

Studies of Extra-Dimensional Models with the ATLAS Detector

Matthew John Palmer

Sidney Sussex College



A dissertation submitted to the University of Cambridge
for the degree of Doctor of Philosophy
January 2005

Studies of Extra-Dimensional Models with the ATLAS Detector

Matthew John Palmer

Abstract

This thesis describes how two extra dimensional scenarios may be investigated using the ATLAS detector that is under construction at the European Organization for Nuclear Research (CERN).

In some extra dimensional scenarios, microscopic black holes may be created in the ATLAS detector. These black holes would quickly explode and be easy to detect. However, measuring their properties would be extremely difficult due to many experimental and theoretical systematic effects. These effects are discussed and the ability of the ATLAS detector to accurately measure black hole events is investigated. Methods of determining their properties are presented including a new technique which is used to demonstrate how well the number of extra dimensions could be measured by ATLAS.

In the second scenario, heavy gravitons may be produced. Monte Carlo studies are presented that show how well two of its couplings could be measured. Measurements like these would be important in demonstrating that the graviton has universal coupling which is one of its characteristic features.

In addition, there is a summary of software that has been developed for the calibration and assembly of the semi-conductor tracker modules that comprise one of ATLAS' sub-detectors. Due to the module design, calibration involves collecting and analysing large amounts of data using several different algorithms — calibration cannot be performed after physics data has been taken. Detailed tests of the performance and algorithmic correctness demonstrate that this software will be able to accurately calibrate the thousands of semi-conductor tracker modules using a modest amount of computer hardware.

DECLARATION

This dissertation is the result of my own work, except where explicit reference is made to the work of others, and has not been submitted for another qualification to this or any other university. The main body of this dissertation consists of less than 60,000 words.

Matthew John Palmer

ACKNOWLEDGEMENTS

I would like to thank the Particle Physics and Astronomy Research Council (PPARC) for providing financial support during this degree (under PPARC student grant number PPA/S/S/2001/03194). I am also very grateful for financial support from the Stokenchurch Educational Charity and Sidney Sussex College.

Much of the work in this thesis has used ATLAS software including the physics analysis framework and tools which are the result of collaboration-wide efforts.

I take this opportunity to thank the rest of the SctRodDaq software team: Alan Barr, Bruce Gallop, Dave Robinson and Alessandro Tricoli and also to those who helped in other ways, particularly: Janet Carter, John Hill, Gareth Moorhead and Peter Phillips. I would also like to thank my collaborators in the Cambridge Supersymmetry Working Group, especially: Ben Allanach, Chris Harris, Christopher Lester, Ali Sabetfakhri, Bryan Webber and in particular, my supervisor, Andy Parker. I owe a special word of thanks to Rachael who helped so much and had to answer too many questions about boring things but who is still wrong about the histograms. Finally, I thank my parents for their invaluable support.

Many animals have been harmed in the making of this thesis. They include cows, sheep, pigs, chicken, fish (many varieties), a goose, partridge, deer, turkeys, ducks and many others. I am deeply grateful to you all.

This thesis is copyright © 2005 Matthew John Palmer.

Thesis quotes are from the Bible, New International Version © International Bible Society.

Chapter quotes are from Monty Python's Holy Grail © National Film Trustee Company Limited.

*The heavens declare the glory of God;
the skies proclaim the work of his hands.
Day after day they pour forth speech;
night after night they display knowledge.
There is no speech or language
where their voice is not heard.
Their voice goes out into all the earth,
their words to the ends of the world.*

CONTENTS

Preface	1
1 Theory and Motivations	5
1.1 The Standard Model	5
1.1.1 Higgs mechanism	6
1.1.2 W–W scattering	8
1.1.3 Hierarchy problem	10
1.2 Extra dimensions	11
1.2.1 Kaluza-Klein towers	12
1.2.2 ADD models	13
1.2.3 Randall-Sundrum models	19
1.3 Summary	23
2 ATLAS	25
2.1 The Large Hadron Collider	25
2.2 The ATLAS detector	26
2.2.1 Magnet systems	29
2.2.2 Inner detector	31
2.2.3 Calorimeters	33
2.2.4 Muon spectrometer	36
2.2.5 Trigger and data acquisition	38
2.3 Simulation software	38
2.3.1 Full simulation	40
2.3.2 Fast simulation	40
2.4 Current status	41
3 The SCT and its Readout Systems	43
3.1 System design	43
3.2 Detector modules	44
3.3 Readout systems	46
3.3.1 Optical readout system	46
3.4 Module characterisation and calibration	50

4	SctRodDaq	53
4.1	Aims	53
4.2	Overview	55
4.2.1	Functional overview	56
4.2.2	ATLAS online software	58
4.2.3	Data flow	61
4.3	Design and implementation	63
4.3.1	General principles	63
4.3.2	External software	64
4.3.3	I/O system	64
4.3.4	Threading	66
4.4	Fitting Service	68
4.4.1	Algorithm	68
4.4.2	Optimisation	72
4.5	Analysis Service	74
4.5.1	Framework	74
4.5.2	Digital tests	74
4.5.3	Strobe delay	75
4.5.4	Trim algorithm	76
4.5.5	Response curve	78
4.5.6	Noise occupancy	79
4.5.7	Time walk	81
4.6	Testing systems	82
4.6.1	Unit tests	82
4.6.2	High-level tests	84
4.6.3	Low-level tests	84
4.6.4	System tests	84
4.6.5	Other tests	85
4.7	Correctness testing	85
4.7.1	Threshold scan comparisons	86
4.7.2	Noise occupancy	88
4.7.3	Response curve	89
4.7.4	Strobe delay	90
4.7.5	Trim calibration	91
4.7.6	Time walk	93
4.7.7	Digital tests	94
4.8	Performance studies	94
4.8.1	Fitting Service	94
4.8.2	Analysis Service	96
4.8.3	System performance	97
4.9	Macro assembly	100
4.10	Future directions	102
4.11	Summary	103

5	Studies of Graviton Decays to Heavy Vector Bosons	105
5.1	Background	105
5.2	Event generation	106
5.2.1	Signals	106
5.2.2	Backgrounds	108
5.3	Reconstruction	109
5.3.1	$W \rightarrow \ell \nu_\ell$ reconstruction	109
5.3.2	$W/Z \rightarrow jj$ reconstruction	111
5.3.3	Graviton reconstruction	112
5.4	Cuts	112
5.4.1	Mass and p_T cuts	113
5.4.2	Top reconstruction and jet veto	116
5.4.3	Summary	116
5.5	Background subtraction	118
5.6	Results	119
5.6.1	Determination of errors	120
5.6.2	Contours of $\sigma \cdot B$	121
5.7	Wider context	123
5.8	Conclusions	123
6	Exploring Higher Dimensional Black Holes	125
6.1	Black hole production and decay	125
6.1.1	Black hole production	126
6.1.2	Black hole decay	127
6.2	CHARYBDIS: A black hole event generator	128
6.2.1	Black hole production cross section	130
6.3	Model uncertainties	131
6.3.1	Production cross section	131
6.3.2	The first stages of decay	132
6.3.3	Energy deposition on the brane	132
6.3.4	Back reaction and the effect of the kinematic limit	132
6.3.5	Remnant decay	133
6.3.6	Time-variation and black hole recoil	134
6.4	Event generation and detector simulation	136
6.5	Corrections to the full simulation	138
6.5.1	Muon \cancel{p}_T adjustment	139
6.5.2	EM–Jet correction	140
6.5.3	Jet energy $ \eta $ dependence	140
6.6	Performance of the full simulation	141
6.6.1	Electron and photon performance	142
6.6.2	Muon performance	143
6.6.3	Jet performance	145
6.6.4	\cancel{p}_T resolution	145
6.7	Characteristics of black hole decays	147

Contents

6.7.1	\cancel{p}_T distribution	150
6.7.2	Black hole charge	150
6.8	Measurement of the black hole mass	152
6.9	Previous attempts to determine the model parameters	154
6.10	Correlation Analysis	155
6.11	Analysis Strategy	158
6.12	Measurement of the black hole temperature	159
6.12.1	Kinematic limit	159
6.12.2	Energy Distributions	168
6.12.3	Discussion	172
6.13	Measurement of the parton-level cross section	173
6.14	Determination of the model parameters	175
6.15	Conclusions	175
7	Summary	179
	Abbreviations	181
	References	185

LIST OF FIGURES

1.1	The Higgs mass as predicted from precision experiments	9
1.2	W–W scattering diagram.	9
1.3	Diagrams that contribute to the correction of the bare Higgs mass.	10
1.4	DØ limits on RS models	22
1.5	The allowed space for the RS model	23
2.1	Schematic of the LHC layout.	26
2.2	A schematic diagram of the ATLAS detector	27
2.3	Higgs discovery potential with a total luminosity of 30 fb^{-1}	28
2.4	The magnet system in ATLAS	30
2.5	The inner detector	31
2.6	ATLAS calorimeters	34
2.7	Geometry of the ECAL	35
2.8	End view of the muon spectrometer	37
2.9	Trigger overview	39
2.10	The ATLAS cavern	42
3.1	The SCT layout	43
3.2	Photograph of a barrel module	45
3.3	Simplified block diagram of the ABCD3TA chip.	46
3.4	Block diagram of the optical readout system for the SCT	47
3.5	Charge distribution schematics for SCT modules	51
4.1	Schematic overview of the SctRodDaq software.	56
4.2	A screenshot of the SctRodDaq GUI	59
4.3	The analysis sub-system data chain.	62
4.4	Description of the SctRodDaq I/O system	65
4.5	A schematic diagram of the threading model	67
4.6	Example of strobe delay data and fit	70
4.7	Example of threshold scan data and fit	72
4.8	Example digital scan data	76
4.9	Mean threshold before and after trimming	77
4.10	Some typical response curve plots	80
4.11	Example noise occupancy data	81
4.12	Comparison of S-curve fitting showing the mean threshold	87
4.13	Comparison of the S-curve fitting showing the width	88

List of Figures

4.14	Comparison of the noise occupancy algorithm	89
4.15	Comparison of the response curve algorithm	90
4.16	Comparison of the strobe delay algorithm for different SCTDAQ versions	91
4.17	Performance of the original trim range algorithm	92
4.18	Performance of the new trim range algorithm	93
4.19	Time walk comparison between SCTDAQ and SctRodDaq	94
4.20	CPU scaling of the Fitting Service	95
4.21	CPU scaling of the Analysis Service	96
4.22	System performance graphs as functions of the number of modules	98
4.23	Noise (in ENC) versus position for barrel 3	100
4.24	The completed barrel 3	101
5.1	Comparison of two methods for setting the neutrino z momentum	110
5.2	Reconstructed $W \rightarrow jj$ mass	112
5.3	Graviton mass resolution in the $G \rightarrow W^+W^- \rightarrow \ell\nu_\ell jj$ channel.	113
5.4	Graviton mass resolution in the $G \rightarrow Z^0Z^0 \rightarrow \ell\ell jj$ channel.	113
5.5	The signal reconstruction efficiency	117
5.6	Background shapes after all cuts	118
5.7	Background subtraction schematic	119
5.8	Background subtraction example	120
5.9	Contours showing the error in $\sigma.B$ for the $G \rightarrow W^+W^- \rightarrow \ell\nu_\ell jj$ channel	122
5.10	Contours showing the error in $\sigma.B$ for the $G \rightarrow Z^0Z^0 \rightarrow \ell\ell jj$ channel	122
5.11	Contours showing the statistical precision other graviton decay channels	124
6.1	Energy of the generator level decay products in the black hole rest frame	134
6.2	Photon energy distributions	135
6.3	Temperature versus black hole mass for fixed and time dependent temperatures	136
6.4	Full simulation \cancel{p}_T bias before and after correction	139
6.5	Jet p_T resolution before correction	141
6.6	Full simulation electron and photon resolutions	142
6.7	Full simulation muon resolution	144
6.8	Jet performance graphs for the full simulation.	146
6.9	Jet performance graphs for the fast simulation	146
6.10	\cancel{p}_T resolution comparison between full and fast simulations	147
6.11	Black hole event multiplicity	149
6.12	$\sum p_T$ distributions for black hole events	150
6.13	\cancel{p}_T distribution for black hole events	151
6.14	Average charge of electrons and muons in black hole events	151
6.15	Black hole mass resolution	152
6.16	Full simulation black hole mass resolution	154
6.17	Correlation plots	156
6.18	p_t of generator level decay particles	157
6.19	Temperature variation for fixed normalisation	158
6.20	Schematic of the energy spectrum near the kinematic limit	161

6.21	Upper and lower bounds on p	162
6.22	Bias in E_{max} measurement	164
6.23	Kinematic limit analysis for different models	165
6.24	Kinematic limit analysis for different n	166
6.25	Kinematic limit analysis for full simulation data	167
6.26	Temperature measurement for the test case	168
6.27	Electron energy distribution with several Planck distributions	169
6.28	Energy distributions with fits to the Planck spectrum	171
6.29	The parton-level cross section as a function of M_{PL} for different n	174
6.30	Model parameters for the test case	175

LIST OF TABLES

1.1	Fermionic particle content of the Standard Model.	7
1.2	Masses of the Standard Model particles.	8
5.1	Details of the signal samples for the W and Z channels	107
5.2	Background sample information	109
5.3	Effects of the cuts for the $G \rightarrow W^+W^-$ channel	115
5.4	Effects of the cuts for the $G \rightarrow Z^0Z^0$ channel	115
5.5	The efficacy of the top cuts	117
6.1	Particle emission probabilities	129
6.2	Black hole production cross sections at the LHC from CHARYBDIS.	130
6.3	Black hole mass resolution and efficiency	153
6.4	Temperatures from electron, muon and photon energy distributions for different n	171
6.5	Energy distribution results for different models	172

PREFACE

Extra-dimensional models, introduced in chapter 1, are some of the more recent and exciting ideas for extensions beyond the Standard Model. These models could give a way of understanding the hierarchy problem (see section 1.1.3) and typically predict strong gravitational effects in the TeV energy range. Such a prospect is very exciting as it would allow quantum gravity to be probed at the next generation of high energy colliders.

Of course, whilst such models are intellectually interesting, to be physically meaningful, they must be confirmed or refuted with experiments. The LHC that is currently under construction at the European Organization for Nuclear Research (CERN) will collide two 7 TeV beams of protons together in four large experiments. One of these is ATLAS which is a general purpose detector that aims to discover and measure whatever physics there may be in the TeV energy range. The construction of this experiment, introduced in chapter 2, (and the other general purpose detector, CMS) is clearly of great importance to high-energy physics (HEP). To this end, I have aided in the assembly of one of the tracking components of ATLAS, the semiconductor tracker (SCT) — described in chapter 3. I am one of five authors^a who wrote the software package SctRodDaq that provides data acquisition and calibration for SCT detector modules. Chapter 4 gives an overview of this software and a description of the components I was most involved with. A series of correctness and performance tests are presented that demonstrate the software is fit-for-purpose. SctRodDaq is now in use at three macro-assembly sites and also at CERN where the entire SCT will be assembled.

The last two chapters return to extra-dimensional models. In particular, whilst it is important and interesting to show how well ATLAS would be able to discover new particles and models,

^aThe others are: A. J. Barr, B. J. Gallop, D. Robinson and A. Tricoli.

Preface

it is also important to determine how well ATLAS would be able to measure the properties of new particles and phenomena. Chapter 5 is a study of how well two of the couplings of massive, narrow graviton resonances, predicted in Randall-Sundrum (RS) models, could be measured. This is particularly pertinent for a graviton resonance as gravitons are expected to have universal couplings as a model-independent feature. Such model-independent features can be very important in these models, as the theoretical details have generally not been fully worked out. Combining measurements of the graviton's couplings to W^+W^- and Z^0Z^0 bosons with measurements of other couplings would allow a test of this prediction and some discrimination against other possible massive neutral resonances. The work in chapter 5 has been published in an abbreviated form in [1].

Although extra-dimensional models are interesting, their relative newness means that they are not fully developed theoretically. In part this is related to the need for a quantum theory of gravity which is not available currently (but their sensitivity to this also makes them very interesting!). These theoretical difficulties are very apparent in the topic of chapter 6 — microscopic black holes. The production of microscopic black holes is one of the possibilities of large extra-dimensional modes, first suggested by N. Arkani-Hamed, S. Dimopoulos and G. R. Dvali (ADD). However, there are many theoretical uncertainties relating to how they would be produced and how they would decay. In fact, it turns out that discovering black holes should be reasonably straightforward, but going beyond a discovery and making any sort of measurement is significantly harder. Chapter 6 discusses the theoretical uncertainties and considers a number of different ways (some of which are novel) of measuring the properties of black holes and includes realistic detector effects. This study is particularly pertinent as some studies have been published which do not consider theoretical uncertainties or experimental effects and are thus wildly optimistic [2]. For one new technique, it is shown that a measurement of the number of extra dimensions and the Planck mass is possible accounting for some of the theoretical effects and also the experimental effects. On the experimental side, there was some concern that the fast simulation used to model the detector would not deal well with black hole events as they tend to be rather extreme. So two

small samples of fully simulated events have been generated and they are used to back up the fast simulation results. The results of chapter 6 have been submitted to the Journal of High Energy Physics (JHEP) [3].

Theory and Motivations

*Supposing two swallows carried it together?
No, they'd have to have it on a line.
Well, simple! They'd just use a strand of creeper!
What, held under the dorsal guiding feathers?
Well, why not?*

1.1 The Standard Model

The Standard Model is one of the great successes of 20th century physics. Its development began shortly after the end of World War II and the bulk of the Standard Model was complete by the late 1970's. The next two decades saw a period of confirmation and of precision measurements of the Standard Model parameters. It has been extensively tested, in some cases to extremely high precision.

The only unconfirmed part of the Standard Model is the description of mass. Recent observations have confirmed, via neutrino oscillations, that neutrinos *do* have mass [4, 5]. The mechanisms via which their tiny masses (the electron neutrino has a mass less than 3 eV) are generated is still controversial. Their masses are, however, so small (and unknown) that neutrinos will be taken to be massless in this thesis. The mechanism via which the other Standard Model particles get their masses has also not been confirmed although there are strong theoretical prejudices for believing that this occurs via the Higgs mechanism, briefly explained in section 1.1.1.

The Standard Model describes all matter and interactions as due to point-like particles. Matter consists of particles called fermions that have internal spin-1/2. Forces between fermions are mediated by integer spin particles, bosons, which allow the Standard Model Lagrangian to be invariant under local gauge transformations. These symmetries are an important and intrinsic part of the

Theory and Motivations

Standard Model. There are three main symmetries.

There is the U(1) hypercharge symmetry with the associated gauge boson field, B . The conserved charge is hypercharge, Y . Next there is SU(2) isospin symmetry which couples only to left-handed fermions. There are three associated gauge fields: W_1 , W_2 and W_3 and the conserved charge is I_3 : the third component of the isospin. These first two symmetries are described by the electroweak theory. Finally there is the SU(3) colour symmetry described by quantum chromodynamics (QCD) which has three conserved colour charges ('red', 'green' and 'blue'). There are 8 gauge bosons, gluons, which carry the force, called the strong force and these also have colour.

During electroweak symmetry breaking, the W_3 and B fields mix to give the photon and Z^0 bosons which are the observed particles. The charge of the remaining U(1) symmetry is the standard electromagnetic charge, $Q = Y + I_3$.

The particle content consists of six quarks which are the only particles to feel the strong force. There are also six fermions which do not feel the strong force. Three carry electromagnetic charge and the other three are the neutrinos which interact only very weakly with matter.

The particle content and their gauge charges are given in table 1.1. An interesting feature of the Standard Model is that the fermionic particles come in three generations, with each generation similar but with higher masses. Normal matter is made up of only the lightest generation: the u and d quarks and electrons.

1.1.1 Higgs mechanism

Particle masses cannot be introduced into the Standard Model directly because mass terms would break the gauge symmetries. The Higgs mechanism [6, 7, 8] was introduced by Higgs, Englert and Brout as a way of generating masses whilst maintaining the high energy symmetry in the Lagrangian. This is achieved by adding two complex scalar fields (thus with four degrees of freedom) to the Standard Model. At low energies these fields have a non-zero vacuum expectation value and Higgs, Englert and Brout showed that this gives rise to electroweak symmetry breaking. This process causes the B and W_3 fields to mix giving the photon and Z^0 bosons and also causes

1.1 The Standard Model

Fermions	Colour	Isospin	I_3	Y	Q
$\begin{pmatrix} u \\ d \end{pmatrix}_L, \begin{pmatrix} c \\ s \end{pmatrix}_L, \begin{pmatrix} t \\ b \end{pmatrix}_L$	$\mathbf{3}$	$1/2$	$1/2$ $-1/2$	$1/6$	$2/3$ $-1/3$
u_R, c_R, t_R	$\bar{\mathbf{3}}$	0	0	$2/3$	$2/3$
d_R, s_R, b_R	$\bar{\mathbf{3}}$	0	0	$-1/3$	$-1/3$
$\begin{pmatrix} \nu_e \\ e \end{pmatrix}_L, \begin{pmatrix} \nu_\mu \\ \mu \end{pmatrix}_L, \begin{pmatrix} \nu_\tau \\ \tau \end{pmatrix}_L$	$\mathbf{1}$	$1/2$	$1/2$ $-1/2$	$-1/2$	0 -1
e_R, μ_R, τ_R	$\mathbf{1}$	0	0	-1	-1

Table 1.1: Fermionic particle content of the Standard Model.

three of the degrees of freedom to be ‘eaten’ by the W^+ , W^- and Z^0 bosons which acquire mass as a result. The remaining degree of freedom becomes the massive Higgs scalar. The fermions then acquire mass through their interactions with the Higgs field, their coupling being proportional to their mass. Note that although neutrinos can also acquire mass through this mechanism, it is not clear whether this is the case or not. Table 1.2 lists the masses of the Standard Model particles.

The Higgs mass is not predicted by the Standard Model, rather it is a parameter. However, it enters many electroweak observables through radiative corrections. Recently $D\mathcal{O}$ updated their measurement of the top mass by re-analysing their data using a new technique. This gave a much improved error and a slightly higher mass [10]. As a result the world average of the top mass has increased to 178.0 ± 4.3 GeV. The new electroweak precision fit is given in figure 1.1; this fit gives the most probable Higgs mass as 117 GeV, up from previous best fit value of 96 GeV, which has been experimentally excluded. This new measurement therefore gives hope that the Higgs particle will indeed be found at the LHC and ATLAS.

(a) Quarks		(b) Leptons		(c) Bosons	
Particle	Mass (MeV)	Particle	Mass (MeV)	Particle	Mass (GeV)
u	1.5–4.5	e	0.511	γ	0
d	5–8.5	μ	106	W	80.4
s	80–155	τ	1780	Z	91.2
c	1.0–1.4 GeV	ν_e	<3 eV	g	0
b	4.0–4.5 GeV	ν_μ	<0.19	H	>114
t	178±4 GeV	ν_τ	<18		

Table 1.2: Masses of the Standard Model particles. The Higgs particle has not been discovered, the mass given reflects the current limit from direct discovery experiments. Masses are taken from [9] and are given to 3 s.f.; errors are not given where they are smaller than this.

1.1.2 W–W scattering

In the Standard Model with no Higgs boson, the amplitude for W–W scattering increases with energy until at about 1.2 TeV it violates unitarity. Since this cannot be true, there must be some deviation from what has already been experimentally determined in the energy range of the LHC. Whilst measuring the W–W scattering cross section is difficult, it can be done at ATLAS in processes such as in figure 1.2. This has been studied for some very general scenarios [11] and the results suggest that ATLAS should be able to measure an important new experimental fact regardless of whether any of the theories that have been put forward are correct^a.

Of course, W–W scattering also provides strong theoretical grounds for favouring the Higgs mechanism. Introducing the Higgs boson corrects the scattering amplitude so that the unitarity bound is not violated.

^aAssuming that ATLAS and the LHC are built to specification.




Figure 1.1: The Higgs mass as predicted from precision experiments. Includes the result of the updated top mass, which moves the most probable mass out of the excluded region. From [10].




Figure 1.2: W–W scattering diagram.

1.1.3 Hierarchy problem

The hierarchy problem may be expressed in two ways. The first, known as the aesthetic hierarchy problem, notes that there appear to be two fundamental scales in particle physics: the electroweak scale (mass ~ 100 GeV) and the Planck scale (mass $\sim 10^{19}$ GeV). Why then are these two scales so different? Alternatively, this question could be phrased as: “Why is gravity so weak compared to the other forces?”.

The second hierarchy problem, the technical hierarchy problem, stems from the Higgs mass. Since the Higgs mechanism is responsible for electroweak symmetry breaking, it would seem natural that the Higgs mass be at the electroweak scale (note that the Higgs mass is a parameter of the Standard Model). Indeed the precision electroweak data presented in the previous section suggest that the Higgs is of this scale. However the bare Higgs mass is modified by corrections due to one-loop corrections such as those in figure 1.3. These corrections are quadratically divergent and must be regulated by a cutoff, Λ , that limits the integral over the loop momentum.




Figure 1.3: Diagrams that contribute to the correction of the bare Higgs mass.

The resulting correction for fermion loops is of the form

$$\delta M_{\text{H}}^2 = \frac{|\lambda_f|^2}{16\pi^2} (-2\Lambda^2 + 6m_f^2 \ln(\Lambda/m_f) + \dots) \quad (1.1)$$

where λ_f is the coupling to the fermion and m_f is the fermion mass. The bosonic correction is similar but with the opposite sign for the Λ^2 term. If the Standard Model is valid up to Planck scale energies then the natural scale for Λ would be the Planck scale. Therefore, to achieve a Higgs mass of order 100 GeV, the bare mass would have to be finely tuned to the level of 1

part in 10^{16} . This is the technical hierarchy problem.

1.2 Extra dimensions

Physics beyond the Standard Model can be motivated (at least, from this perspective) by the hierarchy problem or attempts to find a common description of gravity and the Standard Model. One of the most widely studied models beyond the Standard Model is supersymmetry which solves the hierarchy problem.

Recall that the quadratic divergences in the Higgs mass for fermions and bosons have opposite sign. Supersymmetry uses this to solve the hierarchy problem. It says that every fermion has a supersymmetric bosonic partner and vice versa. The quadratic divergence for each pair of particles will therefore cancel^b. So far, so good, however, supersymmetry must be broken otherwise the supersymmetric partners would have the same mass as the ‘normal’ particles — and we have not yet observed any supersymmetric particles^c. This supersymmetry breaking is not predicted and worse, introduces ~ 100 new parameters to the theory. This is a heavy price to pay for the solution of the hierarchy problem. Also, note that supersymmetry does not solve the aesthetic hierarchy problem — it is still an effective theory that would need to be replaced at very high energies.

However, there is another way of solving the hierarchy problem. The Higgs mass can be stabilised if the cutoff that is introduced into the integrals is not at the Planck scale. If instead the scale of new physics is reduced to ~ 1 TeV then the hierarchy problem is solved. Extra dimensional theories that attempt to solve this problem^d do so essentially by showing that the Planck scale as we observe it is not the fundamental scale of gravity. Rather, gravity only appears to be weak due to the existence of one or more extra dimensions. The scale at which quantum gravity becomes important is thus reduced and both forms of the hierarchy problem are solved^e.

^bThis still leaves a logarithmic divergence, but this is much less problematic.

^cThe bosons in the Standard Model cannot be the supersymmetric partners of (some) of the fermions due to conflicts with precision electroweak data.

^dThere are extra dimensional theories that do not attempt to solve the hierarchy problem.

^eSome would argue that new questions arise such as “Why are there extra dimensions?” and “Why this geometry?”. It can equally be argued that these questions are of the same order as the question: “Why is our universe so big?”, and no current theory provides any sort of answer to that question.

There are many such extra dimensional models now, however, they have some common features. The extra dimensions are normally compactified so that their size is quite small. This is the reason why we do not observe the extra dimensions on everyday scales^f. Typically the Standard Model fields are restricted to a $(3 + 1)$ dimensional slice of the whole space, referred to as a brane. This is not always the case, however, and sometimes they are allowed to extend into the whole space, or bulk. Gravitons can always propagate into the bulk.

One of the most exciting aspects of these theories is that, if true, they would be testable at the LHC. Not only would the extra dimensional aspect be testable, but quantum gravity would also be observed (although predictive theories regarding that are somewhat lacking). Indeed, they open up the possibility that *trans*-Planckian energies would be available in a collider. One of the implications of this possibility is the subject of chapter 6.

Two of these extra dimensional models are discussed below in more detail. First however, it is useful to discuss a generic feature of compactified extra dimensions: Kaluza-Klein towers.

1.2.1 Kaluza-Klein towers

The idea of using extra dimensions to unify physical theories is not new. In 1919, Theodor Kaluza introduced a fifth dimension in an attempt to unify Einstein's theory of gravity with electromagnetism. However, his fifth dimension was essentially identical to the other spacial dimensions with no explanation for why it is not observed. Oskar Klein solved this by modifying Kaluza's idea in 1926 to make the extra dimension compacted. Ultimately their attempts to unify gravity and electromagnetism failed, but the idea of extra, compactified dimensions remained and has been revived recently in the models that will be presented shortly.

One interesting feature of compactified extra dimensions is they inevitably lead to a 'tower' of masses for any field that is allowed to propagate in the extra dimensions. To see how this arises it is useful to follow the example in [12]. Consider a 5 dimensional scalar field, $\Psi(t, x, y, z, \phi)$ where ϕ is the co-ordinate in the extra dimension. If the extra dimension is compactified on a circle of

^fExtra dimensions as small as 10^{-18} m could be observed in HEP experiments if the Standard Model fields could propagate into them.

radius, R , then $\phi = \phi + 2\pi R$. The 5 dimensional field equation is

$$(\partial_5^2 + m^2)\Psi = 0 \quad (1.2)$$

where $\partial_5^2 = \frac{d^2}{dt^2} - \frac{d^2}{dx^2} - \frac{d^2}{dy^2} - \frac{d^2}{dz^2} - \frac{d^2}{d\phi^2}$. Since solutions must be periodic in ϕ , the solution must be of the form

$$\Psi = \sum_{k \in \mathbb{N}} \Phi_j(t, x, y, z) \exp\left(i \frac{k\phi}{R}\right). \quad (1.3)$$

The field equation is thus

$$\sum_{k \in \mathbb{N}} \left(\partial^2 + m^2 + \frac{k^2}{R^2} \right) \Phi_k = 0 \quad (1.4)$$

which looks very much like a standard 4 dimensional field equation, but with a tower of mass states with a separation determined by the radius of compactification.

It is often the case that the 3-brane is rigid in extra dimensional models. In this case, since the brane breaks translational invariance in the extra dimension, there is no need to conserve momentum in that direction. Thus it is possible to have signatures like $gq \rightarrow Gq$ where the final quark remains on our brane but the graviton moves off into the extra dimension. Achieving such a rigid brane is normally part of the detail of the implementation of the model.

1.2.2 ADD models

In 1998, Arkani-Hamed, Dimopoulos and Dvali (ADD) [13, 14, 15] revived interest in extra dimensional models by proposing a model with n extra dimensions all compacted on circles with radius R . Only the gravitational field propagates into the extra dimensions, the Standard Model fields are restricted to the normal 3-brane[§].

Their motivation was to solve the hierarchy problem by reducing the Planck scale to the weak scale. How this is achieved can be seen by considering the gravitational potential when the distance between two masses is very small or very large compared with the size of the extra dimensions. If

[§]The mechanism via which this occurs depends on the way the model is built.

Theory and Motivations

two masses, m_1 and m_2 , are separated by r with $r \ll R$ then the potential is

$$V(r) \sim G_{4+n} \frac{m_1 m_2}{r^{n+2}} \quad (1.5)$$

where G_{4+n} is the gravitational constant in $(4+n)$ dimensions. However, if $r \gg R$ then the potential will be the normal Newtonian one

$$V(r) \sim G_4 \frac{m_1 m_2}{r^2}. \quad (1.6)$$

Requiring that these match at the boundary suggests that $G_{4+n} \sim G_4 R^n$ and therefore, given that $G_{4+n} \sim M_{\text{PL}}^{-(n+2)}$, that

$$M_{\text{PL}}^2 \sim M_{\text{PL}(n+4)}^{n+2} R^n. \quad (1.7)$$

Thus, if there are enough large extra dimensions (large compared to the $(4+n)$ dimensional Planck scale) then the fundamental $(4+n)$ dimensional Planck scale can be reduced from 10^{19} GeV to the weak scale. The weakness of gravity is thus seen as a result of the extra dimensions in which it can propagate rather than any inherent weakness.

By requiring that the fundamental Planck scale be at the electroweak scale, the size of the extra dimensions can be estimated. This gives $R \sim 10^{13}$ cm for $n = 1$ which is clearly ruled out. However, for $n = 2$, R is ~ 0.1 – 1 mm and for higher n they are even smaller. Direct tests of gravity have now reached down to the $100 \mu\text{m}$ level and thus extra dimensions of this type cannot be directly ruled out for $n > 2$. Experimental limits on ADD models will be covered in section 1.2.2.2.

Since gravity can propagate into the bulk, a Kaluza-Klein tower of graviton states emerges in much the same way as presented in section 1.2.1. In this case, the fundamental particle, the spin-2 graviton, is massless and there is a tower of excited modes with masses $m_k = k/R$, which are normally referred to as gravitons. Since R is typically quite large in ADD models, these masses are closely spaced and can be considered to form a continuum.

It is interesting to consider the rate of typical processes which create gravitons. From a 4 dimensional perspective, each graviton will have a coupling $\sim 1/M_{\text{PL}}$. Although this coupling is

very small, there are many possible gravitons to couple to, so the combined effect can be large. If there is an energy, E , available to the graviton, then the number of graviton modes that can interact is $(ER)^n$ and thus the interaction rate to all gravitons is

$$\Gamma \sim \frac{(ER)^n}{M_{\text{PL}}^2} \quad (1.8)$$

which gives

$$\Gamma \sim \frac{E^n}{M_{\text{PL}(4+n)}^2} \quad (1.9)$$

using equation 1.7. This is not surprising, since from a $4+n$ dimensional perspective, there is only one graviton, but it will have a coupling $\sim 1/M_{\text{PL}(4+n)}$ giving the same rate as above. This is a good example of a general principle: that processes may be considered in either the effective 4 dimensional theory, in which case there is a tower of graviton modes; or in the $4+n$ dimensional theory in which case there is a single, strongly coupled, massless graviton.

Of course, the discussion above glosses over many of the technical details, however it does illustrate the most important aspects of the theory. It is also not enough merely to postulate these extra dimensions, but models must be built to show that this can indeed be done. This has been done in several different ways for ADD models but the details are not important here since the most important phenomenology of ADD models does not depend on these details. Limits on ADD models will be discussed shortly, but first a word on conventions.

1.2.2.1 Conventions

There are many different conventions for the Planck mass, even in four dimensions. A useful review of the standard conventions can be found in [12, section 1.5.2]. An important point is that the conventions are not all related by constants of order 1 — some of the relations are n -dependent. In particular one common convention by Dimopoulos and Landsberg [2] is

$$M_{\text{DL}}^{n+2} = \frac{1}{G_4(2\pi R)^n}. \quad (1.10)$$

Theory and Motivations

Another common convention used by ADD [15] and Giddings and Thomas [16] is

$$M_{\text{ADD}}^{n+2} = \frac{1}{8\pi G_4 R^n}. \quad (1.11)$$

The relationship between the two conventions is therefore

$$M_{\text{ADD}}^{n+2} = 2^{n-2} \pi^{n-1} M_{\text{DL}}^{n+2} \quad (1.12)$$

which is increasingly different from M_{DL} , varying from $M_{\text{ADD}} = 1.3M_{\text{DL}}$ for $n = 2$ to $M_{\text{ADD}} = 2.9M_{\text{DL}}$ for $n = 6$. Since most studies (including the ones presented here) have been done at fixed M_{PL} whilst varying n , care should be taken when comparing results where different conventions have been used. The convention used throughout this thesis is that of Dimopoulos and Landsberg which is hereafter referred to as M_{PL} (which thus refers to the $(4 + n)$ dimensional Planck scale). Where appropriate, limits and other results have been converted to this convention.

1.2.2.2 Experimental limits

More detailed overviews of the experimental limits can be found in [9, 12, 17]. The main limits are updated and summarised here using the convention that is used throughout this thesis.

Direct limits from short-scale gravity experiments: Until recently, there were no direct tests of the inverse square nature of gravity below about 1 mm. However a modified version of Cavendish's torsion experiment has been carried out that is sensitive below this level. The EOT-WASH collaboration found an upper limit on the size of the extra dimensions of $150 \mu\text{m}$ [18, 19]. This corresponds to a limit of $M_{\text{PL}} > 1.5 \text{ TeV}$ for $n = 2$. For $n > 2$, the extra dimensions are too small to be detected in mechanical experiments.

Collider searches: The main constraints from colliders come from searches for direct production of gravitons. In these signals, the gravitons are not detected but are observed as missing

energy. These signatures can be distinguished from other models with missing energy (such as supersymmetry) since the gravitons do not have a single mass, but instead form a continuum.

Currently the most competitive measurements from LEP come from the L3 collaboration [20], although the results are similar for the other experiments [21, 22, 23, 24]. Their limits are: $M_{\text{PL}} > 1.3 \text{ TeV}$, $M_{\text{PL}} > 0.46 \text{ TeV}$, $M_{\text{PL}} > 0.25 \text{ TeV}$ for $n = 2, 4$ and 6 respectively.

The limits from the two Tevatron experiments are not as good for low values of n , but are competitive for high values [25, 26]. The limits from CDF are $M_{\text{PL}} > 0.95 \text{ TeV}$, $M_{\text{PL}} > 0.4 \text{ TeV}$, $M_{\text{PL}} > 0.28 \text{ TeV}$ for $n = 2, 4$ and 6 respectively; the limits from $D\bar{D}$ are slightly worse.

Finally, it is possible to place limits on M_{PL} from the effects of virtual gravitons on precision measurements. However these limits are always model dependent and it is normally possible to construct the model in such a way as to avoid them. Typical values for these limits are $M_{\text{PL}} \gtrsim 1 \text{ TeV}$. These limits are summarised in [27].

Cosmic rays: Limits from cosmic rays arise from the non-observation of black holes and the accompanying reduction in neutrino events. These put limits of: $0.4\text{--}0.6 \text{ TeV}$ for $n = 5$ [28] (depending on the assumptions made) with similar results for larger n . However, there is some disagreement with how reliable these limits are, for instance Ahn [29] reduces the limit to $0.2\text{--}0.4 \text{ TeV}$ for $n = 5$.

Cosmological constraints: Limits from cosmology are the most constraining of all limits. Since the lowest mass gravitons are very light, they can be produced in many cosmological processes. A sample of some of the constraints is given. More details can be found in the reviews [9, 12, 17].

Cooling of the supernova SN1987A by graviton emission into the bulk put limits on M_{PL} [30]. They are: $M_{\text{PL}} \gtrsim 45 \text{ TeV}$, $M_{\text{PL}} \gtrsim 3 \text{ TeV}$, $M_{\text{PL}} \gtrsim 1 \text{ TeV}$ for $n = 2, 3$ and 4 . There are many uncertainties associated with this analysis, but it seems that $n = 2$ is strongly disfavoured whilst $n = 3$ or 4 is still possible.

Cosmological processes that generate many gravitons are also susceptible to limits from decays of gravitons into Standard Model particles. For low mass gravitons the dominant mode is

Theory and Motivations

$G \rightarrow \gamma\gamma$ thus limits arise from carefully measured photon energy distributions. This puts a limit of 1% on the total amount of energy emitted into gravitons by all supernovae in the history of universe. This gives $M_{\text{PL}} \gtrsim 60 \text{ TeV}$ and $M_{\text{PL}} \gtrsim 3 \text{ TeV}$ for $n = 2$ and 3 [31]. A similar study into the photon flux from neutron star haloes gives $M_{\text{PL}} \gtrsim 450 \text{ TeV}$ for $n = 2$ and $M_{\text{PL}} \gtrsim 17 \text{ TeV}$ for $n = 3$ [32].

Limits can also be obtained by requiring that gravitons do not cause the universe to become matter dominated too early which would give a lifetime of the universe that is too low. The results depend on the assumption that QCD effects are negligible below a temperature of $\sim 100 \text{ MeV}$ (a conservative assumption) [33]. These give $M_{\text{PL}} \gtrsim 65 \text{ TeV}$, $M_{\text{PL}} \gtrsim 4 \text{ TeV}$ and $M_{\text{PL}} \gtrsim 0.7 \text{ TeV}$ for $n = 2, 3$ and 4.

Finally, the strongest cosmological limits come from requiring that neutron stars are not excessively heated by the decay products of gravitons trapped in them. Slightly higher temperatures than are expected have been observed, but the heating due to graviton decay can be much larger. The limits are $M_{\text{PL}} \gtrsim 1300 \text{ TeV}$ for $n = 2$ and $M_{\text{PL}} \gtrsim 30 \text{ TeV}$ for $n = 3$ [32].

Future constraints ATLAS will be able to perform studies similar to those which CDF and DØ have done, *i.e.* searching for 1 jet + missing energy. This will allow M_{PL} to be measured if it is between $M_{\text{PL}} = 3.6\text{--}8.1 \text{ TeV}$, $M_{\text{PL}} = 2.8\text{--}4.4 \text{ TeV}$, $M_{\text{PL}} = 2.5\text{--}3.0 \text{ TeV}$ for $n = 2, 3$ and 4 with 100 fb^{-1} of luminosity [34]. The lower limits arise because if M_{PL} is too low, ATLAS will probe Planck scale physics which could give the same signature. If extra dimensions were observed, it would be possible to run at a lower energy to reduce these lower limits.

Of all of these limits, the cosmological ones are the most constraining. Although they are all subject to some assumptions and uncertainty, they are strong enough to rule out the straightforward ADD model for $n = 2$ and 3. However, of those bounds, the ones that rely on decay back to our brane (such as the limits from diffuse cosmic gamma radiation) are somewhat model dependent and can be relaxed (for instance, see [35]). All of the cosmological limits rely on the prediction of ADD that there are a large number of very light graviton modes, say, less than 100 MeV . Models

have been constructed that move the lightest mode up to ~ 10 GeV [36] by changing the geometry of the extra dimensions. In this case, all cosmological constraints are avoided.

1.2.3 Randall-Sundrum models

The Randall-Sundrum (RS) scenario [37] was introduced recently and has generated much interest. Like the ADD model, it uses extra dimensions to try to explain the hierarchy problem, but, rather than adding flat extra dimensions, they add a curved, or warped, extra dimension.

The basic RS model adds one extra dimension with a non-factorisable geometry. The extra dimension is compactified on a half circle, *i.e.* $\phi \rightarrow 2\pi + \phi$ and $\phi \rightarrow -\phi$, where ϕ is the co-ordinate in the extra dimension. The metric is

$$ds^2 = e^{-2kr_c|\phi|} \eta_{\mu\nu} dx^\mu dx^\nu + r_c^2 d\phi^2 \quad (1.13)$$

where k is a scale, order of the Planck scale, which determines the curvature of the extra dimension; r_c is the radius of compactification; x^μ are the usual co-ordinates in four dimension and $\eta_{\mu\nu}$ is the four dimensional Minkowski metric. The RS model is a solution of Einstein's equations with two 3-branes; one at $\phi = 0$ and the other at $\phi = \pi$. Our universe is on the 3-brane at $\phi = \pi$.

It is important to note the exponential factor, sometimes called the warp factor, in front of the normal 4 dimensional metric. This implies that moving in the extra dimension causes exponential changes in the length scale of the other four dimensions. This exponential factor generates large differences in scale for small inputs. On our brane it generates a new scale, Λ_π

$$\Lambda_\pi = M_{\text{PL}} e^{-kr_c\pi}. \quad (1.14)$$

For the RS model to solve the hierarchy problem, Λ_π should be $\sim \text{TeV}$; normally the limit $\Lambda_\pi < 10 \text{ TeV}$ is chosen. This requires $kr_c \approx 12$ and thus the size of the extra dimensions is small (but still larger than $1/k$) compared to the fundamental scale. In these equations, M_{PL} is the reduced 4

Theory and Motivations

dimensional Planck scale

$$M_{\text{PL}}^2 = \frac{M^3}{k} [1 - e^{-2kr_c\pi}] \quad (1.15)$$

and M is the 5 dimensional, fundamental, Planck scale. Note that M_{PL} is not much different from the 5 dimensional scale (since $kr_c > 1$) and as will be mentioned later, is related to the observed strength of (normal) gravity. Thus, in this scenario, $M_{\text{PL}} \sim 10^{-19}$ GeV.

