Searching for Supersymmetry with the ATLAS Detector

Martin John White
of Churchill College

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

Searching for Supersymmetry with the ATLAS Detector

Martin John White

This thesis presents a new method by which one may use data from the ATLAS detector of the Large Hadron Collider at CERN to measure the parameters of the theory of supersymmetry (SUSY). The technique uses a Markov Chain Monte Carlo sampling algorithm to combine measurements of exclusive variables, in the form of kinematic endpoints that arise in the invariant mass distributions of leptons and jets given out in sparticle decay chains, with inclusive data, in the form of the cross-section of events passing a missing transverse energy cut. This improves the precision of sparticle mass measurements (beyond that obtained using exclusive data alone), whilst also enabling experimental uncertainties to be handled in an intuitive fashion. The method is demonstrated on an mSUGRA benchmark model, and is also used to constrain a model with a greater number of parameters. Throughout, an attempt is made to break some of the unrealistic assumptions that characterise current SUSY search techniques, and to this end it is successfully demonstrated that one can use the Markov Chain method to obtain precise results even if it is not possible to precisely determine which sparticles produced the endpoints in the invariant mass distributions. This decay chain ambiguity is extended by looking at a sample SUSY model with non-universal Higgs masses, in which cascade decays featuring three body decays become more prominent. The positions of the associated endpoints are calculated for a squark cascade decay, and are subsequently studied with the aid of a benchmark model. In addition, this thesis presents work relating to the development of an online monitoring package for the ATLAS Semi-Conductor Tracker (SCT). A tool for the calculation of noise occupancies is developed and tested against existing calibration data, prior to its application to various data sets obtained during the SR1 cosmic commissioning tests at CERN in the summer of 2006. It is found that the modules of the SCT included in the test have an average noise occupancy of less than 5×10^{-4} per strip per event, and thus meet the design specification.

Declaration

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

Martin White

Acknowledgements

Preface

The theory of supersymmetry (SUSY) proposes a special relationship between fermions and bosons, and was first introduced in the context of particle physics in 1973. Subsequent years have seen much theoretical development, and it has become apparent that many of the problems of the Standard Model are solved by making the theory supersymmetric. Unfortunately, despite the many hours of high energy experiment performed in the last 30 years, there is no hard evidence to suggest that supersymmetry is indeed a valid theory of Nature.

This thesis nevertheless picks up the narrative at a particularly exciting time. As of next year, the largest particle accelerator ever built will be colliding protons at CERN in Geneva, Switzerland, and observing the results with sophisticated, multi-purpose detectors. The Large Hadron Collider (LHC) has two detectors designed to search for new physics, ATLAS and CMS, and these will collect large amounts of data over 10 years, at a centre of mass energy of 14 TeV. It has long been suspected that the new particles required by supersymmetry have masses of the order of 1 TeV, thus making them visible at the LHC.

The purpose of this thesis is to explore a method by which the data from ATLAS can be used to constrain the parameters of the supersymmetric theory. Current analyses in the field tend to rely on unrealistic assumptions, either by fixing most of the 105 possible parameters in order to make calculations easier, or by making false assumptions related to the identification of particles in decay chains. Here, a Markov Chain Monte Carlo method is used to combine exclusive and inclusive variables from the detector in order to provide enough information for these assumptions to be relaxed.

Chapter 1 reviews the Standard Model of particle physics, and explains the theoretical motivation for supersymmetry. This is followed in Chapter 2 by a discussion of the ATLAS experiment at the LHC, thus presenting the equipment with which experimenters hope to discover supersymmetry. The thesis is then divided into two distinct components.

The first covers work that is relevant to the operation of the detector itself. A Semi-Conductor Tracker (SCT) is used, amongst other things, to track particles in the ATLAS inner detector, and the design of this subsystem is presented in Chapter 3. The development of an online monitoring system for the SCT is described in Chapter 4, and the resulting software is applied to various data sets from the recent cosmic ray commissioning tests at CERN in order to investigate the noise occupancy of the detector strips.

The second original section of this thesis returns to the theory of supersymmetry, and applies the Markov Chain Monte Carlo sampling method to the supersymmetric parameter space using simulated ATLAS data. In order to motivate the need for a more general strategy for SUSY parameter measurement, Chapter 5 reviews the current approach to sparticle mass measurement using kinematic endpoints, including two original studies performed by the author. The chapter also introduces the Markov Chain Monte Carlo (MCMC) algorithm used later on. Having identified the key issues that need to be addressed (mainly those related to the low number of parameters probed by current studies, and decay chain ambiguity), Chapter 6 uses the MCMC method to combine a cross-section measurement with kinematic endpoint information in order to make definite statements on the impact of decay chain ambiguity, and the possibility of extending the number of parameters used in SUSY searches. Finally, Chapter 7 investigates a class of SUSY model where the Higgs masses are not unified at the GUT scale, and examines an interesting extension to the decay chain ambiguity considered in chapter 6.

Contents

Chapter 1

Theory

What is the world made of? This is the central question at the heart of particle physics, the study of the fundamental particles of nature and the interactions between them. Though the idea of a natural world comprised of a space filled with an almost infinite number of indivisible particles can be traced back to Greek philosophers, it was only in the last century that we really began to observe and understand Nature on the smallest scales. Indeed, nature appears to be far from the simple empty void full of atoms that Democritus described, instead being composed of families of fundamental objects whose interactions can only be understood through the often counter-intuitive ideas of quantum theory.

During the second half of the last century, experiments at the high energy frontier led to the development of the Standard Model (SM) of particle physics. The theory has proved to be astonishingly successful in predicting a wide range of phenomena with great accuracy, but, as we will see shortly, is deficient in certain key areas, most notably its inability to deal with gravity. For this reason, it is assumed that the SM is an effective theory valid only at low energy scales, and hence there must be some physics that has yet to be found. This forms the motivation for the next round of particle experiments, starting with the commencement of data taking at the Large Hadron Collider in 2007.

Here, a brief review of the SM and its problems is given, after which supersymmetry is introduced as a well-motivated extension of the current theory.

1.1   The Standard Model

1.1.1   Overview

The SM was developed in the early 1970s, and is a quantum field theory consistent with two of the greatest ideas of twentieth century physics- quantum mechanics and special relativity. Matter and forces are described in terms of point-like particles[1] which carry internal angular momentum. The quantum number called 'spin' which parameterises this internal property gives rise to two distinct types of particle; fermions (with half-integer spin) and bosons (with integer spin). All matter in the SM is in the former category (and hence must obey the Pauli exclusion principle), and all forces are mediated by particles that fall into the latter category.

It is believed that four forces are sufficient to describe all phenomena in nature, these being the strong force, the electromagnetic force, the weak force and gravity. The SM is a combination of gauge field theories that explain the first three of these, but does not currently incorporate gravity. The electromagnetic and weak force are given a unified description by 'electroweak theory', whilst the strong force is described by 'quantum chromodynamics' (QCD).

A particular subtlety of the SM is the treatment of mass. It is obvious from observation that most SM particles are massive, and yet putting mass terms into the SM Lagrangian directly breaks gauge invariance and results in a theory which is nonrenormalizable. The most popular solution to this problem is the Higgs mechanism which requires spontaneous symmetry breaking, and which is reviewed in more detail in 1.1.3 after a short description of the unbroken SM. A more complete description can be found in reference[1].

1.1.2   Matter and forces in the SM

One can divide matter in the SM into two categories; particles that feel the strong force (quarks) and particles that do not (leptons). The quarks and leptons in the SM are grouped into three families, distinguished by their successively greater masses. The common household atom is made only of particles in the first family, as particles in the higher families decay rapidly into the those of the first generation, and can only be observed over short timescales in high energy experiments. The origin of this peculiar family structure remains a great mystery.

---

[1]In reality, although the theory treats particles as 'point-like', the only possible empirical statement one can make is that the fundamental particles of the SM have no spatial extent above the current level of observation of $10^{-18} - 10^{-19}$ m.

Families	Color$_{L,R}$	T_L^3	Y_L	T_R^3	Y_R	$Q = T^3 + Y$
$\begin{pmatrix} u \\ d \end{pmatrix}, \begin{pmatrix} c \\ s \end{pmatrix}, \begin{pmatrix} t \\ b \end{pmatrix}$	$\mathbf{3}, \bar{\mathbf{3}}$	$\frac{1}{2}$ $-\frac{1}{2}$	$\frac{1}{6}$ $\frac{1}{6}$	0 0	$\frac{2}{3}$ $-\frac{1}{3}$	$\frac{2}{3}$ $-\frac{1}{3}$
$\begin{pmatrix} \nu_e \\ e^- \end{pmatrix}, \begin{pmatrix} \nu_\mu \\ \mu^- \end{pmatrix}, \begin{pmatrix} \nu_\tau \\ \tau^- \end{pmatrix}$	$\mathbf{1}, \mathbf{1}$	$\frac{1}{2}$ $-\frac{1}{2}$	$-\frac{1}{2}$ $-\frac{1}{2}$	0 0	0 -1	0 -1

Table 1.1: The matter content of the SM, including the gauge quantum numbers of the various species. T is the weak isospin, with third component T^3. Y is the $SU(1)_Y$ hypercharge, and Q is the electric charge. The subscripts L and R denote chiral components.

A force is introduced into the SM by requiring the Lagrangian that describes the matter fields to be invariant under a particular local (i.e. space-time dependent) transformation, called a "gauge transformation", and the precise form of the transformation relevant to each force has been motivated by experiment. In QED[2], for example, we impose invariance under a local U(1) transformation of the electron field, under which the electron field ψ transforms to $e^{i\theta(x)}\psi$, thus corresponding to a local rotation of the phase of the electron field. In order to preserve gauge invariance, an interaction term is introduced with a massless gauge boson (the photon). We also obtain a coupling constant to describe the strength of the interaction.

The weak force is formulated in a similar fashion, except that the transformation is more complicated. All fundamental fermions can be grouped in multiplets, each of which shares similar behaviour under the weak interaction (see Table 1.1) with, for example, u-type quarks always producing d-type quarks in decays [3], and vice versa. The analogue of 'electric charge' is 'isospin', and the weak force is introduced to the SM by requiring that the Lagrangian describing the fermion fields is invariant under SU(2) transformations in the vector space of weak isospin. There are three vector bosons associated with it as a consequence of the general rule that the number of gauge bosons associated with a force is the same as the number of generators of the group describing the transformation. In addition, since SU(2) is non-abelian (not all of the elements commute), we obtain self-interaction terms for the gauge bosons of the weak force.

All fermions in the unbroken SM are massless fields that can be expressed as a sum of left- and right-handed parts (whose behaviour is completely independent in this massless limit).

---

[2]The theory of electromagnetism that will later be unified with the description of the weak force.

[3]In fact, the $W^\pm$ bosons of the weak interaction do not couple directly to the quark fields, but instead to linear combinations of the quark fields.

These interact differently with the bosons of the weak force and this is described in the SM by only allowing the left-handed components of particles (the right-handed components of antiparticles) to transform under SU(2). It is possible to unify the electromagnetic and weak interactions by imposing invariance under the gauge group $SU(2)_L \times U(1)_Y$, where the subscript Y denotes a quantum number called hypercharge, and is intended to indicate that the U(1) transformation of the unified "electroweak" group is not the same as the QED gauge group referred to earlier[2–4]. The subscript L represents the fact that only the left-handed components of particles transform under SU(2).

The description of the strong force involves a Lagrangian invariant under SU(3) rotations in 'colour' space. Quarks are described by a field ψ_i, where the subscript i runs over the three possible colours. Eight gauge bosons, or "gluons", mediate the force, and the non-Abelian nature of SU(3) once more leads to self-interaction terms in the Lagrangian, meaning that gluons themselves carry colour charge. Colour isn't observed directly in nature, and thus it is argued that quarks are confined to colourless states known as hadrons by a potential that increases toward large distances.

1.1.3 Mass and the Higgs mechanism

Mass terms in a renormalizable theory

So far, we have only considered massless fermions and gauge bosons. For fermions, this is because we wanted to describe the parity violation of the weak interactions by decoupling the left and right-handed components of the Dirac field, and for bosons it is because gauge invariance of the unbroken SM can be shown to imply that spin-1 gauge bosons have zero mass if they are the only bosons in the theory. Whilst this second point is correct for photons and gluons, it is clearly not the case for the heavy bosons of the weak interaction, whose non-zero mass was measured to great precision at CERN in the 1980s.

A first naive approach to this problem would be to simply add mass terms for the weak gauge bosons to the Lagrangian, thus explicitly breaking the gauge symmetry. This gives us a theory which is non-renormalizable, however, and whose infinite number of divergences is catastrophically undesirable. Instead, one can invoke the method of Peter Higgs[5], in which the Lagrangian preserves gauge invariance whilst the vacuum state of the theory is no longer a singlet under the action of the gauge group. This "spontaneous symmetry breaking" generates massive gauge bosons without compromising renormalizability, and is achieved in the SM through the inclusion of a scalar field which is a doublet under SU(2).

Spontaneous symmetry breaking

In order to demonstrate the Higgs mechanism, consider a general SU(2) gauge theory for a complex doublet of scalar fields Φ^i, $i = 1, 2$, with the Lagrangian density:

$$\mathcal{L} = -\frac{1}{4}F^a_{\mu\nu}F^{a\mu\nu} + |\mathbf{D}_\mu\Phi|^2 - V(\Phi) \tag{1.1}$$

where $F^a_{\mu\nu}$ is the field strength tensor, and:

$$\mathbf{D}_\mu\Phi = \partial_\mu\Phi + igW^a_\mu\mathbf{T}^a\Phi \tag{1.2}$$

The terms here require some explanation. g is a coupling constant for the weak interaction, that parameterises the strength of the force. The matrices $\mathbf{T}^a$ are the generators of the SU(2) group, the gauge boson fields are denoted by W^a_μ and the superscript $a = 1, 2, 3$ runs over the generators of the group. The indices μ and ν of the field strength tensor are summation indices that run over the four rows and columns of the matrix, and the superscript a runs over the three gauge fields.

Now suppose that the potential $V(\Phi)$ specified above is given by:

$$V(\Phi) = -\mu^2\Phi^\dagger_i\Phi^i + \lambda(\Phi^\dagger_i\Phi^i)^2 \tag{1.3}$$

This has a local maximum rather than a minimum at $\Phi = 0$, and hence there is no longer a zero vacuum expectation value (VEV), but instead an infinite number of degenerate minima given by $\Phi^\dagger_i\Phi_i = \frac{1}{2}\mu^2/\lambda$. Thus there is a degenerate vacuum and symmetry breaking occurs when we make a choice of one of these minima to represent the true vacuum, which here, for simplicity, is selected to be:

$$\langle\Phi\rangle = \frac{1}{\sqrt{2}}\begin{pmatrix} 0 \\ v \end{pmatrix} \tag{1.4}$$

with $v = \mu/\sqrt{\lambda}$. We now expand Φ^i about its VEV:

$$\langle\Phi\rangle = \frac{1}{\sqrt{2}}\begin{pmatrix} \phi_1 + i\phi_2 \\ v + H + i\phi_3 \end{pmatrix} \tag{1.5}$$

Here, a general complex field has been added to the VEV, but one of the components has been

called H in anticipation of the significance it will acquire later on. Note that the fields ϕ_i, $i = 1...3$ and H all have a zero VEV, and that Equation 1.5 can be inserted into Equation 1.3 in order to obtain the potential in terms of the extra fields. In doing so, a term of the form $\mu^2 H^2$ is obtained, which is a mass term for the H field corresponding to a mass of $M_H = \sqrt{2}\mu$, but there are no analogous terms for the ϕ_i fields. Hence, although H is a field for a massive boson known as the Higgs boson, the ϕ_i fields describe massless fields called Goldstone bosons. One can always work in a certain gauge in which all three ϕ_i are set to zero[4]. Working in this gauge (the unitary gauge), one can write $\mathbf{D}_\mu\Phi$ as:

$$\mathbf{D}_\mu\Phi = \frac{1}{\sqrt{2}}\left[\partial_\mu\left(\begin{array}{c} 0 \\ v + H \end{array}\right) + i\frac{g}{2}\left(\begin{array}{cc} W_\mu^3 & \sqrt{2}W_\mu^- \\ \sqrt{2}W_\mu^+ & -W_\mu^3 \end{array}\right)\left(\begin{array}{c} 0 \\ v + H \end{array}\right)\right] \tag{1.6}$$

where the following notation has been introduced:

$$W_\mu^\pm \equiv (W_\mu^1 \pm iW_\mu^2)/\sqrt{2} \tag{1.7}$$

In addition, the following explicit form for the generators of SU(2) in the 2×2 representation is used:

$$\mathbf{T}^1 = \frac{1}{2}\left(\begin{array}{cc} 0 & 1 \\ 1 & 0 \end{array}\right), \mathbf{T}^2 = \frac{1}{2}\left(\begin{array}{cc} 0 & -i \\ i & 0 \end{array}\right), \mathbf{T}^3 = \frac{1}{2}\left(\begin{array}{cc} 1 & 0 \\ 0 & -1 \end{array}\right) \tag{1.8}$$

It is now possible to write explicitly the term $|\mathbf{D}_\mu\Phi|^2$ in Equation 1.1:

$$\begin{aligned} |\mathbf{D}_\mu\Phi|^2 = {} & \frac{1}{2}\partial_\mu H\partial^\mu H + \frac{1}{4}g^2v^2(W_\mu^+W^{-\mu} + \frac{1}{2}W_\mu^3 W^{3\mu}) \\ & + \frac{1}{4}g^2 H^2(W_\mu^+W^{-\mu} + \frac{1}{2}W_\mu^3 W^{3\mu}) + \frac{1}{2}g^2 vH(W_\mu^+W^{-\mu} + \frac{1}{2}W_\mu^3 W^{3\mu}) \end{aligned} \tag{1.9}$$

The first term is a kinetic term for the Higgs boson, whilst the third and fourth terms (counting those in brackets as one term) are both interaction terms between the Higgs field and the weak bosons, specifying quartic interactions and Yukawa interactions respectively. The sec-

---

[4]The Goldstone bosons still play an important part in the theory, however, as the quantum degrees of freedom associated with them must remain even after a choice of gauge is made. To see what happens to them, consider that a massless vector boson has only two degrees of freedom whereas a massive boson has three possible helicity values. This third degree of freedom arises from a Goldstone boson, and hence the gauge bosons "eat" the Goldstone bosons in order to become massive. For the case of the SU(2) group considered here, there are three Goldstone bosons, and each is eaten by one of the three gauge bosons of SU(2). Hence the introduction of the four degrees of freedom of a scalar doublet has given us the Higgs field and the three longitudinal components of the now massive vector bosons.

ond term, meanwhile, is a mass term for the weak bosons. Each of the three gauge bosons associated with the SU(2) group has acquired a mass:

$$M_W = \frac{gv}{2} \tag{1.10}$$

The Higgs mechanism

The Higgs mechanism in the SM proceeds exactly as in the previous section. In order to obtain massive gauge bosons for the $SU(2)_L \times U(1)_Y$ group, one can introduce a scalar 'Higgs' field as a doublet under SU(2), as given by Equation 1.5. The field is a singlet under SU(3). Because the kinetic term $(D_\mu \Phi)^\dagger (D^\mu \Phi)$ now contains contributions from the $U(1)_Y$ boson, one finds that the fields arising from SU(2) mix with that arising from $U(1)_Y$ and hence the physical gauge bosons $W^\pm$, Z and γ must be superpositions of the original gauge fields.

In addition, one can obtain massive fermions by adding gauge invariant terms of the form:

$$\mathcal{L}_{\text{Yukawa}} = -G_e \bar{e}_L{}^i \Phi_i e_R + \text{h.c.} \tag{1.11}$$

where h.c. denotes the Hermitian conjugate of the first term. This is called a 'Yukawa' term, and is an interaction between the left-handed electron doublet $\bar{e}_L{}^i$, the right-handed electron e_R and the scalar doublet Φ_i. Analogous terms exist for muons and taus. G_e is a Yukawa coupling constant that parameterises the strength of the interaction. In the unitary gauge, one may write equation 1.11 as:

$$\mathcal{L}_{\text{Yukawa}} = -\frac{G_e}{\sqrt{2}} \begin{pmatrix} \bar{\nu}_L \\ \bar{e}_L \end{pmatrix}^T \begin{pmatrix} 0 \\ v+H \end{pmatrix} e_R + \text{h.c.} \tag{1.12}$$

The part of this proportional to the vacuum expectation value of the Higgs field is:

$$-\frac{G_e v}{\sqrt{2}} (\bar{e}_L e_R + \bar{e}_R e_L) = -\frac{G_e v}{\sqrt{2}} \bar{e}e \tag{1.13}$$

This is a mass term for the electron, which has acquired a mass that is proportional to the vacuum expectation of the scalar field and to the Yukawa coupling G_e. G_e may be written in terms of the weak coupling constant g, the electron mass m_e and the W mass M_W as:

$$G_e = g \frac{m_e}{\sqrt{2} M_W} \tag{1.14}$$

In addition, equation 1.12 has a term expressing the coupling between the electron and the Higgs field:

$$-g\frac{m_e}{2M_W}\bar{e}He \tag{1.15}$$

This reflects the general feature that the couplings of the Higgs boson are proportional to the masses (or squares of masses) of the particles it is interacting with, which is an important factor in the phenomenology of Higgs searches.

The Higgs mass and the hierarchy problem

Although its appearance in accelerators is hotly anticipated, the Higgs boson is the only particle in the SM that has yet to be observed. Furthermore, we have seen that the Higgs acquires a mass $M_H = \sqrt{2}\mu$, equivalent to $M_H \propto \lambda v^2$, and this unfortunately means that we cannot predict M_H in the SM, as we do not know the value of the quartic coupling λ.[5] One can, however, use experimental constraints along with the structure of the theory to place bounds on the Higgs mass.

One such constraint arises from considering the elastic scattering process $W^+W^- \rightarrow W^+W^-$, which can be described in terms of a matrix S which must be unitary. This implies an upper bound on M_H of $\left(\frac{8\pi\sqrt{2}}{3G_F}\right)^{\frac{1}{2}} = 1008$ GeV[6] and means that, if it exists at all, one would expect to see the Higgs boson at the LHC. One can also consider the fact that λ rises with energy for large Higgs masses, and hence the theory will eventually become non-perturbative. The demand that this does not occur below some energy scale Λ implies an upper bound on the Higgs mass[7, 8]. Finally, a lower bound can be obtained theoretically by studying quantum corrections to the SM, and by imposing the requirement that the effective potential be positive definite[9]. These arguments imply that, if the SM is to be consistent up to the GUT scale $\Lambda \approx 10^{16}$GeV, one expects M_H to lie in the range 130 to 190 GeV. The observation of a Higgs boson with a mass below 130 GeV would imply the existence of new physics below the GUT scale, and hence the Higgs mechanism is intimately connected with searches for new physics.

In addition to theoretical data, one can use experimental results. Indirect bounds are obtained through the use of precision fits to electroweak data, and these currently predict that $M_H < 219$ GeV at the 95% confidence level[10]. There are also direct lower bounds on the SM Higgs mass resulting from the failure to observe it in previous collider experiments; the most substantial of these to date is the LEP2 limit which gives $M_H > 114.4$ GeV at the 95%

---

[5]The VEV of the Higgs field is, however, fixed by the Fermi coupling, and is given by $v = 247$ GeV.
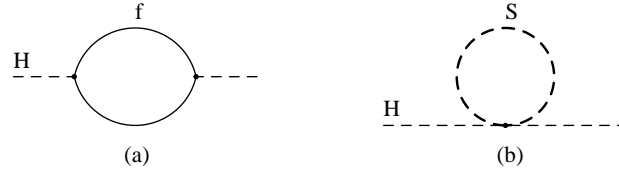
Figure 1.1: Loop corrections to the SM Higgs mass.

confidence level[11].

The fact that all available evidence points toward a light Higgs mass unfortunately causes problems in the SM, which predicts that M_H will be subject to large radiative corrections due to loop diagrams of the form shown in Figure 1.1. A fermion of mass m_f and Higgs coupling λ_f gives the following contribution to the square of the Higgs mass:

$$\delta M_H^2 = \frac{|\lambda_f|^2}{16\pi^2}\left[-2\Lambda^2 + 6m_f^2\ln(\Lambda/m_f) + ...\right] \tag{1.16}$$

where Λ is the cut-off scale of the theory, beyond which new physics is expected to enter. Any scalars resulting from new physics, meanwhile, give a contribution:

$$\delta M_H^2 = \frac{\lambda_s}{16\pi^2}\left[\Lambda^2 - 2m_s^2\ln(\Lambda/m_s) + ...\right] \tag{1.17}$$

Summing these contributions results in divergences quadratic in Λ, and one sees immediately that setting Λ to the unification scale or the GUT scale (both reasonable assumptions) gives a huge Higgs mass. The only way to obtain a Higgs with a mass within the range specified previously is to introduce what may be considered as a high level of fine-tuning, and this is known as the 'hierarchy problem'. We will see shortly that supersymmetry provides us with an elegant solution.

1.1.4   Problems of the SM

Any review of the SM is incomplete without a discussion of the problems of the theory. We have already seen that gravity is not incorporated, and this leads to the conclusion that we have an effective theory only valid below some energy cut off scale Λ. We have also seen that such a cut-off scale causes problems when the Higgs mass is considered, and this leads to a hierarchy problem.

In addition to these, it is noted that the SM does not contain explanations for the particular pattern of quantum numbers we obtain, and neither does it explain the family structure. Furthermore, it features at least 19 arbitrary parameters.

It is when pondering such facts that one is forced to conclude that, despite its successes, there is highly likely to be new physics beyond the SM, and it is at this point that we turn our attention to one of the most well-motivated extensions; supersymmetry (SUSY).

1.2 Supersymmetry

In this section, a brief review of the most relevant topics in supersymmetry is given. More comprehensive discussions can be found in references[12–14].

1.2.1 Theoretical motivation

We have already seen the importance of symmetry in the formulation of the SM, having used the principle of invariance under gauge transformations to construct the Lagrangian of the theory. This naturally leads one to consider the question of what is the most general possible group of symmetries.

The SM is in fact invariant under internal symmetries such as gauge symmetries and a group of transformations known as the Poincaré group, which consists of translations along with Lorentz boosts and rotations. Coleman and Mandula produced a no-go theorem[15] which states that a non-trivial extension of the Poincaré group[6] is not possible, though there is a get out clause in the case of the SM. The only way to extend the group is to make it supersymmetric; that is, to add an operator that transforms bosonic states to fermionic states, and vice versa (this is the Haag-Lopuszanski-Sohnius theorem[16]). A theory is then said to be supersymmetric if it is invariant under this expanded Poincaré group, and the fact that nature appears to respect the other symmetries in the group convinces some theorists that there is no reason why it should not also respect supersymmetry.

The generators of the supersymmetric transformation, Q, change the spin of a single particle state by $\pm\frac{1}{2}$, and the immediate implication is that each bosonic degree of freedom in the theory has a corresponding fermionic degree of freedom. It follows that each of the quarks and

---

[6]i.e. an extension where the generators of the new transformations have non-zero commutation relations with the Poincaré generators.
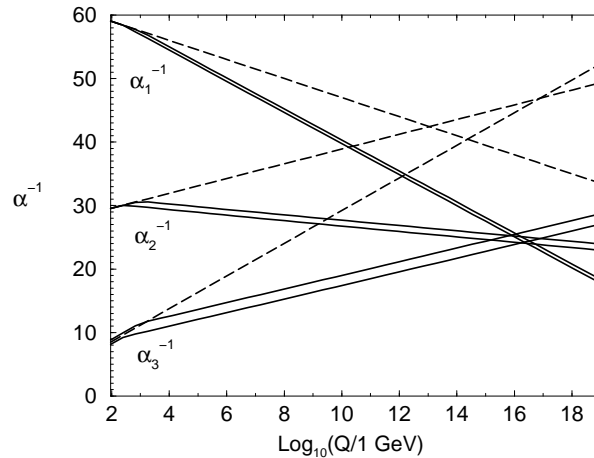
Figure 1.2: Renormalisation group running of the inverse gauge couplings α_i^{-1} in the SM (dashed lines) and the MSSM (solid lines). In the MSSM case, the sparticle mass thresholds are varied between 250 GeV and 1 TeV, whilst $\alpha_S(m_Z)$ is varied between 0.113 and 0.123. Two-loop effects are included[13].

leptons in the SM is accompanied by two complex scalars (one for each chiral component), and this leads to a cancellation of the Λ^2 terms in Equations 1.16 and 1.17 if one sets $\lambda_S = |\lambda_f|^2$, thus reducing the divergence from quadratic to logarithmic. Furthermore, invariance under supersymmetry ensures that this success persists to higher orders, and so one obtains an elegant and natural solution to the hierarchy problem.

There is another piece of theoretical evidence in support of SUSY related to grand unification. The gauge coupling strengths of the SM vary with the energy scale of interactions, and many theorists postulate that they may become equal at a high scale, thus uniting all of the forces of the SM into a single interaction characterised by a single coupling constant at the unification mass. This is equivalent to embedding the SU(3) × SU(2) × U(1) group in a simpler group which is spontaneously broken below the unification mass. The relevance to SUSY is that the running of the gauge coupling strengths is dependent on the particular particles that are available over a given energy range, and hence introducing new particles will affect the differential equations that describe the evolution from the GUT scale to the weak scale. It turns out that adding SUSY particles with masses at the electroweak scale improves gauge unification (see Figure 1.2).

1.2.2   The Minimal Supersymmetric Model

In this section, consideration is given to the Minimal Supersymmetric Model (MSSM); that is, the supersymmetric theory that requires a minimum of extra particle content beyond the SM.

The single-particle states of a supersymmetric theory fall into irreducible representations of the SUSY algebra, called *supermultiplets*. Each supermultiplet contains both boson and fermion states, which are referred to as *superpartners* of each other, and one can use the commutation relations for Q and other operators to show that members of the same supermultiplet must have the same mass, electric charges, weak isospin and colour degrees of freedom.

One can immediately ask what the simplest possibilities for supermultiplets are. It turns out that in models with only one distinct copy of the supersymmetry generators[7] the various possible combinations of particles in a supermultiplet can always be reduced to combinations of chiral supermultiplets (containing a single Weyl fermion and two real scalars that are represented by a single complex scalar field), and gauge supermultiplets (containing a massless spin-1 gauge boson and a massless spin-1/2 Weyl fermion). These both satisfy a rule stating that the number of bosonic degrees of freedom in a supermultiplet must equal the number of fermionic degrees of freedom and, in addition, it can be shown that the fermions in a gauge supermultiplet must have the same gauge transformation properties for left-handed and right-handed components. Thus, the particles of the SM alone cannot form the entire particle content of the MSSM, as all of the quarks and leptons of the SM have left- and right-handed parts that transform differently under the gauge group, and hence they must all be in chiral supermultiplets. New particles are needed to fill the fermionic degrees of freedom in the gauge supermultiplets, and there is also a need for new particles that fill the bosonic degrees of freedom in the chiral multiplets.

The supermultiplets of the MSSM are shown in Tables 1.2 and 1.3. Note that the left-handed and right-handed components of the SM quarks and leptons are separate two-component Weyl fermions, and hence each must have its own scalar partner. These 'squarks' and 'sleptons' are denoted by putting a tilde over the corresponding SM partner, and the subscripts L and R refer to the handedness of the SM partner (the sparticles have spin 0 and hence are neither right- nor left-handed). The gauge interactions of the sparticles are the same as their SM partners.

The Higgs sector of the MSSM requires some explanation, as it is more complicated than

---

[7]Models that do not satisfy this choice are known as "extended" supersymmetric theories, and they cannot easily reproduce the chiral fermions or parity violation observed in the SM. Hence the discussion here will only consider $N = 1$ supersymmetry, where N specifies the number of supersymmetries.

that of the SM. It is clear that the Higgs boson of the SM must exist in a chiral supermultiplet given that it has spin 0, though a less obvious fact is that one actually requires two Higgs doublets to complete the MSSM. One of the reasons for this results from the general structure of supersymmetric theories, in which it can be shown that only a $Y = +1/2$ Higgs chiral supermultiplet can have the Yukawa couplings necessary for charge +2/3 quarks, and only a $Y = -1/2$ Higgs can give the right couplings for charge -1/3 quarks and charged leptons. This gives us the two complex $SU(2)_L$ doublets H_u and H_d shown in Table 1.3, containing eight real, scalar degrees of freedom. When electroweak symmetry breaking occurs in the MSSM three of them form Goldstone bosons which become the longitudinal modes of the Z^0 and $W^\pm$ vector bosons, and the remaining five give us Higgs scalar mass eigenstates consisting of one CP-odd[8] neutral scalar A^0, a charge +1 scalar $H^+$and its conjugate H^-, and two CP-even[9] neutral scalars h^0 and H^0. Whilst the masses of A^0, H^0 and $H^\pm$ can be arbitrarily large, one can set an upper bound on the h^0 mass; assuming that none of the MSSM sparticles have masses exceeding 1 TeV, and that all of the couplings in the theory remain perturbative up to the unification scale, one obtains an approximate upper bound of 150 GeV[17, 18], and thus we see that at least one of the MSSM Higgs bosons must be light.

The states shown in Tables 1.2 and 1.3 mix in the MSSM to produce the following physical eigenstates:

$$H_u^0, H_d^0, H_u^+, H_d^- \rightarrow h^0, H^0, A^0, H^\pm \text{ (Higgs)}$$

$$\tilde{t}_L, \tilde{t}_R, \tilde{b}_L, \tilde{b}_R \rightarrow \tilde{t}_1, \tilde{t}_2, \tilde{b}_1, \tilde{b}_2 \text{ (stop/sbottom)}$$

$$\tilde{\tau}_L, \tilde{\tau}_R, \tilde{\nu}_\tau \rightarrow \tilde{\tau}_1, \tilde{\tau}_2, \tilde{\nu}_\tau \text{ (stau)}$$

$$\tilde{B}^0, \tilde{W}^0, \tilde{H}_u^0, \tilde{H}_d^0 \rightarrow \tilde{\chi}_1^0, \tilde{\chi}_2^0, \tilde{\chi}_3^0, \tilde{\chi}_4^0 \text{ (neutralinos)}$$

$$\tilde{W}^\pm, \tilde{H}_u^+, \tilde{H}_d^- \rightarrow \tilde{\chi}_1^\pm, \tilde{\chi}_2^\pm \text{ (charginos)}$$

where it is noted that the mixing in the squark and slepton sectors is typically proportional to the mass of the associated fermion, and thus is assumed to be largest for the third family.

