Exploring Electroweak Symmetry Breaking
with Jet Substructure at the ATLAS Experiment

Adam Robert Davison

University College London

Submitted to University College London in fulfilment of the requirements for the
award of the degree of Doctor of Philosophy, 17th September 2010

I, Adam Davison, confirm that the work presented in this thesis is my own. Where information has been derived from other sources, I confirm that this has been indicated in the thesis.

Abstract

An important unsolved problem in physics is the nature of electroweak symmetry breaking in the Standard Model. The ATLAS experiment aims to gain insight by studying proton-proton collisions at $\sqrt{s} = 14$ TeV. In order to differentiate between different theoretical models it is important to measure processes where hadrons are produced, such as the hadronic decay of a W, Z or a Higgs boson. However, these decays produce extremely complex signals in the detector which must be analysed carefully. Jet substructure techniques are presented as a novel approach to analysing hadronic signatures relevant to electroweak symmetry breaking. The potential performance of these techniques is evaluated in detail using simulated ATLAS data. Additionally material related to the use of visualisation software to explore ATLAS data is presented.

Summary of Contributions

Modern particle physics experiments tend to be conducted by large collaborations. Thousands of people work on the LHC experiments. As such it is necessary to be especially clear about which parts of this thesis are my own.

Generally speaking this thesis builds upon a great body of research into the fundamental nature of the Universe, especially the Standard Model of particle physics. Chapters 2 and 4 discuss this work of which only the interpretation is mine. The LHC and ATLAS collaborations provide the physical machinery required to make the measurements described in this thesis. The ATLAS collaboration also provides a software framework which enables access to the experimental data, described in Chapter 3, which is not my work.

The contents of Chapter 5 were produced in close collaboration with Erkcan Özcan. I contributed to all aspects of the work presented including software development and analysis. Chapter 6 contains a great deal of insight from Jon Butterworth and Gavin Salam however the plots and numbers presented were produced entirely by me. All material in Chapters 7 and 8 was produced entirely by me.

Finally in Chapter 9, Atlantis is part of the ATLAS software and was written by the Atlantis development team (of which I was a member during this PhD). However all the new work described in Section 9.4 and onward was implemented solely by me.

It should also be considered implicit that all of the work in this thesis contains the guidance of my supervisor Jon Butterworth.

Acknowledgements

I would like to thank Jon Butterworth for his excellent supervision throughout this PhD. I would also like to thank my parents, without whom none of this would have been possible.

I would also like to thank the many people who have assisted the work in this thesis both professionally and personally. Including Zuzanna Wojciechowska for always supporting me; Erkcan Özcan for teaching me so many useful things early in my PhD; Catrin Bernius for all her support and friendship and Gavin Salam, Mathieu Rubin, Giacinto Piacquadio, Christian Weiser and Karl Jakobs for their collaboration.

I would also like to thank the STFC for their financial support; the entire Atlantis development team; the UCL High Energy Physics group and everyone from the ATLAS and LHC projects. Finally I would like to thank my examiners Mark Lancaster and Tony Doyle for their constructive feedback.

Contents

List of Figures

10

List of Tables

Chapter 1

Introduction

The Standard Model of particle physics has been hugely successful in describing a wide range of phenomena. It accurately represents the behaviour of all the known components of matter under the action of the strong, weak and electromagnetic forces. The Standard Model has been found to agree with many high precision measurements [1, 2, 3]. However, the mechanism of electroweak symmetry breaking in the Standard Model has not yet been probed directly. At the LHC, ATLAS will observe collisions at energies well above the electroweak scale. This provides an opportunity to directly test the nature of electroweak symmetry breaking.

Many of the processes which could shed light on electroweak symmetry breaking involve the production of hadrons. These interactions are complex and as such a variety of techniques is required in order to measure these processes. In recent years there have been simultaneous advances in both the theoretical understanding of hadronic processes and the experimental technologies for their measurement. This has made possible a family of new "jet substructure" based approaches to some of the processes which hold promise for understanding electroweak symmetry breaking at the LHC.

This document will begin with a brief discussion of the relevant theory and experimental design, followed by some background on reconstructing hadronic signals at collider experiments. It will then detail my work in evaluating the effectiveness of several new techniques for measuring hadronic signals. Finally there will be some discussion of the use of visualisation software within ATLAS and the ways in which modern graphics technologies can be utilised by a collider experiment.

Chapter 2

The Standard Model

The Standard Model of particle physics is a unified mathematical description of the fundamental components of all visible matter and their interactions through the strong, weak and electromagnetic forces. Formally, the Standard Model is an $SU(3) \times SU(2) \times U(1)$ gauge quantum field theory. It consists of a Lagrangian which describes the behaviour of a set of quantised fields corresponding to the 12 fermionic particles and 4 bosons known to exist. Each particle is represented by degrees of freedom in a field or mixture of fields as shown in Table 2.1.

Many observable properties of the Universe are successfully modelled by the mathematical form of the Standard Model [3]. The fields correctly represent the properties of all known fundamental particles (with the exception of neutrino mass). The interactions of the fields accurately describe all known physical processes except gravity. The symmetries of the theory lead to conservations of quantities such as energy, angular momentum and electric charge. Overall it is a remarkably complete description of many complex phenomena. The electromagnetic interaction encapsulates the cohesion of atoms and molecules and the effects of macroscopic electric and magnetic fields. The weak interaction describes many decay processes, such as the decay of radioactive nuclei, neutrons and heavy quarks. Finally the strong force describes the binding of quarks and gluons into particles such as protons, neutrons and atomic nuclei.

Although the Standard Model may be considered directly in terms of fields, there is an important alternative representation in the form of Feynman diagrams. In the Feynman approach, it is demonstrated that calculations in the field theory are mathematically equivalent to a series of integrals over discrete interactions between point-like

Field		SU(3)	SU(2)	Particles
Left-handed quark	Q_L	Triplet	Doublet	u_L and d_L
Left-handed up anti-quark	$\overline{u}_L$	Triplet	Singlet	$\overline{u}_L$
Left-handed down anti-quark	$\overline{d}_L$	Triplet	Singlet	$\overline{d}_L$
Left-handed lepton	L_L	Singlet	Doublet	e_L and ν_L
Left-handed anti-lepton	$\overline{e}_L$	Singlet	Singlet	$\overline{e}_L$
Weak gauge	B_μ	Singlet	Singlet	$W^\pm$, Z and γ
	W_μ	Singlet	Triplet	
Gluon field	G_μ	Octet	Singlet	g
Higgs field	H	Singlet	Doublet	$H^\dagger$

† Not yet observed

Table 2.1: Fields and Particles in the 1$^\text{st}$ generation of the Standard Model. The 2$^\text{nd}$ and 3$^\text{rd}$ generations introduce additional copies of the first 6 fields. Right-handed particles arise through CPT symmetry.

particles [7] as depicted in Figure 2.1. In the field picture, a potential created by the photon field can influence the evolution of an electron wave-function. Whereas in the Feynman picture an electron undergoes multiple scattering processes of type $e + \gamma \rightarrow e + \gamma$.

This approach reveals the rich particle phenomenology of the Standard Model; particles interacting with each other through the exchange of further particles. The complete set of particles represented by the Standard Model in this way can be seen in Table 2.2. The full phenomenology is far too rich to describe here in full, each particle interacts in certain ways, some are stable, some are unstable and decay. The particles have masses ranging from the eV scale for the neutrinos [8] to 175 GeV for the t or top-quark.

This diagrammatic picture is also of practical use. In many situations these series of integrals converge quickly. Computing the first few terms of the series then gives a fast calculation of an otherwise complex field theory problem. Physical processes for which the series of Feynman calculations is convergent are known as "perturbative". Processes involving the weak and electromagnetic forces are usually perturbative and therefore calculable to a high degree of precision. The best example of the application

Figure 2.1: An electron scattering from an electromagnetic potential, when viewed as an interaction between an electron field, shown here in blue and a potential field (left) and as an infinite series of Feynman diagrams (right).

Particles		Spin	Charge		
			Electric	Weak	Strong
Up-type Quarks	u, c, t	$\frac{1}{2}$	$+\frac{2}{3}$	Yes	Yes
Down-type Quarks	d, s, b	$\frac{1}{2}$	$-\frac{1}{3}$	Yes	Yes
Leptons	e, μ, τ	$\frac{1}{2}$	$+1$	Yes	No
Neutrinos	ν_e, ν_μ, ν_τ	$\frac{1}{2}$	0	Yes	No
Weak Bosons	$W^\pm$	1	± 1	Yes	No
	Z	1	0	Yes	No
Gluon	g	1	0	No	Yes
Photon	γ	1	0	No	No
Higgs Boson[†]	H	0	0	No	No

[†] Not yet observed

Table 2.2: The set of particles described by the Standard Model. The top 3 rows of fermions also have corresponding anti-particles. Neutrinos have anti-particles if they are Dirac particles, otherwise they are their own anti-particle (see Section 2.2).

of this technique is probably the anomalous magnetic moment of the electron which has been calculated and measured down to a precision of better than one part in a trillion [1, 2].

2.1 The Strong Force and QCD

The strong force becomes very strong at low energy scales. Below the GeV scale calculations involving the strong force start to become non-perturbative and the corresponding series do not converge. As such analytic calculations involving the strong force are challenging. However the strong force has a unique phenomenology which merits study. The study of the behaviour of strongly interacting particles is called "Quantum Chromodynamics" (QCD), so named because the SU(3) symmetry of the strong interaction gives rise to a "colour" charge which appears to have 3 degrees of freedom.

At experimentally obtainable energies, the strong force has a relatively large coupling constant compared to the other forces. Additionally the strong force carrier, the gluon, carries colour charge and can self-interact (i.e. processes like $g \to gg$ and $gg \to gg$ occur). These two features combined mean that processes involving quarks and gluons tend to produce large numbers of radiated additional particles compared to electromagnetic processes.

The strong force also exhibits interesting distance behaviour. All the other forces become weaker at larger distances. However, the self-interaction of the gluon field means the strong force effectively becomes stronger as the separation between charges increases.

One consequence of this distance behaviour is "confinement". Quarks and gluons, which interact via the strong force, are said to be "confined". They cannot be observed in isolation. This manifests itself in different ways at different energy scales. At low energies, the force exerted by the gluon field leads to a rich spectrum of "hadrons", bound states of quark anti-quark pairs or 3 quarks are observed, known as "mesons" and "baryons" respectively. The proton, which composes a significant part of all everyday matter is an example of a hadron which results from this confinement and is also the only stable hadron. Other hadrons display a wide range of decay behaviours with

lifetimes ranging from several minutes to less than 10^{-16} s. Although bound states involving more quarks or some mixture of gluons are not excluded, none have yet been observed.

At high energies, the strong force is no longer able to bind quarks into hadrons. In interactions above the GeV scale, the quarks start to behave increasingly like free charges, a concept known as "asymptotic freedom". However, although the gluon field is unable to bind the high-energy quarks, they are still confined. As the separation between a pair of strong charges increases, the distance behaviour of the strong force means the potential stored in the connecting gluon field rises rapidly. Eventually the potential is sufficient that additional quarks and gluons are radiated, screening the original charges from each other. This process of generating additional strongly interacting particles in high energy interactions is known as "fragmentation". The system will continue to fragment until the resulting particles have masses below the GeV scale, at which point the strong force will again bind them into hadrons in a process known as "hadronisation".

One variation on this sequence which is notable for this thesis is the production of heavy quarks in QCD processes. The weak eigenstates of the quark fields differ from the flavour eigenstates by a small rotation, known as the "CKM matrix". The consequence of this is that quarks can change to a flavour from a different generation in some weak interactions (specifically those involving W bosons). Because this process exists, heavy quarks are unstable. Top or t-quarks are the heaviest flavour and decay extremely rapidly, before hadronisation can take place, primarily to a W boson and the next heaviest flavour, a bottom or b-quark. Bottom quarks have longer lifetimes $O(10^{-12}$ s$)$ and do hadronise. At collider experiments b-hadrons can be produced with significant Lorentz boosts. This moderately long lifetime in combination with time dilation from being boosted means that some b-hadrons can travel macroscopic distances $O(1$ mm$)$.

2.2   Limitations of the Standard Model

The Standard Model is exactly as the name suggests, a model; it is not a fundamental description of the mechanics of the Universe. It is constructed largely empirically and as such its descriptive power is generally limited by the available experimental data.

One of the most glaring limitations of the Standard Model is the lack of a description of gravity. There is no known renormalisable quantum field theory which describes gravity in a way compatible with general relativity. The unification of gravity with the other three forces will probably require a fundamentally different treatment.

In the shorter-term the description of the neutrino sector is currently challenged by experimental evidence. Neutrinos are abundant but only interact via the weak force and are therefore very hard to detect. In the Standard Model they are defined as massless, however recent experimental evidence has suggested that they do indeed have masses, albeit far smaller than the other particles. The exact mechanism by which they obtain mass is not known, they may be Dirac particles like the other fermions. However the small mass scale could be a hint that neutrinos have Majorana mass [9] which would mean they are their own anti-particle.

The theoretical appeal of Majorana masses comes from the Seesaw Mechanism [10, 11] where the neutrinos naturally have Dirac masses similar to the charged leptons but a Majorana mass close to the Planck scale. This scenario naturally generates eV-scale left-handed neutrinos and while placing the unobserved right-handed neutrinos at masses far above current experimental limits [12]. Searches are currently underway for a process called neutrinoless double-beta decay, which is only possible if neutrinos have Majorana mass.

2.3   Electroweak Symmetry Breaking

The mechanism of "electroweak symmetry breaking" within the Standard Model is also currently an open problem. Excluding the Higgs field, the formulation of the Standard Model implies that both the weak and electromagnetic forces have similar properties. Both appear to be mediated by massless gauge bosons which in turn implies that they should have infinite range. In reality this symmetry between the electromagnetic and weak forces does not exist. While the electromagnetic force actually has infinite range and a massless gauge boson, the weak force is mediated by massive gauge bosons and has a correspondingly short range.

The electroweak symmetry cannot be broken explicitly in the definition of the forces without compromising gauge-invariance. Therefore it must be broken in some other

more indirect way. In the Standard Model this is achieved via the Higgs mechanism [13, 14, 15]. This postulates the addition of a further SU(2) doublet field ϕ to the Standard Model with an associated Lagrangian:

$$\mathcal{L}_{Higgs} = (D_\mu\phi)^\dagger(D_\mu\phi) - \mu^2\phi^\dagger\phi - \lambda(\phi^\dagger\phi)^2$$

$$\text{where} \quad D_\mu = \delta_\mu + \frac{1}{2}ig\vec{\tau}\vec{W}_\mu + \frac{1}{2}ig'YB_\mu.$$

The first term of the Lagrangian is the kinetic part, in terms of the covariant derivative D_μ. The second and third terms define the form of the Higgs potential which is quadratic in $\phi^\dagger\phi$ with the real coefficients μ^2 and λ. In order to obtain a potential with minima at a finite value of $\phi^\dagger\phi$ greater than zero, we must have $\lambda > 0$ and $\mu^2 < 0$.

By expanding the definition of the covariant derivative, we can see that the kinetic part of the Lagrangian contains terms with forms like $\phi^\dagger\phi B_\mu^\dagger B_\mu$ and we have just declared that in the vacuum $\phi^\dagger\phi$ is a non-zero constant. The fields corresponding to the bosons have been given what is effectively a mass term. Explicit couplings to the field are also introduced to give masses to fermions in a similar manner.

However, in introducing the ϕ field to the Standard Model, we have added additional degrees of freedom. Since the field interacts with massive particles, it will be perturbed from the vacuum minimum when mass is present. Exploring the effects of a small perturbation of ϕ on $\mathcal{L}_{Higgs}$ we find that we have introduced terms which correspond to a new massive scalar boson with mass $\sqrt{2\mu^2}$. This is the Higgs boson. The discovery of such a new scalar boson at the LHC would be very strong evidence for the correctness of the Higgs mechanism.

Theoretically, masses for a Higgs boson of $O(10^2$ GeV$)$ around the electroweak scale are preferred [3, 16]. If the Higgs boson has a mass much above 1 TeV then processes such as vector boson scattering begin to violate unitarity [17] and the boson itself has such a short lifetime that it ceases to behave as a particle. Experimentally, the LEP experiments have directly excluded the presence of a Higgs boson with mass of less than 114.4 GeV with 95% confidence [18]. The Tevatron experiments are also working to expand their current limit, to date having excluded Higgs masses between 158 GeV

and 175 GeV at the 95% confidence level [19].

Further hints as to the likely mass of a Higgs boson come from existing precision electroweak measurements. A theoretical Higgs boson would introduce measurable higher-order corrections to the masses of the top quark and W and Z bosons. The latest measurements of these quantities suggest a Higgs mass below the current experimental limit of 114 GeV and limit the mass of a possible Higgs boson at 186 GeV at 95% confidence [3]. This also implies that a Higgs boson is most likely to be found near the lower experimental limit. In fact the LEP experiments did observe a small excess of events, consistent with a Higgs boson with mass of 115 GeV, however the result was not statistically significant [18].

Of course all of these limits include some degree of assumptions about the properties of a possible Higgs boson. For example the introduction of further particles in Supersymmetry (SUSY) models can change phenomenology at the electroweak scale drastically [20, 21]. In SUSY the Higgs sector is expanded to include multiple Higgs bosons, each with different masses and couplings to fermions. Depending on the SUSY parameters, this can allow one or more of the bosons to have masses in the regions excluded by current experimental searches for a Standard Model Higgs. More generally, any model which introduces additional massive particles at or below the Electroweak scale will alter the phenomenology of the Higgs sector to some degree.

The one certainty is that the Standard Model Higgs mechanism must manifest itself at some energy below 1 TeV or be falsified. Therefore in order to test the nature of electroweak symmetry breaking it is extremely desirable to perform experiments at TeV energy scales.

Chapter 3

The ATLAS Experiment

Since the physics required to explain electroweak symmetry breaking only manifests itself at energy scales $O(10^2 \text{ GeV})$ and above, we must find some way to observe particles interacting at these high energy scales.

3.1   The LHC and Accelerator Complex

The Large Hadron Collider (LHC) at the European Centre for Nuclear Research (CERN) in Geneva, Switzerland takes the relatively direct approach of accelerating groups, or "bunches" of protons up to high energies and colliding them. In these collisions the quarks and gluons which compose the protons interact.

At the LHC a pair of proton beams is provided by a chain of accelerators. Firstly, LINAC2 is a linear accelerator using a duoplasmatron ion source to deliver pulses of 50 MeV protons. These protons are further accelerated by a set of circular synchrotron accelerators, the Proton Synchrotron Booster (PSB), Proton Synchrotron (PS) and Super Proton Synchrotron (SPS) to energies of 1.4 GeV, 26 GeV and 450 GeV respectively [22].

The LHC itself is also a synchrotron accelerator, with a circumference of 27 km, capable of accelerating protons to an energy of 7 TeV. The main technical difference between the LHC and earlier synchrotrons is that in order to achieve such high energies with reasonable power consumption, superconducting bending magnets are required.

The bunches of protons inside the LHC circulate continuously in opposite directions, crossing at a set of fixed points around the ring. Collisions take place at these points

allowing experiments to be constructed to observe the resulting interactions. ATLAS is one such experiment [23] depicted in Figure 3.1.

A typical proton-proton collision event is a complex process. As mentioned previously, protons are composite objects with complex internal structures. To first order, it can be said that a pair of individual quarks or gluons within the incoming protons exchange momentum in a "hard process". However in addition to this the remnants of the protons are still present close to the direction of the beam and all the incoming and outgoing quarks and gluons can also emit additional radiation. The particles from these additional processes are known as the "underlying event" which obscures the hard process being studied.

Many of the processes which give some insight into new physics happen extremely rarely in proton collisions. The rate of processes is described in terms of cross-sections which are a measure of effective scattering target area. The number n of events expected from a particular process with cross-section σ is dependent on the instantaneous luminosity of the collider L:

$$\frac{dn}{dt} = L \times \sigma = f n_b \frac{n_p^2}{A} \sigma.$$

Where instantaneous luminosity is defined for a machine like the LHC with symmetric beams in terms of f the rotation frequency of the beams, n_b the number of bunches per beam, n_p the number of protons per bunch and A the transverse area of the beam. Cross-sections are frequently defined in units of barns, which correspond to 10^{-24} cm^2. The instantaneous luminosity L is in units of inverse barns per second (b^{-1}s^{-1}). Generally accelerators aim for the highest possible luminosity as this enables the study of rarer processes with smaller cross-sections.

When trying to increase luminosity, high numbers of protons per bunch are extremely desirable due to the fact that $L \propto n_p^2$. Similarly small values of A are desirable. However, when n_p^2/A is very large, each bunch crossing tends to produce multiple inelastic proton-proton scatters. These additional interactions are termed "in-time pile-up". Again, high values of n_b are desirable as $L \propto n_b$. However, once again there is a cost because as n_b rises, the time between collisions at an interaction point falls. For values of n_b as high as used at the LHC particles from other bunch crossings can contaminate

observations of a collision. This is known as "out-of-time pile-up". Both effects produce additional particles which can obscure interesting physics.

In the currently planned run, the LHC aims to deliver up to 100 fb^{-1} of collisions at a centre-of-mass energy of 14 TeV. This will enable the study of processes with cross-sections as small as a few fb at these energies. In order to achieve this the beams circulating within the LHC will ultimately be composed of bunches of approximately 10^{11} protons. When colliding, these bunches will be squeezed to a transverse diameter of around 20 μm. As a result at full luminosity each interesting event at the LHC is expected to be accompanied by $O(20)$ in-time pile-up events. In the ultimate configuration the beams have 2808 bunches each. Due to the high energies the beams circulate at $0.999999991c$, leading to a revolution frequency of around 11kHz. This allows delivery of instantaneous luminosities of up to 10^{34} cm^{-2}s^{-1} or 10 nb s^{-1}. Beams circulate in the LHC for around 12 hours at a time, ultimately limited by the amount of scattering taking place in collisions. Running the machine for a period of several years will then allow the required integrated luminosities to be achieved.