Although the discussion in section 1.2.1 is not valid (because the geometry is non-factorisable), a Kaluza-Klein tower of graviton states is still created [38, 39]. The masses of the graviton resonances are given by

$$m_n = kx_n e^{-kr_c\pi} = x_n (k/M_{\text{PL}}) \Lambda_\pi \quad (1.16)$$

where x_n are the roots of the Bessel function ($x_n = 3.8, 7.0, 10.2$ for $n = 1, 2, 3$). The graviton resonances are therefore also \sim TeV. There are two free parameters in RS which are normally chosen to be either m_n and Λ_π or m_n and k/M_{PL} ; they are trivially related by equation 1.16. On our brane the massive graviton resonances couple with strength $1/\Lambda_\pi$ to the energy-momentum tensor. This universal coupling is typical of gravitational interactions and will motivate the analysis in chapter 5. An interesting feature of this model is that the zero mode graviton, the massless state, couples with strength $1/M_{\text{PL}}$, thus cosmological gravity is still weak as required. This situation is reversed for the other 3-brane — there the zero mode is strong and the excited modes are weak. In performing the derivations of these equations, it is necessary for the bulk curvature not to be too large. This requires that $k/M_{\text{PL}} \ll 1$ and normally $k/M_{\text{PL}} \lesssim 0.1$ is used [40]. RS models can also be motivated from string theory. There the Planck mass can be related to the string scale and this leads to the suggestion of low values of k/M_{PL} , typically ~ 0.01 [40].

The RS scenario is therefore very distinct from the ADD scenario. Rather than the continuum of weakly interacting gravitational modes in ADD, there are massive strongly interacting modes. The width of the graviton resonances is given by

$$\Gamma_n = \rho m_n x_n^2 (k/M_{\text{PL}})^2 \quad (1.17)$$

where ρ is a constant that depends on the number of open decay channels. For the preferred, small values of k/M_{PL} , the width can therefore be very narrow.

The distinctive phenomenology of the standard RS model is therefore of massive ($\sim\text{TeV}$) graviton resonances with narrow widths and which are widely separated. Since the gravitons can be produced directly on resonance, they will decay back to Standard Model particles (assuming that there are no other decay routes). Again, this is quite different from the ADD scenario where, below M_{PL} , graviton emission into the bulk is the most promising probe of the model. This feature would allow the graviton couplings to the Standard Model particles to be measured, which would be an important test of whether the resonance is indeed a graviton and also of the universal coupling that would be expected.

There are many variations on this model. For instance, although the basic RS model restricts the Standard Model field to the brane, models have been created where the gauge fields propagate into the bulk and also the fermion fields. Supersymmetry may be added and the second brane can be moved to infinity and may additionally have positive tension. For a review, see [41]. For the interests of this thesis however, only the basic RS model is considered.

1.2.3.1 Experimental limits

Unlike the ADD scenario, RS models are not constrained by cosmological measurements. This is because the first excited graviton state is $\sim\text{TeV}$ and is thus too heavy to be produced in most cosmological situations. Instead, the most stringent limits come from collider experiments and precision electroweak observables.

Collider limits: Limits on RS models come from the $D\bar{O}$ and CDF collaborations operating at the Tevatron collider. Currently the best limits are from $D\bar{O}$ who recently released preliminary results which are shown in figure 1.4 [42]. CDF has similar but slightly worse preliminary results [43]. Prior to these, the best limits were obtained by adapting the results of the CDF and $D\bar{O}$ searches for Z' resonances [44, 45] (this was the approach taken in [1]).




Figure 1.4: DØ limits on RS models. Adapted from [42].

Electroweak observables: Gravitons enter into precision electroweak variables through higher-order loop diagrams, but these are, in principle, dependent on the full theory. Since the RS model is an effective theory, only valid up to the Planck scale, a cut-off must be introduced into the integral over the graviton momentum. The limits obtained are then dependent on this cut-off however they suggest that low mass gravitons and low values of k/M_{PL} are excluded [40].

Future limits: Studies have shown that ATLAS will be able to explore all of the allowed parameter space and would be able to detect gravitons of mass up to 2 TeV even for the worst case with $k/M_{\text{PL}} = 0.01$. The spin-2 nature of the graviton could be determined up to a mass of 1.7 TeV [46]. This study is closely related to the material in chapter 5 and is discussed further there.

These limits, combined with the theoretical requirements that $\Lambda_{\pi} \lesssim 10 \text{ TeV}$ and $0.01 \leq k/M_{\text{PL}} \leq 0.1$, give an allowed space for the basic RS model that is closed. This is illustrated in figure 1.5.




Figure 1.5: The allowed space (white) for the RS model. The limits come from $\Lambda_\pi \lesssim 10$ TeV (dark blue), $0.01 \leq k/M_{\text{PL}} \leq 0.1$ (light blue), electroweak observables (green) and Tevatron limits (brown).

1.3 Summary

This chapter has introduced two classes of extra-dimensional models. The first, ADD involves comparatively large, flat dimensions. One exciting possibility introduced by this model (by virtue of the lowered Planck mass) is that microscopic black holes may be created at the LHC. These black holes are the topic of chapter 6. The other class of model is that of L. Randall and R. Sundrum (RS). This model, with its warped extra dimension, predicts that massive graviton resonances should be found in the TeV energy range. The couplings of these massive resonances is the subject of chapter 5.

However, studies of such models would be a purely intellectual exercise if it were not for building and running of experiments. The experiments that are being constructed at the LHC are the largest, most complex, and most expensive that have been built. My contribution to one of them, ATLAS, is the subject of the next three chapters.

ATLAS

We're an anarcho-syndicalist commune. We take it in turns to act as a sort of executive officer for the week...

2.1 The Large Hadron Collider

The Large Hadron Collider (LHC) is the accelerator, currently under construction at CERN in Geneva, that will provide the high energy protons for collision in the ATLAS detector. It will be able to accelerate two beams of protons in opposite directions around a 27 km circumference ring to an energy of 7 TeV. This is over seven times greater than the current highest energy accelerator in the world, the Tevatron at Fermilab. The two beams of protons are caused to collide at four points around the ring where the main experiments are situated (figure 2.1). The total centre-of-mass energy is 14 TeV.

The LHC will not only provide very high energy collisions, it will also provide a very high total luminosity. The design luminosity is $1 \times 10^{34} \text{ cm}^{-2}\text{s}^{-1}$ with an initial period of low luminosity running at $0.12 \times 10^{34} \text{ cm}^{-2}\text{s}^{-1}$. This is achieved by having very good focusing (the beam transverse size will be $16.6 \mu\text{m}$) and a very high interaction frequency of 40 MHz. Achieving such a high luminosity will be a challenge — it is worth noting that the Tevatron currently achieves peak luminosities of about $10^{32} \text{ cm}^{-2}\text{s}^{-1}$ (although this is limited by the Tevatron's use of anti-protons).

These machine parameters mean that the experiments will have large numbers of very high energy events that can be analysed to find new physics. It does also provide a number of challenges: notably the high interaction rate puts constraints on event selection (see section 2.2.5) and offline data storage and analysis. The high luminosity also comes with a price: there are ~ 23 interactions




Figure 2.1: Schematic of the LHC layout.

per bunch crossing which gives rise to pile-up — low energy events which can mask the high energy events that are of interest. Nevertheless, it is up to the experiments to show that they can deal with these problems and make the most of the opportunities that the LHC provides.

2.2 The ATLAS detector

ATLAS, which ‘stands for’: A Toroidal LHC ApparatuS, is one of two general purpose experiments currently being built at the LHC. ATLAS is the largest particle physics detector ever assembled, with a length of 44 m and a diameter of 22 m^a; a schematic overview is given in figure 2.2.

ATLAS is designed for a wide-range of physics signals, but the most important are:

Higgs boson: Finding the Higgs boson and measuring its properties is one of ATLAS’ principle aims. ATLAS can discover the Higgs boson with 30 fb⁻¹ of data over the complete expected mass range, from 100–1000 GeV. The discovery potential is shown in figure 2.3.

New physics: ATLAS should be sensitive to as much new physics as possible in the TeV energy

^aIt isn’t the heaviest though, CMS — the other general purpose experiment — has that distinction, with a weight of 12,000 tons.

2.2 The ATLAS detector




Figure 2.2: A schematic diagram of the ATLAS detector. From [47].

range. This includes supersymmetry, extra dimensions and any other possible extensions to the Standard Model.

Electroweak physics: Measurements of the precision electroweak observables, such as the W mass, will be important for calibrating and understanding the detector. ATLAS also expects to make competitive measurements of the top mass and to investigate gauge boson couplings.

Strong interactions: ATLAS will investigate many QCD parameters, but perhaps most importantly, it should make new measurements of the parton distribution functions (PDFs) which are essential for many other measurements. QCD processes also form important backgrounds to many new physics signals and will therefore need to be well understood.

In addition to the above list, there are also active investigations into B-physics, heavy-ion physics

and forward physics.




Figure 2.3: Higgs discovery potential with a total luminosity of 30 fb^{-1} . From [48].

In common with many detectors at hadron-hadron machines, ATLAS has a cylindrical geometry. The usual co-ordinates are therefore (R, ϕ, z) where R is the transverse radius from the beamline, ϕ is the azimuthal angle^b and z is the distance parallel to the beamline with 0 at the nominal interaction point.

The true rapidity, $y = 1/2 \ln[(E + p_z)/(E - p_z)]$, of a Lorentz vector is useful quantity at hadron-hadron colliders, where the initial z -momentum is unknown, since differences in rapidity are invariant under longitudinal Lorentz boosts. However, it is difficult to measure since it requires

^b ϕ is defined as the angle from the x -axis which points from the interaction point to the centre of the LHC ring.

that the particle mass be known. Instead, the pseudo-rapidity, $\eta = -\ln[\tan(\theta/2)]$, which is a good approximation to the true rapidity in the relativistic limit, is used. Particles are often described with the parameters (p_T, η, ϕ) and gaps between particles are given in the η - ϕ plane with a separation $\Delta R = \sqrt{\Delta\eta^2 + \Delta\phi^2}$. The central, cylindrical part of ATLAS is called the ‘barrel’ whilst the ends of ATLAS are referred to as ‘end-caps’ and the parts of the detector at very high $|\eta|$ are in the ‘forward’ region.

To achieve the above physics goals, ATLAS is designed to have excellent electron, photon, muon and jet energy resolution over the range 10–1000 GeV with high granularity in η - ϕ . It aims to provide excellent measurement of the missing momentum, \cancel{p}_T , caused by neutrinos or other non-interacting particles (such as the lightest supersymmetric particle). It has a central inner detector that provides tracking of charged particles and accurate momentum measurements from 1–100 GeV. Outside that are the calorimeters which measure electromagnetic and hadronic energy and direction. They have coverage to high $|\eta|$ to give good \cancel{p}_T measurements and can be useful for ‘tagging’ interesting physics events. The outermost part of the detector is the muon spectrometer which provides muon triggering and momentum measurements up to 6 TeV.

It is only possible to provide a short summary of the main components of ATLAS in this thesis. For more information, the reader is referred to the various ATLAS technical design reports (TDRs) [47, 48, 49] and also to the Final Design Reports which detail the changes since the TDRs.

2.2.1 Magnet systems

The magnet systems are an extremely important part of ATLAS — they even give it its name! There are two magnets: the central solenoid which surrounds the inner detector and the toroid system which provides a magnetic field for the muon spectrometer, see figure 2.4. Both magnet systems use NbTi superconductors which are cooled to 4.5 K. This choice of magnet system gives a strong central field in the inner detector with a minimum of material. The toroidal system covers the large volume of the muon spectrometer with a strong field and a light and open structure.

The central solenoid is a 2 Tesla magnet, 2.5 m in diameter, 5.3 m long and only 45 mm thick.

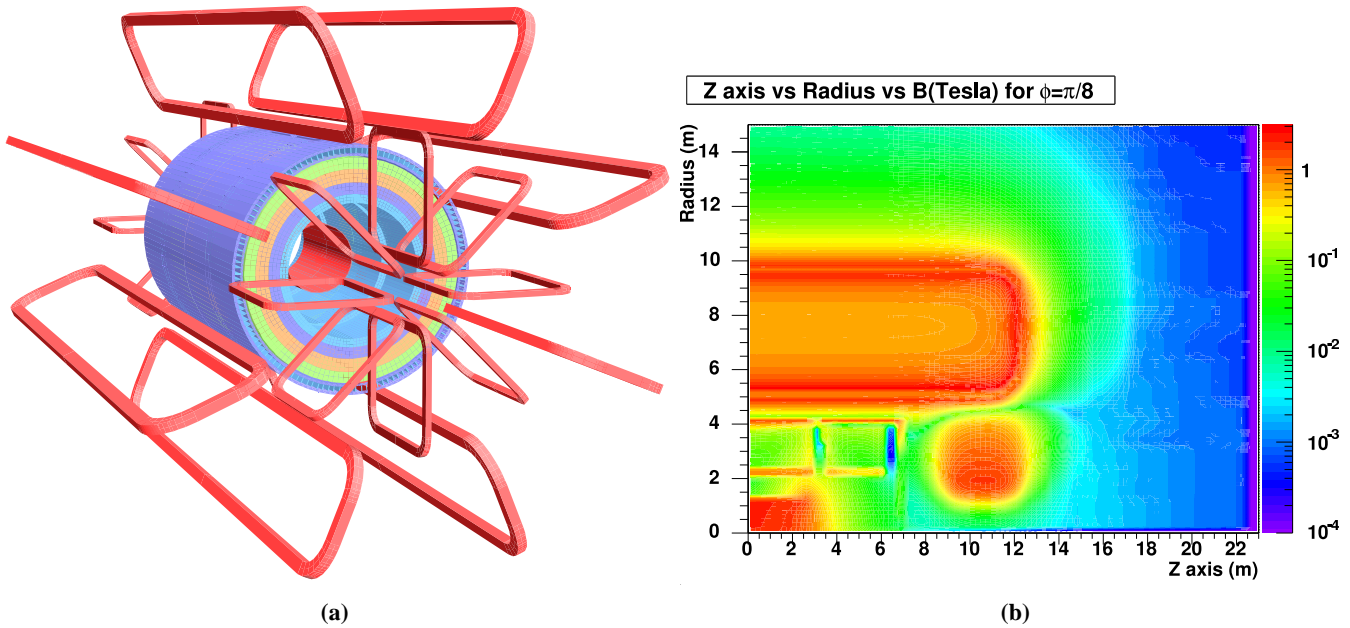


Figure 2.4: The magnet system in ATLAS. **(a)** shows the the magnet system, the hadronic calorimeter acts as the return yoke for the central solenoid. **(b)** shows a part of the magnetic field map through a barrel toroid. Adapted from [50], updated from [47].

The hadronic calorimeter acts as the return yoke. Its thin design is to minimise the amount of material introduced in front of the calorimeters, but still allows it to provide enough of a magnetic field for the inner detector to accurately measure the momentum of charged particles up to 100 GeV.

The toroid system provides a magnetic field in the azimuthal direction which is thus perpendicular to most hard muon tracks. It is generated by 8 ‘race track’ superconducting magnets in the barrel region and 8 similar, smaller magnets in each of the two end-caps. The barrel toroids are immense at 25 m long and 5 m wide. Cooling the 1300 tons of magnet system takes nearly 40 days and when on, the whole system stores 1600 MJ of energy. The average field in the muon spectrometer is 0.6 T which, together with its large size, allows muon momenta of up to 6 TeV to be measured.

2.2.2 Inner detector

The inner detector aims to measure the tracks of charged particles. These measurements are very important, allowing particle momentum (below 100 GeV), interaction vertexes (important for b-tagging and tau reconstruction) and charge to be measured whilst helping with particle identification. The inner detector must achieve this using the minimum amount of material. At high luminosity, there will be many charged particles crossing the inner detector, so it aims to make many precise space-point measurements over the precision region ($|\eta| < 2.5$) with high granularity. The required spacial accuracy varies considerably with radius — very high precision measurements are needed near the beam pipe, whilst further out, the amount of material, and cost, become more significant. The inner detector is therefore split into three separate parts: the pixel detector and the SCT which provide high resolution measurements and the TRT which provides continuous track-following. An overview is given in figure 2.5. Significant restrictions in the design come from the high radiation that the detectors must be able to withstand and from the readout requirements which necessitate the readout electronics being placed within the tracking volume thus increasing the total amount of material. The three components have been matched so that none dominates the momentum resolution.




Figure 2.5: The inner detector as described in the inner detector TDR. From [51].

2.2.2.1 Pixel detector

The pixel detector is closest to the beam pipe. Its active material is depleted silicon which is in the form of 16.4×60.8 mm detectors, each subdivided into 46,080 pixels of size 50×400 μm . The more accurate measurement, in $R\phi$, is to provide a better measurement of the sagitta of the particle tracks giving a better momentum measurement. The pixel detectors are arranged into 3 barrels at radii: 50.5 mm, 88.5 mm and 122.5 mm and 5 discs in each end-cap. The innermost barrel, referred to as the B-layer, is removable and is physically part of the beam pipe. This layout is somewhat different from the one presented in the TDRs, an update can be found in [52] and the references therein. The readout electronics are mounted directly on the detector and compare the electrical signal to a threshold to give a binary output. Due to cost over-runs the middle pixel layer will not be installed in the initial ATLAS machine, it will be added at a later date.

2.2.2.2 Semiconductor tracker

At larger R , the greater area (63 m^2 compared to 2 m^2 for the pixel detector) and readout requirements lead to the use of silicon strip detectors. The SCT uses the same detector technology as the pixel detector, but the detectors are 63.6×64 mm wafers with $80 \mu\text{m}$ pitch strips. Four of these wafers are arranged into modules with two pairs of wafers being connected end-to-end to give an effective length of 123.2 mm. The two pairs of wafers are placed back-to-back with a small stereo angle, 40 mrad, between them to give some measurement along the length of the strip.

As in the pixel detector, the modules are arranged to give greatest accuracy in $R\phi$. There are 4 barrels covering $|\eta| < 1.1$ and 9 discs in each end-cap providing coverage out to $|\eta| = 2.5$. As with the pixel detector, the readout electronics compares the charge collected for each strip with a threshold giving a binary output. The SCT is the subject of chapter 4 and a more complete introduction, covering the modules, readout system and software is given in chapter 3.

2.2.2.3 Transition radiation tracker

The transition radiation tracker (TRT) is a drift tube detector made up of hundreds of thousands of ‘straws’. Each of these is 4 mm in diameter with a central gold-plated tungsten wire which is the anode. The inside of the straw is coated with aluminium and acts as a high-voltage cathode; in between is gas mixture of Xe/CO₂/O₂. Each wire is split in two at the centre and read out at each end to reduce the occupancy. Each of these channels gives a drift-time measurement which gives a spatial accuracy of 170 μm per straw and two independent thresholds.

The lower threshold detects the charge from a minimally ionising particle. On average 36 such measurements are expected for each particle with $p_T > 0.5$ GeV. The higher threshold detects the transition radiation photons that are emitted whenever relativistic particles cross a boundary between materials with different dielectric constants. To increase the amount of such radiation, the gaps between the straws are filled with a polypropylene/polyethylene fibre radiator. This transition radiation helps identify particles, in particular, the pion rejection is expected to be ~ 100 for an electron efficiency of 90%.

The TRT is arranged in a barrel that covers $|\eta| < 0.7$ and 18 wheels in each end-cap that extend to $|\eta| = 2.5$. The current staged scenario envisages only 14 wheels in each end-cap with the rest to be added later. As the staged wheels have the highest $|\eta|$ coverage, this decision means that initially the TRT will only extend to $|\eta| < 2.0$.

2.2.3 Calorimeters

The purpose of the calorimeters is to provide accurate measurements of the energy of electrons, photons and jets. In addition, by measuring the transverse energy, it provides an important contribution to the \cancel{p}_T measurement (the other contributor is the muon spectrometer).

Electrons and photons interact quite differently with matter compared to hadrons. They penetrate much less deeply, with a length characterised by the radiation length, X_0 , and have narrow profiles. Hadronic showers are characterised by the nuclear interaction length which is typically an order of magnitude greater than X_0 . ATLAS therefore, like many general purpose detectors,

has two calorimeter systems: the electromagnetic calorimeter (ECAL) and the hadronic calorimeter (HCAL). The layout of these is shown in figure 2.6.




Figure 2.6: ATLAS calorimeters. From [53] (adapted from [54]).

2.2.3.1 Electromagnetic calorimeter

The ECAL is a lead-liquid argon (LAr) sampling calorimeter with accordion geometry covering the region $|\eta| < 3.2$. It has a slightly complex geometry which is summarised in figure 2.7 and provides complete ϕ symmetry without azimuthal cracks. There is, however, a crack at $|\eta| = 1.5$ where the barrel to end-cap transition occurs and a small crack at $\eta = 0$. In order to provide good particle identification (via shower shape) and shower direction, the ECAL has three longitudinal sections, samplings, with high granularity which allows for excellent π^0/γ and e/π separation.

The ECAL has a thickness of at least $24 X_0$ in the barrel and $26 X_0$ in the end-caps which should keep leakage from high energy showers ($E > 500 \text{ GeV}$) to an acceptable level. As




Figure 2.7: Geometry of the ECAL. From [54].

a sampling calorimeter, its performance is dominated by sampling fluctuations giving $\sigma/E = 8.6\% / \sqrt{E(\text{GeV})} \oplus 0.5\%$ for electrons at $\eta = 0.3$ [54].

2.2.3.2 Hadronic calorimeter

The hadronic calorimeters cover the range $|\eta| < 4.9$ using a variety of different technologies. For $|\eta| < 1.7$, barrel and extended barrel tile calorimeters are used while in the end-cap ($1.5 < |\eta| < 3.2$) LAr calorimeters are used, identical to those in the ECAL. In the very forward region a special high-density LAr calorimeter is used: the forward calorimeter (FCAL). The thickness of the HCAL is about 10 interaction lengths which should allow good containment of even the highest energy jets and also keep punch-through to the muon system to a minimum. The reach to very high $|\eta|$ is necessary to allow the use of tag jets and to provide good p_T measurements.

The tile calorimeter is a sampling calorimeter made from iron as the absorber and scintillating tiles as the active material. The calorimeter has three longitudinal layers and the granularity is 0.1×0.1 in $\Delta\eta \times \Delta\phi$ (0.2×0.1 in the last layer). The FCAL consists of three sections, all using

LAr as the active material. One of the sections uses copper whilst the other two use tungsten to provide even higher densities. In the end-cap and forward calorimeters beyond $|\eta| = 2.5$, the granularity is 0.2×0.2 .

It is important to note that both the ECAL and HCAL are non-compensating, that is, they respond differently to EM and hadronic showers. This difference must be calibrated and accounted for in the reconstruction software. The expected jet energy resolution is $\sigma/E \approx 50\% / \sqrt{E(\text{GeV})} \oplus 1.8\%$ whilst the \cancel{p}_T performance should be $\sigma = 0.39 \times \sqrt{\sum p_T(\text{GeV})}$ where $\sum p_T$ is the total transverse energy in the calorimeters [47].

2.2.4 Muon spectrometer

The muon spectrometer is designed to accurately measure the momentum of high energy muons (low energy muons will be measured by the inner detector) and to trigger on them. These considerations, the large size of the spectrometer and cost necessitate the use of several different detection technologies. Figure 2.8 gives an overview of the layout of the muon system. The two main detectors are the monitored drift tubes (MDTs) which are used for precision measurements and the resistive plate chambers (RPCs) which are used for triggering. At high $|\eta|$, cathode strip chambers (CSCs) and thin gap chambers (TGCs) are used for these tasks, further details can be found in the muon TDR [55].

The MDTs are 30 mm diameter drift tubes arranged into units called chambers. These incorporate a novel optical monitoring system (providing the ‘M’ in the name) to measure any physical deformations. The single-wire resolution is expected to be $80 \mu\text{m}$. The MDTs are arranged to be at approximately constant η .

The drift time in the MDTs is longer than the 25 ns event time, so they cannot be used for triggering. Instead RPCs are used which provide a space–time resolution of $1 \text{ cm} \times 1 \text{ ns}$. An RPC is a pair of parallel plates separated by a narrow gas gap with a high electric field. Primary ionisation electrons from the passage of muons cause an avalanche between the plates. The signal is read out by two sets of strips in orthogonal directions. These measurements are used in the




Figure 2.8: End view of the muon spectrometer. From [55].

level 1 trigger and also to provide a measurement in the orthogonal direction to the MDTs.

The muon spectrometer should provide three precision measurements over the range $|\eta| < 2.7$ (triggering for $|\eta| < 2.4$) except where there are cracks (most notably at $\eta = 0$ where there is a large crack for cabling). This allows the muon momentum to be measured to 2% at 20 GeV with good acceptance. At 1 TeV the momentum can be measured to 10% although the acceptance is rather low (50%), but rises to nearly 90% for a 20% momentum measurement.

2.2.5 Trigger and data acquisition

A full ATLAS event is ~ 1 Mb compressed. Given the event rate of 40 MHz, it is clear that storing all the data ATLAS will produce is grossly infeasible. The trigger and data acquisition (DAQ) system is designed to reduce the rate of events that will be stored to ~ 100 Hz. Even at this rate, ATLAS will generate several petabytes of data per year. Analysing all of this data will take new techniques and systems which are being addressed through the development of the Grid [56, 57, 58, 59, 60].

Of course, the trigger and event filter system must allow the ‘interesting’ events to pass, whilst filtering out the less interesting background and low energy events. It would be disastrous if new physics was not discovered because the events were not recorded. To achieve the reduction in rate of 10^6 , requires several different levels of filtering, summarised in figure 2.9.

The first level trigger takes place in custom hardware and uses information from single detectors only. Only the calorimeters and the fast muon trigger chambers can generate a level 1 accept: tracking information cannot be used as track reconstruction would take longer than the $2.5 \mu\text{s}$ allowed. At level 2, the whole detector is read out but processing of the event is limited to regions of interest (RoI) seeded by the level 1 accept. These RoI may include inner detector tracking information. Finally, the third level, called the event filter, occurs in a processor farm and the whole event is used to make the accept decision. For the first time, full calibration and geometry information is available to improve the reconstruction algorithms.

Part of the trigger system will be staged which has impacts on the range of physics processes that can be accepted. However, this should not significantly affect high- p_T physics which is the subject of this thesis.

2.3 Simulation software

Simulation software is an important part of ATLAS — it is essential for understanding the detector during the research and design phases as well as allowing physics studies to have reasonably

realistic estimates of the performance. It will also be important once ATLAS starts taking data as simulation will be used to help understand acceptance and selection efficiencies. Currently ~100 people are working on software putting it at about the same level as one the detectors.

The current version of the software is written in C++ using object-oriented techniques. The software forms a framework, Athena [62], within which the various stages of simulation and

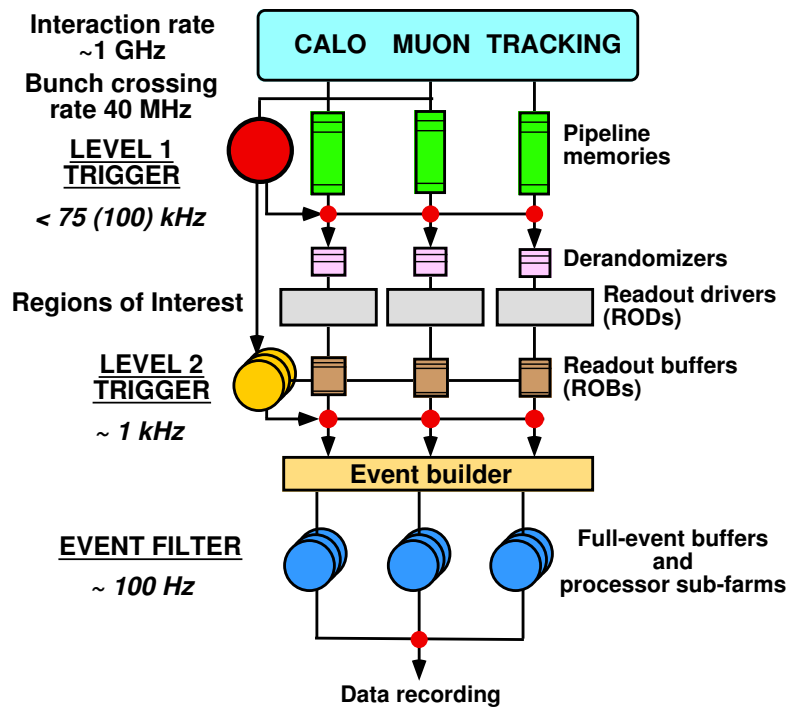


Figure 2.9: Trigger overview. From [61].

reconstruction are run. This is based on the Gaudi framework [63] used by both ATLAS and LHCb. This software supports both full simulation, where simulated particles are tracked through a realistic description of the detector and as many interaction effects as possible are modelled; and fast simulation where the detector's response to particles is parameterised.

2.3.1 Full simulation

The full simulation requires several steps and is very time consuming. First, the event is generated, normally by an external generator such as HERWIG [64] or PYTHIA [65]. This is then passed to the GEANT3 [66] program which is used to model the ATLAS detector. Recently the transition has been made to GEANT4 [67], however, the studies in chapter 6 used GEANT3.

The next step is to turn the GEANT information into the output of the detector, this is known as 'digitisation'. During this step it is possible to add electronic and other noise. The last step is reconstruction which takes the detector output and attempts to determine the physics objects such as tracks, electron, jets etc. This then is the starting point of physics analyses.

2.3.2 Fast simulation

The fast simulation in ATLAS is called ATLFAST. It is available as part of the Athena framework although an older, FORTRAN, version was used for the studies in chapter 5. ATLFAST works by taking the output of the generator and then selecting muons; these are smeared with the expected energy and position distributions and output. All other (visible) particles are deposited on a 0.1×0.1 η - ϕ grid. A clustering algorithm is then run and electron and photons are identified with clusters and smeared with the appropriate resolutions. The remaining clusters are marked as jets and also smeared. The total energy in the map is used with the smeared muon momenta to determine \cancel{p}_T . This rather simplistic treatment is considerably faster than full simulation, achieving many events per second rather than requiring many minutes per event. The resolutions in ATLFAST were obtained at the time of the Detector and Physics Performance TDR and were mainly based on full simulation responses to Standard Model and Higgs events.

2.4 Current status

Both the LHC and ATLAS are now in their construction phases and several important milestones have been passed for both. Currently component delivery for the LHC is on target for a turn-on in summer 2007 although there are concerns regarding the cryogenic distribution lines. A recent picture from the ATLAS cavern is shown in figure 2.10. As can be seen, installation is well underway with the two barrel toroids in place. The central barrel of the tile calorimeter with the LAr calorimeter at its centre can be seen at the far end. ATLAS is aiming to finish installation by the end of 2006 to leave 2007 free for commissioning and calibration. There are a number of potential concerns for achieving this but these issues are being addressed and ATLAS will certainly be ready for the LHC turn-on. As already mentioned, there are a number of parts of the ATLAS detector that will be staged due to budget overruns. These parts of the detector will be added later when funds are available, but certainly before high luminosity operation starts. During the first run it is hoped that 10 fb^{-1} of luminosity will be delivered which should be enough for the Higgs boson to be found if the ATLAS and CMS data is combined.




Figure 2.10: A recent picture of the ATLAS cavern. Two barrel toroids are visible in the foreground and the barrel LAr calorimeter with part of the tile calorimeter surrounding it is visible in the background.

The SCT and its Readout Systems

Strange women lying in ponds distributing swords is no basis for a system of government. Supreme executive power derives from a mandate from the masses, not from some farcical aquatic ceremony.

The SCT was briefly introduced in section 2.2.2.2. In this chapter, a more complete overview of the SCT, including the module design and readout, is given so that the rôle and operation of the SctRodDaq software can be understood. This will still be a brief introduction however; for more detailed information see the SCT TDR [68] and the references therein.

3.1 System design

Figure 3.1 shows the layout of the SCT. It is divided into 4 barrels and 18 discs, 9 on each side. The barrels are placed at radii from 300 mm to 520 mm and the largest discs cover radii 259–560 mm.




Figure 3.1: The layout of the SCT. All 18 discs are clearly visible, but only the outermost (barrel 6) of the four barrels can be seen.

The individual detector units are known as ‘modules’ and there are 4088 in the whole SCT with 2112 in the barrel and 1976 in the end-caps. The innermost barrel, barrel 3^a, has 384 modules

^aThe barrels are numbered from 3 to 6. The rationale for this is based on the pixel detector with the three pixel barrels called the B-layer, barrel 1 and barrel 2.

whilst the outermost, barrel 6, has 672. The discs have between 52 and 132 modules each.

3.2 Detector modules

The individual detector modules^b are made from four silicon wafers arranged into two back-to-back pairs. Each wafer has 768 strips, 80 μm pitch and 61.6 mm long. The effective strip length is doubled by daisy-chaining each pair of wafers with wire-bonds together giving a length of 123.2 mm. One pair of wafers is aligned with the beam direction whilst the other pair is offset at a 40 mrad angle^c to provide some measurement in the z -direction. A technical description of the detector wafers can be found in [70].

Readout of the detectors is handled by the ABCD3TA application specific integrated circuits (ASICs) [71] which are manufactured using the radiation hard DMILL process [72, 73]. Each ABCD3TA chip can read out 128 channels thus six are needed to read out each side of the detector. They are mounted on a wrap-around electronics hybrid. This thermally decouples the ASICs from the wafers which will be cooled to -7°C to reduce radiation damage and the amount of heat generated. Figure 3.2 shows a picture of a barrel module in which the main features can be seen. End-cap modules are geometrically different but are electrically similar (they have more redundant links).

A block diagram of the ABCD3TA ASIC is shown in figure 3.3. The most important thing to note is that the readout is binary: the charge collected on each strip in each event is compared to a threshold value by a discriminator circuit. If the charge is greater than the threshold, then a hit is recorded and stored in the first-in-first-out (FIFO) buffer. If a trigger level 1 accept command is received for that event, then the hit will be read out. Each side of the detector has a separate readout link but they share a common command link.

Each ABCD3TA has a number of configurable parameters. The most important is the discriminator threshold value. In addition, it was found that there were significant channel-to-channel

^bThe details presented here correspond to barrel modules; some of these details are modified for some of the end-cap modules.

^cCompare this to the much larger angles in [69].




Figure 3.2: A photograph of a barrel module. The main features can be seen including the 40 mrad stereo angle and the hybrid with six ABCD3TA chips on it. Just to the left of the hybrid, the small gap between the two wafers can be seen. At the bottom of the picture is the connector via which power, commands and data readout are provided.

variations in the threshold. This has been corrected for by the addition of trim digital to analogue converters (DACs). For each channel a 4 bit number is stored that multiplies a chip-wide step value that is determined by the trim range setting — four trim ranges are available. For each channel this trim DAC output is added to the threshold DAC output. Thus channel-to-channel variations can be corrected for.

To help calibrate the detector, each channel has a 100 fF calibration capacitor. The output amplitude is controlled by an 8 bit DAC allowing a range of charges from 0 to 16 fC to be injected. The calibration capacitors are themselves calibrated during ASIC production; these correction factors are stored in the SCT Production Database. A 5 bit strobe delay register is also provided that offsets the timing of the discriminator relative to the clock signal over a 50 ns window. Finally,

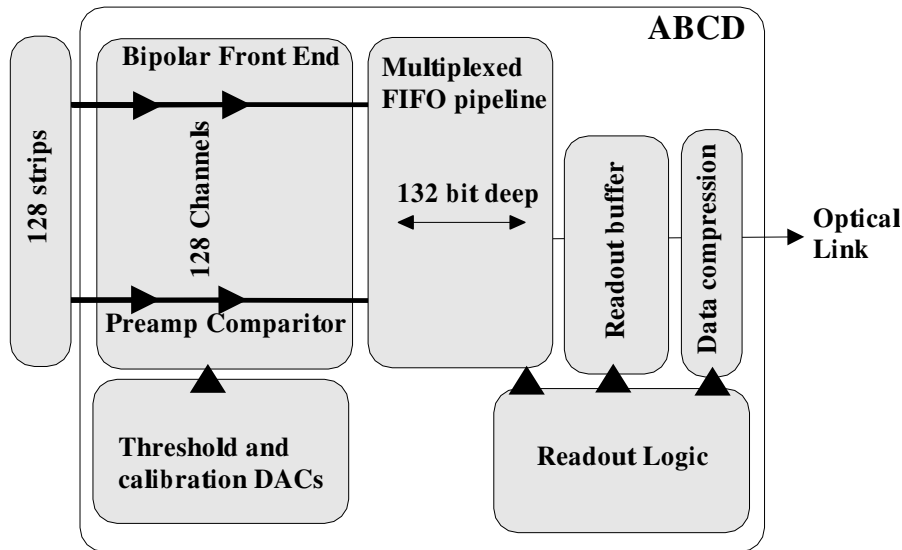


Figure 3.3: Simplified block diagram of the ABCD3TA chip. From [53], adapted from [71].

there is a mask register that can be used to prevent the readout of malfunctioning channels and two 5 bit DACs for adjusting the pre-amp bias current and the shaper current.

3.3 Readout systems

Two readout systems are available for SCT modules. The first, the electrical readout system [74, 75, 76], was designed for use during module and detector research and development. Calibration and control using this system is provided by the SCTDAQ software [77]. The optical readout system is designed for full ATLAS. By using optical signals, it avoids the need for complicated earthing and other electrical arrangements over the >100 m that the signals must travel. It also supports much higher data rates, has less material and better radiation tolerance.

3.3.1 Optical readout system

An overview of the SCT off-detector electronics and its connections to the ATLAS trigger and DAQ system is shown in figure 3.4. A brief description of the main components is given followed

by an explanation of the system operation during data taking.

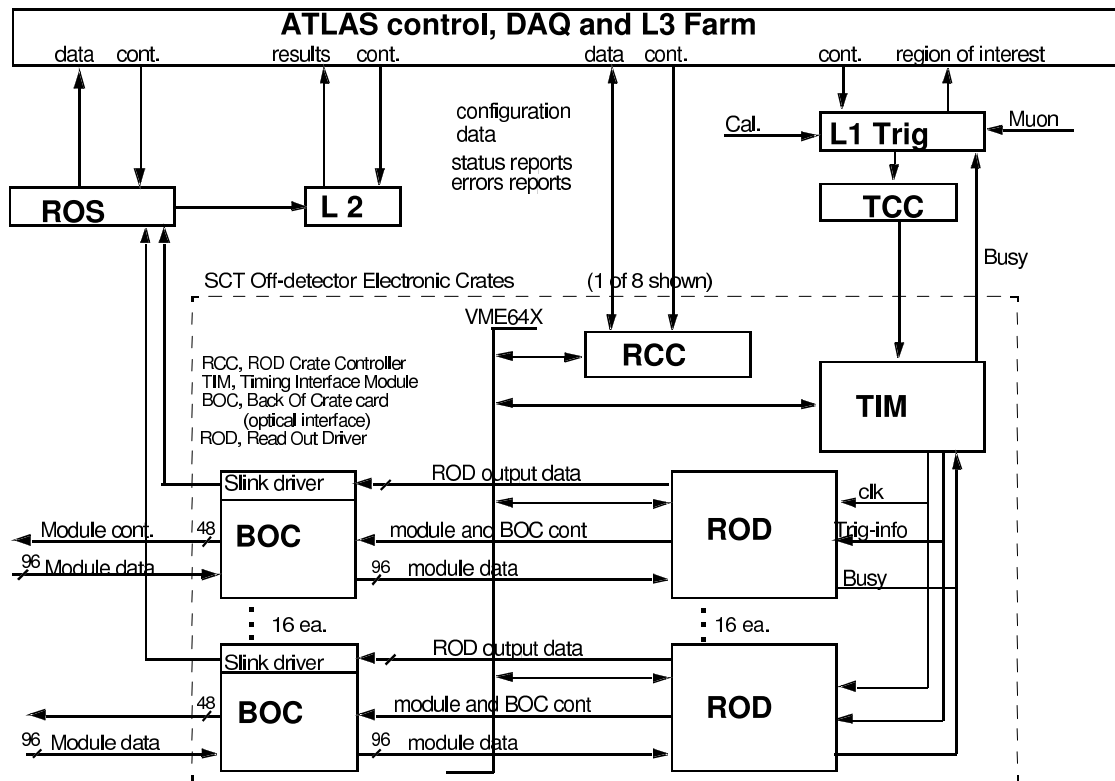


Figure 3.4: Block diagram of the optical readout system showing the connection to the ATLAS DAQ system. The ATLAS specific components are shown in the dashed box which denotes the contents of one electronics crate (of eight total). The modules are connected to the BOCs in the bottom left-hand corner. From [78].

3.3.1.1 BOC

The back of crate card (BOC) handles all the optical interfaces for 48 modules. Each BOC is partnered with a ROD which it communicates with via the crate's custom backplane. The BOC is also the connection to the ATLAS level 2 trigger and DAQ. This connection is provided by the S-Link [79, 80] which is the ATLAS custom high-speed optical data link.

3.3.1.2 ROD

The readout drivers (RODs) handle most of the computational tasks involved in readout and calibration. Each ROD has 5 digital signal processors (DSPs) and over 1 Gb of RAM that can be used for these tasks. It also has a number of field programmable gate arrays (FPGAs) which are used for data decoding and event building. Each ROD receives data from the modules via its BOC. It decodes event data then builds it into an ATLAS standard event fragment. This is then returned to the BOC for transmission up the S-Link. The RODs also receive triggers and pass these onto the modules, again, via the BOC. The DSPs are particularly useful during calibration as they can be used to scan over module parameters and histogram the results. During ATLAS running, they will be useful as data monitoring tools.

3.3.1.3 TIM

The TTC information module (TIM) receives triggers from the ATLAS trigger, timing and control (TTC) system and distributes them to the RODs via the backplane. It can also generate internal synchronised triggers for use during calibration and testing. The TIM also receives the clock from the TTC system thus ensuring that all components are synchronised or can generate an internal clock if no TTC system is available.

3.3.1.4 RCC

The ROD crate controller (RCC), which is a single board computer (SBC), is a standard PC with a VME interface that is mounted in the crate. It is used for initialising and configuring the other off-detector electronics. During calibration, it is used for controlling the operation of the electronics and modules. During data taking it can be used for data monitoring where a fraction of events are delivered to it from the RODs via the VME bus. It is externally connected via a 100 Mbit ethernet connection.

3.3.1.5 Crate

Each crate contains 1 RCC, 1 TIM and up to 16 RODs and BOCs providing readout and control for up to 768 modules. The components are all connected via the backplane using custom connections with a VME bus providing control to the RCC. In total there will be eight crates.

3.3.1.6 TTC

The trigger, timing and control (TTC) system is the ATLAS system that provides the clock and level 1 triggers to the sub-detectors. The TCC is the interface between the TIM and the TTC system.

3.3.1.7 ROS

The readout system (ROS) is part of the ATLAS trigger and DAQ system. It receives event fragments via the S-Links and passes the data to the level 2 trigger processors based on the regions of interest (RoI). If the event passes the level 2 trigger, the fragments are built into a whole event and passed onto the 3rd level event filter.

3.3.1.8 System operation

When the system is switched on, the RCC is used to initialise and configure the other electronic hardware. It also loads the module configurations from external storage and passes these to the RODs. When data taking has started, level 1 triggers are received in each crate by the TIM. This fans the signal out to the RODs which issue the trigger to the modules via the BOCs. The modules return their data to the BOCs which pass them back to the RODs. The RODs decode the hit data and reformat it into event fragments. These are sent back to the BOCs for transmission over the S-Link to the ROS and the level 2 trigger/DAQ system. A fraction of these events can be passed to the DSPs or RCC where they can be analysed to monitor the performance and report statistics.

3.4 Module characterisation and calibration

Before discussing characterisation and calibration, it will be useful to understand a little about how a silicon detector actually works. The silicon is depleted of charge carriers by reverse biasing it. As a charged particle passes through it releases a small amount of energy according to a Landau distribution [81]. This creates electron-hole pairs which are separated by the same electric field that biases the detector and are then collected on the strips; this is the input to the ABCD3TA readout chips.