---

[8]CP is the product of two symmetries, charge conjugation and parity. The parity transformation involves inverting the spatial coordinates of a matter field, and charge conjugation involves turning a particle into its antiparticle. 'CP odd' means that a field ψ transforms to $-\psi$ under a CP transformation.

[9]'CP even' means that a field ψ transforms to itself under a CP transformation.

1.2.3 SUSY breaking

If SUSY were an exact symmetry of nature, the sparticles introduced above would have the same masses as their SM counterparts and would have been seen already in collider experiments. Sadly, this isn't the case, and there remains no direct experimental evidence for supersymmetry. Any valid supersymmetric theory must therefore introduce a mechanism for supersymmetry breaking, and one can introduce spontaneous symmetry breaking in a way directly analogous to the treatment of electroweak symmetry breaking in the SM.

Recall the previous assertion that, if supersymmetry is to solve the hierarchy problem, one requires that $\lambda_S = |\lambda_f|^2$. This must persist in any broken SUSY theory, and hence only 'soft' supersymmetry breaking is allowed; that is, the effective Lagrangian of the MSSM can be written in the form

$$\mathcal{L} = \mathcal{L}_{\text{SUSY}} + \mathcal{L}_{\text{soft}} \tag{1.18}$$

where $\mathcal{L}_{\text{SUSY}}$ preserves supersymmetric invariance, and $\mathcal{L}_{\text{soft}}$ violates supersymmetry but only contains mass terms and couplings with positive mass dimension. This restriction is necessary to ensure that the breaking terms are consistent with Poincare and SM gauge invariance, and do not reintroduce quadratic divergences for scalar particles.

Many models of spontaneous symmetry breaking have been proposed, and, although all of them involve adding new particles and interactions at very high mass scales, there is no general consensus on which is the correct mechanism. We may, however, take the traditional route in physics of simply parameterising our ignorance, accomplished in this case by explicitly adding soft SUSY breaking terms to the Lagrangian. The most general $\mathcal{L}_{\text{soft}}$ that can be constructed that is compatible with gauge invariance and R-parity conservation[10] is

$$
\begin{aligned}
\mathcal{L}_{\text{soft}}^{\text{MSSM}} =\ & -\frac{1}{2}\left(M_3\tilde{g}\tilde{g} + M_2\tilde{W}\tilde{W} + M_1\tilde{B}\tilde{B}\right) + \text{c.c.} \\
& -\left(\tilde{\bar{u}}\mathbf{a_u}\tilde{Q}H_u - \tilde{\bar{d}}\mathbf{a_d}\tilde{Q}H_d - \tilde{\bar{e}}\mathbf{a_e}\tilde{L}H_d\right) + \text{c.c.} \\
& -\tilde{Q}^\dagger\mathbf{m_Q^2}\tilde{Q} - \tilde{L}^\dagger\mathbf{m_L^2}\tilde{L} - \tilde{\bar{u}}\mathbf{m_{\bar{u}}^2}\tilde{\bar{u}}^\dagger - \tilde{\bar{d}}\mathbf{m_{\bar{d}}^2}\tilde{\bar{d}}^\dagger - \tilde{\bar{e}}\mathbf{m_{\bar{e}}^2}\tilde{\bar{e}}^\dagger \\
& -m_{H_u}^2 H_u^* H_u - m_{H_d}^2 H_d^* H_d - (bH_uH_d + \text{c.c.})
\end{aligned}
\tag{1.19}
$$

H_u, H_d, $\tilde{Q}$, $\tilde{L}$, $\tilde{\bar{u}}$, $\tilde{\bar{d}}$ and $\tilde{\bar{e}}$ are chiral 'superfields' corresponding to the chiral supermultiplets listed in Table 1.2, and the adjoint representation indices on the wino and gluino fields, and

---

[10]see Section 1.2.4.

the gauge indices on all of the chiral supermultiplet fields, have been suppressed. M_1, M_2 and M_3 are the gluino, wino and bino mass terms, whilst the second line in Equation 1.19 contains (*scalar*)^3 couplings (each of $\mathbf{a_u}$, $\mathbf{a_d}$, and $\mathbf{a_e}$ is a complex 3×3 matrix with dimensions of (mass)). Line three contains squark and slepton mass terms (each of the $\mathbf{m_i^2}$ matrices is a 3×3 hermitian matrix in family space) and the final line lists the supersymmetry-breaking contributions to the Higgs potential. Clearly, the fact that nothing is known about the mechanism of SUSY breaking has given us many new parameters of the theory, and a careful count reveals that the terms of Equation 1.19 contain a total of 105 masses, phases and mixing angles on top of those already present in the SM. Although it is possible to heavily constrain some of these parameters using existing experimental data, the difficulty of performing a general analysis in such a large parameter space remains a pressing issue, and will be a recurring theme of this thesis.

Having obtained the most general breaking terms allowed in the MSSM, it is necessary to briefly consider how SUSY breaking is postulated to work in practise. It is expected that the MSSM soft terms arise indirectly or radiatively rather than from tree-level renormalizable interactions, and so SUSY breaking must occur in a 'hidden sector' of particles which have little or no direct couplings to the chiral supermultiplets. Of course, some interactions between the hidden and visible sectors of the theory are required in order to mediate SUSY breaking, and these appear as the calculable soft terms. There are various suggestions for what these mediating interactions might be, and the main candidates are reviewed below. In addition, the difficulty involved in working with the large number of parameters of the MSSM means that the vast majority of phenomenological studies are performed in simplified models where many of the parameters are chosen to be degenerate.

Supergravity

In the supergravity (SUGRA) class of models, the mediating interactions are gravitational; that is they result from new physics which includes gravity and enters at the Planck scale. In the simplest formulation of the model (minimal supergravity, or mSUGRA) the gaugino masses are unified to some mass $m_{1/2}$, and the matrices of $\tilde{q}$ and $\tilde{l}$ masses are set to $m_0^2 \mathbf{1}$, where m_0 is some common scalar mass. Similarly, the Higgs masses m_{H_u} and m_{H_d} are set to m_0. Assuming the couplings of sparticles are flavour blind, the coupling matrices are all set to $A_0 \mathbf{y}$, where $\mathbf{y}$ are the Yukawa couplings. Finally, the b mass term is expressed as $b = B_0 \mu$, and one is left with five parameters for the mSUGRA model:

$$m_{1/2}, m_0, A_0, b, \mu$$

It is conventional to replace b and μ by $\tan\beta$ and $\text{sgn}(\mu)$, where β is the ratio of the H_u and H_d vacuum expectation values. It should also be noted that the model contains a gravitino with mass $m_{3/2}$.

The mSUGRA mechanism is elegant, as it requires only the use of existing fields and interactions such as gravity, although one has to fine tune the squark and slepton mass matrices to avoid the introduction of unobserved large flavour-changing neutral current effects. mSUGRA models are very attractive phenomenologically, however, and there is a great deal of effort currently being invested in observing the phenomenology at different points in the parameter space.

Anomaly mediated SUSY breaking

Anomaly Mediated Supersymmetry Breaking (AMSB) is a special case of the gravity mediation scenario, in which there is no direct tree level coupling that transmits the SUSY breaking from the hidden sector to the observable one. In this case, SUSY breaking is communicated to SUSY particles via loop effects; gaugino masses are generated at one-loop level, and scalar masses are generated at two-loop level. The mechanism solves the flavour problem of mSUGRA and all of the low energy phenomenology may be obtained by adding only two extra parameters and one sign to the SM. Unfortunately, this minimal AMSB framework leads to negative squared masses for the sleptons at the electroweak scale, and thus it is necessary to add one additional parameter representing an extra contribution to the SUSY breaking, and this is normally provided in the form of an extra universal scalar mass parameter.

Thus, the parameter set for an AMSB model would comprise of the gravitino mass $m_{3/2}$, the ratio of the vacuum expectation values of the Higgs fields $\tan\beta$, the common scalar mass parameter m_0, and the sign of the Higgs term $\text{sgn}(\mu)$.

Gauge mediated SUSY breaking

In the Gauge Mediated Supersymmetry Breaking (GMSB) scenario, *messenger* particles interact with a SUSY breaking VEV in the hidden sector and communicate with the MSSM particles via ordinary gauge interactions. These messenger particles could be states from a larger gauge group than those found in the SM, e.g. SU(5).

GMSB models have six parameters: F_m (the SUSY breaking scale), M_m (the messenger scale), N_5 (the number of messenger multiplets), $\tan\beta$ and $\text{sgn}(\mu)$ (as in the mSUGRA framework), and C_{grav} (the coupling for decays into gravitinos). In GMSB models, the gravitino is not related to m_{soft} and is expected to be very light; hence, all final states contain gravitinos. Weak coupling between the gravitino and other SUSY particles means that the next to lightest SUSY particles (NLSPs) may have a long lifetime, and thus GMSB models give rise to a very distinctive phenomenology.

1.2.4   R-parity

When one writes down the superpotential for the MSSM, it is possible in general to include terms which are gauge invariant and analytic in the superfields given in Table 1.2, but which violate baryon number (B) and lepton number (L). Both of these processes must cause very small effects in nature in order to reproduce experimental data, with the most obvious constraint being the non-observation of proton decay[19] (which violates both baryon and lepton number by 1 unit). One cannot simply forbid these terms by allowing B and L to be fundamental symmetries of nature, as it is known that they are necessarily violated by non-perturbative electroweak effects (though these effects are negligible for experiments at ordinary energies). However, one can introduce a multiplicatively conserved quantum number called matter parity, defined as:

$$P_M = (-1)^{3(B-L)} \tag{1.20}$$

for each particle in the theory. It is then possible in principle to impose an exact and fundamental symmetry which states that all terms in the Lagrangian must contain a combination of fields that has a P_M product of +1. This forbids the explicit B and L violating terms in the proposed MSSM Lagrangian.

For phenomenological applications, it is often useful to recast matter parity in an alternative form known as 'R-parity':

$$P_R = (-1)^{3(B-L)+2s} \tag{1.21}$$

It is stressed that conservation of R-parity is completely equivalent to conservation of matter parity, but using R-parity conveniently means that all SM particles have even R-parity ($P_R = +1$), whilst all of the new SUSY particles have odd R-parity ($P_R = -1$). Thus, exact R-

Names		spin 0	spin 1/2	$SU(3)_C, SU(2)_L, U(1)_Y$
squarks, quarks	Q	$(\tilde{u}_L \; \tilde{d}_L)$	$(u_L \; d_L)$	$(\mathbf{3}, \mathbf{2}, \frac{1}{6})$
(×3 families)	$\bar{u}$	$\tilde{u}_R^*$	$u_R^\dagger$	$(\bar{\mathbf{3}}, \mathbf{1}, -\frac{2}{3})$
	$\bar{d}$	$\tilde{d}_R^*$	$d_R^\dagger$	$(\bar{\mathbf{3}}, \mathbf{1}, \frac{1}{3})$
sleptons, leptons	L	$(\tilde{\nu} \; \tilde{e}_L)$	$(\nu \; e_L)$	$(\mathbf{1}, \mathbf{2}, -\frac{1}{2})$
(×3 families)	$\bar{e}$	$\tilde{e}_R^*$	$e_R^\dagger$	$(\mathbf{1}, \mathbf{1}, 1)$
Higgs, higgsinos	H_u	$(H_u^+ \; H_u^0)$	$(\tilde{H}_u^+ \; \tilde{H}_u^0)$	$(\mathbf{1}, \mathbf{2}, +\frac{1}{2})$
	H_d	$(H_d^0 \; H_d^-)$	$(\tilde{H}_d^0 \; \tilde{H}_d^-)$	$(\mathbf{1}, \mathbf{2}, -\frac{1}{2})$

Table 1.2: The chiral supermultiplets of the MSSM.

Names	spin 1/2	spin 1	$SU(3)_C, SU(2)_L, U(1)_Y$
gluino, gluon	$\tilde{g}$	g	$(\mathbf{8}, \mathbf{1}, 0)$
winos, W bosons	$\tilde{W}^\pm \; \tilde{W}^0$	$W^\pm \; W^0$	$(\mathbf{1}, \mathbf{3}, 0)$
bino, B boson	$\tilde{B}^0$	B^0	$(\mathbf{1}, \mathbf{1}, 0)$

Table 1.3: The gauge supermultiplets of the MSSM.

parity conservation not only forbids mixing between SM and SUSY particles, but leads to the following important phenomenological results:

1. In collider experiments, such as the LHC, one must produce sparticles in pairs.

2. The lightest supersymmetric particle (LSP) must be absolutely stable. We will see later that, if the LSP is electrically neutral, it interacts only weakly with ordinary matter, and hence is an excellent candidate for dark matter.

3. All sparticles other than the LSP must decay into states featuring an odd number of sparticles. In practise, this number is usually one.

1.2.5 Supersymmetry and dark matter

As mentioned above, in R-parity conserving supersymmetric models, the lightest supersymmetric particle (LSP) is stable and is therefore an ideal candidate for non-baryonic cold dark

matter. By assuming that the LSP is responsible for the dark matter present in the universe, one can use cosmological measurements of the dark matter density to constrain supersymmetry models and attempt to reveal more about the mechanism of SUSY breaking.

The recent data from the WMAP satellite[20] have allowed the matter density of the Universe to be quantified with greater precision than ever before, whilst also strongly disfavouring warm dark matter. With a total reported matter density of $\Omega_m h^2 = 0.135^{+0.008}_{-0.009}$ and a baryon density of $\Omega_b h^2 = 0.0224 \pm 0.0009$, one can infer the following 2-σ range for the density of cold dark matter: $\Omega_{CDM} h^2 = 0.1126^{+0.0161}_{-0.0181}$.

Within a given SUSY breaking scenario, it is possible to calculate the mass spectrum and decay widths of the SUSY particles, and use this information to obtain the relic density of the LSP. Past studies in the context of the minimal supergravity (mSUGRA) class of models have identified regions of the five dimensional mSUGRA parameter space in which the relic density of the LSP (usually the lightest neutralino $\tilde{\chi}^0_1$) is consistent with dark matter constraints[21], and recent studies carried out post-WMAP have narrowed these regions further[22]. There is a general tendency for SUSY models to produce too much dark matter in the form of LSP's, and thus the regions of parameter space that are consistent with WMAP are those in which one or more annihilation processes reduce the number of LSP's to within the acceptable limit. Examples of the regions allowed in the $m_0, m_{1/2}$ plane for certain configurations of the other parameters are shown in Figures 1.3 and 1.4, along with other constraints resulting from existing data that are less relevant to the current discussion.

There is now a lot of interest in examining the phenomenology of points in parameter space that satisfy the WMAP relic density constraint, in the hope that this will provide a useful guide for obtaining the correct parameter set should supersymmetry be observed at the LHC. Two such examples will be seen in Chapter 5.

1.3 Summary

A brief review of the state of the art in the quest to understand Nature on her smallest scales has been given. The SM describes a wide range of phenomena, but suffers from certain technical problems which can only be solved by the introduction of new physics at higher energy scales. Work at CERN is currently producing a particle accelerator whose arrival is anticipated with great excitement.

Supersymmetry has been introduced as a well-motivated extension of the current theory,
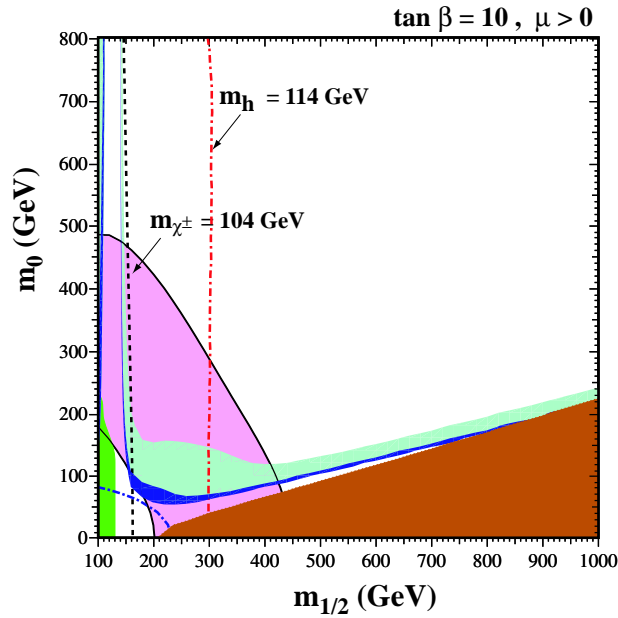
Figure 1.3: The $m_0, m_{1/2}$ plane of the mSUGRA parameter space showing the region consistent with the WMAP data[23], where the other mSUGRA parameters are fixed at the values shown (with $A_0 = 0$). The light blue region shows the region consistent with dark matter relic density measurements before WMAP, whilst the dark blue region is based on WMAP data. The region shaded in brown is disallowed due to the fact that the LSP is charged ($m_{\tilde{\tau}_1} < m_{\tilde{\chi}_1^0}$). The branching ratio of the process $b \rightarrow s\gamma$ receives additional loop corrections if SUSY exists, and hence measurements of the value can be used to exclude regions of the parameter space that are incompatible; these are shown in green. Measurements of the anomalous magnetic moment of the muon $g_\mu - 2$ favour some regions of the parameter space, and these are shown in pink. Finally, the dot-dashed line shows the LEP constraint on the $\tilde{e}$ mass whilst the black line shows the contour $m_{\tilde{\chi}^\pm} = 104$ GeV.
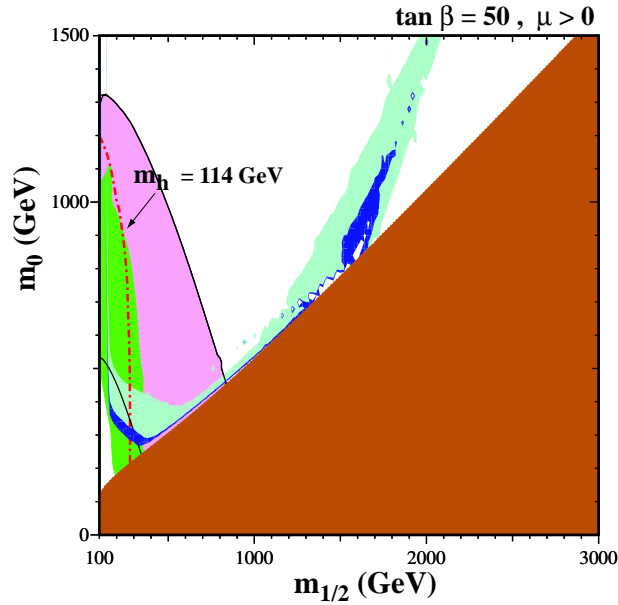
Figure 1.4: Similar to Figure 1.3, but with a higher value of tanβ. This gives rise to a so-called 'rapid annihilation funnel' in which we can satisfy the WMAP constraint through enhanced annihilation to third generation fermions via a heavy Higgs boson. This is enhanced when the heavy Higgs mass m_A is almost twice that of the $\tilde{\chi}_1^0$.

with a wide range of supporting theoretical arguments. Although these are not currently matched by experimental observation, it is expected that, if SUSY really is the natural answer to the problems of the SM, the sparticle masses will be around 1 TeV or less, and this strongly motivates the development of techniques for the measurement of SUSY particle masses and model parameters at the LHC.

SUSY is a broken symmetry, and our woeful ignorance of the SUSY breaking mechanism adds 105 parameters to the 19 parameters of the SM. Practically all phenomenological studies to date have been performed in simplified models with 5 or 6 parameters. Though it is important to use these to develop experimental techniques, there is no reason why Nature should be so kind as to heavily constrain the MSSM parameters in this way, and hence it is absolutely vital to develop techniques that are capable of exploring much more general models. This is the guiding principle of this thesis, in which an attempt is made to move toward a general approach to the problem of constraining the MSSM.

Chapter 2

ATLAS and the LHC

It has been shown that the current state of our knowledge of particle physics can be considered both comprehensive and frustratingly incomplete. It is against such a background that physicists are currently putting the finishing touches on the most advanced particle accelerator yet constructed- the Large Hadron Collider (LHC), located in Geneva, Switzerland. This chapter gives a brief summary of the LHC, with a specific focus on the ATLAS detector, within whose framework all analysis in this thesis will be based.

2.1   The LHC

Upon the commencement of data taking in 2007, the LHC will be largest operational particle accelerator and collider in the world. Protons will eventually be accelerated by superconducting magnets to an energy of 7 TeV, and will collide head on at four points around the 27km circumference tunnel previously used by the LEP electron-positron collider. Rather than having continuous beams, the collider operates by having protons bunched together into approximately 2800 bunches per beam so that interactions between the two beams occur at discrete intervals. At each of the four collision points, activity will be monitored by one of the four primary LHC experiments shown in Figure 2.1: ATLAS, CMS, LHCb and ALICE. The largest of these- ATLAS and CMS- are specifically designed to find new physics beyond the SM, and it is the ATLAS detector that forms the basis for the analysis in this thesis. It is worth noting, however, that the search reach for both experiments is similar, and there is no reason why the physics analysis presented here could not be repeated within the framework of CMS.

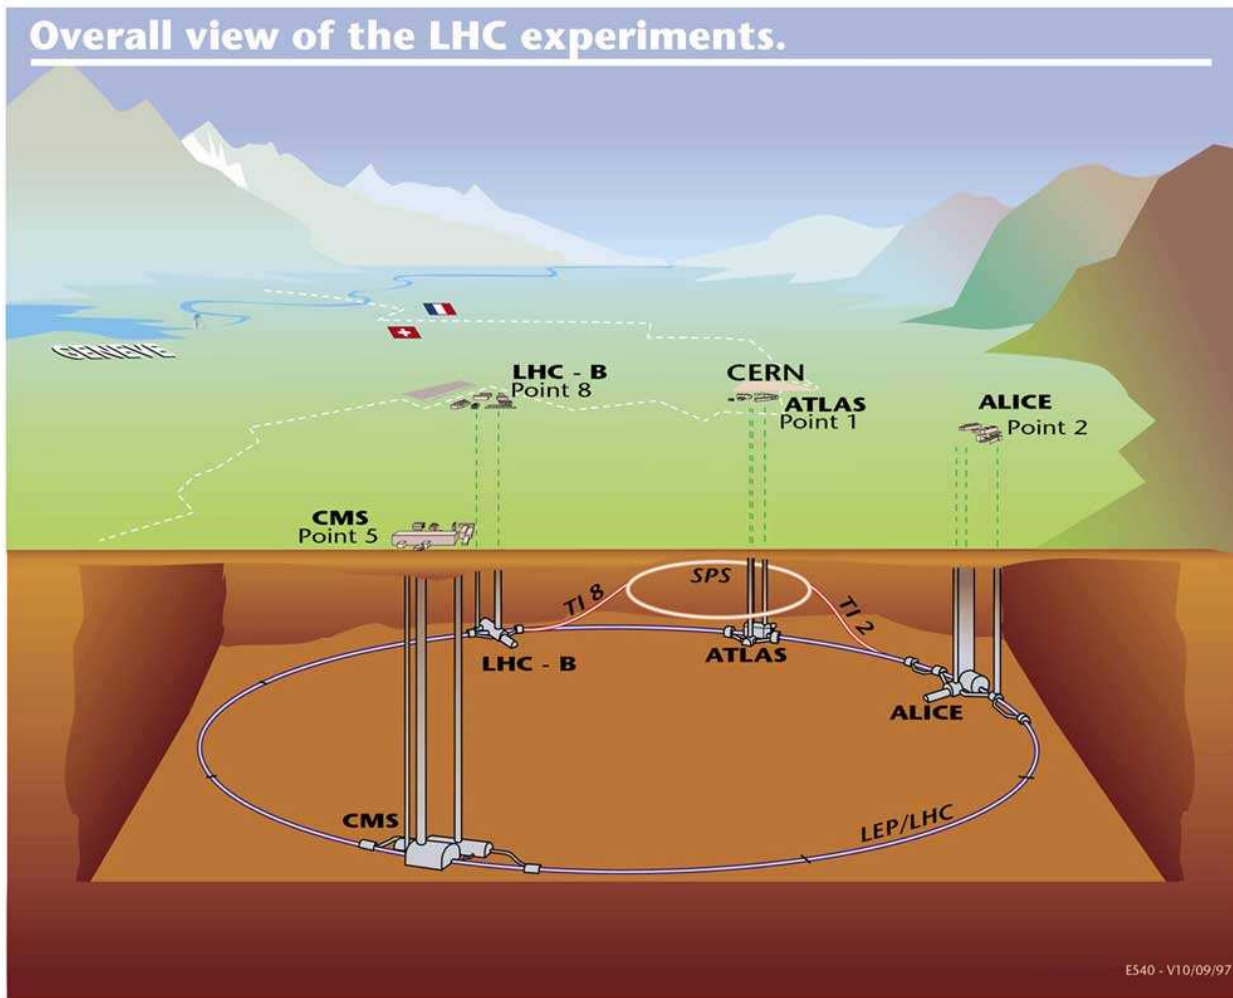At 14 TeV, the centre of mass energy of the LHC is over seven times greater than the previ-

Figure 2.1: The LHC above and below ground.

Energy at collision	7.0 TeV
Energy at injection	0.45 TeV
Machine circumference	26658.833 m
Dipole field at 7 TeV	8.33 T
Time between bunches (at $L = 10^{34}$ cm^{-2}s^{-1})	25 ns
Number of particles per bunch	1.15×10^{11}
Number of bunches per beam	2808
RMS bunch length	7.55 cm
Circulating beam current	0.582 A
Dipole magnet temperature	1.9K
Number of dipole magnets	1232
Number of quadrupole magnets	$\approx$600
Number of corrector magnets	$\approx$7000

Table 2.1: Parameters for the LHC running in proton-proton collision mode at design luminosity.

ous record held by the Tevatron at Fermilab[24]. This is a great step forward, but is of little use unless one can also guarantee a plentiful supply of events, as characterised in particle experiments by the luminosity. Starting in a low luminosity mode running at 1×10^{33} cm^{-2}s^{-1}, the LHC aims for a final design luminosity of 1×10^{34} cm^{-2}s^{-1}, thus comfortably exceeding the Tevatron peak luminosity of approximately 10^{32} cm^{-2}s^{-1}. Other important parameters of the LHC are given in Table 2.1. It is important to remark that the high luminosity leads to immense challenges in, amongst other things, event selection and offline data storage, and the successful solution of these problems will itself be a tremendous achievement. On average, 23 interactions occur in each bunch crossing, and though the majority of these are well-understood low energy processes, care must be taken that these "pile-up" events do not mask more interesting physics.

Owing to the great complexity of the LHC, there have been several delays during the course of the project. Nevertheless, at the time of writing it is possible to give a coherent account of the first run schedule which will see the beam pipe closed at the end of August 2007, followed within weeks by collisions at the injection energy (450 GeV in each beam) until the end of 2007. A subsequent shutdown will allow the remaining machine sectors to be commissioned without beam to full energy in preparation for the first physics run at 14 TeV. It is hoped that
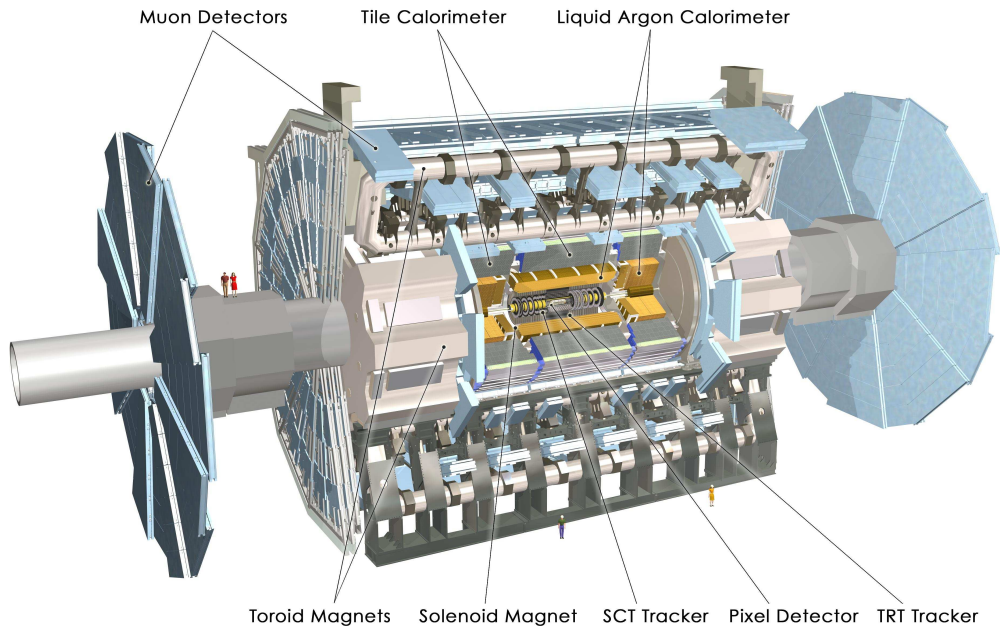
Figure 2.2: Labelled diagram of the ATLAS detector.

a substantial integrated luminosity will be obtained by the end of 2008.

2.2 The ATLAS experiment

2.2.1 Basic detector layout and geometry

The ATLAS detector of the LHC is shown in Figure 2.2. The beam direction defines the z-axis, whilst the positive x-axis points from the interaction point to the centre of the LHC ring and the positive y-axis points upward.

Given that the detector is almost exactly cylindrical, the most convenient coordinate system to use is (R, ϕ, z), where R is the transverse radius from the beamline, ϕ is the azimuthal angle (measured from the x-axis), and the origin of z is at the interaction point. One can also refer to the polar angle θ measured from the z-axis. In addition to these, it is useful to refer to a couple of other quantities that arise from a particular complication of hadron-hadron colliders. Hadrons are composite objects, and one does not know the original z-momentum of the constituents (or "partons") that interact in a given collision. For this reason, it is useful to define the true rapidity, $y = \frac{1}{2}\ln[(E + p_z)(E - p_z)]$, of a Lorentz vector, since differences in

this variable are invariant under longitudinal Lorentz boosts. This is particularly convenient if the mass of a particle is known; if not, one can refer to the pseudo-rapidity, $\eta = -\ln[\tan(\theta/2)]$, which closely approximates the true rapidity in the relativistic limit.

Particles are often described by the parameters (p_T, η, ϕ), where p_T is the transverse momentum with respect to the beam axis. Particles in the $\eta - \phi$ plane are then separated by a distance $\Delta R = \sqrt{\Delta \eta^2 + \Delta \phi^2}$. Particles travelling toward higher $|\eta|$[1] would be described as heading for the 'forward' region.

2.2.2 Physics programme and detector concept

As one of the two LHC "generic" physics detectors, the ATLAS experiment was designed to accommodate as broad a programme of investigation as possible. Work will encompass both precision tests of the standard model, and searches for physics beyond it, and the rich promise of ATLAS in both areas is best summarised in the two volumes of the *ATLAS Detector and Physics Performance Technical Design Report*[25]. The brief highlights are as follows:

1. **The Higgs Boson:** We saw in Chapter 1 that the SM, like the LHC, is currently incomplete. The introduction of particle masses requires interaction with a scalar field whose associated boson has yet to be observed. The Higgs boson is therefore one of the most sought after particles at the LHC, and is one of the chief goals of the experiment. Fortunately, ATLAS is expected to be capable of detecting an SM Higgs boson with a high significance ($> 5\sigma$) in the entire mass range allowed by the theory. If a Higgs is not observed, it is hoped that ATLAS will observe another mechanism of electroweak symmetry breaking such as Technicolour.

2. **CP Violation:** Matter and antimatter do not behave in the same way, and the difference in their behaviour is known as CP-violation. Current experiments designed to probe this phenomenon, such as Babar and Belle, have not yet observed a sufficiently strong effect to explain the apparent absence of antimatter in our observable universe. New physics observations may shed light directly on this problem, and measurements of the decays of B mesons will provide indirect constraints on potential new theories.

3. **The top quark and other particles:** Though long predicted in the SM, the top quark was discovered only as recently as 1995 at Fermilab, and measurements of its properties

---

[1] 'High' in this instance means that the particles head toward the end-cap detectors rather than the barrel region, and thus $|\eta|$ is approximately greater than 1.

lack the precision obtained in other areas of particle physics. The cross-section for $t\bar{t}$ pair production at the LHC has been calculated as 833^{+52}_{-39} pb[26], and thus the LHC will produce a very large number of $t\bar{t}$ pairs, allowing ATLAS to obtain much more precise measurements of its mass and interactions than those currently available. For example, a day of running at low luminosity will produce about 10 pb^{-1} of data, and thus over 8000 top quarks. It is hoped that this extra precision will allow us to more rigorously examine inconsistencies in the SM. The same logic applies to other known particles such as the *W* boson, whose mass is to be sought with a precision of 20 MeV.

4. **New physics searches:** Chapter 1 introduced supersymmetry as a viable extension of the SM and the use of ATLAS as a powerful tool for the measurement of SUSY parameters is the primary topic of this thesis. Other possibilities for observable new physics include the discovery of extra spatial dimensions, which may or may not be associated with mini black hole production; technicolour; compositeness and the so-called 'little Higgs' models.

5. **Strong interactions:** QCD provides a good description of the basic principles behind the nuclear strong force. However, interactions of interest are often outside the applicable range of perturbation theory and are thus very difficult to quantify. Measurements of the strong coupling constant α_S and parton density functions will therefore be of great use.

In order to maximise the potential for successfully completing the tasks listed above, the basic design criteria of the ATLAS experiment include very good electromagnetic calorimetry for measurements of photons and electrons, and full-coverage hadronic calorimetry for jet and missing transverse energy measurements. Muon momentum measurements can be carried out to relatively high precision, and at high luminosity one can obtain a muon momentum resolution of $\approx$10% for momenta of 1 TeV using the external muon spectrometer alone. Efficient tracking is expected at both high and low luminosity, thus fulfilling an important requirement for high-p_T lepton-momentum measurements, electron and photon identification and heavy flavour identification. The detector has a large acceptance in η and almost full coverage in ϕ, whilst also being able to trigger at low p_T thresholds (this is important for many physics processes).

There follows a brief review of the various subsystems of the ATLAS detector, summarising all basic points relevant to the analysis that follows. For more detailed discussions, a good starting point is reference [25], from which one can also find all previous technical design reports on the ATLAS experiment.

2.2.3   The magnet systems

ATLAS contains two magnet sub-systems- a central solenoid that surrounds the inner detector, and a toroid system, comprising one barrel toroid and two end-cap toroids, that generates a magnetic field for the muon spectrometer.

The central solenoid provides a field of 2 T (along the beam axis) for the inner detector, and the coil is designed to be as thin as possible, without compromising performance and reliability, in order to reduce the amount of material obscuring the calorimeters. Thus, despite a diameter of 2.5 m and a length of 5.3 m, the thickness is a mere 45 mm. The system is designed to allow accurate measurements of charged particles in the inner detector up to a momentum of 100 GeV.