In 2010, the LHC provided proton-proton collisions for the first time. Due to concerns over the safety of the magnet systems, the beams were only accelerated to 3.5 TeV. At this energy the LHC delivered approximately 48 pb^{-1} of integrated luminosity in 2010.

3.2 The ATLAS Detector

The interactions of interest occur over very small physical scales, of the order of 10^{-20} s in time[1] and 10^{-25} m spatially[2]. There is currently no viable detector technology capable of directly observing either the spatial or temporal evolution of the particle interactions themselves at these scales. The most powerful microscopes available today [24] are only capable of exploring structure at the atomic scale (around 10^{-10} m) and even then only for relatively static structures.

As a result indirect observation techniques are applied. As described in Chapter 2 some of the particles produced in the collisions are fairly stable and can travel distances of the order of 10 m through a detector. Others are less stable and undergo further

---

[1]Lifetime of a 120 GeV Standard Model Higgs boson is 2.2×10^{-22} s

[2]de Broglie wavelength for a probe with 100 GeV momentum is 4.1×10^{-26} m

Figure 3.1: The ATLAS detector.

processes resulting in stable particles. By either mechanism, the high energy interaction at a very small distance scale produces a set of particles which are easier to detect and in many cases contain information about the fundamental processes that took place.

Generally speaking, the aim is to measure the energy, momentum and type of as many of the particles produced in a collision as possible. However, these particles have widely varying properties; some are charged, while others are neutral; some only travel a few millimetres in a vacuum, others can penetrate several metres of lead.

As a result the ATLAS detector is a large construction [23]. It is cylindrical, approximately 44 m long and 25 m in diameter, weighing around 7,000 tonnes. It is also complex, composed of many different subsystems designed to help observe different types of particle. There are three main parts of ATLAS. The innermost part is the inner detector, which samples the position of charged particles at several points as they pass through under the influence of a magnetic field. Outside this sits the calorimetry which stops most classes of particles and in the process measures their energy. Finally, the muon spectrometer samples the position of muons, again under a magnetic field.

ATLAS is typically described in Cartesian coordinates with the z-axis aligned along the beam, positive z pointing towards Geneva, negative z towards the Jura mountains. The y-axis is vertical while the x-axis covers the remaining transverse dimension. Often

positions within the detector are referred to using polar and azimuthal angle from the beam, denoted θ and ϕ respectively. It is frequently useful in collider experiments to use the quantity "rapidity", which depends on the amount of energy and momentum in the z-axis a particle possesses:

$$y = \frac{1}{2}\ln\left(\frac{E + p_z}{E - p_z}\right).$$

The related quantity "pseudorapidity" is also frequently used:

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

For massless objects, the two quantities are identical. Pseudorapidity maps uniquely to polar angle θ and is therefore preferable when exploring detector effects. Rapidity on the other hand is more physically motivated. Many distributions of particle production are approximately flat in y. Differences in rapidity are also invariant under Lorentz boosts in z. Rapidity is therefore preferable for many physics uses. For an example comparison of the two variables, see Figure 8.3.

Since the beam is along the z-axis, the colliding particles have virtually no momentum in the x-y plane. The presence of particles with large amounts of momentum in this plane therefore implies a significant interaction took place. It is common to discuss momentum in this transverse plane, defined as:

$$p_\mathrm{T} = \sqrt{p_x^2 + p_y^2}.$$

The inner detector (Figure 3.2) is composed of three subsystems designed to detect the position of highly ionising charged particles. The innermost is the Pixel detector, composed of three cylinders and six rings of flat, silicon-based position detectors. The spatial resolution is an extremely precise 10 by 115 μm. The Pixel detector begins approximately 4.5 cm from the beam and extends to a radius of around 24 cm and z of 65 cm. Being extremely precise and extremely close to the beam assists in the detection of particles, such as b- and c-hadrons which have a measurable decay length of the order of a few mm.

Figure 3.2: The ATLAS inner detector.

This is surrounded by the "semiconductor tracker" (SCT). This is a silicon microstrip detector and therefore has excellent resolution along one axis, with an intrinsic accuracy of 17 by 580 μm. Alternating layers of the silicon are angled slightly with respect to each other such that particles can be located more precisely than the long axis of the strips themselves. The SCT consists of 4 such pairs of cylindrical layers in the central region and 9 disks at either end. It is approximately 0.5 m in radius and extends to 3 m in z.

The final component of the inner detector is the "transition radiation tracker" or TRT, composed of gas drift tubes aligned along the beam axis in the central region. As such they provide no z information but have an intrinsic R-ϕ resolution of approximately 130 μm. At the ends the tubes are arranged in a fan layout, providing ϕ-z information only. The TRT coverage extends to around 1 m in radius and 3 m in z.

The TRT also provides some degree of particle identification. Different gases circulate inside and outside the drift tubes. As particles transition between the two different media, they emit "transition radiation", the quantity of which is strongly dependent on the Lorentz boost of the particle. This allows some degree of discrimination between electrons and heavier pions.

The entire inner detector is embedded in a superconducting solenoid magnet pro-

viding a 2 T field aligned with the beam axis. Pattern recognition techniques can be used to find tracks in the sets of hits measured in the inner detector. These tracks have curvature due to the magnetic field which can then be used to measure the momentum and charge sign of the charged particles produced in a collision. Tracks with a common starting point can be used to identify points where a process of interest took place, known as a "vertex".

The next ATLAS detector system outside the inner detector is the calorimeter (Figure 3.3). As the name suggests it attempts to measure the energy of particles produced in a collision. This is done by sampling the energy produced when the particles shower in a medium. Hadronic particles such as neutrons penetrate much more material but shower more broadly than electromagnetic particles like photons or electrons. As a result the central part of the calorimeter system is composed of two main components. The inner component is the LAr (liquid argon) calorimeter, which has excellent angular granularity of down to 0.025 by 0.025 in $\delta\eta - \delta\phi$ for precisely locating the position of electromagnetic showers. This is surrounded by the Tile calorimeter, constructed from alternating steel and scintillating plates. Although the Tile calorimeter has significantly less granularity (0.01×0.01) than the LAr calorimeter it is much thicker and therefore suited to measure deep hadronic showers.

Outside the central region ($|z| \gtrsim 6.4$ m or $|\eta| \gtrsim 1.5$) the calorimeter is composed purely of LAr calorimetry. The electromagnetic and hadronic LAr systems reach as far as $|\eta| \simeq 3.1$. Even closer to the beam ($3.1 \lesssim |\eta| \lesssim 5.0$) is the forward calorimeter or "FCal", using the same LAr technology but with special geometry designed for the high radiation environment close to the beam.

The calorimetry systems of ATLAS form approximately a hollow cylinder, with outer radius of $\simeq 4.25$ m from the beam and extending to $\simeq 6.1$ m in z. Outside this is the outermost component of ATLAS, the muon spectrometer (Figure 3.4). Muons are extremely weakly ionising particles and as such they usually pass through the ATLAS inner detector and calorimeter without interacting significantly. The muon system, like the inner detector, records the position of charged particles traversing the detector elements under the influence of a toroidal magnetic system. However since muons ionize so weakly, the muon system covers a much larger volume, extending to a radius of $\simeq 12$ m and $z \simeq 20$ m.

Figure 3.3: The ATLAS calorimeter.

The muon system is composed of four technologies, Monitored Drift Tubes (MDT), Cathode Strip Chambers (CSC), Resistive Plate Chambers (RPC) and Thin Gap Chambers (TGC). The MDT and CSC provide precision tracking, primarily in the η coordinate; while the RPC and TGC provide triggering and additional ϕ measurements. Again, pattern recognition techniques can be used to identify tracks consistent with the passage of a muon. The muon system is also embedded within a magnetic field, so charge and momentum information can be extracted from the curvature of the tracks.

3.3   Particle Identification

In addition to measuring energy, momentum and charge, the combination of all the ATLAS systems allows for the identification of all stable Standard Model particles. Electrons can be identified by the presence of a track pointing at energy deposited in the electromagnetic calorimeter, supplemented in some cases by the capabilities of the TRT outlined above. Photons are identified as energy in the electromagnetic calorimeter without a pointing track. Muons are identified as tracks in the muon

Figure 3.4: The ATLAS muon system.

spectrometer. Hadrons are detected by the hadronic calorimeter using a clustering procedure, which will be covered in the next chapter.

The presence of neutrinos can be inferred by summing all the energy in the calorimeter. Neutrinos virtually always leave the detector without interacting. Normally the sum of all energy in an event should be zero in the transverse $(x - y)$ plane due to conservation of momentum and the fact that the beams have essentially zero transverse momentum. If the sum is not zero, this missing E_T can be used to infer the presence of a neutrino and measure its transverse momentum.

As discussed in Chapter 2, hadrons containing b-quarks, or "b-hadrons" can have relatively long lifetimes and can travel distances on the order of 1mm before decaying. The presence of b-hadrons can therefore be inferred if tracks are found in the inner detector which originate around this distance from the beam axis.

3.4   Trigger and Computing

With bunch crossings occurring at rates of 40 MHz, managing the large amounts of data produced by the detector is extremely challenging. Most of the collisions do not produce events of interest experimentally. A multi-level trigger system is used to reduce the rate of events being written to disk to around 200 Hz of the most interesting ones [4].

The ATLAS trigger system is composed of three levels. Level 1 is based on dedicated hardware which identifies regions of interest in the calorimeter or muon system. The inner detector is not used at level 1. If the event passes the level 1 trigger, then the regions of interest are passed to level 2, which consists of software running on general purpose PCs. The level 2 trigger has access to the full detector information, including the inner detector and runs reduced versions of the full reconstruction procedure in the regions of interest. Events identified as being of interest at level 2 are passed on to the third level of the trigger, the "Event Filter", also software running on PCs. The Event Filter runs a set of algorithms very close to the full reconstruction and makes a final decision about whether to keep the event or not.

If the trigger system decides data from a collision should be kept, then the raw detector data is written to disk in a format known as Raw Data Object (RDO). The ATLAS offline software package "Athena" is then used to process this data, the output is called Event Summary Data (ESD). This data then undergoes a second processing step where some of the data not needed for analysis is removed, to produce smaller more manageable files known as Analysis Object Data (AOD). These files contain high level objects such as electron or muon candidates and are used for final analysis.

Huge quantities of data are produced by the ATLAS detector and as such a huge computing resource is required. ATLAS is served by a grid computing system spanning hundreds of institutes around the world. Data is moved from a central farm at CERN (termed Tier-0 of the grid) to around ten large scale computing farms in various nations (termed Tier-1 sites). From here data is propagated to a variety of computer farms at individual institutes, or Tier-2 sites. The aim then being that ATLAS users process data on many Tier-2 computer farms simultaneously, thus allowing storage and processing resources to be shared within ATLAS on a global scale.

3.5   Monte Carlo

The Standard Model enables many calculations of the values and differential distributions of various observable quantities. However collider experiments, by their nature, do not directly measure distributions. Instead each collision can be thought of as a single sampling. It can be useful to replicate this approach in theory, both for comparison with experimental expectations and to understand the response of the detector. To this end software packages known as "Monte Carlo event generators" (frequently just "Monte Carlo" or "MC") have been produced. These packages take our analytical understanding of the Standard Model and sample the appropriate distributions to produce individual hard interactions. Generally they also combine these with models of non-analytic processes such as proton structure, underlying event, fragmentation and hadronisation tuned to available experimental data. In this way, Monte Carlo event generators are able to produce events which are representative of the best guess of the events a collider will really produce. An especially powerful feature of Monte Carlo packages is the ability to change the underlying equations, to introduce some new physics for example and then observe the type of collisions which result.

Herwig [25, 26] and Pythia [27] are general purpose Monte Carlo packages capable of simulating a wide variety of different physical processes. They provide good descriptions of simple $2 \rightarrow 2$ scattering processes where two incoming particles interact to produce two outgoing particles. Proton structure is modelled by "Parton Distribution Functions" or PDFs which describe the particle and momentum content of a proton. Both generators provide models of underlying event, although Herwig is frequently used in conjunction with Jimmy [28], which provides a more realistic model than the built-in one. Parton showers assume that a tree of successive $1 \rightarrow 2$ DGLAP splittings [29, 30, 31] combined with an ordering procedure [26, 27] forms a good approximation for the fragmentation process. This approximation allows the packages to simulate the fragmentation of quarks and gluons produced in a hard process. Pythia implements the Lund String model [32] to describe the hadronisation process, while Herwig uses an alternative "clustering" model [33].

The main limitations of these packages tends to be that they assume all additional particles after the hard $2 \rightarrow 2$ scatter are modelled by the parton shower. While the parton shower model has been shown to be a good approximation for particles which are

soft or collinear with one of the two outgoing hard partons, this tends to break down for events where high transverse momentum particles are produced at large angles. This can be important when studying some processes and must be kept in mind when using these packages. More specialised Monte Carlo packages, such as ACERMC [34] and MADGRAPH [35] directly simulate $2 \to 3$, $2 \to 4$ or even higher order processes. This approach can give a more accurate description of the relevant events. However, it relies on complex theoretical calculations which are not always available and often has a cost in terms of processing time.

3.6   Detector Simulation and Performance

An important tool for understanding the expected performance of ATLAS is the AT-LAS detector simulation package [36]. Test events can be generated using one of the Monte Carlo packages described in Section 3.5. The ATLAS detector simulation uses GEANT4 [37, 38] and a very detailed description of the detector geometry to simulate the passage of produced particles through the detector. GEANT4 is a complete simulation of the detector material. It is capable of simulating many different low-energy physical processes such as the interaction of particles with detector elements and the decay of long-lived particles.

Although the full ATLAS detector simulation is extremely detailed and precise, it is extremely slow. Typically it can take $30 - 60$ minutes to simulate the response of the detector to a single event using a single $\sim$2 GHz processor. Most of this time is used simulating the calorimeter. This is due to the fact that particles arriving at the calorimeter produce large showers of secondary particles each of which must be traced through the detector. However the showering behaviour of particles is by now fairly well understood. An alternative ATLAS detector simulation package, known as Atlfast-II attempts to achieve better performance. It accomplishes this by parameterising the response of each calorimeter cell to incoming particles [39]. The inner detector and muon system are still handled by GEANT4 but the calorimeter no longer needs to be fully simulated. This approach reduces the required time per event to around $30 - 60$ seconds, dominated by the time required to fully simulate the inner detector. Of course this speed-up comes at the cost of some degree of realism but extensive testing has

found the parameterised calorimeter to be accurate to the percent level [40].

Finally there is the extremely fast Atlfast-I [41], which applies a parameterised smearing to the Monte Carlo particles. This is extremely fast, of the order of a few ms per event. However, unsurprisingly it is somewhat less realistic than either of the other two simulations.

Based on studies of simulated data, the ATLAS detector is expected to deliver comparable or better performance than the previous generations of experiments at the Tevatron, HERA and LEP accelerators [4, 42]. There is a huge variety of impressive performance figures, only a few relevant to this thesis will be mentioned here. The calorimeter provides jet energy resolution of $\simeq 60\%/\sqrt{E(\text{GeV})}$ which as mentioned previously is coupled with excellent granularity of down to 0.025×0.025 in $\eta - \phi$. Ultimately ATLAS aims to set an absolute jet energy scale with an uncertainty of 1% or lower. Energy resolution for electrons is predicted to be $\simeq 10\%/\sqrt{E(\text{GeV})}$.

Anticipated transverse momentum resolution for tracks can be written as:

$$\sigma_{1/p_\text{T}} = 3.4 \times 10^{-4} \text{ GeV}^{-1} \times (1 \oplus 44 \text{ GeV}/p_\text{T}).$$

For a track with p_T of 1 GeV this corresponds to a resolution of 15 MeV, degrading with increasing p_T. The inner detector hardware allows for vertex location with resolution better than 100 μm in all dimensions. Secondary vertices from b-hadron decays should ultimately be identified with an efficiency of 60% or better with a fake rate of only 1% out to $|\eta| < 2.5$.

The muon reconstruction offers p_T resolution of 3% for muons with transverse momentum of 100 GeV in the central region. This degrades to 10% for 1 TeV muons.

Chapter 4

Jets and Jet Algorithms

Many interesting physical processes involve the production of strongly interacting particles. As described in Chapter 2, quarks and gluons produced at high energies first undergo the largely perturbative process of "fragmentation", creating a further spray of quarks and gluons. These further particles then undergo a non-perturbative "hadronisation", becoming properly confined hadrons. This process is extremely complex and has aspects which cannot be calculated analytically. Any experiment attempting to probe strong processes at these energy scales must somehow understand and measure these quark and gluon production processes.

As described in Chapter 3, this task is further complicated by effects like underlying event and pile-up. At LHC energies, deposits of energy of order 20-30 GeV from underlying event will not be uncommon. At the design luminosity of the LHC the mean number of pile-up interactions per bunch crossing is around 20.

The individual hadrons themselves provide little information. To perform measurements related to quarks and gluons, the hadrons must be organised in some way which provides information about the original hard physical process. This is usually done with a "jet algorithm". A jet algorithm is a clustering algorithm which takes a set of particles and groups them into "jets". Properties of these jets such as energy and momentum can then be measured. In some sense this can be thought of as attempting to undo the effects of the fragmentation process to measure the original quark or gluon. However, in practice this is often not possible or even physically well defined. It is better to consider a jet algorithm as a way of usefully viewing the large amount of information in an event. Particles produced in high energy collisions can easily number

in the thousands and the process of clustering them into jets is not trivial.

4.1   A Brief History of Jets

Historically the idea of jets originates from observations of high energy cosmic rays in the 1940s and 50s. When sufficiently energetic cosmic rays interacted with nuclei in an emulsion plate, a shower or "star" of particles was recorded. As the momentum of the cosmic particles being observed increased, the showers became more complex and increasingly directional. It was in this context the term jet was first published in 1952 to describe this type of physics [43].

These experiments operated by finding tracks in emulsion plates by eye, aided by scanning machines. All tracks originating from a single observed interaction point were then grouped as a jet. Although very precise, the use of emulsion plates limited observations to a small number of events.

The next set of significant jet measurements required the development of multi-GeV colliders in the 1970s. Jet-like hadron production was observed in a collider for the first time at SPEAR in 1975 [44] leading to the discovery of the gluon in 3-jet events at PETRA in 1979 [45]. Both SPEAR and PETRA were electron-positron collision experiments. In electron-positron colliders, the number of hadrons in an event is usually small. These experiments classified jets by looking for hadron production in events with particular shape characteristics.

As collider energy and intensity rose so too did the complexity of hadronic events. The event shape picture of jets did persist [46, 47]. However it struggled to cope with events containing an arbitrary number of jets. It also had difficulties with the more challenging environment of hadron colliders at a time when these machines were becoming increasingly relevant.

Around 1980, driven by these difficulties, the first true "jet algorithms" were constructed [48, 49, 50, 51]. From these initial formulations, two distinct classes of algorithms emerged.

"Clustering" algorithms combine particles based on some distance measure until the remaining distances between objects exceed some threshold. The notable example in this class is the k_T algorithm [52]. Meanwhile "cone" algorithms group all particles

within some radius. A cone class algorithm was used in the discovery of the top quark in 1995 [53].

The top quark discovery also marked a shift in the focus of jet physics. Modern collider experiments are well above the hadronisation scale and the basic structure of perturbative QCD is considered well known. Jets are now studied primarily as part of other physics signatures and precision has become increasingly important. Simultaneously, precise higher-order theoretical calculations are now available for many physical processes. This has enabled some very precise tests of QCD; some examples being measurements of proton structure [54, 55] and jet shapes [56] at HERA. Usability in theoretical calculations has also become a serious concern for jet algorithms. If an algorithm cannot be easily implemented algrebraically and evaluated then it is difficult to make theoretical predictions.

While event-shape based approaches are rarely used in modern jet measurements, both the cone-type and clustering-type classes of algorithms are in active use today. The current state of the art cone-type algorithm is SISCone [57], which extends the cone algorithms JetClu [58] and Midpoint [59] used at the Tevatron to be both infrared-safe and seedless albeit at a moderate cost in performance. The main clustering-type algorithms in active use are related to the k_{T} algorithm including k_{T} itself, Cambridge-Aachen [60, 61] and anti-k_{T} [62].

These clustering-type algorithms are defined by two functions d_i and d_{ij} acting on a set of 4-vectors:

$$d_i = p_{\mathrm{T}i}^{2\alpha}$$

$$d_{ij} = min(p_{\mathrm{T}i}^{2\alpha}, p_{\mathrm{T}j}^{2\alpha}) \frac{\delta R_{ij}^2}{R^2}$$

where $\delta R_{ij}^2 = \sqrt{(\delta\phi_{ij}^2 + \delta y_{ij}^2)}$, $p_{\mathrm{T}i}$ is the p_{T} of the i^{th} 4-vector and R is a parameter.

Clustering proceeds by calculating d_i and d_{ij} for all i and j. These values are then sorted to find the smallest. If the smallest value is d_i, which can be considered as a distance from the beam, then the 4-vector i is declared to be a jet and removed from the clustering procedure. If the smallest value is d_{ij} then a new 4-vector is defined $k = i + j$

and replaces i and j. It is worth noting that while this algorithm would appear to take $O(n^3)$ time to evaluate, application of computational geometry techniques can reduce this to $O(n\log(n))$ [63].