Imperfect charge collection and other sources of noise (such as electronic noise in the preamp circuits) are roughly Gaussian so the charge distribution seen by the threshold discriminator circuit is an ‘improved Landau’ distribution: a Landau convolved with a Gaussian. These noise sources can also cause a hit to be registered even if there is none. This is known as *noise occupancy* and is a key measure of the detector’s performance.

The occupancy, ρ , for a given threshold can be easily calculated

$$\rho = \int_t^{\infty} P(Q) dQ \quad (3.1)$$

where $P(Q)$ is the charge probability distribution and t is the threshold. This is shown in figure 3.5a for a minimally-ionising particle (MIP) and in figure 3.5b for the calibration charge which, since it provides a fixed charge, will have a Gaussian charge distribution determined by the noise. Typical values for a MIP are a most probable value of 3.6 fC with a width parameter of 0.17 fC and noise around 0.2 fC. The noise occupancy charge distribution has a mean just above 0 fC with a width of 0.15 fC.

The specification [68] for the SCT is that the detectors should have an efficiency greater than 99% with a noise occupancy less than 5×10^{-4} . This leads to a nominal operating threshold of about 1 fC. Setting such a threshold that is constant across all the channels is one of the most important tasks of the calibration software.

The basic unit of much of the characterisation and calibration is the ‘threshold scan’. This




Figure 3.5: A schematic showing the charge distributions (thin, solid lines) for (a) a MIP and (b) a calibration pulse. The thick solid lines show the expected average occupancies as a function of threshold for these charge distributions (calculated from equation 3.1). The noise occupancy charge distribution (dashed line) is also shown.

involves using the calibration capacitor to inject a known charge and sending a large number of level 1 triggers for each value of the threshold DAC. The average occupancy then has a shape like the bold line in figure 3.5b. By fitting the occupancy data with a complementary error function, the mean threshold (*i.e.* the 50% occupancy point), t_V , and the width (often referred to as output noise), n_V , can be found. At the mean threshold, the charge is equal to the injected charge, *i.e.* $t = t_V$ at $Q = Q_{\text{cal}}$.

The threshold is not normally set as a charge, but as the voltage output by the threshold DAC (one DAC bit corresponds to a nominal voltage of 2.5 mV). The relationship between the voltage at the discriminator and the charge seen by the ABCD3TA input is initially unknown. This relationship is known as the response curve and is typically slightly non-linear. Each threshold scan gives one point on the response curve — the mean threshold corresponds to the injected charge. Therefore, by performing a series of threshold scans with different injected charges, the response curve can be determined. Once the response curve has been found, the threshold DAC can be set to correspond to a given charge (such as the nominal operating charge of 1 fC).

The SCT and its Readout Systems

The calibration must also determine the trim DAC settings. This is a complex procedure that involves performing many threshold scans for a fixed injected charge (normally 1 fC) for different trim DAC and range settings. The analysis then determines what settings of the trim DACs and trim ranges (these can be set independently for different chips) would give the same measured mean charge for all the channels in a chip or even for a whole module.

Another key performance measure of a module is its time walk. Time walk is defined as the maximum difference in the time the discriminator registers a hit for calibration charges between 1.25 fC and 10 fC. The time walk must be less than 16 ns so that, given the pre-amp rise-time of 10 ns, there will be an unambiguous association of a hit with an event (the beam crossing separation time is 25 ns). Time walk is measured using the strobe delay scan. The strobe delay scan involves measuring the average occupancy for different values of the strobe delay, which is the time delay between the clock and the discriminator readout. By performing strobe delay scans for different calibration charges the time walk can be measured.

Whilst doing this, the calibration must check for channels that are performing incorrectly. The very worst channels are masked off (so that they don't add to the average occupancy). A full characterisation therefore involves running a number of additional tests to check that all aspects of the modules are performing correctly. These include checking the electrical and optical connections to the module; testing the redundant links and looking for defects in the pipeline memory and de-randomising buffers.

This characterisation and calibration is a complicated task. It is especially difficult due to the large numbers of modules that must be tested. Testing must be done frequently, during macro-assembly to ensure that the modules have been mounted correctly and not damaged and during ATLAS running the modules will change with time due to radiation damage. The software to perform this task, SctRodDaq, is the subject of the next chapter.

SctRodDaq: Data Acquisition and Calibration Software for SCT Macro-Assembly

*What... is the air speed velocity of an unladen swallow?
What do you mean? An African or European swallow?*

This chapter presents the SctRodDaq software [82] which provides data acquisition and calibration for the SCT using the optical readout system. A high-level overview of the design and implementation of SctRodDaq is given followed by the results of tests that prove the software is fit-for-purpose. This software was written together with Alan Barr, Bruce Gallop, Dave Robinson and Alessandro Tricoli and many other people contributed in various ways. They are listed in the Acknowledgements. This chapter concentrates on the areas in which I contributed significantly, namely the analysis and control sub-system (see figure 4.1). For detailed descriptions of the other components see, for instance, [83]. The design and studies presented here relate to SctRodDaq version 3.0. Version 3.01 is the version in current use at the macro-assembly sites in Oxford, Liverpool and Nikhef and differs from 3.0 primarily in some features used in the hardware interface. This part of the code is not discussed in detail so the conclusions and detail of this chapter are not materially altered.

4.1 Aims

For many years, the SCT community has used the SCTDAQ [77] software for calibration and data-acquisition of detectors, modules, hybrids and at testbeam. However, it is intimately connected to the electrical readout system. For macro-assembly and final ATLAS, it was therefore neces-

sary that new software be developed to perform calibration and data-acquisition using the optical readout system.

The aims of SctRodDaq software are to:

1. Provide data readout and calibration control of modules via the optical readout system:

This is the essential requirement that necessitates writing a completely new software system.

2. Allow full calibration of the modules using the tests in the Electrical Test Specifications [84]:

This is the primary use-case for macro-assembly. It is essential to check that when modules have been mounted on a barrel or a disc that they still work and that there is no degradation in performance compared to the reception tests^a.

3. Scale to at least a full barrel using a reasonable amount of computer hardware:

The largest single unit during macro-assembly will be barrel 6 with 672 modules. A full characterisation sequence for this barrel would generate as much as 17 Gb of data which would all have to be analysed in a reasonably short time (see further discussion in section 4.8.3). Whilst it is not clear what a ‘reasonable’ amount of hardware is, clearly a room full of PCs would be unacceptable.

4. Integration with the ATLAS online software:

The system developed should be capable of developing into a system for final ATLAS. It is therefore essential to integrate with the ATLAS online software. This should also allow maximum re-use of existing development.

5. Robustness:

Macro-assembly will be operating under extremely tight time constraints^b so the software must be stable and dependable.

^aModules are tested when they arrive at the assembly institution using SCTDAQ. These tests include a standard characterisation sequence and are referred to as reception tests.

^bIt is currently one of the items on the critical path for ATLAS installation.

6. Code quality and maintainability^c:

SctRodDaq must be well designed and written so that it will remain usable for many years despite the major revisions that will be required to keep track of developments within ATLAS. Hence designing for change by using techniques such as modularity and interfaces is important. Also, since three of the five authors are students, it is extremely likely that its development and maintenance will be passed onto others. Therefore code quality, readability, maintainability and documentation are of the highest importance.

Of course, it is not enough to merely claim that the software meets these requirements. Developing the testing systems to prove that these aims have been met was an important part of the development. The results of this testing are presented in sections 4.7 and 4.8.

4.2 Overview

SctRodDaq is a large system; it currently comprises over 100,000 lines of C, C++, Java and Interface Definition Language (IDL) code^d which represents about 4 person-years of work. The scalability requirement immediately implied that the software had to be distributed across multiple computers rather than the single computer design of SCTDAQ. This of course requires that SctRodDaq be split into multiple processes, which fitted well with the multi-author, geographically distributed nature of the development team! Figure 4.1 gives a high-level overview of the design with the arrows indicating the main direction of interactions. The SctRodDaq boxes (in blue) correspond to individual processes which will be explained in the following sections.

^cMany authors have written on the importance of this. Consider for instance these quotes from [85]: “60% of software’s dollar is spent on maintenance, and 60% of that maintenance is enhancement” and “understanding the existing product consumes roughly 30% of the total maintenance time.”

^dThese figures were estimated from the ‘sclc’ script [86] and exclude blank lines and comments.

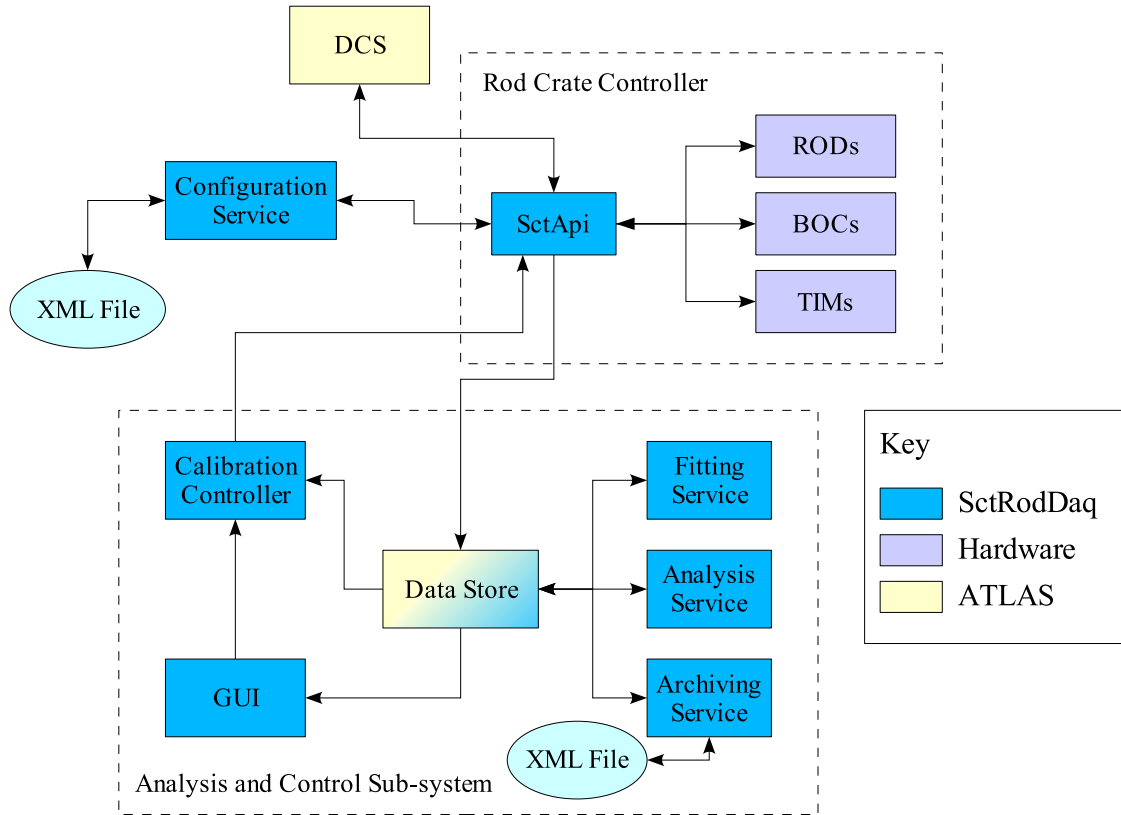


Figure 4.1: Schematic overview of the SctRodDaq software.

4.2.1 Functional overview

4.2.1.1 SctApi

SctApi is the interface between the software system and the hardware and thus is responsible for communications with the BOCs, RODs and TIMs and via them, the actual modules. In its current state, it can control a single crate and runs on the RCC, communicating with the various electronics via the VME bus. It provides an external application programming interface (API) that can be used to scan over various variables in order to characterise module performance. It outputs the data in a standard form to the analysis sub-system. SctApi provides many other features such as interactions with the detector control system (DCS) and scans over BOC parameters that are not relevant to this thesis. A complete description can be found in [83].

4.2.1.2 Configuration Service

The Configuration Service reads in the configuration of the electronics (RODs and BOCs) and the modules from a set of eXtensible Markup Language (XML) files. This configuration is then queried by SctApi which actually sets the configurations of the various components and is also used by the graphical user interface (GUI) to display the module layout. Also contained in the configuration are mappings to the various geographical and channel layouts (a module can be identified in at least five different ways!).

4.2.1.3 Calibration Controller

SctApi provides an interface for performing scans, however to fully characterise a module it is necessary to do groups of scans with some variable modified between each. These are referred to as ‘tests’. The Calibration Controller is responsible for performing tests and for updating the module configurations on the basis of the tests. It also informs the rest of the analysis system of what the user requested so that the appropriate analyses can be performed when the data is available. Finally, it currently includes the RunController which is the external interface to the ATLAS online software.

4.2.1.4 Fitting Service

The Fitting Service takes the raw scan data output by SctApi and fits functional forms to it. For instance, scans over discriminator threshold have a typical ‘S-curve’ shape that can be fitted with a complementary error function. This data is used in many of the analyses. The Fitting Service is described in more detail in section 4.4.

4.2.1.5 Analysis Service

The Analysis Service analyses the data based on what the user requested, using raw scan data and/or the output of the Fitting Service. It supports the analyses in the Electrical Test Specifications [84], but tries to match the most recent SCTDAQ version, 3.42. The analyses check for defects

which may be present and also measure the performance of the modules. The output is returned to the Calibration Controller which uses it to update the configurations. The Analysis Service also outputs summary data that is suitable for being uploaded to the SCT Production Database. The Analysis Service is described in more detail in section 4.5.

4.2.1.6 Archiving Service

The Archiving Service collects all the data that the other services produce and stores it to disk in a form suitable for long-term storage. It currently writes an XML-style^e file that should be robust to version changes and, if necessary, is human editable.

4.2.1.7 Graphical User Interface

The graphical user interface (GUI) presents the data in a simple way allowing the user to see at a glance all the modules on a barrel or disc (see the screenshot in figure 4.2). If desired, any of the data produced may be viewed as can the hardware configurations. The GUI is also the interface to the SCT Production Database allowing both uploads and downloads. The GUI can display DCS information using the same display allowing easy identification with modules. It also has scripting support to allow direct access to all the components in the system.

4.2.2 ATLAS online software

As discussed in the aims, it was important (as well as easier!) to use the ATLAS online software. The online software has many components and much more information can be found on their web pages [87]. Here I present a brief overview of the components that were used in SctRodDaq.

^eI refer to it as ‘XML-style’ since the file is not generated by a compliant XML tool, there is therefore no guarantee that the files produced are valid XML. The Archiving Service is under review for future versions and this aspect of its operation may well be improved.

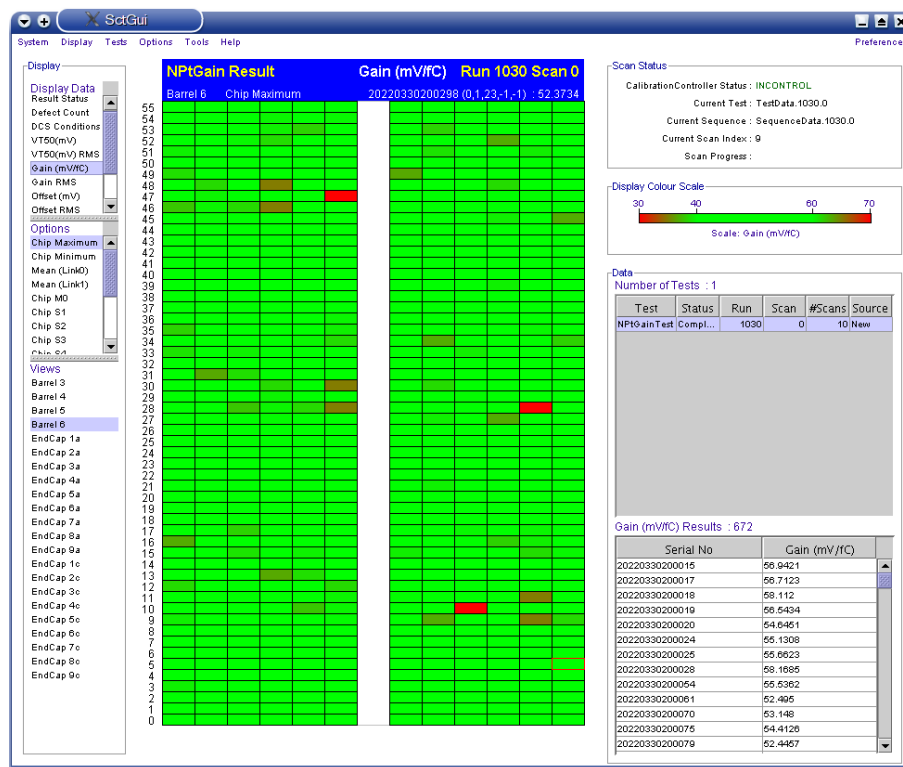


Figure 4.2: A screenshot of the SctRodDaq GUI. The central portion shows a (simulated) fully populated barrel 6 which has been unwrapped. The colours in the screenshot indicate values of the gain measured during a response curve test (see section 4.5.5). This view highlights those modules that may have problems.

4.2.2.1 Inter-Process Communication

This is the lowest level component of the online software, providing for the most important and basic requirement for a distributed system — namely that processes can communicate easily and efficiently. Inter-Process Communication (IPC) is a wrapper around a Common Object Request Broker Architecture (CORBA) implementation. CORBA is an open standard for cross-platform inter-process communication which has been written by the Object Management Group (OMG), an open industry-wide consortium. As it is just a standard, an implementation must be obtained of which there are many. The advantage is that since CORBA is an open standard, CORBA-compliant implementations are guaranteed to work together.

A process, the server application, makes available an external interface which is defined in

IDL. The CORBA implementation contains a ‘stubber’ that converts the IDL into a target language such as C++; this hides all the complexity of actually performing the inter-process communication. A client application can access the interface on the server by using a copy of the IDL file which is turned into code by a stubber in much the same way (note that the CORBA implementation used by the client and the server do not have to be the same). Clients and servers are contacted via an identifier that is world-unique. This means that processes can communicate, via the internet, to anywhere in the world if this is required. A complete description of CORBA is given in the specifications [88]; more pedagogical introductions can be found at the OMG and CORBA websites [89, 90] or in any of the many books on the subject.

Up to version 00-20-00, the online software used the CORBA implementation from Xerox; the Inter-Language Unification system (ILU) [91]. In 1997, when this decision was made, ILU was one of the few freely available CORBA implementations. It isn’t strictly CORBA compliant, rather it understands IDL and translates it internally into its own language. In particular the output code is not CORBA compliant which has important implications for the future of SctRodDaq since, from version 00-21-00, the online software has moved to a new implementation, omniOrb [92], which is CORBA compliant. This is discussed further in section 4.10.

All of the main components in SctRodDaq provide interfaces using IPC. These allow the user or the GUI to find out status information or to modify their behaviour (such as the analytic threshold scan fitting discussed in section 4.4.1.2). A number of utility programs are available which use these interfaces, such as the ‘Re-Analyser’ which uses them to re-analyse old data.

4.2.2.2 Information Service

The Information Service (IS) acts rather like a message board. Information providers can add, change or remove information from the board and interested clients can read the information that is there. Unlike a message board, clients can subscribe to receive notifications about changes that they are interested in. The IS server acts as the message board and there can be many of these (distinguished by different names). The server maintains the information until it quits but can also

save its state to disk.

IS passes the information around as XML-style strings — in fact the IS server treats these as entirely opaque blocks of data. IS includes client-side libraries that can convert C++ classes to and from these XML-style strings. This has important performance implications, discussed in section 4.3.3.

The Information Service is used extensively within SctRodDaq for two main purposes. IS is used as an event mechanism to inform services such as the Fitting, Analysis and Archiving Services that new data is available. This is discussed further in section 4.2.3. IS is also used by processes such as the Calibration Controller to publish their status information.

4.2.2.3 Message Reporting Service

The Message Reporting Service (MRS) provides a mechanism for processes to report what is happening with a variety of different levels of verbosity. This information can then be presented to the user and is also logged. The programs in SctRodDaq use MRS to report informative, diagnostic and error messages.

4.2.2.4 Configuration databases

The Configuration databases (not to be confused with the Configuration Service in SctRodDaq) provide a way to describe the software, the processes and the computers they must run on. This information is then used by the online software to start the processes on the appropriate machines.

4.2.3 Data flow

The flow of data through SctRodDaq is described in figure 4.3. The process starts when the user requests a ‘test’ using the GUI. The GUI forwards this to the Calibration Controller which breaks the test down into a sequence of scans. It records the request so that the rest of the analysis system can discover what the user wanted.

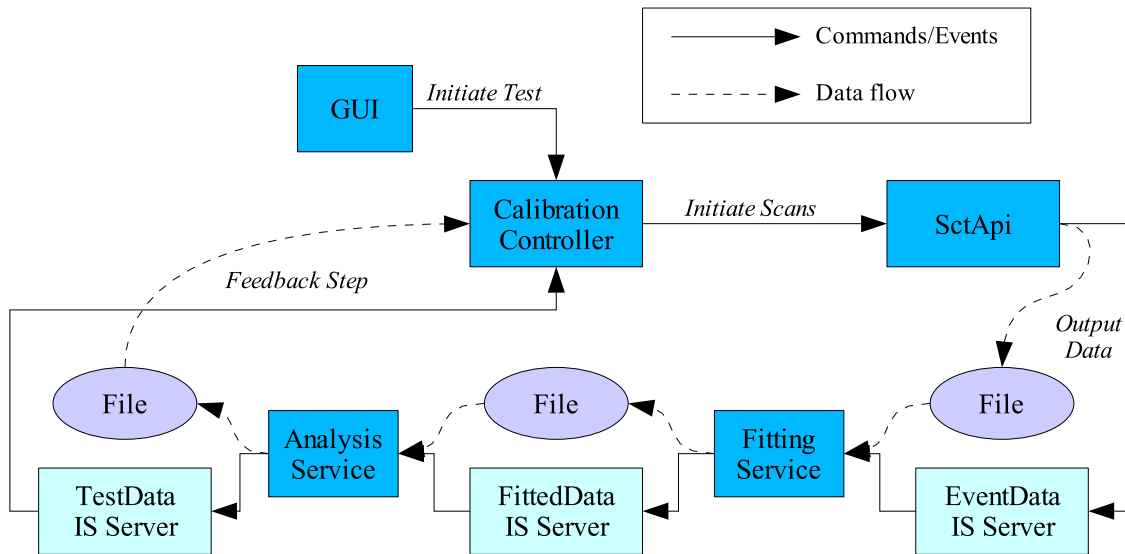


Figure 4.3: The analysis sub-system data chain.

The Calibration Controller then instructs SctApi to perform the scans. When each scan finishes, SctApi writes out the raw scan data to a file; one file for each module in the scan. Simultaneously, a record of this file is sent to IS (the implementation of this uses the I/O system described in section 4.3.3).

The IS server sends a message to the Fitting Service. If appropriate, the Fitting Service reads the data^f and does the fitting, writing the data to file (again, one per module) and recording that event in a different IS server.

The Analysis Service is notified of both raw scan data and fitted data arriving in the IS servers. When it has all the data it needs to perform the analysis for one module, it reads in all the data and performs the analysis. The data is written to a file with a notification sent to a third IS server.

When the Calibration Controller receives an event saying that the Analysis Service has finished analysing a module, it reads in the output and updates the module's configuration appropriately. When the data for all modules have been analysed, the Calibration Controller ends the test and returns control to the user.

^fThe files are stored in an network file system (NFS) directory which handles the transferring of the data between different computers. This is discussed in section 4.3.3.

The Archiving Service responds to events from all of the IS servers and saves the data to long-term storage. The GUI also has access to all the IS servers and can use this to show the test progress and to display plots to the user.

4.3 Design and implementation

4.3.1 General principles

As mentioned in the aims (section 4.1), good and readable code was an important consideration and this affects the design of the software. Throughout the development of SctRodDaq we have tried to use industry best-practise. Two books which guided much of the design are “Design Patterns” [93] and “C++ FAQs” [94]. Also the wealth of information and discussion in the Portland Pattern Repository [95] has been extremely helpful.

A brief introduction to *design patterns* is useful to understand some of what follows. Design patterns attempt to capture good solutions to common problems in a way that can be applied to many different scenarios. They were introduced first by C. Alexander who noticed that within architecture and town planning the same types of problems had to be solved many times, but each time with a slightly different solution. Alexander said, “Each pattern describes a problem which occurs over and over again in our environment, and then describes the core of the solution to that problem, in such a way that you can use this solution a million times over, without ever doing it the same way twice” [96]. This concept was introduced into computer science in [93] and the patterns included in it have become part of the language of object-oriented programming. Design patterns are important in two ways; firstly they help find a good solution to a problem at an architectural or design level. Secondly they provide a language that can aid comprehension of the code. For instance, names such as ‘Worker’, ‘Builder’, ‘Factory’ are now well known in object-oriented programming and immediately give a person trying to understand the code, an idea of what the problem is, what the solution is and how the code is trying to achieve that. Some examples of design patterns in use will be given in the following sections.

As is clear from section 4.2.3, the analysis sub-system operates on each module independently. This decision was taken early on to ensure that the system could scale to the large numbers of modules that it must handle. It would be a relatively simple exercise to allow the performance to be improved by running multiple Fitting Services or Analysis Services (see section 4.10). This principle does mean that the system cannot currently search for any problems which may cross modules, for instance correlated noise or cross-talk. It would be possible to add a new service to do this or else to perform these types of analysis offline.

4.3.2 External software

Much external software has been used whilst developing SctRodDaq to reduce development and increase robustness, portability and performance. A good example of this is the use of the Boost libraries [97]. The ‘smart pointer’ library was used to ensure that objects are always deleted[§] and the threading and testing libraries have also been used; these are briefly discussed later (see sections 4.3.4 and 4.6).

The ROOT libraries [98] have been used to provide data structures and display for graphs and histograms and the NAG C libraries [99] were used for high-performance curve fitting (see section 4.8.1). Several other libraries have also been used, particularly in the GUI to provide features such as scripting and graph plotting.

4.3.3 I/O system

Since different processes deal with different parts of the analysis, there must be a mechanism for moving data from process to process and machine to machine. This mechanism must be efficient and robust. The complexity of the in-memory data model and the need for efficiency suggested that a simple CORBA implementation was not possible. Instead, the solution taken was to use the Atomiser pattern [100] (see also the online version of [94]) combined with the Late Creation pattern [101].

[§]This greatly simplifies memory management, particularly in complex situations involving error handling. It also enables advanced object management schemes to be used — see the footnote on page 73 for an example.

A full description of this system is not appropriate here, but the basic principle is to split serialisation^h into two parts. The first part is to define a streamer class for each class in the in-memory data model. This class is a helper class that converts the in-memory data format into a *logical* format or vice-versa. The second part is to define an implementation, called the back-end, that writes out the logical format, converting it into a specific *output* format. This is illustrated in figure 4.4.

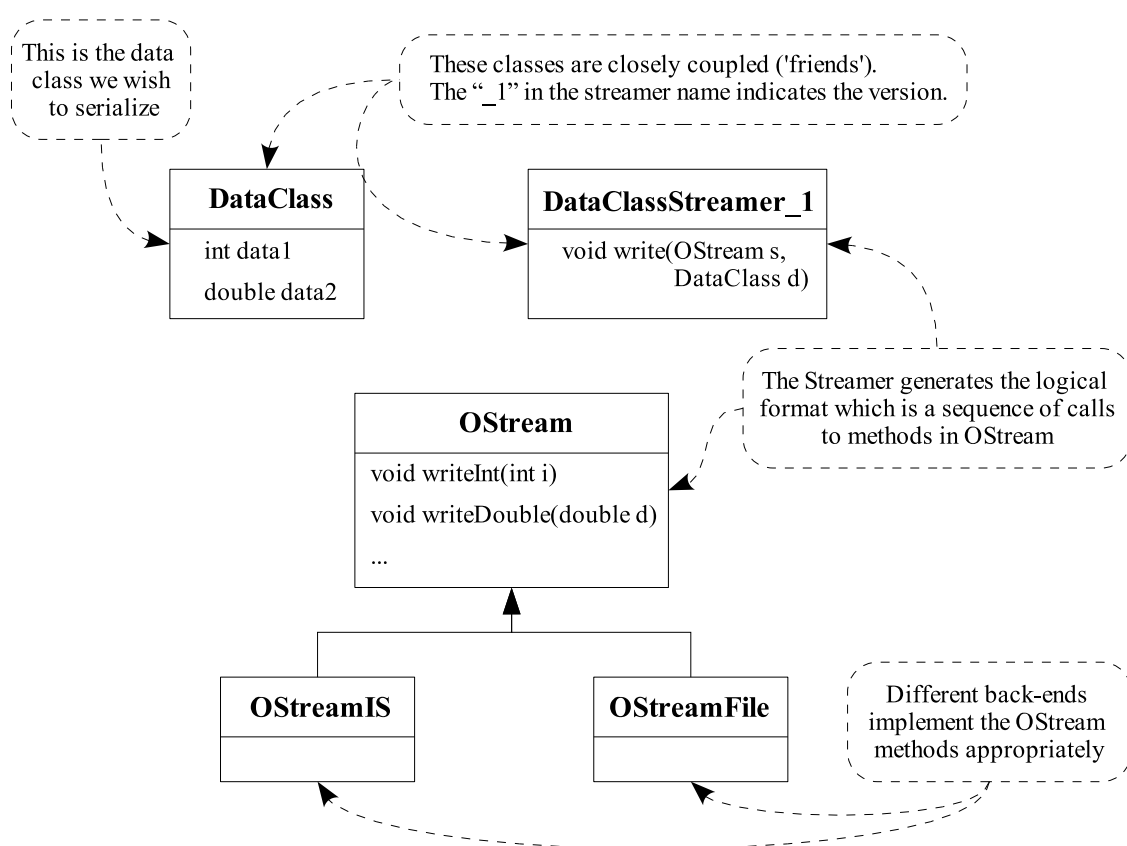


Figure 4.4: An illustration of the I/O system showing some of the output classes. The solid lines and arrow indicate that `OStreamIS` and `OStreamFile` are concrete implementations of the (abstract) `OStream` class. A symmetrical set of classes and methods exists for input.

The first back-end wrote its output to IS. However, it was realised that although IS had a high enough data transfer rate, it required a large amount of CPU time to generate and decode the XML-

^hSerialisation is the process of converting a complex set of objects into a stream of 1's and 0's. A symmetric de-serialisation process can then recreate the complex set of objects.

style strings that it uses. A replacement back-end was written which outputs a binary format to file. This is significantly faster. Since the data flow (see section 4.2.3) relies on information being added to IS, a proxy is written which simply contains the name of the file that contains the real data. When the proxy is read from IS, it automatically opens the file and reads the data from there. This change was transparent to both the data producers (the data model and associated streamers) and also the data receivers.

The binary file still needs to be transferred between different computers. This is achieved using NFS. NFS allows a directory on a remote server to be mapped to a location in the file system of a client computer. Transferring data between computers is thus entirely transparent and looks like simply reading and writing to files. There may however be issues regarding the use of NFS, see sections 4.8.3 and 4.10.

Binary formats have many disadvantages for long-term data storage, so the Archiving Service uses an alternative back-end that writes XML-style files. This back-end also includes version information. Whenever the data model is changed, the associated streamer needs to be modified and thus the logical format changes, rendering all previously saved data useless. This is solved by having multiple streamers, one for each version. It is up to the streamers to deal with the conversion of old data into the new data model. Since the data model should be reasonably constant, this should give stored data a long lifetime. The decision to use a text based format for the long-term storage of data means that it can be read by humans if necessary.

4.3.4 Threading

In a multi-process system, it is essential that all processes are available to be communicated with all of the time. This is particularly true if those processes must undertake time-consuming tasks. The danger of not taking steps to deal with this issue is that events will be dropped which could lead to data loss. It can also give the user the false impression that the system is not working.

There are two standard methods for dealing with this. The first is to regularly check for incoming messages during any long task. This approach is simple but leads to maintenance issues as

long passages of analysis code have to be punctuated with unrelated calls to check for messages. We implemented the other alternative which is to use threads. The patterns used to implement this are the Listener/Worker, Worker Pool and Job Queue patterns which are well described in [102].

The basic principle is to have a ‘listener’ thread which responds to incoming requests or events. It adds any work that needs doing to a queue. There are also one or more ‘worker’ threads. When a worker thread is idle, it checks the queue to see if there is more work that needs to be done. Since the listener thread has relatively little to do, it is always available to respond to new events. Another advantage of this pattern is that is simple to take advantage of multiple CPUs or to process multiple tasks at once (this is only profitable if the tasks are not CPU bound but are perhaps I/O bound). Figure 4.5 illustrates how this process works.

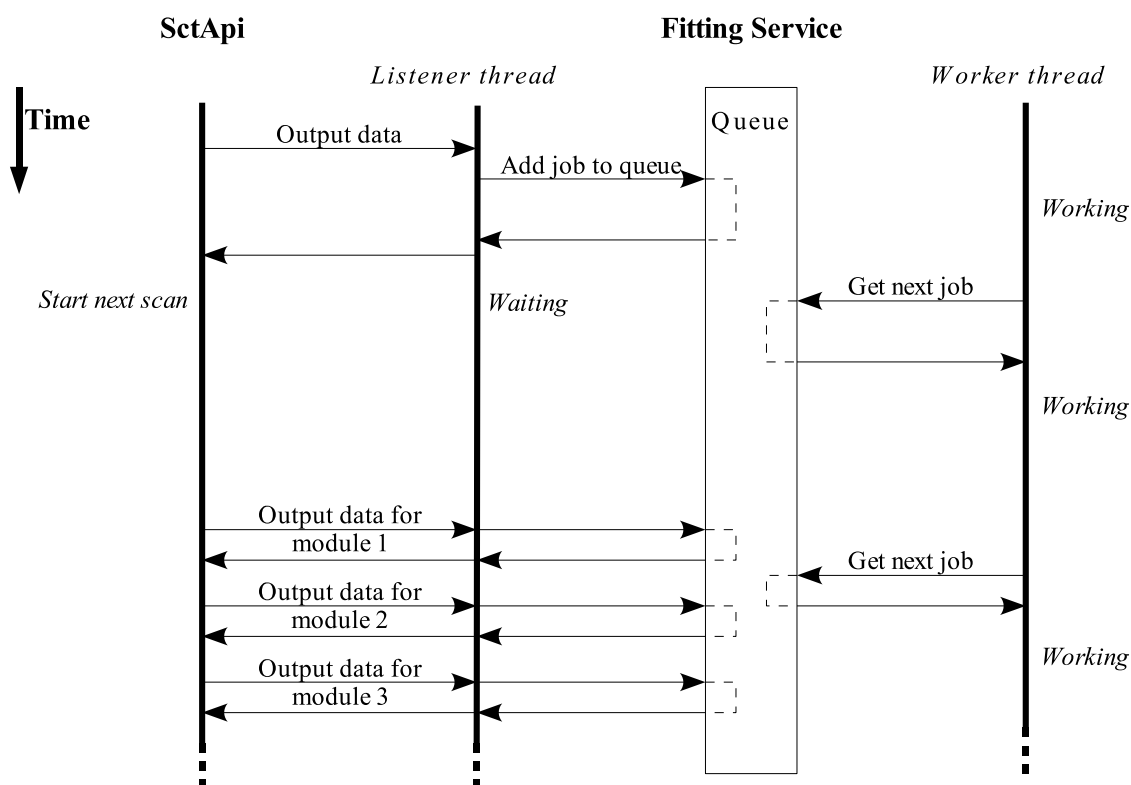


Figure 4.5: A schematic diagram (sequence diagram) of the threading model. The solid vertical lines represent the various threads (left the SctApi, centre and right the Fitting Service) with time increasing downwards. The horizontal arrows represent messages and function calls between the threads.

The model described is implemented and used in each of the main services: the Fitting Service, the Analysis Service and the Archiving Service. The implementation is greatly simplified by the use of the Boost thread library. We have also taken steps to ensure that our code is thread-safe. However, ROOT is not thread-safeⁱ so although the number of worker threads can be trivially increased, this mechanism cannot currently be used. Note that because the listener thread is very simple and doesn't use ROOT, this does not affect the basic operation of this system.

4.4 Fitting Service

The Fitting Service is an important part of the analysis sub-system. The fit results form the basis of several of the analyses (including the strobe delay, response curve and trim analyses) and it dominates the CPU load of the analysis sub-system. This section concentrates on the algorithmic details; details of the implementation are available in the documentation included in SctRodDaq [82].

4.4.1 Algorithm

The Fitting Service provides a relatively simple function, but must provide it quickly. The algorithm has therefore been generalised and applies to all fits, whether they are done to chips or channels and regardless of the type of data. The outline of the general algorithm is:

1. Read data from disk and determine the appropriate algorithm.
2. Fit channel data.
3. Veto serious defects and masked channels.
4. Fit chip data (if requested).
5. Output data to disk and send event notification to IS.

ⁱThere are some aspects which are thread-safe, but the most important: object creation and destruction isn't. This makes ROOT at best slow (since multiple threads cannot construct objects simultaneously) and in practise useless for multi-thread applications. See also comments in section 4.10.

The data that the Fitting Service expects to deal with is scan data from the RODs. This normally takes the form of a 2-d histogram as shown in figure 4.6a. The channel is shown on the x -axis with the variable being scanned over on the y -axis and the number of hits is the bin content (the number of triggers sent is recorded elsewhere and is always the same for all the channels, but may be different for different bins). Data may be fitted for individual channels, also chip data may be fitted. In this case a projection across all the channels in that chip is formed, ignoring masked channels and channels with serious defects which gives a result like that shown in figure 4.6b, showing the occupancy as a function of the bin for the required channel or chip. The procedure for fitting data from a single channel or chip is as follows:

1. First an occupancy projection of the histogram across the required channels is created. The occupancy for bin i , ρ_i , and its error, σ_i , are given by

$$\rho_i = \frac{z_i}{n_i}, \quad \sigma_i = \sqrt{\frac{(z_i + 1)(n_i - z_i + 1)}{n_i^3}} \quad (4.1)$$

where $z_i = \sum_{c \in \text{Channels}} z_{ci}$ and $n_i = \sum_{c \in \text{Channels}} n_{ci}$

and z_{ci} and n_{ci} are the number of hits and triggers respectively. The errors calculated in this way are approximately binomial with the addition of ‘softening factors’ to prevent the error being zero at an occupancy of 1 or 0. This is the approach taken in parts of SCTDAQ.

2. A pre-check is made to see if there are any obvious defects. If there are any serious defects, then fitting is abandoned.
3. An initial guess of the fit parameters is made. Quick, accurate guessing can speed up fitting dramatically.
4. The algorithm returns a function that is used to fit the data using an abstract interface to the fitting and minimisation code.
5. Finally there is a post-check for defects.

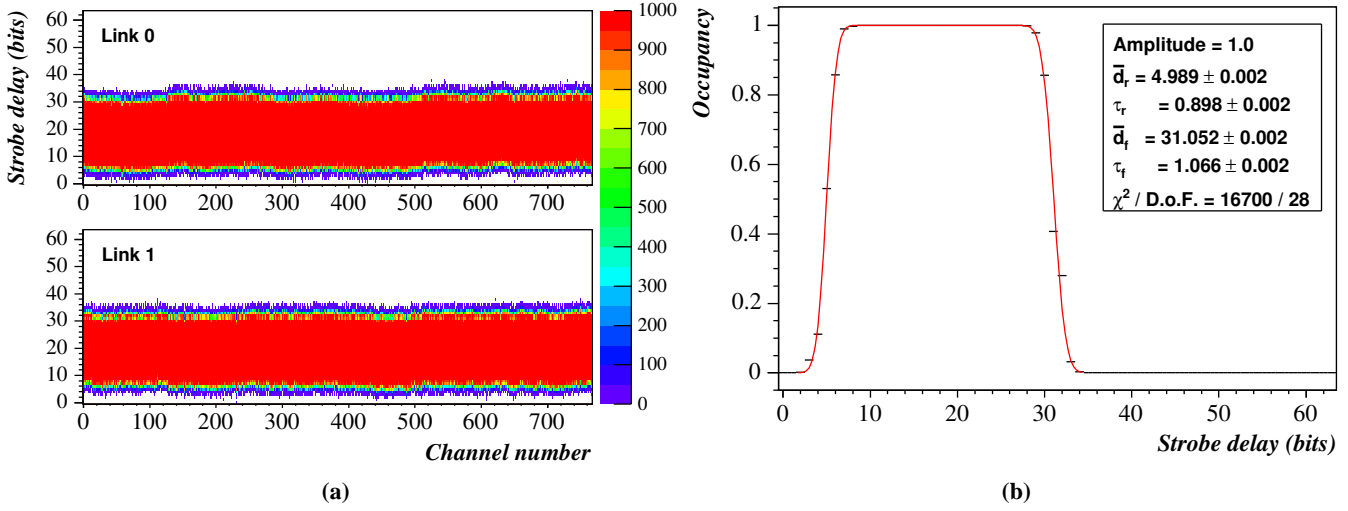


Figure 4.6: Example of strobe delay data and fit. **(a)** shows the raw data with the hits shown on the colour scale (white indicates no hits). There were 1000 triggers for each channel and bin. **(b)** shows an occupancy projection for one chip with a fit (red line).

There are only two types of analogue data currently defined: strobe delay scans which have a ‘top hat’ shape and threshold scans which have an ‘S-curve’ shape (see section 3.4). There are two specialised fitting algorithms for each of these cases that set up the fit function, perform the necessary checks for defects and guess the initial parameters.

4.4.1.1 Strobe delay scans

The occupancy, ρ , as a function of strobe delay, d , is fitted with the product of an error function and a complementary error function which gives a smoothed top hat shape

$$\rho = \frac{A}{4} \left[1 - \operatorname{erf} \left(\frac{\bar{d}_r - d}{\sqrt{2}\tau_r} \right) \right] \left[1 + \operatorname{erf} \left(\frac{\bar{d}_f - d}{\sqrt{2}\tau_f} \right) \right] \quad (4.2)$$

where A is the maximum amplitude, \bar{d}_r is the position of the rise, τ_r is the rise time, \bar{d}_f is the position of the fall and τ_f is the fall time. An example is given in figure 4.6. This differs slightly from the implementation in SCTDAQ which fits an error function to the leading edge and a separate error function to the trailing edge. This allows the possibility of dips between the two edges being

missed and this has to be checked for explicitly. In SctRodDaq there is a χ^2 check that should catch this type of problem. Since the strobe delay is set chip-by-chip, only chip data is fitted.

During the parameter guessing stage, the rise (fall) position is set to the first (last) bin with occupancy greater than 50%. The times are guessed by finding the 82% occupancy bins and subtracting this from the rise/fall positions. The fit range is set to the bins with occupancy of more than 1% with an extra bin on either side.

4.4.1.2 Threshold scans

The threshold algorithm fits occupancy data as a function of threshold (measured in mV), t_V , with a complementary error function:

$$\rho = \frac{A}{2} \left[1 + \operatorname{erf} \left(\frac{\bar{t}_V - t}{\sqrt{2}n_V} \right) \right] \quad (4.3)$$

where A is the amplitude, \bar{t}_V is the mean threshold, and n_V is the width. Recall from section 3.4 that \bar{t}_V corresponds to the injected charge and n_V is the output noise in mV at that charge. An example is given in figure 4.7.

The mean threshold is guessed to be the position of the first bin with less than 50% occupancy and the width is the difference between this bin and the first with less than 82% occupancy. The fit range is set to the bins with occupancy between 99% and 1% with an extra bin on either side.

An alternative threshold fit algorithm is available which uses the analytic properties of the error function to determine \bar{t}_V and n_V . This is called the analytic threshold algorithm and works by calculating the difference of the occupancy data which should be (approximately) Gaussian. The mean and standard deviation of this data then correspond to \bar{t}_V and n_V respectively. This method is of course very fast since it does not need to evaluate any functions or do any minimisation. However, threshold scan data is not always well described by an error function, so this algorithm is available for testing purposes only.

