TDAQ	Trigger and DAQ
TGC	Thin Gap Chamber
TIM	Timing Interface Module
TLA	Three Letter Acronym
TRT	Transistion Radition Tracker
TTC	Trigger, Timing and Control
Tx	optical transmission
VCSEL	Vertical Cavity Surface Emitting Laser
VDC	VCSEL Driver Chip
VHDL	VHSIC Hardware Description Language
VHSIC	Very High-Speed Integrated Circuit
VME	Verse Module European
ZDC	Zero-Degree Calorimeter

List of Figures

1.1	Feynman diagram of the weak charged current interaction.	7
1.2	Feynman diagram of the WWZ triple gauge boson interaction.	9
1.3	Feynman rule for the anomalous ZZV vertex.	10
1.4	Anomalous coupling unitarity limits.	12
1.5	Production of an anomalous coupling vertex via heavy fermions.	12
1.6	Leading order Feynman diagrams for ZZ production at the LHC.	13
1.7	Feynman diagrams for ZZ production via gluon-gluon fusion.	13
1.8	Cartoon of a ZZ event.	16
1.9	ZZ Cross Section	19
1.10	ZZ branching fractions.	21
1.11	Leading order Feynman diagrams for Z and W production.	22
1.12	Leading order Feynman diagrams for $t\bar{t}$ production.	22
1.13	Leading order Feynman diagrams for $Zb\bar{b}$ production.	22
1.14	ZZ cross section with anomalous couplings.	23
1.15	Combined ZZ results from the LEP experiments.	26
1.16	$ZZ \rightarrow lll$ events observed at the Tevatron.	27
1.17	Predicted ZZ kinematics in ATLAS at $\sqrt{s} = 14$ TeV.	28
2.1	A cartoon showing the layout of the LHC underground areas.	32

2.2	The ATLAS Detector.	33
2.3	A diagram of the ATLAS Inner Detector.	37
2.4	The ATLAS calorimeter system.	38
2.5	The ATLAS muon spectrometer.	39
2.6	Block diagram of the ATLAS TDAQ subsystems, from[49].	41
3.1	Quarter-section of the ATLAS inner detector.	47
3.2	Photographs of SCT barrel and end-cap modules.	47
3.3	The silicon layers used in the construction of an SCT module.	49
3.4	The components of the ABCD3T chip.	51
3.5	Off-detector electronics of the SCT.	53
3.6	The SCT ROD data path.	56
3.7	The structure of the SCT DAQ software.	58
3.8	Example calibration plots for barrel sector module 20220170200123 . . .	60
3.9	SCT noise characteristics from late 2009.	63
3.10	SCT spacepoints during a beam splash event.	65
3.11	Atlantis event displays highlighting SCT performance.	66
3.12	SCT performance during collisions.	67
4.1	The SCT barrel and end-cap sectors.	70
4.2	Photos of the barrel sector during construction.	71
4.3	Photograph of the end-cap disk.	74
4.4	The SCT sectors as seen by the DAQ and DCS software.	76
4.5	Characteristic calibration parameters for 47 modules from the SR1 sectors.	77
5.1	Noisiest module hits against event number (run 89335).	82

5.2	The SCT formatter simulation block.	83
5.3	Example of a 19-bit Linear Feedback Shift Register.	87
5.4	A software-implemented 19-bit XNOR LFSR.	88
5.5	Simulator occupancy.	91
5.6	Monitoring hit-map with the simulator enabled.	92
5.7	Event fragment size against occupancy.	93
5.8	ABCD3T chip numbering scheme.	96
5.9	Noise occupancy of the SCT from the ROD monitoring.	102
5.10	Noise occupancy comparison at nominal SCT conditions.	103
5.11	Ratio of measured and predicted coincidence rates for noise.	105
5.12	The coincidence rate per module in the SCT.	105
5.13	Coincidence rate map with cosmic triggers.	107
5.14	Beam splash event as recorded by the ROD monitoring	108
5.15	Spacepoint rates in the SCT for 7 TeV collisions.	109
5.16	Spacepoint map in the SCT for 7 TeV collisions.	110
5.17	Derived particle flux against time.	111
5.18	The fraction of hits with 01X time structure for noise and cosmic triggers.	112
5.19	SCT Collision timing plots.	113
6.1	Re-weighting from 10 to 7 TeV.	121
6.2	Electron reconstruction efficiency.	130
6.3	Efficiency, purity and resolution plots for muons in the analysis.	133
6.4	$ZZ \rightarrow e^+e^-\mu^+\mu^-$ event display.	137
6.5	$ZZ \rightarrow lll$ kinematic Variables	141
6.6	$ZZ \rightarrow e^+e^-\nu\bar{\nu}$ event display.	145