Choosing $\alpha = 1$ yields the k_{T} algorithm. Choosing $\alpha = 0$ removes the p_{T} dependence to give purely angular ordered clustering (Cambridge-Aachen algorithm). While $\alpha = -1$ produces the anti-k_{T} algorithm, which yields surprisingly cone-like jets.

4.2   Jet Substructure

Overall, modern jet algorithms provide useful ways of grouping and combining a large amount of information into a small list of aggregate objects. However, reducing such a complex dataset in this way can destroy more subtle features which can be of use. Specifically, the flow of energy inside a jet can contain interesting physics. Any physics taking place well above the hadronisation scale but well below the p_{T} scale of the hard processes will tend to be lost by a jet algorithm but can be studied by looking at the internal structure of the jet.

At the LHC for the first time there will be many events where the electroweak scale is below the p_{T} scale. Heavy particles like W and Z bosons with masses around 100 GeV will be produced with p_{T} several times higher. If these particles decay to hadrons, they will be boosted close together in angle and will often be clustered into a single jet. However, the energy flow inside the jet contains additional information. This is therefore one situation where looking inside the jet can help to identify interesting physics.

This is the "jet substructure" approach, viewing jets not just as fixed objects but as complex composite entities. One of the main aims of this thesis is to demonstrate the capability of jet substructure methods to outperform traditional jet techniques in several important situations.

As mentioned in Section 4.1 jet shapes averaged over many events have been explored at previous colliders [56]. However the concept of making jet-by-jet decisions based on the internal configuration of a jet is novel.

4.3   Jets at ATLAS

Real calorimeters often cannot measure the individual hadrons produced in an inter-
action. A significant fraction of the hadrons are π^0s, which decay to photons before
reaching the detector. Others are scattered or absorbed in uninstrumented material.
Even then, the calorimeter is not an individual particle detector, it has finite spatial
resolution. If two particles arrive very close together, the calorimeter cannot distin-
guish their energy deposits as separate. To further complicate matters, the calorimeter
is not a perfect measuring device, it is subject to electronic noise and has a finite energy
resolution. The calorimeter responds differently to the dense showers initiated by elec-
tromagnetic particles like photons than the broader showers from hadronic particles.
This is all in addition to the experimental issues discussed above such as underlying
event and pile-up.

In this difficult environment it is important to provide a reliable input for use in jet
algorithms. The main technique used by ATLAS for this purpose is an algorithm known
as "topological clustering" [64]. This algorithm searches through all the cells until it
finds one with $E_{\mathrm{cell}} > 4\sigma_{\mathrm{noise}}$. Where σ_{noise} is the standard deviation of the distribution
of noise in that cell. Cells passing this cut are declared to be the "seed" for a cluster. At
this point all neighbouring cells with $E_{\mathrm{cell}} > 2\sigma_{\mathrm{noise}}$ are added to the cluster iteratively
until no more neighbouring cells qualify. One layer of all neighbouring cells are then
added to the cluster regardless of energy. Finally some tests are performed to see if
the distribution of energy within the cluster contains multiple local maxima. If this is
found to be the case, the cluster is split into multiple clusters.

Topological clustering aims to give the closest possible approximation to individual
particles with a calorimeter. It has the added benefit of automatically suppressing
uncorrelated calorimeter noise.

These clusters are calibrated in one of two ways. The first and simplest option is
to calibrate each cell based on the energy and density of energy within the cell. The
second option, targeted at long-term physics use, is to classify each cluster as either
electromagnetic or hadronic and then apply an appropriate energy calibration.

The clusters can then be built into jets with a jet algorithm. Further calibrations
are usually applied to the jets themselves to fix an absolute energy scale. There are
multiple approaches but the exact nature of these depends on the physics analysis. The

jet algorithm for most ATLAS analyses will be anti-$k_{\rm T}$ with $R = 0.4$ or 0.6.

4.4   Tagging b-hadrons in Jets

Identifying b-hadrons in jets (or "b-tagging") is an important part of many analyses at ATLAS. As mentioned in Chapter 2, b-hadrons have observable lifetimes and typically travel a few mm before decaying. Any charged particles produced in the decay create tracks not originating at the primary vertex. Analyses generally proceed by locating tracks which have momentum pointing within some $\delta R \simeq 0.4$ of the jet axis. There are then several approaches to identifying tracks originating from b-hadron decay:

- **Impact Parameter** - If many of the tracks do not cross the jet axis at the primary vertex, this implies the presence of a displaced decay.

- **Secondary Vertexing** - If several tracks intersect each other at a point other than the primary vertex this implies a displaced decay.

- **Lepton Tagging** - b-hadrons sometimes produces leptons when decaying. The presence of a lepton within a jet can sometimes therefore be used to identify a b-hadron.

- **Track Counting** - Jets containing b-hadron decays tend to contain a higher number of charged particles. Although crude this method is simple and therefore useful in very early data when other algorithms may not be so well understood.

Typically the best performance is achieved by using a combination of an impact parameter algorithm and a secondary vertexing algorithm. The algorithms can be combined either by defining an explicit function or training a neural network.

The exact b-tagging algorithms used at ATLAS will depend a great deal on experience with real data. However, the anticipated default b-tagging approach is to match tracks to anti-$k_{\rm T}$ jets found with an R-parameter of 0.4. These tracks are then processed by a 3-dimensional impact parameter algorithm and a secondary vertexing algorithm [42]. Commissioning the b-tagging algorithms at ATLAS will be an ongoing challenge; precise descriptions of the alignment and dead material in the detector simulation are required as are good understanding of efficiencies and fake rates. In 2011

ATLAS might expect to obtain an efficiency and fake rate of 60% and 1% respectively, with systematic uncertainties of 5% and 10% on these values [4].

Chapter 5

Vector Boson Scattering

5.1 The Electroweak Chiral Lagrangian

As discussed in Chapter 2, the mechanism of electroweak symmetry breaking in the Standard Model should manifest itself at the LHC. Additionally, unitarity constraints tell us that this mechanism must have a significant effect on vector boson scattering processes at the TeV scale. Studying vector boson scattering at the LHC is therefore guaranteed to provide some insight into this physics. Particularly here the focus will be the case where a new high-mass (> 300 GeV) particle or particles are introduced.

Since the nature of electroweak symmetry breaking is not yet known, it is useful to explore the sensitivity of different analyses in a generic model-independent way. The "Electroweak Chiral Lagrangian" (EWChL) is one such generic mechanism, a low-energy effective theory describing the scattering of chiral states of the electroweak W and Z bosons. It parameterises new physics as higher-order effective terms in a Lagrangian [65].

The lowest-order terms which are not constrained by existing measurements of trilinear gauge boson couplings [66] are the dimension-4 terms:

$$\mathcal{L}^{(4)} = a_4(\langle D_\mu U D^\nu U^\dagger \rangle)^2 + a_5(\langle D_\mu U D^\mu U^\dagger \rangle)^2 \tag{5.1}$$

where D_μ and D^μ are covariant derivatives of the field U.

These terms give rise to effective quartic couplings of the gauge bosons. Although higher-order terms may change the character of new physics somewhat, the parameters

a_4 and a_5 can be used to broadly model a wide variety of different scenarios [17].

However, such a lowest-order perturbation is only physically valid at energy scales well below that of any new physics. At the LHC we hope to probe up to the energy scale of electroweak symmetry breaking. Therefore a unitarisation procedure, in this case Padé [67], is applied which ensures that the theory is at least unitary at these energies.

The phenomenology of this theory at collider experiments is the scattering of (primarily longitudinally polarised) vector boson pairs. Applying the Padé unitarisation scheme results in an invariant mass spectrum of the scattered pairs which may or may not contain resonances, depending on the choice of a_4 and a_5, as shown in Figure 5.1.

There are a variety of possible pairs of W and Z bosons and their decay modes. The WW case is shown in Figure 5.2. Experimentally, modes where both bosons decay hadronically have the highest rates but are very challenging to extract from the QCD backgrounds. Modes where both bosons decay leptonically are very clean but have very low rates. In between is the semi-leptonic case, where events contain one leptonically decaying W or Z boson and one hadronically decaying W or Z boson. Particularly this analysis concentrates on the case where the signal contains a leptonically decaying W boson and one hadronically decaying boson. These events result in the detection of a single lepton, missing energy and jets.

The feasibility of this channel has been previously explored at the hadron-level [17] where it was found to compare favourably with non-substructure approaches and with the Atlfast-I fast simulation [68, 69] but never with a full detector simulation. Version 12.0.6 of the ATLAS software was used throughout this chapter.

5.2 Jet Substructure with the k_{T} Jet Algorithm

A key element of this analysis is that vector bosons produced at high invariant mass, such as from the decay of a high mass resonance will have significant transverse momentum. These bosons therefore have large Lorentz boosts in the lab frame when they decay, causing the decay products to be observed close together.

Normally a hadronically decaying vector boson would be identified as a pair of jets with an invariant mass close to that of the original boson. However, as described in

Figure 5.1: Resonances introduced by Padé unitarisation in different regions of parameter space. Taken from ATLAS CSC Book [4].

Chapter 4, when a highly boosted massive particle decays to hadrons in a detector like ATLAS, traditional jet reconstruction techniques start to fail. The separation between the decay products is small and they are often reconstructed as a single jet. An alternative and innovative approach to this analysis is presented in [17]. The key element of this analysis is to accept that weak boson decays will be reconstructed as single jets. The internal structure of these jets can then be examined to help distinguish between signal and background.

The first such technique is to calculate the invariant mass of the weak boson candidate jets. If a jet encapsulates all the decay products of a weak boson then by conservation of momentum and excepting experimental effects, the jet should have the same mass as the boson itself. Jets produced in purely QCD processes by contrast tend to have much lower masses. This variable therefore provides discrimination between signal and background.

Further information about the substructure of the jet may be extracted by analysis of the jet clustering procedure itself. The jets in this analysis are clustered using the k_T algorithm with an R-parameter of 0.6. As described in Section 4, k_T is a sequential

recombination algorithm where clustering is ordered in relative p_T, with the lowest first. Since the combinations are ordered in this way, the k_T distance of the last combination to take place represents the hardest splitting in the jet. From this splitting, a variable known as the y-scale can be calculated as shown in Figure 5.3. This scale is expected to be $O(m_{boson}/2)$ in jets originating from the decay of a weak boson but lower for jets from pure QCD processes.

The LHC experiments will be the first to make extensive use of these variables in a physics analysis. It was therefore important to evaluate the resolution of the ATLAS detector for such quantities. In the absence of sufficient data, this was done using simulated events. Samples of semi-leptonic WW signal events were simulated using both the full ATLAS detector simulation software and the Atlfast-I fast simulation. These events by definition only include one high-p_T hadronic object, the hadronically-decaying W boson produced in the scattering. As a result, selecting the highest p_T jet in each event yields a reasonably clean sample of highly boosted W bosons.

In each simulated event, the highest p_T jet was examined, if it had $p_T > 300$ GeV then the jet mass and y-scale were calculated. The same analysis was also calculated in hadron-level Monte Carlo events, providing a measure of the "true" values without any experimental effects. By comparing the values of the variables before and after simulation it was possible to determine the experimental resolution introduced by the detector as shown in Figure 5.4.

The presence of clear peaks in these plots suggests that ATLAS is indeed capable of extracting these variables from events. The peaks of these curves were fitted with a Gaussian distribution. The standard deviation (σ) parameter of this fit in the full ATLAS simulation was found to be 7.4±0.2 GeV for the jet mass and 12.3±0.3 GeV for the y-scale. This implies that for jets with p_T of around 200 GeV and mass of around m_W the ATLAS detector is capable of measuring mass with a precision of $\sim 10\%$ and y-scale with $\sim 20\%$ precision.

5.3 Analysis Procedure

Given that the efficacy of these variables had been established, an analysis making use of them could be constructed (again based on [17]). In this analysis, jets are

Figure 5.2: Feynman diagram showing the WW scattering signal process. The central vertex is the effective vertex described by the EWChL.

$$p_{\mathrm{T}c}^2 \times y_2 = \min(E_a^2, E_b^2)\theta_{ab}^2$$
$$y_{\mathrm{scale}} = p_{\mathrm{T}c}\sqrt{y_2}$$

Figure 5.3: Calculation of a y-value for a parent object c breaking into two children, a and b. y_2 is the dimensionless y-value. The y-scale has dimensions of energy.

Figure 5.4: Resolution of jet masses and y-scales as determined from simulation of the ATLAS detector.

again found with the k_{T} algorithm with an R-parameter of 0.6. A hadronic W boson candidate is defined as any jet with mass between 68.4 GeV and 97.2 GeV and a hadronic Z candidate between 68.7 GeV and 106.3 GeV. Hadronic boson candidates must additionally have a y-scale between 30 GeV and 100 GeV. The mass cut values are chosen as 2σ either side of the boson masses based on the previous fits.

A leptonic W boson candidate is defined by taking the highest p_{T} lepton in the event and combining with the $E_{\mathrm{T}}^{\mathrm{miss}}$ in the event. The z-component of the missing energy is unknown but can be constrained by assuming the system has the invariant mass of the W boson. This requires the solution of a quadratic equation. The W boson candidate is only accepted if the quadratic is soluble. If there is more than one solution, one is chosen at random.

A characteristic feature of vector boson scattering is the presence of "tag" jets, which represent the quarks that originally radiated the incoming vector bosons (see Figure 5.2). Presence of two tag-jets in an event is defined as a pair of jets, other than the hadronic boson candidate, with $p_{\mathrm{T}} > 10$ GeV, $E > 300$ GeV and a distance between them in pseudorapidity $|\delta\eta| > 5$.

Since the W and Z bosons in vector boson scattering processes carry no strong charge, these events tend to contain less additional radiation in the central region than QCD backgrounds. Therefore to additionally suppress such backgrounds, central jets

were used to veto events, with a central jet defined as any jet with $p_\mathrm{T} > 30$ GeV detected between the tag-jets in pseudorapidity except hadronic W boson candidate jets.

Finally, a top quark candidate was defined as being present in the event if any jet could be combined with a W boson candidate to form an invariant mass near the top mass, the range being masses between 130 GeV and 240 GeV.

Given these definitions, the analysis can be defined as follows:

1. Require 1 hadronic boson candidate with $p_\mathrm{T} > 200$ GeV and $|\eta| < 2$.

2. Require 1 leptonic W boson candidate with $p_\mathrm{T} > 200$ GeV and $|\eta| < 2$.

3. Require 2 tag-jets.

4. Require no top candidates.

5. Require no central jets.

6. Plot invariant mass of boson candidate pair.

5.4   Results

The efficacy of this analysis was evaluated by applying it to a number of samples of Monte Carlo events, with full ATLAS detector simulation applied. A list of these samples is given in Table 5.1. The signal samples are generated using a version of PYTHIA modified to incorporate the EWChL a_4 and a_5 scattering amplitudes. Several sets of amplitudes were chosen and the resulting resonances are shown in the table. The exception to this is the sample labelled PYTHIA-73, which is generated using PYTHIA process 73 (longitudinal WZ scattering) with MSTP(46)=5. This is a Technicolour motivated QCD-like model which uses Padé unitarisation to introduce a resonance which decays like a q' of the chosen mass.

The results of applying the above selection to a number of these samples of events can be seen in Tables 5.2 and 5.3. No significant differences in efficiency were found between the WW and WZ samples or the models with vector and scalar resonances. The invariant mass spectra of the vector boson pairs was plotted and can be seen in

Sample name	Generator	$\sigma \times$ B.R. (fb)
$qqWZ \to qql\nu jj$ 500 GeV	Pythia-73	83.9
$qqWZ \to qql\nu jj$ 800 GeV	Pythia-ChL	35.2
$qqWZ \to qql\nu jj$ 1.1 TeV	Pythia-ChL	12.3
$qqWW \to qql\nu jj$ 499 GeV (s)	Pythia-ChL	66.5
$qqWW \to qql\nu jj$ 821 GeV (s)	Pythia-ChL	27.5
$qqWW \to qql\nu jj$ 1134 GeV (s)	Pythia-ChL	17.0
$qqWW \to qql\nu jj$ 808 GeV (v)	Pythia-ChL	29.8
$qqWW \to qql\nu jj$ 1115 GeV (v)	Pythia-ChL	17.9
$qqWW \to qql\nu jj$ non-resonant	Pythia-ChL	10.0
$qqWZ \to qql\nu ll$ bckg	MadGraph	96
$qqZZ \to qq\nu\nu ll$ bckg	MadGraph	123
		σ (pb)
$W^+ +$ 4 jets, QCD diagrams	MadGraph	163.3
$W^+ +$ 4 jets, EW diagrams	MadGraph	1.76
$W^+ +$ 3 jets, QCD	MadGraph	6.08
$W^+ +$ 3 jets, EW	MadGraph	0.219
$t\bar{t}$	MC@NLO	833

Table 5.1: List of Monte Carlo samples with generators used and cross-sections. The bracketed s and v indicate samples with scalar or vector resonances.

Figure 5.5 for various different signal cases. Errors shown are due to Monte Carlo statistics and are taken as Gaussian, except for $t\bar{t}$ where a 1σ equivalent Poisson confidence interval is drawn. The spectra are not corrected for detector effects and therefore do not represent a physical cross-section measurement, such a correction is possible but beyond the scope of this work.

Unfortunately for technical reasons, the available Monte Carlo statistics are limited. This is especially true for the $t\bar{t}$ sample, where only a limit on the cross-section after selection could be set. The limited statistics also determines the wide binning of the plots. However, the results are still sufficiently firm to make some broad conclusions about sensitivity.

Cut	500 GeV Scalar Resonance		800 GeV Scalar Resonance		1.1 TeV Vector Resonance	
	Efficiency (%)	σ (fb)	Efficiency (%)	σ (fb)	Efficiency (%)	σ (fb)
Starting sample	–	66	–	28	–	18
$\equiv 1$ Hadronic W	32.1 ± 0.5 (34)	21 (23)	40.0 ± 0.5 (45)	11 (13)	39.5 ± 0.7 (43)	7.1 (7.8)
$\equiv 1$ Leptonic W	45.4 ± 0.9 (54)	9.6 (12)	48.5 ± 0.8 (57)	5.4 (7.1)	48.8 ± 1.2 (55)	3.5 (4.3)
p_T (Had. W) > 200 GeV	57.6 ± 1.3 (69)	5.5 (8.5)	88.2 ± 0.7 (90)	4.8 (6.4)	86.6 ± 1.1 (88)	3.0 (3.8)
$|\eta|$ (Had. W) < 2	91.9 ± 0.9 (93)	5.1 (7.9)	95.3 ± 0.5 (95)	4.6 (6.1)	93.4 ± 0.9 (92)	2.8 (3.5)
p_T (Lep. W) > 200 GeV	43.8 ± 1.8 (42)	2.2 (3.3)	91.3 ± 0.7 (89)	4.2 (5.4)	92.4 ± 1.0 (89)	2.6 (3.1)
$|\eta|$ (Lep. W) < 2	95.5 ± 1.1 (94)	2.1 (3.1)	95.3 ± 0.6 (95)	4.0 (5.1)	92.8 ± 1.0 (93)	2.4 (2.9)
$\equiv 2$ tag jets	32.0 ± 2.6 (37)	0.7 (1.1)	42.4 ± 1.3 (49)	1.7 (2.5)	43.7 ± 2.0 (55)	1.1 (1.6)
$\equiv 0$ top candidates	50.0 ± 5.0 (40)	0.3 (0.5)	52.0 ± 2.1 (41)	0.9 (1.0)	51.4 ± 3.0 (44)	0.5 (0.7)
Central jet veto	100.0 ± 0.0 (98)	0.3 (0.4)	96.7 ± 1.0 (97)	0.8 (1.0)	91.6 ± 2.3 (93)	0.5 (0.7)

Table 5.2: Efficiencies of the cuts for three different $qqWW \to qq\ell\nu qq$ signal samples defined relative to the previous cut. The numbers in brackets are from the fast simulation. Starting sample is after generator cuts.

Cut	Non-resonant Signal		$t\bar{t}$ Background		W+jets Backgrounds	
	Efficiency (%)	σ (fb)	Efficiency (%)	σ (fb)	Efficiency (%)	σ (fb)
Starting sample	–	10	–	450000	–	21365
$\equiv 1$ Hadronic W	38.0 ± 0.7 (41)	3.8 (4.1)	18.9 ± 0.1 (19)	85000 (84000)	8.3 ± 0.1 (9)	1760 (1820)
$\equiv 1$ Leptonic W	48.2 ± 1.1 (55)	1.8 (2.3)	22.1 ± 0.2 (29)	19000 (25000)	23.3 ± 0.7 (31)	410 (570)
p_{T} (Had. W) > 200 GeV	82.1 ± 1.3 (86)	1.5 (1.9)	16.8 ± 0.4 (20)	3200 (5000)	34.4 ± 1.7 (43)	140 (240)
$|\eta|$ (Had. W) < 2	94.4 ± 0.8 (94)	1.4 (1.8)	90.3 ± 0.7 (90)	2900 (4500)	80.1 ± 2.4 (77)	110 (190)
p_{T} (Lep. W) > 200 GeV	90.4 ± 1.1 (87)	1.3 (1.6)	34.5 ± 1.3 (29)	990 (1300)	48.5 ± 3.3 (40)	55 (75)
$|\eta|$ (Lep. W) < 2	96.0 ± 0.8 (96)	1.2 (1.5)	94.6 ± 1.0 (90)	930 (1200)	80.4 ± 3.9 (79)	44 (59)
$\equiv 2$ tag jets	45.1 ± 2.0 (54)	0.6 (0.8)	8.1 ± 1.3 (10)	76 (120)	13.9 ± 3.5 (22)	6 (13)
$\equiv 0$ top candidates	56.5 ± 3.0 (47)	0.3 (0.4)	7.9 ± 4.4 (2)	5 (2)	60.5 ± 13.1 (23)	4 (3)
Central jet veto	91.1 ± 2.3 (94)	0.3 (0.4)	< 50 (< 25)	< 5 (< 1)	84.9 ± 13.7 (91)	3 (3)

Table 5.3: Efficiencies of the cuts for one $qqWW \rightarrow qq l\nu qq$ signal sample and the backgrounds defined relative to the previous cut. The numbers in brackets are from the fast simulation. Starting sample is after generator cuts.