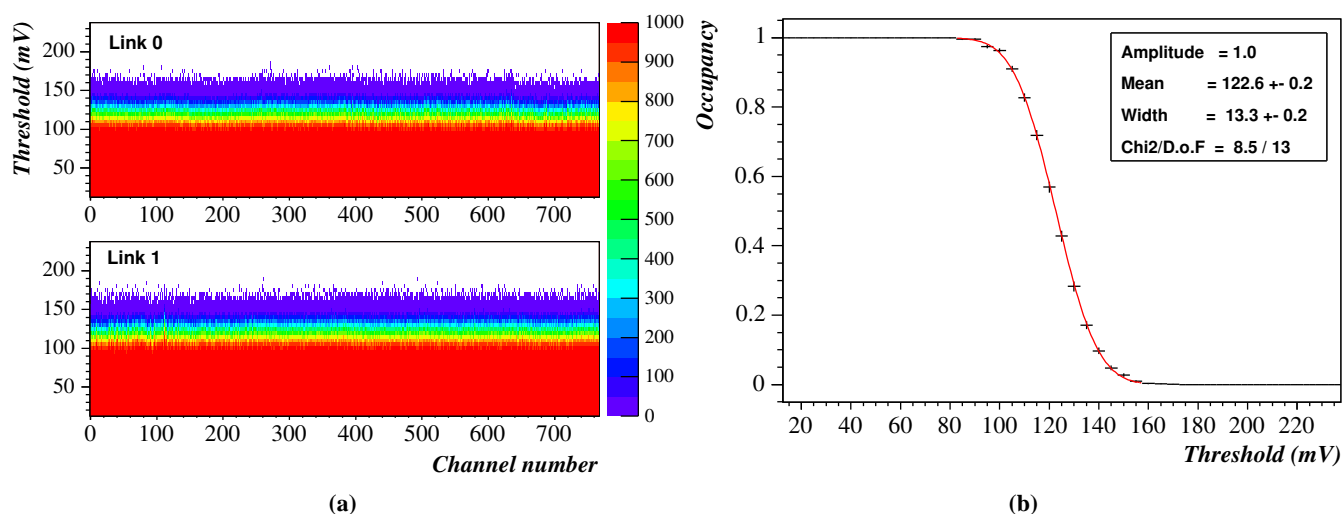


Figure 4.7: Example of threshold scan data and fit. **(a)** shows the raw data with the hits shown on the colour scale (white indicates no hits). There were 1000 triggers for each channel and bin. **(b)** shows an occupancy projection for one channel with a fit (red line).

4.4.2 Optimisation

In SCTDAQ, fitting S-curves is the slowest step and dominates the time taken to characterise a module. In a full characterisation of barrel 6, over 27,000 threshold scans must be analysed. Fitting this much data on a single PC is a demanding task so optimisation of the Fitting Service was necessary. As already mentioned, changes were made to the I/O system to ensure that it did not take up much CPU time. Several steps were taken to make sure the fitting itself was as fast as possible:

- The fit ranges are reduced during the ‘guess initial parameters’ stage. This leads to a large reduction in the number of times the fit function is called and thus a significant speed up in fitting time. It can potentially alter results, particularly in the threshold scan case, if the normalisation is allowed to vary.
- The normalisation is not allowed to vary but was determined by the maximum value. This prevents the problem mentioned above and speeds up the fitting by reducing the number of parameters.

- Both the threshold fit function and the strobe delay fit function involve many calls to evaluate the error function (or its complement). These were removed and replaced with a cached function which pre-calculates an array of values. When asked to evaluate the error function, it interpolates^j between the stored values.
- Each fit is stored in a 'FitObject' class. Since there can be up to 1548 fits per modules, this could lead to a lot of objects being briefly created and then destroyed when the data had been written. This is particularly slow since FitObjects use ROOT classes and creating and deleting ROOT classes is very slow. The solution is to use a pool of pre-created FitObjects that are recycled. Instead of creating FitObjects, free objects are returned from the pool. Whenever a FitObject is no longer needed, it is not deleted, but returned to the pool^k.
- The actual fitting is implemented through an abstract interface. This allows multiple fitting strategies (another design pattern) to be provided. The original implementation used ROOT, but it was noticed that this was very slow, so a new strategy that uses the NAG C library was developed. This uses a custom χ^2 routine along with the NAG minimisation code and gives a significant speed increase. The fitting strategy that is actually used can be chosen at runtime using the Fitting Service's IPC interface; the default is the NAG strategy.

The code optimisation was focused on those areas that really took the most CPU time. This was achieved by using the valgrind [103] and KCachegrind [104] tools to frequently profile the software.

^jInterpolation is used so that the first derivative of the function is continuous which is a requirement of most minimisation routines.

^kThe implementation of this was greatly simplified by the use of Boost smart pointers and the Factory Method pattern. The Factory Method pattern delegates creation of objects to a special method (rather than using the new operator). This method was modified to use the pool and the smart pointers returned are configured so that, instead of deleting the object, they return it to the pool. The use of this pool was thus invisible outside of the Factory Method.

4.5 Analysis Service

In this section a brief description of the logic and algorithms used by the Analysis Service is given. Only those that are used to calibrate and characterise SCT modules are discussed although others are available.

4.5.1 Framework

As in the case of the Fitting Service, it was important that the design of the Analysis Service deal with modules independently, recall that both scan data and fitted data are produced for single modules. The algorithm for handling data is:

1. When a new test is started, the Calibration Controller records what has been requested in IS. The Analysis Service reads this and finds the appropriate algorithm for analysing the data.
2. Whenever scan data is produced by SctApi or fitted data by the Fitting Service, the Analysis Service records that this data is available.
3. If all the data needed for the appropriate algorithm is available (some algorithms need fitted data, some scan data, some both) then it is run.
4. All of the data that is needed for that algorithm and that module is read using the I/O system.
5. The analysis is performed and the result is written to disk with a notification sent to IS.
6. A summary file is generated in a form that can be sent to the SCT Production Database and is stored in IS.

4.5.2 Digital tests

There are three digital tests. The first, the 'NMask' test is a test of the readout system. In this test, no calibration charge is injected, instead, a pulse output command is sent shortly before the level 1 trigger which causes all the discriminators to record a hit. Since the module mask still applies,

the data returned should be the contents of the mask registers. The NMask test uses this to check the readout system, stepping through different settings of the mask register to generate different events.

The pipeline test aims to test the FIFO pipeline buffers in the ABCD3TA chips — it is possible for parts of these buffers to become dead or stuck. It uses the same technique as the NMask test of using the pulse output command. In the pipeline test though, a soft reset is sent followed by a variable delay before sending the pulse output command and the level 1 trigger. The delay is scanned over and ensures that each of the cells in the FIFO buffer are tested. Two scans are performed: one with all the channels on and one with all the channels masked off. This allows dead and stuck cells to be identified.

The full bypass test is a test of all the redundant links in the ABCD3TA chips. These are available to improve the ABCD3TA's robustness to radiation. High radiation doses can cause some chips to stop functioning. Due to the way readout occurs (recall that there is one data fibre for 6 chips and one control fibre for a whole module), this could prevent an entire module being used. The redundant links can be used to bypass the affected chip, restoring readout from the others.

All of the digital tests are simple to analyse as they have fixed output patterns. The expected output data is shown in figure 4.8 (pipeline scans are not shown as they should result in either occupancy of 1 or 0 everywhere).

4.5.3 Strobe delay

The strobe delay test aims to set a value for the strobe delay that will ensure that all of the other calibration tests can be carried out. It therefore sets the calibration charge to 2.0 fC and does a strobe delay scan. The strobe delay algorithm is trivial — any strobe delay in the main part of the top hat should be fine. The algorithm loads the fit of the strobe delay data, produced by the Fitting Service, and sets the optimum strobe delay to a fraction of the gap between the leading and trailing edges:

$$d_{\text{opt}} = f(\tau_f - \tau_r) \tag{4.4}$$

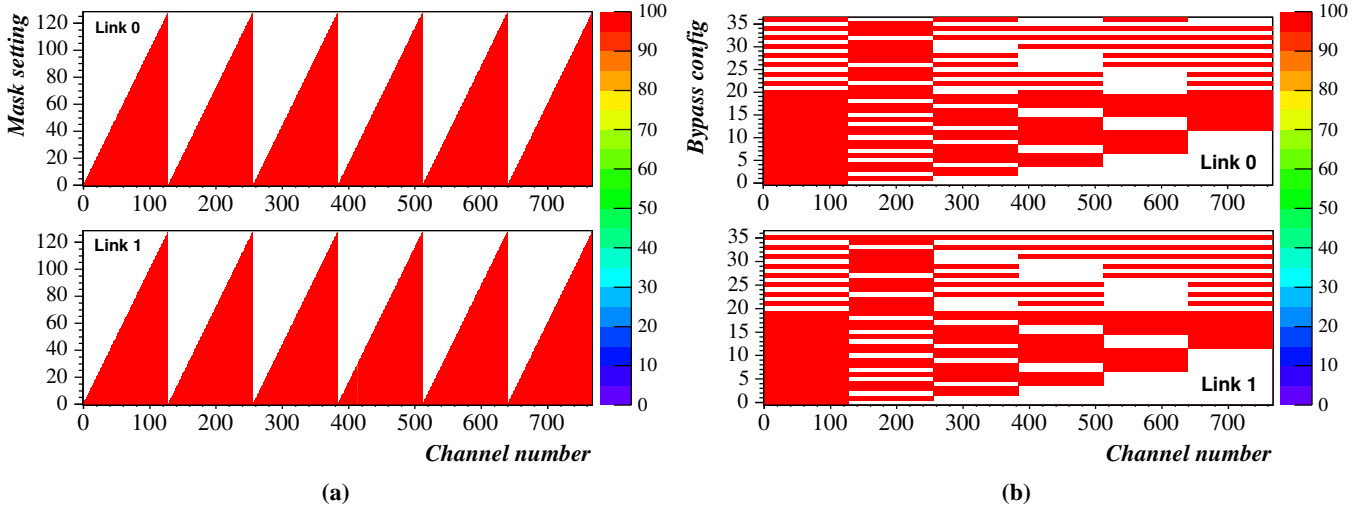


Figure 4.8: Example digital scan data for (a) NMask test and (b) full bypass test. Both scans had 100 triggers for each bin.

where the fraction, f , is 25% in SctRodDaq version 3.0. In SCTDAQ this fraction has recently been increased to 40% to account for some hardware effects and SctRodDaq followed suite in version 3.01. The algorithm checks that the strobe delay is within the range $28 < d_{\text{opt}} < 35$.

4.5.4 Trim algorithm

As mentioned in section 3.4, the channel trims must be calibrated to minimise the channel-to-channel variations in the mean threshold. Figure 4.9 shows the mean threshold before and after trimming. The key performance measurement of the trim algorithm is therefore the standard deviation of the mean threshold, σ_t .

The data that is collected consists of 28 threshold scans, all with the calibration charge set to 1.0 fC. The first 16 are for trim range 0 with one for each of the 16 possible trim DAC settings for each channel (possible trim DAC settings are 0–15). To speed up calibration time, only 4 thresholds scans are done for each of the other ranges, with the trim DACs set to 3, 7, 11 and 15 for all the channels.

The analysis proceeds as follows. For each trim range and each channel, a graph is formed




Figure 4.9: Mean threshold with a calibration charge of 2.0 fC (a) before and (b) after trimming.

of the mean threshold against trim DAC setting. This graph is approximately linear and a straight line is fitted to it (using least-squares regression). This line can then be used to calculate the trim DAC setting for a required mean threshold.

The algorithm then considers a range of possible trim targets. The trim target is the mean threshold that the algorithm will try to set each channel to. The range considered is 2.5–302.5 mV in 2.5 mV steps (the threshold DAC has a step size of 2.5 mV). For each target, the trim DAC setting for each channel to get that target is calculated from the straight line fit. If this value is from -1 to 16 inclusive¹, then that channel is trimmable. The total number of trimmable channels is stored.

The algorithm then chooses the trim ranges and DAC settings that will minimise σ_t . There are three possible modes: both trim range and trim target are allowed to vary from chip-to-chip; the trim range can vary but the trim target is the same for all chips or both the trim range and trim target must be the same for all chips in the module. The default is for the trim target to be the same for all chips, but to allow the range to vary from chip-to-chip. The best range and target are determined using the following criteria:

¹For an explanation of this, see section 4.7.5.

1. The chosen range and target should maximise the total number of trimmable channels.
2. The range should be minimised. This is done by summing the trim ranges for all chips and choosing targets which minimise this number. Note that this means that one chip in range 3 (and the rest in range 0) is preferred to 4 chips in range 1 with the rest in range 0.
3. If there is an range of possible trim range and target values, the chosen values should lie in the middle of the possible range. The way this is currently implemented means the algorithm is likely to fail for modules which have two peaks in the optimum trim range and target settings.

This data is then stored and defects are created for any channels that cannot be trimmed. Checks are also made on the straight line fits. If the values are outside normal ranges, then defects are recorded.

4.5.5 Response curve

The response curve aims to determine how voltage measured at the comparator in the ABCD3TA relates to the charge that was deposited in the strip. This is achieved by running several threshold scans. At the mean threshold, the voltage is equivalent to the calibration charge that is injected, *i.e.* $V = \bar{t}_V$ at $Q = Q_{\text{cal}}$. Thus each threshold scan gives one point on the response curve. By performing several threshold scans for different calibration charges, the response curve can be determined.

Two versions of the response curve are available: the full version does threshold scans at 10 calibration charges (0.5, 0.75, 1.0, 1.25, 1.5, 2, 3, 4, 6, 8 fC) whilst the quick version, called a three-point gain, does only three points: 1.5, 2.0, 2.5 fC. Both types are handled identically except that the three-point gain data is always fitted with a straight line. The response curve is normally fitted with the equation

$$V = a + \frac{b}{1 + e^{-Q/c}} \quad (4.5)$$

where a , b and c are fit parameters. A quadratic and the Grillo function^m are also available. It is possible for the threshold DAC to saturate if the input charge is greater than 5 fC. This leads to an occupancy histogram that does not drop to zero leading to fitting problems. This is known as the ‘8 fC effect’ and is dealt with by ignoring any affected points. In this case, a comment is included in the output.

The response curve is summarised by three parameters: the offset, $o = V(0)$, the gain, $G = \frac{\partial V}{\partial Q}$ and the noise, $n_E = 6250 \times n_V / G$ where n_V is the noise in mV from the threshold scan at 2.0 fC (if available). This equation converts the noise in mV to equivalent noise charge (ENC)ⁿ.

This procedure is followed for each channel and chip. The fit parameters for the chips are stored in the output module configuration – they can then be used to set a discriminator threshold corresponding to a required charge accurately. The channel fits are done as they allow a number of potential problems to be diagnosed. For instance, a low noise can indicate a bonding failure or a break in the detector strip and low gain can indicate problems in the ABCD3TA ASIC. Figure 4.10 shows some typical response curve data.

4.5.6 Noise occupancy

The noise occupancy with threshold set to 1.0 fC is a key performance measurement for SCT modules — the specification is that it should be less than 5×10^{-4} . The noise occupancy test uses a threshold scan with no injected charge. However, because the occupancy at 1.0 fC is so low, the scan is modified so that more triggers are sent for points which are at higher charge.

This algorithm uses both the S-curve fit from the Fitting Service and the raw data. The mean threshold from the S-curve is called the offset and, since it corresponds to 0 fC injected charge, it should be the same as the offset from the response curve test. The mean noise occupancy, $\bar{\rho}_n$, is the average occupancy at 1.0 fC for all channels in a chip and the standard deviation, σ_ρ , are also key measurements of this analysis.

^mThe Grillo function is $V = a + \frac{bcQ}{\sqrt{(bQ)^2 + c^2}}$ where a , b and c are fit parameters.

ⁿENC is defined as: $ENC = Q/e$.




Figure 4.10: Some typical response curve plots. (a) shows a response curve with a fit through the points. (b) shows the calculated input noise across a module. The dashed lines indicate chip boundaries and the chip structure is clearly visible.

The raw data is used to form a graph of $\log \rho$ against t_V^2 for occupancies below 50%. When the occupancy is small, the shape of the S-curve is expected to be approximately

$$\rho \approx A \exp \left[-\frac{t_V}{2n_V^2} \right]. \quad (4.6)$$

Thus a graph of $\log \rho$ against t_V^2 should be approximately straight with the gradient proportional to n_V . This graph is calculated for each chip, using a projection across the channels formed in the same way as in the Fitting Service (see section 4.4.1). Errors are asymmetric errors calculated directly from the approximately binomial errors in the projection. A straight line fit is made and the noise in ENC is calculated as $n_{ENC} = 6250 / \sqrt{-1/2a}$ where a is the gradient. An example fit is shown in figure 4.11. All of this data is stored and a check of the noise occupancy is made for all the channels. Any with a noise greater than the specification (5×10^{-4}) are masked off.

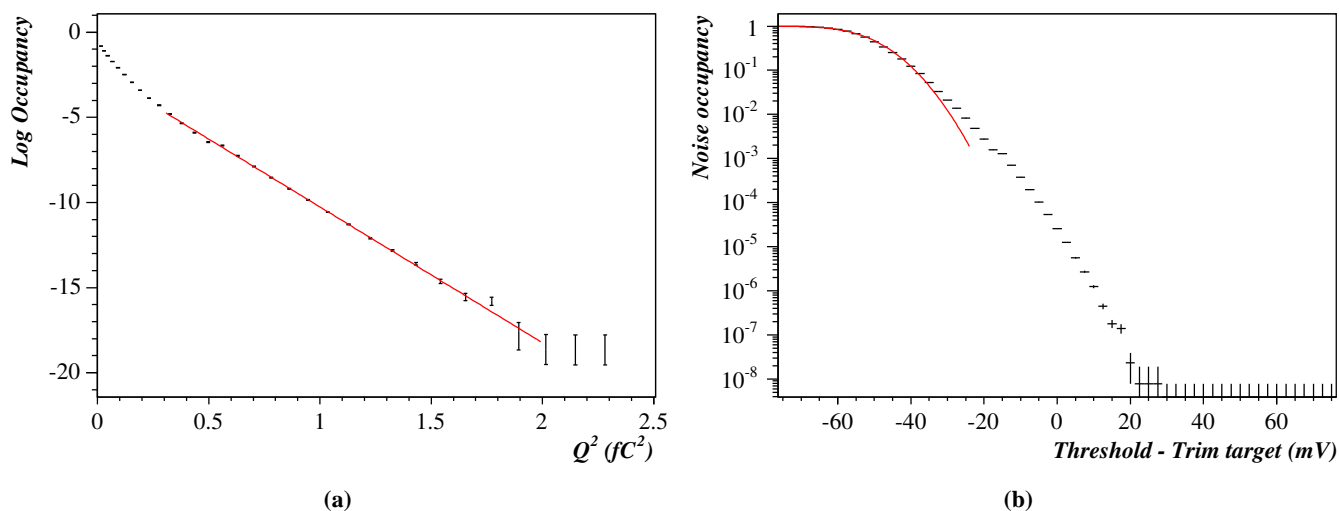


Figure 4.11: Example noise occupancy data. (a) shows a graph of $\log \rho$ against t_V^2 with the straight line fit. (b) shows the threshold S-curve (relative to the trim target, thus 0 mV indicates 1.0 fC) with the fit from the Fitting Service (red line). Note that the fit is increasingly bad at low occupancy.

4.5.7 Time walk

Recall from section 3.4 that the time walk is defined as the maximum difference in time that the discriminator registers a hit for charges between 1.25 fC and 10.0 fC. The time walk can be measured by performing a series of strobe delay scans with the calibration charge between 1.25 fC and 10.0 fC. Before the time walk can be measured, the strobe delay DAC must first be calibrated. This is achieved by putting the modules in edge sensing mode and using the 10.0 fC strobe delay scan. At 10.0 fC, the input to the discriminator circuit will rise very sharply giving a sharp low edge to the strobe delay data. A sharp high edge is given by being in edge sensing mode. In this mode, the module will only return a hit in a given time bin if there was no hit in the previous time bin. Since the shaper circuit has a rise time of 50 ns, the input to the discriminator will be higher than threshold for longer than 25 ns. Therefore, once the strobe delay is increased by more than 25 ns after the low edge, the discriminator output will be high in the time bin before the trigger as well as the one containing the trigger. Thus no hit will be returned. This means, that the width of the strobe delay must be 25 ns. This is used to calibrate the strobe delay DAC.

The time walk is measured by taking the maximum difference between the positions of the trailing edges and converting it to nanoseconds using the calibration factor. This is slightly different to the method used in SCTDAQ which fits a Gaussian to the strobe delay data and measures the movement of the peak position. A defect is recorded if the measured time walk, Δt is outside the range $5 \leq \Delta t \leq 15$.

4.6 Testing systems

Testing software is extremely important — it is only with comprehensive test suites that it can be shown that the software satisfies its criteria and is (comparatively) bug-free. Also, tests can give confidence that when changes are made, as frequently happens, nothing has been broken. This confidence can greatly increase productivity and reduce the time to implement new features. SctRodDaq has several testing systems that operate at different levels and are described in the following sections.

4.6.1 Unit tests

Unit tests^o are tests that operate on a code unit, check it behaves properly and run automatically. These are used extensively within SctRodDaq to check the internal operation of much of the C++ and Java code. A good example of this is the ‘ModuleElement’ class which is part of the data model. The ModuleElement class represents a continuous range of channels in a module. It is therefore defined by a start channel and an end channel (the range is inclusive). It is one of the basic building blocks of the data model and so must be known to operate correctly. It also provides some non-trivial but useful methods. This is a simplified section of the class definition:

```
class ModuleElement {  
    ModuleElement(int firstChannel, int lastChannel);  
    ...  
}
```

^oSee discussions in [95, 105].

```
//return true if this element is contained in or equal to the argument
bool isWithin(ModuleElement m);

//return true if this element contains or is equal to the argument
bool contains(ModuleElement m);

...
};
```

To be certain these methods work as intended, a set of unit tests is created which look like:

```
ModuleElement element1(562, 589);
ModuleElement element2(570, 580);

BOOST_CHECK(element2.isWithin(element1));
BOOST_CHECK(!element2.contains(element1));
```

where `BOOST_CHECK` is a macro that checks that its argument is true. As can be seen, this technique provides a useful external check of a class. If the implementation of `ModuleElement` is changed and a bug introduced, then the tests should catch it^P.

This technique is increasingly used in industry; some programming methodologies [105] even recommend writing the tests *before* writing the code. This way, the tests initially fail, but when the code is written, they pass. An important part of this testing is that the tests are carried out automatically every time the code is built; tests should also be quick and easy to write. This functionality has been provided by the Boost testing framework in the C++ code and by JUnit [106] in the Java code.

^PIn case the reader thinks this is unlikely, it is worth pointing out that these tests did in fact catch a bug in the logic of one of the more complex methods of the `ModuleElement` class.

4.6.2 High-level tests

Whilst the unit tests can test the internal logic of the code, it is difficult to test algorithms like the ones in the Fitting and Analysis Services — they need data to operate on. Consequently a test system was written to convert data taken from modules during reception tests and convert them to the appropriate format for analysis. This system consists of a set of scripts and programs that locates the reception test data and publishes it in a format that corresponds to the output of SctApi (see figure 4.3). It can therefore be used to test the Fitting Service, the Analysis Service, the Archiving Service and some aspects of the GUI. It cannot test the operation of the Calibration Controller or of SctApi however — this is left to the low-level testing system.

This system has been used to produce the automated correctness tests that will be discussed in section 4.7. It was also used to test the performance of the individual services in section 4.8.

4.6.3 Low-level tests

As mentioned above, the high-level test system cannot test the Calibration Controller or SctApi. To reduce the need for testing these components with real hardware (although this is needed as well), a low-level test system was devised. This extends SctApi and overrides the methods dealing with its interaction with hardware. It provides a dummy interface that, when asked to do a scan, attempts to match what was requested with data it has available from the module reception tests. Although there are some limitations to the different tests this system can provide, it has proved a useful testing tool, particularly for ensuring that the logic of the Calibration Controller is correct. This tool has also been used to stress the whole system. Simulations have been done with up to 672 modules to ensure that the system can indeed handle that amount of data (see section 4.8.3).

4.6.4 System tests

The system tests are a loose collection of tests that are too complex to be made into unit tests, but should still be run automatically. There are many tests currently, from tests to ensure that the data model is the same in Java and C++, to tests of aspects of SctApi and the Configuration Service as

well the correctness tests that will be discussed in section 4.7. These tests are run automatically when a release is built — anything failing them should be picked up at the beta or release candidate stage. This represents a good compromise between the execution time of the tests and picking up problems early on.

4.6.5 Other tests

Producing a system of the size of SctRodDaq is not trivial and the only way to be sure that the software is acceptable is to extensively test it. As mentioned earlier, this also gives confidence that changes can be made without breaking everything which in turn encourages practises such as re-factoring⁹ and leads to better, more understandable, more maintainable code.

Many other tests have been made on a more ad-hoc basis. These include memory leak testing with valgrind [103], user evaluation and performance testing, which will be discussed later.

4.7 Correctness testing

A considerable amount of effort has gone into ensuring the correctness of the algorithms used in SctRodDaq, particularly in the analysis sub-system. In this, it has been extremely helpful to have SCTDAQ, and the vast amount of data generated using it, to compare to. As mentioned in section 4.6.2 a test system was created that could take SCTDAQ data and run it through the analysis sub-system. This has been used to do correctness testing which used the reception test data available at Cambridge. This consisted of full characterisation sequences of up to 194 modules.

The test procedure was to run the data through SctRodDaq, then to download the SCTDAQ test data from the SCT Production Database. As many as possible of the variables were compared and significant differences were understood. Specific differences are discussed in the following sections. Typically the difference between the SctRodDaq and SCTDAQ measurements of some parameter have been plotted. This is always given as the SctRodDaq parameter minus the SCTDAQ

⁹See for example the page ‘RefactorMercilessly’ in [95].

parameter. Often fractional differences have been shown, in which case, they are fractions of the SCTDAQ measurement.

This testing had many benefits besides testing SctRodDaq. For instance, two modules were found to have been transposed in the SCT Production Database. Several bugs were also found in SCTDAQ as well as differences between SCTDAQ and the Electrical Test Specifications [84].

4.7.1 Threshold scan comparisons

The basis of many of the analogue scans is the fitting of threshold scan data (see sections 3.4 and 4.4.1) — so called ‘S-curves’. The algorithm for fitting S-curves is repeated three times in SCTDAQ. The algorithm presented in section 4.4.1 is the same as two of them and it was believed that this was the algorithm used when SctRodDaq was being implemented. It now appears that SCTDAQ uses the third routine for S-curve fitting in most of the analyses relevant to this thesis. In this third routine, SCTDAQ uses a log-likelihood technique to fit the binomial data, thus there is no need for the error calculation in equation 4.1. It is therefore inevitable that there will be some differences (aside from the usual differences expected due to different implementations, floating point numbers and random number generators).

The response curve data for 194 modules was fitted in SctRodDaq using the high-level test system (see section 4.6.2). This amounts to 1940 threshold scans with nearly 3 million individual channel fits. All of these except for those scans affected by the 8 fC effect (see section 4.5.5) were compared with the SCTDAQ fits. Any with a difference in the mean threshold of more than 2% or a difference in the width of more than 20% were individually checked.

The comparison of the mean threshold is shown in figure 4.12 for both the normal and analytic algorithms. Note that the standard deviation for the normal algorithm is consistent with the error on the mean which is typically 0.2%.

Of the channels with a mean difference of more than 2%, the majority were due to one module which seemed to have a problem at some points in the threshold scans which affected SctRodDaq and SCTDAQ differently. This problem was not flagged by either program and has been referred to

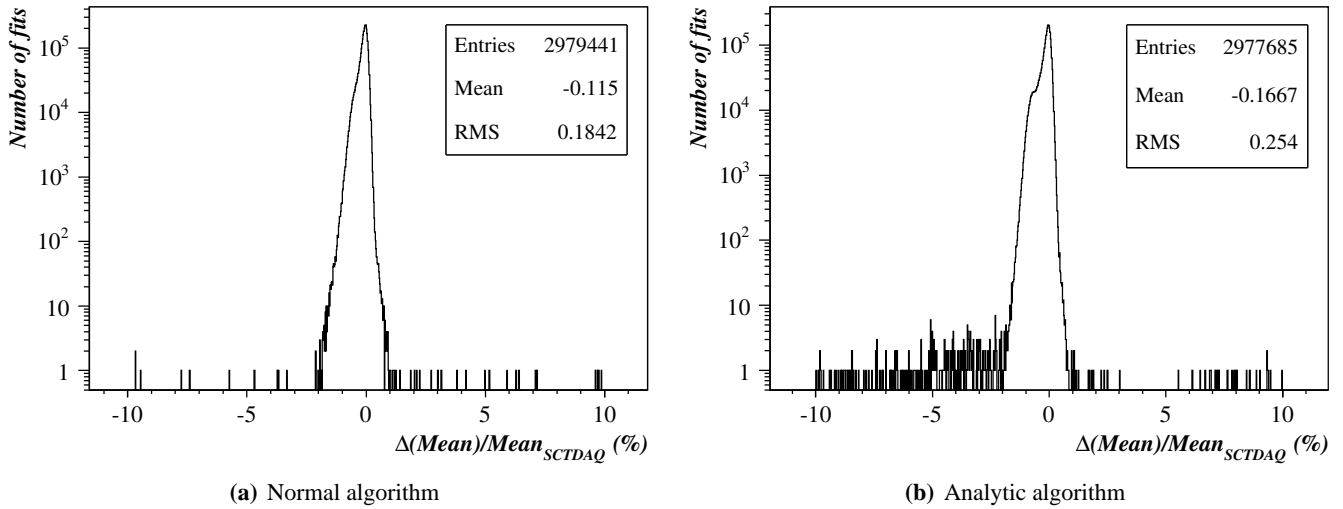


Figure 4.12: Comparison of S-curve fitting showing the fractional difference in the \bar{i}_V in percent for the two algorithms. Note the log scale.

the module experts [107]. Only three other channels had a difference in the mean greater than 2%; of these, two had slightly unusual data — several data points did not fit the error function shape and increased the error on the mean. The final channel appears to be an outlier. These results show that the mean occupancy as measured by SctRodDaq agrees with SCTDAQ to better than 2% in all but less than 1 in 10^6 cases.

Figure 4.13a shows a similar plot for relative differences in the width of the S-curve. Note that the width is much more poorly determined but the standard deviation for the normal algorithm is again consistent with the typical error of 1.5%. All channels which had a difference in the width of more than 20% were investigated and all were found to be due to data that caused the error on the fit to be higher than usual.

Figures 4.12b and 4.13b, for the analytic threshold algorithm, show that although the agreement is still very good (better than 2% for all but 1 in 10^4 channels), it is clearly not as good as for the normal fit method and has larger systematic tails. More investigation of this technique is needed, but it may be a useful method for the ROD to perform or could form a first step that is improved with the normal fitting technique if the χ^2 value is high. The analytic method is more

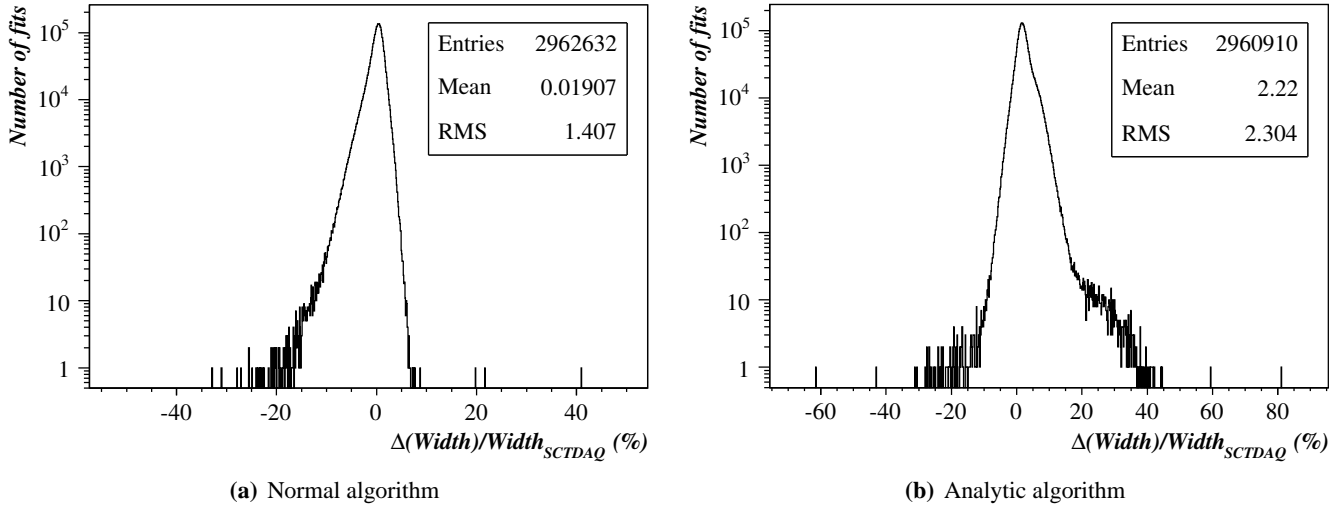


Figure 4.13: Comparison of S-curve fitting showing the fractional difference in the width, n_V , in percent for the two algorithms. Note the log scale.

than three times faster than the normal method (although it currently does not calculate a χ^2 value) and could probably be made even faster with optimization.

4.7.2 Noise occupancy

The noise occupancy algorithm was tested with data from 179 modules. The main output of the noise occupancy algorithm is the occupancy at 1.0 fC which is required to be less than 5×10^{-4} and this is compared in figure 4.14.

The main cause of differences between SCTDAQ and SctRodDaq is the way of determining the threshold voltage that corresponds to 1.0 fC. The noise occupancy test is run on trimmed modules, so the trim target will correspond to 1.0 fC. However, it is also possible to use the response curve to determine the threshold that corresponds to 1.0 fC and slight differences, of up to 2.5 mV (1 bit in the threshold discriminator DAC) can arise. SCTDAQ uses the trim target whereas SctRodDaq uses the response curve. Since the noise occupancy is steeply falling, this can give a difference in noise occupancy of up to 1×10^{-4} . This illustrates the inherent uncertainty in the noise occupancy measurement and is not considered to be a problem.




Figure 4.14: Comparison of the noise occupancy algorithm showing the difference in the mean noise occupancy at 1.0 fC between SctRodDaq and SCTDAQ.

4.7.3 Response curve

Response curve data for 187 modules has been used to test the analysis, amounting to nearly 300,000 individual fits. In general agreement was very good with SCTDAQ and SctRodDaq giving identical lists of problem channels (modulo some differences due to different versions of SCTDAQ). Comparisons of the gain and offset are presented in figure 4.15. The input noise at 2.0 fC was also compared but is strongly correlated to the gain and the 2.0 fC threshold scan fit. Although the vast majority of fits are in very good agreement, there are some clear systematic tails in figure 4.15a. Figure 4.15b shows that the high side is due to the correlation with the offset. The small low side peak is due to a difference in the handling of the 8 fC effect between SCTDAQ and SctRodDaq. It appears that SCTDAQ trims points more aggressively than SctRodDaq. The channels in that peak have not had points trimmed by SctRodDaq which leads to a higher gain since the 8 fC effect tends to increase the mean threshold of affected points. A review of these channels showed that they are borderline cases for trimming, but this could be investigated further. The extreme outliers have been investigated and were found to be due to issues during the threshold scan fitting which were covered in section 4.7.1.

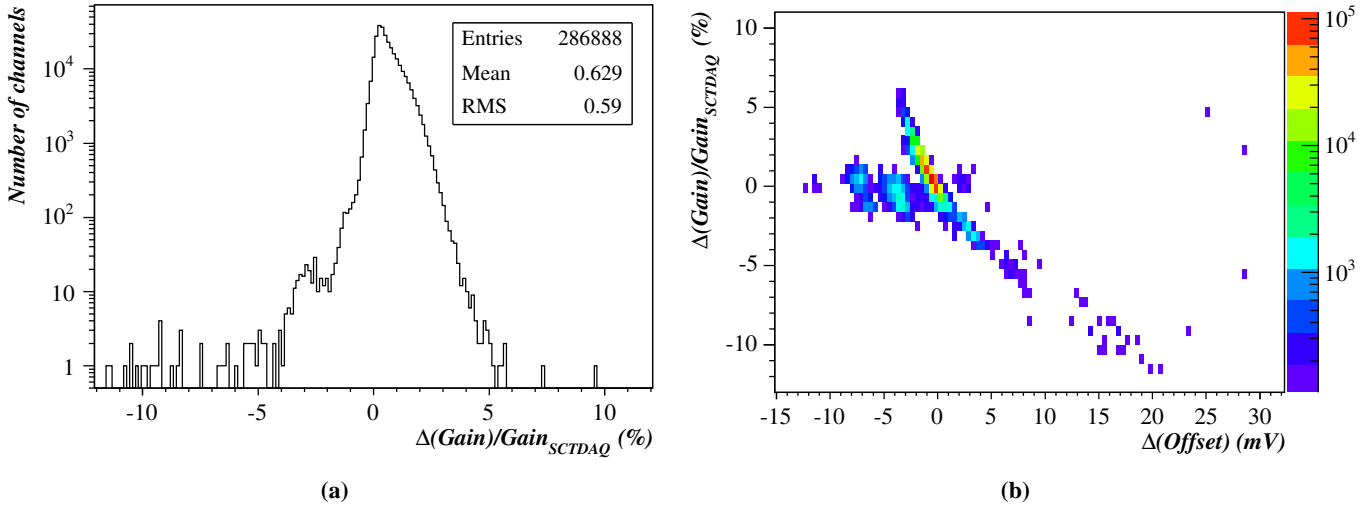


Figure 4.15: Comparison of the response curve algorithm. **(a)** shows the fractional difference in the gain between SctRodDaq and SCTDAQ. **(b)** shows the fractional difference in the gain against the difference in the offset. Note how they are correlated.

4.7.4 Strobe delay

The strobe delay algorithm has been compared to SCTDAQ using the data from 191 modules. It was found that the agreement depended strongly on the SCTDAQ version and this is shown in figure 4.16.

For most SCTDAQ versions, the maximum width of the distribution is 1. This is expected since SCTDAQ reports strobe delays to the nearest integer whereas SctRodDaq reports it to 1 decimal place^r. However it can be seen that in versions prior to 3.41, the distribution is not centred on 0, but is between 0 and 1. This is because SCTDAQ had a bug whereby it truncated the strobe delay rather than rounding it.

Version 3.42 seems to have two separate distributions, one around 0 and one around -4 . This is because, part way through generating data, the SCTDAQ version in Cambridge was modified so that the strobe delay fraction was increased from 0.25 to 0.4 (see section 4.5.3). In SctRodDaq version 3.01, this fraction has also been set to 0.4. It is currently unclear why there are so many differences with SCTDAQ version 3.41.

^rThe strobe delay must, in any case, be set to an integer, so this difference is not important.




Figure 4.16: Comparison of the strobe delay algorithm showing the difference in strobe delay against SCTDAQ version. The colour scale shows the number of chips with white indicating zero.

4.7.5 Trim calibration

The trim algorithm depends very sensitively on the details of algorithm and the parameters of the fits. Although in most cases the output of the algorithm should be the same as in SCTDAQ this will not always be true. Consider, for instance, a chip where one channel is on the border line of being trimmable in a certain trim range for a given trim target. A slight variation in the threshold scan fitting could lead to that channel being considered trimmable or untrimmable. If it is considered untrimmable, then that trim range and trim target could be rejected in favour of other trim ranges and targets which would potentially change all of the trim settings for a module. It is therefore necessary to consider the statistical properties of the trim algorithm such as the average trim range or the standard deviation of the mean threshold across a chip, σ_t (see section 4.5.4).

The original trim algorithm followed the Electrical Test Specification and considered a channel to be trimmable if a trim setting between 0 and 15 gave the required trim target. The performance of the original algorithm is compared to SCTDAQ in figure 4.17. The average trim range is higher as is the average σ_t , showing that SCTDAQ outperforms SctRodDaq.

Understanding the causes of this difference led eventually to a bug in SCTDAQ. This bug caused

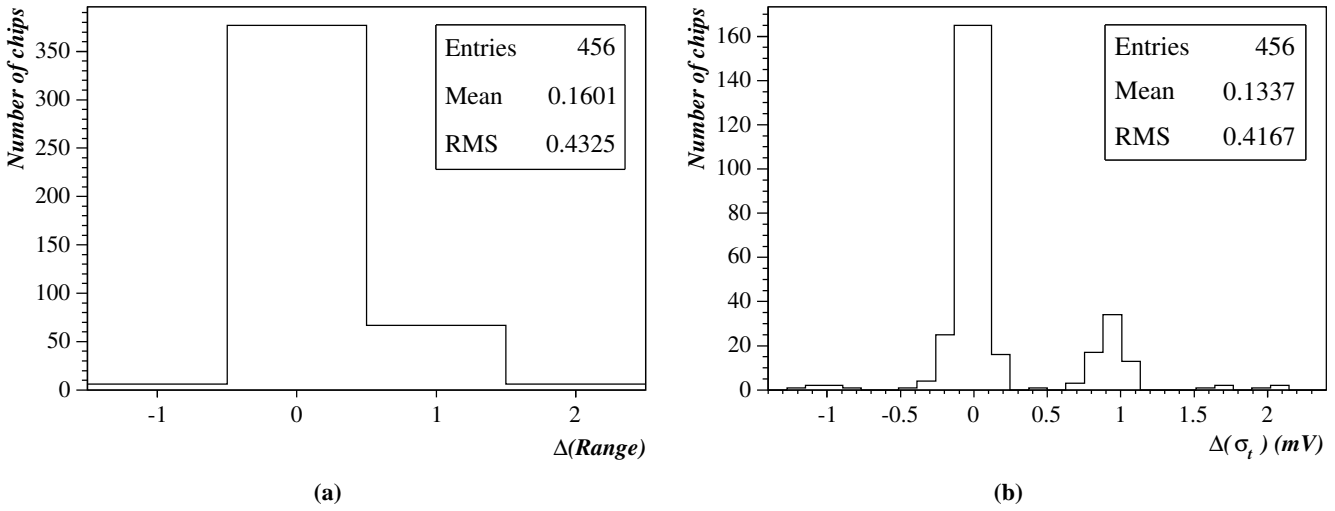


Figure 4.17: Performance of the original trim range algorithm. **(a)** The difference in the trim range between SctRodDaq and SCTDAQ and **(b)** the difference in σ_t between SctRodDaq and SCTDAQ.

channels that were only just untrimmable to be considered trimmable and have their trim set to 0. A careful consideration of why this ‘bug’ caused improved performance showed that the trim range algorithm could in fact be improved further. The reason is that often it is a single channel that just untrimmable and leads to the rejection of the lowest trim ranges. Increasing the trim range increases σ_t since the step between trim DAC values is larger. There is therefore a small untrimmable region where increasing the trim range actually reduces the performance of the algorithm.

The trim algorithm was therefore changed to consider channels trimmable if their trim setting is between -1 and 16 but it sets their trims to 0 and 15 respectively^s. This improved algorithm is compared to SCTDAQ in figure 4.18. The average trim range is decreased as is the average σ_t with no significant high tails. This change has since been implemented in SCTDAQ.

Unfortunately the sensitivity of the trim range algorithm makes a more detailed comparison with the modified algorithm difficult. Nevertheless, these studies suggest that there are no serious

^sThis actually decreases the performance for these channels slightly, but these are in the minority and it increases the performance for all the other channels that would otherwise have to be at a higher trim range.

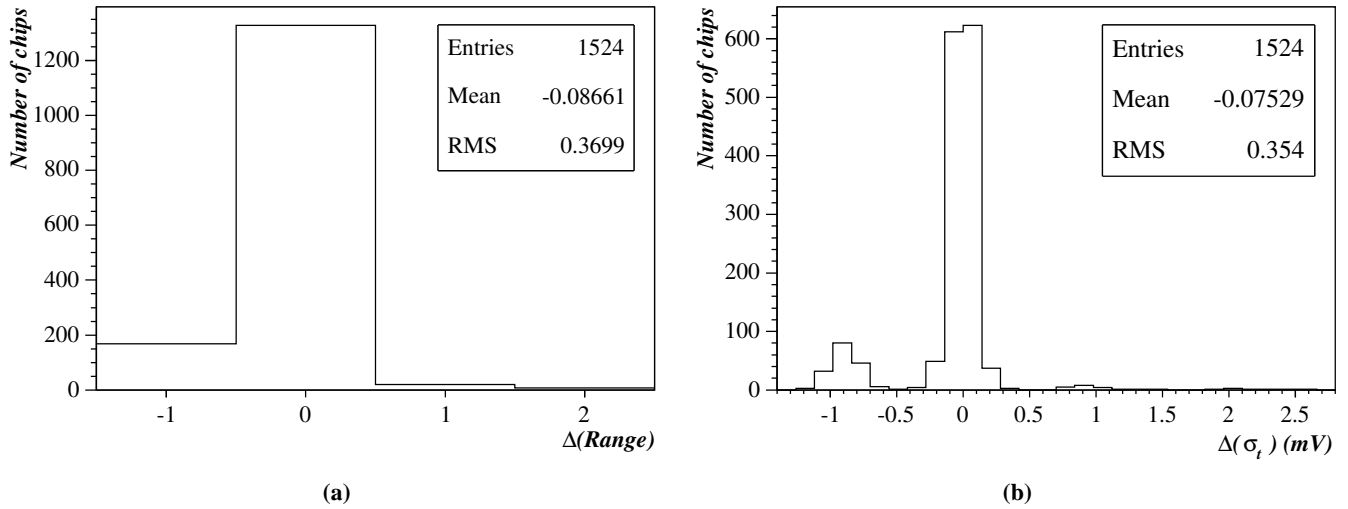


Figure 4.18: Performance of the new trim range algorithm. (a) The difference in the trim range between SctRodDaq and SCTDAQ and (b) the difference in σ_t between SctRodDaq and SCTDAQ.

bugs in the SctRodDaq implementation of the trim range algorithm.