Each of the three toroids consists of eight coils assembled radially around the beam axis, with the 8 'race track' magnets of the barrel toroid mimicked by smaller versions in each of the two end-caps. The immense size of the barrel toroids gives rise to one of the most distinctive visual features of the detector, as the magnets are a staggering 25 m long and 5 m wide. The end-cap toroids are a more modest 5 m in length, and are rotated by 22.5 deg with respect to the barrel toroid coil system in order to provide radial overlap and to improve bending power in the interface regions of the coil systems.

The 1300 tonnes of the ATLAS magnet system present a considerable engineering challenge. Cooling the system to its operational temperature of 4.8 K will take approximately 30 days, and the system will store 1600 MJ of energy when operational. The barrel toroid is mostly operational over a range $|\eta| \leq 1.0$, whilst the end-cap magnets dominate in the range $1.4 \leq |\eta| \leq 2.7$. In the intermediate range, bending is provided by both systems and the whole configuration is such that the field is mostly orthogonal to muon trajectories. The average field in the muon spectrometer is 0.6 T, allowing muons with momenta of up to 6 TeV to be measured.

2.2.4   Inner Detector

Accurate observations of particles in accelerator experiments rely on precise measurements of a particle's momentum and charge, whilst also requiring information such as whether a particle was produced at the primary interaction vertex or at a secondary vertex (the latter is important for b-tagging). The ATLAS inner detector is designed to perform all of these measurements by observing the tracks of particles in the 2 T magnetic field of the central solenoid. The
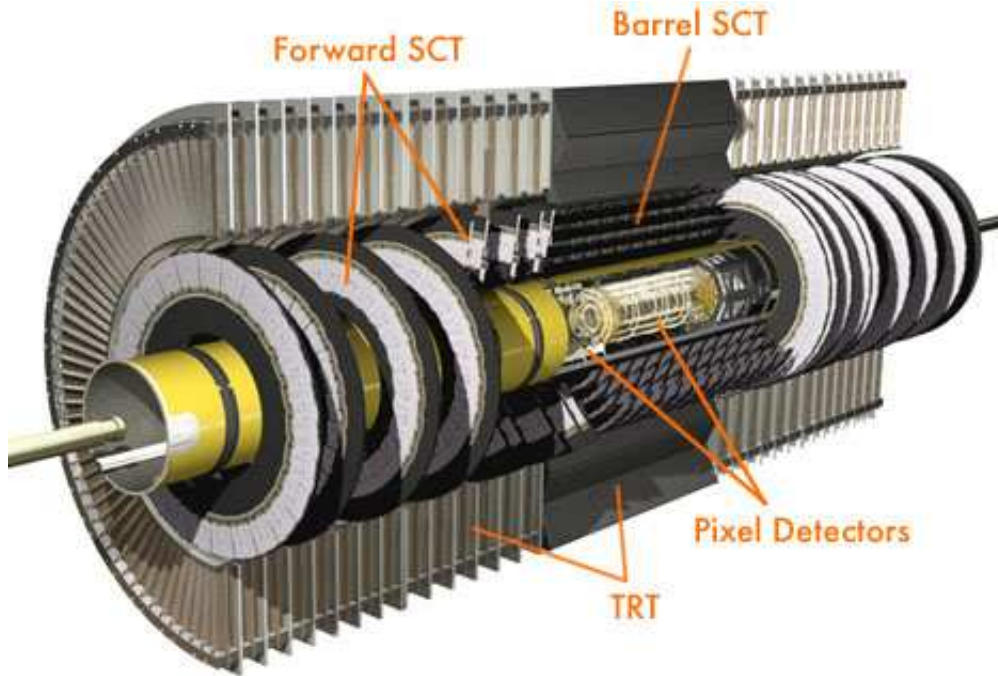
Figure 2.3: The ATLAS inner detector[28].

field is along the beam axis and hence particle tracks are bent according to their transverse momentum. The momentum resolution of the inner detector can be parameterised as[27]:

$$\sigma(1/p_T) = 0.00036 \oplus \frac{0.013}{p_T \sqrt{\sin\theta}} [\text{GeV}^{-1}] \tag{2.1}$$

Obtaining high precision results in both momentum and vertex resolution requires fine-granularity detectors, particularly in an environment with as many tracks as the high luminosity LHC. Semiconductor detectors can, and will, be used in the ATLAS inner detector to meet this criterion though they are not used alone due to several factors. Cost is one consideration, and one must also take note of the fact that the inner detector is placed inside the calorimeters (which measure particle energy), and that these must be placed nearer to the vertex for optimum performance. Thus the SCT must use the minimum possible material.

The ATLAS inner detector (shown in Figure 2.3) is divided into three separate parts: a semiconductor pixel detector that provides high granularity near the vertex region, a semiconductor tracker (SCT) that utilises silicon micro-strip technology, and a cheaper straw tube tracker (TRT) that provides continuous track-following with much less material per point. The detector consists of a barrel region and two end-cap regions, and the design is such that

no single component dominates the momentum measurement.

The aim here is to give a brief introduction to the workings of the ATLAS inner detector, concentrating on the details most relevant to physics analysis. Much more complete information is given in the Technical Design Reports (TDRs) of the inner detector[28], pixel detector[29], and overall detector and physics performance[25].

The Pixel detector

Situated closest to the beam pipe, the pixel detector is built from modules that contain 46,080 pixels of size 50×400 μm, providing high-precision measurements with high granularity as close to the interaction point as possible. It therefore bears the brunt of both the high particle flux density at the LHC, and the extreme radiation environment, and the permanent components must be capable of dealing with the equivalent of 5×10^{14} neutrons per cm^2 over a 10 year lifetime. The pixels are smaller in the $R\phi$ direction in order to improve measurement of the sagitta (deviation from a straight line) of particle tracks, thus enhancing the momentum measurement. The detector mostly determines the impact parameter resolution of ATLAS, and the ability to find B hadrons in the inner detector.

The detector consists of three barrel layers, located at radii of 50.5 mm (B-layer), 88.5 mm (1st layer) and 122.5 mm (second layer), and six disk layers (three on each end-cap). The 'B-layer' is designed to be replaced every few years as a result of the received radiation dose being roughly four times higher than that received by the 1st layer. The modules, which form the basic building blocks of the detector, are identical in all regions, and each comprises of a silicon pixel array bump-bonded to 16 front-end integrated read out circuits. These read out chips each serve an array of 18×160 pixel diodes, and their basic function is to compare the electrical signal in the diodes to a threshold in order to provide a binary output.

It should be noted that, due to delays in the provision of the radiation-hard integrated-circuit electronics, the layout of the pixel detector has changed since its inception, with the notable changes recorded since the TDR detailed in reference[30]. At the time of writing, the pixel detector still lies on the critical path.

The SCT

The SCT uses the same basic technology as the pixel detector, but can afford to use fewer readout channels owing to the lower track density at a larger value of R. Thus, the detectors

of the SCT are strips rather than pixels, and the modules are arranged into four barrel layers (each with two sides) and 2 end-cap sections each containing 9 disks.

Each module of the barrel SCT consists of four 63.6×64 mm wafers with 80 μm pitch strips, with two wafers placed end-to-end on each side to give an effective length of 123.2 mm. The two pairs of wafers are placed back-to-back with a stereo angle of 40 mrad between them to give some measurement along the z-direction, whilst the modules are arranged to give an $R\phi$ precision per layer of the order of 17 μm. The end-cap modules are similar, but have one set of tapered strips aligned radially. The output from each module is binary as in the pixel detector.

A more complete description of the SCT is given in Chapter 3. Changes since the TDR are noted in the Final Design Reports[31].

The Transition Radiation Tracker

The Transition Radiation Tracker (TRT) is a drift tube system, with a design that incorporates "straw" detectors. Each of these is a small cylindrical chamber of 4 mm diameter filled with a gas mixture of Xe, CO_2 and O_2, in which the aluminium coated inner wall acts a cathode whilst a central gold-plated tungsten wire acts as an anode. Charged particles passing through ionise the gas, and the resulting ionisation cluster is amplified by a factor of $\approx 2.5 \times 10^4$ whilst drifting through the electric field in the straw. The wires are split in half at the centre and read out at each end, and each channel provides a drift time measurement. The expected spatial accuracy is 170 μm per straw.

The space between the straws is filled with a polypropylene/polyethylene fibre radiator which increases the amount of transition radiation photons produced in the detector. These are produced when relativistic particles cross a boundary between materials with different dielectric constants, and the threshold above which radiation is produced is dependent on $\gamma = (1 - \frac{v^2}{c^2})^{-1/2}$ where v is the particle velocity and c is the speed of light. The Xe in the straw gas presents a high interaction cross-section to these photons and a signal is produced which has a higher amplitude than the signal arising from minimally ionising particles. There are thus two different categories of signal that one wishes to detect in each straw, and for this reason each channel has two independent thresholds. The lower threshold detects the tracking hits, of which an average of 36 are expected for each particle with $p_T > 0.5$ GeV, and the higher threshold is designed for the transition radiation photons. This higher threshold aids particle identification, as, for example, electrons start producing transition radiation when their
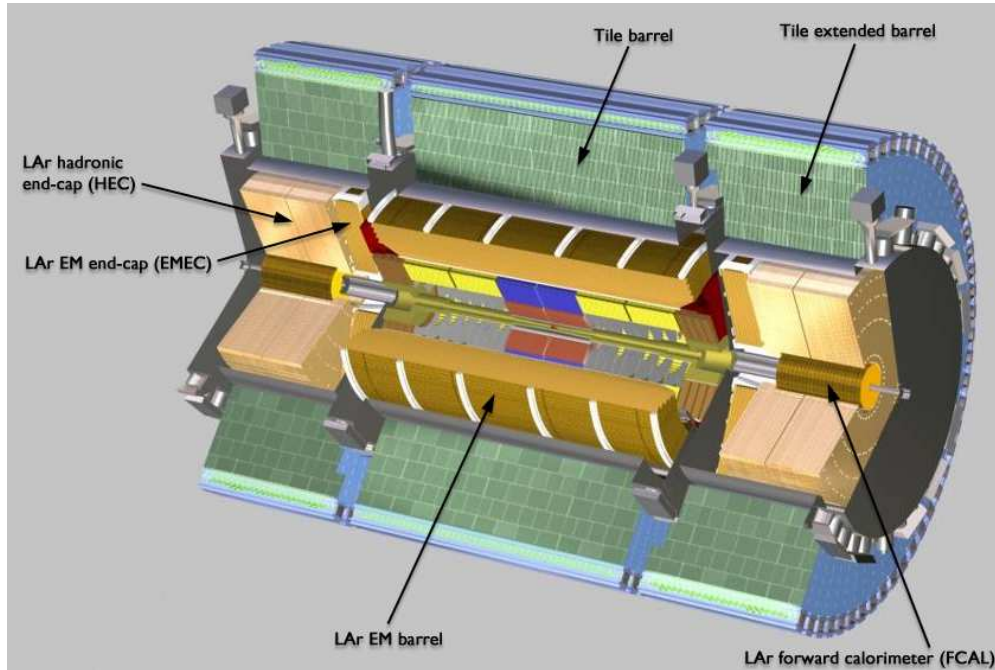
Figure 2.4: Labelled diagram of the ATLAS calorimeters.

momentum is close to 1 GeV, whilst pions start to radiate only when their momentum is close to 100 GeV. At ATLAS, the pion rejection is expected to be ≈100 for an electron efficiency of 90%.

The TRT is intrinsically radiation hard, and provides a large number of measurements at relatively low cost. A barrel region covers $|\eta| < 0.7$ and 18 wheels in each end-cap extend coverage to $|\eta| = 2.5$. In the current schedule, only 14 wheels of the end-cap will be present at start up, reducing the initial coverage to $|\eta| < 2.0$. For a recent review, see[32].

2.2.5 Calorimeters

The ATLAS calorimeters (see Figure 2.4) are situated outside the central solenoid that surrounds the inner detector, and their primary job is to precisely measure the energy of electrons, photons and jets. High energy particles entering a calorimeter produce a cascade of secondary particles known as a 'shower'. The incoming particle interacts via either the electromagnetic or strong interaction to produce new particles of lower energy which react in a similar fashion, producing very large numbers of particles whose energy is deposited and measured. Calorimeters also play an important role in determining the position of particles, measuring the missing

transverse momentum per event, identifying particles and selecting events at the trigger level.

There are two calorimeters in ATLAS, due to the fact that the interaction with matter for electrons and photons is different from that of hadrons. Electrons and photons penetrate much less deeply than hadrons, and produce narrower showers. Energy loss occurs predominantly via bremsstrahlung[2] for high energy electrons (which for most materials means energies greater than ≈ 10 MeV), and high energy photons lose energy via the related process of e^+e^- pair production. The characteristic amount of matter traversed by a particle before undergoing one of these interactions is described by the 'radiation length' X_0, and is entirely set by the properties of the material being traversed. The expectation values of the energy of an electron $E(x)$ and the mean number of photons $N_\gamma(x)$ as a function of the distance x into the material are given by:

$$\langle E(x) \rangle = E_0 e^{-\frac{x}{X_0}} \tag{2.2}$$

$$\left\langle N_\gamma(x) \right\rangle = N_0 e^{-\frac{7}{9}\frac{x}{X_0}} \tag{2.3}$$

Hadrons interact with the nuclei of the calorimeter material via the strong force, and the resulting showers are characterised by a nuclear interaction length λ which is typically an order of magnitude greater than X_0. This length is a function of both the energy and type of incoming particle, since it depends on the inelastic cross-section for nuclear scattering. Longitudinal energy deposition profiles have a maximum at:

$$x \approx 0.2\lambda \ln(E_0/1\text{GeV}) + 0.7 \tag{2.4}$$

where x is the depth into the material and E_0 is the energy of the incident particle. The depth required for containment of a fixed fraction of the incident particle energy is also logarithmically dependent on E. The energy deposit in a hadronic cascade consists of a prompt EM component due to π^0 production, followed by a slower component due to low-energy hadronic activity. These two different types of energy deposition are usually converted to electrical signals with different efficiencies, the ratio of which is known as the intrinsic e/h ratio.

Both ATLAS calorimeters are sampling calorimeters; that is the material that produces the particle shower is distinct from that used to measure the energy and it is thus through periodic sampling of the particle shower that an energy measurement is obtained.

---

[2]Bremsstrahlung (which translates as 'braking radiation') is produced by the acceleration of a charged particle after deflection by another charged particle.

The Electromagnetic Calorimeters

The Electromagnetic CALorimeter (ECAL) uses lead and stainless steel to absorb energy, with liquid argon as the sampling material. This LAr sampling technique is radiation resistant and combines good energy resolution with other attractive features such as long-term stability of the detector response and relatively easy detector calibration. The geometry of the detector describes a complex accordion shape, shown in Figure 2.5, chosen to provide complete ϕ symmetry with no azimuthal cracks. As in the inner detector, the calorimeter has a barrel region (covering the range $|\eta| < 1.475$) and end-cap regions (covering the range $1.375 < |\eta| < 3.2$), and the lead thickness in the absorber plates has been optimised as a function of η to maximise the energy resolution performance.

No discussion of the ATLAS ECAL is complete without mention of two points relevant to physics analysis: there is a crack in the coverage at $|\eta| = 1.5$ due to the barrel/end-cap transition, and a small crack at $|\eta| = 0$ arising from the fact that the barrel calorimeter is constructed of two identical half barrels separated by 6 mm at $z = 0$. The ECAL is 'non-compensating'; that is, it responds differently to electromagnetic and hadronic showers as the intrinsic e/h ratio is not equal to unity. This difference is corrected for in the reconstruction software.

We have already seen that the inner detector sits between the interaction point and the calorimeters, and the total material seen by an incident particle before it reaches the front face of the calorimeter is $\approx 2.3X_0$ at $\eta = 0$. This amount increases with η because of the particle angle. A presampler consisting of an active LAr layer of thickness 1.1 cm (0.5 cm) in the barrel (end-cap) region is used to correct for the energy lost by electrons and photons en route to the ECAL in the region $|\eta| < 1.8$. In the transition region between the barrel and end-cap, this is supplemented by a scintillator slab, as the amount of obscuring material reaches a localised maximum of $\approx 7X_0$.

The total thickness of the ECAL is above 24 X_0 in the barrel, and above 26 X_0 in the end-caps. As shown in Figure 2.5, the system is divided into three longitudinal sections ("samplings"). The overall design utilises high granularity to provide precise position measurements, and the information obtained by the presampling and the first sampling layer provides excellent γ/π^0 and e/π separation by measuring the transverse energy profile of the preshower. Table 2.2 gives the resolution in the position and angle measurements for photons of E_T=50 GeV at $\eta = 0.3$, assuming that the readout electrodes of the EM calorimeter are aligned in z with the expected precision of 400 μm. The π^0 rejection factor under the same conditions is 3.50 ± 0.17 at $\eta = 0.3$ for 90% photon efficiency.
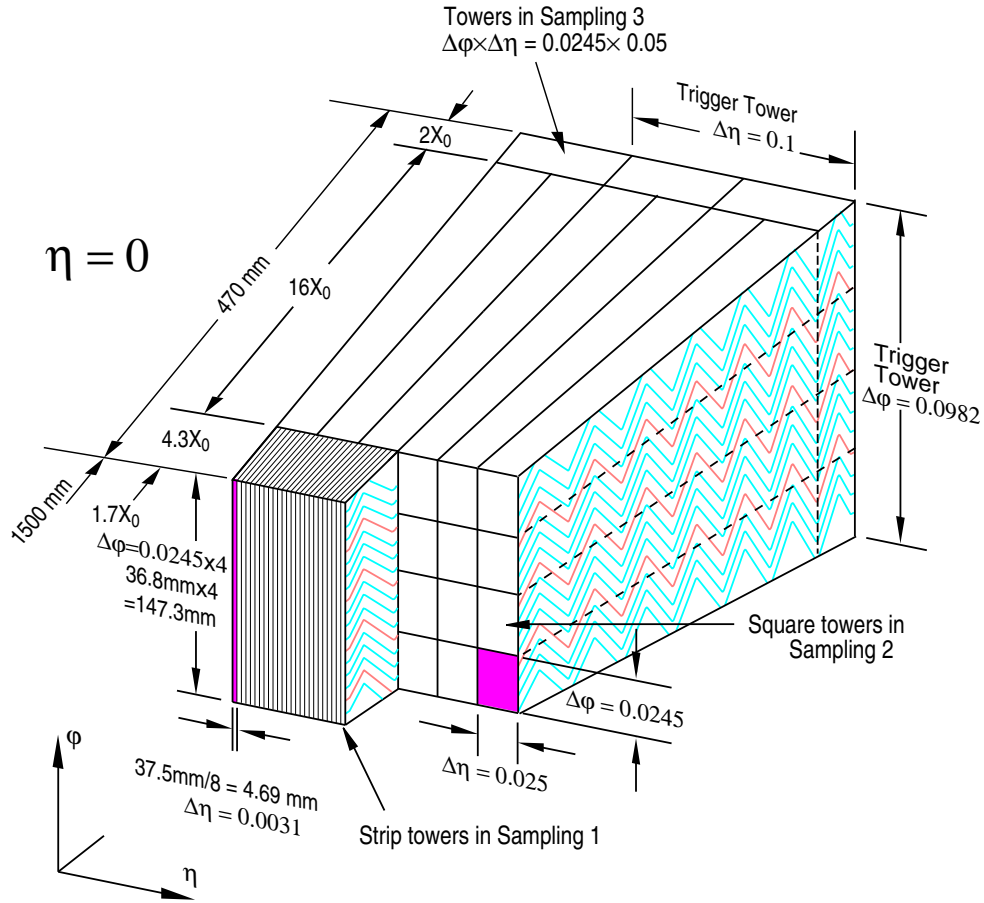
Figure 2.5: Geometry of the electromagnetic calorimeter[35].

The energy resolution σ_E/E of a calorimeter can be parametrised as $a/\sqrt{E} \oplus b \oplus c/E$. The first term is a stochastic term representing statistical fluctuations arising in both the shower development and the sampling procedure. The constant term b arises from detector non-uniformity and calibration uncertainty, and can in principle be increased by radiation damage of the active medium. This is not expected to be a problem at ATLAS[33]. The final term is generated by electronic noise in the readout chain[34]. The energy resolution of the ATLAS ECAL is dominated by sampling fluctuations, and is close to $\approx 0.10 \; \mathrm{GeV}^{\frac{1}{2}}/\sqrt{E(\mathrm{GeV})}$ for the entire range of η coverage, with a constant term of less than 0.007. The c term is expected to be 0.245 GeV$/p_T$.

σ_η strip section	$(0.215 \pm 0.008) \times 10^{-3})$
σ_η middle compartment	$(0.652 \pm 0.021) \times 10^{-3})$
σ_θ (mrad)	(7.180 ± 0.260)

Table 2.2: Resolution of the η and θ measurements for photons of $E_T = 50$ GeV at $\eta = 0.3$, assuming that the readout electrodes of the EM calorimeter are aligned in z to within 400 μm.

The Hadronic Calorimeters

The hadronic calorimeters cover the range $|\eta| < 4.9$, and must use different techniques over this range to cope with the variation in the performance requirements and radiation environment. A tile calorimeter is used for $|\eta| < 1.7$, using iron as the absorber and scintillating tiles as the active material. The system consists of one barrel and two extended barrels, and is longitudinally segmented in three layers. For $\approx 1.5 < |\eta| < 4.9$ LAr calorimeters (similar to those in the ECAL) are used, with the system comprising of a Hadronic End-Cap calorimeter (HEC) extending to $|\eta| < 3.2$ and a high density forward calorimeter (FCAL) covering the range $3.1 < |\eta| < 4.9$. The FCAL is divided into three sections, one of which uses copper (as does the HEC) and two of which use tungsten for its higher density. A summary of design parameters for the hadronic calorimeters is provided in Table 2.3.

The hadronic calorimeter must be thick enough to provide containment for hadronic showers and to keep punch-through into the muon system to a minimum. A thickness of about 10 interaction lengths has been shown to be sufficient to reduce punch-through to a sufficient level, and this feature combines with the large η coverage to help obtain a good measurement of the missing transverse energy in each event. This is particularly important for experimental SUSY searches, and is thus highly relevant to this thesis. The fact that good calorimeter coverage is essential has been shown by a study of $A \rightarrow \tau\tau$ events with $m_A = 150$ GeV, performed at the particle level. The resolution of each component of the E_T^{miss} vector was shown to degrade from 2.3 GeV to 8.3 GeV if the calorimeter coverage is reduced from $|\eta| < 5$ to $|\eta| < 3$. This significantly increases the width of the measured A mass peak.

The HCAL, like the ECAL, is non-compensating. Chapter 9 of the ATLAS Detector and Physics Performance TDR lists an expected jet energy resolution for the full calorimeter (i.e. ECAL and HCAL) of $\sigma/E = (48.2 \pm 0.9\%)/\sqrt{E} \oplus (1.8 \pm 0.1\%)$ for $|\eta| = 0.3$ (barrel region), and $\sigma/E = (55.0 \pm 2.5\%)/\sqrt{E} \oplus (2.2 \pm 0.2\%)$ for $|\eta| = 2.45$ (end-cap region).

HADRONIC TILE	Barrel	Extended Barrel
Coverage	$|\eta| < 1.0$	$0.8 < |\eta| < 1.7$
Longitudinal segmentation	3 samplings	3 samplings
Granularity($\Delta\eta \times \Delta\phi$)		
Samplings 1 and 2	0.1×0.1	0.1×0.1
Sampling 3	0.2×0.1	0.2×0.1
HADRONIC LAr		End-cap
Coverage		$1.5 < |\eta| < 3.2$
Longitudinal segmentation		4 samplings
Granularity($\Delta\eta \times \Delta\phi$)		0.1×0.1 for $1.5 < |\eta| < 2.5$
		0.2×0.2 for $2.5 < |\eta| < 3.2$
FORWARD CALORIMETER		Forward
Coverage		$3.1 < |\eta| < 4.9$
Longitudinal segmentation		3 samplings
Granularity($\Delta\eta \times \Delta\phi$)		$\approx 0.2 \times 0.2$

Table 2.3: Design parameters for the hadronic calorimeters[25].

2.2.6   The Muon Spectrometer

The ATLAS muon spectrometer is essentially a huge straw tracker, with two primary tasks. The first is to precisely measure the momentum of high energy muons that are not measured well by the inner detector, and this is performed using monitored drift tubes (MDTs) over most of the η range, and cathode strip chambers at large η and at close proximity to the interaction point. These detectors are all referred to as the 'precision chambers'. The second task is to trigger on high energy muons (for a description of the trigger, see Section 2.2.7), and this is done using resistive plate chambers (RPCs) in the barrel and thin gap chambers (TGCs) in the end-caps. Triggering requires less precise position data, but a response time better than the LHC bunch spacing of 25 ns, thus explaining the use of a separate technology for the purpose. The trigger system covers the pseudo-rapidity range $\eta \leq 2.4$. A schematic view of the layout of one quadrant on the muon spectrometer, showing the position of the various detector technologies, is shown in Figure 2.6.

The MDTs are aluminium tubes of 30 mm diameter and 400 μm wall thickness, with a 50 μm diameter central W-Re wire, and are commercially available. The single wire resolution is expected to be 80 μm. These are arranged into chambers, each of which contains
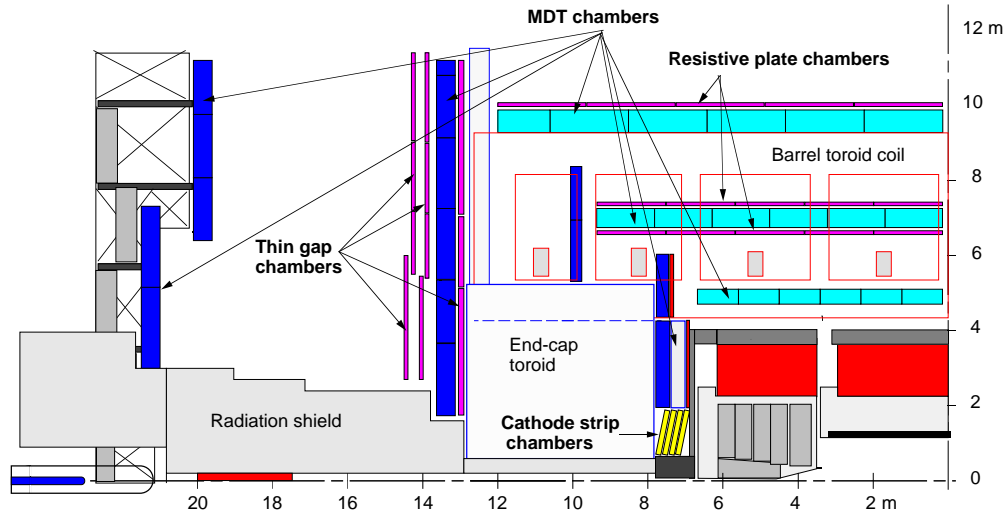
Figure 2.6: A view of one quadrant of the ATLAS muon spectrometer[36].

multilayer pairs of either three or four mono-layers of tubes. The CSCs are fast multiwire proportional chambers, with parallel high voltage wires strung in a gas volume and closed by conducting planes at 0 V. Precise position measurements along the wires are obtained by determining the centre of gravity of the charge induced on the strips of the conducting plates. Position resolutions of better than 60 μm are expected, and the high granularity of the system ensures that it will cope with the demanding rate and background conditions at high η. Meeting the momentum resolution requirements in the muon spectrometer requires that the relative positioning of chambers matches the intrinsic resolution of the precision chambers. For this reason, an optical alignment system will be used to constantly monitor chamber deformations and displacements.

Over most of the muon spectrometer, the magnetic field is essentially in the ϕ direction, and thus muons bend in the $R-z$ plane. The MDTs and CSCs provide precision measurements in this plane.

The RPCs of the trigger system are made of a pair of parallel plates separated by a narrow gas gap, with a high electric field applied. Muons passing through release ionisation electrons, and these form an avalanche between the plates whose signal is read out by two sets of strips in orthogonal directions. The expected space-time resolution is 1 cm × 1 ns, and the measurements will be used both to complement the MDT tracking data (the RPC measurement is in an orthogonal direction), and for the level 1 trigger. The TGCs in the end-caps are similar in design to multi-wire proportional chambers, though the anode wire pitch is larger.

The muon spectrometer should give three precision measurements for $|\eta| < 2.7$, except in crack positions (such as at $|\eta| = 0$). A 20 GeV muon will give a momentum measurement at 2% precision with good acceptance. This rises to 10% as the energy approaches 1 TeV.

2.2.7 Trigger and Data Acquisition System

It has already been noted that the number of interactions at the LHC will be very large indeed, and that a majority of these are of little interest to those chasing the frontier of particle physics. Thus, a detector such as ATLAS does not record every interaction but instead is designed to trigger on interesting processes. Limitations on the amount of data that can be stored require an initial bunch-crossing rate of 40 MHz to be reduced to a rate of selected events of 100 Hz, and the challenge is to do this without missing any of the rare new physics processes that motivate the entire experiment!

The ATLAS trigger and data-acquisition system (DAQ) , shown pictorially in Figure 2.7 is based on three levels of online event selection, with each trigger level refining the decision made at the previous level through the use of additional selection criteria. The process will start with the level-1 (LVL1) trigger which uses only reduced-granularity information from a subset of detectors, identifying combinations of the presence or absence of objects such as high p_T muons, electrons and photons, jets, hadronic τ-lepton decays and large missing and total transverse energies. The time from a proton-proton collision until the availability of the LVL1 trigger decision at the front-end electronics must not exceed 2.5 μs. During this time, information for all detector channels is stored in 'pipeline' memories placed on or close to the detector in harsh radiation environments. The pipeline lengths must be kept as short as possible for reasons of cost and reliability and hence this imposes constraints on the LVL1 latency. One must also consider the time required for signals to leave the detector via cabling. The LVL2 trigger system will use region of interest (RoI) information, provided by the LVL1 trigger. This permits the LVL2 system to have access to the full detector information (i.e. full granularity and inner detector tracking information which is excluded at LVL1), since it need only read out a fraction of the whole detector. Finally, the last stage of the trigger process is the event filter (EF), which makes use of fully reconstructed events in order to decide which should be committed to permanent storage for subsequent offline analysis.
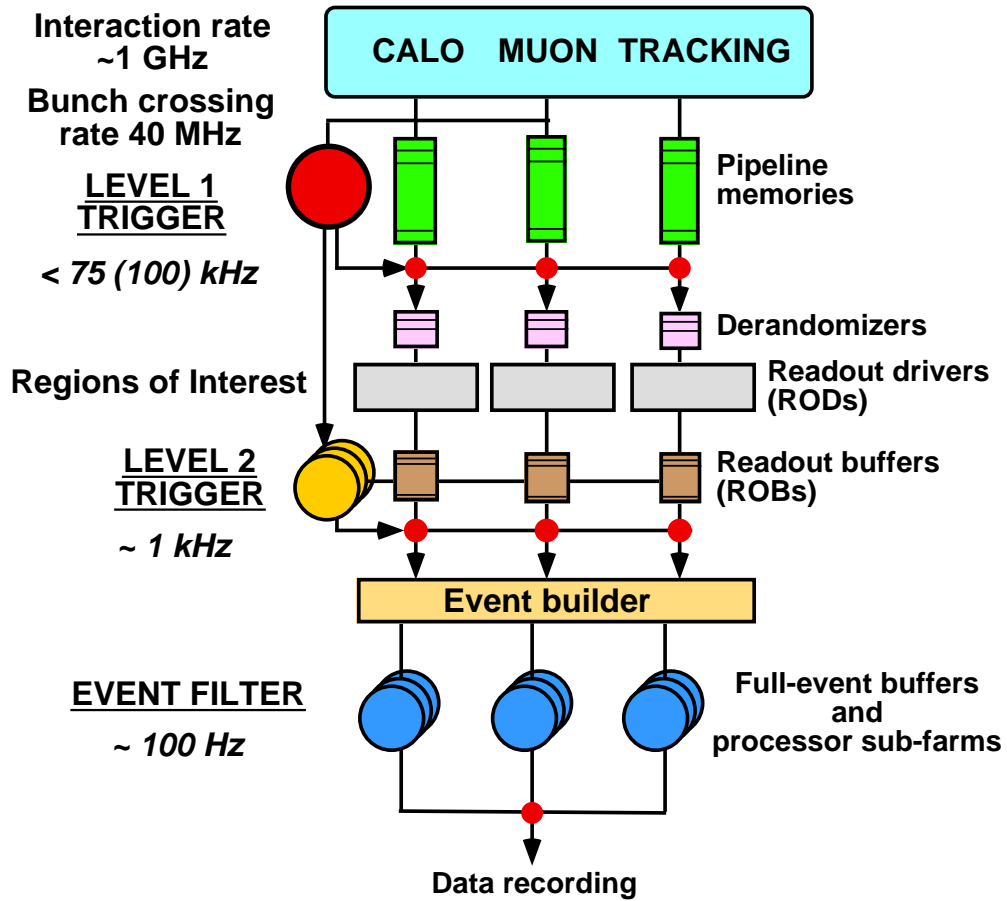
Figure 2.7: Block diagram of the ATLAS trigger/DAQ system[25].

2.3   Summary

A short description of the LHC and ATLAS has been provided, illustrating not only the tremendous challenges involved in getting the experiment up and running, but also the strong potential for seeing genuinely new and exciting physics. With less than a year to completion (not counting staged components!) there is now a great deal of excitement amongst both theorists and experimentalists, and we are undoubtedly approaching one of the great periods in particle physics.

The next chapters of this thesis will cover part of the ATLAS detector in more detail, through a description of software work performed in conjunction with the SCT monitoring group.

Chapter 3

The ATLAS Semiconductor Tracker

The ATLAS detector was described in Chapter 2, along with the ATLAS Semi-Conductor Tracker (SCT). The purpose of this chapter is to cover the design and operation of the SCT in sufficient detail to understand the monitoring software developed in Chapter 4. The focus throughout will be on the SCT barrel detector, and the discussion will concentrate mostly on factors that are relevant to noise occupancy.

A thorough overview of the design of the SCT is given in the SCT TDR, located in volume 2 of the ATLAS Inner Detector TDR [28], whilst the barrel modules are covered in reference[37].