6.7	$ZZ \rightarrow ll\nu\bar{\nu}$ kinematic Variables	147
6.8	Predicted number of signal and background events.	153
7.1	Examples of Z p_T spectra including anomalous couplings.	159
7.2	The cross section against anomalous coupling strength.	161
7.3	Fitted anomalous coupling parameters.	162
7.4	ZZ selection efficiency against $p_T(Z)$	164
7.5	Z boson reconstruction resolution.	165
7.6	Shape of the Z p_T spectrum for signal and background.	165
7.7	Number of events against coupling strength.	167
7.8	Example Toy Monte Carlo datasets.	168
7.9	Best fitting and 95% confidence limits for 1000 toy datasets.	170
7.10	Z boson p_T against best fitting coupling.	171
7.11	Best fitting f_4^Z coupling against the input value.	172
7.12	Anomalous coupling limits as a function of integrated luminosity.	175
7.13	Gluon-gluon fusion $p_T(Z)$ spectrum.	180
8.1	Cumulative integrated luminosity against time.	183
8.2	$Z \rightarrow ll$ candidate event displays.	187
8.3	Lepton kinematics in Z candidate events.	188
8.4	Dilepton invariant mass and Z transverse momentum.	189
8.5	\cancel{E}_T and number of vertices of Z candidates.	191
A.1	Evolution the page, line, word and figure counts in this thesis.	201
A.2	Fourier transforms of the word rate.	202

List of Tables

1.1	95% confidence limits on anomalous triple gauge couplings.	29
2.1	Performance goals of the ATLAS detector.	36
2.2	Intrinsic measurement accuracies of the Inner Detector.	37
3.1	SCT dimensions and distribution of modules.	48
3.2	Active length of silicon sensors in the SCT.	49
3.3	The state of the SCT configuration as of May 2010.	62
5.1	Bits of the link header word.	81
5.2	Formatter slice usage with the simulator.	84
5.3	The SCT formatter simulation configuration register.	85
5.4	The location of the simulator registers on the SCT ROD.	86
5.5	The trigger and fast commands of the SCT Module command protocol.	86
5.6	Number generated with a 19-bit XNOR LFSR.	88
5.7	Coincident and non-coincident chips.	96
5.8	Bits in the ATLAS L1A event type.	99
5.9	Comparison of noise-corrected coincidence rates.	107
6.1	Signal and background simulated data samples at $\sqrt{s} = 10$ TeV.	119
6.2	Signal and background simulated data samples at $\sqrt{s} = 7$ TeV.	123

6.3	Cross sections for $ZZ \rightarrow lll$ ($l = e, \mu, \tau$) production.	124
6.4	Cross sections for $ZZ \rightarrow ll\nu\bar{\nu}$ ($l = e, \mu$).	125
6.5	Summary of electron pre-selection cuts.	131
6.6	Summary of muon pre-selection cuts.	134
6.7	Expected yields in 1 fb^{-1} at 10 TeV, for the $ZZ \rightarrow lll$ channel.	140
6.8	Re-weighted yields in 1 fb^{-1} at 7 TeV, for the $ZZ \rightarrow lll$ channel	140
6.9	Expected yields in 1 fb^{-1} at 7 TeV, for the $ZZ \rightarrow lll$ channel.	140
6.10	Expected yields in 1 fb^{-1} at 10 TeV, for the $ZZ \rightarrow ll\nu\bar{\nu}$ channel.	149
6.11	Reweighted yields in 1 fb^{-1} at 7 TeV, for the $ZZ \rightarrow ll\nu\bar{\nu}$ channel	149
6.12	Expected yields in 1 fb^{-1} at 7 TeV, for the $ZZ \rightarrow ll\nu\bar{\nu}$ channel.	149
6.13	trigger efficiency after selection cuts.	151
6.14	Expected diboson signal and background yields.	154
6.15	Selection uncertainties for ZZ events.	155
7.1	Expected 95% confidence limits, with $\Lambda_{FF} = 1.2 \text{ TeV}$	173
7.2	Expected 95% confidence limits, with $\Lambda_{FF} = 2.0 \text{ TeV}$	174
7.3	Systematic errors on predicted coupling limits.	178
8.1	Number of events passing $Z \rightarrow ee$ selection cuts.	186
8.2	Number of events passing $Z \rightarrow \mu\mu$ selection cuts.	186
8.3	Fits to the dilepton invariant mass.	189
A.1	Summary of statistics for this thesis.	199
A.2	Frequency of selected words.	200