4.7.6 Time walk

Figure 4.19 shows comparisons for the time walk algorithm, generated from data for 191 modules. The calibration factor is calculated from the difference between the leading and trailing edges of a strobe delay scan. A difference of up to 2 bits is therefore to be expected.

One source of difference between the SCTDAQ and SctRodDaq calculations of the time walk is that SCTDAQ fits a Gaussian to each strobe delay scan to find the peak. The difference in peak position between scans with 1.0 fC injected charge and 10.0 fC injected charge is then the time walk. SctRodDaq however, uses the maximum difference in the position of the trailing edge to estimate the time walk. One bug discovered during this study is that the calibration factor (which is used to convert strobe delay bits into nanoseconds) is not used to calibrate the time walk in SctRodDaq. This does not have much effect since 1 bit is approximately 1 ns, however, it probably explains the strange shape of figure 4.19b. This analysis is currently being studied further.




Figure 4.19: Time walk comparison between SCTDAQ and SctRodDaq. (a) shows the difference in the calibration factor and (b) the difference in the time walk.

4.7.7 Digital tests

The digital tests are somewhat easier to compare with SCTDAQ. Both the full bypass and pipeline tests were compared using data from 190 modules with no deviations although the process of achieving this led to several bug fixes and the discovery of some differences between SCTDAQ and the Electrical Test Specifications.

4.8 Performance studies

4.8.1 Fitting Service

The performance of the Fitting Service has been investigated. This was done by publishing response curve data from 10 modules and using the internal timing statistics of the Fitting Service to measure the amount of time it spent fitting^t. The Fitting Service and the publishing program were the only programs running on the machine and the data files were stored locally. The Fitting Service was always near to 100% CPU usage showing that it was CPU limited. Figure 4.20 shows

^tThe internal timing statistics record only the CPU time used. These results are thus only achievable if the CPU usage is 100%.

the performance (in number of threshold scans/second) as a function of the CPU speed.




Figure 4.20: CPU scaling of the Fitting Service. All CPUs were Pentium 4's.

At 3 GHz, the Fitting Service would take about 35 minutes to fit all the data from a response curve on 672 modules. This puts a lower limit on a characterisation sequence of about 2 hours which is certainly acceptable^u. All of the CPUs in figure 4.20 had a cache size of at least 512 Kb (the 3.2 GHz CPU had a 1 Mb cache). Further tests were performed with a 2 GHz CPU with 256 Kb cache. This CPU showed a 50% drop in speed. It is notable that the 1 Mb cache of the 3.2 GHz CPU did not improve its performance beyond what would be expected for its speed. It is likely that the reason for these results is that having a large cache lets the table of error function values to be permanently stored in the cache rather than having to re-fetched from main memory continually.

The Fitting Service does not show any dependence in speed with the number of modules; this is due to its design which treats each module entirely independently. However, there are aspects of system performance which can depend on the number of modules being scanned and this is discussed in section 4.8.3.

^uSCTDAQ currently takes about 3 hours to characterise four modules although this is on a slower machine.

4.8.2 Analysis Service

The performance of the Analysis Service has been investigated in a similar way to the Fitting Service. Like the Fitting Service, the Analysis Service records its own performance data which is a fairly accurate measure of its CPU usage (see section 4.8.3 for independent timing studies which include factors other than algorithm speed). The only analysis which takes a significant amount of time is the response curve since this involves fitting a non-trivial function for every channel in a module (it is therefore of a similar complexity to the Fitting Service). Of course, the performance of the Analysis Service is not as critical as the Fitting Service, since algorithms will only be run once for each module in a test.

The test setup used a 2.66 GHz Xeon for most processes including the Fitting Service. The high-level test system was used to generate response curves for 20 modules with data taken from reception tests. The Analysis Service was run on a range of systems, all with Pentium 4 CPUs and 512 Kb cache (except for the 3.2 GHz CPU which had a 1 Mb cache). All machines had at least 512 Mb of RAM. The scratch space used for transferring data via NFS was set to an independent machine. The total time to perform the analyses was obtained from the Analysis Service internal timing routine and is shown in figure 4.21 as the number of analyses per second.




Figure 4.21: CPU scaling of the Analysis Service. The data show the number of response curve analyses per second as a function of CPU speed.

There is some evidence of a slight non-linearity in the scaling, which is likely to be due to the time taken to read in all the fits — this is estimated to account for 13% of the total time. A similar test was done on a 2.0 GHz CPU with a 256 Kb cache. In this case, the slow down in performance was 20% which shows that the Analysis Service is not as susceptible to cache size as the Fitting Service but that it is still important (note also that the 3.2 GHz CPU was not notably faster).

The overall speed of analysis is somewhat below the Fitting Service. This is important in terms of system performance since the analysis can only begin when the Fitting Service has completed its work. However, since there are many modules, the Analysis Service can be analysing one module whilst waiting for the Fitting Service to fit the last threshold scan of the next module. Nevertheless, the achieved performance suggests that the Analysis Service would add about 3 minutes to the time to do a response curve for 672 modules if both the Fitting Service and Analysis Service were running on separate 3 GHz CPUs. This is less than a 10% increase and probably does not justify further optimization. In any case, the analysis for all other tests is significantly faster than the Fitting Service, so the total time added to a characterization sequence is minimal.

4.8.3 System performance

A series of system performance tests have been carried out. For these tests, the low-level test system was used to generate response curve data. Most of the processes including the GUI, IS servers, Calibration Controller and the test system were started on a 2.66 GHz Xeon system. The Fitting Service was run on a 3 GHz system and the Analysis Service on a different 3 GHz system. The scratch space, used to transfer the data files between machines, was hosted on the machine the Fitting Service was running on. This choice is optimal for system performance with the current configuration. All of the machines had no other CPU-intensive processes running on them. The Xeon system used to host the programs that were not being timed was rarely at 100% usage and only affected the performance of the system in that the time taken to publish all of the raw scan data was somewhat dependent on it — this is an insignificant effect, however.

The results are presented in figure 4.22 which show the system performance. Figure 4.22a

shows the times that important events occurred relative to the start time of the response curve. Note that the time for a single scan was always 60 seconds, but that publishing the scan data took a varying amount of time. This is why there is a slight upward trend in the ‘last raw scan data’ line. The start and finish of fitting and analysis were used to calculate the real performance of the Fitting and Analysis Services in figure 4.22b in operations per second. The Fitting Service spends much time waiting for new data at low numbers of modules and this is the reason for the apparently low performance. An attempt was made to estimate the amount of time it spent waiting and was used to produce the ‘corrected’ line. This graph corresponds to those generated in the individual tests of performance in sections 4.8.1 and 4.8.2.




Figure 4.22: System performance graphs as functions of the number of modules. **(a)** shows times relative to the beginning of the response curve test. **(b)** shows the performance of the Fitting and Analysis Services measured in operations per second (an operation is either a fit to the scan data from a module or an analysis of the data for a module); see text.

The graphs show two clear regions, with a switch between them occurring somewhere between 250 and 270 modules. For small numbers of modules, the Fitting Service is fast enough to fit all of the data from one scan before the next scan finishes. However, around 250 modules, this changes and thus, by the time the last scan is finished, the Fitting Service is still catching up with previous scans.

Figure 4.22b also shows that there is another reason for the steeper increase in time above 270 modules; which is that the Fitting Service performance decreases. It also shows that the Fitting Service performance is at least 10% less than that expected from section 4.8.1. It was also observed that the Fitting Service CPU usage started out at about 90–95% but dropped to 75–80% after a short period (with spikes to 95%). The remaining CPU was partly used by the NFS service, but typically 10–15% was idle.

This can be explained by considering disk caching strategies. Data transferred in files over NFS, such as the raw scan data, will not be immediately written to disk, but will be cached until all available memory is used up. The normal strategy is then to write out the oldest data from cache to disk, retaining the most recently accessed files. In the case of the Fitting Service, the queue will get very large, so large that the data will be larger than available memory, so the first data will be written to disk. However, the Fitting Service uses a FIFO queue so this data is the very data that it will ask for next. The file will therefore have to be read back from disk which has is slow and has high latency (important since there is only one worker thread).

This would also explain why performance is highest for 250 modules. There is a brief period at the beginning of each scan when the raw scan data is cached and the Fitting Service can use most of the CPU. However, above 250 modules, the queue is never cleared, so this happens only once, at the beginning of the test. This discovery suggests that the Fitting Service queue should be changed to a first-in-last-out (FILO) queue which could give a 10–15% performance improvement.

The performance of the Analysis Service does not show a similar trend. In fact the Analysis Service CPU usage was consistently around 70–75%. This is likely to be because it must transfer 10 fits from the Fitting Service machine (note that the NFS server is hosted on the Fitting Service machine) and this is slow with high latency. This effect could be mitigated by running with multiple threads, if that were possible, although the ‘spare’ CPU could certainly be used by, for instance, the Archiving Service, which was excluded from these tests.

These results show that for a full barrel 6, the software performance would account for more than 50% of the total time, assuming that a single scan can be done in 1 minute and all the data

for 672 modules published in about 30 seconds. However, if the scan and publish time was in fact about 4.5 minutes, then the software would contribute only about 10% of the total time which is more acceptable. Whilst these figures may suggest that the software is a large contribution to the total time, the current hardware system is significantly slower than this. Currently a single scan would take at least 6 minutes on a fully populated ROD. At this speed, the software would add at most a few minutes to the total time.

4.9 Macro assembly

SctRodDaq is in current use at three macro-assembly sites: Oxford, Liverpool and Nikhef. Assembly is progressing well with barrel 3 having been recently completed at Oxford [108, 109, 110]. It is now undergoing large scale testing (a significant test of the software too) before being shipped to CERN. A plot of the input noise of the modules is shown in figure 4.23. This data was collected with SctRodDaq during module mounting.




Figure 4.23: Noise (in ENC) versus position for barrel 3. The data was taken from mounting tests. From [111], update of older version [112].

As mentioned in the previous section, the data from Oxford shows that currently a scan takes

about 6 minutes for a fully populated ROD although ways are being investigated to improve this performance, possibly significantly. However, reading out all of the histogram data from the RODs is likely to be significant for large numbers of modules since it cannot be parallelised. Estimates currently suggest that the time taken to read out the data is about 0.3 s/module which would be about 3 minutes for 672 modules^y. Another important area for system performance is updating the module configurations based on the analysis results. This is currently very slow, ~ 1 s/module, but there are suggestions for ways to improve this to a few tenths of a second.




Figure 4.24: The completed barrel 3. From [113].

^yThere are several possibilities for improving this including compressing the data on the RODs (since occupancy data is mostly either full or no occupancy, it would compress quickly and easily) and moving the fitting onto the RODs.

4.10 Future directions

A piece of software like SctRodDaq is never finished, but undergoes constant improvement whilst it is still useful and there are certainly many areas where SctRodDaq can be improved. For instance, speed increases could come through removing ROOT and using FILO rather than FIFO queues. More importantly, new features such as new tests specifically for macro-assembly or better long-term archival could be added. For final ATLAS, features such as monitoring of physics data and continual calibration will be needed. However, the most important changes are likely to be those required to keep in step with external software. For instance, the latest versions of the ATLAS online software have switched to using the omniOrb CORBA implementation (which is fully compliant) and this implies significant changes in SctRodDaq. SctRodDaq must also move to the new calibration and conditions databases as they become available. Finally some administrative changes may be needed and this may force a move away from using NFS as the data transfer mechanism.

The success and ease with which these changes can be made will be the measure of whether SctRodDaq fulfils one of its aims: for well-written, maintainable code. Whilst it is difficult to be quantitative, it is encouraging that C. Lester, who recently joined the SctRodDaq team, was able to implement the change to omniOrb reasonably easily.

An important part of being able to make changes in the future is how well the code is documented, and indeed, good documentation was an aim. Several sources of documentation for SctRodDaq are available: the code is reasonably well documented with the doxygen tool [114] — about 20% of non-blank lines are comments although this is somewhat off the generally accepted aim of one third. There is also the SctRodDaq Wiki [82] which contains a large amount of collected knowledge about the software, the beginnings of a manual for the GUI [115], B.J. Gallop's thesis [83], which has a description of SctApi and the Configuration Service, and of course, this chapter. However, it is certainly the case that some parts of the code are not as well documented as others, and that more could be done to improve this.

4.11 Summary

SctRodDaq, a system for data-acquisition and calibration of SCT modules during macro-assembly has been developed. It has been shown in this chapter that the analysis sub-system is fit-for-purpose by demonstrating that the algorithms are correctly implemented and that it is fast enough to handle the large quantities of data that it must deal with. It is currently in use at three macro-assembly sites where it is aiding module mounting and there are no reports of problems with the software. It will be further developed in the future to become a system for integration of the whole SCT and finally for use during assembly of ATLAS and data taking.

SctRodDaq has certainly satisfied its functional aims — it has provided a stable, robust system for macro-assembly. The extent to which it has satisfied the aim of being a good quality, maintainable, well-documented piece of software is harder to judge, however there are encouraging signs that it has at least adequately fulfilled this aim and that SctRodDaq will continue to be a useful tool throughout the lifetime of ATLAS. The next chapters will consider just how ATLAS can be used if certain extra dimensional models turn out to describe nature.

Studies of Graviton Decays to Heavy Vector Bosons

Oh, let me have just a little bit of peril?

5.1 Background

The Randall-Sundrum (RS) model was introduced in section 1.2.3. One of its key predictions is the existence of a tower of widely spaced, narrow massive graviton resonances. The lowest energy resonance should be accessible to ATLAS if the RS model is to solve the hierarchy problem. Since the second and higher resonances are not normally accessible, the lowest mass resonance is referred to as the graviton. Its mass is given by $m_G = x_1 \Lambda_\pi k / M_{\text{PL}}$ where $x_1 \approx 3.8$, the first root of the Bessel function. Recall from section 1.2.3 that k/M_{PL} and Λ_π are parameters of the model with k/M_{PL} expected to be between 0.01 and 0.1 and $\Lambda_\pi \lesssim 10$ TeV.

An earlier study [46] has shown that the graviton could be discovered in the $G \rightarrow e^+e^-$ channel if its mass is less than 2 TeV for the most pessimistic case of $k/M_{\text{PL}} = 0.01$ (note that the cross section rises with $(k/M_{\text{PL}})^2$). One key property of the graviton is its universal coupling — this is quite different from other massive neutral resonances such as heavy Higgs and Z' bosons. This study aims to determine how well the coupling of the graviton to heavy vector bosons could be measured.

The $G \rightarrow W^+W^-$ and $G \rightarrow Z^0Z^0$ channels are considerably more challenging than the discovery $G \rightarrow e^+e^-$ channel. Not only is the branching ratio smaller, but the reconstruction is harder and there are sizable Standard Model backgrounds to consider. In both these cases, the

semi-leptonic decay modes have been considered, as the lepton(s) are useful for triggering and background rejection whilst retaining a reasonable branching fraction. The signal channels are therefore $G \rightarrow W^+W^- \rightarrow \ell\nu_\ell jj$ and $G \rightarrow Z^0Z^0 \rightarrow \ell\ell jj$ where ℓ stands for either an electron or muon (not tau) and j represents a jet. This convention is followed throughout this chapter.

5.2 Event generation

All events were generated using HERWIG 6.301 [64, 116] and passed through the ATLAS fast simulation software, ATLFast 2.60 [117]. ATLFast default settings were used except that the high luminosity option was turned on and the jet finder algorithm used was Mulguisin with $\Delta R = 0.2$ (see section 5.3.2). As ATLFast does not include electron or muon identification efficiencies, these were added afterwards. The electron efficiency was set to 90% and the muon efficiency 85% based on the expected values [47, sections 7–8]. It should be noted that HERWIG does not include interference terms with Standard Model processes. It is therefore only valid if the experimental resolution is larger than the width of the graviton which is the case in this study.

5.2.1 Signals

Signal samples at a selection of graviton masses in the range 500–3000 GeV were generated with $k/M_{\text{PL}} = 0.05$ (0.01 at 500 GeV). Data corresponding to an integrated luminosity of about 100 fb^{-1} or a minimum of 20,000 events were generated at each point. Table 5.1 gives details about the $G \rightarrow W^+W^-$ and $G \rightarrow Z^0Z^0$ signal samples. In the $G \rightarrow Z^0Z^0$ case, one Z^0 was forced to decay leptonically; this has been accounted for in the cross sections reported, which were based on the HERWIG output.

The $G \rightarrow W^+W^- \rightarrow \ell\nu_\ell jj$ signal is characterised by:

- A high p_T lepton in the central region.
- Missing transverse energy due to the neutrino.
- Two high p_T central jets with small separation.

Mass (GeV)	k/M_{PL}	$G \rightarrow W^+W^-$		$G \rightarrow Z^0Z^0 (Z^0 \rightarrow \ell\ell)$	
		Cross section (fb)	Events	Cross section (fb)	Events
500	0.01	975	120,000	66.7	60,000
800	0.05	2730	100,000	182	160,000
1000	0.05	917	100,000	61.3	50,000
1300	0.05	240	40,000	16.2	20,000
1500	0.05	111	20,000	7.41	20,000
1650	0.05	65.5	20,000	4.38	20,000
1800	0.05	39.8	20,000	2.64	20,000
2000	0.05	24.1	20,000	1.42	20,000
2200	0.05	11.9	20,000	0.790	20,000
2500	0.05	5.20	20,000	0.346	20,000
2750	0.05	2.63	20,000	0.181	20,000
3000	0.05	1.47	20,000	0.00969	20,000

Table 5.1: Details of the signal samples for the W and Z channels showing the model parameters, cross section and number of events generated for each data sample. Note that one of the Z^0 bosons was required to decay leptonically — this is included in the cross sections. All numbers are given to 3 s.f.

This signal is experimentally similar in some respects to heavy Higgs decays which have been studied in some detail in ATLAS [118, 119]. There are however significant differences; the production mechanism for these heavy Higgs particles is normally via vector boson fusion which allows the possibility of using very high $|\eta|$ jets to tag the events and of vetoing on hadronic activity when the tag jets and the central hard process [120]. The production mechanism for the gravitons considered here is quark or gluon fusion, so these techniques cannot be used. However, the graviton signal does have the advantage of being a very narrow resonance (unlike a heavy Higgs) which will aid the separation of the signal and background.

The $G \rightarrow Z^0Z^0 \rightarrow \ell\ell jj$ signal is similar except that there are two central leptons and no missing energy. This means that the leptonic Z^0 is easier to reconstruct and the mass resolution

will be better, but the Z^0Z^0 channel is hampered by lower statistics.

5.2.2 Backgrounds

The main Standard Model backgrounds for the $G \rightarrow W^+W^- \rightarrow \ell\nu_\ell jj$ signal are:

- $t\bar{t} \rightarrow W^+W^- b\bar{b} \rightarrow \ell\nu_\ell jj b\bar{b}$.
- W+2 jets where the W decays leptonically and the two jets fake a W.
- Standard Model W^+W^- production.

Of these, the most important are the reducible backgrounds: $t\bar{t}$ and W+2 jets, as both of these will have very high cross sections at the LHC. The irreducible Standard Model W^+W^- background, by contrast, is negligible, but has been included for completeness.

The backgrounds for the $G \rightarrow Z^0Z^0 \rightarrow \ell\ell jj$ signal are similar except that there is no equivalent of the $t\bar{t}$ background. The dominant background is therefore Z+2 jets production; again Standard Model Z^0Z^0 production is negligible.

Approximately 100 fb^{-1} of data was generated for each of the backgrounds where this was possible. The W/Z+2 jets backgrounds were generated using the HERWIG W/Z+jet processes with the second jet coming from initial or final state radiation. It is known that these processes underestimate the cross section for multi-jet final states, so the cross sections were multiplied by a factor

$$1.6 \times \left[\frac{m_W^2 + (p_T^{thr})^2}{m_W^2} \right]^{2(N_{\text{jet}}-1)} \quad (5.1)$$

where N_{jet} is the number of jets (2 in this case), and p_T^{thr} is the jet transverse energy threshold which was set to 15 GeV. This correction brings the HERWIG cross section into better agreement with tree-level matrix element calculations [121].

The W+jets sample was split into two with different p_T ranges. The high p_T sample was generated with all decays, but in the low p_T sample the W was forced to decay to an electron or muon to reduce the number of events that needed to be generated. Even so, it was still not possible

to generate the full 100 fb^{-1} of data. In principle, the lack of $W \rightarrow \tau \nu_\tau$ could bias the result as $\tau \rightarrow \ell \nu_\ell \nu_\tau$ is an additional background. To compensate for this, the scale factor was increased by the branching ratio for this decay (17%). It was confirmed in a small sample that this was the increase in acceptance between $W \rightarrow \ell \nu_\ell$ and $W \rightarrow \text{all}$. It should be noted that this only has an impact on signals below $\sim 800 \text{ GeV}$.

The $t\bar{t}$ cross section was increased by a factor 1.84 to bring the HERWIG cross section into agreement with the NLL cross section calculated in [122]. Table 5.2 gives further details on the background samples generated.

Background	p_T range (GeV)	Cross section (pb)	Scale factor	Events ($\times 10^6$)
$t\bar{t}$	> 80	246.4	1.84	50.5
W+jets	> 200	125.3	1.71	24.5
W+jets, ($W \rightarrow \ell \nu_\ell$)	$80 < p_T < 200$	530.4	2.016	24.5
W^+W^-	> 80	9.34	1	1.0
Z+jets	> 80	65	1.71	14.8
Z^0Z^0	> 100	0.1616	1	0.25

Table 5.2: Background sample information showing the p_T range which indicates the p_T cut placed on the hard process at generator level, cross sections as reported by HERWIG, the scale factor (see text) and the number of events generated.

In total over 150 million events were generated. This was made possible by using the UK e-Science grid and in particular machines at Cambridge, Imperial College London, RAL and Southampton. For further information see [123].

5.3 Reconstruction

5.3.1 $W \rightarrow \ell \nu_\ell$ reconstruction

In the $G \rightarrow W^+W^-$ case, the leptonic W needs to be reconstructed from the lepton and \cancel{p}_T . This is done by using the W mass as a constraint and gives a quadratic for the z component of the

Studies of Graviton Decays to Heavy Vector Bosons

neutrino's momentum, p_z^ν . This constraint effectively fixes the angle between the neutrino and the lepton but this, of course, leads to a degeneracy in which side of the lepton the neutrino is. Since the W is highly boosted, the discriminant in the quadratic is typically close to zero and in about 30% of cases is measured to be negative. To increase the signal efficiency, these cases were allowed, with the discriminant set to zero.

For the cases where there are two solutions to the quadratic, one must be chosen (or both used with the events given a weight of 0.5). For the method of choosing the solution randomly, the graviton mass resolution is shown in figure 5.1a which has a characteristic shape from two Gaussians — one narrow one corresponding to choosing the correct solution and a much wider one from the wrong solution. For this study, an alternative method was used that reduces the size of the tails. p_z^ν was set to the average of the two solutions which corresponds to choosing the pseudorapidity of the neutrino to be the same as that of the lepton. This is illustrated in figure 5.1b. Note that although the Gaussian is slightly wider, the tails are much smaller. This is slightly beneficial when the background is taken into account, although the difference is small.




Figure 5.1: Comparison of graviton mass resolution for two methods of setting p_z^ν with $m_G = 1$ TeV. In (a) p_z^ν was randomly chosen to be one of the two solutions whereas in (b) p_z^ν was set to the average of both. See text.

5.3.2 W/Z → jj reconstruction

Reconstructing the hadronic W or Z is made difficult by the high boost it typically receives from the very massive graviton. This makes the separation between the two jets very small. For instance, the jet separation, $\Delta R = \sqrt{\Delta\phi^2 + \Delta\eta^2}$, peaks at $\Delta R \approx 0.35$ for a 1 TeV graviton and at $\Delta R \approx 0.2$ for a 2 TeV graviton. ATLFast models the whole calorimeter as a grid of 0.1×0.1 cells in the central region, so the smallest average separation that can be detected is $\Delta R = 0.2$. In practise, it may be possible to do rather better than this — recall that the ECAL geometry has a higher granularity in some regions (see section 2.2.3 and figure 2.7) and tracking information could also be useful. Studies of these possibilities would need full simulation.

Reconstructing highly boosted W or Z bosons remains a topic of interest [11, 124, 125]. The technique used here was to use the Mulguisin algorithm [124] with $\Delta R = 0.2$. This gave the highest reconstruction efficiency of the available methods (cone, k_T and Mulguisin) [126].

The W is reconstructed by using the highest p_T jet and searching for another jet that gives a mass in the range $65 < M_{jj} < 85$ GeV. The reconstructed W mass is shown in figure 5.2 for two example signals. Note that the peak is somewhat below the true W mass; this is probably due to the small ΔR parameter used in the jet reconstruction — some energy will be lost outside of the jets. Analysis of boosted hadronic W decays could be used to re-calibrate the jet energy scale and to correct for this. In the Z → jj case, the mass range was $75 < M_{jj} < 100$ GeV.

At very high graviton masses, resolving the two jets becomes extremely difficult and the reconstruction efficiency decreases. A number of techniques have been discussed for improving reconstruction in such cases, such as using sub-jet or heavy jet techniques [118, 119]. It is possible that such techniques would be helpful at high mass where the Standard Model background is also much smaller (reconstruction of the hadronic W/Z is an important rejection of the W/Z+jets backgrounds). A full study of this would require full simulation.

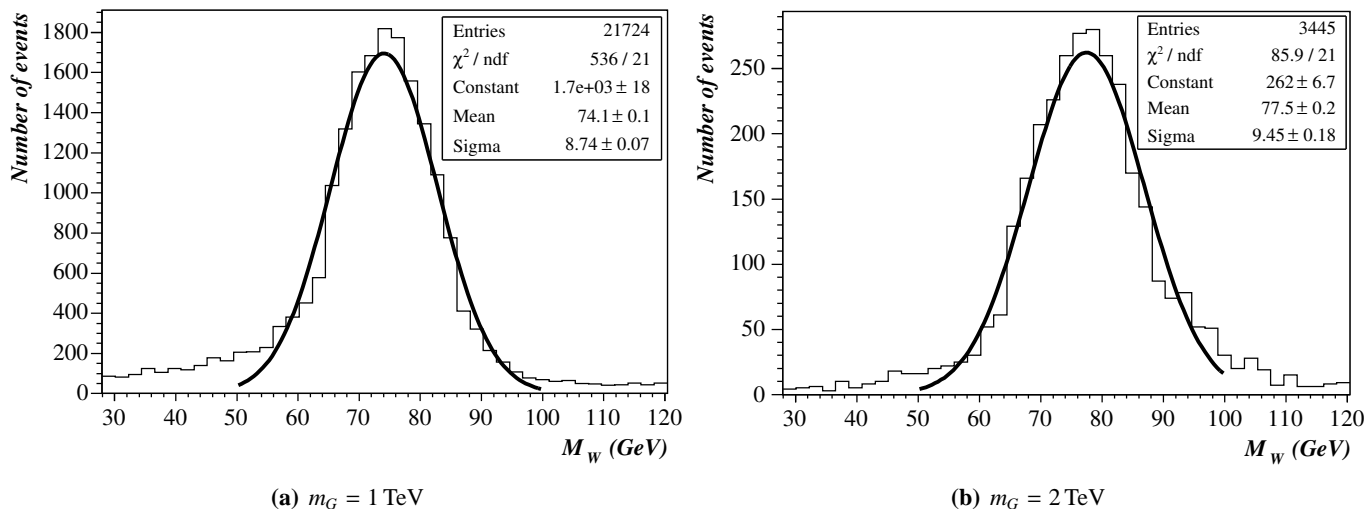


Figure 5.2: Reconstructed $W \rightarrow jj$ mass for (a) $m_G = 1$ TeV and (b) $m_G = 2$ TeV.

5.3.3 Graviton reconstruction

The graviton is trivially reconstructed from the W^+W^- or Z^0Z^0 pair. The mass resolution is shown in figure 5.3 for the $G \rightarrow W^+W^-$ case and figure 5.4 for the $G \rightarrow Z^0Z^0$ case. The mass resolution is about 6% for $G \rightarrow W^+W^-$ at high mass rising to 9% for $m_G = 500$ GeV. The resolution is a constant 3% for $G \rightarrow Z^0Z^0$ due to the presence of two leptons. In all cases, the mass is slightly underestimated by up to 10 GeV — this effect was evident in the W and Z mass reconstruction (figure 5.2).

5.4 Cuts

The aim of the cuts is not to simply to maximise S/\sqrt{B} — discovery of the graviton is better achieved in other channels, rather the aim is to measure the couplings as accurately as possible. The cuts should therefore reduce the large Standard Model backgrounds as much as possible whilst leaving a smooth shape under the signal. This is necessary so that the backgrounds can be fitted in the sidebands and subtracted. In particular, it is possible to increase S/\sqrt{B} by increasing the cuts (particularly the p_T cuts) in the following sections, but this causes the background to peak under




Figure 5.3: Graviton mass resolution in the $G \rightarrow W^+W^- \rightarrow \ell\nu_\ell jj$ channel.

the signal.

5.4.1 Mass and p_T cuts

The graviton signals will produce high p_T decay products since their mass is large. The Standard Model backgrounds by contrast typically produce particles with low p_T . The cuts applied in the




Figure 5.4: Graviton mass resolution in the $G \rightarrow Z^0Z^0 \rightarrow \ell\ell jj$ channel.

Studies of Graviton Decays to Heavy Vector Bosons

$G \rightarrow W^+W^-$ case were:

- Lepton $p_T > 40$ GeV.
- $\cancel{p}_T > 40$ GeV.
- Jet $p_T > 40$ GeV for both jets.
- W $p_T > 80$ GeV for both Ws.
- $65 < M_{jj} < 85$ GeV.

The effects of these cuts are summarised in table 5.3. The most powerful are the cuts on the jet and lepton p_T s and on \cancel{p}_T . The W p_T cuts have been omitted from the table as they have only a small effect (they are effectively implied by the other p_T cuts). The cut on the reconstructed W mass, M_{jj} is very important in reducing the backgrounds.

The cuts for the $G \rightarrow Z^0Z^0$ case are somewhat similar:

- Lepton $p_T > 50$ GeV for the highest p_T lepton.
- Jet $p_T > 50$ GeV for both jets.
- Z $p_T > 100$ GeV for both Zs.
- $85 < M_{\ell\ell} < 100$ GeV.
- $75 < M_{jj} < 100$ GeV.

The p_T cuts can be slightly higher as the higher resolution allows the peak position to be moved up. The effects of these are shown in table 5.4. Again, the M_{jj} cut is important in reducing the Z+jets background.

It is possible to slightly improve the final results by tuning these cuts for a given graviton mass. If this is done, care has to be taken to ensure that the background does not peak under the signal and such an optimisation would be somewhat sensitive to the details of how the background subtraction is carried out.

5.4 Cuts

Cuts	Signals, m_G				Backgrounds			
	0.5 TeV	1 TeV	1.5 TeV	2 TeV	$t\bar{t}$	W+jets	W+jets, low p_T	W ⁺ W ⁻
Events with 1 lepton and ≥ 2 jets	34528	31313	6256	6057	15.3×10^6	3.85×10^6	10.3×10^6	211×10^3
Lepton $p_T > 40$ GeV	73%	85%	89%	92%	55%	75%	67%	66%
$\cancel{p}_T > 40$ GeV	55%	75%	84%	87%	39%	60%	33%	39%
Jet $p_T > 40$ GeV	39%	62%	71%	72%	34%	39%	13%	17%
$65 < M_{jj} < 85$ GeV	13%	29%	33%	26%	5.0%	3.6%	0.56%	7.2%
Total passing p_T cuts	4331	9165	2093	1592	763×10^3	139×10^3	57.7×10^3	15.1×10^3
Equiv. for 100 fb^{-1}	3532	8426	1165	170	685×10^3	122×10^3	252×10^3	14.1×10^3

Table 5.3: Effects of the cuts for the $G \rightarrow W^+W^-$ channel, given as a percentage of the events passing the selection criteria for signals at various graviton masses and the main backgrounds.

Cuts	Signals, m_G				Backgrounds	
	0.5 TeV	1 TeV	1.5 TeV	2 TeV	Z+jets	Z^0Z^0
Events with 2 leptons and ≥ 2 jets	30476	26167	8847	7776	8.0×10^6	141×10^3
Lepton $p_T > 50$ GeV	98%	100%	100%	100%	61%	64%
$85 < M_{\ell\ell} < 100$ GeV	89%	89%	87%	84%	51%	54%
Jet $p_T > 50$ GeV	49%	64%	66%	64%	14%	16%
$75 < M_{jj} < 100$ GeV	15%	31%	35%	31%	0.52%	7.4%
Total passing p_T cuts	4608	8003	3077	2379	41708	10496
Equiv. for 100 fb^{-1}	506	982	114	17	31338	687

Table 5.4: Effects of the cuts for the $G \rightarrow Z^0Z^0$ channel, given as a percentage of the events passing the selection criteria for signals at various graviton masses and the main backgrounds.

After these cuts, the backgrounds are still large. Ways of further reducing the $t\bar{t}$ background will be discussed shortly. Further reduction of the W/Z +jets backgrounds is difficult however. In [118], an attempt was made to use the jet energy profiles but, although there was a difference between their signal and backgrounds, once they had applied their other cuts, this difference disappeared. This technique would require full simulation for it to be applied in this case and it is not clear how well the jet energy profiles could be measured for such highly boosted jets.

5.4.2 Top reconstruction and jet veto

After the p_T cuts the Standard Model backgrounds for $G \rightarrow W^+W^-$ are still extremely large. The $t\bar{t}$ background was further reduced by the use of two cuts:

- An attempt was made to reconstruct a top quark. This was done by looping over all the remaining jets in the event and forming the Wj invariant mass for each W . If $150 < M_{Wj} < 200$ GeV then the event was assumed to contain a top quark and was rejected. b -tagging was not used for this as it could only reduce the power of this cut.
- A cut on the number of central jets is made. If there are more than four (including those used to reconstruct the W) with p_T greater than 15 GeV then the event is rejected.

The effect of these cuts is shown in table 5.5. As can be seen, they greatly reduce the $t\bar{t}$ background although at the cost of some reduction in signal efficiency. In the $G \rightarrow Z^0Z^0$ case there is no equivalent of the $t\bar{t}$ background so these cuts were not employed.

5.4.3 Summary

The overall signal reconstruction efficiencies are shown in figure 5.5. As can be seen, the efficiency peaks at $m_G \approx 1500$ GeV. At lower masses the efficiency is impacted by the p_T cuts whereas at high m_G the reconstruction efficiency is dominated by the ability to accurately reconstruct two close jets. It is possible that some improvement in signal reconstruction could be achieved at high m_G by relaxing the M_{jj} cut, however this is likely to increase the signal resolution.

Cuts	Signals, m_G				Backgrounds			
	0.5 TeV	1 TeV	1.5 TeV	2 TeV	$t\bar{t}$	W+jets	W+jets, low p_T	W+W ⁻
Events passing p_T cuts	4331	9165	2093	1592	793×10^3	139×10^3	57.7×10^3	15.1×10^3
No top quarks	94%	87%	85%	83%	66%	88%	99%	99%
< 5 central jets	81%	68%	65%	60%	34%	71%	92%	96%
Total passing top cuts	3522	6207	1352	950	270×10^3	99.1×10^3	53.2×10^3	14.5×10^3
Equiv. for 100 fb^{-1}	2872	5707	752	101	242×10^3	86.9×10^3	232×10^3	13.5×10^3

Table 5.5: The efficacy of the top cuts, given as a percentage of the events passing the mass and p_T cuts (top line).




Figure 5.5: The signal reconstruction efficiency measured as a fraction of the estimated number of signal events (see table 5.1) and counted in a window of $m_G \pm 2\sigma$.

Studies of Graviton Decays to Heavy Vector Bosons

The backgrounds are rejected by an average factor of 304 however, because of their extremely large cross section (in total 1747 pb^{-1} including scale factors), they still dominate the signal. The shape of the backgrounds after all the cuts is shown in figure 5.6 together with an example signal. As can be seen, both backgrounds have a smooth, steeply falling shape. The method for estimating and removing the background under the signal peak is dealt with in the following section.




Figure 5.6: Background shapes after all cuts for (a) the $G \rightarrow W^+W^- \rightarrow \ell \nu_\ell jj$ channel and (b) the $G \rightarrow Z^0Z^0 \rightarrow \ell \ell jj$ channel. A sample signal with $m_G = 800 \text{ GeV}$ and $k/M_{\text{PL}} = 0.05$ is also shown.

5.5 Background subtraction

In this section, it is assumed that the graviton mass is known which would come from the discovery channel, $G \rightarrow e^+e^-$. It is also assumed that the signal resolution at the graviton mass is known which would be determined by studies like this. It is therefore not necessary to establish a high signal to noise ratio before attempting to measure the coupling.

One of the constraints on the p_T cuts discussed in section 5.4.1 was to leave a smooth background under the signal peak so that the background could be estimated from the sidebands. The method used to do this is illustrated in figure 5.7.




Figure 5.7: Background subtraction schematic. The background fit is made over regions 1 and 3 simultaneously. This is then extrapolated into region 2 and subtracted to leave the signal. The p_T cuts were set to ensure that region 1 did not extend below the peak position + 100 GeV.

A fit was made to the background using one of several functions (power law or exponential — the choice did not affect the results but was used to estimate the systematic error on the fit). The fit region was taken to be $m_G \pm 6\sigma$ with the central $m_G \pm 2\sigma$ region excluded. The fit function was then used to subtract the background from under the signal peak. An example of this is shown in figure 5.8 for $m_G = 1.5$ TeV and $k/M_{\text{PL}} = 0.05$.

5.6 Results

The branching ratio \times cross section, $\sigma.B$, can be measured by counting the number of signal events if the acceptance and luminosity are known. In this study, $\sigma.B$ is not actually calculated, rather the error on such a measurement is estimated. These errors are then plotted in terms of the model parameters, Λ_π and m_G to allow easy comparison with other models.




Figure 5.8: Background subtraction example for $m_G = 1.5$ TeV and $k/M_{\text{PL}} = 0.05$. **(a)** shows the fit to the background. The arrows indicate where the central excluded region is and the fit is also dashed in this region. **(b)** shows the result of the subtraction with a Gaussian fit.

5.6.1 Determination of errors

The estimated number of signal events, \widehat{N}_S , is given by

$$\widehat{N}_S = N_S + N_B - \widehat{N}_B \quad (5.2)$$

where N_S and N_B are the number of signal and background events and \widehat{N}_B is the estimated number of background events from the background subtraction. The error on the number of signal events is therefore

$$\Delta \widehat{N}_S^2 = \Delta N_S^2 + \Delta N_B^2 + \Delta \widehat{N}_B^2. \quad (5.3)$$

Since N_S and N_B will be Poisson distributed, $\Delta N_S^2 = N_S$ and $\Delta N_B^2 = N_B$. $\Delta \widehat{N}_B$ is more difficult, but was estimated in two ways. Three different functions were fitted to the background and the differences between them were used to estimate the error. This was compared with the value $1/\sqrt{N_{\text{fit}}}$ which should give the fractional error on \widehat{N}_B where N_{fit} is the number of events in the fit region. These methods were in broad agreement giving errors on \widehat{N}_B between 3 and 6%. For the

final plots, the first method was used; typically this error is small compared to the others.

This calculation gives the statistical error. However systematic errors will arise from the acceptance, luminosity and the various reconstruction efficiencies. The luminosity is expected to be measured to 5–10% using conventional methods [47] although methods are being investigated to reduce this to 1–2% [127, 128]. The errors on the acceptance and reconstruction efficiencies should be small but will depend on the detector and reconstruction code and are thus beyond the scope of this study. These systematic errors are therefore very uncertain and have not been included, in any case, the two channels studied here are generally statistically limited.

5.6.2 Contours of $\sigma.B$

The equations in the previous section were inverted to give the number of signal events required for a given statistical error. This was calculated for each signal and compared to the actual number of signal events and from this the theoretical parameters were determined (recall that the cross section scales with $(k/M_{\text{PL}})^2$). This technique is valid so long as the width of the graviton is less than the experimental resolution, which is always the case in this study. The results of this method are shown in figures 5.9 and 5.10. The contours have been extended into the theoretically disfavoured region above $\Lambda_\pi \sim 10$ TeV as it may be applicable for other models.

At lower values of m_G the measurement of $\sigma.B$ would be entirely dominated by the background. This is why the $G \rightarrow Z^0 Z^0$ channel outperforms the $G \rightarrow W^+ W^-$ channel — the backgrounds are lower there. In particular it is interesting to note that the reach in Λ_π decreases below $m_G = 800$ GeV in the $G \rightarrow W^+ W^-$ channel due to the steeply rising backgrounds. At higher masses the measurement would be dominated by signal reconstruction efficiencies. The comparatively low branching ratio explains why the $G \rightarrow Z^0 Z^0$ channel drops off faster than the $G \rightarrow W^+ W^-$ one.

In the basic RS scenario, $\Lambda_\pi \gtrsim 10$ TeV is disfavoured so that the hierarchy problem is solved. Much of the allowed parameter space is therefore covered by these results. A measurement of the $G \rightarrow W^+ W^-$ coupling could be made for $\Lambda_\pi < 10$ TeV up to $m_G = 1.8$ TeV and a measurement of




Figure 5.9: Contours showing the statistical precision in $\sigma.B$ for the $G \rightarrow W^+W^- \rightarrow \ell\nu_\ell jj$ channel and an integrated luminosity of 100 fb^{-1} . The brown region is excluded by the Tevatron (see section 1.2.3.1) and the blue region indicates where $k/M_{\text{Pl}} > 0.1$. Lines (dashed) of constant k/M_{Pl} are also shown.

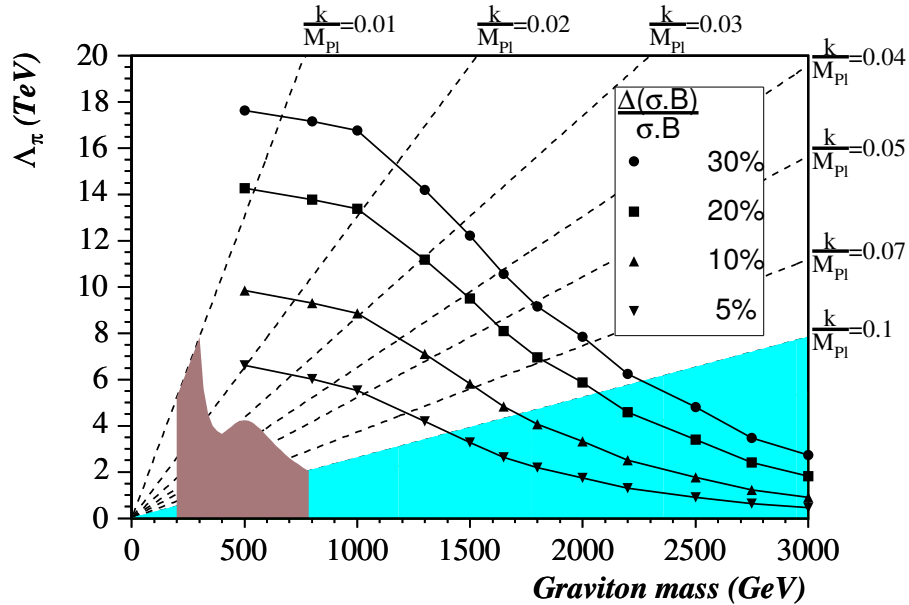


Figure 5.10: Contours showing the statistical precision in $\sigma.B$ for the $G \rightarrow Z^0 Z^0 \rightarrow \ell\ell jj$ channel. As for figure 5.9.

the $G \rightarrow Z^0 Z^0$ coupling up to $m_G = 1.6$ TeV.