3.1   Layout of the SCT barrel region

The SCT barrel region is divided into four separate layers, referred to as barrels 3-6 (with the pixel layers providing barrels 0, 1 and 2). Each layer is assembled from identical modules, mounted in a tile arrangement to minimise dead regions. In order to minimise the adverse effects of radiation damage, chiefly that due to anti-annealing of the silicon (see reference [38]), the module sensors will be operated at $\approx -7°$ C. Table 3.1 summarises important barrel parameters whilst Figure 3.1 displays the geometry of the modules for one quadrant of the detector. The barrel region covers a pseudo-rapidity range $|\eta| < 1.4$, and has 3.2×10^6 channels in total.

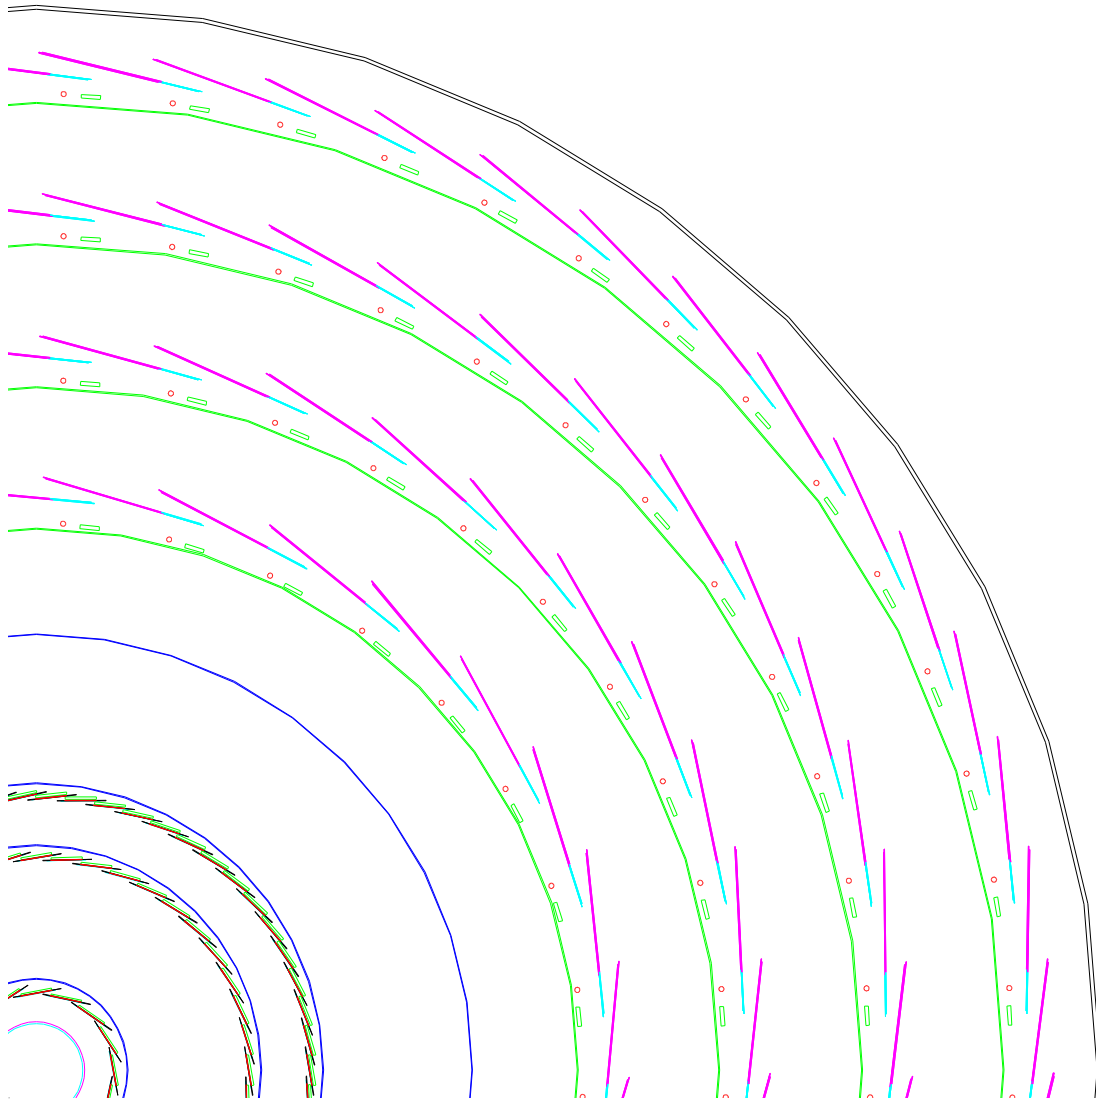**Figure 3.1:** Transverse view of a quadrant of the ATLAS inner detector (without the TRT). The outer four layers show the four barrels of the SCT, with support structures shown in green, lumped power cables and cooling shown as red circles, active silicon shown in pink, and electronics boards shown in pale blue[28]. The inner three layers show the pixel detector. The distance from the centre to the outermost layer shown is 514 mm.

Barrel cylinder	Radius (mm)	Full length (mm)	Number of modules
Barrel 3	299	1492	384
Barrel 4	371	1492	480
Barrel 5	443	1492	576
Barrel 6	514	1492	672
Total			2112

Table 3.1: SCT barrel parameters[37].

3.2 SCT Module Design

3.2.1 Silicon Detector Basics

The basic component of an SCT module is a silicon detector strip, formed from a reverse-biased p-n junction. In a semiconductor, there is a small gap (3.6 eV) between the highest filled energy level (the 'valence band') and the next available energy level (the 'conduction band'). A charged particle passing through the semiconductor will release electron-hole pairs which, under the influence of an applied electric field, will drift toward opposite electrodes. However, in pure silicon, there are many more free charge carriers than those produced by the through passage of a charged particle, and the electron-hole pairs quickly recombine. Hence one must find a way to deplete the material of charge carriers before one can usefully apply the technique to measure the position of a particle.

The solution lies in doping. In an 'n-type' semiconductor, donor ions of group V are added that introduce energy levels close to the lower end of the conduction band, thus creating a surplus of electrons in the material. In a 'p-type' semiconductor, acceptor ions of group III are added that introduce energy levels close to the top of the valence band, which absorb electrons from the valence band and create a surplus of holes. When these two types of doped semiconductor are brought together, a gradient of electron and hole densities results in the diffuse migration of majority carriers across the junction. The ionised donors now have positive charge, whilst the ionised acceptors acquire negative charge, and the interface region becomes depleted of carriers. There is a potential difference across this 'depletion region', which can be increased through application of an electric field (adding a 'reverse bias voltage'). This increases the width of the depletion region. Any electron-hole pairs produced by a charged particle passing through the region will drift along the field lines to opposite ends
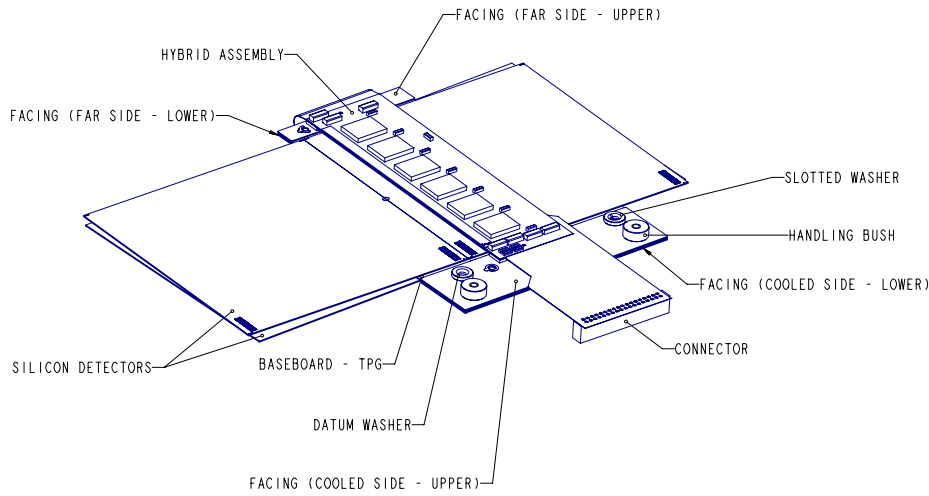
Figure 3.2: 3D view of an ATLAS barrel SCT module[41].

of the junction, and if the p-n junction is made at the surface of a silicon wafer, a prototype silicon strip detector is obtained. Charge is collected on the strips, and is amplified before readout.

Information on the strips used in the ATLAS SCT can be found in reference[39].

3.2.2   ATLAS Module Design

Module layout

Each SCT barrel module is made of four single-sided silicon wafers arranged as two back-to-back pairs, with each wafer containing 768 strips of 80μm pitch and active length 61.6 mm. One side of the module is aligned precisely along the beam direction, whilst the other is rotated by 40 mrad in order to provide a measurement in the z-direction. Technical information on the wafers is given in references[37] and[40], and a 3D view of the module layout is shown in Figure 3.2.

Readout of the sensor strips is performed by dedicated Application Specific Integrated Circuits (ASICs), known as `ABCD3TA` chips[42]. Each of these reads out 128 strips and is a single chip implementation of a binary readout (i.e. each strip registers either 'hit' or 'no hit'), manufactured using the radiation hard DMILL process[43,44]. There are 12 chips per

module, and these are mounted on a wrap around electronics hybrid that thermally decouples them from the wafer. Each side has six chips and a readout fibre, or 'opto-link'. Information is transported through the module by having each chip pass a token to the next chip. Only one chip, called the 'master', is connected to the link. If a single chip fails, it can be bypassed so that information from the remaining five chips can still be read out. If the failure occurs in the master chip (or the link) the other chips on the problematic side can be read out via the link on the other side of the module.

ABCD3TA ASIC design

A block diagram of the ABCD3TA ASIC is shown in Figure 3.3, showing the various functions performed during processing and readout. The front-end section of each chip is an analogue signal processing chain involving charge integration, pulse shaping and amplitude discrimination. The basic function of the discriminator is to register a hit if the amount of charge collected in the strip is above a pre-defined threshold value, and this threshold can be set using an internal programmable digital-to-analogue converter (DAC). At the start of each clock cycle, the chip samples from the discriminator output and stores the binary 'hit' or 'no hit' output in a pipeline until a level-1 trigger signal is received. If a trigger is received, the corresponding set of values is passed to the readout buffer along with their neighbours in time, thus providing information on the next and previous bunch crossing in addition to the current one (referred to as the three 'time bins'). This information is compressed before transmission off the chip through the link.

If one injects a known charge Q_i into a strip and then scans through different thresholds for that strip, one can extract the threshold, t_{50} at which 50% efficiency is achieved for pulses of magnitude Q_i. By varying Q_i and plotting the change in t_{50}, one can determine the threshold response of the module. The threshold at a charge $Q_i = 0$ C should be 0 V, and any difference is termed the 'threshold offset' of that channel. This offset varies from strip to strip and hence a 'TrimDAC' of 4-bit resolution is available for each channel that allows a threshold correction to be applied to each channel individually (this is in addition to the overall chip threshold DAC). The TrimDAC has four selectable ranges ('trim ranges') in order to cope with the extra spread of threshold offsets that occurs after radiation damage. To enable calibration of the detector, each channel has a calibration capacitor of 100 fF connected to its input to allow the injection of test charges in a range from 0 to 16 fC (set by an 8 bit DAC). These capacitors are themselves calibrated during production, and the information is stored in the SCT Production Database. There is also the capability to offset the timing of the discriminator with respect to
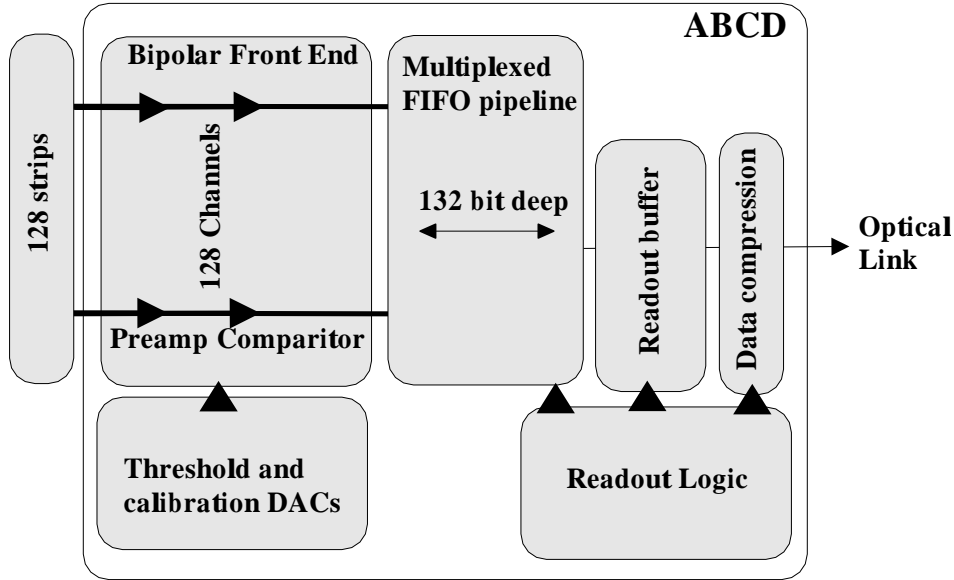
Figure 3.3: Block diagram of the ABCD3TA readout chip[45].

the clock signal over a 50 ns range, and the ability to mask channels that are not performing correctly.

3.3 Sources of Noise

It was observed in Section 3.2.1 that a silicon strip detector measures the presence of electron-hole pairs created by the passage of a charged particle through the strip. However, any practical device will be subject to stray capacitance effects, leakage currents, thermal noise and other effects, and these can cause a strip to register a 'noise' hit in the absence of any incident charged particle (with leakage currents being the main offender in the case of silicon). An important measure of detector performance is the average number of noise hits in a given strip per event, a quantity referred to as the 'noise occupancy'.

Given a charge probability distribution $P(Q)$ (either from noise or from legitimate sources), one can calculate the occupancy ρ for a given threshold t as follows:

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

The distribution $P(Q)$ is a Landau distribution for the charge released by a minimally ionis-

ing particle (MIP). Noise, meanwhile, has an approximately Gaussian distribution, and so, in practise, the charge distribution seen by the discriminator circuit of the ASIC in response to a MIP is actually a Landau distribution convolved with a Gaussian- the so-called 'improved Landau' distribution (see Figure 3.4). Typical numbers for a silicon detector are a most probable value of $\approx$ 3 fC, and a Landau width of 0.17 fC (before convolution with a Gaussian), whilst the noise distribution has a width of around 0.15 fC. The Landau distribution has a low energy tail which extends to even lower energies after noise broadening. Figure 3.4 demonstrates that one can reduce the noise occupancy by increasing the threshold of the module, and ideally one would position the threshold so as to lie to the right of the tail of the Gaussian noise distribution. However, the thick black curve on the same figure indicates that, by moving to the right, the efficiency of the module decreases. Thus, care must be taken to balance the competing demands of low noise occupancy and high efficiency.

The ATLAS inner detector is required to reconstruct a large number of tracks, whilst having access to only 6 silicon clusters per track and the TRT hits. For this reason, it is vital to reduce noise occupancy as much as possible, as sensible results will only be obtained by the pattern recognition software if the detector activity is dominated by legitimate hits. It is particularly important to remember that the SCT has only a binary output, so the effect of noise hits could be severe if too many modules are noisy.

The SCT specification requires that the detectors should be over 99% efficient for a noise occupancy of less than 5×10^{-4}, with most strips operating with a threshold setting of 1 fC. This noise occupancy condition arose from a combination of the experience of prototype silicon module performance gained from test beam measurements (see Figure 11-67 of reference [28]), and investigations of the track reconstruction efficiency obtained in the pattern recognition software as a function of noise (see page 77 of reference [25]). During the lifetime of the LHC, radiation damage of the modules will decrease the signal collection and increase the noise at the front-end of the ASICs [1], and the range of threshold settings for which one can still obtain the design performance will be narrower. For this reason, the initial settings are such that there is still room to manoeuvre when required later on.

---

[1]The main consequences of radiation damage to silicon devices are an increased leakage current due to the creation of defects in the strips and a change in the effective doping of the bulk silicon from n- to p-type followed by an increase of the full depletion voltage. The first of these effects increases the front end noise, in addition to which there may be radiation effects on the readout chip.
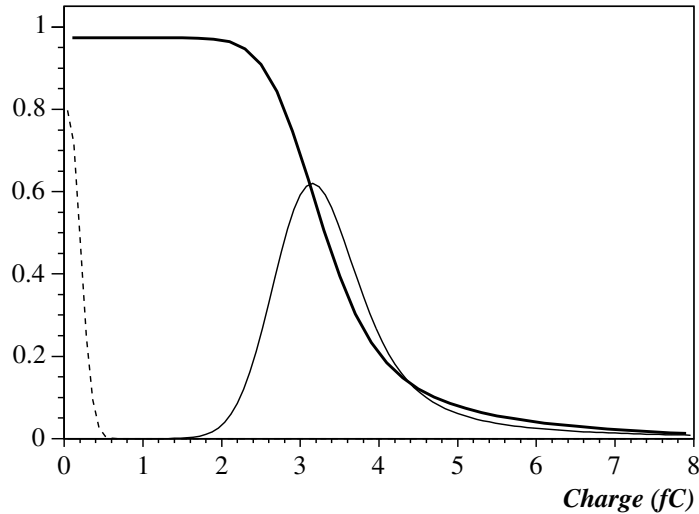
Figure 3.4: Schematic diagram showing an improved Landau distribution (pale black line), noise distribution (dashed line) and the efficiency vs threshold (thick black line) for a typical silicon detector[46].

3.4 Module calibration

3.4.1 Overview

In addition to possessing the various configurable parameters detailed in Section 3.2.2, each module has mechanisms for performing a series of calibration tests in order to tune the settings for optimum performance. Since Chapter 4 will deal with the development of software tools for the continuous monitoring of SCT module noise occupancy, it is instructive at this point to briefly review the calibration tests that exist on modules for the purpose of characterising noise occupancy. Results from these tests will later be compared with the results from cosmic ray commissioning tests.

The main goal of calibration tests is to ensure that the threshold is set at the right level for each strip, and that the thresholds are uniform over the detector. This involves finding the optimum common chip threshold DAC for each chip, and the channel correction TrimDAC settings for all strips within that chip.

Calibration of the detector is aided by the internal capacitor present for each strip, and various tests may be performed in order to determine strip behaviour. These were initially controlled by use of the SCTDAQ software package[47], now superseded by the SctRodDaq

package[48] which uses the optical readout system. This current package was designed to perform the same functions as SCTDAQ (these basic functions are described in[49]), and its development is described in considerable detail in references [46] and [50]. Conveniently brief reviews of the SctRodDaq functionality are provided in[51] and[52]. The aim of this section is to cover only those tests which are of relevance to the noise occupancy calculations presented in Chapter 4.

3.4.2 Calibration tests

The basic method of calibration is to perform a series of occupancy measurements while a chip parameter (usually the threshold) is varied in steps within a specific range. The relevant analogue tests that can be performed in this fashion are summarised below:

1. **Trim Range Test:** The chips used to build the final SCT modules were selected such that the threshold offsets (as defined in Section 3.2.2) of all channels can be brought into line by adjustment of the TrimDAC settings using trim range 0 or 1 (where 3 is the most coarse trim range). Trim ranges 2 and 3 will be required at a later date once radiation damage sets in. It is therefore necessary to determine the initial TrimDAC characteristic, for each strip, for all of the trim range settings. This is done using a trim range test.

   For trim range 0, a threshold scan is performed for each of the 16 possible TrimDAC settings (i.e. set the TrimDAC, take a new scan), with an injected charge of 1 fC. This yields a mean threshold value t_{50} defined in Section 3.2.2. The purpose of using all 16 possible TrimDAC settings is to check that each TrimDAC step can be selected for the particular channel of interest. A plot of mean threshold t_{50} vs TrimDAC setting is found to have a linear shape, and by fitting a straight line to this plot one obtains the formula for working out the TrimDAC setting that should be used for a required value of t_{50}.

   The test is then repeated for the other trim ranges, though with only 4 threshold scans per trim range, rather than 16, in order to save time. For each of these tests, a straight line fit once again gives the formula that tells us which TrimDAC setting to use for a required mean threshold t_{50}.

   The test now asks the question: "If I wanted to set my channel threshold to a particular t_{50} value, is there a TrimDAC setting that I can use to do it?". A range of possible t_{50} values to which one might want to set the threshold for a given channel is considered (referred to as 'trim targets'). For each hypothetical t_{50} value, the TrimDAC setting required to set that particular threshold on the channel is calculated using the straight

line equation. If this value lies within the permitted range of TrimDAC settings, then we would be able to obtain the t_{50} value if we so wished, and the channel is termed 'trimmable'. This test is done on all channels, and the total number of trimmable channels for each hypothetical threshold is stored. This is repeated for each of the four trim ranges. The aim for the detector is to have as much control over the channels as possible- we therefore want most of them to be 'trimmable'.

Finally, the TrimDAC and trim range settings at which each channel should be operated are selected, and this is done in such a way as to minimise the spread in the mean thresholds t_{50} that the channels are set to. Any defective channels are noted and stored. The goal is to choose settings so as maximise the number of trimmable channels whilst minimising the sum of the trim ranges of chips (thus meaning that most channels are operated with low trim ranges). There is a choice of three ways in which to do this optimisation. One can either choose to vary both the trim range and trim target from chip-to-chip, or have the trim target fixed for all chips in a module but the trim range variable or, finally, to have the trim range and the trim target fixed for all chips in a module. Modules with their TrimDAC settings set to the mean thresholds resulting from this optimisation test are said to be running with their 'trim target settings'.

2. **Three Point Gain Test:** A series of threshold scans is performed for three different injected charges (e.g. Q_i = 1.5, 2.0 and 2.5 fC). For each injected charge, a complementary error function is fitted to the data, the mean of which corresponds to the threshold t_{50} at which 50% efficiency is achieved for pulses of magnitude Q_i (defined originally in Section 3.2.2, and the RMS of which is a measure of the output noise (in mV). The gradient of a straight line fit to the graph of t_{50} vs Q_i gives the gain of each channel, and one can then obtain the input noise for each channel by dividing the output noise at the 2.0 fC data point by the gain. The threshold voltage corresponding to a charge of 0 fC obtained from this fit gives the threshold offset, which is corrected for using the 4-bit TrimDAC referred to in Section 3.2.2.

3. **Response curve test:** This is an extended version of the three point gain test, where threshold scans are run at ten different injected charge values (Q_i = 0.5, 0.75, 1.0, 1.25, 1.5, 2, 3, 4, 6 and 8 fC). The aim is to determine more accurately how the voltage measured at the comparator in the ABCD3TA relates to the charge that was injected from the calibration capacitor. Rather than fitting the graph of t_{50} vs Q_i with a straight line as before, the response curve is normally fitted with the following empirical function

which is found to give a good fit to the data:

$$V = a + \frac{b}{1 + e^{-Q/c}} \tag{3.2}$$

where a, b and c are fit parameters. Examples of fits are given in Figure 3.5. One can extract the gain ($\frac{\partial V}{\partial Q}$), offset (voltage at $Q = 0$) and noise as before. One can set the TrimDAC settings for channels so as to compensate for the offset encountered, and these are termed the 'response curve' settings.

4. **Noise occupancy test:** A threshold scan is performed *without* charge injection in order to determine the noise occupancy of each module as a function of threshold. This is often done after the response curve obtained in the previous test has been analysed, and the aim is to test more precisely the noise occupancy near the 1 fC nominal working point. This can be taken directly from the plot of threshold (t) vs noise occupancy (N), and is an important statistic. However, the total effect of noise depends not only on the noise voltage/current in the strip, but also on the amplification and pulse shaping that occur in the front-end processing chain. The cumulative effect of these factors is described by the 'equivalent noise charge' (ENC) which describes a noise pulse in terms of the charge pulse at the detector needed to create the same output. To evaluate this for a strip, one can make use of the following relationship between the noise occupancy and the threshold, which is a close approximation to the complementary error function at low occupancy:

$$N \propto \exp\left[-\frac{t^2}{2\sigma^2}\right] \tag{3.3}$$

where σ is the ENC of the strip[53]. A plot of $\ln(N)$ vs t^2 therefore gives a straight line whose gradient is determined by σ (for an example, see Figure 4.9).

3.5   Summary

A review of the module design and readout of the SCT barrel has been given, along with a summary of the calibration tests most relevant to noise occupancy. These tests are vital both to ensure that the detector is operating correctly, and to establish what action needs to be taken with regard to the changing of DAC settings, etc, should a channel be operating too far from its nominal settings.
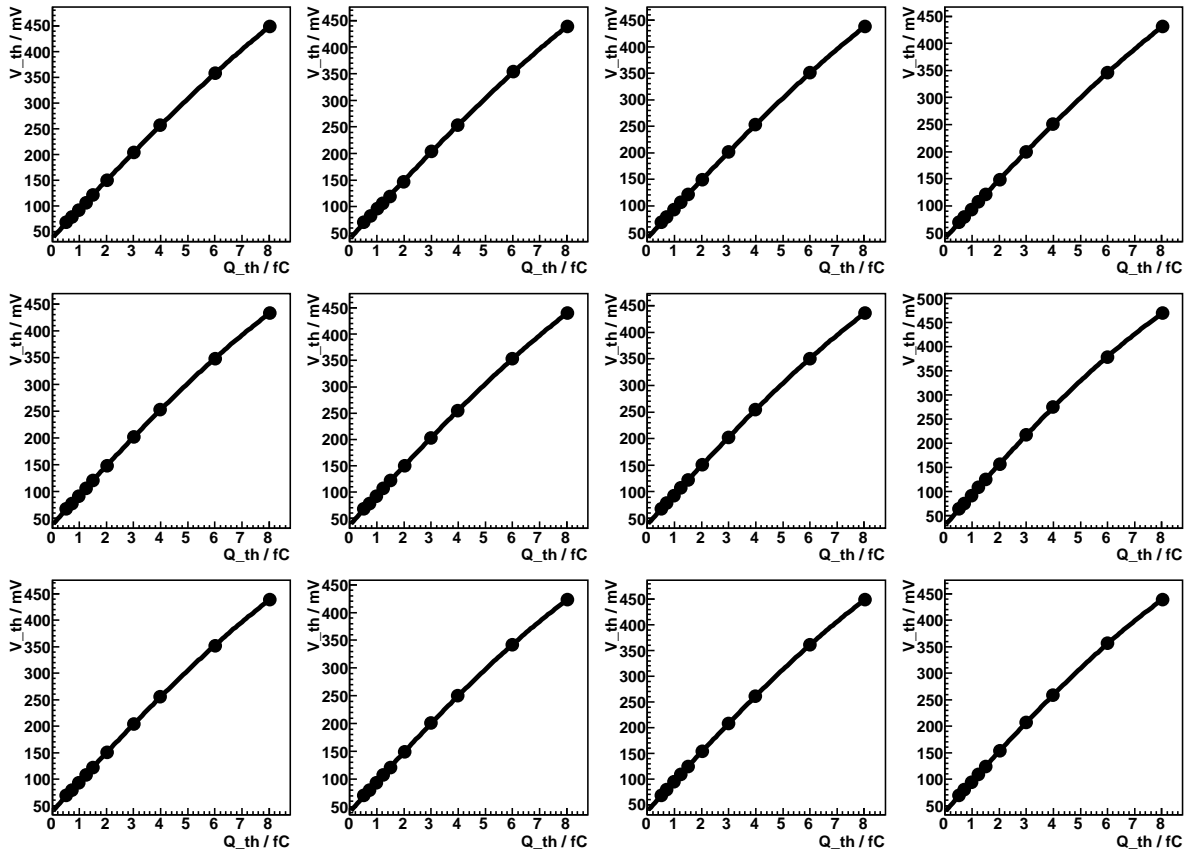
Figure 3.5: Examples of response curves, for the 12 chips of one module of the SCT, obtained in a calibration run taken during the SR1 cosmics commissioning period.

During its lifetime in the ATLAS detector, the SCT will be monitored constantly in order to track changes in the detector response due to radiation damage and/or component malfunctions. This will be performed using dedicated software that will use ATLAS event data to calculate module efficiencies and noise occupancies, which was recently commissioned using cosmic ray tests at CERN. The SCT monitoring software is the subject of the next chapter.

Chapter 4

Development of the SCT Monitoring Package

From next year, the ATLAS semiconductor tracker will find itself in the centre of an unprecedented maelstrom of radiation, facing a constant barrage of damaging particles. Although this is partly the point of having it there in the first place, it is impossible to avoid detector damage in such an environment, and it is vitally important to track changes in the detector response and adjust calibration settings to maintain optimum performance.

As described in Chapter 3, one effect of radiation damage is an increase in the noise occupancy of affected strips, and this can be addressed by changing the voltage threshold settings if the location of problem strips can be identified. Unfortunately, this is not the only type of problem with the detector, and there are other effects that will be relevant even in the earliest days of low luminosity running. For example, if the high voltage supply of a module trips off, one registers large numbers of spurious hits and hence can no longer trust the output of a module. In an environment where there are lots of tracks to be reconstructed from a small number of hits, pattern recognition is already tough even if noise is entirely absent. Excellent control must be exerted over noise hits in order to obtain good quality data for physics analysis.

It is clear, then, that the successful running of the SCT demands a system capable of continuously monitoring the detector in order to observe changes in strip behaviour. The system should run online during data taking and provide visual output for the tired eyes of shift workers, as well as communicating with the databases that store module conditions and configurations.

This chapter details the development of software tools for the online monitoring of the

SCT, written in C++ within the `Athena` software framework of the ATLAS collaboration (see Section 4.1). This package has existed for some time prior to the work in this thesis, and is the result of ongoing work within the ATLAS collaboration[54]. The detailed analysis in this chapter will focus on tools for measuring noise occupancy and DAQ errors.

During the development of the noise monitoring tool, large amounts of data were obtained at CERN by running a sub-section of the SCT barrel and observing the passage of cosmic rays through the detector (these are referred to as the 'SR1' commissioning tests, after the building in which they were performed). This data was used to debug the tool in comparison with other offline analyses, and the tool was subsequently used on later data sets to investigate several different noise effects.

Section 4.1 contains a brief explanation of the ATLAS software framework, thus defining some of the jargon used in later discussions. This is followed by an overview of the SCT Monitoring `Athena` package as a whole, and an introduction to the SR1 cosmic tests. Original work performed for the development of this package is contained in Section 4.4 in order to provide a clear distinction between the work of the collaboration and the work of the author. Finally, Section 4.5 presents noise occupancy results obtained during the SR1 cosmic testing at CERN in the summer of 2006.

4.1   Overview of the ATLAS Software Framework

4.1.1   The Athena Framework

ATLAS will produce a very large amount of data which needs to be processed at both the sub-detector and detector level in order to form events. These must then be analysed by dedicated algorithms. The sheer complexity of this process demands a large amount of varied software, and any software project of such a size requires a firm organisational principle in order to guarantee firstly that it functions at all, and secondly that it is able to keep pace with new developments.

To this end, the software for the ATLAS experiment is developed within a framework called `Athena` which is itself an implementation of the `Gaudi` framework developed by the LHCb experiment. This provides a skeleton C++ application in which users can develop their code, and supplies the tools for communication between different parts of the software and for any common functionality that is required. `Athena` provides the functionality for offline

analysis algorithms and packages and can also be run online at the event filter stage of the triggering process.

Each release of the ATLAS software is divided into 'packages', each conforming to a specified structure. A package can be built and then configured at run-time using a Python script called a 'job option' file. Once developed, code is stored in a central CVS repository, where it is available to the rest of the collaboration.

4.1.2 The Event Filter and Athena Processing Tasks

The Event Filter (EF) is the third and final stage in the ATLAS triggering process (see Figure 2.7), receiving at its input fully built ATLAS events with an input rate of the order of 1 kHz. The EF will exist as a farm of computers that each run 'event Processing Tasks' (PTs) in order to provide trigger decisions in addition to performing event analysis and reconstruction, and 'Event Filter DataFlow Managers' (EFDs,) that move events from the LVL2 trigger to the PTs and to mass storage. The combination of reduced data rate and full event information means that it is possible to use `Athena`-based algorithms at the event filter stage, and these are implemented as PTs. The SCT Monitoring package is one example of a process that may be run as an `Athena` PT.

4.2 The SCT Monitoring Athena Package

4.2.1 Overview

The SCT Monitoring package is written in the ATLAS offline `Athena` framework, but is able to run online as an `Athena` PT. The package originally existed as a combined test beam tool, but has recently been updated to provide monitoring facilities during the 2006 cosmic commissioning tests (see Section 4.3.) The package remains in development, and is currently being extended further.

The package was designed to be able to process either simulated data or data from the bytestream that is generated from the DAQ in real time. This data passes through several stages of processing in the `Athena` software chain. Simulated data is 'digitised' in a tool that recreates the action of the SCT readout by comparing the simulated signals to a given threshold for each strip to give a binary output. The results are SCT 'hits' which are stored as

Raw Data Objects (RDOs). On the other hand, bytestream data do not need to be digitised, and are merely converted to RDOs by a conversion algorithm. After digitisation/conversion, the data passes through a clusterisation algorithm which groups the signals in neighbouring strips into 'clusters' of hits. These are characterised by a mean position and width.

The SCT monitoring tool must have access to reconstructed tracks online, and must also be able to read and write to the database that stores module conditions (implemented as a COOL database[55].).

4.2.2   Online running

When running as an `Athena` PT, the monitoring package periodically sends updated histograms to an Online Histogram server (OHS) via an interface to the High Level Trigger software known as the `MonAthenaEF` service. Histograms are then collected from the OHS by an interactive GUI called the Online Histogram Presenter (OHP) with a frequency set by the user. Note that the update-to-server frequency is quite separate from the retrieve-from-server frequency; the former is a property of the `Athena` algorithm, and the latter a property of the display GUI. A configuration file sets which histograms on the server are subscribed to, and which of these subscribed histograms are displayed in the OHP. Subscribing to more histograms slows the process due to slow retrieval, and any number beyond about 3000 is unreasonable. This sets limits on the design of the package, and a final design strategy is still under discussion.

4.2.3   Package contents

The SCT Monitoring package contains a series of tools with a common interface that are driven by a top level algorithm. Each tool has an *initialize*, *execute* and *finalize* method, and these are called sequentially by the top level algorithm for each of the tools in turn. All histograms are booked in the *initialize* method of the top algorithm. The *execute* method of the top algorithm performs analysis and fills histograms, and periodically updates the OHS. The *finalize* method can be used to provide any information that is required on output.