Figure 5.5: WW (top 4) and WZ (bottom 3) invariant mass spectra in the $\ell\nu\,j(j)$ semi-leptonic channel, showing the total W+jets and $t\bar{t}$ backgrounds and the signal for the three resonant signal samples and the continuum sample. The error bars reflect the uncertainty from the Monte Carlo statistics.

Given the presence of certain amounts of signal and background in the invariant mass plots it is possible to estimate the significance for a range of luminosities and therefore calculate the luminosity required for ATLAS to observe these signals as shown in Table 5.4. The significances have been calculated using the naïve measure $S/\sqrt{B}$, where S and B are the number of signal and background events respectively in a three bin window around the resonance peak.

These results should be considered with the limitations of the study in mind, chiefly the limited Monte Carlo statistics. It is however, possible to conclude from these results that for a 3σ observation of the most optimistic of these resonances in vector boson scattering will require of the order of 30 fb^{-1} of recorded luminosity at ATLAS.

No systematic effects have been considered here. The main systematic uncertainty is probably understanding of background cross-sections. From Monte Carlo alone there are uncertainties of the order 50% which would need to be reduced using data. However this is beyond the scope of this analysis and it is not clear what level these would ultimately reach. Other lesser systematics will include signal efficiency, luminosity and jet energy scale. However these are likely to be at the few percent level.

5.5 Conclusions

The overall significance estimates suggest that this is a relatively high luminosity measurement but that some models may be excludable with around 30fb^{-1} of integrated luminosity. The lowest mass resonances at around 500 GeV are harder to extract due to higher backgrounds. The sensitivity to TeV mass resonances is limited by the cross-

Process	Cross Section (fb)		Luminosity (fb^{-1})	
	Signal	Background	For 3σ	For 5σ
$WW/WZ \to \ell\nu\, jj$, 500 GeV	0.31 ± 0.05	0.79 ± 0.26	74	206
$WW/WZ \to \ell\nu\, jj$, 800 GeV	0.65 ± 0.04	0.87 ± 0.28	19	52
$WW/WZ \to \ell\nu\, jj$, 1.1 TeV	0.24 ± 0.03	0.46 ± 0.25	72	200

Table 5.4: Approximate signal and background cross sections expected after the analyses. Approximate values of the luminosity required for 3σ and 5σ significances are shown. The uncertainties are due to Monte Carlo statistics only.

section and ultimately ATLAS has greatest sensitivity to resonances at around 800 GeV with this analysis. A fuller determination of sensitivity would require higher Monte Carlo statistics and some treatment of systematic effects. Comparisons at the hadron-level [17] imply that this will be a greater reach than possible with non-substructure approaches.

Although the ability to establish sensitivity was somewhat limited, this study was the first use of jet substructure techniques on fully simulated ATLAS Monte Carlo. The resolution results show clearly that ATLAS has the ability to measure jet structure quantities. It also shows that jet structure can form part of an effective analysis strategy.

Chapter 6

Boosted HW/HZ with $H \to b\bar{b}$

As described in Chapter 2, in the Standard Model electroweak symmetry breaking is achieved via the Higgs mechanism which predicts the existence of a scalar Higgs boson. Discovering a scalar boson at a mass consistent with existing electroweak data would be a very strong indicator of the reality of the Higgs mechanism.

6.1 Low Mass Higgs Searches at ATLAS

The Higgs boson prefers to couple to heavier fermions. However due to kinematic constraints it is generally forbidden from decaying to particles with $m > m_H/2$. As a result, the phenomenology of a Higgs boson varies significantly with mass. Here we will focus on the low Higgs mass (115-130 GeV) part of the non-excluded parameter space. This is also the region favoured by current experimental evidence. In this range, the dominant decay is to pairs of b-quarks ($\sim 70\%$ of the time) with secondary decay modes including pairs of virtual W bosons ($\sim 13\%$), light hadrons ($\sim 10\%$), τ leptons ($\sim 7\%$) and rarely but cleanly to a pair of photons via a top quark loop ($\sim 0.2\%$) [16].

Circa 2007, the ATLAS approach to low mass Higgs discovery focussed on the $H \to \tau\tau$ and $H \to \gamma\gamma$ search channels. Both providing relatively clean leptonic signatures at the expense of low rates. Attempts to search for the Higgs boson in $H \to b\bar{b}$ decays had been tried and found to be uncompetetive both in the HW/HZ production mode [5] and in the $t\bar{t}H$ mode [4]. However, there are good reasons for wanting to observe $H \to b\bar{b}$, the simplest being that it accounts for the largest fraction of low mass Higgs events and therefore has the potential to enhance overall experimental sensitivity. However,

it is also important to note that a measurement of the fraction of Higgs bosons which decay to $b\bar{b}$ (the $H \to b\bar{b}$ branching ratio) will be essential for determining that any new resonance discovery is indeed a Higgs boson [70].

The original ATLAS approach to observing HW/HZ [5] failed due to a combination of poor signal-to-noise ratio and difficult to control systematics (see Figure 6.1). The analysis searched for an isolated lepton and a pair of b-tagged jets. It also applied vetoes on additional jets and leptons. The assumptions involved in the analysis were aggressive. As an example the jet veto required no additional jets with $p_\mathrm{T} > 15$ GeV and $|\eta| < 5.0$ which is extremely sensitive to effects like pile-up and models of underlying event. The analysis also assumed a Higgs mass of 100 GeV, which enhances the production cross-section. Although such a mass was possible at the time, it has since been ruled out by LEP2 [18]. Even given these assumptions, the analysis found:

> "... the extraction of a signal from $H \to b\bar{b}$ decays in the WH channel will
> be very difficult at the LHC, even under the most optimistic assumptions
> ..."

The main difficulty is control of systematics, the authors suggest the $Wb\bar{b}$ background shape would probably have to be determined from Monte Carlo. Although the statistical significance of observing a signal in this analysis is a respectable 3.3σ with 30 fb^{-1}, the introduction of an optimistic 5% uncertainty on background shapes reduces this to 1.9σ. Since this result there has only been limited study of HW/HZ production at ATLAS until the work presented below.

6.2   A New Approach

An alternative and novel approach is to study only the high p_T part of the HW/HZ process. By placing a cut requiring both the Higgs boson and weak boson to have at least 200 GeV in p_T, 95% of the cross-section is lost, however the remaining events have a number of useful properties:

- **Within Acceptance** - the detector only has b-tagging coverage out to $|\eta| < 2.5$, most high p_T events fall within this region.

Figure 6.1: Result of HW, $H \to b\bar{b}$ analysis for 30 fb^{-1} and $m_H = 100$ GeV showing background (dashed) and signal (solid) [5].

- **Simpler Topology** - the bosons are quite boosted and therefore their decay products tend to be close together and well separated from the rest of the event. Reconstructing these events relies on identifying pairs of decay products so this simplifies the analysis.

- **Reduced Background** - the background cross-sections fall somewhat faster with p_{T} than the signal.

- $Z \to \nu\bar{\nu}$ **Mode** - at high p_{T} it becomes possible to extract decays of the Z to missing energy.

6.3 $H \to b\bar{b}$ Candidate Identification

The key component of this analysis is identifying the hadronic decay of a boosted Higgs boson to two b-quarks. As mentioned above, the b-quarks tend to be separated by a small angle. A jet finding procedure is constructed which searches for deposits of energy with structure matching that expected from a Higgs decay.

The Cambridge-Aachen algorithm is used for the substructure analysis here. This algorithm clusters based on angular distance between objects only. With the k_{T}-based procedure described in Chapter 5, the final splitting was the hardest which allowed us to simply break the jet into two to explore substructure. However it meant that any soft radiation from underlying event or pile-up was included in our final heavy particle candidate. The angular ordering of Cambridge-Aachen means that when undoing the clustering the radiation furthest from the center of the jet will be removed first. This radiation tends to be exactly the soft radiation we do not want to keep. The Cambridge-Aachen algorithm therefore has useful properties for this kind of analysis. The downside is that we cannot simply take the first splitting but must somehow search the splittings for substructure.

First, the event is clustered with the Cambridge-Aachen algorithm with an R-parameter of R_{initial}, which will be around 1. Then the following procedure is applied to each jet j:

1. Undo the last clustering step of j to get two subjets j_1 and j_2 ordered such that $m_{j1} > m_{j2}$. If j cannot be unclustered (i.e. it is a single particle) it is not a suitable candidate, so stop.

2. If the splitting has $m_{j1}/m_j < \mu$ (large drop in mass) and $y_2 > y_{2\mathrm{cut}}$ (fairly symmetric) then go to step 4.

3. Otherwise redefine $j = j_1$ and go back to step 1.

4. Recluster the constituents of j with the Cambridge-Aachen algorithm with an R-parameter of $R_{b\bar{b}} = \min(0.3, \delta R_{j1,j2})$ finding n new subjets $s_1, s_2 \ldots s_n$.

5. Redefine $j = \displaystyle\sum_{i=1}^{\min(n,3)} s_i.$

Where $y_2 = \frac{\min(p_{tj1}^2, p_{tj2}^2)}{m_j^2} \delta R_{j1,j2}^2$ and $\delta R_{j1,j2} = \sqrt{\delta y_{j1,j2}^2 + \delta \phi_{j1,j2}^2}$.

Steps 1 to 3 attempt to identify a hard structure in the distribution of energy in the jet which would imply the decay of a heavy particle. The y_2 cut further helps by suppressing very asymmetric decays of the type favoured by QCD splittings. Steps 4 and 5 filter out some of the particles in the candidate jet, the aim being to retain particles relevant to the hard process while reducing the contribution from effects like underlying event and pile-up. Although this only has a small effect on background rejection it narrows the signal peak significantly. In this analysis, μ is taken to be 2/3, while $y_{2\text{cut}} = 0.09$. These values were chosen from a scan of the available parameter space [71].

Tagging of b-quarks is handled by counting how many of s_1, s_2 and s_3 contain b-mesons. A b-tagging efficiency b_{eff} and fake-rate b_{fake} are defined. b_{eff} is the probability of correctly tagging a b-meson experimentally, while b_{fake} is the probability of incorrectly tagging a jet as containing a b-meson when in fact none is present. In this analysis we will require that two of s_1, s_2 and s_3 contain b-mesons but rather than reject Higgs candidates probabilistically, they will be assigned a weight w_H based on the probability of being b-tagged correctly in an experiment:

$$
w_H = \begin{cases} b_{\text{fake}}^2 & \text{if none of } s_i \text{ contain a } b\text{-meson} \\ b_{\text{fake}} \times b_{\text{eff}} & \text{if 1 of } s_i \text{ contain a } b\text{-meson} \\ b_{\text{eff}}^2 & \text{if 2 or more of } s_i \text{ contain } b\text{-mesons} \end{cases}
$$

The correlations between the b-tagging of the separate subjets s_i are ignored. Although in reality they will be correlated for both physical and experimental reasons this is not easy to model without a realistic detector simulation.

6.4   Leptonic Vector Boson Identification

The leptonic part of the analysis proceeds by selecting events matching one of three lepton-based criteria representing respectively, the decay of a Z boson to stable leptons, the decay of a Z boson to neutrinos and the decay of a W boson to a single stable lepton and a single neutrino. Stable leptons are defined as electrons or muons produced with

a $p_\mathrm{T} > 30$ GeV with $|\eta| < 2.5$. Then events are categorised as passing one of three selections:

(a) An opposite sign stable lepton pair with an invariant mass between 80 and 100 GeV and $p_\mathrm{T} > p_\mathrm{Tcut}$.

(b) Missing transverse momentum $> p_\mathrm{Tcut}$.

(c) Missing transverse momentum > 30 GeV and a stable lepton consistent with a W boson with $p_\mathrm{T} > p_\mathrm{Tcut}$.

Typical values for p_Tcut are 200 - 300 GeV. These selections are designed to be achievable experimentally, being above likely relevant trigger thresholds.

6.5   Analysis Procedure

Given these definitions of hadronic and leptonic selection criteria, the complete analysis is as follows:

1. Require a leptonic vector boson candidate **(a)**, **(b)** or **(c)**.

2. Reject events with stable leptons (as defined in 6.4) other than those used in step 1.

3. For each hadronic $H \to b\bar{b}$ candidate:

    (i) Require candidate has $p_\mathrm{T} > p_\mathrm{Tcut}$.

    (ii) Require candidate has $|\eta| < 2.5$.

    (iii) For each additional light jet in the event with $p_\mathrm{T} > 50$ GeV modify $w_H = w_H * (1 - b_\mathrm{fake})$.

    (iv) For each additional b-tagged jet in the event with $p_\mathrm{T} > 50$ GeV modify $w_H = w_H * (1 - b_\mathrm{eff})$.

    (v) If event passed **(c)** require no additional jets with $p_\mathrm{T} > 25$ GeV.

    (vi) Fill a histogram with the mass of the candidate, weighted by w_H.

The second cut has the effect of forcing the three different leptonic selections to be exclusive, since they then require exactly 2, 0 and 1 stable leptons respectively. The cut on candidate η reflects the fact that experimental b-tagging capability tends not to extend beyond $|\eta| < 2.5$.

The modifications to the candidate weight w_H are equivalent to applying a veto on additional jets containing b-hadrons with $p_T > 50$ GeV. This helps to suppress backgrounds with additional hard b-hadrons such as $t\bar{t}$ at only a small cost in terms of signal efficiency.

The analysis was tested on Monte Carlo data samples generated using the HERWIG generator. A full set of samples is listed in Table 6.1. Statistics corresponding to 30 fb^{-1} of data were generated for all samples except dijets, where the large cross-section made this impractical. In this case, it was assumed that the efficiency of the hadronic and leptonic selections is uncorrelated.

Two analyses were tried, one took $R_{\text{initial}} = 0.7$ and $p_{\text{Tcut}} = 300$ GeV, the other used $R_{\text{initial}} = 1.2$ and $p_{\text{Tcut}} = 200$ GeV. The remaining $H \to b\bar{b}$ candidates passing each of the selections **(a)**, **(b)** and **(c)** can be seen in Figures 6.2 and 6.3 and also broken down into the individual samples in Figures 6.4 and 6.5. The overall significance is calculated as $S/\sqrt{B}$ based on the number of signal and background events within a given window. It is found to be somewhat better in the 200 GeV/$R = 1.2$ analysis where it approached 6σ. The stricter cuts of the 300 GeV/$R = 0.7$ analysis reduce the available signal but also strongly reduce some of the backgrounds leaving a significance of 5.5σ. Most notably $t\bar{t}$ is significantly suppressed in selection **(c)** by the higher p_T cut. If $t\bar{t}$ should prove to be more problematic than predicted, this stricter version of the analysis may be more useful.

In both of these analyses, it was assumed that $b_{\text{eff}} = 70\%$ and $b_{\text{fake}} = 1\%$. These were aggressive but not unrealistic choices based on the current understanding of ATLAS performance from simulation [4]. However, the effect of changing these parameters was also tested as can be seen in Figure 6.8(a). Unsurprisingly a decline in b_{eff} or an increase in b_{fake} both affect the significance negatively. The final analysis plots with less optimistic choices of $b_{\text{eff}} = 60\%$ and $b_{\text{fake}} = 2\%$ can be seen in Figure 6.6.

As mentioned previously, all these analyses assume a Higgs boson mass of 115 GeV. The effect of higher Higgs masses can be seen in Figure 6.8(b). At higher masses the

Process	Generator $p_{\mathrm{T}min}$	σ(pb)	Sample Size (fb^{-1})
HW ($m_H = 115$GeV)	none	1.190	1590
HW ($m_H = 120$GeV)	none	0.980	1530
HW ($m_H = 130$GeV)	none	0.623	2410
HZ ($m_H = 115$GeV)	none	0.680	2200
HZ ($m_H = 120$GeV)	none	0.562	2670
HZ ($m_H = 130$GeV)	none	0.358	4190
WW	none	65.9	75.8
WZ	none	25.7	195
ZZ	none	10.2	489
W+jets	150 GeV	390	89.5
Z+jets	150 GeV	160	248
$t\bar{t}$	150 GeV	429	46.1
Single top	none	161	61.5
Dijets	180 GeV	88600	0.598

Table 6.1: Table of Monte Carlo samples. All samples were generated with HERWIG with all decay modes active and a beam centre-of-mass energy of 14 TeV.

Higgs production cross-section falls. At the same time, the branching ratio $H \to b\bar{b}$ is falling. These two factors combine to reduce the effectiveness of the analysis at higher masses. The analysis plots with a Higgs mass of 130 GeV can be seen in Figure 6.7. For Higgs masses above 130 GeV it is expected that the analysis becomes ineffective rapidly.

6.6   Conclusions

Overall, it is possible to surmise that this is an extremely promising approach, offering significances of up to 6σ after 30 fb^{-1} of luminosity. Of course there are several aspects of the analysis that have not been fully treated here. Tagging the b-hadrons experimentally is a complex process, as is reconstructing the jet substructure. All of this will be dealt with more fully in the next chapter.
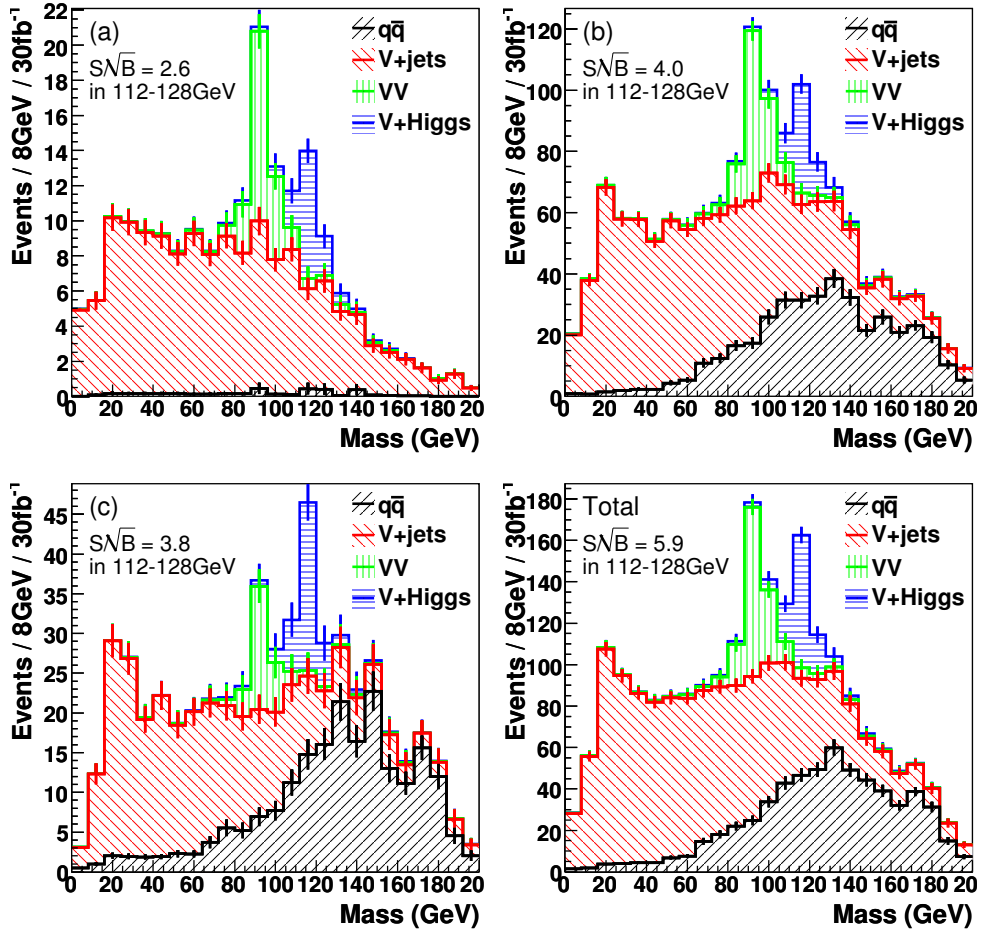
Figure 6.2: Mass of heavy particle candidates after jet substructure analysis with $R_{\mathrm{initial}} = 1.2$ and $p_{\mathrm{Tcut}} = 200$ GeV on events passing a leptonic selection with samples grouped. Errors are the Gaussian uncertainties due to Monte Carlo statistics.
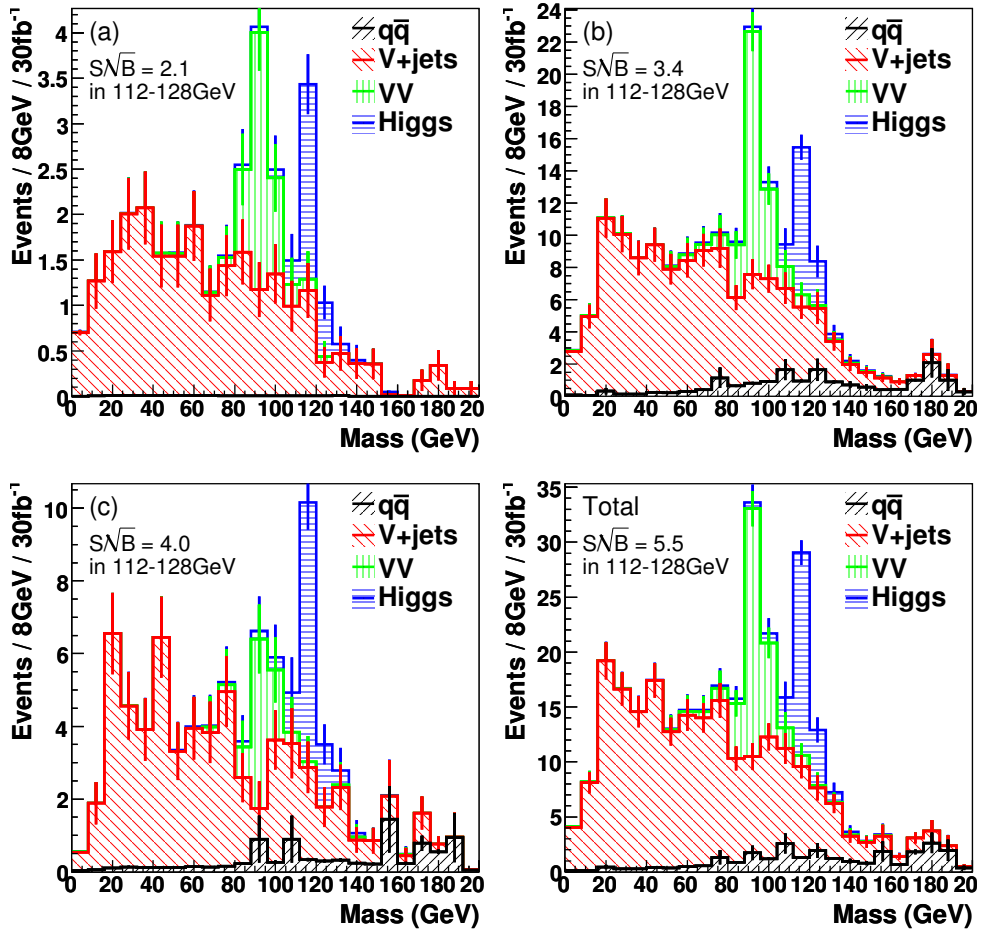
Figure 6.3: Mass of heavy particle candidates after jet substructure analysis with $R_{\mathrm{initial}} = 0.7$ and $p_{\mathrm{Tcut}} = 300$ GeV on events passing a leptonic selection with samples grouped. Errors are the Gaussian uncertainties due to Monte Carlo statistics.