5.7 Wider context

This study was not conducted in isolation. M. A. Parker and A. Sabetfakhri have studied the $G \rightarrow \mu^+ \mu^-$, $G \rightarrow \gamma\gamma$ and $G \rightarrow jj$ channels building on the earlier work in [46] which investigated the $G \rightarrow e^+ e^-$ discovery channel. Their results are presented in figure 5.11.

Taken together these results show that the graviton's couplings could be measured in a wide range of channels for a significant part of the allowed parameter space — this would then be an excellent test of the universal coupling that is a model-independent feature of the graviton and should allow rejection of heavy Higgs and Z' hypotheses which do not have such universal couplings. A measurement of the $G \rightarrow Z^0 Z^0$ coupling would be particularly useful in rejecting a pseudo-scalar hypothesis (such as an unmixed Z') since such a coupling violates parity conservation.

5.8 Conclusions

- The $G \rightarrow W^+ W^-$ and $G \rightarrow Z^0 Z^0$ couplings have been shown to be measurable to 30% or better for $m_G < 1.8$ TeV in the $G \rightarrow W^+ W^-$ case and $m_G < 1.6$ TeV in the $G \rightarrow Z^0 Z^0$ case for $\Lambda_\pi < 10$ TeV. This covers much of the allowed parameter space for the original RS model.
- The reach at high m_G is limited by the reconstruction efficiency due to the inability to distinguish the jet pairs.
- Taken with the couplings in other channels, these would provide confirmation of the universal nature of the couplings and would help rule out other potential candidates *e.g.* Z' and heavy Higgs bosons.

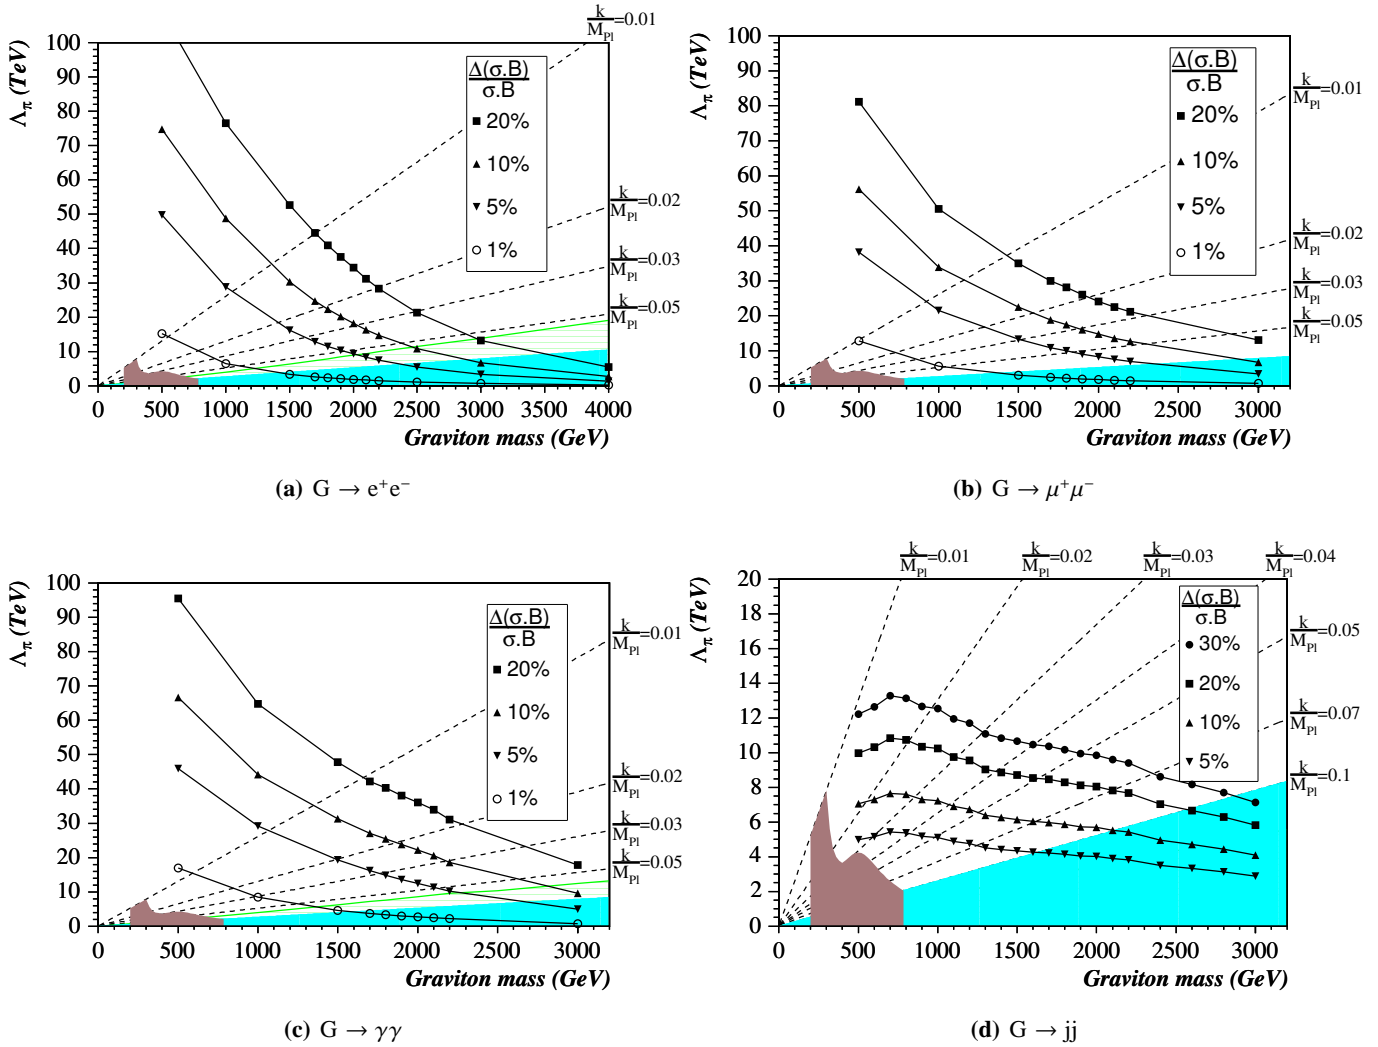


Figure 5.11: Contours showing the statistical precision other graviton decay channels. From [1] with updated Tevatron limits. As for figure 5.9. The green striped region indicates where the graviton width is greater than the experimental resolution and thus where the analysis would need modification.

Exploring Higher Dimensional Black Holes

*It's not a question of where he grips it! It's a simple question of weight ratios!
A five ounce bird could not carry a one pound coconut.*

One intriguing possibility of the extra-dimensional scenarios presented in section 1.2 is that if they are realised, then it may be possible to create quantum-scale black holes at the LHC. If they are created then observing them and measuring their properties will be an important and exciting task for the LHC. One of the main focuses for this chapter is the uncertainty in the black hole theory. This has led to the development of a new technique for measuring the black hole temperature that tries to account for many of the uncertainties. This has been applied to a test case and the results used to estimate how well the Planck mass and the number of extra dimensions could be measured.

6.1 Black hole production and decay

There are several scenarios in which the fundamental Planck scale can be as low as a TeV; some were presented in chapter 1. In any theory where trans-Planckian energies are accessible, black hole production is theoretically possible. So far, black holes in extra dimensions have been most studied in the earliest and simplest scenario, that of Arkani-Hamed, Dimopoulos and Dvali (ADD) and this model has been used throughout this chapter. A more detailed overview of black hole theory can be found in [16] and the references therein.

Physics at or near the Planck scale will necessarily be described only by a quantum theory of gravity. However, significantly above this scale, it is assumed that semi-classical general relativity can be used to describe black hole production and decay. To be within this regime, it is necessary

for the mass of the black hole, M_{BH} , to be much greater than the Planck mass, M_{PL} and also for the black hole temperature to be much lower than the Planck mass. This assumption is known as the semi-classical approximation. Throughout this chapter, as in chapter 1, the convention for the Planck mass according to Dimopoulos and Landsberg [2] has been used *i.e.* the Planck mass, M_{PL} is

$$M_{\text{PL}}^{n+2} = \frac{1}{G_4(2\pi R)^n} \quad (6.1)$$

where n is the number of extra dimensions, R is the size of the extra dimensions (assumed to have the same size) and G_4 is Newton's constant in 4 dimensions (see section 1.2.2.1).

6.1.1 Black hole production

If the centre of mass energy, $\sqrt{\hat{s}}$, is much greater than the Planck scale then it is expected that black holes will be well described by the semi-classical approximation. In this approximation, if two particles approach within the event horizon given by their centre of mass energy, then a black hole must form regardless of how violent the interaction is. In this case then, the cross section is geometric [129] and given by

$$\sigma \approx \pi r_{\text{BH}}^2 \left(\sqrt{\hat{s}} \right) \quad (6.2)$$

where r_{BH} is the effective radius of the event horizon. For a spherically symmetric, non-spinning, uncharged black hole, the radius is

$$r_{\text{BH}} = \frac{1}{\sqrt{\pi} M_{\text{PL}}} \left(\frac{M_{\text{BH}}}{M_{\text{PL}}} \right)^{\frac{1}{n+1}} \left(\frac{8\Gamma\left(\frac{n+3}{2}\right)}{n+2} \right)^{\frac{1}{n+1}} \quad (6.3)$$

where $M_{\text{BH}} = \sqrt{\hat{s}}$. It is assumed that the difference between uncharged, non-spinning black holes and more realistic black holes is only incremental. A more careful consideration of the production shows that the cross section is modified by a form factor which is of order unity [130].

The cross section therefore *increases* with centre of mass energy. This is markedly different to perturbative physics where cross sections decrease with $\sqrt{\hat{s}}$. Since it cannot matter what inter-

actions occur inside the event horizon, black hole production suppresses perturbative physics and thus marks the end of short distance physics. This also means that the dominant contribution to the black hole cross section comes when the impact parameter is at its largest which means that black holes will typically be formed with high angular momenta.

Note that the cross section depends only on the centre of mass energy, geometric and kinematic factors. It does not depend on the particle type or flavour. Therefore, although this parton-level cross section must be convoluted with the PDFs at the LHC, black holes can be formed from valence quarks and thus much higher masses are accessible than is normally the case for new physics signals. Some cross sections which include PDFs are given in section 6.2.1.

6.1.2 Black hole decay

Black hole decay in extra dimensions is based on analogy with results and simulations of decays of astrophysical black holes in 4 dimensions. Note that none of the aspects of astrophysical black hole decay, including Hawking radiation, have been experimentally verified.

Once formed, a black hole is expected to decay quickly (although this depends on the details of the extra dimensional model). Estimates of the typical lifetime vary from 10^{-26} s to 10^{-17} s [12, 131, 132]. The decay goes through three major phases:

1. The *balding phase* in which the asymmetries and moments in the event horizon ('hair') due to the violent production are lost.
2. A Hawking *evaporation phase* which begins with a brief *spin-down phase* in which the angular momentum of the black hole is shed followed by a *Schwarzschild phase* which accounts for the majority of the energy loss.
3. Finally a *Planck phase* which occurs at the end of the decay when the black hole mass or temperature reach the Planck scale. At this stage the black hole is often referred to as the 'remnant'.

Of these stages of decay, only the evaporation phase is at all understood although work has

been done on the others using well-informed guesses. They are discussed further in section 6.3. The evaporation phase occurs via Hawking radiation [133] which is characterised by a temperature, T_H ,

$$T_H = \frac{n+1}{4\pi r_{\text{BH}}} \quad (6.4)$$

which gives the important relation

$$\log T_H = \frac{-1}{n+1} \log M_{\text{BH}} + \text{constant} \quad (6.5)$$

where the constant depends only on M_{PL} and geometric factors. The energy spectrum for Hawking radiation from an uncharged, non-spinning black hole of fixed temperature in its rest frame is approximately black-body and is given by

$$\frac{dN}{dE} \propto \frac{E^2}{\exp(E/T_H) \mp 1} \frac{\gamma}{T_H^{n+6}} \quad (6.6)$$

where γ are the ‘grey-body’ factors that are spin and angular momentum dependent. The denominator includes a spin statistics factor that is -1 for bosons and $+1$ for fermions.

The relative emissivities for different particles can be calculated from the grey-body factors and from the numbers of degrees of freedom for each particle. The relative emission probabilities are given in table 6.1 for $n = 3$.

6.2 CHARYBDIS: A black hole event generator

The black hole event generator CHARYBDIS [134, 135] has been used to generate Monte Carlo event samples. It is designed to simulate the production and decay of black holes in hadron collider experiments. The generator is interfaced, via the Les Houches accord [136], to HERWIG 6.5 [64, 137] which performs the parton shower evolution of the partons produced in the decay and their hadronization.

The CHARYBDIS event generator models black hole production and decay using the theory

Particle	Probability (excl. grey-body)	Probability (inc. grey-body)
Quarks	0.6102	0.5964
Gluons	0.1356	0.1412
Charged leptons	0.1017	0.0994
Neutrinos ^a	0.0508	0.0497
Photon	0.0169	0.0177
Z boson	0.0254	0.0283
W bosons	0.0508	0.0567
Higgs boson	0.0085	0.0107

Table 6.1: Particle emission probabilities excluding and including grey-body factors for $n = 3$.

presented in section 6.1, *i.e.* it assumes that the semi-classical approximation is valid and that black holes are well described as non-spinning and uncharged. CHARYBDIS models only the Schwarzschild phase and includes the grey-body factors in the energy spectrum of decay products.

The decay particles are chosen according to the relative emission probabilities calculated for the evaporation phase. The decay of black holes is subject to the constraint that baryon number and charge must be conserved^b.

CHARYBDIS includes many options to modify its behaviour, which can be important in investigating some of the systematic effects on black hole decays which are dealt with in detail in section 6.3. Since black hole production is not understood near the Planck mass, it generates black holes within a mass range that can be controlled. The cross section thus cuts off abruptly at the lower edge of the mass range. CHARYBDIS has options controlling the particles it decays into and whether the black hole temperature is allowed to vary; it can also control some aspects of the remnant decay. However, it does not model emission into the bulk and the remnant cannot be stable.

^aThis assumes that right-handed neutrinos do not exist.

^bBaryon number may not be a conserved quantity but it is necessary for technical reasons. This should not be experimentally observable.

Nevertheless, it is the most complete black hole event generator available.

6.2.1 Black hole production cross section

The cross section reported by CHARYBDIS is given in table 6.2 for some indicative values of n and M_{BH} . The default HERWIG 6.5 PDF set has been used^c. Note that, even though the cross section does not include the form factors calculated in [130] which tend to reduce the n dependence, the cross section does not vary much with n (see also section 6.13). Also it is clear that if M_{PL} is of the order of a TeV, black holes will be produced copiously at the LHC, indeed it would qualify as a black hole factory, with even extremely massive black holes being produced in observable numbers. This is quite unlike the normal situation with exotic signals and with such a wealth of data it should be straight-forward to discover black holes (see section 6.7). Measuring their properties may be much harder however, as the next section discusses.

Topology		Total Cross Section (fb)
$M_{\text{BH}} \geq 5 \text{ TeV}$	$n = 2$	55,500
	$n = 4$	33,000
	$n = 6$	30,100
$M_{\text{BH}} \geq 8 \text{ TeV}$	$n = 2$	560
	$n = 4$	300
	$n = 6$	260
$M_{\text{BH}} \geq 10 \text{ TeV}$	$n = 2$	6.9
	$n = 4$	3.5
	$n = 6$	3.0

Table 6.2: The black hole production cross sections at the LHC for $M_{\text{PL}} = 1 \text{ TeV}$ as given by CHARYBDIS. Note that CHARYBDIS does not include the form factors mentioned in section 6.1.1.

^cThe default HERWIG PDF is the average of the MRST98 central and higher gluon leading order fits [138].

6.3 Model uncertainties

There are a large variety of different theoretical models and variations that could be regarded as model uncertainties. However, for the studies in this chapter, I have considered only the ‘standard’ models in which black hole decay occurs approximately as discussed in section 6.1 within the standard ADD scenario. Nevertheless, there are still many uncertainties and assumptions. Whilst it may be possible to measure the model parameters, typically M_{PL} and n , given some set of assumptions, it is generally much harder to verify these assumptions. Worse, a different set of assumptions may give very different values for the same model parameters. These assumptions and uncertainties thus result in large systematic errors. This section discusses these so that analyses can be developed which are as widely applicable as possible. This section also shows how some of them can be manifest as large systematic errors.

6.3.1 Production cross section

As with many aspects of black hole production, there is no clear consensus on the validity of the semi-classical calculation of the cross section. Several effects that could suppress black hole production have been considered [139, 140] whilst other authors have refuted these or confirmed the semi-classical calculations [130, 141, 142, 143]. The current state of the literature seems to suggest^d that either black holes will be seen at the LHC, in which case the semi-classical approximation is likely to be good, or the semi-classical approximation is bad, in which case black holes will not be produced.

The form factor that relates the left- and right-hand sides of equation 6.2 is somewhat uncertain, although they have been calculated using numerical simulations [130]. More importantly, the transition from the parton-level cross section to the hadron-level cross section is based on factorisation and it is unclear whether this would remain valid into the trans-Planckian regime. There are also issues regarding the PDFs which are discussed in section 6.13.

^dAssuming the ADD scenario is realised.

6.3.2 The first stages of decay

CHARYBDIS does not model the initial ‘balding’ or spin-down phases of the black hole decay. The amount of energy emitted from the black hole during these phases is expected to be small (some estimates give the fraction as 16–25% [16, 144]) so such an omission should not be too significant. However, it is probable that the energy spectrum will be modified at low energies.

6.3.3 Energy deposition on the brane

Estimates vary as to how much energy is expected to be emitted into the bulk by emission of gravitons, but it could be significant. One estimate suggests that the fraction of energy emitted into the bulk is around 20% for $n = 2-4$ rising to nearly 50% for $n = 7$ [12]. Any energy emitted into the bulk will make an accurate measurement of the mass more difficult. Although this effect could in principle be observed as a change in the expected shape of the cross section as a function of black hole mass, determining this would be experimentally challenging. In these studies it has been assumed that all the energy would be deposited on our brane. A modified generator and further study would be necessary to understand the full impact of this assumption although it would necessarily impact the mass resolutions presented in section 6.8.

There have also been suggestions that radiation into the bulk could be much greater than previously thought and that this could cause the black hole to recoil off of our brane thus seeming to disappear [145].

6.3.4 Back reaction and the effect of the kinematic limit

The energy distribution for emission from a black hole, the grey-body spectrum, extends to infinite energy. However, the maximum energy that is kinematically allowed is half the black hole mass. The distribution must therefore be modified near and above this kinematic limit. Of course, for very massive black holes this is not an issue — the probability of an energy greater than the kinematic limit being chosen is negligible. But, for the black holes considered here, which are near to the validity of the semi-classical approximation, this limit can be significant. Also, if

T_H is allowed to vary, as a black hole decays and gets lighter, the temperature increases and this effect must become important. Of course, it can be argued that at this point the semi-classical approximation is not valid and the decay should move to the Planck phase. However, there is an intermediate, transition region and it is not clear what should be done in these cases.

Figure 6.1 shows the energy of the primary generator level decay products in the rest frame of the black hole. which shows how the kinematic limit affects many decays. The generator implements two options when it samples from the grey-body distribution. In the first case (KINCUT=FALSE), if an unphysical decay is chosen, it is thrown away and a new one is chosen. This process continues until the black hole has a mass less than the Planck mass when the decay moves to the final, remnant, stage. In the other option (KINCUT=TRUE), when an unphysical decay is chosen, the black hole decay moves straight to the final stage. The final stage of the decay is dealt with in the next section. It should be noted that for high temperature black holes, where the probability of an unphysical decay is large, the different choices implemented in the generator will lead to a large difference in the multiplicities and will have a significant impact on the energy distributions.

The energy emission spectra have also been calculated without allowing for the back reaction from the metric. This amounts to the assumption that $E \ll M_{BH}$ or alternatively $T_H \ll M_{PL}$ [12]. Again, at higher energies a full quantum theory would be needed to calculate the probabilities. It has been assumed that any effects due to this are small.

6.3.5 Remnant decay

At the end of the evaporative phase, a Planck-scale black hole or remnant remains. Some authors have argued that this remnant could be stable, due to perhaps the Information Paradox [131, 132]. In this case, the remnant would carry away much of the energy, however this could possibly be handled in the same way as emission into the bulk, *i.e.* by studying the parton-level cross section variation with measured black hole mass.

Alternatively, it is normally assumed that the remnant decays to several particles. CHARYBDIS




Figure 6.1: Energy of the generator level decay products in the rest frame of the black hole for a 5 TeV black hole and 1000 events. The colour scale indicates the number of particles in each bin. **(a)** for $n = 2$ the kinematic limit ($E = M_{\text{BH}}/2$, black lines) constricts the energy distribution at low masses. **(b)** for $n = 4$ the kinematic limit clearly affects the energy distribution at all masses.

implements this remnant decay as an isotropic decay into 2–5 bodies (the number is an option). When the remnant decay occurs depends on the option chosen for handling the kinematic limit as described in section 6.3.4. It should be noted therefore, that the uncertainty here can easily affect the multiplicity and energy spectra. One example of an affected experimental observable is the photon energy spectrum. Figure 6.2 shows the photon energy distributions (of all photons in the event) for 2-body and 4-body remnant decays for two values of n . Even for $n = 2$ there is a noticeable effect, but for $n = 4$, the effect is very large.

6.3.6 Time-variation and black hole recoil

It has been argued [2] that due to the speed of the decay, the black hole does not have enough time to equilibrate between emissions and therefore that the time variation of the temperature can be ignored. Therefore, the initial Hawking temperature might be measured by fitting the grey-body distribution for a fixed temperature (equation 6.6) to the energy spectrum of the decay products for different bins in the initial black hole mass. Using equation 6.5 the number of dimensions can




Figure 6.2: The photon energy distributions for (a) $n = 2$ and (b) $n = 4$. The black and red lines are for 2-body and 4-body remnant decays respectively.

then be extracted.

C. M. Harris has performed this procedure in a very informative study [12] for two different cases: T_H fixed and T_H time-varying. Each case was generated with $n = 2$ in 500 GeV bins and M_{BH} between 5000 and 10000 GeV without grey-body factors. For each mass bin the (fixed T_H) black-body spectrum was fitted to the generator level electron energy. The Hawking temperature from each fit was then plotted against the black hole mass. In both cases the fit was reasonable although there was some degradation when fitting the spectrum to the time-varying case^e.

Figure 6.3a shows the result of this for fixed T_H together with a fit using equation 6.5 from which the number of extra dimensions was determined to be $n = 1.7 \pm 0.3$. Figure 6.3b shows the result of the same procedure and the same test case but with T_H time dependence turned on. In this case n was determined to be $n = 3.8 \pm 1.0$ which is well away from the model value. The conclusion is therefore that this is a systematic effect with a strong impact on any measurement of n .

One effect that has not been taken into account in previous studies is the recoil of the black

^eFor instance in the 8 TeV bin the χ^2 value went from 146 for the data generated with fixed T_H case to 173 for the data with time-varying T_H . The number of degrees of freedom was 134 in both cases.




Figure 6.3: The plot of $\log(T_H)$ versus $\log(M_{BH})$ for $n = 2$ and $M_{PL} = 1$ TeV, (a) with a fixed Hawking temperature, and (b) with time dependent Hawking temperature. The statistics used correspond to 30 fb^{-1} . From [3].

hole. When a particle is emitted from the black hole, the black hole recoils against it. Therefore the next emission is in a boosted frame. Even in the case of a fixed temperature decay, the effects of recoil become more significant as the decay progress as the black hole gets lighter and lighter. This is exacerbated in the time varying case since the black hole also gets hotter as it decays. Any analysis which makes use of the energy spectrum should therefore aim to be insensitive to this.

6.4 Event generation and detector simulation

The CHARYBDIS event generator described above was used for the event generation. Unless otherwise mentioned, the CHARYBDIS options were set as follows:

- Time variation of the black hole temperature was on (TIMVAR=TRUE).
- Grey-body effects were on (GRYBDY=TRUE).
- The black hole was allowed to decay to all Standard Model particles including Higgs bosons (MSSDEC=3).
- Kinematic cut-off was turned off (KINCUT=FALSE).

- The number of particles in the remnant decay was 2 (NBODY=2).

This set of options together with a Planck mass set of 1 TeV is called the ‘test case’. The number of dimensions is always specified. If a mass is given, then the generator was forced to produce black holes with a fixed mass, otherwise the mass range was set to 4500–14000 GeV. In the case that a range was used, only reconstructed masses greater than 5 TeV were used in order to stay within the semi-classical approximation (this requirement is necessary but not sufficient as high temperatures can also cause the semi-classical approximation to break down). The lower limit of 4500 GeV was set to partially account for mass resolution effects (see section 6.3.4 for an example of where this is important) although without an understanding of how the black hole production cross section varies near threshold, the magnitude of such effects is unknown.

The event generation and detector simulation were carried out using the ATLAS offline software framework, Athena version 6.0.4. Using Athena, CHARYBDIS was interfaced to HERWIG 6.5 for hadronization and showering. For the majority of the studies the ATLAS fast simulation software, ATLFAST was used. Standard options were used for ATLFAST with the high luminosity option off and jet reconstruction using the cone algorithm with $R = 0.4$. Details of individual data sets are given where they are used. Typically, where a range of black hole masses was generated, it was not possible to generate 30 fb^{-1} of data, but as most analyses are not statistically limited, this is not significant.

ATLFAST was developed and tuned for low multiplicity, low energy Standard Model and Higgs events therefore it may be unreliable for the rather extreme black hole events considered here. In particular the many very high p_T overlapping jets may not be well modelled by ATLFAST which could also affect the p_T measurement. The ATLAS full simulation software has been used to give an improved simulation of the ATLAS detector and to back up the ATLFAST results. The design layout of the ATLAS detector was simulated using GEANT3 and most options were left at their default. The GEANT output was passed through the Athena standard digitisation and reconstruction. The full reconstruction was mostly performed with standard options. The main points to note are:

Exploring Higher Dimensional Black Holes

- Noise was turned off: it is known that several algorithms are not optimised for noise in version 6.0.4.
- The ATLAS calorimeter has a different response to electro-magnetic and hadronic activity. Weights were applied to calorimeter cells on the basis of their energy density. This is commonly referred to as H1-style weighting [146, 147] and was used for both jet and \cancel{p}_T reconstruction.
- The cone jet algorithm was used with $R=0.4$.

Full simulation is a time consuming exercise — simulating the events took about 45 mins/event and reconstructing them 5–15 mins/event. Two samples of ≈ 1000 events each were fully simulated and reconstructed which represents nearly 3 months of CPU usage. The samples were generated for the test case with a range of black hole masses from 4500–14000 GeV. The first sample was for $n = 2$, for which 900 events were generated and the second sample for $n = 4$, for which 860 events were generated.

6.5 Corrections to the full simulation

Although Athena 6.0.4 is validated for physics use, it is not a finished product. Improving its performance and realism is one of the many tasks being undertaken in the run up to the turning on of ATLAS in 2007. Consequently there are many effects which have not been simulated (principally noise) and many algorithms that have not yet been finalised. It is hoped that any loss in performance due to effects not yet simulated will be cancelled out by improved algorithms.

For these studies, incremental modifications and improvements were avoided, but there were clearly a few critical problems that had to be fixed. These are detailed in the following sections.

6.5.1 Muon \cancel{p}_T adjustment

The \cancel{p}_T algorithm only takes account of energy imbalances in the calorimeter. In particular, it doesn't account for energy lost by muons. This was adjusted for by using the correction

$$\cancel{p}_T^{corr} = \cancel{p}_T - \sum_{\text{muons}} (\mathbf{p}_{T_i} - \mathbf{p}_{T_{lost}}) \quad (6.7)$$

where $\mathbf{p}_{T_{lost}}$ is the energy lost by the muon in the calorimeter if it is available (see section 6.6.2). The \cancel{p}_T bias and resolution were determined by comparing the reconstructed \cancel{p}_T with the sum of the transverse momenta of all invisible particles in the generator final state. The bias, $\Delta(\cancel{p}_T) = \cancel{p}_T^{\text{recon}} - \cancel{p}_T^{\text{true}}$, is shown before and after applying the correction in figure 6.4. The \cancel{p}_T resolution before correction also shows a similar trend where there is a significant increase when there is a muon in the event.




Figure 6.4: The \cancel{p}_T bias as a function of the number of muons (a) before and (b) after correcting with equation 6.7. Note the different scales for the two graphs.

6.5.2 EM–Jet correction

It was noticed that there is a tendency for high energy photons and electrons to be reconstructed both as a photon/electron and also as a jet with almost identical energy at the same place. Clearly there is some problem in the reconstruction, possibly relating to how leakage from the EM calorimeter into the hadronic calorimeter is handled. For this study, this was corrected for using the following algorithm for each photon and electron:

1. Find the nearest jet to the photon/electron.
2. If it is further away than $\Delta R=0.4$ then do nothing.
3. Otherwise perform 4-vector subtraction: $p_{jet}^{corr} = p_{jet} - p_{EM}$.
4. If the remaining jet is closer than $\Delta R=0.1$ and has $E_{jet}^{corr} < 10\% \times E_{EM}$ then remove the jet.

This algorithm attempts to account for the fact that jets will have been calibrated at the hadronic scale rather than the EM scale and that other hadronic energy could have been merged into the jet.

6.5.3 Jet energy $|\eta|$ dependence

As has been noted elsewhere [148], there is a strong $|\eta|$ dependence in the jet energy resolution. A simple factor was applied to reduce the effect of this, but such a correction should really be applied at the cell level by the reconstruction. The uncorrected jet p_T resolution is shown in figure 6.5 as a function of $|\eta|$ and this data was used to determine the correction factor. The true p_T was determined by applying the same jet reconstruction algorithm to the generator level particles. Jets were only included if the reconstructed and true jets were close enough so that $\Delta R < 0.2$. The correction factor used is:

$$F = \frac{1}{1 - 0.07|\eta| + 0.0074|\eta|^2}. \quad (6.8)$$

The result of this correction can be seen in section 6.6.3. Although the difference between the corrected energy and the original energy could be used to adjust the p_T measurement, this was not done as it was found to decrease the p_T resolution. Performing this additional correction would

in any case treat cells inconsistently for the p_T measurement and is presumably the reason for this effect.




Figure 6.5: The jet p_T resolution as a function of $|\eta|$ before correction showing a quadratic fit that equation 6.8 was based on.

6.6 Performance of the full simulation

In this section the performance of the full simulation is reviewed and compared to the fast simulation which was run on the same events and also to that expected from the TDR.

Many other studies have also performed with the full simulation. Some of these have higher statistics and they cover a wide variety of physics channels. Studies relating to Athena 6.0.4 can be found on the web pages of the Athens Physics Workshop [149] but are rarely formally written up. Automated studies with more recent version of the software may be found on the Validation group web pages [150].

6.6.1 Electron and photon performance

Electrons and photons are reconstructed together in Athena using the calorimeter as the primary source. Once EM clusters have been identified, they are matched to tracks to give e/γ separation. This is performed by the EGamma package which uses either of the two available track reconstruction packages: iPatRec [151] and xKalman [152]. There is little difference between them [153] and for this study the E/p ratio from iPatRec was used.

The e/γ energy, eta and phi were all taken from the calorimetry, which should be optimal for the high energy particles which are most interesting here. The energy resolution is expected to be about 1.6% or better for 50 GeV photons with a constant term of 0.7% or better with similar results for electrons. The results show that energy resolution is in agreement with the TDR for both photons and electrons. The tails in the distributions are primarily due to electrons/photons that are reconstructed in the vicinity of a jet. In these cases, overlap and misassignment of energy can decrease the energy resolution.




Figure 6.6: The ratio of reconstructed p_T to true p_T for (a) electrons and (b) photons. The true particles were required to have $p_T > 100$ GeV and the reconstructed and true particles were required to be matched to $\Delta R < 0.1$.

The efficiency of electrons and photons has been investigated for high p_T particles that are well

isolated from jets ($\Delta R > 0.7$) in both the true and the reconstructed events. Typically reconstructed electrons/photons were well matched to true electrons and photons with good energy resolution. However, of the 142 true electrons in both data sets, only 108 were reconstructed (76%). Of the ones that were not reconstructed, 21 (15%) were reconstructed as photons (probably due to conversion), 8 (6%) were at very high energy (greater than 1 TeV) and were reconstructed as jets, 2 were electrons that overlapped and for 4 (3%) there was no obvious reason why they were reconstructed as jets rather than electrons. These figures are in broad agreement with the TDR figures which suggest around 6% of electron candidates are misidentified as photons and this increases with energy.

The efficiency figures are similar for photons although there were significantly fewer. The primary cause of failure to reconstruct a photon is conversion which causes an electron to be reconstructed instead (up to 30% of photons convert to e^+e^- pairs). Excluding these, the reconstruction efficiency was 84%.

Conversion is a significant cause of electron/photon misidentification. A package is under development to identify cases where conversion occurs, but was not used for this study. In any case, for the events considered here, conversion will not be significant for the determination of mass or for the kinematic limit analyses. It will however impact (slightly) on the energy distributions. At very high energies there seems to be a problem reconstructing EM particles.

6.6.2 Muon performance

Two muon reconstruction packages are available within Athena: `Muonbox` [154] and `Moore` [155]. Both reconstruct the track in the muon spectrometer which is used to determine the muon momentum. `Muonbox` then extends the track back through the calorimeter to the interaction vertex, estimating how much energy would have been lost in the calorimeter (this is used in the \cancel{p}_T correction, see section 6.5.1). `Moore` only provides the track parameters at the entrance to the muon spectrometer. Neither package uses the inner detector or the calorimeter although algorithms are under development to do this [156, 157].

Both packages were used and compared in a similar manner to the electrons and photons. The track parameters extrapolated back to the interaction vertex were used for `Muonbox`, which was found to have a low efficiency for reconstructing muons: nearly 60% of well isolated muons with energies above 100 GeV were not reconstructed. This, of course, also has a significant impact on the p_T measurement. It is not clear why `Muonbox` showed such poor performance although it could be related to the track extrapolation.

`Moore` on the other hand, had an overall efficiency of 95% (127 reconstructed muons out of 134 total) with 5 of the missed muons at $\eta \approx 0$ where there is a crack in the muon spectrometer (enlarged since the TDR). Due to its much better performance, `Moore` was used for all further studies. The TDR results showed that very high energy muons can fail to be reconstructed due to showers which can spoil the tracks. There the reconstruction efficiency for 1 TeV muons was found to be 90%. The p_T resolution is shown in figure 6.7 which shows a resolution of 4.6%. There is a slight bias to low energies although this is likely to be due to energy loss in the calorimeters which is typically a few percent. This is in broad agreement with the TDR, although a detailed comparison is difficult as the resolution drops significantly as p_T increases. In the TDR, the resolution is given as 3% for p_T of 100 GeV dropping to 5% at 500 GeV and 8% at 1 TeV.




Figure 6.7: The ratio of reconstructed p_T to actual p_T for muons reconstructed with `Moore`. The true particles were required to have $p_T > 100$ GeV and the reconstructed and true particles were required to be matched to $\Delta R < 0.1$.

6.6.3 Jet performance

As mentioned in section 6.4, jets were reconstructed using the cone algorithm with $\Delta R = 0.4$. The provided H1 style weights were used to calibrate the calorimeter. True jets were generated using the same algorithm but applied to the generator level final state particles. They were matched to the reconstructed jets in η - ϕ with the requirement $\Delta R < 0.2$.

Figure 6.8 shows a number of performance plots for jet reconstruction. The jet reconstruction efficiency is good for jet $p_T > 30$ GeV with notable drops at $|\eta| \approx 1.5$ and $|\eta| \approx 3.2$ where transitions from barrel to forward to end-cap calorimeters occur. This is also visible in the jet resolution against $|\eta|$ in figure 6.8a. Even aside from these gaps, it is worth noting that the resolution is not at all flat with $|\eta|$ despite the correction that has been applied. As mentioned earlier, cell level corrections need to be improved to tackle this. The jet resolution linearity with p_T is shown in figure 6.8b. This shows that at high energies the jet p_T tends to a low value whilst at lower energies the the jet p_T is too high. It is not clear why this behaviour occurs other than noting that the H1-style weights applied depend only on the cell energy. There is no $|\eta|$ based correction (other than the correction applied afterwards) and there is no calibration at the tower or jet stage.

For comparison, similar plots for the fast simulation are shown in figure 6.9. These are generated from the same events as those in figure 6.8 and show that there is essentially no $|\eta|$ dependence in the fast simulation jet reconstruction. The jet resolution linearity with p_T broadly agrees with the full simulation at high energies, but differs markedly at low energies.

6.6.4 \cancel{p}_T resolution

Missing p_T is reconstructed in a similar way to jets. The same H1-style weighting scheme is applied to the cell energies and the total imbalance is returned. Figure 6.10 shows the resolution for both the full and fast simulations. The full simulation resolution is significantly worse than expected: the TDR gives $\sigma(p_{x,y}^{miss}) = 0.39 \times \sqrt{\sum p_T}$ which would suggest the resolution for 5 TeV black holes should be 20–30 GeV. Instead, the \cancel{p}_T resolution was found to be approximately flat over a range of over 3 TeV at about 160 GeV. Other studies have investigated $\sum p_T$ as high as 2 TeV

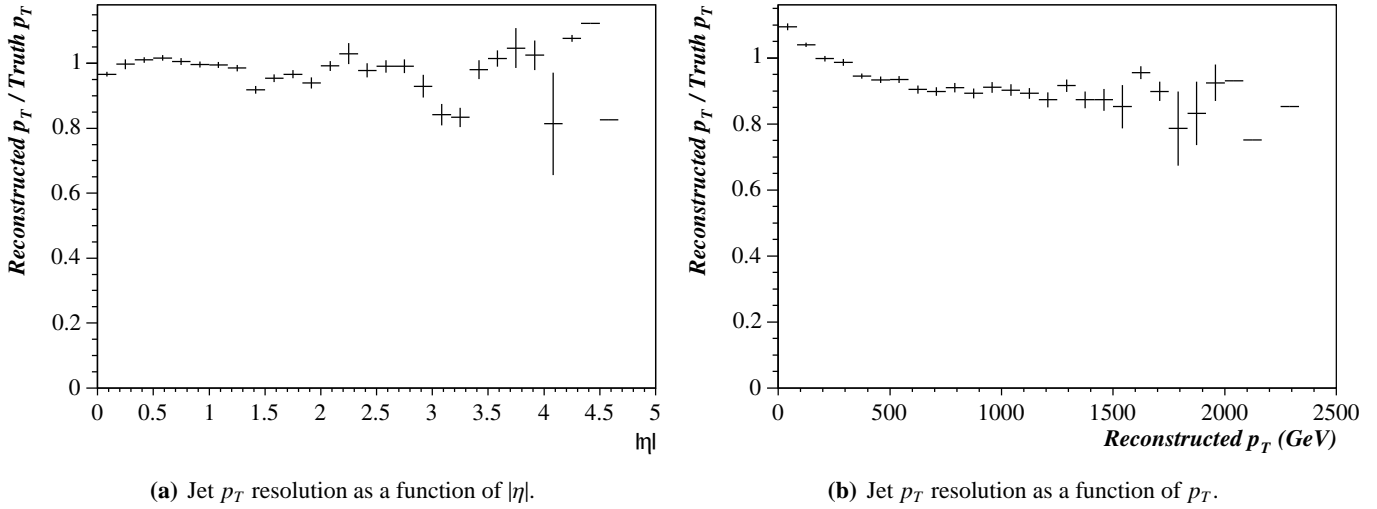


Figure 6.8: Jet performance graphs for the full simulation.

(which is towards the bottom of the range here — black hole events are likely to be the worst case of \cancel{p}_T reconstruction) and they also found a deviation from the $\sigma \propto \sum p_T$ relationship. This effect is still under study in the \cancel{p}_T reconstruction group. Hopefully the performance will be improved by the time ATLAS starts, but in these studies, the decrease in \cancel{p}_T resolution must be considered




Figure 6.9: Jet performance graphs for the fast simulation from the same events as used in the full simulation.

when comparing the full and fast simulations.

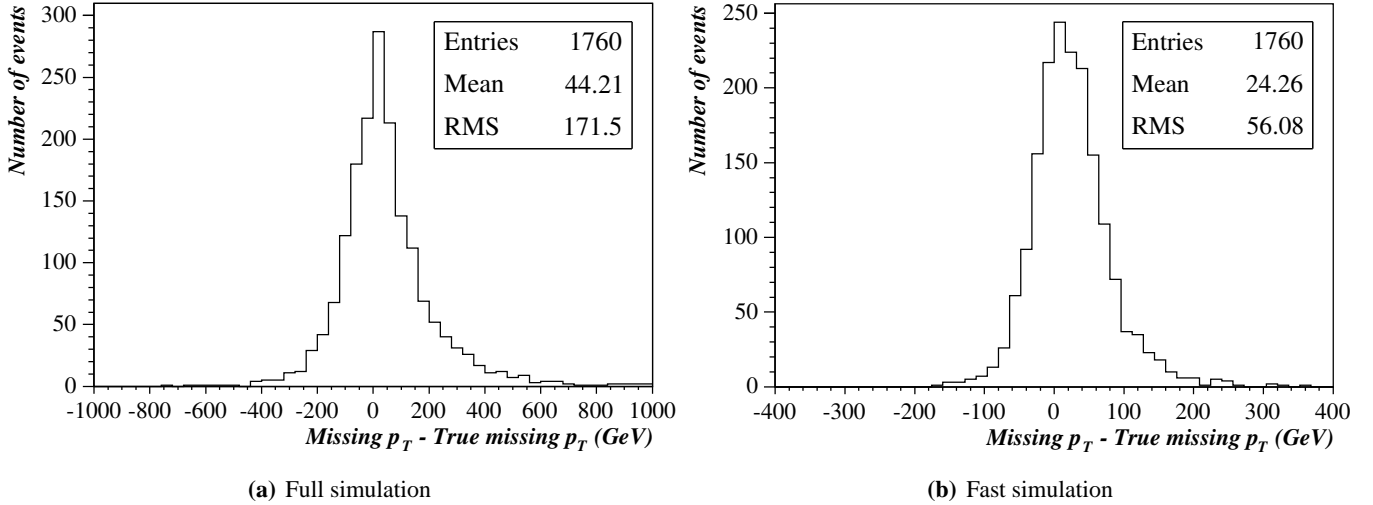


Figure 6.10: \cancel{p}_T resolution for (a) full simulation and (b) fast simulation. Both fully simulated data sets have been combined. Note the different horizontal scale.

6.7 Characteristics of black hole decays

The most notable characteristic of black holes is their access to extremely high energies. As can be seen in the cross sections in section 6.2.1, since all quarks and gluons can take part in the production, black hole masses as high as 10 TeV are accessible at a reasonable rate although this is strongly dependent on M_{PL} . Since most of this energy will be released hadronically, black hole events would be readily observable with a very large total transverse energy in the calorimeters ($\sum p_T$). Black hole events are also quite dissimilar from other processes in that, due to their non-perturbative nature, their cross section rises with energy. Determining that this is happening will be strongly dependent on a good understanding of the PDFs and on the luminosity measurement but such an observation would be a smoking gun for black holes and non-perturbative physics.

The nature of black hole decays depends greatly on the multiplicity^f. Several authors have

^fThe event multiplicity is defined depending on the context. In a theoretical or generator context it is the number of particles emitted in the *evaporation phase* of the black hole decay. In an experimental context it is the total number of jets, leptons and photons.