The tools contained in the package are summarised below:

SCTEffMonTool

The efficiency monitoring tool makes histograms of module efficiencies and track-fitting efficiencies. For each reconstructed track in an event that remains after a quality cut, the algorithm takes the perigee parameters [1] and extrapolates the track, ultimately returning a vector of places where the track should cross the various layers of the SCT, in the form of Track State On Surface (TSOS) objects. For cosmic events, this extrapolation should simply be that of a straight line through the detector. It then compares the TSOS positions to the positions of SCT clusters in the event, and to the SCT clusters that lie on tracks. If the TSOS is in the same position as a cluster lying on a track then both the module and the track reconstruction are efficient. If not, then the tracking is noted as inefficient. However, this does not exclude the possibility that a cluster was observed in the position expected; there could be a cluster object at that location that, for whatever reason, was not associated to a track. Thus, the algorithm loops through all clusters to determine if a cluster exists in the location predicted in the TSOS, and the module is deemed efficient if such a cluster is found, and inefficient if not.

Note that this algorithm is biased since for each layer it uses hits in that layer to determine the efficiency. A more rigorous method that removes hits in each layer when investigating that layer exists in the offline code, but is too slow to be run online.

The output of the efficiency tool can be displayed either as 2D $\eta - \phi$ maps of the module/track efficiency per module (for each layer and side of the barrel) or as 1D histograms showing the average efficiency in each layer.

SCTTrackMonTool

The track monitoring tool makes histograms of tracking quality information. The `Athena` track fitting algorithms fit tracks to the clusters in the SCT that pattern recognition software has identified as tracking clusters. There are then a series of quantities to which one may refer in order to judge the quality of the fitted track. The 'residual', r, of a cluster is defined as follows:

$$r = p_{\text{track}} - p_{\text{clus}} \tag{4.1}$$

---

[1] The perigee of a track is the point of closest approach to the nominal interaction point (and thus the origin of the ATLAS coordinate system). The perigee is specified by the following five parameters: the distance of closest approach from the track to the z-axis d_0, the z value at the point of closest approach z_0, the azimuthal angle of the particle momentum at the point of closest approach ϕ_0, the polar angle of the momentum θ and the charge over momentum magnitude q/p.

where p_{track} is the predicted position of a cluster on the track, and p_{clus} is the observed cluster position. The residuals are calculated using the local track parameters in the plane representation[2] and are calculated using the local X variable.

The 'pulls' are the residuals divided by the quadrature sum of the track prediction position error and the actual hit position error:

$$p = \frac{p_{\text{track}} - p_{\text{clus}}}{\sqrt{\sigma^2_{\text{track}} + \sigma^2_{\text{clus}}}} \tag{4.2}$$

The pulls and residuals are both retrieved and plotted by the `SCTTrackMonTool`, and one can view both 2D $\eta - \phi$ maps and 1D summary histograms for each layer, in addition to 2D $\eta - \phi$ maps of the RMS values of the residual and pull distributions. The tool also provides a histogram of the track fit χ^2 (divided by the number of degrees of freedom) and the number of SCT hits per track.

SCTHitsMonTool

The hit monitoring tool makes histograms of hits in the SCT, as well as providing cluster information. One can view 2D $\eta - \phi$ hit maps for each layer and side of the SCT, as well as 1D summary histograms. The tool generates plots of the cluster size for each layer and side, and also provides plots of the number of hits vs the trigger type. For more detailed checks, one can view hit information at the strip level for each module.

In order to look for anomalous hit patterns, the `SCTHitsMonTool` generates correlation plots showing the η of hits in one layer vs the adjacent layer, and for one side of a module vs the other side of a module. The same is done for ϕ, and one can also view self-correlation η and ϕ plots for each layer and side.

SCTErrMonTool

The data coming from each optical link in the SCT is received by dedicated off-detector read out boards called 'read out drivers' (RODs) that convert the information from the SCT to

---

[2]The track parameters may be given in several representations, with the local 'at a plane' parameters being ($locX$, $locY$, ϕ, θ, q/p), where $locX$ and $locY$ are position parameters in an arbitrarily oriented plane. The last three variables refer to the momentum of the particle producing the track; ϕ is the azimuthal angle, θ is the polar angle and q/p is the charge to momentum ratio.

a common ATLAS event format. This occurs in several stages, the first of which uses a 'formatter' to read and decode the information from the SCT. At this stage, there may be errors in the bytestream, and these are flagged. The information is then passed to an 'event fragment builder', that, amongst other things, compares the bunch crossing ID and LVL1 trigger ID of the module information to the central values governing the detector as a whole. Any discrepancies are flagged.

There are thus a variety of errors that may arise, and the `SCTErrMonTool` plots these in 2D $\eta - \phi$ maps for each layer and side, and as 1D summary histograms. Possible errors include:

1. **BCID:** The bunch crossing ID in the data header does not match the detector central value.

2. **LVL1ID:** The level-1 trigger ID in the data header does not match the central value.

3. **First Hit:** The first hit in a cluster has an error. There are a variety of data taking modes for the SCT, due to the fact the information is available for the previous-current-next bunch crossing. Thus, one may specify a readout mode of the form ZZZ, where the three Zs of the 'hit pattern' denote the three successive bunch crossings, and are referred to as 'time bins'. Possible selections for each time bin are 0 (no hit in that crossing), 1 (a hit in that crossing) and X (either a hit or no hit in that crossing). If a hit is detected but its hit pattern does not match that specified by the data taking mode, a hit error will be registered. Persistent hit errors indicate miscalibration of the SCT timing with respect to the bunch crossing.

4. **Second Hit:** The second hit in a cluster has an error.

5. **Preamble:** The event data did not begin with the standard preamble '11101'.

6. **Formatter:** Formatter errors are those detected in the data stream by the formatter, before the event information is sent to the event fragment builder. Examples include errors in the header or trailer of the stream, and errors in the link readout.

SCTNoiseMonTool

The SCT noise monitoring tool generates plots of the noise occupancy of modules in the SCT, as 2D $\eta - \phi$ maps. The development of the tool is covered in detail in Section 4.4.

4.3   SR1 cosmic commissioning tests

The main aim of the current round of improvements to the SCT Monitoring package was to provide an effective monitoring system for use during the SR1 cosmics commissioning period at CERN. These tests provided data for both the debugging of the tools and the final results obtained from them.

4.3.1   Overview

The SCT was inserted into the TRT at CERN on 17 February 2006[56], after which work commenced on a planned series of tests in the SR1 building at the CERN surface. There were several aims, chief of which was simply to gain experience of running the detector in a setup similar to that which will be used in the cavern from 2007. Getting the detectors operational tested not only the components themselves but the services, offline software chain, and the readout and trigger systems. Specific detector performance goals included testing the combined system for cross talk and noise resulting from synchronous operation of the SCT and TRT together, as well as checking issues related to grounding, heating and trigger rate. Furthermore, the passage of cosmic rays through the detector provides an opportunity to test the alignment of components. A picture from the ATLAS event display `Atlantis` is given in Figure 4.1, showing a track passing through the top and bottom sectors of the SCT.

4.3.2   SR1 Layout

The SR1 cosmic tests involved the use of only a subset of the total detector modules. The final setup is shown in Figure 4.2, and involved approximately one quarter of the SCT barrel (468 of the 2112 modules) and one eighth of the TRT barrel. Readout of the SCT was performed using 12 RODs, and the detector was kept in a dry thermal enclosure throughout. Triggering for cosmic runs was performed using the three scintillators shown in the figure. The scope of the work was such as to include both standalone runs for the SCT and TRT, and combined tests of both systems, and a concise summary is provided in reference[57].
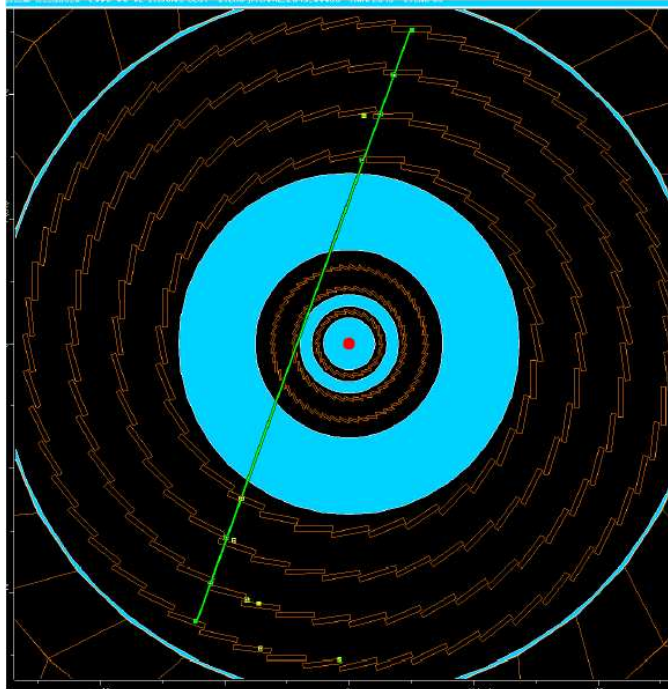
Figure 4.1: View from the ATLAS event display program `Atlantis`, showing a cosmic muon passing through the top and bottom sections of the SCT that were connected during the SR1 cosmic commissioning tests in the summer of 2006.
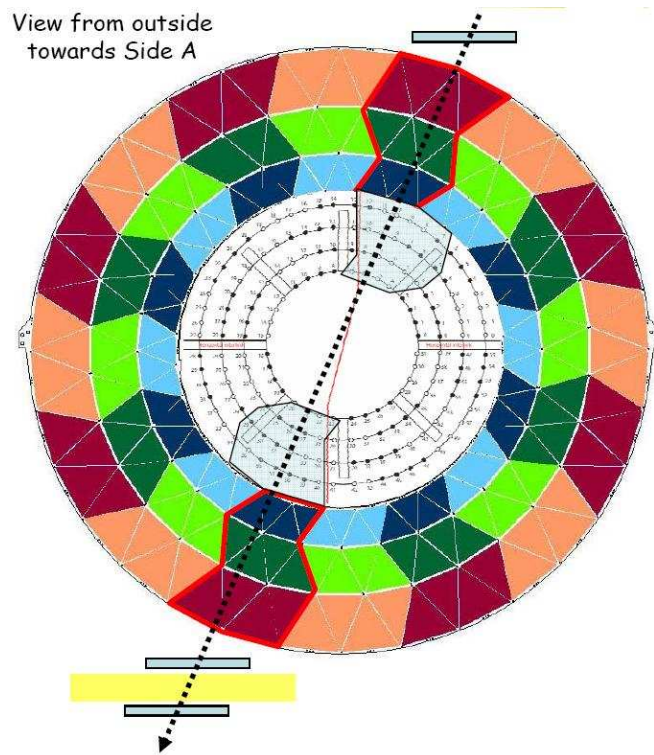
Figure 4.2: The SR1 cosmic setup, showing the two regions (top and bottom) of the barrel that were commissioned by the end of the test period[57].

4.3.3   Modes of operation

The tests performed at SR1 include both the online calibration tests referred to in Chapter 3 and so-called 'physics' runs in which all modules in the barrel are triggered simultaneously by either a pulse at fixed frequency or the passage of a muon through the scintillators.

Of the physics runs, those triggered by the scintillators would be termed 'cosmic' runs, and those triggered by a pulse would be referred to as 'noise' runs. A cosmic run will feature a combination of tracking and noise hits, whilst the probability of catching a muon passing through the detector during a noise run is sufficiently small that all hits seen in such a run can be labelled as noise hits. Most of the analysis presented in this chapter involves the use of noise runs rather than cosmic runs.

It is important to note that the 'physics' noise runs were originally taken with each channel set to what was assumed to be a 1 fC threshold, using data from a previous SCTDAQ configuration file. This information turned out to be incorrect however, and subsequent runs were taken using 'trim target' settings until a new response curve could be taken with a full 10 point gain test (a new trim range test was not performed). The initial incorrect data has not been used for analysis in this thesis.

4.4   Development of Noise Occupancy Tool

4.4.1   Overview and design

The noise occupancy monitoring tool functions in a similar way to the hits occupancy tool, with the exception that it aims to produce maps from the noise hits per module rather than the total number of hits. There are several ways to define a noise hit, and the choice made here is to select hits that are not associated to any tracks in the event.

For each event, the tool firstly retrieves the collection of hits that were associated to tracks in the event. These are known as the Reconstruction Input Objects (RIOs) on the track. After running on all of the events, a 1D histogram is obtained for each module, containing the number of RIOs in each strip of the module, and a 2D $\eta - \phi$ map with a bin for each module is filled for each layer and side of the SCT. The tool then fills analogous histograms for all 'X1X' hits in the event, where the three characters denote the previous-present-next bunch crossing, as defined earlier. All noise occupancy histograms are filled by taking the difference

between these histograms and dividing by the number of events. 2D 'module map' histograms are further divided by the number of strips in one side of a module in order to give an average noise occupancy per strip for each module.

4.4.2   Debugging

The noise occupancy tool has been extensively debugged by comparing results to various offline analyses performed by other members of the ATLAS collaboration, and by comparing with the results of the online calibration tests described in Chapter 3.

Offline analyses

During the SR1 cosmic testing phase at CERN, many data sets were obtained under different conditions in order to both gain experience of running the detector and to review the detector performance. The large amounts of data generated were analysed by several members of the ATLAS collaboration, thus enabling the author to make some simple cross checks of the results of the noise occupancy tool. Of course, this does not exclude the possibility that all of the noise analyses were wrong, and hence a comparison was also made to the data obtained in the online calibration tests, as described in Section 4.4.2.

Several of the SR1 cosmic runs were analysed by Reisaburo Tanaka[58], and the results were found to be consistent with those generated by the `SCTNoiseMonTool`. In addition, Ola Oye analysed several of the data runs, and his noise occupancies were found to be in good agreement with those generated by the noise monitoring tool.

Calibration tests

A more rigorous test of the `SCTNoiseMonTool` was obtained by comparing the output to the results of the online noise occupancy calibration tests described in Section 3.4. Many of these calibration tests were performed during the SR1 cosmic commissioning period, and the results were stored in a database that is accessible via an interactive web display system.

In order to compare the `SctNoiseMonTool` results with the online results, one needs to use a 'physics' noise run where the channel thresholds match those for which the online calibration tests apply. There is thus a question as to whether to compare the online results with a trim target run or an RC run.

In order to resolve the question, it is necessary to consider the fine details of the calibration noise occupancy scan. This sets up the threshold scan in units of mV relative to the trim target. This voltage is then converted to a threshold in fC and this is done through the response curve. At the end of the test, the threshold scan point nearest the 1.0 fC response curve point is picked in order to obtain the noise occupancy at 1.0 fC. However, if the trim target settings are slightly off (which they may well be in this case, since the trim target test was not rerun after a new response curve was taken), that point might not be the closest point to 1 fC, as the threshold scan points are generated using trim target information. Thus, one obtains the noise occupancy at a threshold that is not quite 1 fC, and this should not be compared to a 'physics' run taken after the new response curve was obtained (where the thresholds *are* 1 fC). In other words, the online settings do not correspond exactly to a 1 fC threshold due to an error that occurred during the data collection, and the physics runs whose thresholds match the calibration thresholds most closely are the trim target runs. Note that this only applies to runs taken during the SR1 testing period- comparison between SR1 physics data and pre-SR1 online data sets is valid.

For this reason, a comparison was made between noise run number 2900 and the calibration run number 2981. 2900 was taken on 12th June 2006, with the thresholds set to trim target and the SCT in any hit mode, meaning that a hit in any of the three time bins will be recorded. The noise software then selects all X1X hits. Calibration run 2981 was taken on the 16th June, and involved all 4 barrels though only a subset of the modules were read out. This was the closest relevant noise run (in time) whose data was available on the DAQ web display. The `SctRodDaq` software which implements the calibration test also selects X1X hits.

Figures 4.3 to 4.6 show the modulus of the difference between the physics noise occupancy and the calibration noise occupancy, divided by the RMS of the calibration noise occupancy. A plot is given for each barrel, and the occupancies shown are module averages. Thus, one is essentially viewing deviations from the calibration data in units of the calibration RMS.

Excellent agreement is obtained for most modules, though there remain three modules with a deviation greater than 3σ[3]. This may simply be due to statistical fluctuations, though it is thought more likely that it results from a difference in the way that the calibration and

---

[3]It is noted that the numbers in Figures 4.3 to 4.6 are not distributed as one would expect for unit standard deviation, with most of the deviations being between -0.5 and 0.5 (approximately). This is of little consequence here, given that the main aim of the test is to identify outlying modules. A suggested explanation for the fact that the variance is not unity is that the numbers are affected by the modulation of the noise occupancy across each chip (visible in Figure 4.7). This modulation is clearly larger than the intrinsic variation in noise occupancies (which is the thickness of the line in Figure 4.7), and thus the errors on the individual measurements are overestimated by using the overall RMS.
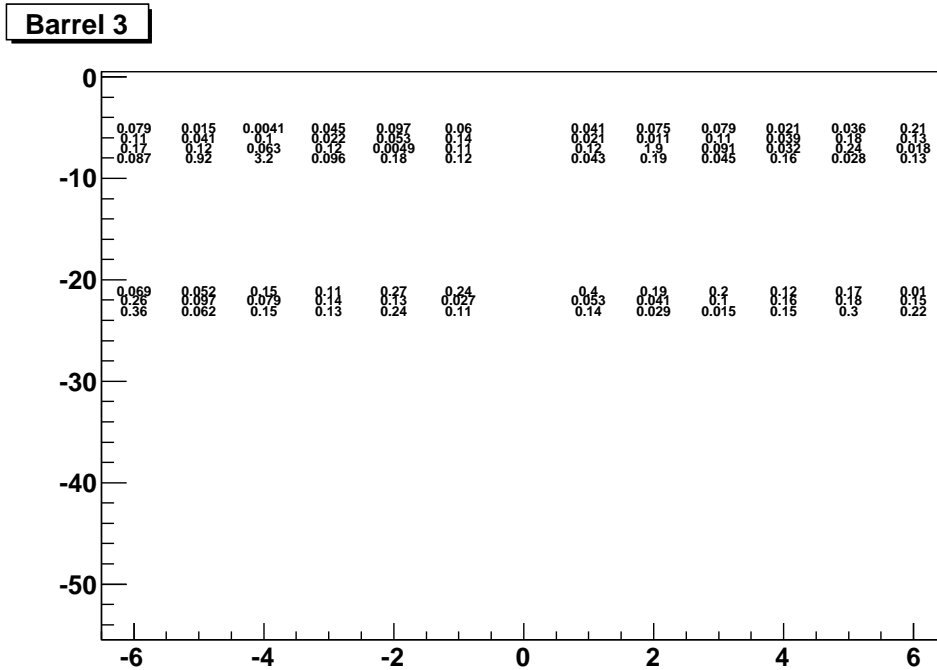
Barrel 3

```
      0 ┌
         │   0.079   0.015   0.0041   0.045   0.097   0.06            0.041   0.075   0.079   0.021   0.036   0.21
         │   0.11    0.041   0.1      0.022   0.053   0.14            0.021   0.011   0.11    0.039   0.18    0.13
         │   0.17    0.12    0.063    0.12    0.0049  0.11            0.12    1.9     0.091   0.032   0.24    0.018
    -10 ─│   0.087   0.92    3.2      0.096   0.18    0.12            0.043   0.19    0.045   0.16    0.028   0.13
         │
    -20 ─│   0.069   0.052   0.15     0.11    0.27    0.24            0.4     0.19    0.2     0.12    0.17    0.01
         │   0.26    0.097   0.079    0.14    0.13    0.027           0.053   0.041   0.1     0.16    0.18    0.15
         │   0.36    0.062   0.15     0.13    0.24    0.11            0.14    0.029   0.015   0.15    0.3     0.22
    -30 ─│
         │
    -40 ─│
         │
    -50 ─│
         └──┼──────┼──────┼──────┼──────┼──────┼──────┼
           -6     -4     -2      0      2      4      6
```

Figure 4.3: Histogram showing the absolute difference between the noise occupancy from a physics run (2900) and the noise occupancy from a calibration run (2981), divided by the RMS of the calibration run noise occupancy for barrel 3 (layer 0) of the SCT. The x-axis shows η, whilst the y-axis shows $-\phi$. The module at $\eta = -4$ and $\phi = 8$ has a 3.2σ discrepancy.

physics tools treat the data. In the calibration noise occupancy test, any strips that were not successfully fitted with an S-curve are ignored for the purposes of the noise occupancy calculation, thus essentially vetoing very noisy strips. There is no way to replicate this procedure in the physics code, as there is no access to threshold scan information for each strip (the data is obtained at one threshold only). Hence, a simpler method of vetoing very noisy strips was implemented- any strip with an occupancy of greater than 500×10^{-5} was ignored when filling the noise histograms. This reduced the number of discrepancies to the present level, and it is assumed that the remaining difference in behaviour between the calibration and physics results is due to an imperfect correlation between ignored strips. This is difficult to check given that there is no way to recreate the online masking exactly.

Given the close agreement of the vast majority of modules, however, it is argued that this comparison provides comprehensive evidence that the `SCTNoiseMonTool` is performing as expected, and is reporting accurate module noise occupancies for the SCT barrel region.
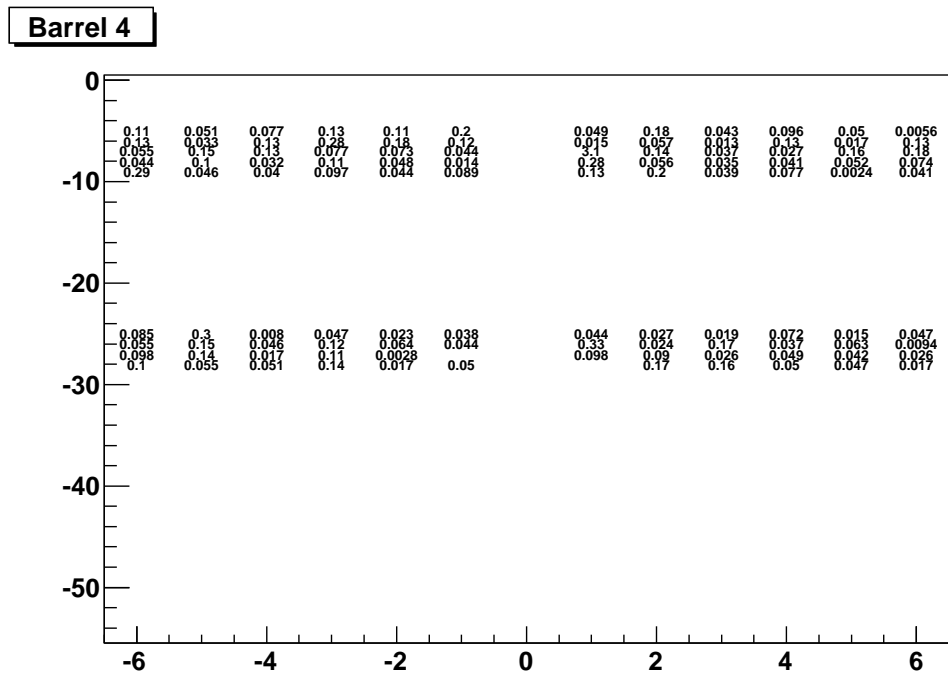
Barrel 4

0.11	0.051	0.077	0.13	0.11	0.2		0.049	0.18	0.043	0.096	0.05	0.0056
0.13	0.033	0.13	0.28	0.18	0.12		0.015	0.057	0.013	0.13	0.017	0.13
0.055	0.15	0.13	0.077	0.073	0.044		3.1	0.14	0.037	0.027	0.16	0.18
0.044	0.1	0.032	0.11	0.048	0.014		0.28	0.056	0.035	0.041	0.052	0.074
0.29	0.046	0.04	0.097	0.044	0.089		0.13	0.2	0.039	0.077	0.0024	0.041

0.085	0.3	0.008	0.047	0.023	0.038		0.044	0.027	0.019	0.072	0.015	0.047
0.055	0.15	0.046	0.12	0.064	0.044		0.33	0.024	0.17	0.037	0.063	0.0094
0.098	0.14	0.017	0.11	0.0028			0.098	0.09	0.026	0.049	0.042	0.026
0.1	0.055	0.051	0.14	0.017	0.05			0.17	0.16	0.05	0.047	0.017

Figure 4.4: As Figure 4.3 but for barrel 4 (layer 1). The module at $\eta = 1$ and $\phi = 7$ has a 3.1σ discrepancy.
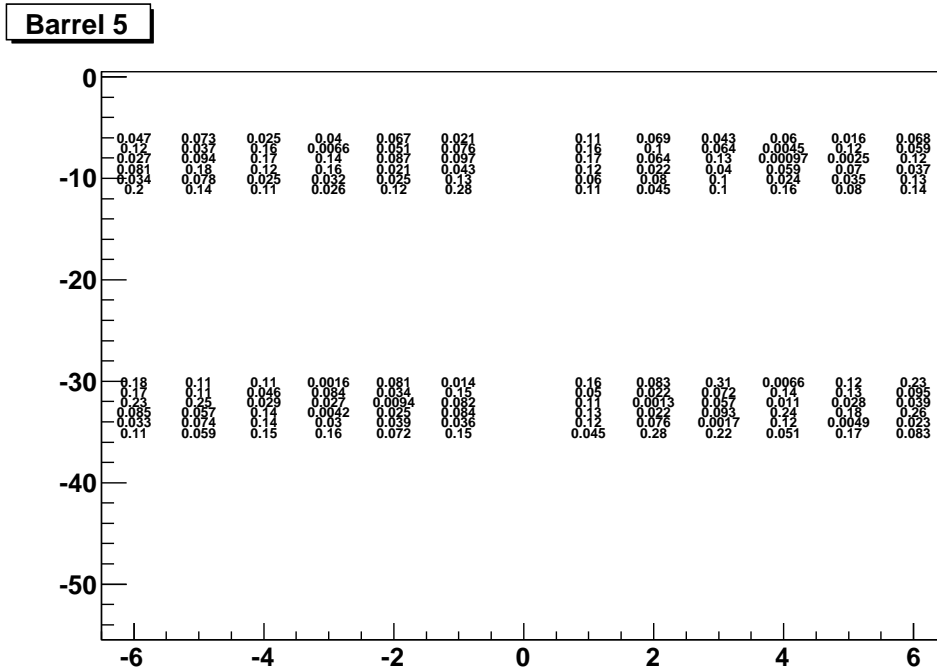
Figure 4.5: As Figure 4.3 but for barrel 5 (layer 2).

4.5   SR1 cosmic commissioning results

4.5.1   Overview

Having been thoroughly tested, the `SCTNoiseMonTool` was then applied to various SR1 data sets. A quick summary of the noise behaviour of the barrel SCT can be seen in Figure 4.7, which shows the strip number for all hits in the first 10,000 events of data set 2900. This was a noise run taken in 'physics' mode, with both the SCT and TRT powered and read out. Both sub-detectors were triggered using a pulse of frequency 50 Hz and the thresholds of the SCT modules were set to the trim target settings.

Figure 4.7 shows several features of the noise occupancy. Firstly, although most strips have noise occupancies in a narrow range, there are a few strips that are vastly noisier. These have been known about for some time, as shown by their presence in the calibration data, seen in Section 4.4.2. Secondly, one can clearly identify the six chips of a wafer as the strip number increases. Finally, there is a modulation in noise occupancy within one chip, with the noise being higher toward the middle of the chip. This has been observed since the module
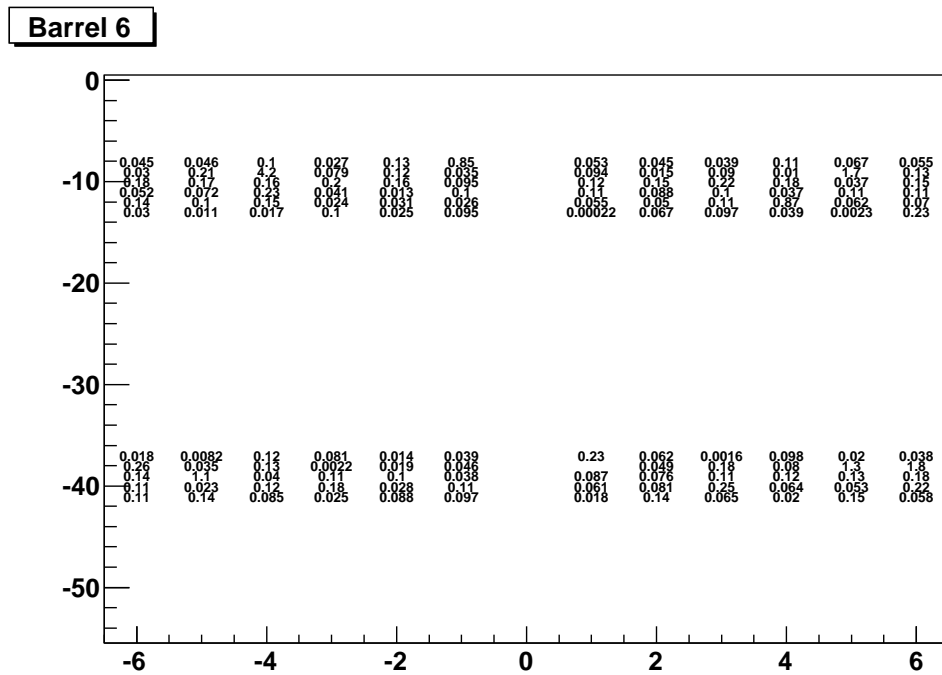
Figure 4.6: As Figure 4.3 but for barrel 6 (layer 3).The module at $\eta = -4$ and $\phi = 9$ has a 4.2σ discrepancy.
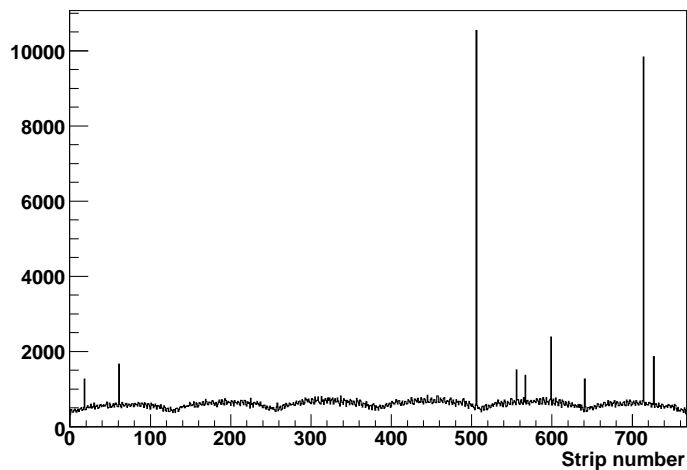
Figure 4.7: The strip number distribution for all hits in the first 10,000 events of data set 2900. The offline strip identifier goes from 0 to 767, and represents all strips in one wafer (one side of a module).

prototype production era. Although the precise cause is unknown it is thought to be due to the impedance of the power lines or ground lines not being sufficiently low. The chip edge has wire bondings to the hybrid lines, but the chip centre can potentially be more greatly affected from other channels switching on and off leading to a higher noise occupancy than that observed at the edge.

Several tests were performed during the last weeks at SR1 in order to investigate specific noise effects. These are summarised below, and the results of the noise tool are given for each.

4.5.2 Trigger rate scan

Run numbers 3064 to 3067 were performed under similar conditions, but with a different trigger rate, in order to see whether the noise occupancy changed as a function of trigger rate. The runs were all done with the new response curve settings, with modules set to a threshold of 1 fC. The TRT was clocked but not triggered. A summary of the run details is given in Table 4.1.

The `SCTNoiseMonTool` was used to obtain noise occupancy values for the data sets 3065 to 3067, with 3064 unfortunately lost due to a unfixable bug that occurred when uploading

Run number	Number of events	Conditions
3064	20K	50 Hz trigger rate
3065	42K	500 Hz trigger rate
3066	90K	5 kHz trigger rate
3067	70K	50 kHz trigger rate

Table 4.1: Details of the noise runs taken at different trigger rates, for the purposes of checking the variation in noise occupancy with trigger rate. All runs were taken with a 1 fC threshold obtained using the response curve settings.

		Run number	
	3065	3066	3067
Layer 0 Side 0	4.87(3.13)	4.94(3.14)	4.92(3.17)
Layer 0 Side 1	4.94(3.28)	5.01(3.28)	4.97(3.23)
Layer 1 Side 0	4.36(1.83)	4.33(1.81)	4.33(1.80)
Layer 1 Side 1	4.92(2.39)	4.89(2.37)	4.84(2.33)
Layer 2 Side 0	4.26(2.16)	4.26(2.12)	4.22(2.15)
Layer 2 Side 1	4.89(2.51)	4.90(2.49)	4.88(2.44)
Layer 3 Side 0	4.90(3.34)	4.89(3.31)	4.86(3.30)
Layer 3 Side 1	4.75(2.93)	4.78(2.96)	4.78(2.98)

Table 4.2: The mean (RMS) noise occupancy (scaled by 10^5) for the trigger rate scans detailed in Table 4.1.

conditions data to the database for that data set[4]. The analysis was performed using the first 10,000 events of each run, and without the noisy strip veto introduced in Section 4.4.2. Table 4.2 contains a summary of the results for each layer and side. It is concluded that the trigger rate does not affect the noise occupancy.

4.5.3 Noise threshold scan

Noise run numbers 3069 to 3074 were taken at different thresholds (using the new response curve settings), in order to check the response of modules as a function of threshold. The

---

[4]Technically speaking, the data set is fine, but it can no longer be read.
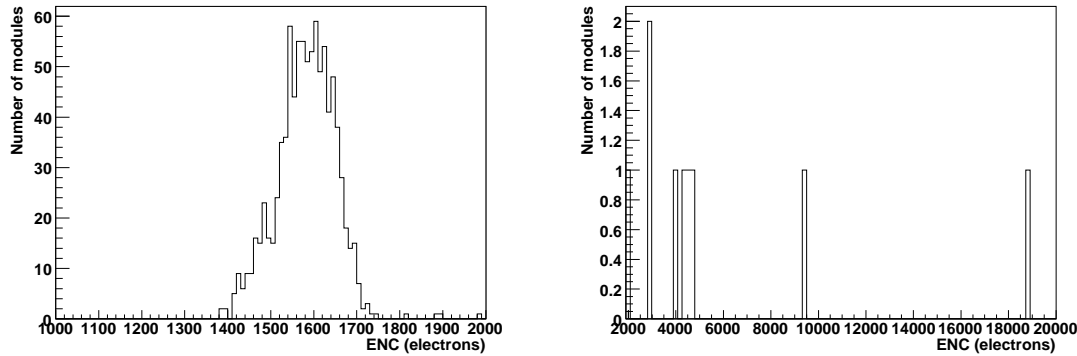
Figure 4.8: Histograms showing the distribution of the ENC values (in number of electrons) obtained for all active wafers in the SR1 data runs 3069-3074. Most wafers fall below the acceptable limit of 1900 electrons, whilst 9 are above the required value, as shown on the right.

triggers came from a pulse at a frequency of approximately 500 Hz, and the TRT was clocked but not read out. Details are given in Table 4.3, and the analysis once more used only the first 10,000 events of each run. No veto was applied to noisy strips.