Figure 6.4: Mass of heavy particle candidates after jet substructure analysis with $R_{\mathrm{initial}} = 1.2$ and $p_{\mathrm{Tcut}} = 200$ GeV on events passing a leptonic selection. Errors are the Gaussian uncertainties due to Monte Carlo statistics.
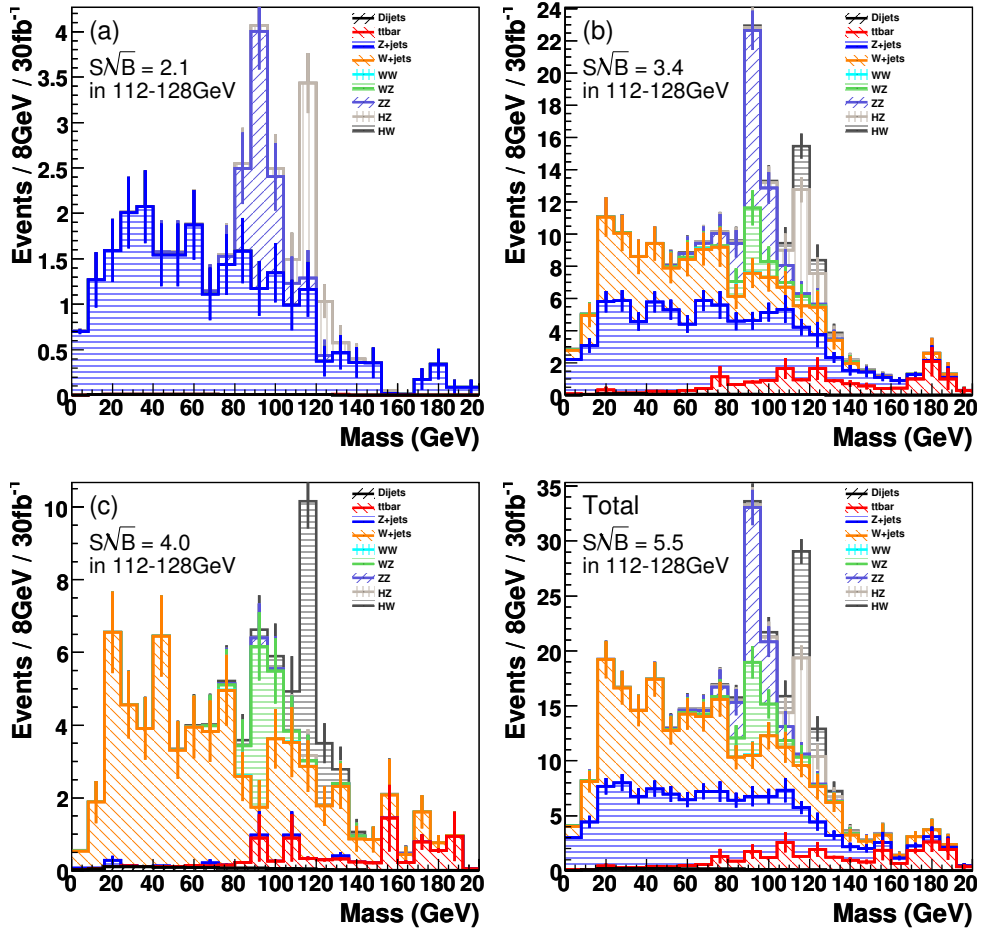
Figure 6.5: Mass of heavy particle candidates after jet substructure analysis with $R_{\text{initial}} = 0.7$ and $p_{\text{Tcut}} = 300$ GeV on events passing a leptonic selection. Errors are the Gaussian uncertainties due to Monte Carlo statistics.

Figure 6.6: Mass of heavy particle candidates as in Figure 6.2 but with the more pessimistic b-tagging choices of $b_{\text{eff}} = 60\%$ and $b_{\text{fake}} = 2\%$.

Figure 6.7: Mass of heavy particle candidates as in Figure 6.2 but with $m_H = 130$ GeV.
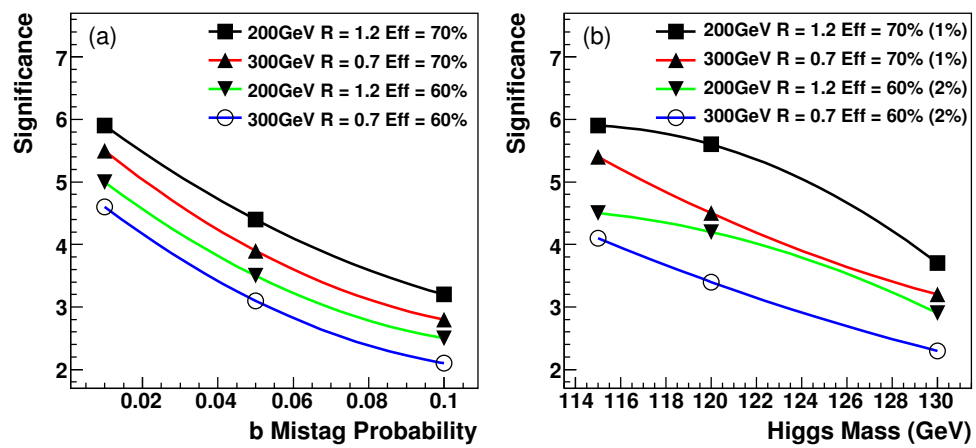
Figure 6.8: Effect of parameter choices on analysis significance for (a) b-tagging parameters and (b) Higgs mass.

Chapter 7

Observing Boosted HW/HZ at ATLAS

In order to fully evaluate the effectiveness of the analysis presented in Chapter 6, experimental effects must be taken into account. This is accomplished by producing samples of Monte Carlo events similar to those used in Chapter 6 and applying a realistic simulation of the ATLAS detector. Release 14.1.0 of the ATLAS software was used throughout this chapter.

A number of Monte Carlo samples were simulated in order to evaluate the effectiveness of the analysis at the detector level, see Table 7.1. The main difference from Chapter 6 is that because simulation is so resource intensive, filters have been applied to ensure only events with relevant final states were simulated. Most of the samples used were generated with the HERWIG generator with JIMMY for the underlying event. The single-top sample was generated with AcerMC and HERWIG. All samples assume a center-of-mass energy of 14 TeV. These samples were then simulated with the Atlfast-II detector simulation package, as described in Chapter 3.

Due to the relative difficulty of simulating sufficient quantities of QCD dijet events and their expected low contribution based on the hadron-level analysis, these processes were not considered here. The requirement of a leptonically decaying vector boson should be sufficient to suppress these backgrounds however this is a topic which will require further study in any real analysis.

Process	Generator p_{Tmin}	σ(pb)	Filter	Filter efficiency	Sample Size (fb^{-1})
HW	none	0.953 pb	$p_{TH} > 150$ GeV, $p_{TW} > 100$ GeV, $p_{Te,\mu} > 15$ GeV	$4.38 \pm 0.04\%$	624
HW	none	0.953 pb	$p_{TH} > 150$ GeV, $E_T^{miss} > 100$ GeV	$4.39 \pm 0.04\%$	622
HZ	none	0.545 pb	$p_{TH} > 150$ GeV, $p_{TW} > 100$ GeV, $p_{Te,\mu} > 15$ GeV	$3.51 \pm 0.02\%$	1900
HZ	none	0.545 pb	$p_{TH} > 150$ GeV, $E_T^{miss} > 100$ GeV	$2.34 \pm 0.03\%$	3760
WW	150 GeV	2.059 pb	$p_{Te,\mu} > 15$ GeV	$40.7 \pm 0.4\%$	57.9
WW	150 GeV	2.059 pb	$E_T^{miss} > 100$ GeV	$30.8 \pm 0.4\%$	76.7
ZZ	150 GeV	0.440 pb	$p_{Te,\mu} > 15$ GeV	$61.2 \pm 0.2\%$	830
ZZ	150 GeV	0.440 pb	$E_T^{miss} > 100$ GeV	$47.9 \pm 0.5\%$	610
WZ	150 GeV	0.96 pb	$p_{Te,\mu} > 15$ GeV	$33.6 \pm 0.2\%$	76.5
WZ	150 GeV	0.96 pb	$E_T^{miss} > 100$ GeV	$33.3 \pm 0.5\%$	77.0
$t\bar{t}$	150 GeV	112.7 pb	$p_{Te,\mu} > 20$ GeV	$47.5 \pm 0.2\%$	31.6
$t\bar{t}$	150 GeV	112.7 pb	$E_T^{miss} > 100$ GeV	$21.5 \pm 0.2\%$	32.8
$t\bar{t}$	150 GeV (p_{Tmax})	298.7 pb	$p_{Te,\mu} > 20$ GeV	$39.8 \pm 0.5\%$	19.6
Z+jets	150 GeV	160.3 pb	$p_{Te,\mu} > 15$ GeV	$13.2 \pm 0.2\%$	144
Z+jets	150 GeV	160.3 pb	$E_T^{miss} > 100$ GeV	$21.6 \pm 0.2\%$	44.6
W+jets	150 GeV	384.5 pb	$p_{Te,\mu} > 15$ GeV	$21.1 \pm 0.1\%$	29.8
W+jets	150 GeV	384.5 pb	$E_T^{miss} > 100$ GeV	$19.5 \pm 0.2\%$	32.3
Wt	none	57.89 pb	$p_{TW} > 150$ GeV, $p_{Tt} > 100$ GeV, $p_{Te,\mu} > 15$ GeV	$9.76 \pm 0.09\%$	92.6
Wt	none	57.89 pb	$p_{TW} > 150$ GeV, $p_{Tt} > 100$ GeV, $E_T^{miss} > 100$ GeV	$9.72 \pm 0.08\%$	93.0

Table 7.1: List of Monte Carlo samples simulated as part of the analysis. $m_H = 120$ GeV in all samples.

7.1 $H \to b\bar{b}$ Candidate Identification

Jets are found by taking the locally calibrated topological clusters (see Chapter 4) and applying the Cambridge-Aachen algorithm with an R-parameter of 1.2, as in Chapter 6. The calibration of these jets was known to be incorrect at the time when this analysis was performed and a simple Monte Carlo based calibration is applied to account for this. See Section 7.6 for more details.

The identification of Higgs candidates within these jets is that defined in Chapter 6, taking $\mu = 1/\sqrt{3}$ and $y_{2\mathrm{cut}} = 0.1$. This is slightly different from the values used in Chapter 6, a change which was found to cause no significant differences (see Section 7.6). A key difference in this analysis is the tagging of b-hadrons. In Chapter 6 this was done by searching for b-hadrons in any two of the three filtered subjets and applying a flat efficiency and fake-rate. This had a number of limitations such as assuming that the identification of the two b-hadrons is uncorrelated and that efficiency is constant in η. It also ignored the effects of c-hadrons, which can appear as fake b-hadrons much more frequently than light jets.

As briefly outlined in Chapter 4, tagging b-hadrons proceeds by identifying tracks from the simulation within $\delta R < 0.4$ of each of the subjets. The tracks are then processed by a number of tagging algorithms which look for signs of b-hadron decay. Specifically here, the algorithms used are a 3-dimensional impact parameter algorithm [72] and JetFitter [73]. JetFitter extends the idea of secondary vertex finding and simultaneously fits the primary, secondary b-hadron and tertiary c-hadron vertices in the tracks provided. Incorporating additional understanding of the decay physics in this way has been shown to offer improved performance compared to simpler secondary vertexing algorithms. The results of these two algorithms are combined by a neural network to provide a final estimate of the probabilities p_u, p_c and p_b that a jet is light, contains a c-hadron or contains a b-hadron respectively. Here the b-tagging "weight" is defined as:

$$ w = \ln \left(\frac{p_b}{p_u} \right) $$

A Higgs candidate after the splitting and filtering procedure has been applied is considered to be b-tagged if the two highest p_T subjets are both found to have $w > 1.0$.

There is no attempt to take advantage of possible correlations between the weights of the two subjets which could potentially provide better discrimination.

7.2 Events containing $Z \to l\bar{l}$

The goal of this selection is the identification of $H \to b\bar{b}$ candidates produced in association with a Z boson which then decays to either electrons or muons. Electrons from Z bosons are identified according to the ATLAS "loose" requirements. This corresponds to cuts on the shape of the shower in the electromagnetic calorimeter and also on the amount of energy which leaks into the hadronic calorimeter. A very loose match to a track in the inner detector is also required. An additional isolation requirement is applied on top of the standard selection, requiring that 90% of the energy in a cone of $R = 0.2$ around the electron is contained in the electron cluster.

Muons from Z bosons are identified according to ATLAS "combined" requirements, which means a track in the muon spectrometer successfully matched to a track in the inner detector. Again an additional isolation requirement is imposed, here that 75% of the energy in a cone of $R = 0.4$ around the muon is contained in the muon track.

A $Z \to l\bar{l}$ candidate is defined as a pair of opposite sign electrons or opposite sign muons with an invariant mass between 80 GeV and 100 GeV. The analysis is then as follows, for each hadronic $H \to b\bar{b}$ candidate H:

1. Require $p_{\mathrm{T}H} > 200$ GeV

2. Require $|\eta_H| < 2.5$

3. Require one Z candidate composed of two leptons l_1 and l_2 where $p_{\mathrm{T}l_1} > p_{\mathrm{T}l_2}$

4. Require $p_{\mathrm{T}Z} > 200$ GeV

5. Require $p_{\mathrm{T}l_1} > 25$ GeV

6. Require $p_{\mathrm{T}l_2} > 20$ GeV

7. Require $d\phi_{Z,H} > 1.2$

8. Require that for the two highest p_{T} subjets in H, $w > 1.0$

Some kinematic distributions relevant to the identification of $Z \to l\bar{l}$ candidates are shown in Figures 7.1 and 7.2. From these histograms it is clear that Z boson decays are indeed identified at the expected mass. The additional lepton p_{T} requirements, designed to avoid problems with triggering and identification of low p_{T} leptons, do not appear to have a significant effect on analysis efficiency.

Figure 7.3 shows the distance in ϕ between the leptonic Z and hadronic H candidates in events. It is quite easy for the leptonic Z candidate to pass the hadronic splitting and filtering procedure but very difficult for it to pass the b-tagging requirements. Conversely and also very rarely soft leptons could be present in the H candidate, leading to a fake Z candidate. Obviously for a jet to be identified as both the hadronic and leptonic candidate implies a failure of the analysis. This "overlap" cut guards against this by requiring separation between the candidates. The figure is made before the b-tagging weight requirement, so many pairs do indeed overlap. Excluding these, the events have the expected characteristic that the Z and H candidate are primarily back-to-back in the detector.

The number of events passing each of the cuts can be seen in Table 7.2. Having applied all the cuts, the final mass distribution of $H \to b\bar{b}$ candidates with 30 fb^{-1} of integrated luminosity can be seen in Figure 7.4. The mass resolution is somewhat worse than previously observed in Chapter 6 however, the twin Z and H peaks are still clearly visible above the background. The purely statistical significance of this result can again be calculated from the number of signal and background events within a mass window (in this instance 104 to 136 GeV) and is found to be 1.5σ. This is somewhat but not severely reduced compared to the hadron level analysis which obtained 2.6σ. The mass resolution is again somewhat although not seriously degraded, which also accounts for the wider window for significance calculation.

7.3   Events Containing $E_{\mathrm{T}}^{\mathrm{miss}}$

$E_{\mathrm{T}}^{\mathrm{miss}}$ is defined as the negative of the sum of all calibrated physics objects in the detector and remaining calorimeter clusters not included in any physics object. The goal here is primarily to identify $H \to b\bar{b}$ production in association with a $Z \to \nu\bar{\nu}$ which results in large $E_{\mathrm{T}}^{\mathrm{miss}}$. As such a leptonic vector boson candidate is defined by

Figure 7.1: Mass and p_T spectra of $Z \to l\bar{l}$ candidates.
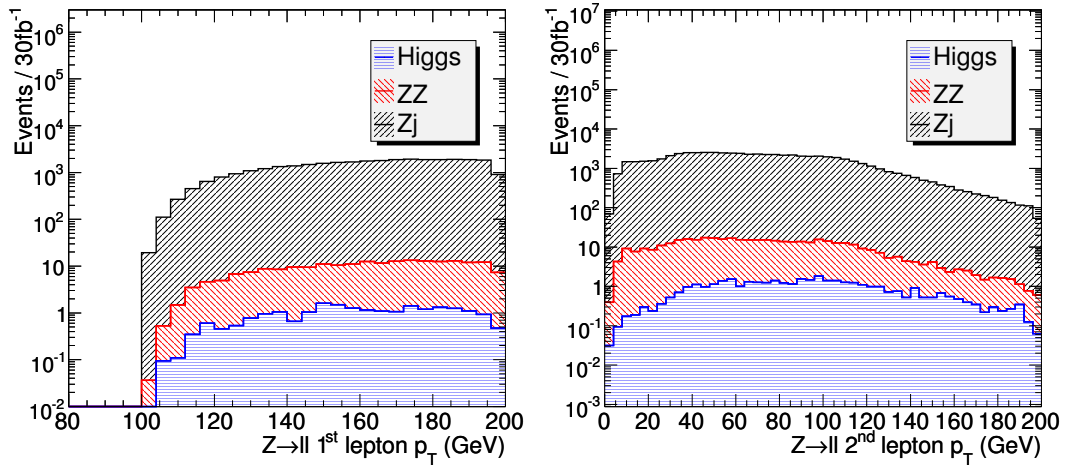
Figure 7.2: p_T spectra of $Z \to l\bar{l}$ candidate decay leptons.

requiring the presence of E_T^{miss} in the event.

As previously, the analysis is defined for each hadronic $H \to b\bar{b}$ candidate H:

1. Require $p_{TH} > 200$ GeV

2. Require $|\eta_H| < 2.5$

3. Require one Z candidate defined as equal to the E_T^{miss} in the event
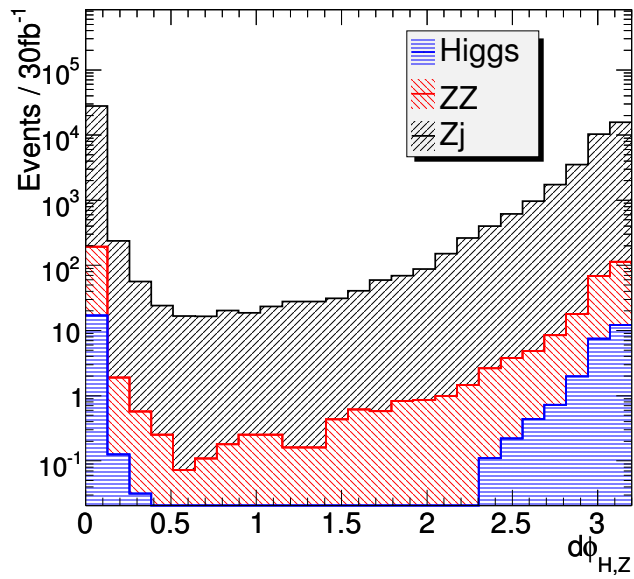
4. Require $p_{TZ} > 200$ GeV

Figure 7.3: The distribution of $\delta\phi_{Z,H}$ in each of the different samples in the $Z \to l\bar{l}$ selection.

Figure 7.4: Final invariant mass distribution of H candidates in the $Z \to l\bar{l}$ channel. Errors represent Gaussian uncertainty due to Monte Carlo statistics.