Exploring Higher Dimensional Black Holes

commented on the expected characteristics if the multiplicity is very large [12, 16, 158]. With a multiplicity ~ 100 , black hole events would be isotropic (due to their mass, black holes would have low boosts) and spherical. It would be difficult to distinguish individual jets and isolating leptons would also be hard. Experimental tests of Hawking radiation and other features would therefore have to rely predominantly on hard muons and possibly hard electrons and photons. Also there would be many neutrinos per event (the probability of a single emission being a neutrino is $\sim 5\%$) so their missing momentum would largely average out giving a systematic offset in the measurement of the black hole mass.

At lower multiplicities, ~ 20 , the events would be cleaner and experimentally easier to handle. Individual jets should be recognisable and there would be no problems isolating electrons, muons and photons. Also since the number of neutrinos would be ~ 1 , it would be possible to exclude events with large \cancel{p}_T to improve mass resolution. As the mass of the black holes approaches the limit of the semi-classical approximation, the multiplicity drops and the events become less spherical. These events would be harder to understand as they would become dominated by quantum gravity effects.

Black hole events have an energy distribution that is mainly derived from the evaporation phase and thus approximately black body. However, this will be modified at low energy by the initial stages of the decay (see section 6.3.2) and at higher energies by the black hole remnant. If the typical energy of the remnant decay products is much higher than $\sim 2T_H$, then it may be possible to disentangle any remnant effects. However, the remnant may merge with the grey-body spectrum in which case the remnant would be unimportant so long as the multiplicity is high. Finally, the energy spectrum may be modified by recoil effects. Again these are more important for lighter black holes.

Given the LHC energy of 14 TeV and limits on the Planck mass, it is likely that if black hole production occurs, it will be on the edge of validity of the semi-classical approximation. Therefore multiplicities are likely towards the low end of expectations and disentangling the evaporation phase from the others is likely to be difficult.

6.7 Characteristics of black hole decays

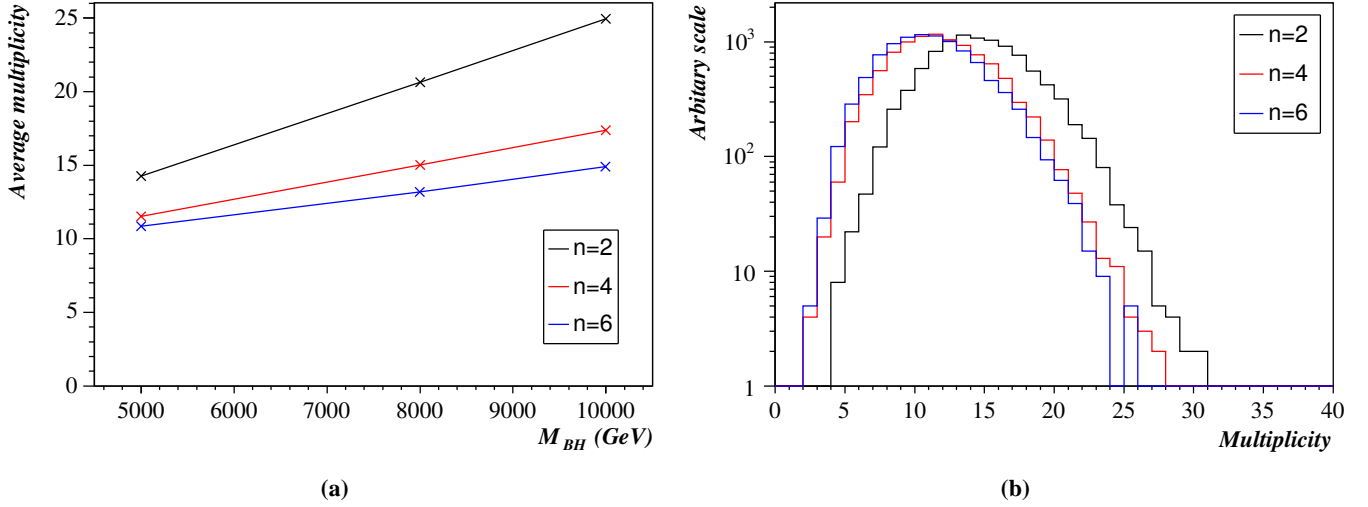


Figure 6.11: Multiplicity (total number of jets, photons, electrons and muons); **(a)** shows the average multiplicity for different mass black holes; **(b)** shows the multiplicity distributions for 5 TeV black holes.

A black hole decay is also characterised by a large total transverse energy (figure 6.12) which increases as the black hole mass increases. Due to the large multiplicity and moderate black hole boost, the events are rather spherical. These characteristics are very different from standard model and SUSY events. Therefore, selecting events with high $\sum p_T$, high multiplicity (> 4) and relatively high sphericity should give a pure set of black hole events. In addition, it should be noted that the already small Standard Model background would be suppressed by the black hole production [2]. This means that discovering black holes and measuring their properties are experimentally quite different. The discovery potential for ATLAS was discussed in [159]. However, this chapter is concerned with measuring black hole properties and in this case the Standard Model backgrounds are negligible (and suppressed!) and have therefore been ignored.

There are two further characteristics which will be interesting to measure and confirm the nature of the events: the \cancel{p}_T distribution and the black hole charge.




Figure 6.12: The distribution of $\sum p_T$ (a) for 5 TeV and (b) for 8 TeV black holes.

6.7.1 \cancel{p}_T distribution

Although not all black hole decays contain neutrinos, some have one or more with energies that can be as high as half the black hole mass. In contrast, most of the Standard Model processes tend to have lower missing transverse momenta – the missing energy can be even larger than for much of SUSY parameter space. Figure 6.13 presents the distribution of \cancel{p}_T for two data samples. Depending on the cross section, it is clear that events with \cancel{p}_T as high as 2 TeV or more may be detected.

6.7.2 Black hole charge

Black holes are typically formed from valence quarks, so it is expected that the black holes would be charged. The actual charge is somewhat energy and PDF dependent, but should be about $+2/3$. The rest of the charge from the protons is expected to disappear down the beam pipes or at high $|\eta|$. The average black hole charge, $\langle Q_{\text{BH}} \rangle$, can be measured by determining the average charge of the electrons and muons, $\langle Q_\ell \rangle$, which should be equal to the black hole charge times the probability of emitting a charged lepton. Figure 6.14 shows such a measurement for the test case with $n = 2$ which gives $\langle Q_\ell \rangle = 0.133 \pm 0.002$ and thus $\langle Q_{\text{BH}} \rangle = 0.69 \pm 0.01$ using the expected charged lepton




Figure 6.13: The distribution of \cancel{p}_T for $n = 2$ (black line) and 6 (red line) for the test case.

emission probability of 0.1936. It is possible that if the remnant is stable, its average charge is non-zero, in which case, such a measurement would not be possible.




Figure 6.14: The average charge of the electrons and muons for the test case with $n = 2$ for approximately 1fb^{-1} of data.

6.8 Measurement of the black hole mass

The black hole is reconstructed simply by summing the 4-momenta of all the particles in the event. This gives good measurements of the black hole mass and its momentum (important later — see, for example, section 6.12.1).

Several cuts were made to improve the mass resolution. The missing momentum was limited to 100 GeV to remove the majority of events with neutrinos in them. Also, it was observed that events with jets at high $|\eta|$ were much more likely to be reconstructed poorly and in particular, to have a very high mass. A cut was therefore made to reject events with any jets with $|\eta| > |\eta|_{cut}$. This effect is illustrated in figure 6.15 which gives mass resolution plots for $n = 2, 5$ TeV black holes for two choices of $|\eta|_{cut}$. As can be seen the efficiency is very dependent on this parameter, but so are the tails, particularly the high mass tail. Presumably this tail is caused by underlying event fragments at high $|\eta|$ giving an apparently high mass. Since a good mass resolution is generally important in the following analyses, the cut was set to reject events if there were any jets outside the tracking detector ($|\eta|_{cut} = 2.5$, jets were required to have a minimum p_T of 10 GeV).




Figure 6.15: Mass resolution for $n = 2, 5$ TeV black holes with (a) $|\eta|_{cut} = 2.5$ and (b) $|\eta|_{cut} = 3.2$. Note the significantly large high tail which meant the fit range had to be reduced.

The Gaussian mass resolution, bias and efficiency are presented in table 6.3 for selected mass and n points. The bias and resolution are generally good for such massive objects, however, it is clear that the bias increases slightly with multiplicity which is probably due to multiple neutrino emissions which cancel to give low \cancel{p}_T . The efficiency is also lower for higher mass black holes. Again, this can be understood as being due to the greater probability of neutrino emission which will tend to give $\cancel{p}_T > 100$ GeV and thus fail the cuts.

Topology		Mass Resolution (GeV)	Bias (GeV)	Efficiency (%)
$M_{\text{BH}} = 5$ TeV	$n = 2$	176	-70	17
	$n = 4$	153	-19	20
	$n = 6$	171	-3	21
$M_{\text{BH}} = 8$ TeV	$n = 2$	183	-237	7.9
	$n = 4$	183	-108	12
	$n = 6$	210	-56	13

Table 6.3: The reconstructed Gaussian mass resolution and the overall signal efficiency after the mass resolution cuts.

In addition, the mass resolution was checked using the full simulation samples. Since the full simulation samples were generated with a range of masses, this check was only done at 5 TeV with the true M_{BH} for the fully simulated sample restricted to $4500 < M_{\text{BH}}^{\text{true}} < 5500$. Also since the \cancel{p}_T resolution is worse for full simulation, the \cancel{p}_T cut was increased to 200 GeV (see section 6.6.4). This was chosen to give the same efficiency as for the fast simulation, but will naturally result in a broader distribution. Figure 6.16 shows the mass resolution for 5 TeV black holes for the $n = 2$ case alone and also for both $n = 2$ and $n = 4$ added to give greater statistics. These plots should be compared to figure 6.15a. As can be seen, the resolution is broadly comparable and suggests that, at least for the mass measurement, the fast simulation is a reliable indicator of the expected ATLAS performance.




Figure 6.16: Mass resolution using the full simulation for 5 TeV black holes; (a) for $n = 2$ and (b) for $n = 2$ and 4 combined.

6.9 Previous attempts to determine the model parameters

One technique that has been suggested [2] uses the energy spectrum of electrons and photons with energies less than $M_{BH}/2$ to determine the black hole temperature as a function of mass. This was then fitted to the expected variation, equation 6.4; the variation with mass constrains n and the overall normalisation constrains M_{PL} . This study was essentially a theoretical study that included only the evaporation phase of the decay and accounted for statistical effects only. It therefore ignored the likely effects to the low energy spectrum from the initial parts of the decay (section 6.3.2), the effect of the remnant decay (section 6.3.5) and the recoil of the black hole (section 6.3.6) as well as any detector effects or underlying event. The analysis was also only valid for fixed temperature black holes (recall the study in section 6.3.6). Perhaps not surprisingly, it concluded that both n and M_{PL} could be well determined using only the energy distributions.

Another group has written a black hole generator and performed some initial studies with it [159]. Whilst they mainly focused on the details of their generator, they also suggested that the cross section could be used to determine M_{PL} [§]. This technique is discussed further in section 6.13.

[§]They have used the same convention as used here for M_{PL} .

A. Sabetfakhri has studied [3] various kinematic distributions and event shape variables including the energy spectra of electrons and photons. He concluded that there is a small variation with the number of dimensions in these variables that, given the high statistics, would be measurable. Unfortunately, connecting these differences to the number of dimensions is extremely difficult due to the systematic effects and uncertainties. This was consistent with a similar study of a smaller set of distributions considered in [158].

What is clear from the previous studies is that early theoretical attempts to suggest analyses have been over-optimistic and when tested using experimental simulations have been found to have little success. The great uncertainty in the models used to simulate black hole decay has not helped the situation. In the following sections I present several studies which look at how the model parameters could be measured whilst trying to control the systematic uncertainties.

6.10 Correlation Analysis

An initial attempt to determine the model parameters was made along the same lines as that used by Dimopoulos and Landsberg [2] but modified for time varying temperature. In this case there is one $M_{\text{BH}}-T_{\text{H}}$ point per emission, rather than just one per event. This method therefore tries to make maximum use of the event by fully reconstructing it including the order of the emissions. The method is:

1. Reconstruct the black hole from all the particles in the event.
2. Determine the first (next) particle to be emitted.
3. Use the measured properties of this particle to determine the temperature of the black hole.
4. Record this mass–temperature point.
5. Reconstruct the black hole for the next stage using all the particles except for those that have been emitted.
6. Repeat steps 2–5 until there are no particles left.

Exploring Higher Dimensional Black Holes

There are two key parts to this algorithm: determining the order of the emitted particles and using the particle properties to determine the temperature. It was hoped that a method might be found which when averaged over many events would give the expected $M_{\text{BH}}-T_{\text{H}}$ relation.

Since the black hole gets hotter as it decays, the average energy of particles emitted increases. Therefore the order was determined by assuming that the softest particles were emitted first.

The most probably energy for emission is proportional to the temperature, therefore $p_T \propto T_{\text{H}}$. Since the expected $M_{\text{BH}}-T_{\text{H}}$ relationship is a power law, the $M_{\text{BH}}-p_T$ relationship should have the same n dependence and has been plotted in the correlation plots.

The result of applying this method to fast simulation data is shown in figure 6.17. The shape of the graph is very different from that expected, but does show some separation between different numbers of dimensions. Unfortunately this technique is subject to many systematics which means that extracting the number of dimensions is extremely difficult to do.




Figure 6.17: Correlation plots (see text) for $n = 2, 4,$ and 6 for 8 TeV black holes.

To explain the difference between the expected result and the actual result, generator level plots are presented in figure 6.18. These plots show the results of using the algorithm above on the

generator level particles from the black hole decay^h with the following differences:

- The particles are chosen in the correct order, *i.e.* in the order the generator chose them.
- The energy of the particles is given in the rest frame of the black hole.

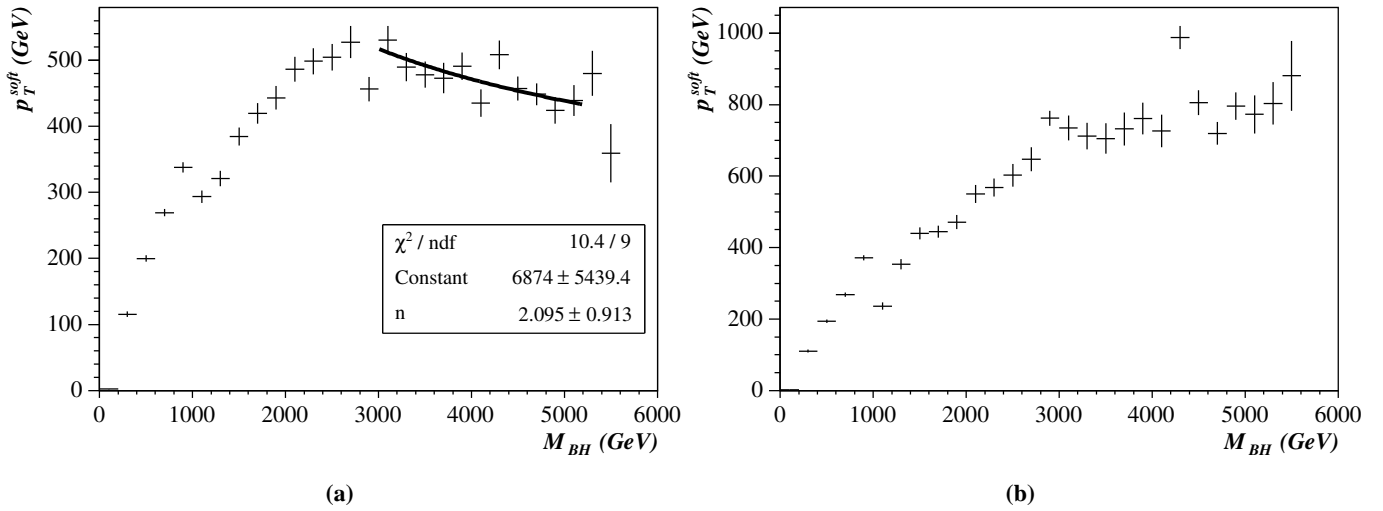


Figure 6.18: The average p_T against black hole mass in the rest frame of the black hole for the generator level decay particles for 5 TeV black holes from 1000 events. (a) $n = 2$. Equation 6.5 is fitted to the right part of the graph and (b) $n = 4$.

As can be seen from figure 6.18a, for $n = 2$ at high masses there is the expected relation (see equation 6.5) between mass and energy for high masses. At around 3 TeV the relationship breaks down which is where the energy is cut off by the kinematic limit (see section 6.3.4). Unfortunately, the $n = 4$ plot in figure 6.18 shows that there is no region that is not affected by the kinematic limit. Therefore it is not possible to fit the distribution to extract the number of dimensions.

One particularly problematic aspect of this method is that if one particle is placed in the incorrect order, then all the later emissions will be incorrect. Later emissions thus have larger errors than earlier ones.

Ultimately, this method has inherent flaws that mean that it is not successful in measuring n .

^hAlthough the data were generated with a range of black hole masses, only masses less than 5.2 TeV are included.

6.11 Analysis Strategy

At this point, it is useful to review the reasons why the correlation analysis failed and why the method of Dimopoulos and Landsberg is so susceptible to systematic effects such as the one presented in section 6.3.6. Both analyses attempt to use the black hole temperature and mass to determine both n and M_{PL} . In particular, the n dependence comes through equation 6.5, however this dependence is very weak. Figure 6.19 shows a plot of equation 6.5 for different n with the normalisation fixed at $M_{\text{BH}} = 5 \text{ TeV}$ to 200 GeV — a typical value. The variation over a 5 TeV black hole mass range is only about 40 GeV for $n = 2$ and 15 GeV for $n = 6$. This implies that the temperature must be measured very precisely with excellent control on the systematics.




Figure 6.19: The variation of temperature with black hole mass with fixed normalisation at $M_{\text{BH}} = 5 \text{ TeV}$. See equation 6.5 and text.

From this plot we can understand why time variation had the effect it did in section 6.3.6 and figure 6.3. A time varying temperature will give an energy distribution with a peak at a slightly higher energy than for the non-time varying model since lighter black holes are hotter. This effect will be larger for more massive black holes because of their higher multiplicity and since they cover a larger range of black hole masses. This effect may not be large. Judging from figure 6.19 the relative shift at 10 TeV compared to 5 TeV in the peak position of the energy distribution would be about 30 GeV (note that the peak energy is about $2T_{\text{H}}$ for fixed temperature) to give a shift from

$n = 2$ to $n = 4$.

I conclude that using the $T_{\text{H}}-M_{\text{BH}}$ relation to determine n is unlikely to be successful unless a better specified model is provided. Even then, this approach will only have success with leptons and photons and will require very high statistics and good control of experimental systematics. This rules out analyses along the lines of those presented in [2] and section 6.10.

The alternative to using the $T_{\text{H}}-M_{\text{BH}}$ relation is to measure the normalisation of the temperature (which after all, doesn't change much with mass!) and then to make a second independent measurement. The next sections discuss ways of measuring the overall temperature in a more model independent way, then look at using the cross section as a second measurement. Finally the measurement of temperature and cross section are brought together to determine n and M_{PL} for a test case.

6.12 Measurement of the black hole temperature

6.12.1 Kinematic limit

In this section a new idea is presented that is strongly dependent on the initial temperature of a black hole, but is valid for many different models.

If a particle is emitted with an energy close to the kinematic limit (i.e. $E \sim M_{\text{BH}}/2$), then that particle is probably the first to be emitted: future emissions of this energy will not be possible since, due to the lower black hole mass, they would be kinematically forbidden. In particular, it is possible to measure the fraction of events, p , where the highest energy particle has an energy, $E_{\text{max}} > E_{\text{cut}}$ where $E_{\text{cut}} = M_{\text{BH}}/2 - E_d$ and E_d defines an acceptably small energy range close to the kinematic limit and should be small compared to M_{BH} . The probability of the first emission being greater than E_{cut} can also be calculated from integrating the Planck spectrum (equation 6.6) thus there is a direct connection between the experimental measurement and theory. It should be noted that $p \rightarrow 0$ as $M_{\text{BH}} \rightarrow \infty$ for any fixed E_d . These equations are therefore only valid when p is small.

This method has many advantages: since it deals with the first emission, the probability of $E_{max} > E_{cut}$ will be the same regardless of whether the black hole temperature is time-varying or not. Also, since we restrict ourselves to black hole with masses much larger than the Planck mass, this measurement will not be affected by Planck scale effects or the remnant decay (unless the total multiplicity is very low). Any effects which modify the low energy part of the spectrum should also have no effect. Finally, the black hole boost can be taken into account by determining E_{max} in the black hole rest frame.

6.12.1.1 Bounds on systematic uncertainty

This technique is however strongly affected by the uncertainty in dealing with the kinematic limit (see section 6.3.4). The Planck spectrum, $P(E)$, is shown schematically in figure 6.20 near the kinematic limit. Approximate upper and lower limits can be placed on p . The upper bound on p can be found by integrating from E_{cut} up to infinity. This corresponds to assuming that real distribution will have all the probability above the kinematic limit shifted to just below that limit. The upper bound is therefore

$$p_{upper} = k \int_{E_{cut}}^{\infty} P(E) dE. \quad (6.9)$$

The lower bound can be calculated by making the opposite assumption: that the spectrum simply cuts off at the kinematic limit. There would need to be some increase in the normalisation or a slight change of shape at lower energies. This lower bound is

$$p_{lower} = k \int_{E_{cut}}^{M_{BH}/2} P(E) dE. \quad (6.10)$$

A third, more realistic, possibility is shown as the dot-dashed line in figure 6.20. That is, that near the kinematic limit, the distribution drops suddenly to zero. This is accompanied by an increased normalisation, as shown. I have assumed that E_{cut} is set low enough for this possibility not to give a prediction for p significantly below the lower bound quoted above. This is equivalent to stating that the dip near the kinematic limit is entirely cancelled out by the increase in




Figure 6.20: A schematic showing the Planck spectrum (solid line) near the kinematic limit ($M_{\text{BH}}/2$). One possibility for how the spectrum could be modified is shown as the dot-dashed line.

normalisation. A further study could investigate how valid this assumption is here given some model for modifying the Planck spectrum near the kinematic limit, for instance, a calculation of the back reaction from the metric during the black hole emission. In the calculations used later on of p_{upper} and p_{lower} , the grey-body factors and the difference between fermions and bosons have been ignored as those effects are small at high energies.

Figure 6.21 shows the upper and lower limits as a function of black hole mass for $M_{\text{PL}}=1$ TeV. As can be seen, there is significant separation between the bands, particularly for low n . At higher n , the bands start to overlap and high statistics at high M_{BH} would be necessary to distinguish different numbers of dimensions.

6.12.1.2 Initial soft emission

For $E_d > 0$, there is the possibility that the black hole will emit a low energy particle before emitting a high energy particle that gives $E_{\text{max}} > E_{\text{cut}}$. The probability of such an initial soft emission increases with E_d . Whilst this is not in itself a problem, it does mean that the time varying and non-time varying cases become more different. It also means that the black hole is




Figure 6.21: Plot of p_{upper} and p_{lower} for $M_{PL}=1$ TeV

less likely to be in the rest frame measured. In the calculation of the upper and lower bounds, I have included a corrective factor for the emission of a first soft particle. This was calculated as:

$$k_{cor} = k \int_0^{xE_{cut}} P(E)dE \quad (6.11)$$

giving $p_{cor} = (1 + k_{cor})p$ where x is a parameter that can be chosen. An initial soft emission can occur up to an energy of $2E_d$ and still allow the next emission to pass E_{cut} . However, taking into account the boost of the black hole, it is possible for an initial emission of *any* energy to be followed by an emission which passes the cut. The probability of this drops steeply after $2E_d$ since it requires that the particle be emitted close to the direction of the black hole boost. For the lower limit, $x = 2$ was chosen and for the upper limit, $x = 3$.

6.12.1.3 Experimental aspects

It is clear that the smaller E_d , the more the upper and lower limits separate. In the rest of this study I have set $E_d = 400$ GeV. This has been chosen somewhat arbitrarily and a more detailed study could investigate the best method for setting it. My initial studies suggest that the optimum setting is somewhat n -dependent. Nevertheless, this setting gives a large enough window for reasonable

statistics to be available and give typical values of p that are not too large. Also, the temperature change between the initial mass and the mass after a soft emission for the lightest black holes considered here (5 TeV) is limited to 10%. This means that results should be applicable to both time-varying and non-time-varying scenarios. The separation between the upper and lower limits is also not too great with this value which gives a reasonable chance for a measurement of the temperature. Finally, the probability of an initial emission can be rather large but, as mentioned above, this is not a problem so long as the temperature does not change much.

Mass reconstruction is important as is clear from the steeply falling shape of the figure 6.21. Consequently the cuts used to determine the mass resolution in section 6.8 have been applied here. Additionally, a cut to reduce low effective multiplicity events was also added. Events were rejected if the three highest energy particles accounted for more than 95% of the total energy in the black hole rest frame. Additionally, events which still pass the cuts, but have a very poorly measured mass will affect the measurement of p . In particular, since the mass resolution tails are quite large (see section 6.8) and the cross section is steeply falling, a significant proportion of high mass black holes will in fact be low mass black holes that have been poorly reconstructed. Since these events cannot possibly pass the cut, they suppress p at high masses. A simple estimate was made of the level of this effect from the tails of the mass resolution in figure 6.15 and the approximate form of the cross section. Consequently, a scale factor of 1.1 was applied to measurements of p . This is, of course, a rather simple correction factor and it probably under-estimates the impact of this effect at high mass.

One experimentally tricky aspect of this measurement is that since the black hole mass is measured by adding all the particles in the event, the maximum energy in the reconstructed black hole rest frame must be less than $M_{\text{BH}}/2$. This introduces a bias towards low energies in the measurement of E_{max} and thus reduces the measurement of p . This has been corrected for by increasing E_d by 100 GeV which is an estimate for the average under-measurement of E_{max} . Figure 6.22a shows the resolution for measuring E_{max} . It should be noted that not boosting into the black hole rest frame gives a very significant over-measurement of p . Figure 6.22b shows the same measure-




Figure 6.22: Bias in measurement of E_{max} for (a) 10,000 fast simulation events and (b) both full simulation data sets.

ment using both sets of full simulation data. The shape is very similar to the fast simulation plot, including the shift to lower energies. This important correction is therefore corroborated by the full simulation.

6.12.1.4 Results

Here I present the results of this analysis applied to several different models and with different numbers of dimensions. Figure 6.23 shows the results for four different models all with $n = 4$. Models which have the kinematic cut off (see section 6.3.4) should be consistent with the lower limit; this is the case for sub-figures a, b and d. However, if the kinematic limit is on and the remnant decay is set to 2-body (sub-figure c), the data is expected to be consistent with the upper limit. This is because if an energy is chosen in the forbidden region (which is the additional region included in the upper limit integral), it will cause the decay to be terminated and the black hole to split into two. One of the two remnant decay particles must pass the cut. As can be seen from the figure, the plots do agree with these expectations.

The results of this analysis as applied to the test case for different values of n are shown in figure 6.24. This emphasises that this technique is sensitive to n whilst being largely model

6.12 Measurement of the black hole temperature

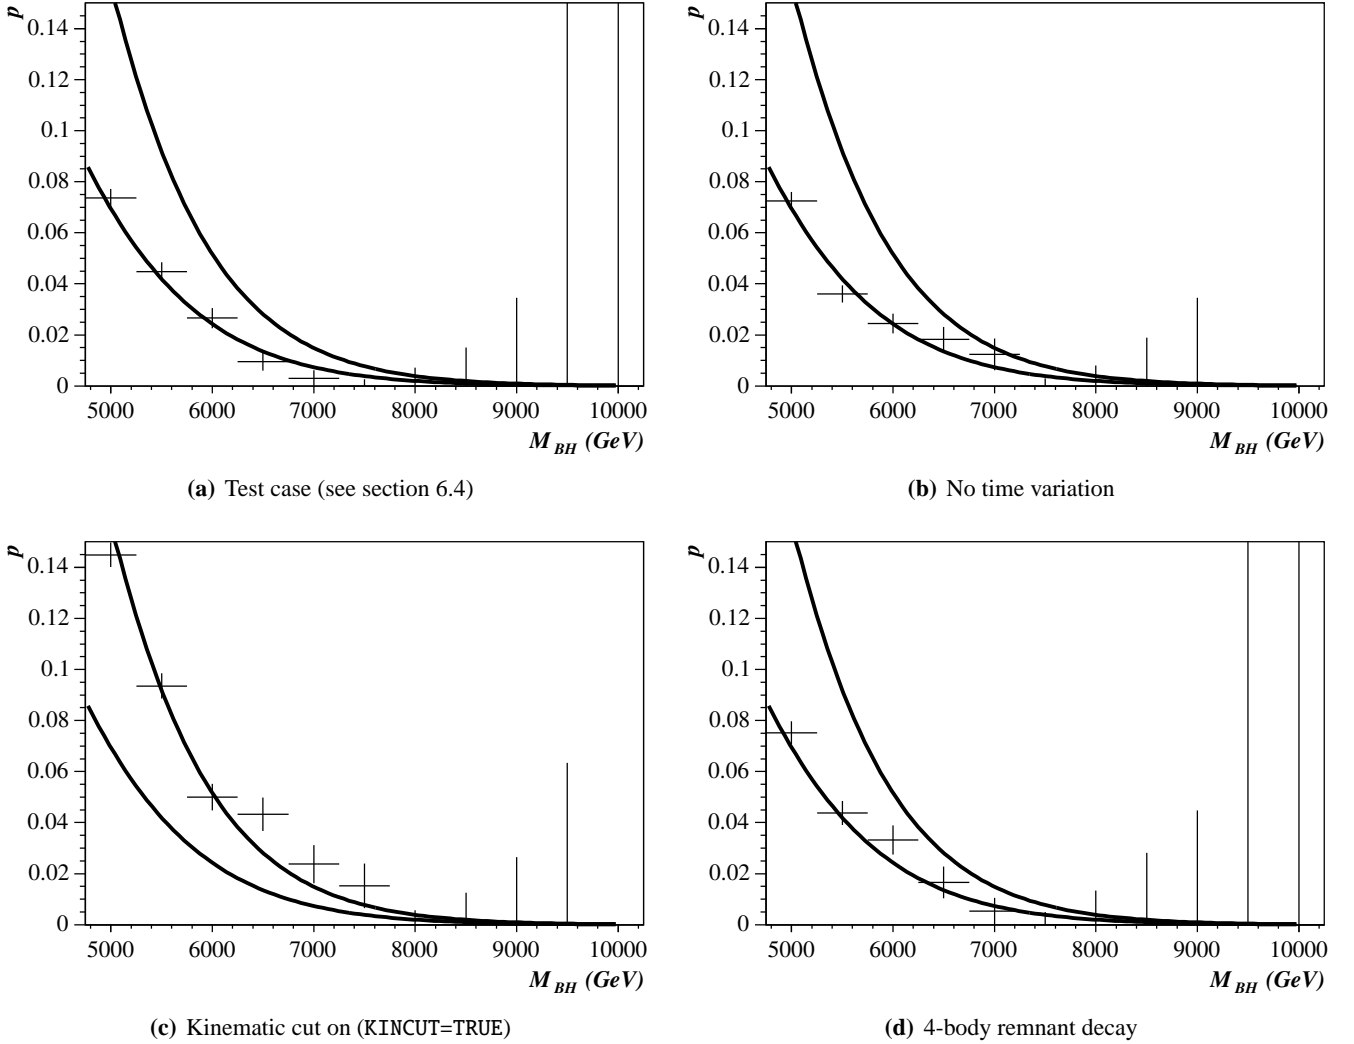


Figure 6.23: Fraction of events passing the cut, p , as a function of M_{BH} for different models all with $n = 4$. Upper and lower limits for $n = 4$ are also shown.

independent. Note that at high n , the data starts to drop below the lower limit. This is due to the high temperature here ($T_H \sim 470$ GeV for $n = 5$) which significantly reduces the multiplicity. This result suggests that this analysis has an upper limit of validity in the region of ~ 450 – 500 GeV given a Planck mass of 1 TeV. In any case, at this temperature the semi-classical limit has broken down. Also, from figure 6.24a, it can be seen that at low temperatures, this analysis will not measure the temperature (unless lighter black holes are seen, or very high statistics are available),

Exploring Higher Dimensional Black Holes

but would instead place an upper limit on it. This may be enough to constrain n , or alternatively, the energy distribution analysis presented in section 6.12.2 could be used here.

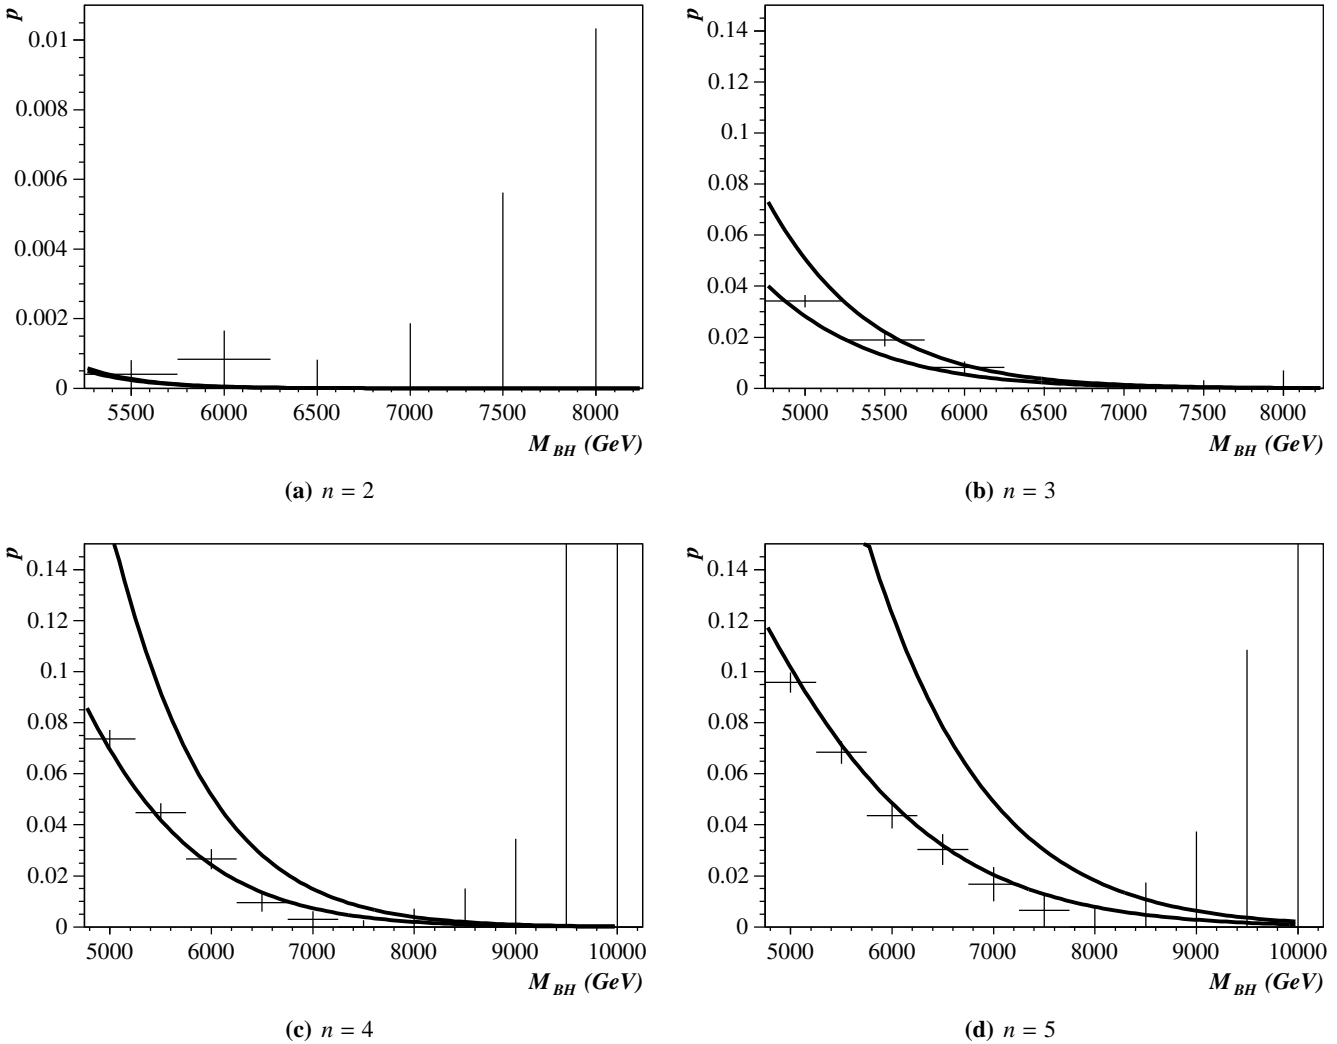


Figure 6.24: Fraction of events passing the cut, p , as a function of M_{BH} for different values of n for the test case. Appropriate upper and lower limits are shown.

This analysis has been applied to the full simulation data sets, the results of which are in figure 6.25. Unsurprisingly the limited statistics mean the results have limited value. However, the $n = 4$ plot seems to correspond to the expected distribution and the change between the two plots is encouraging.

6.12 Measurement of the black hole temperature




Figure 6.25: Fraction of events passing the cut, p , as a function of M_{BH} for full simulation data; (a) for $n = 2$ and (b) for $n = 4$. Appropriate upper and lower limits are shown.

It is of course possible to convert the measurement of p into a measurement of the black hole temperature. Doing this requires assuming a model and it also is dependent of the black hole mass at each point. This has been done for a 30 fb^{-1} data sample for the test case with $n = 4$ in figure 6.26. The data points correspond to choosing the lower limit model. The band represents the possible range of data points that would be obtained by making different model choices and includes a systematic from the bias on the measurement of the black hole mass of $\pm 200 \text{ GeV}$. The true temperature–mass relation is also shown on the plot. Note that the lower edge of the systematic band does not have the correct temperature–mass relationship: the temperature should certainly *decrease* with black hole mass. This could enable the model to be restricted giving some information about the shape of the energy distribution near the kinematic limit.

Note that on the left of figure 6.26, there are high statistics but the systematic error is large whereas on the right, the statistical errors are worse, but the systematic error is better. In the middle, the total error is smallest and using this I have estimated that the temperature is best determined at a mass of 7 TeV and is $340 \pm 30 \text{ GeV}$. The actual temperature at that mass is 355 GeV . This is an important demonstration: for the first time the black hole temperature has been accu-




Figure 6.26: Results for test case, $n = 4$

rately determined in a full event generation and detector simulation study with at least some of the approximations and uncertainties accounted for. This result will be discussed in section 6.12.3, together with the results of the next method of determining the black hole temperature.

6.12.2 Energy Distributions

Whilst the method in [2] was not successful at measuring the *variation* in temperature with mass, it may be successful at measuring the overall scale. This section therefore attempts to repeat their analysis but adding event simulation and detector response. I have also considered the effects of some different model assumptions.

The energy distributions are experimentally simpler and offer higher statistics compared to, for instance, the kinematic limit analysis. However they are harder to interpret theoretically and careful understanding of the systematics involved is necessary. The study presented here is a quick attempt to understand the areas of likely validity of the energy distributions and needs more detailed study to fully understand all of the effects. Nonetheless, there are several interesting

features.

The method is to plot the energy distributions of electrons, muons and photons separately. Once the distributions have been plotted, the temperature must be determined. There are two ways of doing this. The peak of the distribution can be determined or the whole shape can be fitted to some expected form (such as the Planck distribution). Since the statistics will be good, considering the three distributions separately will enable cross-checks to be made to ensure that the data is consistent. In any case, photons must be considered separately since the Planck distribution and grey-body factors differ significantly for gauge bosons.




Figure 6.27: Electron energy distribution for the test case and $n = 2$. The lines are Planck distributions with temperatures of 160, 180 and 200 GeV.

Once black hole events have been selected, it is not clear which cuts should be employed to provide the best energy distributions. Clearly the electrons, muons and photons should be well isolated, so an isolation cut requiring $\Delta R_{jet} > 0.7$ (where $\Delta R = \sqrt{\Delta\eta^2 + \Delta\phi^2}$ and the minimum value for all jets has been considered) has been employed. A cut on $|\eta|$ has been considered in order to reduce the underlying event, however it had little impact (typically 2–5 GeV) on the temperature measurements so has not been applied to these data. Although a range of black hole masses will be produced and thus the distributions for several different temperatures are overlapped, the temperature does not vary much with mass and the dominant contribution will come from the

lowest mass black holes. Therefore expected temperatures have been calculated for 5 TeV black holes.

Ignoring grey-body factors, the peak of the Planck distribution is at $\approx 2.2T_H$ for fermions and at $\approx 1.6T_H$ for bosons. However, actually observing the peak of the distribution may be difficult since there are likely to be many low energy electrons, muons and photons from the underlying event, pile-up and from the non-evaporation phases of the black hole decay (see section 6.3.2). Fortunately it appears that the shape of the high energy tail constrains the temperature of the black hole. This is illustrated in figure 6.27 which shows the electron energy distribution for the test case with $n = 2$ and the Planck distribution for several values of T_H . The temperature has therefore been obtained by fitting the Planck spectrum to the energy distribution for energies above 200 GeV. Energies below this were excluded as they tend to be affected by large numbers of low energy particles from other parts of the event. This low energy cut has been varied over the range 0–1000 GeV and the variation in the fitted temperature is approximately ± 20 GeV.

Table 6.4 shows the results of this procedure for the test case with varying n and figure 6.28 shows the distributions and fits for $n = 2$. The electron and muon results are reasonably consistent with each other and the actual temperature for $n = 2$. For $n > 2$ however, the measured temperatures begin to diverge from the expected temperatures. Note that this is unlikely just to be due to the time variation of the black hole temperature, since this would tend to give a higher average temperature than in the non-time varying case (which using the Planck distribution assumes); rather it is due to the model variation.

It is also notable that the photon distribution always gives a significantly higher temperature than the electron and muon distributions. This is a clear indication of the effect of the grey-body factors. One of the most notable features of the grey-body factors is that they tend to 0 at low energy for gauge bosons for all n (see figure 4.6 of [12]). This can cause a significant increase in the position of the peak of the energy spectrum for low temperatures. An observation of such a difference between the e/μ and g distributions would be an excellent indication of the effect of the grey-body factors and would allow them to be estimated for the spin 1 case.

6.12 Measurement of the black hole temperature

n	Actual T_H (GeV)	e temperature (GeV)	μ temperature (GeV)	γ temperature (GeV)
2	179	183	191	231
3	282	234	249	298
4	380	271	286	339
5	470	287	308	356
6	553	304	318	370

Table 6.4: Temperatures for the test case with varying n for the electron, muon and photon distributions. Approximately 1 fb^{-1} of data were used for each case.




Figure 6.28: Energy distributions for the test case and $n = 2$ with fits to the Planck spectrum; for (a) electrons, (b) muon and (c) photons.

Another important observation is that the low energy part of the photon spectrum is much clearer than the equivalent parts of the electron and muon distributions. This allows the peak of the spectrum to be unambiguously seen.

The results in table 6.4 suggest that this method can only work for lower temperatures. I have therefore investigated this method applied to several different models for $n = 2$. These results are presented in table 6.5.

The variation in these results suggests that the temperature can be measured in a reasonably model independent way using this technique with a systematic error of about $\pm 20 \text{ GeV}$. The effect of the grey-body factors can be seen clearly in the increased temperature (and non-Planckian shape) of the photon energy distribution. Results with a range of temperatures suggest that this

Exploring Higher Dimensional Black Holes

Model	e temperature (GeV)	μ temperature (GeV)	g temperature (GeV)
Test case	183	189	231
No time variation	165	174	214
4 body remnant decay	172	182	220
Kinematic cut on	196	204	264

Table 6.5: Results for different models with $n = 2$. The models considered are the same as those in section 6.12.1.4.

technique would be successful up to a temperature of 200–250 GeV although more studies would be needed to confirm this. It is also likely that at lower temperatures it would become harder to determine the temperature due to low energy particles from the underlying event or other parts of the black hole decay, but again, this would need further study.

This method was also applied to the $n = 2$ full simulation data. However, the poor statistics do not give a very meaningful result as typical fitting errors were ~ 50 GeV (although the results were consistent with the fast simulation results within this error). Also, the full simulation seemed to suffer a lower efficiency (by a factor of 2) compared to the fast simulation. Whilst some reduction is expected (since electron, muon and photon efficiencies are not included in ATLFast), this is surprising. Algorithms for brems recovery and conversion have not been applied, so there would be some improvement here, but the greatest reduction comes from the isolation requirement. Perhaps it would be necessary to use a more complex isolation cut such as requiring that there be no more than a certain amount of energy near the particle.