As stated in Chapter 3, a plot of $\ln(N)$ vs t^2 should be a straight line for each module, and the gradient of each line allows the equivalent noise charge for each module to be extracted. The distribution of ENC values obtained for all wafers (i.e. one side of a module) is given in Figure 4.8, where it is observed that all active modules in the setup give reasonable ENC values, and hence obey Equation 3.3. An example of one of the straight line fits is given in Figure 4.9. The ENC values obtained are roughly equivalent to those reported from calibration runs, though an exact comparison is difficult to make due to differences in the way the data is analysed.

The main question to answer regarding the ENC values is whether they are within acceptable limits. The noise occupancy specification limit of 5×10^{-4} demands that the total equivalent noise charge should not exceed 1900 e$^-$[59]. All of the ENC values presented here are comfortably within this limit with the exception of those obtained for 9 wafers, which are distributed across all of the barrels (with two of the wafers being in one module). On further investigation with the strip level plots from the `SCTNoiseMonTool`, each of these turned out to

Run number	Number of events	SCT threshold
3069	25K	0.9 fC RC
3070	50K	1.1 fC RC
3071	100K	1.2 fC RC
3072	232K	1.0 fC RC
3073	32K	1.05 fC RC
3074	49K	0.95 fC RC

Table 4.3: Details of the noise runs taken at different threshold values.

be due to one or more defective strips. A cross check against known defects in the calibration data revealed that these are all known defects. When the analysis is run with a veto on strips that have a noise occupancy greater than 500×10^{-5}, there are no ENC values greater than 1900 e$^-$.

1. Layer 1, Side 1, $\eta = -4$, $\phi = 9$, wafer has one defective strip

2. Layer 2, Side 0, $\eta = -2$, $\phi = 6$, wafer has one defective strip

3. Layer 2, Side 1, $\eta = 5$, $\phi = 30$, wafer has one defective strip

4. Layer 2, Side 1, $\eta = -3$, $\phi = 33$, wafer has two defective strips

5. Layer 3, Side 0, $\eta = -3$, $\phi = 11$, wafer has one noisy strip

6. Layer 3, Side 0, $\eta = -3$, $\phi = 13$, module has fifteen noisy strips in total

7. Layer 3, Side 0 (and Side 1), $\eta = -4$, $\phi = 41$, module has fourteen defective strips in total

8. Layer 3, Side 1, $\eta = -5$, $\phi = 38$, module has two defective strips

4.5.4   Grounding issues

Run numbers 3084 and 3085 were performed with different grounding schemes. Run 3084 was run under normal conditions, whilst 3085 had cards installed in the power supply crates to short the cable shields to the power supply chassis. These runs were taken with thresholds
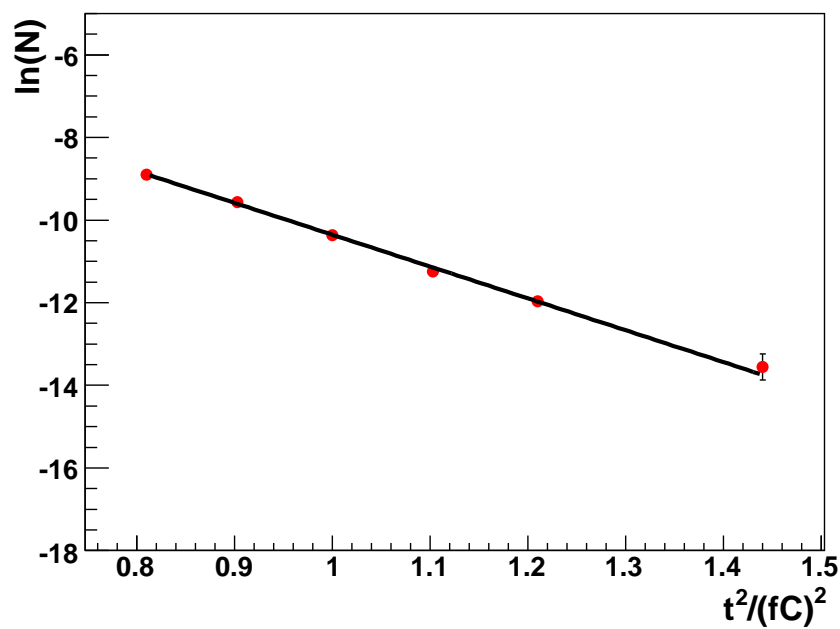
Figure 4.9: A plot of $\ln N$ vs t^2 for one wafer of the barrel SCT. A straight line has been fitted (see Equation 3.3), with a gradient $p_0 = -2.6 \pm 0.2$ $(fC)^{-2}$ and intercept on the x-axis $p_1 = -7.7 \pm 0.2$ $(fC)^2$.

	3084	3085
Layer 0 Side 0	3.52(2.30)	3.51(2.29)
Layer 0 Side 1	3.62(2.38)	3.60(2.42)
Layer 1 Side 0	3.07(1.28)	3.11(1.29)
Layer 1 Side 1	3.51(1.70)	3.54(1.70)
Layer 2 Side 0	3.03(1.52)	3.04(1.56)
Layer 2 Side 1	3.51(1.79)	3.51(1.80)
Layer 3 Side 0	3.52(2.41)	3.50(2.39)
Layer 3 Side 1	3.41(2.14)	3.39(2.05)

Table 4.4: The mean (RMS) noise occupancy (scaled by 10^5) for the grounding runs detailed in the text.

of 1 fC (using the new response curve settings), and both the SCT and TRT were triggered from a pulse with an approximate frequency of 50 Hz.

Average occupancies for each layer and side of the SCT barrel are given in Table 4.4. It is concluded that the grounding boards do not affect the noise occupancy.

4.5.5   BOC Coarse Delay Scan

Since the TRT and SCT were run together for the first time during the SR1 tests, an important point to investigate is whether the TRT induces any noise in the SCT. Runs numbers 2869 and 2870 were devoted to answering this question, and were taken at a trigger rate of 50 Hz with thresholds set to the trim target settings. In run 2869 the TRT was off, whilst in 2870 the TRT was on.

The TRT starts reading out about 3.25 μs after a trigger. The SCT was therefore triggered about 3 μs after triggering the TRT to search for pickup induced in the SCT by the TRT readout cycle. The BOC Tx coarse delay was then scanned in steps of 75 ns in order to add extra delay between the triggering of the SCT and the TRT. Each time a new value was assigned to the coarse delay, the trigger word was set to a new value (updated in steps of 3). If no noise pickup occurs, the noise occupancy should be flat when plotted against trigger word (which is essentially a time variable here), as no extra noise occupancy will be observed when moving around the point at which the TRT is read out.

Plots of the noise occupancy vs trigger word for runs 2869 (TRT off) and 2870 (TRT on) are given in Figure 4.10 for 1XX, X1X and XX1 hits. The analysis was performed with a veto on strips with a noise occupancy greater than 500×10^{-5}. The noise occupancy is flat in each plot, and thus it is concluded that there is no evidence for noise pickup from the TRT.

4.5.6   Final Noise Occupancies

The distribution of noise occupancies for all active wafers in the SR1 tests is shown in Figure 4.11. The data were generated using the first 10,000 events of data set 3072, with the module thresholds set to 1 fC using the new response curve settings. The analysis was performed with a veto on strips with an occupancy greater than 500×10^{-5}.

All occupancies lie below the design specification of 5×10^{-4}, though there is slight increase from previous tests of the barrels operating individually (a mean of 4.5×10^{-5} is quoted in reference[37], compared to 4.76×10^{-5} here). Although it is possible in principle that this small increase will scale badly with the number of modules that are connected, previous noise tests performed with differing numbers of modules on the final SCT barrels showed no evidence that the performance degrades in the large system[37]. It would thus appear that the SCT is running comfortably within the limits set by its design specification.

A particular point of interest is that side 1 is noisier than side 0, as seen in Figure 4.12. This has been observed in earlier test periods, and the precise cause is unknown. One suggestion is that the lower faces of modules are exposed to light leaks originating within the barrel assembly.

A final statement on the noise occupancy of the SCT will be available early next year when the barrel is installed in the pit with its final grounding scheme.

4.6   Summary and future plans

A tool for the online monitoring of SCT noise occupancy has been developed and debugged within the framework of a general online monitoring package. The tool reproduces the results of previous calibration runs, as well as confirming the results of other offline analyses.

Running on the SR1 data sets has revealed that there seems to be no noise pickup from the TRT, that the grounding boards used during testing had no effect on the noise occupancy
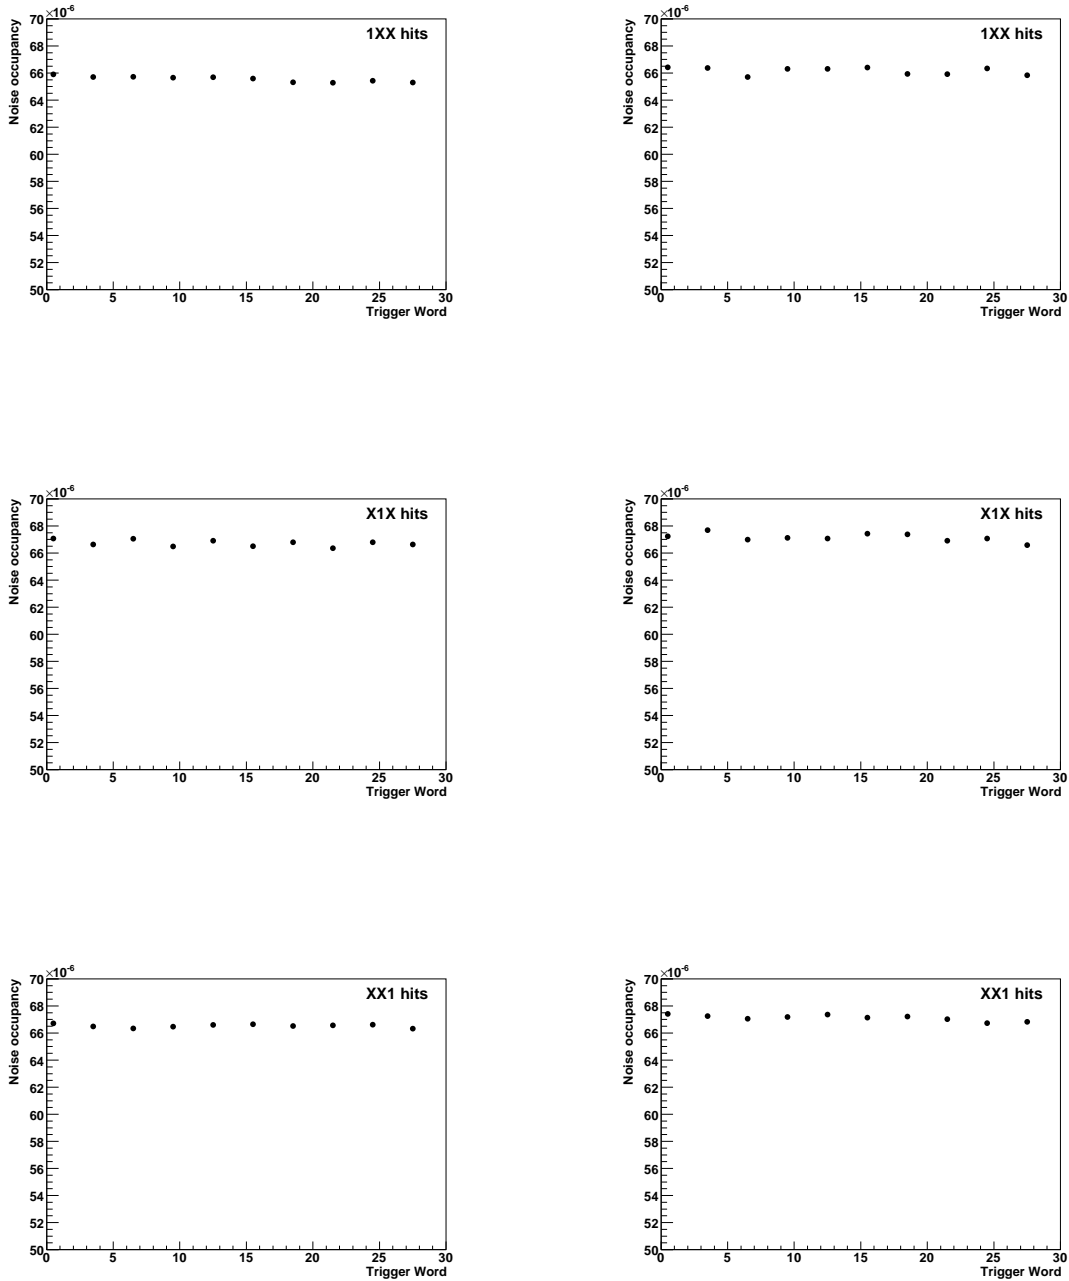
Figure 4.10: Noise occupancy vs trigger word plots for runs 2869 (left) and 2870 (right), with the plots representing, from top to bottom, 1XX, X1X and XX1 hits. The TRT was off during run 2869, and thus the data provide a series of calibration plots with no noise pickup from the TRT. Run 2870 was performed with the TRT on.
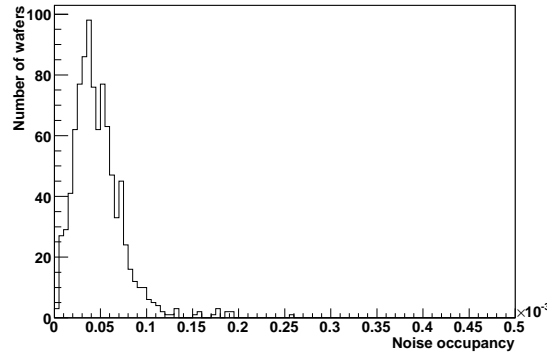
Figure 4.11: The distribution of noise occupancies for all active wafers of the SCT barrel detector, obtained using run number 3072 (the thresholds were set to 1 fC using the new response curve settings). The distribution has a mean of 4.76×10^{-5} and an RMS of 2.76×10^{-5}.
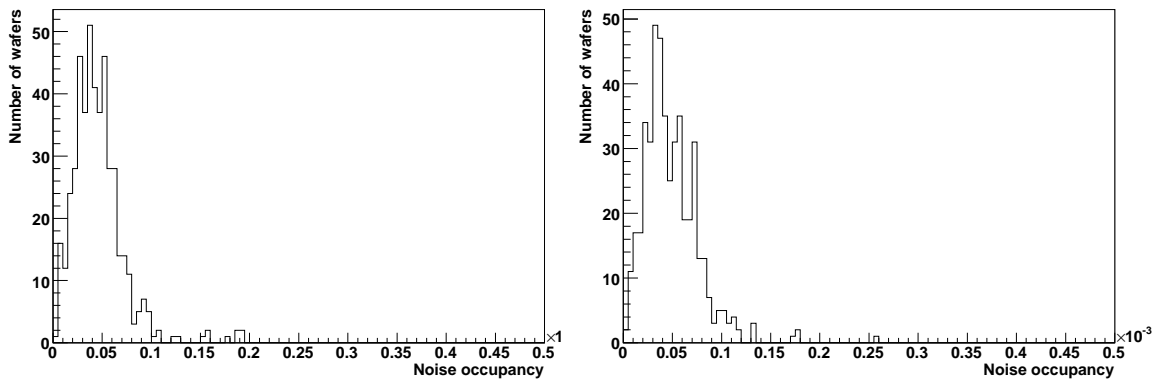
Figure 4.12: The distribution of noise occupancies for the active wafers of the SCT barrel detector for side 0 (left) and for side 1 (right). The plots show data obtained in run number 3072, with the module thresholds set to 1 fC using the new response curve settings. The side 0 distribution has a mean of 4.62×10^{-5} and an RMS of 2.70×10^{-5}, whilst the side 1 distribution has a mean of 4.90×10^{-5} and an RMS of 2.80×10^{-5}.

and that the trigger rate does not affect the noise occupancy. Furthermore, a trigger rate scan demonstrated that all modules obey the expected linear relationship between $\ln(N)$ and t^2, and that the ENC values for each module are within the design specification except for those with known defects. The final noise occupancy values, with a veto on strips with noise occupancy greater than 500×10^{-5}, all lie within the design specification.

The SCT online monitoring package remains in continuous development, and the final version will not be ready until 2007, when cosmic commissioning in the pit is due to start. The most immediate requirement is to extend the package to include the end-caps, and there are still decisions to be made regarding exactly how many histograms to include in the package, and the precise form in which to present the noise occupancy information. It is also important that the package can log defects in the ATLAS conditions database, so that problem strips can be masked off when performing data analysis. Additionally, changes in the noise occupancy over time need to be handled in an intuitive manner, so that shift workers can easily determine sudden changes in strip performance that will need to be corrected before the next data run.

Notwithstanding the need for improvements, the package has most of the necessary algorithms in place, and has already made significant contributions to monitoring during the SR1 cosmic commissioning period at CERN. The SCT barrel is currently being prepared for transport into the ATLAS cavern, to be followed by 11 weeks of cabling. The end-caps are due to be moved in early December 2006.

This chapter concludes the detector work presented in this thesis, and the focus now shifts on to the particle measurements that may be obtained using the ATLAS detector as a whole. The following chapters review the use of ATLAS in the search for supersymmetric particles, and present a new method for using ATLAS data to constrain the parameters of the SUSY Lagrangian.

Chapter 5

Experimental Supersymmetry Searches at ATLAS

Chapter 1 discussed the theoretical motivation for supersymmetry, and considered some of the experimental constraints on the theory. The aim of this section is to review current strategies and methods for (R-parity conserving) supersymmetry searches at the LHC, including original studies that have been performed using existing techniques. By the end of this chapter, the problems in the current approach will become obvious, and these will be addressed in the next chapter.

5.1   General characteristics of SUSY events at ATLAS

It has already been found that, if supersymmetry is to provide a natural solution to the problems of the standard model, one would expect to see sparticles at the LHC (see Chapter 1) and thus much effort has been focused on developing strategies for SUSY searches in ATLAS. Owing to the large number of parameters that arise from the unknown SUSY breaking sector of the theory, it is very difficult to perform general analyses that encompass all phenomenological possibilities, and hence practically all work to date has focused on particular 'reduced' models such as mSUGRA or GMSB, following the strategy laid out in the ATLAS physics Technical Design Report[25].

SUSY events are dominated by the production of gluinos and squarks, which will decay through one or more intermediate particles before producing stable $\tilde{\chi}_1^0$'s that leave the detector unseen. Thus, if weak scale supersymmetry exists one expects to see events with large missing

energy, high jet multiplicity and one or more isolated leptons. These are quite general characteristics, and should be enough to distinguish SUSY from the SM and thus formally 'discover' supersymmetry [1]. Once it is established that ATLAS data are consistent with SUSY, one can start to isolate specific processes in an attempt to measure sparticle masses, and subsequently learn more about the SUSY Lagrangian.

There are essentially two categories of measurement one may make in ATLAS: **inclusive** measurements, which involve variables taken over the whole event sample, and **exclusive** measurements, in which one determines the properties of specific decay processes. Both of these will prove essential in identifying the mechanism of SUSY breaking at the LHC, and will be revisited many times in the following text. Thus, it is worth considering briefly some key examples of each class of variable.

5.1.1   Inclusive measurements

Because inclusive measurements are taken over the whole event sample, they will be the first measures used to try and identify the presence of SUSY in ATLAS events, as one should be able to see deviations from the SM in only a few fb^{-1} of data. Examples include the missing transverse energy in the event (E_T^{miss}), the number of jets, the p_T of the hardest jets, the mean number of leptons, and variables that describe the shape of the event (e.g.sphericity, oblateness and thrust, covered in appendix B). These discriminate between the SM and SUSY, and also provide some sensitivity to the underlying SUSY model.

An important inclusive variable for SUSY discovery is the 'effective mass', M_{eff}, originally defined as follows:

$$M_{\mathrm{eff}} = E_T^{\mathrm{miss}} + p_{\mathrm{T},1} + p_{\mathrm{T},2} + p_{\mathrm{T},3} + p_{\mathrm{T},4} \tag{5.1}$$

where $p_{\mathrm{T},i}$ is the transverse momentum of the ith hardest jet in the event. Note that the definition does not conventionally include leptons. M_{eff} is correlated to the SUSY mass scale, though work in[61] demonstrated that the best model independent measure of the effective SUSY mass scale actually results from extending the jet sum to include all jets in the event.

---

[1]It is noted that there are, of course, other possible extensions to the SM that do not involve supersymmetry, and so the discovery of SUSY will also involve demonstrating convincingly that any non-SM signatures observed are actually due to sparticle decays and nothing else. A possible alternative model that has been considered recently is that of Universal Extra Dimensions (UED), which may mimic SUSY signatures unless the spin of the particles can be measured[60]. Specific comparisons between such possible models are beyond the scope of this thesis.

Hence, all references to M_{eff} in this thesis will refer to the following definition:

$$M_{\text{eff}} = E_T^{\text{miss}} + \sum_i |p_{\text{T},i}| \qquad (5.2)$$

5.1.2 Exclusive measurements

Exclusive measurements are those performed on specific decay processes, such as particular sparticle decays or decay chains. Although weak scale SUSY should first be discovered through the use of inclusive variables, any attempt to understand how SUSY is broken will require the measurement of particular sparticle masses and decay widths, and hence will require the observation of more specific processes. For this reason, the bulk of the work performed to date within the ATLAS collaboration that is relevant to measuring SUSY masses (and subsequently the GUT scale parameters) has involved exclusive signatures, with much attention paid to investigating different points in the reduced parameter spaces of mSUGRA, GMSB, AMSB, etc, with a view to mapping their characteristics.

Although there is undoubtedly some merit in this approach, it is clearly limited by the fact that there is no particular motivation for any of these reduced parameter sets, let alone for one particular point within these model choices. Furthermore, as we shall see, even within one particular model choice it will be difficult to obtain unambiguous mass measurements. There is also a separate issue related to the fact that if exclusive data alone is used, one ignores the wealth of useful information contained in inclusive variables. Nevertheless, the strategy at least enables one to explore the often dramatic changes in phenomenology over the whole parameter space, and the hope is that this will form a useful guide for the interpretation of ATLAS data should SUSY be found.

In this chapter, I present two original examples of the use of exclusive signatures within the mSUGRA parameter space, with the specific goal of measuring sparticle masses. The first of these, performed on a point in the so-called 'co-annihilation region' will be used to demonstrate how kinematic endpoints in invariant mass distributions can be used to reconstruct the masses involved in the decay chain. The second example, performed in the 'funnel region' of mSUGRA space, gives a concrete example of the fact that, even in a reduced parameter space such as mSUGRA, one observes dramatically different phenomenologies as the parameters are varied and hence cannot in principle generalise this kind of analysis to all points without substantial further work. In the next chapter, I will address these problems by introducing a method that combines inclusive and exclusive data.

5.2   Co-annihilation region

5.2.1   Definition of co-annihilation model

Of the regions favoured by WMAP that are displayed in Figure 1.3, the region at low m_0 and extending to high $m_{1/2}$ is the stau co-annihilation region, in which the stau is light enough to have efficient co-annihilation with the lightest neutralino (indeed, the region runs along the boundary that is excluded due to the fact that the stau would become the LSP). A point was chosen for study in this region with the following mSUGRA parameters:

$$m_0 = 70 \text{ GeV}, m_{1/2} = 350 \text{ GeV}$$

$$\tan\beta = 10, A = 0, \mu > 0$$

The masses of the most relevant particles are contained in Table 5.1, whilst branching ratios for some of the most significant decay processes are given in Table 5.2. Although similar to the point 5 analysed in the ATLAS Physics TDR[25], this particular case is somewhat more complicated owing to the small mass differences between the $\tilde{\chi}_1^0$ and the $\tilde{e}_R$ and between the $\tilde{\chi}_2^0$ and the $\tilde{e}_L$. This leads to the production of soft leptons, and these may be missed by the ATLAS detector thus reducing the statistics in any resulting decay chain plot.

5.2.2   Monte Carlo event generation and detector simulation

The first step in analysing a point in the SUSY parameter space is to take the mSUGRA parameters given above and calculate the mass spectrum of SUSY particles at the weak scale by solving the renormalisation group equations. This was accomplished in this instance using `ISAJET v7.69`[62], in conjunction with `ISAWIG`. The latter produces an input file for the `HERWIG 6.5`[63–65] Monte Carlo event generator, which produces simulated events.

These are then passed through the `ATLFAST` detector simulation, which simulates the effects of detector resolution for the ATLAS detector at the LHC. Jets are reconstructed using a cone algorithm with a cone size $\Delta R = 0.4$. Electrons, muons and jets were subject to a minimum p_T cut of 5, 5 and 10 GeV respectively.

Note that the `ATLFAST` reconstruction algorithms affect the ability to reconstruct leptons and jets in close proximity, and this is potentially a source of systematic error in the endpoint

Particle	Mass/GeV
$\tilde{\chi}_1^0$	137
$\tilde{\chi}_2^0$	264
$\tilde{e}_L, \tilde{\mu}_L$	255
$\tilde{e}_R, \tilde{\mu}_R$	154
$\tilde{\tau}_1$	147
$\tilde{\tau}_2$	257
$\tilde{g}$	832
$\tilde{u}_L$	760
$\tilde{u}_R$	735
$\tilde{d}_L$	764
$\tilde{d}_R$	733
$\tilde{b}_1$	698
$\tilde{b}_2$	723
$\tilde{t}_1$	574
$\tilde{t}_2$	749
h	116

Table 5.1: The most important particle masses at the mSUGRA point defined in the text.

Process	Branching Ratio
$\tilde{\chi}_2^0 \rightarrow \tilde{l}_R l$	2%
$\tilde{\chi}_2^0 \rightarrow \tilde{l}_L l$	29%
$\tilde{\chi}_2^0 \rightarrow \tilde{\tau}_1 \tau$	18%
$\tilde{\chi}_2^0 \rightarrow \tilde{\tau}_2 \tau$	2%
$\tilde{\chi}_2^0 \rightarrow \tilde{\chi}_1^0 h$	48%

Table 5.2: Branching ratios for important decay processes occurring at the mSUGRA point defined in the text. $\tilde{l}$ is used to denote selectrons and smuons, whilst l denotes electrons and muons.

observations (particularly in the threshold position). This will also occur in full simulation, though to a lesser extent. A study of these systematic effects in both the fast and full simulation is long overdue, but is sufficiently complicated to warrant a separate study. It is important to remember that this affects all previous endpoint analyses, and is not specific to that considered here.

A sample corresponding to 100fb^{-1} was generated, representing one year of design luminosity.

5.2.3   Search for squark decay

Previous studies (for example, in[25]) have illustrated the procedure of searching for kinematic edges in the various invariant mass distributions resulting from a given event. The procedure is used here in conjunction with the decay:

$$\tilde{q} \to q\tilde{\chi}_2^0 \to ql^{\pm}\tilde{l}_L^{\mp} \to ql^{\pm}l^{\mp}\tilde{\chi}_1^0 \qquad (5.3)$$

This is an excellent starting point for analysis due to the clear signature provided by the two opposite-sign, same-flavour leptons. The left-handed slepton is considered here rather than the right-handed slepton due to the much greater branching ratio $\text{BR}(\tilde{\chi}_2^0 \to \tilde{l}_L e)$.

If one considers being in the rest frame of the squark, the invariant mass of the two decay products is equal to the rest mass of the squark. The same logic applies to each of the sparticles in the chain- the invariant mass of the decay products must have a maximum value. Hence, the invariant masses of different combinations of decay products from the chain (e.g. the two leptons, or the two leptons and the quark) have theoretical maxima whose precise values are calculated by boosting between the rest frames of the sparticles in the chain and working through the kinematic algebra. These maxima appear as edges in the invariant mass plots associated with the decay products, and each edge position is a function of the sparticle masses in the chain. Thus, by obtaining values for at least four edge positions in visible decay distributions, one can measure the masses in the chain.

The following endpoints are expected to be observed in the invariant mass spectra associated with this decay chain ($\tilde{\psi} = m_{\tilde{\chi}_2^0}^2, \tilde{q} = m_{\tilde{q}}^2, \tilde{l} = m_{\tilde{l}_L}^2, \tilde{\chi} = m_{\tilde{\chi}_1^0}^2$)[66]:

$$(m_{ll}^2)^{\text{edge}} = \frac{(\tilde{\psi} - \tilde{l})(\tilde{l} - \tilde{\chi})}{\tilde{l}} \tag{5.4}$$

$$(m_{llq}^2)^{\text{edge}} = \begin{cases} \max\left[\frac{(\tilde{q}-\tilde{\psi})(\tilde{\psi}-\tilde{\chi})}{\tilde{\psi}}, \frac{(\tilde{q}-\tilde{l})(\tilde{l}-\tilde{\chi})}{\tilde{l}}, \frac{(\tilde{q}\tilde{l}-\tilde{\psi}\tilde{\chi})(\tilde{\psi}-\tilde{l})}{\tilde{\psi}\tilde{l}}\right] \\ \text{except when } \tilde{l}^2 < \tilde{q}\tilde{\chi} < \tilde{\psi}^2 \text{and } \tilde{\psi}^2\tilde{\chi} < \tilde{q}\tilde{l}^2 \\ \text{where one must use}(m_{\tilde{q}} - m_{\tilde{\chi}_1^0})^2. \end{cases} \tag{5.5}$$

$$(m_{lq}^2)^{\text{edge}}_{\text{max}} = \max\left[\frac{(\tilde{q} - \tilde{\psi})(\tilde{\psi} - \tilde{l})}{\tilde{\psi}}, \frac{(\tilde{q} - \tilde{\psi})(\tilde{l} - \tilde{\chi})}{\tilde{l}}\right] \tag{5.6}$$

$$(m_{lq}^2)^{\text{edge}}_{\text{min}} = \min\left[\frac{(\tilde{q} - \tilde{\psi})(\tilde{\psi} - \tilde{l})}{\tilde{\psi}}, \frac{(\tilde{q} - \tilde{\psi})(\tilde{l} - \tilde{\chi})}{(2\tilde{l} - \tilde{\chi})}\right] \tag{5.7}$$

$$(m_{llq}^2)^{\text{thres}} = \frac{2\tilde{l}(\tilde{q} - \tilde{\psi})(\tilde{\psi} - \tilde{\chi}) + (\tilde{q} + \tilde{\psi})(\tilde{\psi} - \tilde{l})(\tilde{l} - \tilde{\chi}) - (\tilde{q} - \tilde{\psi})\sqrt{(\tilde{\psi} + \tilde{l})^2(\tilde{l} + \tilde{\chi})^2 - 16\tilde{\psi}\tilde{l}^2\tilde{\chi}}}{4\tilde{l}\tilde{\psi}} \tag{5.8}$$

where "min" and "max" refer to minimising and maximising with respect to the choice of lepton. In addition, "thres" refers to the threshold that appears in the m_{llq} distribution when events are chosen such that the angle between the two lepton momenta exceeds $\pi/2$ in the slepton rest frame, corresponding to $m_{ll}^{\text{edge}}/\sqrt{2} < m_{ll} < m_{ll}^{\text{edge}}$.

5.2.4 Note on endpoint positions

The formulae for endpoint positions presented in Section 5.2.3 take a squark mass as input. In reality not all squarks have the same mass, and so chains containing squarks with different masses will have endpoints at slightly different positions. This effect manifests itself as a smearing of the endpoints in any plots of experimental or simulated data. Plots of this kind

are shown in Section 5.2.5 and when indicating the positions at which endpoints are expected to be found in these plots, one is required to choose a "typical" squark mass for insertion into the relevant endpoint formula. A 750 GeV "typical" squark mass is chosen here for this purpose, although it should be borne in mind that the actual endpoints seen in the plots will be somewhat smeared due the the non-degeneracy of the squark masses contributing to them.

5.2.5   Invariant mass distributions

Cuts

In order to see the above edges clearly, one must apply various cuts to the event data in order to isolate a clean sample of the squark decay chain. Here, one can select events with the Opposite Sign Same Flavour (OSSF) lepton signature described above, and one can also exploit the fact that the required events have a large amount of missing energy (due to the departure from the detector of two invisible $\tilde{\chi}_1^0$'s). Furthermore, one expects to obtain hard jets in SUSY events, resulting from the decay of gluinos and squarks. All plots are obtained through the use of the following cuts:

- $E_T^{\mathrm{miss}} > 300$ GeV;

- exactly two opposite-sign leptons with $p_T > 5$ GeV and $|\eta| < 2.5$;

- at least two jets with $p_T > 150$ GeV;

Although the cuts chosen are similar to those used for point 5 in the ATLAS Physics TDR, there are some exceptions. For example, one needs to impose a p_T cut on the leptons in the event due to the fact that `ATLFAST` is not parametrised properly for low p_T leptons, and yet it is essential to pick up soft leptons due to the small mass differences that crop up in the decay chain. Hence, a compromise between these two factors must be chosen. Some plots are the result of additional cuts, and these are given below.

The SM background for dilepton processes is generally negligible once these cuts have been applied, as can be seen from Figure 20-11 in reference[25] in which the SM background is less than 1 per cent of the SUSY signal, for a representative mSUGRA point [2]. However,

---

[2]Of course, this does not guarantee that the SM background is negligible over the whole parameter space, as, although the SM background does not change with the SUSY parameters, the SUSY signal may be weaker in some areas. This should not affect the point studied here, however, and it is therefore assumed that the dominant source of background is that of other SUSY processes.

the OSSF lepton signature can be produced by SUSY processes other than the decay of the $\tilde{\chi}_2^0$. One would expect these to produce equal amounts of opposite-sign opposite-flavour (OSOF) leptons and hence one can often subtract the dilepton background by producing "flavour subtracted plots" in which one plots the combination $e^+e^- + \mu^+\mu^- - e^+\mu^- - e^-\mu^+$. This is only possible in cases where there are sufficient statistics, and was not possible for every plot below.

m_{ll} plot

As seen in Figure 5.1, a sharp edge is produced in the spectrum at ≈ 58 GeV, and this is a very clear signature. There is, however, a second edge visible at ≈ 98 GeV, resulting from a squark cascade decay featuring a right-handed selectron rather than a left-handed selectron. In practise, it will be very difficult to assign these two edges correctly, and we see here a first example of decay chain ambiguity. It will be observed in Chapters 6 and 7 that this is not merely limited to choice of slepton, but could extend to other possibilities, and we therefore need a more general approach than that used here in order to reconstruct masses.

m_{llq} plot

This is produced by selecting events with exactly two leptons, and forming the invariant mass m_{llq} with each of the two highest p_T jets. Given that the hardest jets in events containing cascade decays of squarks and gluinos usually come from the processes $\tilde{q} \to \tilde{\chi}_2^0 q$ and $\tilde{q} \to \tilde{\chi}_1^0 q$, the lower of the two invariant masses formed in this way should lie below the llq edge defined by Equation (5.5).