	HZ(120 GeV)	ZZ	Z+jets
Generated	575 ± 3	3129 ± 11	647460 ± 367
$p_{\mathrm{T}H} > 200$ GeV	263 ± 2	1186 ± 7	232152 ± 220
$|\eta_H| < 2.5$	259 ± 2	1166 ± 6	228413 ± 218
$80 GeV < m_Z < 100$ GeV	33.6 ± 0.7	372 ± 4	63149 ± 115
$p_{\mathrm{T}Z} > 200$ GeV	29.5 ± 0.7	310 ± 3	51160 ± 103
$p_{\mathrm{T}l_1} > 25$ GeV	29.5 ± 0.7	310 ± 3	51160 ± 103
$p_{\mathrm{T}l_2} > 20$ GeV	28.8 ± 0.7	282 ± 3	46204 ± 98
$d\phi_{Z,H} > 1.2$	23.0 ± 0.6	197 ± 3	33341 ± 83
b-tagged ($w > 1.0$)	8.0 ± 0.4	10.4 ± 0.6	91 ± 4
104 GeV $< m_H < 136$ GeV	5.3 ± 0.3	1.0 ± 0.2	11 ± 2

Table 7.2: Expected number of events in the $Z \to l\bar{l}$ channel with 30 fb^{-1} of data, after each selection criterion is applied.

5. Require no other leptons in the event with $p_{\mathrm{T}} > 30$ GeV

6. Require no other jets in the event with $p_{\mathrm{T}} > 30$ GeV

7. Require $\delta\phi_{Z,H} > 1.2$

8. Require that for the two highest p_{T} subjets in H, $w > 1.0$

Some basic properties can be seen in Figure 7.5, the $E_{\mathrm{T}}^{\mathrm{miss}}$ distribution shows a typical falling structure at high values but is biased at low values by the choice of minimum p_{T}s in the generated samples.

The effects of the veto on additional jets is shown in Figure 7.6. The p_{T} spectrum does not reveal significant differences between the samples. However, the scan of different possible cut values does reveal some structure, with stricter cuts offering better significance. Unfortunately, the rate of production jets with p_{T} of approximately 20 GeV at the LHC is expected to depend somewhat on underlying event properties which were not well constrained at the time this analysis took place. Considering this and the fact that the gains in significance are modest, a cut of 30 GeV was chosen as a safe value. It is also useful to note from this plot that the veto p_{T} cut could be increased somewhat further without a large loss of significance.

The number of events passing each of the cuts can be found in Table 7.3. The final mass spectrum of the candidates after all cuts have been applied and assuming 30 fb^{-1} of integrated luminosity is shown in Figure 7.7. As in Chapter 6 the peaks from both W/Z bosons and $H \to b\bar{b}$ are clearly visible. As in the $Z \to l\bar{l}$ selection above, the significance is calculated by counting signal and background events within a mass window of 104 to 136 GeV. And again, the significance of 1.6σ is somewhat lower than the 4.0σ found in the hadron level analysis.

7.4   Events Containing $W \to l\nu$

Although not undertaken as part of this thesis, the third leptonic selection as described in Chapter 6 has also been analysed [6, 74]. The results are visible in Figure 7.8. As with the other two channels, the significance is reduced somewhat from 3.9σ to 3.0σ but the overall effect of the analysis is retained, with both the weak boson and Higgs peaks clearly visible above the background after 30 fb^{-1} of integrated luminosity. The full details of this analysis will not be described here but it is sufficient to understand that the selection is exclusive and this is therefore an independent result which can be combined with the main results from this chapter.

7.5   Likelihood-based Combination of Channels

The results above represent the discovery potential for each of the three separate channels. In order to interpret the overall sensitivity of ATLAS to HW/HZ, these three results must be combined. In Chapter 6 this was done by simply adding the event selections together. However in this more complete analysis, it is desirable to include some treatment of the effects of systematic uncertainties and correlations.

When this analysis is eventually done in real data, it will probably be necessary to take into account shape information in a more complete fitting procedure. In this process, a large selection of systematics will have to be considered, many with subtle effects. Not only is this a large undertaking, it is not clear how well complex quantities such as these are modelled by detector simulation. As a result the approach here will be to treat the systematics as simply as possible while retaining some information about their broad effects on the significance.

Figure 7.5: $E_\mathrm{T}^\mathrm{miss}$ and lepton p_T spectra in the $E_\mathrm{T}^\mathrm{miss}$ channel, note that $E_\mathrm{T}^\mathrm{miss}$ spectra at low scales are significantly biased by generator level cuts.

Figure 7.6: Left: Scan of possible p_T cut values for jets to be vetoed in the $E_\mathrm{T}^\mathrm{miss}$ channel. Right: p_T spectrum of additional jets that are considered in the jet veto.

Figure 7.7: Final invariant mass distribution of H candidates in the E_T^{miss} channel. Errors represent Gaussian uncertainty due to Monte Carlo statistics.

Figure 7.8: Invariant mass distribution of H candidates in the $W \to l\nu$ channel [6].

	$HZ(120)$	$HW(120)$	WW	WZ	ZZ
Generated	379 ± 2	1254 ± 8	18975 ± 86	9549 ± 61	4650 ± 15
$p_{\mathrm{T}H} > 200$ GeV	133 ± 1	513 ± 5	7321 ± 54	3560 ± 37	1573 ± 9
$|\eta_H| < 2.5$	130 ± 1	505 ± 5	7114 ± 53	3465 ± 37	1537 ± 9
$E_{\mathrm{T}}^{\mathrm{miss}} > 200$ GeV	96.4 ± 0.9	71 ± 2	3403 ± 36	2047 ± 28	1132 ± 7
$p_{\mathrm{T}l} < 30$ GeV	94.8 ± 0.9	25 ± 1	2151 ± 29	1637 ± 25	1063 ± 7
$p_{\mathrm{T}j} < 30$ GeV	62.7 ± 0.7	10.5 ± 0.7	1060 ± 20	914 ± 19	699 ± 6
$d\phi_{E_{\mathrm{T}}^{\mathrm{miss}},H} > 1.2$	62.7 ± 0.7	10.4 ± 0.7	1059 ± 20	914 ± 19	699 ± 6
b-tagged $(w > 1.0)$	20.9 ± 0.4	3.1 ± 0.4	1.2 ± 0.7	14 ± 2	36 ± 1
$104 < m_H < 136$ GeV	14.4 ± 0.3	1.9 ± 0.3	0.0 ± 0.0	0.8 ± 0.6	3.3 ± 0.4

|	$Z + \mathbf{jets}$	$W + \mathbf{jets}$	Wt	$t\bar{t}$
Generated	1031751 ± 833	2249139 ± 1445	168765 ± 233	729000 ± 816
$p_{\mathrm{T}H} > 200$ GeV	266352 ± 423	630230 ± 765	38666 ± 112	331419 ± 550
$|\eta_H| < 2.5$	261090 ± 419	617078 ± 757	38418 ± 111	328448 ± 547
$E_{\mathrm{T}}^{\mathrm{miss}} > 200$ GeV	188447 ± 356	291114 ± 520	11428 ± 61	82062 ± 274
$p_{\mathrm{T}l} < 30$ GeV	185851 ± 354	195853 ± 426	6696 ± 46	51338 ± 216
$p_{\mathrm{T}j} < 30$ GeV	110517 ± 273	92467 ± 293	1344 ± 21	5375 ± 70
$d\phi_{E_{\mathrm{T}}^{\mathrm{miss}},H} > 1.2$	110516 ± 273	92426 ± 293	1343 ± 21	5373 ± 70
b-tagged $(w > 1.0)$	279 ± 14	127 ± 11	32 ± 3	129 ± 11
$104 < m_H < 136$ GeV	32 ± 5	17 ± 4	11 ± 2	41 ± 6

Table 7.3: Expected number of events in 30 fb^{-1} of data in the $E_{\mathrm{T}}^{\mathrm{miss}}$ channel after each selection criterion is applied.

Specifically, here the significance calculation will still be based on event counting within a mass window. In an event counting experiment, the most important systematic is the background rate, which is the only one which will be considered here. The background samples will be divided into three categories:

Group	Samples Included
t	$t\bar{t}$, WW, WZ and ZZ
w	$W + $ jets and Wt
z	$Z + $ jets

It will be assumed that the rate of each of these categories has a Gaussian uncertainty. These rates are considered to be completely uncorrelated between the groups

but the same between the three different analyses. In reality the event counts would be subject to a number of uncertainties, some correlated between channels, some correlated between groups and some with more complicated correlations still. However, this model has around the right number of degrees of freedom and should be sufficiently flexible to at least roughly describe most scenarios. As an additional check, the impact of breaking the samples into alternate groupings, including a set of four groups was tried and found not to be significant.

The likelihood function for a simple event counting experiment i, where n events are observed can be described by a Poisson distribution:

$$L_i(\mu) = \frac{(\mu s_i + b_i)^{n_i}}{n_i!} e^{-(\mu s_i + b_i)}. \tag{7.1}$$

Where b and s represent the expected background and signal counts respectively. The symbol μ parameterises the level of signal present where $\mu = 1$ represents the standard model case.

Since here we need to combine several experiments we need a likelihood form which describes this. The likelihood function for N experiments of the type described above can be expressed as:

$$L(\mu) = \prod_{i=1}^{N} L_i(\mu). \tag{7.2}$$

As mentioned earlier, we also want to introduce the effects of the systematic uncertainty caused by our imperfect understanding of the backgrounds, here assumed to be Gaussian. This requires separation of the background component of the background b_i of each experiment into separate components:

$$b_i = t_i + w_i + z_i. \tag{7.3}$$

Expanding the combined likelihood in this way and adding additional terms for the uncertainties on each quantity yields the following likelihood function:

$$L(\mu, C_t, C_w, C_z) = \prod_{i=1}^{N} \frac{(\mu s_i + C_t t_i + C_w w_i + C_z z_i)^{n_i}}{n_i!} e^{-(\mu s_i + C_t t_i + C_w w_i + C_z z_i)}$$
$$\times Gaus(C_t) \times Gaus(C_w) \times Gaus(C_z)$$

$$\text{where} \quad Gaus(C_x) = \frac{1}{\sigma_x \sqrt{2\pi}} e^{-\frac{(C_x - 1)^2}{2\sigma_x^2}}.$$

The new constants C_x allow the different groups of background to vary in a manner regulated by the Gaussian terms. We can now proceed to calculate a significance. To do this we first define the likelihood ratio and for convenience an associated test statistic:

$$\lambda(\mu) = \frac{L(\mu, \hat{C}_t, \hat{C}_v)}{L(\hat{\mu}, \hat{\hat{C}}_t, \hat{\hat{C}}_v)}$$

$$q(\mu) = -2\ln(\lambda(\mu)).$$

Where the variables $\hat{C}_t$ and $\hat{C}_v$ maximise L for the given μ, whereas $\hat{\mu}$, $\hat{\hat{C}}_t$ and $\hat{\hat{C}}_v$ maximise L overall.

In order to calculate a discovery significance, we attempt to reject the background-only case, where the true $\mu = 0$. Given this we also do not allow $\hat{\mu} < 0$; in such cases, we fix $\hat{\mu} = 0$. This is because although experiments where the observation is less than the expected background-only count may not look consistent with the background prediction, they certainly do not imply the presence of a Higgs signal. First we define a function which represents the distribution of possible outcomes of experiments where only background is present:

$$f(x|\mu_{\text{true}}) = P(q(0) = x|\mu_{\text{true}} = 0). \tag{7.4}$$

The p-value associated with the outcome of an actual experiment can then be defined as the fraction of experiments at least as unlikely as this one:

$$p = \int_{q(0)}^{\infty} f(x|0)dx. \tag{7.5}$$

Which can be associated to a significance using the Gaussian integral:

$$Z = \Phi^{-1}(1-p). \tag{7.6}$$

The biggest computational challenge is determining the form of f with sufficient accuracy. The function could be determined via Monte Carlo but given that a significance of 5σ corresponds to a p-value $\sim 10^{-7}$, the number of pseudo-experiments required is $\gtrsim 10^8$. It is known [4] that the form of f when $\mu > 0$ can be approximated by a chi-square distribution with one degree of freedom. However, for experiments where $\mu < 0$ we are setting $\mu = 0$, so these will always give $q(0) = 0$. Therefore the form of f can be described as follows, the only unknown parameter being w, the fraction of experiments where $\mu > 0$:

$$f(x|0) = wf_{\chi_1^2}(x) + (1-w)\delta(x). \tag{7.7}$$

There are many possible ways to determine the value of w. A simple one used here is to perform a smaller number of Monte Carlo pseudo-experiments ($\sim 10^6$). Each pseudo-experiment consists of generating the constants C_x according to their Gaussian uncertainty. These constants are used to modify the background sample expectation values, which are then in turn used to generate Poisson random event counts. In the single fits, the nominal values for the constants C_x are used, so that their difference with respect to the expectation values used for generation correctly reflect their expected systematic uncertainty. The test statistic $q(0)$ is then calculated for each pseudo-experiment and used to fill a histogram.

If the lowest bin in this histogram starts at $x = 0$ and ends at a small number $x = a$, the contents of this bin are the integral:

$$\int_0^a f(x|0)dx = N_{\exp} \times \left(w \int_0^a f_{\chi_1^2}(x)dx + (1-w) \right). \tag{7.8}$$

This equation is trivially soluble for w since the integral of the chi-square distribution is calculable. Typically the statistical uncertainty on this value of w is small since approximately half of all Monte Carlo experiments have values of $\mu < 0$. This

value of w can then be used to test the validity of the chi-square approximation of f for experiments where $\mu > 0$. In this analysis, the agreement appears to be very good, as seen in Fig. 7.9.

With the form of f determined sufficiently well, many experiments containing signal ($\mu_{\text{true}} = 1$) can be generated and their significance calculated. Out of this distribution of possible outcomes, the median is taken as a reasonable expectation of performance. We can now attempt to produce our combined significance estimate for the three channels: $Z \rightarrow l\bar{l}$, $W \rightarrow l\nu$ and $E_{\text{T}}^{\text{miss}}$. Signal and background expectations are taken from the analyses performed, the values for which are given in Table 7.4. Initially, a perfect understanding of the backgrounds is assumed and the results for 30 fb^{-1} are as seen in Fig. 7.10. Here a median significance of 3.7σ is observed, which is consistent with what would be expected from adding the $S/\sqrt{(B)}$ significances in quadrature. A variety of possible background uncertainties are then tested, the results of which can be found in Table 7.5.

The exact background uncertainties for each channel depend on the mixture of different background samples. A 10% uncertainty in each of the three background samples corresponds to a 9%, 7% and 6% uncertainty on the total background level in $Z \rightarrow l\bar{l}$, $W \rightarrow l\nu$ and $E_{\text{T}}^{\text{miss}}$ channels respectively. These numbers scale linearly with sample uncertainty.

One method of determining the background event rates experimentally is with a fit. This approach is explored in Ref. [74], which implies that using this method alone, with 30 fb^{-1} of intergrated luminosity results in an uncertainty of 15% per background sample. Given that after 30 fb^{-1} of data ATLAS should also have a robust understanding of the background from the data itself we estimate that an uncertainty of 10% or better is realistic.

The significance of 3.7σ in the perfect case is found to be reduced to 3.2σ in the case of a 10% uncertainty and 3.0σ in the case of a 15% uncertainty on the expected level of each background sample. It is useful to note that the combination of the three separate channels with differing background compositions reduces the effects of this systematic uncertainty.

Figure 7.9: Distribution of the test statistic $q(0)$ for background-only experiments, the histogram shows Monte Carlo in excellent agreement with the solid line representing a one degree of freedom chi-squared distribution.

Channel	s_i	t_i	w_i	z_i	$S/\sqrt{B}$
$Z \to l\bar{l}$	5.34	0.98	0.0	11.2	1.5
$W \to l\nu$	13.5	7.02	12.5	0.78	3.0
$E_\mathrm{T}^\mathrm{miss}$	16.3	45.2	27.4	31.6	1.6
Combined					3.7

Table 7.4: Experiment event count expectations used in the combination.

7.6   Additional Material

In addition to the plots produced as part of the main analysis, a number of other short tests were performed for various, generally technical reasons and will be described in this section.

The first set of plots shows the effect on the analysis of choosing one of two different μ parameters. While $\mu = 1/\sqrt{3}$ was used in Chapter 6 and in this analysis, the result in Section 7.4 used $\mu = 2/3$. This brief test was conducted to see if this change was likely to have any significant effect on the outcome of the analysis. The results are shown in Figure 7.11. The conclusion being that a small change in this parameter has very little effect on either shape or efficiency.
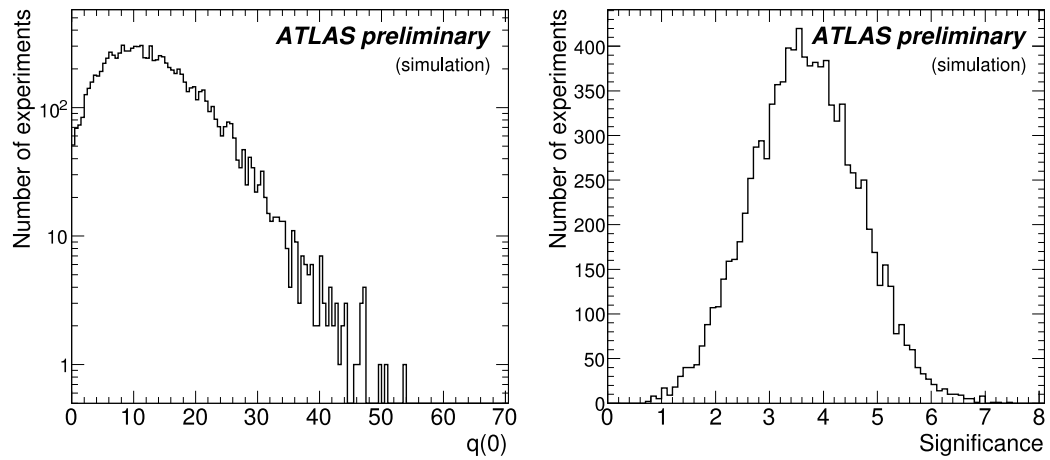
Figure 7.10: Distributions of $q(0)$ and significance for a range of possible signal experiments.

σ_t	σ_w	σ_z	Significance
Perfect	Perfect	Perfect	3.7
5%	5%	5%	3.5
10%	10%	10%	3.2
15%	15%	15%	3.0
20%	20%	20%	2.8
30%	30%	30%	2.5
50%	50%	50%	2.2
75%	75%	50%	2.0
50%	10%	10%	2.8

Table 7.5: Significances for different analysis scenarios with differing background uncertainties. The bottom two rows show scenarios with differing uncertainties amongst the backgrounds.

Figure 7.11: Final mass spectra after $Z \to l\bar{l}$ analysis with different choices of μ parameter in HZ events (left) and Z+jets events (right).

When this analysis was conducted, it was known that the p_{T} scale of jets constructed from locally calibrated topological clusters was not correctly calibrated. In order to check the effects of this a sample of HZ events was analysed. The jet finding and Higgs candidate identification procedure was applied to both the detector level topological clusters and the true Monte Carlo hadrons. Figure 7.12 takes all detector level candidates with $p_{\mathrm{T}} > 150$ GeV which also match a truth jet with $\delta R < 0.3$ and plots the relative difference in p_{T} between the two. The figure shows a peak at approximately -0.06. This is interpreted as meaning that on average a jet reconstructed in the detector will have a p_{T} of 6% less than the equivalent hadron-level jet. Based on this, all Higgs candidate jets are calibrated by multiplying their four-vector by 1.06.

The analysis in Section 7.4 also used a slightly different clustering procedure as compared in Figure 7.13. The comparison is again performed using the HZ signal sample. When reclustering to find the 3 filtered subjets, the analysis in Section 7.4 requires that the 2 highest p_{T} subjets come from opposite halves of the jet as it is after the hard splitting, this strategy is labelled "separate" in the figure. The procedure as described in Chapter 6 makes no such discrimination and simply takes the 3 highest p_{T} objects after the reclustering, labelled "global" in the figure. The figure shows that although apparently some events migrate between bins under the different clustering strategies, the differences are small and unlikely to have a large effect on the overall

Figure 7.12: Distribution of relative p_T scales of reconstructed and hadron-level jets, fitted with a Gaussian distribution. The peak indicated reconstruction underestimated jet energy by approximately 6%.

effectiveness or significance of the analysis. From this we may conclude that this minor difference in analysis procedure does not introduce any significant issues.

7.7   Conclusions

This analysis represents a significantly more comprehensive study of the effectiveness of a high-p_T, jet substructure based analysis for light Higgs discovery at the LHC than Chapter 6. Many subjects which were not treated or incompletely treated at the hadron level have now been more fully addressed including, b-tagging, charm backgrounds, single top, lepton identification and resolution effects. In addition, a basic treatment of the most important systematic, background cross-section has been made. In all areas, the performance of the detector has been found to be comparable to expectations. In particular, some aggressive assumptions about b-tagging efficiency and fake rates were found to be justified. However, overall a number of small degradations in performance and the addition of systematic effects conspire to reduce the overall significance of the analysis somewhat. The significance of around 5σ at the hadron level has been reduced to around 3σ.

This analysis does of course have limitations. Chiefly, it is based purely on Monte Carlo and detector simulation. Much care has been taken to design an analysis which does not depend on poorly understood aspects of these packages. The analysis also operates mainly in the high-p_T region of phase space where QCD is perturbative and
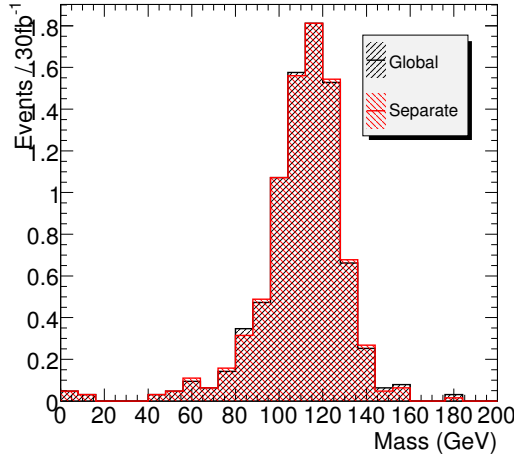
Figure 7.13: Comparison of different clustering strategies. Separate requires that the 2 highest p_{T} subjets come from opposite halves of the jet after the hard splitting stage, global does not.

theory uncertainties tend to be small. However, in spite of all this, large deviations from our current best expectations are possible and would cause correspondingly large changes in the effectiveness of this analysis.