6.12.3 Discussion

Two analyses have been described that were able to measure the black hole temperature: the energy distribution analysis below about 200 GeV and the kinematic limit analysis at 200–450 GeV. Although the two analyses took very different approaches they are actually not as different as might appear. Both are actually most sensitive to the high part of the energy distributions. The energy distribution analysis is experimentally robust, but suffers from model variation and the in-

6.13 Measurement of the parton-level cross section

ability to see the peak of the spectrum. The kinematic limit analysis suffers from requiring a good measurement of the black hole mass (this would be difficult if bulk emission or stable remnants had been included) and from assuming that the back reaction is not important. However it is robust against other model changes and can reach higher temperatures.

The kinematic limit analysis was motivated by searching for a technique in which the time varying and fixed temperature models were the same. However, perhaps it is better to accept this as a source of uncertainty (since it is reasonably easy to work in either model) and instead to require robustness against the less determinable assumptions such as the back reaction from the metric or systematics in the black hole mass measurement.

These two analyses can then point the way forward. In order to achieve the above aims, it might be a good idea to investigate the moderately high part of the energy distributions, say between 800–1500 GeV. This should be low enough for the kinematic limit (and thus the black hole mass) and back reaction to be unimportant whilst being high enough to be robust against most remnant and low energy effects. The method used in the kinematic limit analysis of considering integrals could also be useful for reducing any effects that modify the shape of the distribution. This would be an interesting avenue to explore in future studies.

6.13 Measurement of the parton-level cross section

As suggested previously in [159], the cross section for black hole production is strongly dependent on M_{PL} and only very weakly dependent on n^{\dagger} . The parton-level cross section is shown in figure 6.29 which shows this and includes the form factors calculated in [130].

Assuming that the Planck mass is not too large, there will be ample statistics at the LHC to measure the cross section accurately. The accuracy of such a measurement will therefore depend predominantly on experimental effects such as the luminosity and efficiency and also on how well the PDFs can be determined. In addition, there may be theoretical uncertainties which affect

[†]This is only true in the Dimopoulos and Landsberg convention used here. For instance, in the ADD convention, the *temperature* is approximately independent of n .




Figure 6.29: The parton-level cross section as a function of M_{PL} for different values of n . Form factors are included.

the cross section. ATLAS expects to measure the luminosity to better than 5% and computer simulations should be able to estimate the efficiency well, so experimental errors are likely to be of the order of 7%.

Measurement of the PDFs could however be difficult. Black holes are formed mainly through high Q^2 , large x processes and in these regions, ATLAS would normally expect to measure the PDFs to about 5–10%. However if TeV scale gravity is observed then it is hard to imagine that the evolution equations will not be significantly modified. Also, direct measurements of the PDFs, for instance by comparison to Z^0 production, will not be possible at these scales since the production of black holes means the end of short-scale perturbative physics. For these studies I have conservatively estimated that the parton-level cross section could be measured to 20% which corresponds to a measurement of M_{PL} of about 10%. It is of course possible that measurements of other processes could also enable an independent fix on M_{PL} .

6.14 Determination of the model parameters

Given a measurement of the parton-level cross section and the temperature it is reasonably straightforward to convert these into measurements of M_{PL} and n . This has been done using the temperature measurement obtained in section 6.12.1 ($T_{\text{H}} = 340 \pm 30$ GeV at $M_{\text{BH}} = 7$ TeV) and taking the correct parton-level cross-section together with a 20% error (see section 6.13). The result is the skewed Gaussian shape in figure 6.30 which gives $M_{\text{PL}} = 1029 \pm 150$ GeV and $n = 3.8 \pm 0.75$. If the cross section error could be reduced to 10%, the figures would be $M_{\text{PL}} = 1029 \pm 70$ GeV and $n = 3.8 \pm 0.6$ which confirms that the cross section dominates the determination of M_{PL} .




Figure 6.30: Results for test case, $n = 4$

6.15 Conclusions

- The theory of black hole production and decay in ADD models has been briefly introduced and a number of the theoretical issues that affect these processes have been discussed.

Exploring Higher Dimensional Black Holes

- The ability of the ATLAS detector to measure the black hole mass was investigated making the assumption that all of the energy is emitted onto our brane. In this case, the mass resolution was found to be good: $\sim 3\text{--}4\%$.
- Three analyses have been presented which attempted to determine the temperature of the black holes. Of these, the correlation analysis failed, the energy distribution analysis was found likely to be valid at low temperatures and the kinematic limit analysis was found to be sensitive to the temperature in the range $200\text{--}450$ GeV. The kinematic limit analysis was found to be robust against a number of different models and included estimates of some of the systematic effects.
- Modifications to these analyses to deal with some of the assumptions that they still needed to make have also been discussed.
- I have also considered how the cross section information could be used to find the Planck mass. I then showed that the temperature and cross section measurements could be combined to measure both n and M_{PL} . For the test model with $n = 4$ and $M_{\text{PL}} = 1$ TeV and 30 fb^{-1} of data, the simulated measurements were: $n = 3.8 \pm 0.75$ and $M_{\text{PL}} = 1029 \pm 70$ GeV with strongly correlated errors between the measurements.
- In addition, the ATLAS full simulation software has been used to check the validity of the above results. Before this could be done, several corrections were made to the reconstruction and the overall performance was investigated. It was found that electron, photon and muon reconstructions were generally good and in agreement with the expectations as published in the TDR. The jet reconstruction was less good with significant effort needed to improve jet linearity as a function of p_T and to remove the strong $|\eta|$ dependence on the reconstructed jet energy. The \cancel{p}_T resolution was found to be significantly different from that expected being about 5–6 times worse than expected. This had a negative impact on the black hole mass reconstruction and thus potentially on the other analyses (although the statistics were not good enough to determine this).

Although this chapter has presented many ideas and studies that have advanced our understanding of how black hole properties could be measured, it is clear that this remains an experimentally challenging problem.

Summary

Exactly. So, logically...

If... she... weighs... the same as a duck,... she's made of wood.

And therefore?

A witch!

Whilst the Standard Model of particle physics has been immensely successful, there are strong theoretical reasons for believing that there should be new physics at the TeV scale. Some of the most interesting and exciting possibilities are extra dimensions, for which a number of different models have been developed.

Ultimately, such models must be tested experimentally if they are to have any scientific validity so the construction of a TeV scale collider and detector is of the highest priority. To this end, I have aided in the construction of the ATLAS detector being built at CERN by writing part of the SctRodDaq software to calibrate and characterise SCT detector modules during macro-assembly. I have shown that this software is fit-for-purpose by comparing its algorithms to previous calibration software and by conducting performance tests. SctRodDaq is now in use at the macro-assembly sites and at CERN.

In this thesis I have presented two studies of extra-dimensional models that have both aimed to determine how well ATLAS could measure the model parameters. The first of these studies, of graviton decays to heavy vector bosons, showed that the $G \rightarrow W^+W^-$ and $G \rightarrow Z^0Z^0$ couplings could be measured over much of the allowed parameter space in the basic RS model. By combining these measurements with those from other channels, it would be possible to test the universality of the graviton's couplings, which is a model independent feature.

In the final study, I have investigated the decay of microscopic black holes in ATLAS. This exciting possibility is experimentally extremely challenging due to the complexity of the decay

Summary

and the lack of theoretical understanding. I have considered these theoretical uncertainties and investigated their affect on several analyses. I have also introduced a new technique that is able to account for some of these uncertainties and demonstrated that this technique would be able to determine the Planck mass to about 10% and the number of dimensions to ± 0.75 for a test case with four extra dimensions.

*What does the worker gain from his toil?
I have seen the burden God has laid on men.
He has made everything beautiful in its time.
He has also set eternity in the hearts of men;
yet they cannot fathom what God
has done from beginning to end.
I know that there is nothing better for men
to do good while they live.
That everyone may eat and drink,
and find satisfaction in all his toil
— this is the gift of God.*

ABBREVIATIONS

ADD	N. Arkani-Hamed, S. Dimopoulos and G. R. Dvali
API	Application Programming Interface
ASIC	Application Specific Integrated Circuit
ATLAS	A Toroidal LHC ApparatuS
BOC	Back Of Crate card
CERN	The European Organization for Nuclear Research
CMS	Compact Muon Solenoid
CORBA	Common Object Request Broker Architecture
CSC	Cathode Strip Chamber
DAC	Digital to Analogue Converter
DAQ	Data AcQuisition
DCS	Detector Control System
DSP	Digital Signal Processor
ECAL	Electromagnetic CALorimeter
ENC	Equivalent Noise Charge
FCAL	Forward CALorimeter
FIFO	First-In-First-Out
FILO	First-In-Last-Out
FPGA	Field Programmable Gate Array
GUI	Graphical User Interface
HCAL	Hadronic CALorimeter
HEP	High-Energy Physics

Abbreviations

IDL	Interface Definition Language
ILU	Inter-Language Unification system
IPC	Inter-Process Communication
IS	Information Service
JHEP	Journal of High Energy Physics
LAr	lead-Liquid Argon
LHC	Large Hadron Collider
LHCb	Large Hadron Collider beauty experiment
MDT	Monitored Drift Tube
MIP	Minimally-Ionising Particle
MRS	Message Reporting Service
NFS	Network File System
OMG	Object Management Group
PDF	Parton Distribution Function
p_T	Transverse momentum
\cancel{p}_T	Missing transverse momentum
QCD	Quantum ChromoDynamics
QED	Quantum ElectroDynamics
RAM	Random Access Memory
RCC	ROD Crate Controller
ROD	ReadOut Driver
RoI	Regions of Interest
ROS	ReadOut System
RPC	Resistive Plate Chamber
RS	L. Randall and R. Sundrum
S-Link	S-Link
SBC	Single Board Computer

SCT	SemiConductor Tracker
TCC	Interface between the TIM and the TTC system
TDR	Technical Design Report
TGC	Thin Gap Chamber
TIM	TTC Information Module
TRT	Transition Radiation Tracker
TTC	Trigger, Timing and Control system
VME	Versa Module Europa
XML	eXtensible Markup Language

REFERENCES

URLs have been given for references where they are available. Many can be obtained through the CERN and ATLAS websites and to make the URLs more readable, the abbreviations [CERN] and [ATLAS] have been used, which stand for: <http://doc.cern.ch/archive/electronic> and <http://atlas.web.cern.ch/Atlas> respectively.

- [1] B. C. Allanach, K. Odagiri, M. J. Palmer, M. A. Parker and A. Sabetfakhri, *Exploring small extra dimensions at the Large Hadron Collider*, JHEP **12** (2002) 039, hep-ph/0211205.
- [2] S. Dimopoulos and G. Landsberg, *Black holes at the LHC*, Phys. Rev. Lett. **87** (2001) 161602, hep-ph/0106295.
- [3] C. M. Harris, M. J. Palmer, M. A. Parker, P. Richardson, A. Sabetfakhri and B. R. Webber, *Exploring higher dimensional black holes at the Large Hadron Collider*, hep-ph/0411022.
- [4] **Super-Kamiokande** Collaboration, Y. Fukuda *et. al.*, *Evidence for oscillation of atmospheric neutrinos*, Phys. Rev. Lett. **81** (1998) 1562–1567, hep-ex/9807003.
- [5] **SNO** Collaboration, Q. R. Ahmad *et. al.*, *Direct evidence for neutrino flavor transformation from neutral-current interactions in the Sudbury Neutrino Observatory*, Phys. Rev. Lett. **89** (2002) 011301, nucl-ex/0204008.
- [6] F. Englert and R. Brout, *Broken symmetry and the mass of gauge vector mesons*, Phys. Rev. Lett. **13** (1964) 321–322.

References

- [7] P. W. Higgs, *Broken symmetries and the masses of gauge bosons*, Phys. Rev. Lett. **13** (1964) 508–509.
- [8] P. W. Higgs, *Spontaneous symmetry breakdown without massless bosons*, Phys. Rev. **145** (1966) 1156–1163.
- [9] S. Eidelman, K. Hayes, K. Olive, M. Aguilar-Benitez, C. Amsler, D. Asner, K. Babu, R. Barnett, J. Beringer, P. Burchat, C. Carone, C. Caso, G. Conforto, O. Dahl, G. D’Ambrosio, M. Doser, J. Feng, T. Gherghetta, L. Gibbons, M. Goodman, C. Grab, D. Groom, A. Gurtu, K. Hagiwara, J. Hernández-Rey, K. Hikasa, K. Honscheid, H. Jawahery, C. Kolda, K. Y., M. Mangano, A. Manohar, J. March-Russell, A. Masoni, R. Miquel, K. Mönig, H. Murayama, K. Nakamura, S. Navas, L. Pape, C. Patrignani, A. Piepke, G. Raffelt, M. Roos, M. Tanabashi, J. Terning, N. Törnqvist, T. Trippe, P. Vogel, C. Wohl, R. Workman, W.-M. Yao, P. Zyla, B. Armstrong, P. Gee, G. Harper, K. Lugovsky, S. Lugovsky, V. Lugovsky, A. Rom, M. Artuso, E. Barberio, M. Battaglia, H. Bichsel, O. Biebel, P. Bloch, R. Cahn, D. Casper, A. Cattai, R. Chivukula, G. Cowan, T. Damour, K. Desler, M. Dobbs, M. Drees, A. Edwards, D. Edwards, V. Elvira, J. Erler, V. Ezhela, W. Fetscher, B. Fields, B. Foster, D. Froidevaux, M. Fukugita, T. Gaisser, L. Garren, H.-J. Gerber, G. Gerbier, F. Gilman, H. Haber, C. Hagmann, J. Hewett, I. Hinchliffe, C. Hogan, G. Höhler, P. Igo-Kemenes, J. Jackson, K. Johnson, D. Karlen, B. Kayser, D. Kirkby, S. Klein, K. Kleinknecht, I. Knowles, P. Kreitz, Y. Kuyanov, O. Lahav, P. Langacker, A. Little, L. Littenberg, D. Manley, A. Martin, M. Narain, P. Nason, Y. Nir, J. Peacock, H. Quinn, S. Raby, B. Ratcliff, E. Razuvaev, B. Renk, G. Rolandi, M. Ronan, L. Rosenberg, C. Sachrajda, Y. Sakai, A. Sanda, S. Sarkar, M. Schmitt, O. Schneider, D. Scott, W. Seligman, M. Shaevitz, T. Sjöstrand, G. Smoot, S. Spanier, H. Spieler, N. Spooner, M. Srednicki, A. Stahl, T. Stanev, M. Suzuki, N. Tkachenko, G. Trilling, G. Valencia, K. van Bibber, M. Vincter, D. Ward, B. Webber, M. Whalley, L. Wolfenstein, J. Womersley, C. Woody, O. Zenin and R.-Y. Zhu, *Review of Particle Physics*, Physics Letters B **592** (2004) 1,

<http://pdg.lbl.gov>.

- [10] **DØ** Collaboration, V. M. Abazov *et. al.*, *A precision measurement of the mass of the top quark*, *Nature* **429** (2004) 638–642, [hep-ex/0406031](#).
- [11] J. M. Butterworth, B. E. Cox and J. R. Forshaw, *WW scattering at the LHC*, *Phys. Rev.* **D65** (2002) 096014, [hep-ph/0201098](#).
- [12] C. M. Harris, *Physics Beyond the Standard Model: Exotic Leptons and Black Holes at Future Colliders*. PhD thesis, **University of Cambridge**, 2003.
- [13] N. Arkani-Hamed, S. Dimopoulos and G. R. Dvali, *The hierarchy problem and new dimensions at a millimeter*, *Phys. Lett.* **B429** (1998) 263–272, [hep-ph/9803315](#).
- [14] N. Arkani-Hamed, S. Dimopoulos and G. R. Dvali, *Phenomenology, astrophysics and cosmology of theories with sub-millimeter dimensions and TeV scale quantum gravity*, *Phys. Rev.* **D59** (1999) 086004, [hep-ph/9807344](#).
- [15] I. Antoniadis, N. Arkani-Hamed, S. Dimopoulos and G. R. Dvali, *New dimensions at a millimeter to a Fermi and superstrings at a TeV*, *Phys. Lett.* **B436** (1998) 257–263, [hep-ph/9804398](#).
- [16] S. B. Giddings and S. Thomas, *High energy colliders as black hole factories: The end of short distance physics*, *Phys. Rev.* **D65** (2002) 056010, [hep-ph/0106219](#).
- [17] Y. Uehara, *A mini-review of constraints on extra dimensions*, *Mod. Phys. Lett.* **A17** (2002) 1551–1558, [hep-ph/0203244](#).
- [18] **EOT-WASH** Group, E. G. Adelberger, *Sub-millimeter tests of the gravitational inverse square law*, [hep-ex/0202008](#).
- [19] C. D. Hoyle *et. al.*, *Sub-millimeter tests of the gravitational inverse-square law: A search for ‘large’ extra dimensions*, *Phys. Rev. Lett.* **86** (2001) 1418–1421, [hep-ph/0011014](#).

References

- [20] **L3** Collaboration, P. Achard *et. al.*, *Single- and multi-photon events with missing energy in e^+e^- collisions at LEP*, Phys. Lett. **B587** (2004) 16–32, hep-ex/0402002.
- [21] **LEP** Collaboration, M. Gataullin, *Searches for new physics in photonic final states at LEP*, Eur. Phys. J. **C33** (2004) s791–s793, hep-ex/0311014.
- [22] **DELPHI** Collaboration, J. Abdallah *et. al.*, *Photon events with missing energy in e^+e^- collisions at $s^{1/2} = 130$ GeV to 209 GeV*, hep-ex/0406019.
- [23] **ALEPH** Collaboration, A. Heister *et. al.*, *Single- and multi-photon production in e^+e^- collisions at $s^{1/2}$ up to 209 GeV*, Eur. Phys. J. **C28** (2003) 1–13.
- [24] **OPAL** Collaboration, G. Abbiendi *et. al.*, *Photonic events with missing energy in e^+e^- collisions at $s^{1/2} = 189$ GeV*, Eur. Phys. J. **C18** (2000) 253–272, hep-ex/0005002.
- [25] **CDF** Collaboration, D. Acosta, *Search for Kaluza-Klein graviton emission in $p\bar{p}$ collisions at $s^{1/2} = 1.8$ TeV using the missing energy signature*, Phys. Rev. Lett. **92** (2004) 121802, hep-ex/0309051.
- [26] **DØ** Collaboration, V. M. Abazov *et. al.*, *Search for large extra dimensions in the monojet + missing- E_T channel at DØ*, Phys. Rev. Lett. **90** (2003) 251802, hep-ex/0302014.
- [27] G. Landsberg, *Extra dimensions and more. . .*, hep-ex/0105039.
- [28] L. A. Anchordoqui, J. L. Feng, H. Goldberg and A. D. Shapere, *Updated limits on TeV-scale gravity from absence of neutrino cosmic ray showers mediated by black holes*, Phys. Rev. **D68** (2003) 104025, hep-ph/0307228.
- [29] E.-J. Ahn, M. Cavaglia and A. V. Olinto, *Uncertainties in limits on TeV-gravity from neutrino induced air showers*, hep-ph/0312249.
- [30] S. Cullen and M. Perelstein, *SN1987A constraints on large compact dimensions*, Phys. Rev. Lett. **83** (1999) 268–271, hep-ph/9903422.

-
- [31] S. Hannestad and G. Raffelt, *New supernova limit on large extra dimensions*, Phys. Rev. Lett. **87** (2001) 051301, hep-ph/0103201.
- [32] S. Hannestad and G. G. Raffelt, *Stringent neutron-star limits on large extra dimensions*, Phys. Rev. Lett. **88** (2002) 071301, hep-ph/0110067.
- [33] M. Fairbairn, *Cosmological constraints on large extra dimensions*, Phys. Lett. **B508** (2001) 335–339, hep-ph/0101131.
- [34] L. Vacavant and I. Hinchliffe, *Signals of models with large extra dimensions in ATLAS*, J. Phys. **G27** (2001) 1839–1850.
- [35] R. N. Mohapatra, S. Nussinov and A. Perez-Lorenzana, *Large extra dimensions and decaying KK recurrences*, Phys. Rev. **D68** (2003) 116001, hep-ph/0308051.
- [36] G. D. Starkman, D. Stojkovic and M. Trodden, *Large extra dimensions and cosmological problems*, Phys. Rev. **D63** (2001) 103511, hep-th/0012226.
- [37] L. Randall and R. Sundrum, *A large mass hierarchy from a small extra dimension*, Phys. Rev. Lett. **83** (1999) 3370–3373, hep-ph/9905221.
- [38] H. Davoudiasl, J. L. Hewett and T. G. Rizzo, *Phenomenology of the Randall-Sundrum gauge hierarchy model*, Phys. Rev. Lett. **84** (2000) 2080, hep-ph/9909255.
- [39] A. Pomarol, *Gauge bosons in a five-dimensional theory with localized gravity*, Phys. Lett. **B486** (2000) 153–157, hep-ph/9911294.
- [40] H. Davoudiasl, J. L. Hewett and T. G. Rizzo, *Experimental probes of localized gravity: On and off the wall*, Phys. Rev. **D63** (2001) 075004, hep-ph/0006041.
- [41] J. Hewett and M. Spiropulu, *Particle physics probes of extra spacetime dimensions*, Ann. Rev. Nucl. Part. Sci. **52** (2002) 397–424, hep-ph/0205106.

References

- [42] **DØ** Collaboration, V. M. e. a. Abazov, *Search for Randall-Sundrum gravitons in the dielectron and diphoton channel with 200 pb^{-1} of data with the DØ detector*, DØ note: 4533-Conf, 2004.
- [43] For the **CDF** and **DØ** Collaborations, M. K. Unel, *Searches for extra dimensions at the Tevatron*. Presented at 15th Topical Conference on Hadron Collider Physics (HCP2004), East Lansing, Michigan, 14-18 Jun 2004.
- [44] **DØ** Collaboration, V. M. Abazov *et. al.*, *Search for heavy particles decaying into electron positron pairs in $p\bar{p}$ collisions*, Phys. Rev. Lett. **87** (2001) 061802, hep-ex/0102048.
- [45] **CDF** Collaboration, F. Abe *et. al.*, *Search for new gauge bosons decaying into dileptons in $p\bar{p}$ collisions at $s^{1/2} = 1.8\text{ TeV}$* , Phys. Rev. Lett. **79** (1997) 2192–2197.
- [46] B. C. Allanach, K. Odagiri, M. A. Parker and B. R. Webber, *Searching for narrow graviton resonances with the ATLAS detector at the Large Hadron Collider*, JHEP **09** (2000) 019, hep-ph/0006114.
- [47] **ATLAS** Collaboration, *Detector and physics performance, Volume I*, 1999.
ATLAS-TDR-14,
[CERN]/other/generic/public/cer-0317330.pdf.
- [48] **ATLAS** Collaboration, *Detector and physics performance, Volume II*, 1999.
ATLAS-TDR-15,
[CERN]/other/generic/public/cer-0317331.pdf.
- [49] **ATLAS** Collaboration, *List of ATLAS TDRs*,
[ATLAS]/internal/tdr.html.
- [50] **ATLAS** Collaboration, *Magnetic field group website*,
[ATLAS]/GROUPS/MUON/magfield/mag_page1.html.

-
- [51] **ATLAS** Collaboration, *Inner detector technical design report, Volume I*, July, 1997.
ATLAS-TDR-4,
[CERN]/other/generic/public/cer-0254563.pdf.
- [52] **ATLAS Pixel** Collaboration, C. Gemme, *The ATLAS pixel detector*, Nucl. Instrum. Meth.
A501 (2003) 87–92.
- [53] A. J. Barr, *Studies of supersymmetry models for the ATLAS experiment at the Large Hadron Collider*. PhD thesis, **University of Cambridge**, 2002.
CERN-THESIS-2004-002,
[CERN]/cern/preprints/thesis/thesis-2004-002.pdf.
- [54] **ATLAS** Collaboration, *Calorimeter performance technical design report*. ATLAS-TDR-1,
<http://atlasinfo.cern.ch/Atlas/TDR/caloperf/caloperf.html>.
- [55] **ATLAS** Collaboration, *Muon spectrometer technical design report*. ATLAS-TDR-10,
<http://atlasinfo.cern.ch/Atlas/GROUPS/MUON/TDR/Web/TDR.html>.
- [56] I. Foster, C. Kesselman and S. Tuecke, *The anatomy of the Grid: Enabling scalable virtual organizations*, Lecture Notes in Computer Science **2150** (2001) 1,
<http://citeseer.ist.psu.edu/foster01anatomy.html>.
- [57] **ATLAS DC1** Task force, R. Sturrock *et. al.*, *A step towards a computing grid for the LHC experiments: ATLAS data challenge 1*. CERN-PH-EP-2004-028.
- [58] The Globus alliance,
<http://www.globus.org>.
- [59] GridPP: UK computing for particle physicists,
<http://www.gridpp.ac.uk>.
- [60] European DataGrid,
<http://eu-datagrid.web.cern.ch>.

References

- [61] **ATLAS** Collaboration, *DAQ, EF, LVL2 and DCS technical progress report*. CERN/LHCC 98-16,
[ATLAS]/GROUPS/DAQTRIG/TPR/PDF_FILES/TPR.bk.pdf.
- [62] **ATLAS** Collaboration, *Offline software framework: Athena*,
[ATLAS]/GROUPS/SOFTWARE/00/.
- [63] G. Barrand *et. al.*, *GAUDI - a software architecture and framework for building HEP data processing applications*, *Comput. Phys. Commun.* **140** (2001) 45–55.
- [64] G. Corcella *et. al.*, *HERWIG 6: An event generator for hadron emission reactions with interfering gluons (including supersymmetric processes)*, *JHEP* **01** (2001) 010,
hep-ph/0011363.
- [65] T. Sjostrand *et. al.*, *High-energy-physics event generation with PYTHIA 6.1*, *Comput. Phys. Commun.* **135** (2001) 238–259, hep-ph/0010017.
- [66] GEANT: Detector description and simulation tool. CERN Program Library Long writeup W5013,
<http://wwwasdoc.web.cern.ch/wwwasdoc/pdftdir/geant.pdf>.
- [67] S. Hurtado, M. Garcia-Leon and R. Garcia-Tenorio, *GEANT4 code for simulation of a germanium gamma-ray detector and its application to efficiency calibration*, *Nucl. Instrum. Meth.* **A518** (2004) 764–774,
<http://dx.doi.org/10.1016/j.nima.2003.09.057>.
- [68] **ATLAS** Collaboration, *Inner detector technical design report, Volume II*, July, 1997.
ATLAS-TDR-5,
http://atlasinfo.cern.ch/Atlas/GROUPS/INNER_DETECTOR/TDR/tdr.html.
- [69] R. L. Copley, *Structure, reactivity and molecular mechanisms in the Cu and Ag catalysed selective oxidation of alkenes*. PhD thesis, **University of Cambridge**, 2004.

-
- [70] **ATLAS SCT** Collaboration, *Silicon detector PRR papers*, 2000. ATL-IS-MR-0001, <http://edmsoraweb.cern.ch>.
- [71] **ATLAS SCT** Collaboration, *ABCD3TA ASIC: Requirements and specification v1.3*, 2000, https://edms.cern.ch/file/317413/1/abcd3ta_spec_v1.3.pdf.
- [72] M. F. Dentan *et al.*, *Rd29 final status report*, 1998. CERN-LHCC-98-037.
- [73] M. Dentan *et al.*, *Dmill (durci mixte sur isolant logico-lineaire): A mixed analog-digital radiation hard technology for high energy physics electronics*, Nucl. Phys. Proc. Suppl. **32** (1998) 530–534.
- [74] **ATLAS SCT** Collaboration, *SLOG: SLOW command Generator*, <http://hepwww.rl.ac.uk/atlas-sct/mm/Slog>.
- [75] **ATLAS SCT** Collaboration, *CLOAC: CLOCK And Control*, <http://www.hep.ucl.ac.uk/atlas/sct/cloac/Welcome.html>.
- [76] **ATLAS SCT** Collaboration, *MuSTARD: Multi-channel Semiconductor Tracker ABC(D) Readout Device*, <http://jchill.home.cern.ch/jchill/mustard.html>.
- [77] P. Phillips, L. Eklund, B. Moorhead, G. Gallop and J. Hill, *SCTDAQ*, <http://sct-testdaq.home.cern.ch/sct-testdaq/sctdaq/sctdaq.html>.
- [78] **ATLAS SCT** Collaboration, *SCT back-end electronics FDR/PRR*, <http://agenda.cern.ch/fullAgenda.php?ida=a03748>.
- [79] H. van der Bij, R. McLaren, O. Boyle and G. Rubin, *S-LINK, a data link interface specification for the LHC era*, in *IEEE Transactions on Nuclear Science*, vol. 44, pp. 398–402, ECP division, CERN, IEEE, June, 1997, http://ieeexplore.ieee.org/xpl/abs_free.jsp?arNumber=603679.

References

- [80] O. Boyle, R. McLaren and H. van der Bij, *The S-LINK interface specification*, 1997,
<http://hsi.web.cern.ch/HSI/s-link>.
- [81] L. Landau, *On the energy loss of fast particles by ionization*, Jour. Phys. USSR **8** (1944)
201.
- [82] A. J. Barr, B. J. Gallop, C. G. Lester, M. J. Palmer, D. Robinson and A. Tricoli,
SctRodDaq website,
<http://www.hep.phy.cam.ac.uk/~atlasdaq>.
- [83] B. J. Gallop, *ATLAS SCT Barrel Macro-Assembly testing and design of MAPS test
structures*. PhD thesis, **University of Birmingham**, 2005. In preparation.
- [84] L. Eklund, D. Ferrère, B. Gallop and P. W. Phillips, *Electrical tests of SCT hybrids and
modules*, ATLAS Comm. (2003) ATL-COM-INDET-2003-004,
[\[CERN\]/cern/others/atlnot/Communication/indet/com-indet-2003-004.pdf](http://cern/cern/others/atlnot/Communication/indet/com-indet-2003-004.pdf).
- [85] R. L. Glass, *Facts and Fallacies of Software Engineering*. Addison Wesley, 2003.
0-321-11742-5.
- [86] B. Appleton, *Source code line counter*,
<http://www.cmcrossroads.com/bradapp/clearperl/sclc-cdiff.html>.
- [87] **ATLAS** Collaboration, *Online software website*,
<http://atlas-onlsw.web.cern.ch/Atlas-onlsw>.
- [88] Object Management Group Inc., *Common Object Request Broker Architecture: Core
specification, v3.0.3*, March, 2004,
http://www.omg.org/technology/documents/formal/corba_iiop.htm.
- [89] CORBA website,
<http://www.corba.org>.

- [90] Object Management Group website,
<http://www.omg.org>.
- [91] Xerox Corporation, *Inter-language unification system*,
<ftp://ftp.parc.xerox.com/pub/ilu/ilu.html>.
- [92] omniORB,
<http://omniorb.sourceforge.net>.
- [93] E. Gamma, R. Helm, R. Johnson and J. Vlissides, *Design Patterns — Elements of reusable object-oriented software*. Addison-Wesley, 1995.
- [94] M. Cline, G. Lomow and M. Girou, *C++ FAQs*. Addison-Wesley, second ed., 1999,
<http://www.parashift.com/c++-faq-lite>.
- [95] Portland pattern repository,
<http://c2.com/cgi-bin/wiki>.
- [96] C. Alexander, S. Ishikawa, M. Silverstein, M. Jacobsen, I. Fiksdahl-King and S. Angel, *A Pattern Language*. Oxford University Press, New York, 1977.
- [97] Boost: Peer-reviewed C++ libraries,
<http://www.boost.org>.
- [98] R. Brun, F. Rademakers *et. al.*, *ROOT system*, Nucl. Inst. & Meth. in Phys. Res. A **389** (1997), Proceedings AIHENP'96 Workshop, Lausanne, Sep. 1996 81–86,
<http://root.cern.ch>.
- [99] Numerical Algorithms Group, *NAG C library 6*,
<http://www.nag.co.uk>.
- [100] D. Riehle, W. Siberski, D. Bäumer, D. Megert and H. Züllighoven, *The atomizer—efficiently streaming object structures*, in *Pattern Languages of Programs*,

References

- Washington University Department of Computer Science, 1996. Technical Report WUCS-97-07.
- [101] D. Bäumer and D. Riehle, *Late creation—a creational pattern*, in *Pattern Languages of Programs*, Washington University Department of Computer Science, 1996. Technical Report WUCS-97-07.
- [102] B. Gröne and P. Tabelaing, *A system of patterns for concurrent request processing servers.*, in *Viking Pattern Languages of Programs*, 2003.
- [103] N. Nethercote and J. Seward, *Valgrind: A program supervision framework*, in *Third Workshop on Runtime Verification, Colorado*, vol. 89, Theoretical Computer Science, 2003,
<http://valgrind.kde.org>.
- [104] J. Weidendorfer, *Kcachegrind: Profiling visualization*,
<http://kcachegrind.sourceforge.net>.
- [105] Extreme programming: A gentle introduction,
<http://www.extremeprogramming.org>.
- [106] JUnit: A unit testing framework for Java,
<http://www.junit.org>.
- [107] P. W. Phillips. Private communication.
- [108] PPARC press release, *ATLAS gets its heart*,
http://www.pparc.ac.uk/Nw/Atlas_news.asp.
- [109] The Register, *UK boffins sniff for Higgs boson*,
http://www.theregister.co.uk/2004/11/19/lhc_atlas_detector.
- [110] BBC News Online, *Device to probe limits of physics*,
<http://news.bbc.co.uk/1/hi/sci/tech/4035747.stm>.

-
- [111] A. J. Barr. Private communication.
- [112] ATLAS SCT Collaboration, A. J. Barr, *Front end calibration and control for the ATLAS semiconductor tracker*, in *Nuclear Science Symposium Conference Record*, IEEE, 2004.
- [113] G. Viehhauser, *Barrel macro-assembly photographs*,
<http://www-pnp.physics.ox.ac.uk/~viehhaus/pictures/>.
- [114] D. van Heesh, *Doxygen: A code documentation system*,
<http://www.doxygen.org>.
- [115] D. Robinson, *SctGUI user manual*,
<http://www.hep.phy.cam.ac.uk/~robinson/DisplayGUI.html>.
- [116] G. Corcella *et. al.*, *HERWIG 6.3 release note*, hep-ph/0107071.
- [117] E. Richter-Was, D. Froidevaux and L. Poggioli, *ATLFAST 2.0: A fast simulation package for ATLAS*, ATLAS Notes (1998) ATL-PHYS-98-131.
- [118] P. Savard and G. Azuelos, *The discovery potential of a heavy Higgs ($m_H \geq 800$ GeV) using full GEANT simulation of ATLAS*, ATLAS Note (1998) ATL-PHYS-98-128,
[CERN]/cern/others/atlnot/Note/phys/phys-98-128.pdf.
- [119] D. Froidevaux, L. Poggioli and S. Zmushko, *$H \rightarrow W^+ W^- \rightarrow \ell \nu jj$ and $H \rightarrow Z^0 Z^0 \rightarrow \ell \ell jj$ particle level studies*, ATLAS Note (1997) ATL-PHYS-97-103,
[CERN]/cern/others/atlnot/Note/phys/phys-97-103.pdf.
- [120] V. Cavasinni, D. Costanzo and I. Vivarelli, *Forward tagging and jet veto studies for Higgs events produced via vector boson fusion*, ATLAS Note (2002) ATL-PHYS-2002-008,
[CERN]/cern/others/atlnot/Note/phys/phys-2002-008.pdf.
- [121] W. T. Giele, T. Matsuura, M. H. Seymour and B. R. Webber, *W boson plus multijets at hadron colliders: HERWIG parton showers versus exact matrix elements*. Contribution to

References

- Proc. of 1990 Summer Study on High Energy Physics: Research Directions for the Decade, Snowmass, CO, Jun 25 - Jul 13, 1990.
- [122] R. Bonciani, S. Catani, M. L. Mangano and P. Nason, *NLL resummation of the heavy-quark hadroproduction cross-section*, Nucl. Phys. **B529** (1998) 424–450, hep-ph/9801375.
- [123] M. J. Palmer, *First use of the UK e-Science grid*, Presented to GridPP 5, September, 2002, http://www.gridpp.ac.uk/gridpp5/gridpp5_escience.ppt.
- [124] I. Park, *A new clustering algorithm: Mulguisin.*, ATLAS Comm. (1999) ATL-COM-PHYS-99-055, [CERN]/cern/others/atlnot/Communication/phys/com-phys-99-055.pdf.
- [125] P. Savard, *The $W \rightarrow \text{jet-jet}$ and top mass reconstructions with the ATLAS detectors*, ATLAS Note (1997) ATL-CAL-97-092, [CERN]/cern/others/atlnot/Note/cal/cal-97-092.pdf.
- [126] M. Bosman, I. Park, M. Cobal, D. Costanzo, S. Lami, R. Paoletti, G. Azuelos and K. Strahl, *Jet Finder Library: version 1.0*, ATLAS Note (1998) ATL-SOFT-98-038, [CERN]/cern/others/atlnot/Note/soft/soft-98-038.pdf.
- [127] K. Eggert, K. Eggert, M. Oriunno and M. Bozzo, *TOTEM Technical design report*, CERN Report (2004) CERN-LHCC-2004-002, [CERN]/cern/preprints/lhcc/public/lhcc-2004-002.pdf.
- [128] P. Jenni and M. Nessi, *ATLAS forward detectors for luminosity measurement and monitoring*, CERN Report (2004) CERN-LHCC-2004-010, [CERN]/cern/preprints/lhcc/public/lhcc-2004-010.pdf.
- [129] T. Banks and W. Fischler, *A model for high energy scattering in quantum gravity*, hep-th/9906038.

-
- [130] H. Yoshino and Y. Nambu, *Black hole formation in the grazing collision of high-energy particles*, Phys. Rev. **D67** (2003) 024009, gr-qc/0209003.
- [131] S. Hossenfelder, S. Hofmann, M. Bleicher and H. Stoecker, *Quasi-stable black holes at the LHC*, Phys. Rev. **D66** (2002) 101502, hep-ph/0109085.
- [132] R. Casadio and B. Harms, *Can black holes and naked singularities be detected in accelerators?*, Int. J. Mod. Phys. **A17** (2002) 4635–4646, hep-th/0110255.
- [133] S. W. Hawking, *Particle creation by black holes*, Commun. Math. Phys. **43** (1975) 199–220.
- [134] C. M. Harris, P. Richardson and B. R. Webber, *CHARYBDIS: A black hole event generator*, JHEP **08** (2003) 033, hep-ph/0307305.
- [135] C. M. Harris and P. Kanti, *Hawking radiation from a $(4+n)$ -dimensional black hole: Exact results for the Schwarzschild phase*, JHEP **10** (2003) 014, hep-ph/0309054.
- [136] E. Boos *et. al.*, *Generic user process interface for event generators*, hep-ph/0109068.
- [137] G. Corcella *et. al.*, *HERWIG 6.5 release note*, hep-ph/0210213.
- [138] A. D. Martin, R. G. Roberts, W. J. Stirling and R. S. Thorne, *Parton distributions: A new global analysis*, Eur. Phys. J. **C4** (1998) 463–496, hep-ph/9803445.
- [139] M. B. Voloshin, *Semiclassical suppression of black hole production in particle collisions*, Phys. Lett. **B518** (2001) 137–142, hep-ph/0107119.
- [140] M. Cavaglia, S. Das and R. Maartens, *Will we observe black holes at LHC?*, Class. Quant. Grav. **20** (2003) L205–L212, hep-ph/0305223.
- [141] S. B. Giddings and V. S. Rychkov, *Black holes from colliding wavepackets*, hep-th/0409131.
- [142] V. S. Rychkov, *Topics in black hole production*, hep-th/0410295.

References

- [143] S. D. H. Hsu, *Quantum production of black holes*, Phys. Lett. **B555** (2003) 92–98, hep-ph/0203154.
- [144] P. Kanti, *Black holes in theories with large extra dimensions: A review*, hep-ph/0402168.
- [145] D. Stojkovic, *Distinguishing between the small ADD and RS black holes in accelerators*, hep-ph/0409124.
- [146] **H1** Collaboration, I. Abt *et. al.*, *The tracking, calorimeter and muon detectors of the H1 experiment at HERA*, Nucl. Instrum. Meth. **A386** (1997) 348–396.
- [147] R. Lefèvre and C. Santoni, *A study of the jet energy reconstruction*, ATLAS Note (2001) ATL-PHYS-2001-011, [CERN]/cern/others/atlnot/Note/phys/phys-2001-011.pdf.
- [148] M. Biglietti, F. Brochu, D. Costanzo, F. De, E. Duchovni, A. Sen-Gupta, I. Hinchliffe, C. G. Lester, A. Lipniacka, P. Loch, E. Lytken, H. Ma, J. L. Nielsen, F. Paige, F. Polesello, S. Rajagopalan, D. Schrager, G. Stavropoulos, D. R. Tovey and M. Wielers, *Full supersymmetry simulation for ATLAS in DC1*, ATLAS Note (2004) ATL-PHYS-2004-011, [CERN]/cern/others/atlnot/Note/phys/phys-2004-011.pdf.
- [149] **ATLAS** Collaboration, *Athens physics workshop website*, <http://agenda.cern.ch/fullAgenda.php?id=a031081>.
- [150] **ATLAS** Collaboration, *Athena software validation website*, [ATLAS]/GROUPS/PHYSICS/SOFT_VALID/soft_valid.html.
- [151] R. Clift and A. Poppleton, *iPatRec: Inner detector pattern-recognition and track-fitting*, ATLAS Note (1994) ATL-SOFT-94-009, [CERN]/cern/others/atlnot/Note/soft/soft-94-009.pdf.
- [152] I. Gavrilenko, *Description of global pattern recognition program (XKalman)*, ATLAS Note (1997) ATL-INDET-97-165.

-
- [153] N. C. Benekos, R. Clift, M. Elsing and A. Poppleton, *ATLAS Inner detector performance*, ATLAS Note (2003) ATL-INDET-2004-002, [CERN]/cern/others/atlnot/Note/indet/indet-2004-002.pdf.
- [154] M. Virchaux, *Muonbox: a full 3d tracking programme for muon reconstruction in the ATLAS spectrometer*, ATLAS Note (1997) ATL-MUON-97-198.
- [155] D. Adams, K. Assamagan, M. Biglietti, G. Carlino, G. Cataldi, F. Conventi, A. Farilla, Y. Fisyak, S. Goldfarb, E. Gorini, T. Lagouri, K. Mair, L. Merola, A. Nairz, A. Poppleton, M. Primavera, S. Rosati, J. T. Shank, S. Spagnolo, L. Spogli, G. Stavropoulos, M. Verducci and T. Wenaus, *Track reconstruction in the ATLAS muon spectrometer with MOORE*, ATLAS Note (2003) ATL-SOFT-2003-007, [CERN]/cern/others/atlnot/Note/soft/soft-2003-007.pdf.
- [156] D. Fassouliotis, T. Lagouri, A. Poppleton, S. Robins, G. Stavropoulos and I. Vichou, *Muon identification using the MUID package*, ATLAS Comm. (2003) ATL-COM-MUON-2003-003, [CERN]/cern/others/atlnot/Communication/muon/com-muon-2003-003.pdf.
- [157] T. Lagouri, D. Adams, K. Assamagan, M. Biglietti, G. Carlino, G. Cataldi, F. Conventi, A. Farilla, Y. Fisyak, S. Goldfarb, E. Gorini, K. Mair, L. Merola, A. Nairz, A. Poppleton, M. Primavera, S. Rosati, S. Shank, S. Spagnolo, S. Spogli, G. Stavropoulos, M. Verducci and T. Wenaus, *A muon identification and combined reconstruction procedure for the ATLAS detector at the LHC at CERN*, ATLAS Note (2003) ATL-CONF-2003-011.
- [158] I. Mocioiu, Y. Nara and I. Sarcevic, *Hadrons as signature of black hole production at the LHC*, Phys. Lett. **B557** (2003) 87–93, hep-ph/0301073.
- [159] J. Tanaka, T. Yamamura, S. Asai and J. Kanzaki, *Study of black holes with the ATLAS detector at the LHC*, ATLAS Note (2003) ATL-PHYS-2003-037, [CERN]/cern/others/atlnot/Note/phys/phys-2003-037.pdf.