As shown in Figure 5.2, a clear endpoint is seen in the distribution at ≈ 600 GeV.

m_{llq}^{thres} plot

This is reproduced in Figure 5.3, and is constructed by applying the cuts defined in Section 5.2.5 along with the additional constraint that m_{ll} must exceed $m_{ll}^{\text{max}}/\sqrt{2}$. Furthermore, because a threshold is expected, the higher of the two llq masses is chosen.

A threshold is clearly observed, though it is difficult to state its precise position due to uncertainty in the expected shape of the edge and the manner in which it is modified by detector effects. Also, the extra cut limits the statistics. There is in fact a series of thresholds resulting from the presence of different squarks in the decay chains contained in the event sample (with
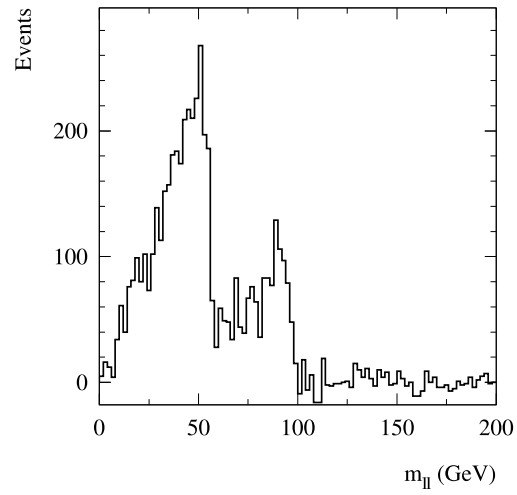
Figure 5.1: The flavour-subtracted dilepton invariant mass plotted with the cuts described in the text.
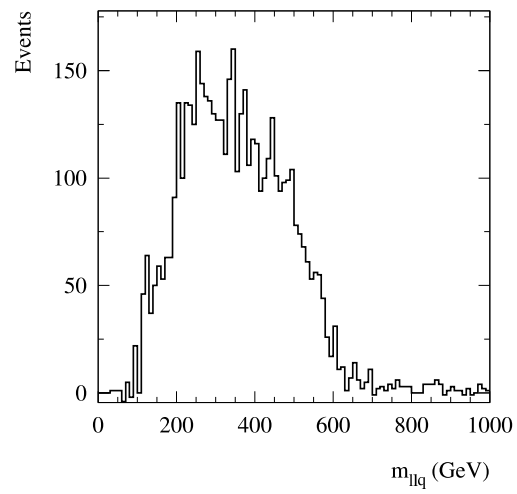
Figure 5.2: The llq invariant mass plot.
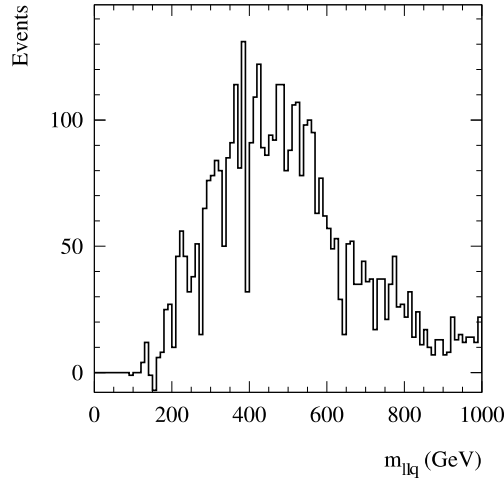
Figure 5.3: The llq invariant mass threshold plot.

different squark masses). This essentially produces a smearing of the threshold, and there has yet to be a satisfactory solution to this problem.

m_{lq}^{max} and m_{lq}^{min} plots

These are plotted by forming llq invariant masses with the two hardest jets in the event. The jet from the lowest mass combination (which is our best guess for the quark emitted in the squark cascade decay) is then used to form the lq invariant mass with each of the leptons in the events. The maximum of these is plotted in the m_{lq}^{max} plot (shown in Figure 5.4) and the minimum is plotted in the m_{lq}^{min} plot (shown in Figure 5.5). For the m_{lq}^{max} plot, events were selected using the cuts defined above, but with the additional constraint that one of the llq invariant masses formed with the two hardest jets must be above the llq endpoint, and the other must be below.

The m_{lq}^{min} plot has one additional cut: the dilepton invariant mass must be less than the value of m_{ll}^{max} observed in Figure 5.1.

Both plots exhibit endpoints, and the edge is particularly abrupt in the m_{lq}^{max} histogram. Although there are events beyond the endpoint in the m_{lq}^{min} plot (due to SUSY background processes), there is nevertheless a convincing edge at ≈ 180 GeV.
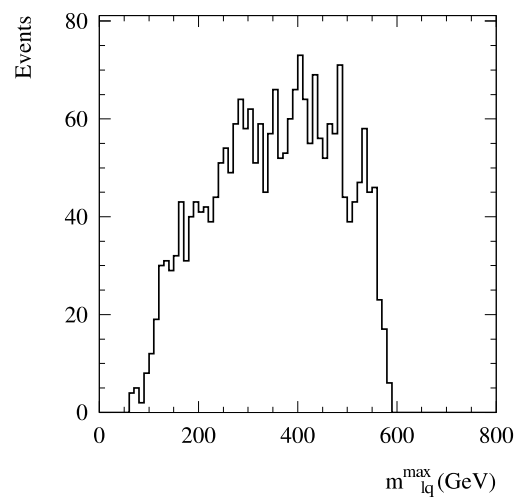
Figure 5.4: The lq max invariant mass plot.
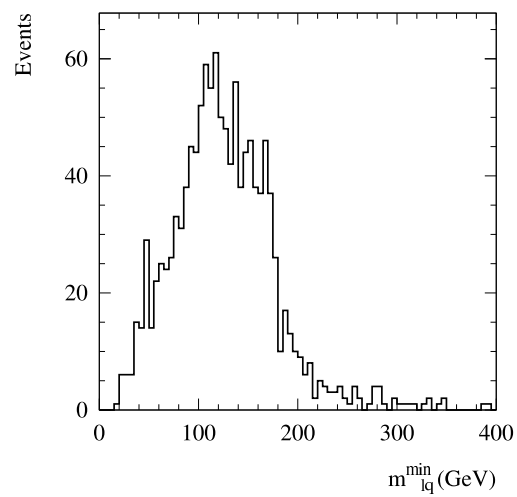
Figure 5.5: The lq min invariant mass plot.

Edge	Predicted (GeV)	Observed (GeV)
ll edge	57.64	57.5±2.5
llq edge	600.1	600±10
llq threshold	134.0	150±30
lq max edge	592.1	590±10
lq min edge	181.7	180±10

Table 5.3: Predicted and observed edge positions for the mSUGRA mass point described in the text. Error estimates have been obtained 'by eye', and reflect a lack of information regarding the precise shapes of the endpoints.

5.2.6 Comparison of observed and predicted edges

The edges predicted by Equations (5.4) to (5.8) are summarised in Table 5.3, where a common squark mass of 750 GeV is used, and all other masses are taken from Table 5.1. The observed positions of the endpoints are also given.

It is common when extracting the observed edges from plots such as those above to fit a function to the endpoint in order to determine both the precision and the accurate position. For the purposes of this analysis, the edges, and their estimated errors, have been determined 'by eye'[3] for several reasons. Firstly, not all edges can be fitted with functions (in the case of the llq threshold, for example, the correct shape is not known). Indeed, recent work highlights the need for caution in applying these functions too readily without first investigating the theoretical shape of the distribution[67], as endpoints can often exhibit tails or 'feet' that will be confused with SUSY background and hence may lead to inaccurate measurements. The shapes of the endpoints for distributions involving quarks vary significantly over the parameter space, introducing a model dependence into the precision with which one may realistically measure endpoint positions and hence masses. Given that the purpose of this discussion is merely to present a typical analysis based on exclusive data in order to discuss the limitations of current techniques, however, it is sufficient here to reflect the current state of ignorance by using the conservative errors given in Table 5.3. For those interested, the fitting of endpoint functions has been done in work leading to reference [68] which contains estimates of the

---

[3]Endpoint positions were obtained by looking for places where the distributions suddenly start to decay into the tail, and the errors were ascertained by looking at the shape of the distribution around this endpoint. Endpoints that are less steep are assigned larger errors. This process, though crude, will not lead to incorrect results given that the errors were assigned very conservatively. In fact, one would expect the final results to improve if more care was taken over the assignment of these errors.

precision expected if one were to take a more optimistic view.

5.2.7   Mass reconstruction

Having measured four endpoints and one threshold, one can in principle now try and reconstruct the masses of the sparticles involved in the squark cascade decay chain. In order to do so, it is necessary to assume that one has correctly identified the edges in the dilepton invariant mass distribution, and that it is certain that the edges observed in each of the distributions are indeed produced by the squark decay process outlined in Equation 5.3. These assumptions lie behind almost all current ATLAS analyses, and will therefore be retained in this chapter despite being somewhat unsatisfactory. An answer to this problem will be one of key developments presented in the next chapter. The endpoint equations will be solved using a Markov Chain Monte Carlo method, details of which may be found in Appendix A. The technique is an excellent way of efficiently exploring high dimensional parameter spaces, and will be a crucial component of the analysis presented in Chapter 6.

The five endpoints observed in the previous section essentially provide an (over-constrained) set of simultaneous equations in the four unknowns $m_{\tilde{q}}$, $m_{\tilde{l}_L}$, $m_{\tilde{\chi}_2^0}$ and $m_{\tilde{\chi}_1^0}$, and these can be solved to determine the masses. Given a set of observed edges $\mathbf{e}^{obs}$, and a set of postulated masses $\mathbf{m}$, the ultimate goal is to evaluate $p(\mathbf{m}|\mathbf{e}^{obs})$ and thus to find the regions of mass space favoured by the data. The best way of doing this is to sample masses $\mathbf{m}$ from $p(\mathbf{m}|\mathbf{e}^{obs})$, subsequently histogramming the samples to reveal directly the shape of the probability distribution.

Using Bayes' Theorem we know that

$$p(\mathbf{m}|\mathbf{e}^{obs}) \propto p(\mathbf{e}^{obs}|\mathbf{m})p(\mathbf{m}). \tag{5.9}$$

The prior $p(\mathbf{m})$ is chosen to be uniform[4] over the mass space considered. This choice seems a good as any other, and has the added benefit that plots of the posterior distribution $p(\mathbf{m}|\mathbf{e}^{obs})$ are also just plots of the likelihood $p(\mathbf{e}^{obs}|\mathbf{m})$, permitting the effects of other priors $p(\mathbf{m})$ to be easily imagined.

---

[4]Some points $\mathbf{m}$ in mass space do not satisfy the hierarchy $m_{\tilde{q}} > m_{\tilde{\chi}_2^0} > m_{\tilde{l}_L} > m_{\tilde{\chi}_1^0} > 0$ required by the chosen decay chain. Under the model, then, these points yield $p(\mathbf{e}^{obs}|\mathbf{m}) = 0$ and are vetoed. While this veto is technically part of the likelihood (given the model) it simplifies later discussion of the likelihood in more complicated scenarios if the veto is pulled out of the likelihood and moved into the prior $p(\mathbf{m})$. In practise then, the effective prior is uniform over all of the region of mass space in which the required hierarchy is present, and zero elsewhere. The effect is the same as if the veto had been left in the likelihood, but the likelihoods will be simpler to describe and define.

One can sample from $p(\mathbf{m}|\mathbf{e}^{obs})$ using the Metropolis Method as follows. First a mass point $\mathbf{m}$ is chosen, and $p(\mathbf{m}|\mathbf{e}^{obs})$ is evaluated using Equation (5.9). For the edges e_1, e_2, e_3, e_4, and e_5, the likelihood $p(\mathbf{e}^{obs}|\mathbf{m})$ is given by the product

$$p(\mathbf{e}^{obs}|\mathbf{m}) = \prod_{i=1}^{5} p(e_i^{obs}|\mathbf{m}),$$

(5.10)

where

$$p(e_i^{obs}|\mathbf{m}) \approx \frac{1}{\sqrt{2\pi\sigma_i^2}} \exp\left(-\frac{(e_i^{obs} - e_i^{pred}(\mathbf{m}))^2}{2\sigma_i^2}\right)$$

(5.11)

in which σ_i is the statistical and fit error associated with the edge measurement of edge e_i, and where e_i^{obs} and $e_i^{pred}(\mathbf{m})$ are respectively the observed and predicted positions of the edge. This probability distribution assigns a weight $p(\mathbf{m}|\mathbf{e}^{obs})$ to each point $\mathbf{m}$ in mass space, including the errors associated with the endpoint measurements. Note that $p(\mathbf{m}|\mathbf{e}^{obs})$ is the equivalent of the $P^*(\mathbf{x})$ defined in the appendix in Equation (A.2), as it is defined only up to an unknown normalisation constant.

So, in order to plot the probability distribution, one follows the following steps of the Metropolis Algorithm:[5]

1. A new mass point $\mathbf{m}^{proposal}$ is suggested on the basis of the current point $\mathbf{m}^{current}$. The mass-space proposal distribution for the Metropolis Algorithm was chosen to be a 4-dimensional Gaussian whose width in each dimension was 5 GeV and whose centre was the position of the current point $\mathbf{m}^{current}$. The width was chosen to match the typical width of the lines in mass space, and should not effect the final result once convergence has occurred.

2. $p(\mathbf{m}^{proposal}|\mathbf{e}^{obs})$ is evaluated at the proposed point.

3. A decision is made on whether to jump to the new point, or remain at the current point on the basis of the ratio of $p(\mathbf{m}^{proposal}|\mathbf{e}^{obs})$ to $p(\mathbf{m}^{current}|\mathbf{e}^{obs})$.

4. If a decision to *not jump* is made, then the next point in the chain $\mathbf{m}^{next}$ is again set equal to $\mathbf{m}^{current}$, otherwise it is set equal to $\mathbf{m}^{proposal}$. When proposals are rejected, therefore, successive points in the chain are duplicates of each other.

---

[5]See Appendix A.3 for discussion of the motivations behind each of these steps, and for definitions of "proposal functions" and the decision mechanism.

5. All steps are repeated until the sampler has sufficiently explored the interesting regions of parameter space.[6]

It is noted that in the real ATLAS detector, one might have a systematic shift of the endpoints due to the jet energy scale error, and this is considered in Section 6.2.

5.2.8 Mass space plots

The Metropolis sampler ensures that points which are more likely are sampled more often. One can observe the shape of the probability distribution by simply histogramming the sampled points. This is a 4 dimensional shape in mass space, which can be viewed as a projection onto each pair of axes. This is done in Figure 5.6, revealing that a lengthy region of parameter space is compatible with the edge data, and extra information is required to constrain this further. Note that the endpoint equations discussed previously are sensitive principally to mass differences, and hence one observes lines in each plane of the mass space, constraining each mass only in terms of the others. Given that the endpoint data does not set the overall mass scale, the system is able to wander to high masses without affecting the position of the endpoints provided that the mass differences remain consistent. This is a limitation of the endpoint approach, and will be addressed later.

Finally, it is noted that the lines are broader in the plots featuring squark masses, and this is due to the fact that the end points were calculated using an average squark mass, whilst the Monte Carlo events feature a range of squark masses. Hence the resolution is smeared somewhat relative to the other masses.

5.3 Funnel region

We have now completed one example of the use of exclusive data to measure sparticle masses in the ATLAS detector, and have met some of the problems associated with the method. It is instructive at this point to look at another point in the mSUGRA parameter space that is also consistent with the WMAP data and yet displays a different phenomenology.

As previously seen in Figure 1.4, one obtains a 'funnel' that is consistent with the WMAP relic density limits as one increases $\tan\beta$. This is due to the presence of a heavy Higgs boson

---

[6]For a discussion of the convergence tests used throughout this thesis, see Appendix A.5.
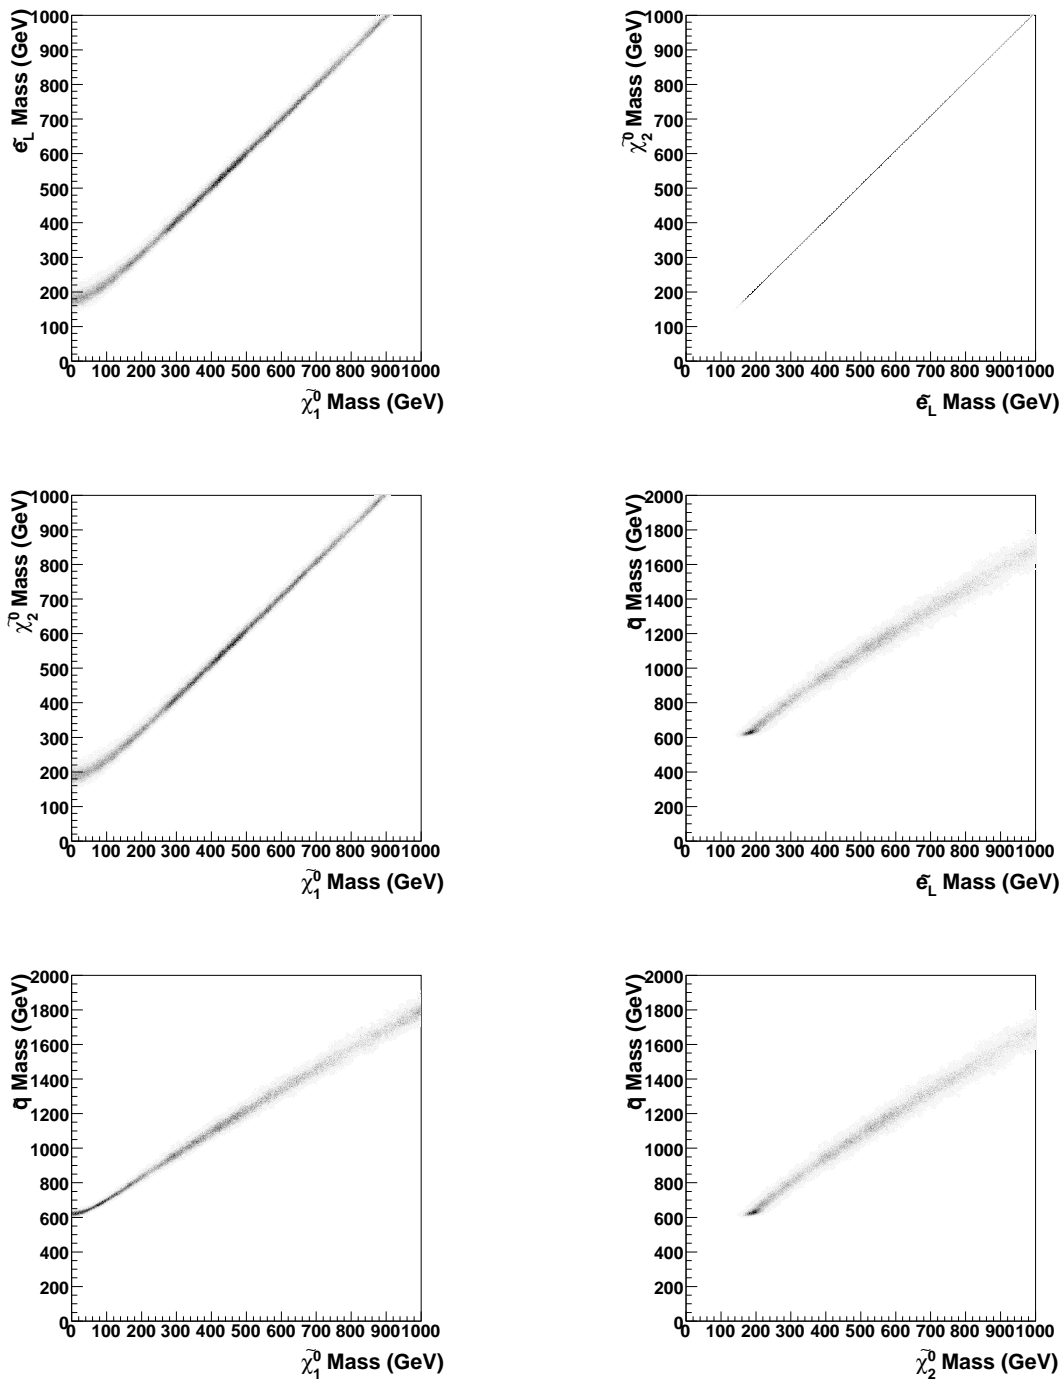
Figure 5.6: The region of mass space consistent with the kinematic edge measurements described in the text, obtained using a Markov chain sampler.

whose mass m_A is roughly twice that of the $\tilde{\chi}_1^0$, leading to enhanced $\tilde{\chi}_1^0$ annihilation. The region is a particularly interesting case for the ATLAS experiment, because the decay chain investigated previously (given in Equation 5.3) can no longer occur with a selectron or smuon present in the chain due to their mass being larger than that of the $\tilde{\chi}_2^0$. However, one can instead observe the following process in order to obtain a measurement of the stau mass:

$$\tilde{q} \rightarrow q\tilde{\chi}_2^0 \rightarrow q\tau^{\pm}\tilde{\tau}_1^{\mp} \rightarrow q\tau^{\pm}\tau^{\mp}\tilde{\chi}_1^0 \tag{5.12}$$

Thus, in order to observe endpoints, one must look at tau-jet distributions which is problematic for two reasons:

1. Tau reconstruction at ATLAS is much less efficient than electron or muon reconstruction, and is approximately of the order of 50% (compared to 90-100% for electrons and muons).

2. Tau decays involve missing energy (due to the production of neutrinos) that leads to distortion of the endpoint shape from the expected distribution.

This section presents plots analogous to those shown in Section 5.2.5.

5.3.1   Definition of funnel region model

A point was chosen for study in the funnel region with the following mSUGRA parameters:

$$m_0 = 320 \text{ GeV}, \, m_{1/2} = 375 \text{ GeV}$$

$$\tan\beta = 50, \, A = 0, \, \mu > 0$$

This has a relic density $\Omega_\chi h^2 = 0.107$ as calculated using the `Micromegas 1.1.1` software code[69], and is thus consistent with WMAP as well as being consistent with a measurement of the branching ratio for the decay $b \rightarrow s\gamma$ and recent measurements of the anomalous magnetic moment of the muon [7]. The masses of the most relevant particles are contained in Table 5.4, whilst branching ratios for some of the most significant decay processes are given in Table 5.5.

---

[7]In addition to these constraints, it is possible to interpret recent data from the EGRET satellite experiment as evidence of a dark matter annihilation signal in the form on an excess of Galactic gamma rays with energy above 1 GeV[70]. This interpretation leads to a range of preferred neutralino masses, but has not been used in this analysis owing to the uncertain nature of the result.

Particle	Mass/GeV
$\tilde{\chi}_1^0$	150
$\tilde{\chi}_2^0$	288
$\tilde{l}_L$	412
$\tilde{l}_R$	351
$\tilde{\tau}_1$	181
$\tilde{\tau}_2$	393
$\tilde{g}$	895
$\tilde{u}_L$	867
$\tilde{u}_R$	842
$\tilde{d}_L$	871
$\tilde{d}_R$	840
$\tilde{b}_1$	717
$\tilde{b}_2$	779
$\tilde{t}_1$	642
$\tilde{t}_2$	798
h	116

Table 5.4: The most important particle masses at the mSUGRA funnel region point defined in the text.

Process	Branching Ratio
$\tilde{\chi}_2^0 \rightarrow \tilde{\tau}_1 \tau$	48%
$\tilde{\chi}_2^0 \rightarrow \tilde{\chi}_1^0 h$	3.6%
$\tilde{\chi}_2^0 \rightarrow \tilde{\chi}_1^0 Z^0$	0.8%
$\tilde{\tau}_1 \rightarrow \tilde{\chi}_1^0 \tau$	100%

Table 5.5: Branching ratios for important decay processes occurring at the mSUGRA funnel region point defined in the text.

HERWIG 6.505 and ISAJET 7.69 were used along with ATLFAST, as in the previous section. One million signal events were generated, corresponding to 194 fb^{-1} of data.

5.3.2   Tau-jet invariant mass distributions

Cuts

The cuts here are defined in a similar fashion as those used in Section 5.2.5, except that we now apply selection cuts to taus rather than the lighter leptons. All plots below are obtained through the use of the following cuts:

- $E_T^{\mathrm{miss}} > 300$ GeV;

- exactly two opposite-sign taus with $p_{T,1} > 100$ GeV, $p_{T,2} > 50$ GeV (as reconstructed by `ATLFAST`) and $|\eta_{1,2}| < 2.5$;

- at least two non-tau jets with $p_T > 150$ GeV;

One can also perform same sign tau subtraction on the plots, and this is done in all cases below.

$m_{\tau\tau}$ plot

The ditau invariant mass is shown in Figure 5.7, though it must be remembered that the neutrinos from tau decays are not detected. An endpoint is expected at 125 GeV, and this is observed, although the distribution is clearly distorted from the triangular phase space distribution that is more closely observed in Figure 5.1. The tail of events past the endpoint results from SUSY background processes, and it is seen that the most significant effect of the distortion of the endpoint shape is to confuse the position of the endpoint, and thus lower the precision with which a measurement can be made. There currently exists no satisfactory solution to this problem, and further work needs to be done to model the distortion effectively.

$m_{\tau\tau q}$ plot

The $m_{\tau\tau q}$ distribution is shown in Figure 5.8. As before, the invariant mass is formed with the two hardest (non-tau) jets in the event, and the lowest of these is histogrammed. Once more we see an endpoint that is consistent with the predicted value (in this case 690 GeV), although the shape is easy to confuse with the tail resulting from background processes.
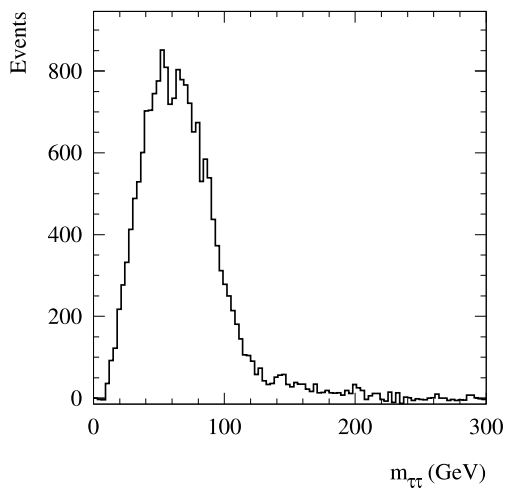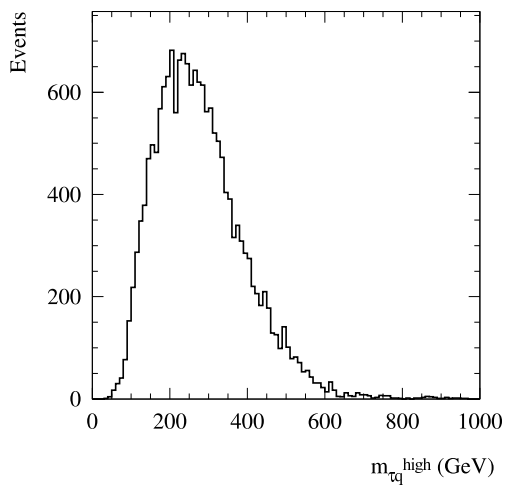
Figure 5.7: The same sign subtracted ditau invariant mass plotted with the cuts described in the text.
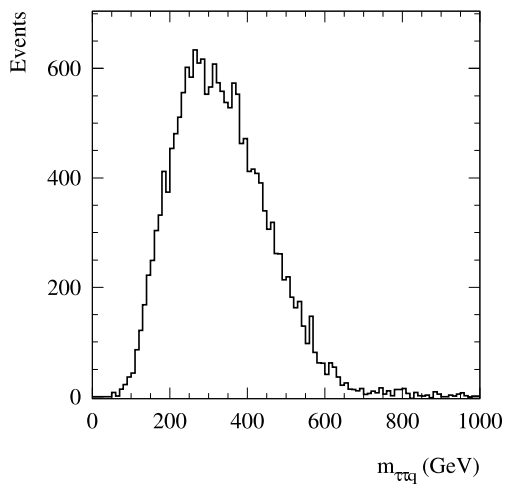
Figure 5.8: The same sign subtracted $m_{\tau\tau q}$ invariant mass plotted with the cuts described in the text.

$m_{\tau q}^{\max}$ **and** $m_{\tau q}^{\min}$ **plots**

These are shown in Figures 5.9 and 5.10, and are both subject to an additional cut specifying that the ditau invariant mass must be less than the endpoint observed in Figure 5.7 (taken to be 127 GeV). $m_{\tau q}^{\max}$ and $m_{\tau q}^{\min}$ endpoints are expected at 640 and 400 GeV respectively, and both plots are consistent with these values though the $m_{\tau q}^{\max}$ plot is more strongly affected by the tail problem described above.

5.3.3   Discussion

The second example of endpoint analysis has demonstrated that, even within the mSUGRA parameter space, one can observe dramatically different phenomenologies, and must in principle develop new techniques in order to use exclusive analysis (or, at the very least, substantially modify existing techniques). The problem of having to observe endpoints involving taus naturally leads to problems resulting from the tau reconstruction, which, coupled with an irreducible tail of SUSY background, makes it much harder to obtain precise endpoint measurements in tau-jet distributions [8]. There are two ways of solving this problem:

1. Do a thorough study of the effect of tau reconstruction on endpoint shapes (itself complicated by the fact that the shapes of the distributions themselves vary considerably over the parameter space before one takes the distortion from tau reconstruction effects into account, as seen in reference[67]).

2. Use other data in conjunction with the endpoint analysis in order to make up for the lack of precision in the endpoint measurements.

Both of these are reasonable strategies, though it is the second that will be developed in this thesis owing to the fact that it will also help to solve the problems encountered in the co-annihilation region; chiefly the problems associated with ambiguous edges and with the lack of constraint from endpoints on the overall mass scale.

---

[8]It is noted for completeness that such endpoints are sometimes visible even when channels involving the lighter leptons are also open, and in such a case, one could improve the stau mass measurement using information obtained through cascade decays featuring sleptons and smuons.

Figure 5.9: The same sign subtracted $m_{\tau q}^{max}$ invariant mass plotted with the cuts described in the text.
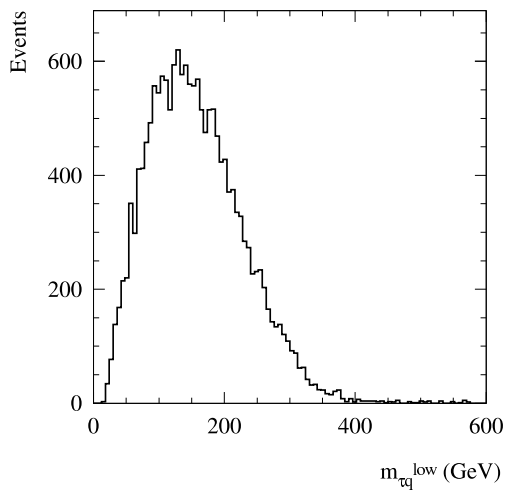
Figure 5.10: The same sign subtracted $m_{\tau q}^{min}$ invariant mass plotted with the cuts described in the text.

5.4 Summary

This chapter has reviewed the major characteristics of experimental supersymmetry searches at ATLAS, and has presented two original examples of exclusive analysis within the mSUGRA parameter space. It was seen that exclusive data alone cannot unambiguously determine SUSY masses, even at a point where we would expect to see cascade decays of a form studied in previous analyses. Furthermore, there are other points at which such processes cannot be observed, and the funnel region example gives an indication of problems that can arise when looking at alternative decays.

The Metropolis sampling algorithm has been introduced as a tool for efficiently exploring parameter spaces of high dimensionality.

Chapter 6

Combination of Inclusive and Exclusive Data

Thus far, it has been assumed that the strategy for SUSY searches in ATLAS will be of the following form:

1. Use inclusive signatures to observe an excess over SM background, and thus formally discover SUSY.

2. Select SUSY-like events, and search for exclusive processes with which to measure sparticle properties.

3. Use these properties to infer details of the SUSY Lagrangian.

We have already seen, however, that, in the R-parity conserving case, exclusive signatures alone are not guaranteed to give precise measurements of SUSY masses; the presence of an invisible $\tilde{\chi}_1^0$ at the end of each decay chain forces one to use techniques like kinematic endpoints that are sensitive to mass differences rather than absolute masses. Furthermore, one may observe more than one endpoint in a given distribution, and there is then a struggle to assign the endpoints to the sparticle decay chains that produced them. Encountering these problems may even be considered a luxury- one could just as easily find oneself in a region of parameter space where there are heavy scalars and thus cascade decays of the form shown previously will not be observed, or in a region where very small mass differences in the chain lead to soft leptons that cannot be detected.

For this reason, this chapter is devoted to developing a method to combine inclusive and exclusive measurements to directly constrain sparticle masses and SUSY parameters. This has rarely been attempted before due to the difficulty of calculating the predicted values of inclusive signatures at many points in parameter space. For example, whilst kinematic endpoint positions are given by simple analytic functions of sparticle masses (such as those given in Equations 5.4 to 5.8), the calculation of a cross-section involves running a Monte Carlo generator such as `HERWIG` in order to obtain an event sample, followed by putting the sample through a suitable detector simulation, all of which takes time. Thus, investigating the variation of inclusive variables over a parameter space is non-trivial, and it is therefore hard to find the regions of parameter space that are consistent with an inclusive measurement.[1]

The existence of a general method to combine exclusive and inclusive data would allow one to simultaneously use all data from the ATLAS detector to constrain SUSY models[2]. If endpoints are not visible, one can instead use other information that is available to try and constrain the parameter space.

The method proposed in this chapter uses the Markov Chain Monte Carlo sampling technique introduced in the previous chapter along with a modification of the FORTRAN version of `HERWIG`. Throughout, the major focus is on the problems encountered in the previous section, namely:

1. The problem of setting the overall mass scale.

2. The problem of decay chain ambiguity when observing kinematic endpoints.

3. The problem of working in restricted model spaces.

In addition, the effect of an uncertainty on the jet energy scale at ATLAS will be examined in detail, as an example of how to model systematic contributions to the experimental error.