An exact treatment of systematics has not been attempted here, although for the HW channel this has been done in some detail [75]. Additionally, the mechanism by which the background shape would be measured has not been explored in detail. One possible approach is to make assumptions about the functional form of the backgrounds and perform a likelihood based fit, again tried in [75]. When this measurement is performed with data it is likely that the background estimate will be a combination of several inputs, such as theoretical calculations, measurements of related processes and fitting the analysis output itself.

On the positive side, the analysis has not been exhaustively optimised and gains are probably still possible. In addition, even some fairly modest multivariate techniques may offer significant performance improvements. One particular example being that the p_{T} of additional jets and the signal candidate are strongly correlated.

Overall, this represents a significant enhancement over the previous ATLAS analysis, which for this Higgs mass of 120 GeV was only able to achieve a significance of 1.7σ compared to the 3σ or better acheived here. Even then this significance was achieved

by assuming a better understanding of systematics than in this chapter in spite of the fact that the systematics in the older analysis are more challenging.

This analysis technique therefore transforms low mass Higgs HW/HZ production from a high-luminosity couplings measurement to a serious contender for Higgs discovery at ATLAS. The effectiveness of Higgs discovery techniques was last evaluated for 10 fb^{-1} of integrated luminosity at an LHC center-of-mass energy of 14 TeV. Simply scaling the significance by luminosity for this analysis implies 2σ ought to be achievable with 10 fb^{-1}. For comparison, with the same amount of luminosity, the other leading discovery channels, $H \to \tau\tau$ and $H \to \gamma\gamma$ both expect to achieve just over 2σ.

Beyond the pure discovery measurement this analysis has further value. In combination with other measurements, this analysis offers the possibility of measuring the $H \to b\bar{b}$ branching ratio. As mentioned previously, recent studies have shown that measuring this ratio is critical to determining that any new resonance is indeed a Higgs boson [70]. This analysis is also the easiest way to access information about the HW and HZ couplings.

Finally, it should be noted that there are theoretical models, notably some SUSY scenarios, which predict enhanced $H \to b\bar{b}$ production [76]. This analysis will be able to place limits on such anomalous couplings with less than 10 fb^{-1} of integrated luminosity, although this has not been explicitly explored here.

Chapter 8

Observing Boosted W and Z Bosons in Early Data

In the early collisions at the LHC, large numbers of W and Z bosons will be produced. Many will be singly produced in W+jets and Z+jets events. Events where the bosons decay leptonically have fairly distinctive features and are easily identified. The events where the bosons decay hadronically however are significantly more difficult to extract. As in Chapters 5 and 6, the heavy particle decay signal is buried in QCD background, made worse by combinatorial effects. Again, one possible solution is to apply a hard p_{T} cut followed by jet substructure analysis.

8.1 Analysis Procedure

The effectiveness of the Cambridge-Aachen based technique from Chapter 6 for identifying W and Z bosons in early ATLAS data will be explored in this chapter. In order to achieve this, events were generated with HERWIG 6.510 [25, 28] using the Rivet [77] framework. Samples of QCD dijets, W+jets and Z+jets events were generated corresponding to the amounts expected for $1~\mathrm{fb}^{-1}$ of integrated luminosity at the LHC. All samples were generated with a centre-of-mass energy of 7 TeV as this is the expected running configuration for early collisions.

Initially we define jets using the Cambridge-Aachen algorithm with a radius parameter $R = 0.7$. This value was chosen based on experience with the high p_{T} selection in Chapter 6. Some basic kinematic variables are shown in Fig. 8.1 and Fig. 8.2. Fig. 8.1a

shows that before applying any jet mass or substructure requirements, the background is around two orders of magnitude above the signal. Fig. 8.1b shows that the jets are back-to-back, that is the structure of the events is dijet-like. In Fig. 8.2 the rapidity and pseudorapidity distributions are shown, after applying a $p_\mathrm{T} > 400$ GeV cut. The jets are central in the detector.

We next apply the procedure as described in [71] with the parameters set to $\mu = 1/3$ and $y_{2\mathrm{cut}} = 0.09$. Significances are calculated by counting the number of jets in a mass window of 75 GeV to 95 GeV. Jets from W+jets or Z+jets events are considered to be signal (S), while jets from QCD dijet events are considered to be background (B).

In performing the subjet analysis we scan across the relevant parameters. First we adjust the p_T requirement, the effects on S/B and $S/\sqrt{B}$ can be seen in Fig. 8.3. Higher p_T cuts tend to offer somewhat better S/B but worse $S/\sqrt{B}$. Tuning this cut offers some ability to reduce statistical uncertainties at the expense of increasing systematic uncertainties and vice versa. Based on this we choose 400 GeV as a reasonable benchmark cut and also sufficiently high as to be relatively free from experimental effects such as trigger limitations. The distribution of heavy-particle candidate masses after a $p_\mathrm{T} > 400$ GeV cut but before any subjet analysis is shown in Fig. 8.4.

The results of tuning μ can also be seen in Fig. 8.3. It is observed that lower cuts (i.e. stricter mass drop requirements) tend to offer better S/B. However at very low values the significance starts to be affected by strongly falling signal statistics. Based on this graph, we choose two benchmark points, one conservative: $\mu = 1/3$ and one somewhat more aggressive: $\mu = 1/5$. These two points offer hadron-level significances of around 5σ with 1 fb^{-1} of LHC data with S/B of around 5% and 13% respectively.

The distributions after the two choices of subjet analysis described above, are shown in Fig. 8.5. Clearly the signal is greatly enhanced compared to Fig. 8.4. The plots are binned in 8 GeV intervals, a value believed (based on full detector simulation [6]) to broadly approximate the experimental resolution effects. In both, the peaks are clearly visible above the QCD background although the background shape is very different between the two. The slightly peaked background for $\mu = 1/5$ may be disadvantageous from a systematic point of view although this may be compensated by the higher S/B.

One further tuning that can be explored is to increase the radius parameter R of the initial jet finding from 0.7 to 1.2. The effects of this change can be seen in Fig. 8.6.

Figure 8.1: Kinematic distributions for Cambridge-Aachen jets with $R = 0.7$, p_T (left) and $d\phi$ between the two leading jets where the leading jet has $p_T > 400$ GeV (right).


Figure 8.2: η (left) and rapidity (right) for Cambridge-Aachen jets with $R = 0.7$ and $p_T > 400$ GeV.

Figure 8.3: Effects of a scan of possible values of p_T cut (left) and μ cut (right).

Figure 8.4: Mass of heavy particle candidates from Cambridge-Aachen R=0.7 jets with $p_T > 400$ GeV where no jet substructure procedure has been applied.

Figure 8.5: Mass of heavy particle candidates after jet substructure analysis on Cambridge-Aachen R=0.7 jets with $p_T > 400$ GeV for two scenarios, $\mu = 1/3$ (left) and $\mu = 1/5$ (right).

Although the signal region is largely unaffected, the shape of the background in the high mass tail is strongly flattened. This ability to choose a flatter background shape may well be useful in a full study. Overall these plots show that by tuning the available parameters there is a great deal of flexibility in terms of background shape.

The shape of the signal distribution can be seen in Fig. 8.7, plotted with both 8 GeV binning as the above plots and with 4 GeV binning. Although experimental resolution will probably not reach 4 GeV the large signal statistics may make it possible at somewhat higher luminosities to obtain information about the relative rates and positions of the W and Z boson peaks.

8.2 Conclusions

Extracting the singly-produced hadronically decaying W and Z bosons in W+jets and Z+jets events is a challenging task which has never been accomplished at a hadron collider in this high p_T case (although unboosted decays have been observed [78]). Subjet techniques such as this offer a possible approach, here showing that with minimal tuning at the hadron-level it is possible to extract a significance of around 5σ within 1 fb^{-1} of LHC luminosity. A full study with detector simulation is required to more accurately

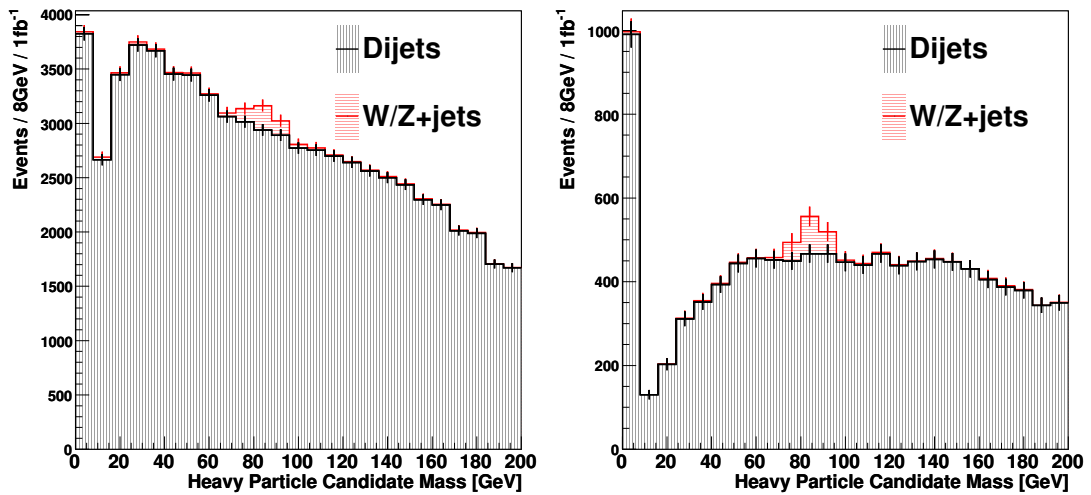Figure 8.6: Mass of heavy particle candidates after jet substructure analysis on Cambridge-Aachen R=1.2 jets with $p_T > 400$ GeV for two scenarios, $\mu = 1/3$ (left) and $\mu = 1/5$ (right).


Figure 8.7: Mass of heavy particle candidates in signal sample only after jet substructure analysis on Cambridge-Aachen R=1.2 jets with $p_T > 400$ GeV and $\mu = 1/3$ in 8 GeV bins (left) and 4 GeV bins (right).

evaluate the sensitivity of this technique but the available evidence and current LHC schedule suggests a promising outlook for this measurement in the 2010/2011 run.

Chapter 9

Visualisation

Modern particle physics experiments like ATLAS produce vast quantities of extremely complex data. Visualisation software is an important tool when attempting to make sense of this data. The ability to display experimental data in a visual manner is also a key communication tool.

The Atlantis software package is a specific type of visualisation tool, known as event display software. Whereas most monitoring and analysis in ATLAS focusses on trends in large samples of events, event display software focusses on the different types of data available within a single event. By exploring ATLAS data in this largely orthogonal manner, problems or interesting correlations can be found which otherwise may go unnoticed.

One of the goals of Atlantis is to be as accessible as possible. As such, Atlantis is developed in the Java programming language [79]. Java is compiled to bytecode which is executed by a virtual machine. Since Java virtual machines are available for all major operating systems, Atlantis is extremely portable. Data such as detector geometry, detector output and reconstruction information is extracted from the ATLAS software by a package of the ATLAS software called JiveXML and stored in an XML format as shown in Figure 9.1. Various methods of transmitting the XML are supported including storage on disk, via a built-in server and via HTTP from a standalone web-server.

Atlantis then displays this data in a number of graphical projections. Since ATLAS data is fundamentally 3-dimensional whereas a computer screen is fundamentally 2-dimensional it is important to find ways of representing data usefully to the end user. For an example of the usual appearance of Atlantis, see Figure 9.2.

Figure 9.1: Atlantis receives data in XML format via the JiveXML package in the ATLAS software.

9.1   Usage During LHC Startup

As mentioned previously, Atlantis is extremely useful for understanding detector output. This is never more critical than during the startup phase of an experiment, when detector response is relatively poorly understood. For this purpose, ATLAS possesses a small farm of "monitoring" computers. A small portion of the raw output from AT-LAS is fed to these computers. The raw data is then decoded directly and processed to produce output for monitoring applications, including Atlantis. This mechanism makes it possible to run Atlantis in the ATLAS control room, showing events seconds after they are observed by the detector (see Figure 9.3). It is also the mechanism by which events are provided to "ATLAS Live" [80].

During both the 2008 startup of the LHC and the 2009 restart, the Atlantis development team were on call 24 hours a day to assist with examining the very first events to come out of the detector (Figure 9.4). Co-operating closely with relevant detector experts, many useful insights were gained into the working of ATLAS on both occasions. The uses to which Atlantis was put during these periods are too numerous to describe exhaustively here. Some highlights with which I had involvement include:

Detector Timing - especially during the 2008 startup, Atlantis was used extensively to explore timing offsets between the different parts of the ATLAS detector.

Identifying Collisions - during the first collisions in 2009, Atlantis was used to both
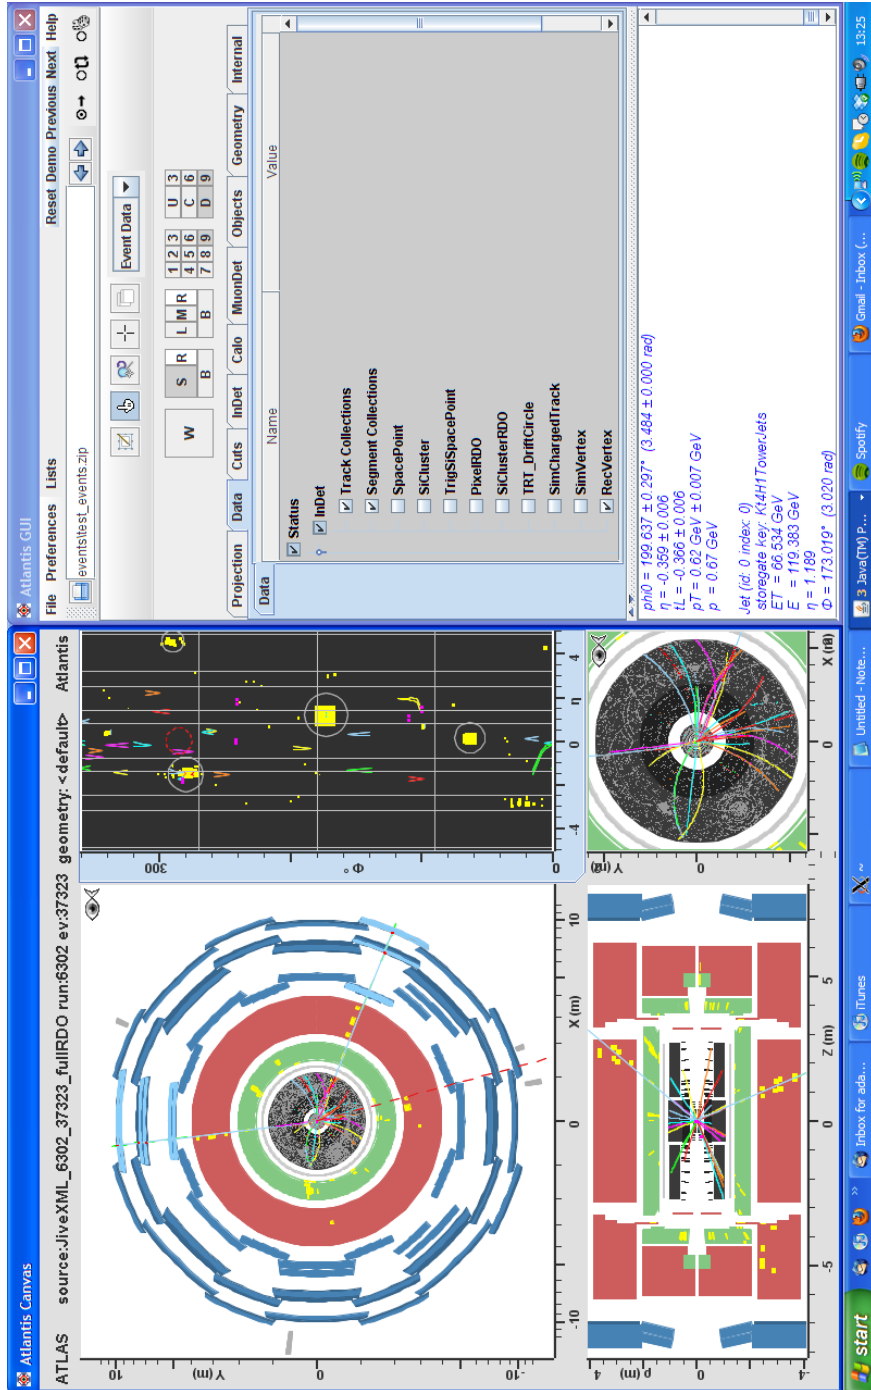
Figure 9.2: Typical usage of the Atlantis application. The left window shows event data and allows the user to interact, while the right window allows the user to control various aspects of the display.

to extract calorimeter timing information and to visually identify the collision vertex. The combination of these two pieces of information allowed ATLAS to confirm the observation of the first LHC collisions (Figure 9.5).

Highest Energy - during machine testing on 8^{th} December 2009, ATLAS detected protons colliding at a centre-of-mass energy of 2.36 TeV, making the LHC the highest energy collider in the world (Figure 9.6).

Outreach - images produced with Atlantis were amongst the very first LHC results to be presented to the general public on both occasions.

As expected, Atlantis played a key role in understanding the ATLAS detector in this crucial startup phase.

9.2   Graphics Code

Atlantis was ported to Java from FORTRAN code based on the DALI event display package, written in the late 1980s for the ALEPH experiment. It is worth noting that graphical user interfaces only came into widespread use in the early 80s. For example the X Window System was designed in 1984. At the time, it was a relatively advanced and innovative piece of software.

Of course rather a lot has changed in the world of computer graphics between the 1980s and today. In that period, graphics hardware generally consisted of dumb framebuffers, a block of physical memory reflected by the state of pixels on a display device. All processing to turn data into pixels on screen was written by hand and executed on the main processor (see Figure 9.7a). This was the environment in which Atlantis was originally developed and the design of the internal graphics code reflects this.

During the 1990s, systems designers noted that in graphics heavy applications, especially games, most processor time was spent on operations such as matrix and vector calculations. These are relatively simple tasks but must be performed a large number of times to draw a scene. As such, graphics processing units (GPUs) which were capable of performing these calculations quickly in specialised hardware began to be added to computers, as shown in Figure 9.7b.

Figure 9.3: Atlantis in use in the ATLAS control room during the 2009 LHC commissioning.

As hardware advanced and graphical applications became more complex, it became clear that there was a desire for more flexibility. Many aspects of graphical applications involve repetitive calculations over large sets of data. Motivated by this, modern GPUs are becoming increasingly programmable. The latest GPUs are essentially fast "Single Instruction Multiple Data" (SIMD) parallel processors. All signs suggest that as time goes on, graphics hardware will look increasingly like a programmable general purpose device. On such a device it will be possible to replace many of the slow elements of an application like Atlantis with efficient "shader" programs running on a GPU, Figure 9.7c.

9.3 Current Graphics Engine

The current graphics engine generates arrays of co-ordinates corresponding to graphics primitives such as lines and polygons. These co-ordinates are then transformed purely in software to calculate where they should be drawn. The results are either drawn to

Figure 9.4: Members of the Atlantis development team working late into the night scanning events in the ATLAS Data Quality Satellite Control Room: Nikos Konstantinidis (left), Mark Stockton (right).

screen using the "Java2D" graphics API or less frequently, converted to the PostScript vector graphics format. It makes no explicit use of any available GPU, although some operations may be indirectly accelerated by Java2D depending on the version of Java and platform.

The process of rendering an image in this graphics engine can be broadly broken down into the following steps. Class names are in *italics*:

Request - a request to repaint the view of the detector is generated, usually by user interaction. The *AWindow* object managing the current view is called which requests that the current projection object (of class *AProjection*) redraw the visible data.

Coordinate Generation - several *AData* and *ADetector* objects are called which generate *ACoord* objects. Each *ACoord* object contains co-ordinates describing graphics primitives in the co-ordinate system of the calling *AProjection*.

Figure 9.5: One of the first proton-proton collisions observed by ATLAS. In this event $\sqrt{s} = 900$ GeV.

Figure 9.6: The first ever observation of proton-proton collisions with $\sqrt{s} > 1.96$ TeV at a collider experiment.
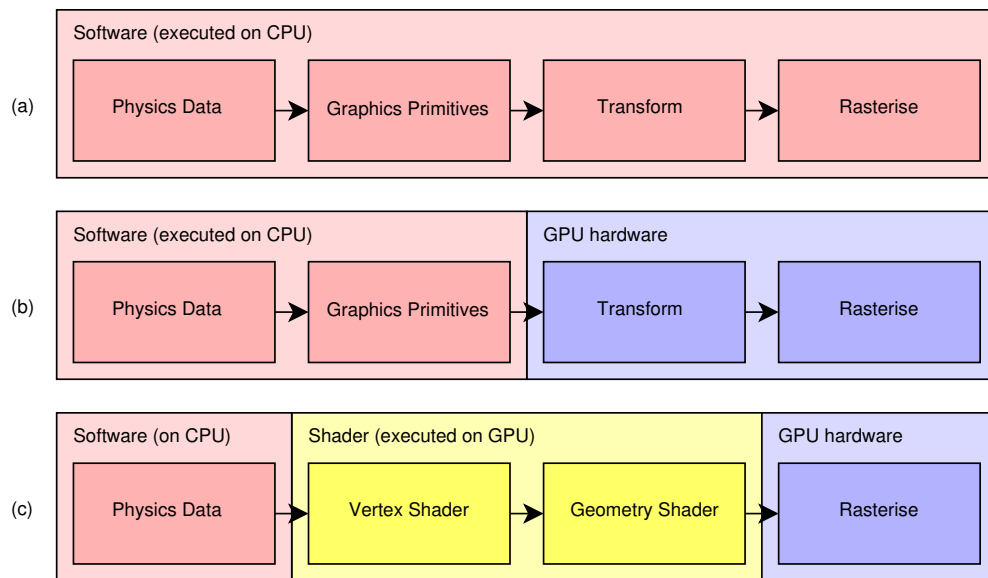
Figure 9.7: Possible graphical rendering pipeline designs from (a) 1980s to early 1990s (b) mid-1990s to mid-2000s (c) late 2000s onward.

Transformation - the *AData* or *ADetector* object calls back into the calling *AProjection* object to apply transformations such as fisheye to the *ACoord* objects.

Rendering - the *AData* or *ADetector* object makes calls to an *AGraphics* object to draw the transformed data.

This process allows the rendering of numerous different data types within Atlantis. However, it is subject to a number of limitations:

Object Separation - the code has been ported to Java from a non-object oriented programming language (FORTRAN in this case) and separation between objects is poor. Subclasses of *AData* for example, frequently contain both the detector data and graphics code. Division of functionality is often inconsistent within class hierarchies. This makes extending the engine extremely difficult.

Recalculation - even though detector data is relatively static, no attempt is made to cache the results of transformation operations. Every repaint causes a complete regeneration of the required graphics primitives and associated transformations.

Hardware Acceleration - as mentioned previously, no explicit use of dedicated GPU hardware is present.

Performance - as a consequence of the above two points, frequently the engine is slow to draw a scene, with frame rates falling to a few frames per second or worse when dealing with large events. This seriously damages the interactivity of the application.

Limited Functionality - the engine cannot satisfy requests for additional functionality or improved graphical quality. In some cases this is due to limitations in the engine design or Java2D. In other cases it is because engine performance is already at breaking point and additional computation is not acceptable. Anti-aliasing for example can currently only be used selectively without causing unacceptably low performance..

Overall there are many pressing concerns with the current engine code. The performance concerns especially are compounded by the fact that as LHC luminosity rises, events will become significantly larger. Simulation events generated as part of SLHC and ATLAS upgrade studies will probably be essentially unusable in Atlantis with the current graphics code. The ideal solution to this would of course be to introduce incremental updates to introduce improvements without breaking any existing functionality. However, the tight coupling of classes makes incremental changes extremely difficult. Furthermore the basic layout of the graphics code does not match well with modern graphics concepts. Given this situation, the only realistic option is a radical redesign.

9.4   A New Design

The design of this new graphics engine is motivated significantly by the limitations of the old code. The goal is to deliver a high-performance graphics engine, which is flexible, extensible and scalable. The requirements are as follows:

Performance - must offer a fluid user experience at all times.

Quality - must render high-quality graphics, using features such as anti-aliasing whenever possible.

Scalability - must be able to maintain performance with events an order of magnitude more complex than those currently being visualised.

Flexibility - must be possible to add new features and visual elements, such as transparency, animation and new projections.

Portability - must run on all platforms.

The only feasible way to satisfy the performance, quality and scalability requirements is to take advantage of dedicated GPU hardware. Therefore an important first choice is the technology and APIs which will be used to access the GPU. There are only two modern APIs for accessing GPU hardware, DirectX and OpenGL. DirectX is a proprietary system designed by Microsoft, only available as part of the Windows operating system. OpenGL by contrast is available on all major operating systems and is therefore the only choice that fulfils the portability requirement above.

Java does not contain a built-in library for using OpenGL. There are however a number of libraries available. Java3D is a library offered by Oracle Corporation (formerly Sun Microsystems) the company responsible for Java itself, however it has been unmaintained for several years and does not offer many modern features. JOGL [81] and LWJGL [82] are both better maintained and offer reasonable access to modern OpenGL features. Ultimately JOGL was chosen because it has better integration with the existing Java user interface framework (which Atlantis uses heavily), whereas LWJGL is more game-oriented.

The final question was which version of OpenGL should the engine target. Newer versions require more recent graphics hardware but provide more functionality. Ultimately it was decided that some notion of abstracted renderers was required which allowed the engine to function adequately on the widest possible range of hardware.

The overall design approach was prototyping-based. First a basic set of features was selected and implemented as a prototype. Then some additional features were selected and implemented, identifying points of tension in the design and refactoring code as necessary. This procedure was then repeated iteratively. At this time, the graphics engine is not complete, however it has reached a sufficiently advanced stage that the core of the design is stable.

Since Atlantis is a relatively stable software package in regular use, care was taken to minimise changes to the existing code. In order to differentiate the classes in the new graphics engine from frequently similarly named (and purposed) classes within

Atlantis, all new classes start with *AN*. The new code is broadly divided into these areas:

Management - *ANManager* and *ANFrameManager* manage graphics buffers and the high-level aspects of the rendering process.

Projection - *ANProjection* classes define projections, which decide which objects should be drawn to the screen.

Objects - *ANObject* classes define things which can be drawn, usually some way of representing a piece of data or geometry.

Caching - *ANCache* and *Token* classes provide a simple transparent API for retaining objects and determining their validity.

Animation - *ANAnimationManager* and *ANAnim* classes provide the skeleton of support for animating projections.

Picking - *ANPick* classes provide a basic implementation of picking.

Rendering - *ANRenderer* classes provide an interface to perform actual drawing operations.

Transform - *ANTransform* classes provide non-linear transformations.

Technically, the prototype engine renders to a hardware accelerated offscreen pixel buffer (known as a PBuffer). The contents of the PBuffer are then copied back into main memory as an image which is then painted into the canvas window using Java2D. Although this introduces the overhead of a slow copy back to main memory followed by another paint, it simplifies integration with the existing code. Notably, the old and new graphics engines can be mixed within an Atlantis canvas with no changes to the old code.

This design provides many advantages over the current graphics code. Firstly classes are clearly separated with well defined roles. The introduction of the "object" classes separates the graphical representation from both the actual data and choice of projection which should facilitate both reuse and flexibility. The rendering code makes it possible to take advantage of GPU hardware.

The caching code makes it possible to avoid recalculating graphics primitives whenever possible. This is achieved through the use of a cache token hierarchy. Any object can be used in conjunction with a token, which can in turn depend on further tokens and so on. When any token is marked as invalid, all tokens which depend on it also become invalid. A simple API is provided to allow projections to use this approach and to avoid recalculating geometry and primitives whenever possible.

Picking, the operation of selecting an object in a scene with a pointing device is implemented in software. An alternative rendering class *ANRendererPick* is provided which performs coordinate transformations in software and calculates how close to the pointer click they would appear.

Currently the implementation contains about 5000 lines of code in around 70 classes.

9.5   Some Examples

The code is sufficiently functional that some tests can be executed. See Figure 9.8 for an example of the new graphics code running inside Atlantis. Although missing some elements, the code is clearly functional and integrated.

All following performance measurements are calculated using a PC with an Intel Core Duo processor clocked at 1.6GHz, with an ATI X1400 graphics card. Manufactured in 2006 it does not represent an especially state-of-the-art configuration.

A simple demonstration of some of the potential of the new graphics code can be seen in Figure 9.9. The main quality difference is in the TRT hits, which is due to the fact that Java2D only accepts integer co-ordinates for drawing, which makes anti-aliasing small objects very difficult while OpenGL has no such restriction.

It takes the old graphics code approximately 500 ms to render this image with Java2D anti-aliasing enabled. The OpenGL image is rendered using 4x FSAA (Full Scene Anti-Aliasing), while the TRT hits are rendered with the use of a compiled display list. This image took only 30 ms to render, an order of magnitude faster. Profiling reveals that the limiting factor is the buffer-copying step of the rendering process, demonstrating how well the new code is able to scale to large numbers of graphics primitives.

It is also possible to explore areas where the current graphics code is difficult to

Figure 9.8: A typical usage scenario, as Figure 9.2 but with the main YX view drawn by the prototype new graphics engine.
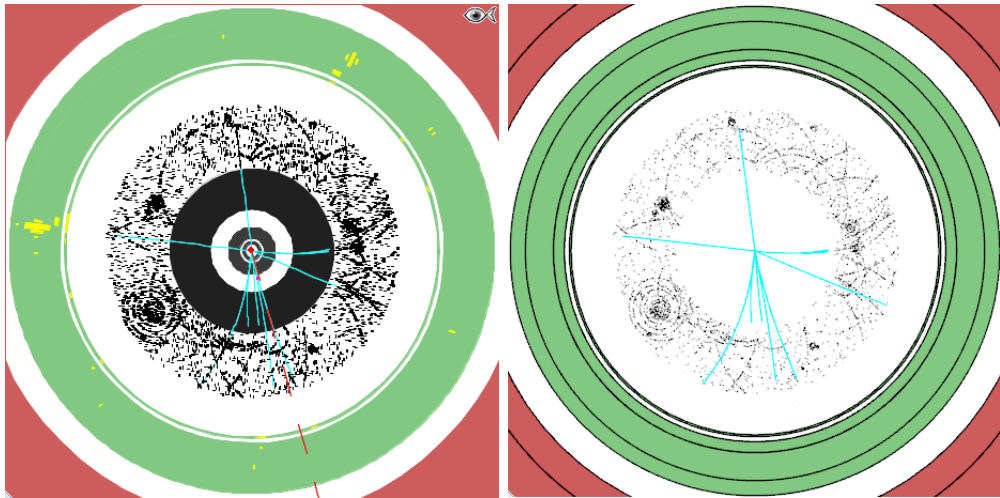
Figure 9.9: Comparison of TRT hits as drawn by the old graphics engine (left) and the new prototype (right).

extend. One frequent request is for transparent objects, which can be used to great effect in some circumstances. Such requests have gone unfulfilled because it was simply too difficult to implement with the current graphics code. Figure 9.10 shows an example of transparent objects being rendered by the prototype graphics engine.

9.6   Conclusions

The work in this chapter has clearly showed that implementing an OpenGL based graphics engine for Atlantis is a viable strategy. It has also demonstrated that such an engine can deliver the order of magnitude increase in performance required in a scalable manner. It has also demonstrated that it is possible to deliver new features which were previously unobtainable.

Assuming the ultimate goal is to replace the current Atlantis graphics code then there is much work still to be done. Although the core of the graphics code is complete, there are some areas where the design would benefit from further work. One particular open topic is how to make use of modern graphics features without excluding older machines.

Overall this work offers a clear path to solving the current problems with the Atlantis graphics engine.

Figure 9.10: Example of transparent objects being rendered. The current Atlantis graphics code does not support any kind of alpha blending.

Chapter 10

Summary and Conclusions

The ATLAS detector has the potential to make many groundbreaking measurements. In this thesis it has been shown that techniques involving jet substructure, which look at the internal structure of jets, have the potential to significantly enhance the ATLAS physics programme. Specifically, it has been shown that jet substructure techniques improve the ATLAS discovery potential for a low mass Higgs boson in the $H \rightarrow b\bar{b}$ decay mode. Previous ATLAS studies [5] had only been able to obtain a significance of 1.7σ with 30 fb^{-1} of integrated luminosity and even then only by making some very aggressive assumptions. The work in Chapters 6 and 7 showed that a jet substructure approach has potential to acheive significances of 3σ or better with 30 fb^{-1}. In Chapter 5 it has also been shown that jet substructure techniques have potential for a WW scattering measurement at ATLAS. In combination, these two results represent a significant enhancement of ATLAS' ability to understand the nature of electroweak symmetry breaking in the Standard Model. Finally in Chapter 8 it has been shown that W and Z boson production presents an opportunity to validate jet substructure techniques in real data.

Visualisation software is a key aspect of our ability to understand and use data from the ATLAS detector. In order to stay relevant, it is therefore essential for Atlantis to make good use of the available graphics technologies. The work presented in this thesis does not claim to be a complete solution to this problem but it is a serious beginning, which addresses many of the key challenges. It is my hope that this code will continue to be developed and will form a key part of Atlantis in the near future.

Bibliography

[1] B. Odom, D. Hanneke, B. D'Urso and G. Gabrielse, Phys. Rev. Lett. **97**, 030801 (2006).

[2] G. Gabrielse, D. Hanneke, T. Kinoshita, M. Nio and B. Odom, Phys. Rev. Lett. **97**, 030802 (2006).

[3] C. Amsler *et al.* (Particle Data Group), Phys. Lett. B **667** (2008, 2009 partial update for the 2010 edition).

[4] The ATLAS Collaboration, G. Aad *et al.*, arXiv:0901.0512.

[5] T. A. Collaboration, Detector and physics performance technical design report, cern/lhcc/99-15, 1999.

[6] The ATLAS Collaboration, CERN Report No. ATL-PHYS-PUB-2009-088. ATL-COM-PHYS-2009-345, 2009 (unpublished).

[7] R. P. Feynman, Rev. Mod. Phys. **20**, 367 (1948).

[8] B. A. Reid, L. Verde, R. Jimenez and O. Mena, JCAP **1001**, 003 (2010), [astro-ph/09100008].

[9] E. Majorana, Nuovo Cimento **14**, 171 (1937).

[10] T. Yanagide, Proc. Workshop on Unified Theory and Baryon Number in the Universe (KEK) (1979).

[11] M. Gell-Mann, P. Ramond and R. Slansky, Supergravity (1979).

[12] A. Faessler and F. Simkovic, J. Phys. **G24**, 2139 (1998), [hep-ph/9901215].

[13] P. W. Higgs, Phys. Rev. Lett. **13**, 508 (1964).

[14] F. Englert and R. Brout, Phys. Rev. Lett. **13**, 321 (1964).

[15] G. S. Guralnik, C. R. Hagen and T. W. B. Kibble, Phys. Rev. Lett. **13**, 585 (1964).

[16] M. S. Carena and H. E. Haber, Prog. Part. Nucl. Phys. **50**, 63 (2003), [hep-ph/0208209].

[17] J. M. Butterworth, B. E. Cox and J. R. Forshaw, Phys. Rev. **D65**, 096014 (2002), [hep-ph/0201098].

[18] LEP Working Group for Higgs boson searches, R. Barate *et al.*, Phys. Lett. B **565**, 61 (2003), [hep-ex/0306033].

[19] The TEVNPH Working Group of the CDF and D0, [hep-ex/10074587].

[20] E. L. Berger, C.-W. Chiang, J. Jiang, T. M. P. Tait and C. E. M. Wagner, Phys. Rev. D **66**, 095001 (2002).

[21] E. Boos, A. Djouadi, M. Mühlleitner and A. Vologdin, Phys. Rev. D **66**, 055004 (2002).

[22] O. S. Brüning *et al.*, LHC Design Report (2004).

[23] The ATLAS Collaboration, G. Aad *et al.*, JINST **3**, S08003 (2008).

[24] J. E. Demuth, R. J. Hamers, R. M. Tromp and M. E. Welland, IBM Journal of Research and Development **30**, 396 (1986).

[25] G. Corcella *et al.*, [hep-ph/0210213].

[26] G. Corcella *et al.*, JHEP **01**, 010 (2001), [hep-ph/0011363].

[27] T. Sjostrand, S. Mrenna and P. Skands, JHEP **05**, 026 (2006), [hep-ph/0603175].

[28] J. M. Butterworth, J. R. Forshaw and M. H. Seymour, Z. Phys. C **72**, 637 (1996), [hep-ph/9601371].

[29] G. Altarelli and G. Parisi, Nucl. Phys. **B126**, 298 (1977).

[30] V. N. Gribov and L. N. Lipatov, Sov. J. Nucl. Phys. **15**, 438 (1972).

[31] Y. L. Dokshitzer, Sov. Phys. JETP **46**, 641 (1977).

[32] T. Sjöstrand, Comp. Phys. Commun. **82**, 74 (1994).

[33] G. Marchesini *et al.*, Comp. Phys. Commun. **67**, 465 (1992).

[34] B. P. Kersevan and E. Richter-Was, [hep-ph/0405247].

[35] F. Maltoni and T. Stelzer, JHEP **02**, 027 (2003), [hep-ph/0208156].

[36] D. Costanzo *et al.*, CERN Report No. ATL-SOFT-PUB-2005-004. CERN-ATL-SOFT-PUB-2005-004. ATL-COM-SOFT-2005-008, 2005 (unpublished).

[37] J. Allison *et al.*, Nuclear Science, IEEE Transactions **53**, 270 (2006).

[38] S. Agostinelli *et al.*, Nuclear Inst. A **506**, 250 (2003).

[39] M. Dührssen, CERN Report No. ATL-PHYS-INT-2008-043. ATL-COM-PHYS-2008-093, 2008 (unpublished).

[40] D. Adams *et al.*, CERN Report No. ATL-PHYS-INT-2009-110. ATL-COM-PHYS-2009-554, 2009 (unpublished).

[41] E. Richter-Was, D. Froidevaux and L. Poggioli, CERN Report No. ATL-PHYS-98-131, 1998 (unpublished).

[42] The ATLAS Collaboration, G. Aad *et al.*, JINST **3**, S08003 (2008).

[43] R. R. Daniel, J. H. Davies, J. H. Mulvey and D. H. Perkins, Phil. Mag. **43**, 753 (1952).

[44] G. Hanson *et al.*, Phys. Rev. Lett. **35**, 1609 (1975).

[45] R. Brandelik *et al.*, Physics Letters B **86**, 243 (1979).

[46] S. Wu and G. Zobernig, Z. Phys. C **2**, 107 (1979).

[47] J. B. Babcock and R. E. Cutkosky, Z. Phys. C **15**, 133 (1982).

[48] G. Sterman and S. Weinberg, Phys. Rev. Lett. **39**, 1436 (1977).

[49] J. Dorfan, Z. Phys. C **7**, 349 (1981).

[50] H. J. Daum, H. Meyer and J. Burger, Z. Phys. C **8**, 167 (1981).

[51] G. Arnison *et al.*, Physics Letters B **123**, 115 (1983).

[52] S. Catani, Y. L. Dokshitzer, M. H. Seymour and B. R. Webber, Nucl. Phys. B **406**, 187 (1993).

[53] CDF, F. Abe *et al.*, Phys. Rev. Lett. **74**, 2626 (1995), [hep-ex/9503002].

[54] ZEUS, S. Chekanov *et al.*, Eur. Phys. J. **C42**, 1 (2005), [hep-ph/0503274].

[55] H1, F. D. Aaron *et al.*, JHEP **01**, 109 (2010), [hep-ex/09110884].

[56] ZEUS, S. Chekanov *et al.*, Nucl. Phys. B **700**, 3 (2004), [hep-ex/0405065].

[57] G. P. Salam and G. Soyez, JHEP **05**, 086 (2007), [hep-ph/0704.0292].

[58] F. Abe *et al.*, Phys. Rev. D **45**, 1448 (1992).

[59] G. C. Blazey *et al.*, [hep-ex/0005012].

[60] Y. L. Dokshitzer, G. D. Leder, S. Moretti and B. R. Webber, JHEP **08**, 001 (1997), [hep-ph/9707323].

[61] M. Wobisch and T. Wengler, [hep-ph/9907280].

[62] M. Cacciari, G. P. Salam and G. Soyez, JHEP **04**, 063 (2008), [hep-ph/08021189].

[63] M. Cacciari and G. P. Salam, Phys. Lett. B **641**, 57 (2006), [hep-ph/0512210].

[64] W. Lampl *et al.*, CERN Report No. ATL-LARG-PUB-2008-002. ATL-COM-LARG-2008-003, 2008 (unpublished).

[65] T. Appelquist and C. Bernard, Phys. Rev. D **22**, 200 (1980).

[66] DELPHI, P. Bambade, [hep-ex/0307056].

[67] A. Dobado and J. R. Pelaez, Phys. Rev. **D56**, 3057 (1997), [hep-ph/9604416].

[68] E. Stefanidis, UCL PhD Thesis, 2007.

[69] S. Allwood, Manchester PhD Thesis, 2006.

[70] R. Lafaye, T. Plehn, M. Rauch, D. Zerwas and M. Duehrssen, arXiv:0904.3866 [hep-ph].

[71] J. M. Butterworth, A. R. Davison, M. Rubin and G. P. Salam, Phys. Rev. Lett. **100**, 242001 (2008), [hep-ph/08022470].

[72] S. Corréard, V. Kostyukhin, J. Lévêque, A. Rozanov and J. B. De Vivie de Régie, Aix-Marseille 2. Cent. Phys. Part. Report No. ATL-PHYS-2004-006, 2003 (unpublished).

[73] G. Piacquadio and C. Weiser, J. Phys. Conf. Ser. **119**, 032032 (2008).

[74] G. Piacquadio, Freiburg PhD Thesis, 2010.

[75] G. Piacquadio, ATLAS Note ATL-COM-PHYS-2009-308.

[76] E. Richter-Was, Int. J. Mod. Phys. A **13**, 1371 (1996).

[77] B. M. Waugh *et al.*, [hep-ph/0605034].

[78] UA2, J. Alitti *et al.*, Z. Phys. C **49**, 17 (1991).

[79] Oracle, Java Programming Language, http://www.oracle.com/technetwork/java/index.html.

[80] The ATLAS Collaboration, G. Aad *et al.*, Latest atlas events, http://atlas-live.cern.ch.

[81] JOGL Developers, JOGL, http://www.jogamp.org.

[82] LWJGL Developers, LWJGL, http://www.lwjgl.org.