Throughout the following chapter, it is necessary to assume a given SUSY model in order to demonstrate the use of the analysis technique. I will make the arbitrary choice of examining the co-annihilation model presented in the previous chapter, given that we already have some idea of the measurements that can be obtained in this model. It is important to realise, however,

---

[1]For an endpoint calculation, it is sufficient to know the mass spectrum of a given point in parameter space- this is obtained by solving the renormalisation group equations with the high scale parameters as input, and is a much quicker process.

[2]In fact, one is not limited to using ATLAS data, but could include any other observables from relevant experiments.

that any model can be used with the following analysis, and, indeed, the next chapter will consider a very different model.

6.1 Combination of signatures in mSUGRA

6.1.1 General remarks

In order to demonstrate how to combine inclusive and exclusive measurements at the LHC, I will combine the endpoint information obtained for the co-annihilation model in the previous chapter with one example of a inclusive signature. It should be possible for the ATLAS detector to measure the cross-section of events with missing p_T greater than 500 GeV. Unlike the endpoint information, this is sensitive to the overall mass scale of the theory; as the masses of sparticles increase, the missing p_T associated with the escaping neutralinos will increase, but the total production cross-section will decrease and hence the high mass solutions encountered in the previous section will lead to missing p_T cross-sections that are lower than the value obtained at the co-annihilation point.

Exclusive and inclusive variables are combined by adding the cross-section information to the definition of the probability function for the Markov Chain sampler that was defined previously, though there is a subtlety here. It should be noted that up to now the analysis has been model independent but, from here on in, some model will have to be assumed in order to draw conclusions from the measurements. This is because endpoint data can be analysed purely in the mass space S_{mass} (hereafter "M") defined by the weak scale masses, but inclusive measurements must be compared to a given scenario (through the use of a suitable Monte Carlo generator) before any conclusions can be drawn, and therefore must be analysed in the space of parameters, S_{model} of that model. This section investigates the constraints imposed by a cross-section measurement on the parameter space S_{mSUGRA} (hereafter "P") of a particular model, mSUGRA, in order to introduce the technique in a familiar context. The limitations of this approach will rapidly become apparent and will be tackled in Section 6.2.

In view of this change of the constrained-space, (from the space of weak-scale masses $\mathbf{m} \in M$ to the space of mSUGRA models $\mathbf{p} \in P$) the description of the Metropolis algorithm in Section 5.2.7 must, in this section, be considered re-written in terms of $p(\mathbf{p}|\mathbf{e}^{obs})$ rather than $p(\mathbf{m}|\mathbf{e}^{obs})$. This is made more explicit in Section 6.2.2 when a *further* enlargement of the constrained-space is made to accommodate uncertainty in the absolute jet energy scale.
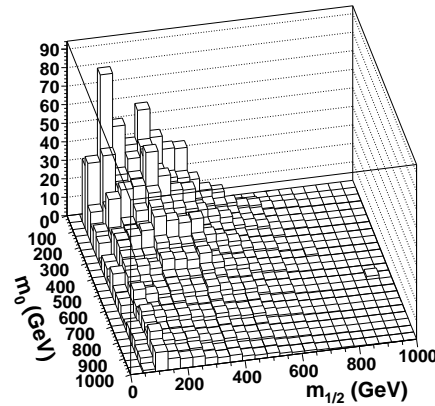
Figure 6.1: Variation of the cross-section of events passing a missing p_T cut of 500 GeV in the $m_0 - m_{1/2}$ plane of mSUGRA

The cross-section in picobarns for events passing a missing p_T cut of 500 GeV, for tan$\beta = 10$ and positive μ, obtained using `HERWIG`. The value at the co-annihilation point is 2.03 pb.

6.1.2   Implementation

The starting point for this study is that the cross-section of events with missing p_T greater than 500 GeV can be measured at ATLAS. One can then pick points in the mSUGRA parameter space S_{mSUGRA}, work out the mass spectrum, generate Monte Carlo events and work out the cross-section of events passing this cut. Only certain points in the parameter space are consistent with this measurement, and these will give a range of masses that are consistent. Naively, the overlap of this region of the mass space with the region consistent with the edge data will give the new region of mass space that is compatible with the ATLAS data. In fact, since the end points are not entirely independent of the cross-section measurement, one needs to include both sets of information in the final analysis in order to obtain the required region. The 'overlap' picture is approximately true, however, given that the measurement of the cross-section is not strongly related to the measurements of the edge positions, and is a useful guide for preliminary investigation before the final analysis is carried out.

A plot of the missing p_T cross-section in the m_0, $m_{1/2}$ plane for fixed tanβ and positive μ is shown in Figure 6.1. As can be seen, there is a strong variation over the parameter space and a measurement of even modest precision will be able to rule out large areas.

The full process of picking mSUGRA points and obtaining the cross-section of events that pass the missing p_T cut has been accomplished by successively running `ISAJET`, `HERWIG` and `ATLFAST`, with 1000 events being generated at each point. This is rather time consuming,

however, and a simple scan of the mSUGRA parameter space is unfeasible if any more than two of the mSUGRA parameters are varied. For this reason, the Metropolis sampling technique introduced in the previous section is utilised once more and, indeed, it is here that the power of the method becomes apparent. The algorithm has been used to map the interesting region of the parameter space with fewer points than would be required in a scan in order to obtain similar performance.

To demonstrate this, consider the following. There are four and a half parameters in the mSUGRA parameter space, though A_0 has been held constant for simplicity.[3] Of the remaining parameters, one is simply a sign (the sign of μ), and hence one sampling run was performed with this sign positive, and another with it negative. In any one application of the software, then, three parameters are varied – m_0, $m_{1/2}$ and $\tan\beta$ – and even a coarse scan requiring 100 points along each axis would require one million points for each sign of μ. The Metropolis algorithm maps the interesting regions of the space in approximately 15,000 points per sign of μ, a dramatic improvement without which the analysis would have taken many months, if not years.

Even with this improvement, it was still necessary to reduce the run time of `HERWIG` significantly through the use of a parallel computer (see appendix C). Although the Metropolis algorithm itself cannot be parallelised, Herwig can be adapted to run on a parallel machine with the use of Message Passing Interface (MPI) code[71], thereby substantially reducing the run time per point. This required a substantial effort, and is one of the main components of making the problem of calculating inclusive variables at many points in parameter space tractable.

Definition of Metropolis quantities for cross-section

The Metropolis algorithm for use with (only) the cross-section data is defined as follows. As in the previous section, one requires the definition of the probability distribution $p(\mathbf{p}|\sigma^{obs})$ from which samples are to be taken, in which σ^{obs} represents the cross section measurement. This necessitates the definition of a suitable prior $p(\mathbf{p})$ on the model space P which again is taken to be flat (but equal to zero for invalid values of any of the model parameters $p_i \in \mathbf{p}$). Finally the Metropolis sampler's proposal distribution must be modified to act on the model space P rather than on the mass space M. The proposal distribution was again chosen to be a multi-dimensional Gaussian centred on the current point $\mathbf{p} \in P$. The widths of the proposal

---

[3]In retrospect there was no compelling reason to hold A_0 fixed, and a later study is expected to look at the effect of allowing A_0 to vary and be measured by the data along with all the other model parameters.

distribution in m_0, $m_{1/2}$ and $\tan\beta$ were respectively usually 25 GeV, 25 GeV and 2, except when both cross-section *and* edge constraints were in use simultaneously (only in Sections 6.1.3 and beyond) in which case a smaller set of widths was used (5 GeV, 5 GeV and 2). The widths were obtained through trial and error and will not affect the results once convergence has occurred.

The sampled probability distribution $p(\mathbf{p}|\sigma^{obs})$ follows a similar definition to that encountered previously for $p(\mathbf{m}|\mathbf{e}^{obs})$. The analogue of Equation 5.10 is then just the single term $p(\sigma^{obs}|\mathbf{p})$ quantifying the cross-section likelihood according to:

$$p(\sigma^{obs}|\mathbf{p}) \approx \frac{1}{\sqrt{2\pi\sigma_{err}^2}} \exp\left(-\frac{(\sigma^{obs} - \sigma^{pred}(\mathbf{p}))^2}{2\sigma_{err}^2}\right), \tag{6.1}$$

where σ_{err} is the error associated with the cross-section measurement and $\sigma^{pred}(\mathbf{p})$ is the value of the cross section expected at the point $\mathbf{p}$ in mSUGRA parameter space P. For the 'measured' cross-section, the value obtained from `HERWIG` using the correct mSUGRA parameters for our mass point was taken, with an assumed error σ_{err} of ten per cent which reflects the likely experimental uncertainty[4].

Certain regions of mSUGRA parameter space P are known to be unphysical – for example there may be no electroweak symmetry breaking or there may be a charged LSP. In both cases, `ISAJET` will detect this and fail to run. Furthermore there are points $\mathbf{p}$ for which `HERWIG` will not run. When any of these problems occur the point $\mathbf{p}$ is taken to be unphysical and the likelihood is multiplied by zero (as unphysical points *cannot* have generated the observed data!).

Results in mSUGRA space (for cross-section information alone)

The results of the Markov Chain in mass space for positive μ can be seen in Figure 6.2, with those for negative μ presented in Figure 6.3. The distributions look very similar in the m_0, $m_{1/2}$

---

[4]After the simulations had been completed, it was pointed out that the choice of σ_{err} almost certainly underestimates the theoretical uncertainty which will increase the error to something like thirty per cent. This will affect all plots featuring nothing but cross-section information, in which the consistent region is certain to expand. This expansion cannot readily be quantified without further simulation, and will have to be the subject of future studies. I do not expect the enlargement to give a significant effect in cases where the parameter space is over-constrained through a combination of cross-section measurement with endpoint data, but would expect the effect to be significant when exploring spaces with fewer constraints; i.e. those with larger numbers of parameters(Figures 6.16 and 6.17), or those without endpoint information (Figures 6.2 and 6.3). On the other hand, ATLAS is now much more optimistic about cross-section measurements, which would tend to reduce σ_{err} by reducing the experimental error.
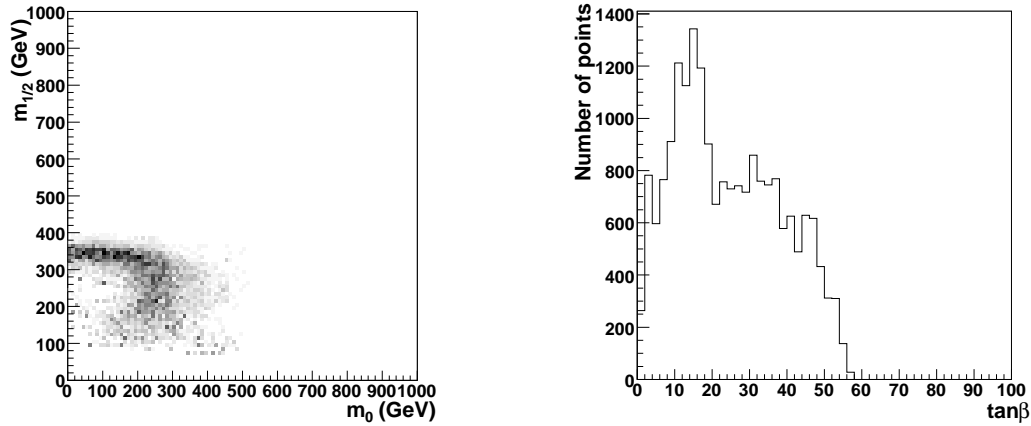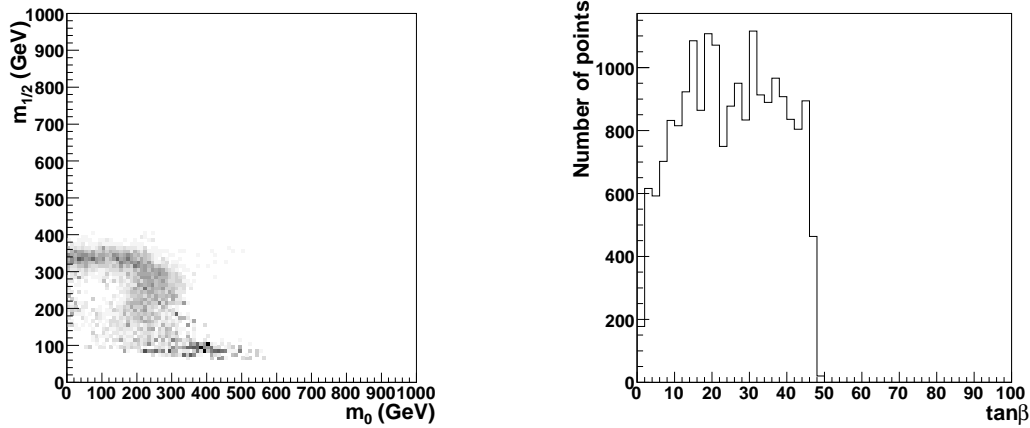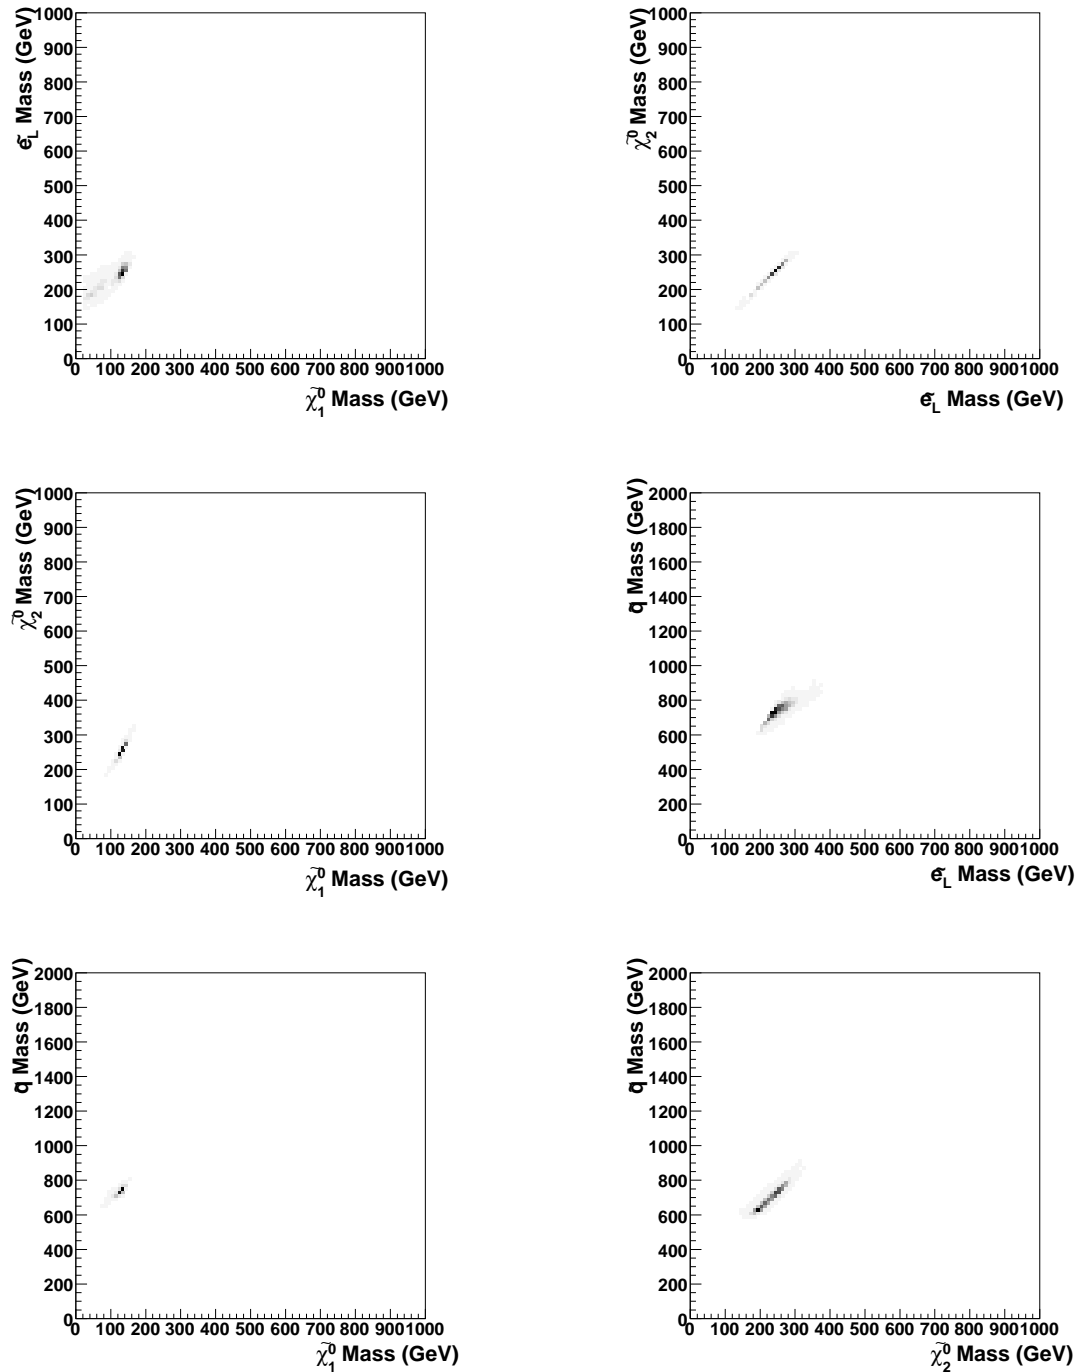
Figure 6.2: The region of mSUGRA parameter space consistent with the measurement of the cross-section of events with missing p_T greater than 500 GeV, for positive μ.

plane, reflecting a lack of sensitivity to the sign of μ. The tanβ distribution is approximately flat for negative μ, whilst there is some insignificant preference for the 'correct' value of tanβ in the positive μ case.

Results in mass space (for cross-section information alone)

It is now instructive to relate the results in Figures 6.2 and 6.3 to the weak scale mass space in which we have already observed the regions consistent with the kinematic edge analysis. The positive μ and negative μ data sets presented in Figures 6.2 and 6.3 have been evolved to the weak scale using `ISAJET` and combined into a single data set by weighting each of the two sets by the average likelihood of that set. This gives a region in mass space which is shown in Figure 6.4, and is dramatically different from that obtained using the edge analysis (shown in Figure 5.6). The overlap between the mass space region obtained from the cross-section information (Figure 6.4) and the mass space region obtained using the edge data (Figure 5.6) is shown in Figure 6.5, and was found by multiplying the previous data sets.

The overlap of the cross-section region with the edge data region has produced much tighter constraints. It is worth noting that the region in Figure 6.5 can be projected onto each pair of axes, and that this projection has a different size in each plane, with the smallest
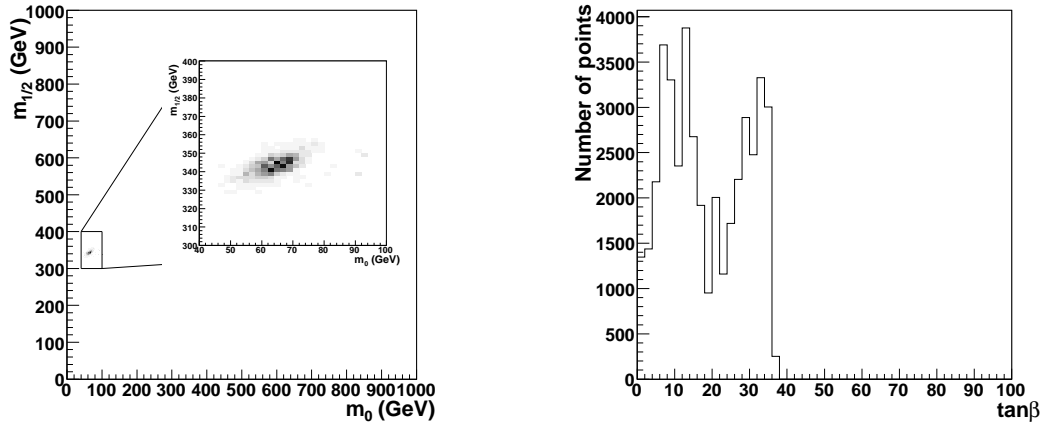
Figure 6.3: The region of mSUGRA parameter space consistent with the measurement of the cross-section of events with missing p_T greater than 500 GeV, for negative μ.

being that in the plane of the neutralino masses. This could be used to remove some of the area shown in the other planes, although the strictly correct procedure (followed in Section 6.1.3) is to run a Markov Chain with the edge and cross-section information implemented at the same time.

6.1.3   Further analysis

The overlap plots presented in the previous subsection give a rough idea of what to expect from the combination of edge and cross-section information, but the approach is only approximately valid given that the cross-section measurement is not independent of the kinematic edge positions. Both measurements are dependent on the particle masses and, although the endpoints are more sensitive to mass differences rather than absolute masses there is still some weak correlation between the cross-section and the endpoints. In order to be fully rigorous, one must run a Markov Chain whose probability density function combines both the cross-section and the edge information at the same time – in other words one must sample this time from $p(\mathbf{p}|\mathbf{e}^{obs}, \sigma^{obs})$.

Accordingly, a Metropolis sampler of $p(\mathbf{p}|\mathbf{e}^{obs}, \sigma^{obs})$ was set to explore the mSUGRA parameter space P.
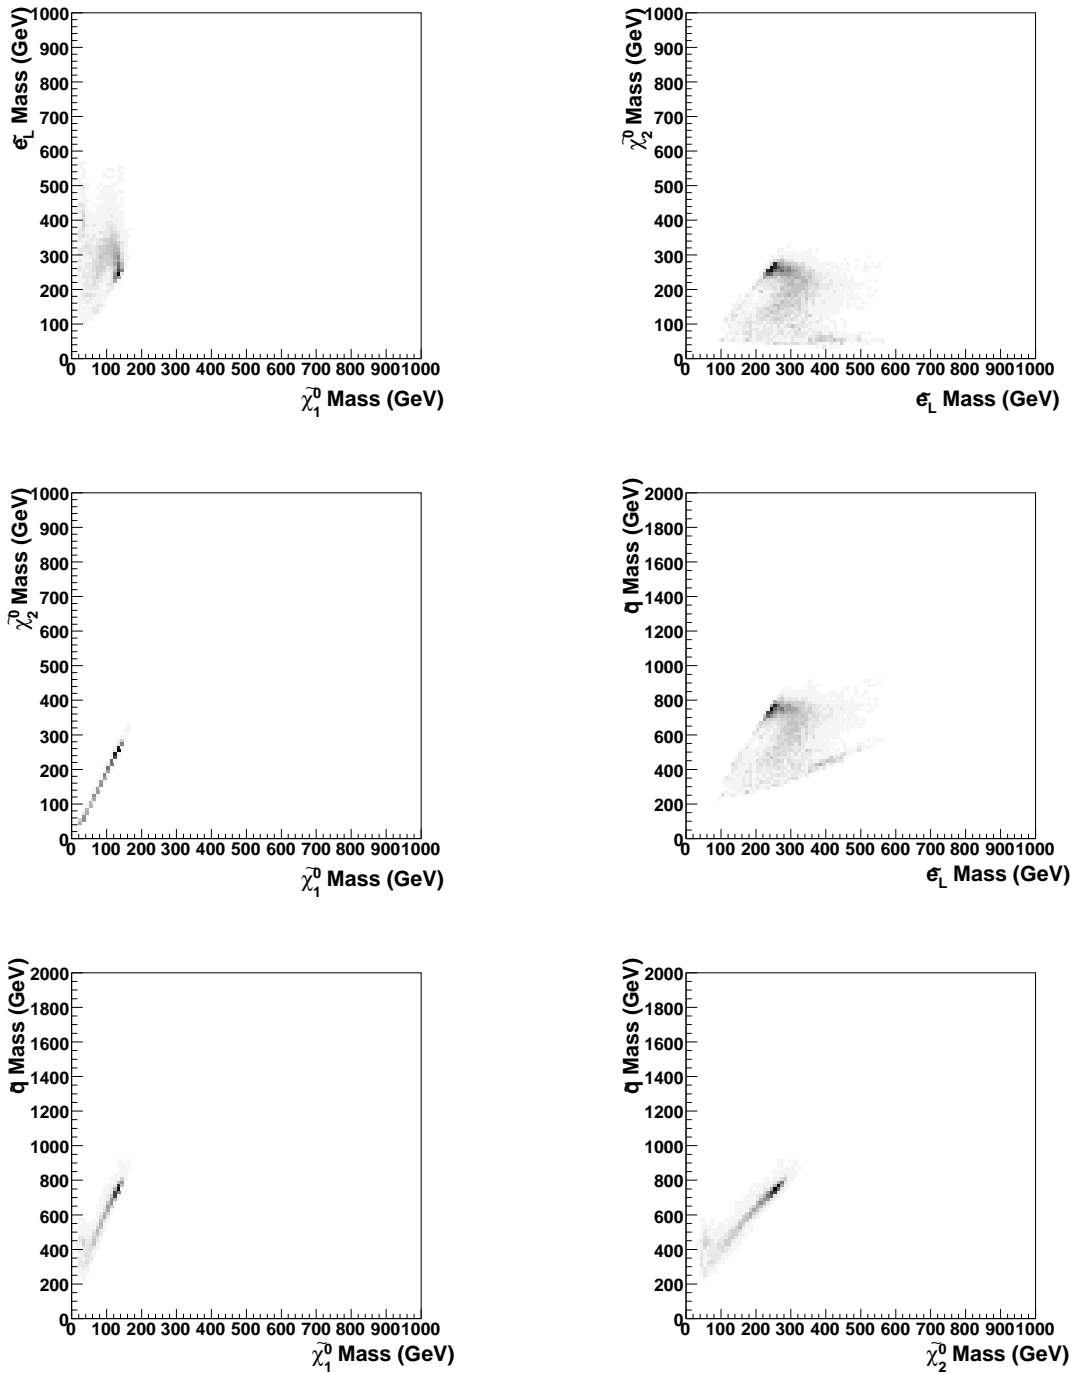
Figure 6.4: The region of mass space consistent with a measurement at 10% precision of the cross-section of events with missing p_T greater than 500 GeV, obtained using a Markov chain sampler.

Figure 6.5: The region of mass space consistent with a measurement at 10% precision of the cross-section of events with missing p_T greater than 500 GeV, overlapped with a measurement of the squark decay kinematic endpoints obtained in Section 5.2.3.

At each point $\mathbf{p} \in P$ the number of events passing the missing p_T cut was obtained from the `ATLFAST` output whilst the `ISAJET` mass spectrum was used to find the expected position of the endpoints. This information was then compared to the 'measured' information (in this case, the endpoints shown earlier, and the cross-section obtained through Monte Carlo simulation of the co-annihilation mass point) in the definition of the probability weight for each point $\mathbf{p} \in P$. The likelihood $p(\mathbf{e}^{obs}, \sigma^{obs}|\mathbf{p})$, the analogue of Equations (5.10) and (6.1), is this time the product of the pair of them:

$$p(\mathbf{e}^{obs}, \sigma^{obs}|\mathbf{p}) = p(\sigma^{obs}|\mathbf{p}) \prod_{i=1}^{5} p(e_i^{obs}|\mathbf{m}(\mathbf{p})). \tag{6.2}$$

The same flat prior $p(\mathbf{p})$ on mSUGRA space P was used as in Section 6.1.2. The likelihood was multiplied by zero if the sparticle masses $\mathbf{m}(\mathbf{p})$ obtained at a point $\mathbf{p}$ were not consistent with the mass hierarchy required for the squark decay chain to exist. The Metropolis algorithm's proposal distribution was the same as that used previously in Section 6.1.2. Chains were run separately for positive and negative μ.

Results for cross-section *and* edge measurements together

The mSUGRA space results for cross-section *and* edge measurements are shown in Figures 6.6 and 6.7, with the results in mass space shown in Figure 6.8. Note that the precision in the $m_0, m_{1/2}$ plane is greatly improved, and this is to be expected given that many points in mSUGRA space are now rejected as they give the wrong mass hierarchy for the decay chain. This leads to superb precision when the results are viewed in the weak scale mass space, a fact that can be explained by realising that the dimensionality of our original parameter space has been reduced. In Sections 5.2.3 and 5.2.7 the end point data were analysed in the weak scale mass space (4 dimensions), and the cross-section data were considered in the mSUGRA space. By shifting the end point data to the mSUGRA space, the fit now occurs in a space where there are only really two parameters that control the Markov chain (as we are mostly sensitive only to m_0 and $m_{1/2}$), and this improves the precision *even before the cross-section measurement is added*. However, the fact that there is an excellent level of constraint immediately implies that we are no longer restricted to looking at the mSUGRA model, and in fact illustrates the power of the technique.

Hence it is observed that we have obtained solutions to the three problems posed at the start of the chapter. The cross-section measurement is sensitive to the mass scale, and thus improves
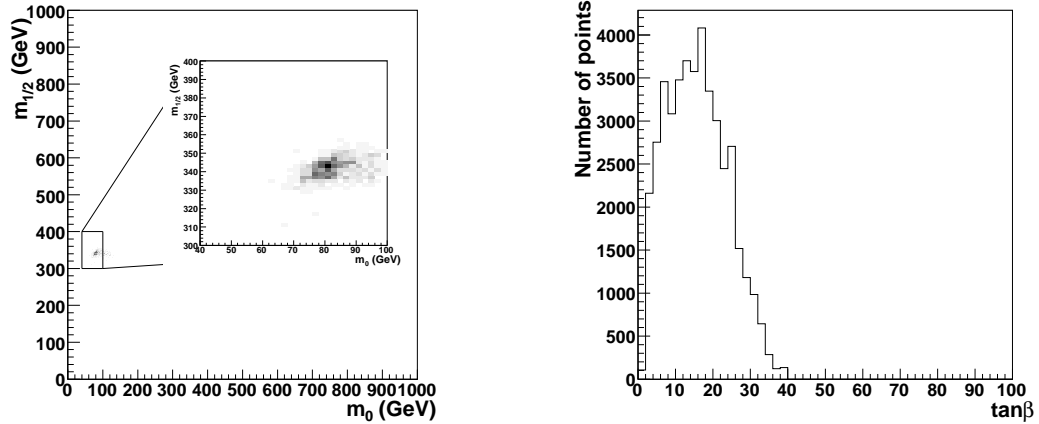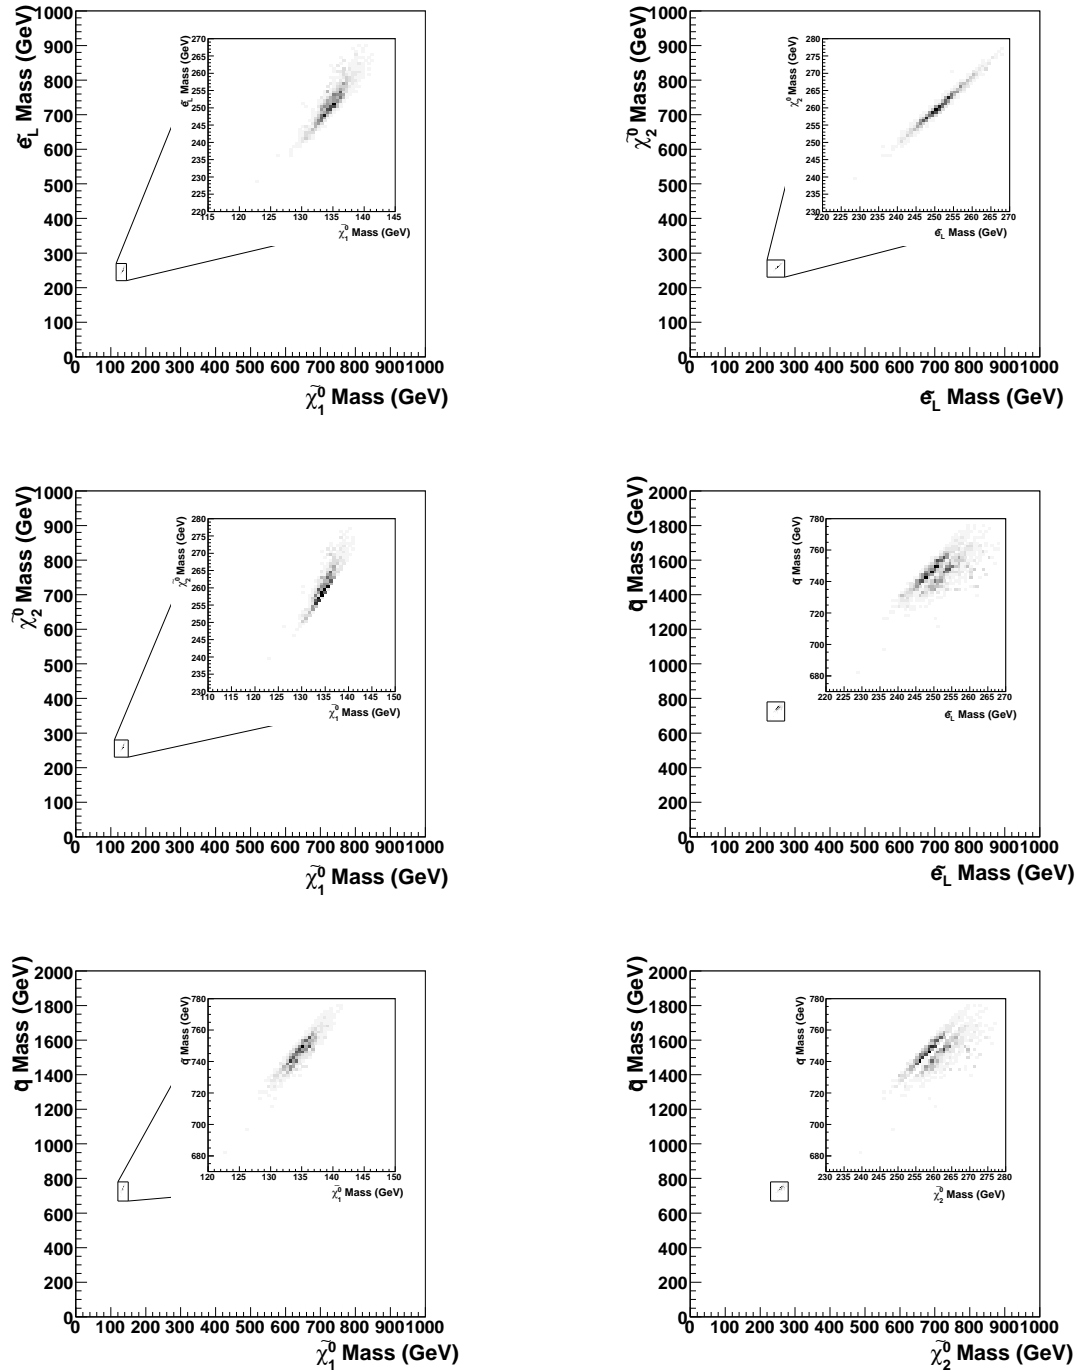
Figure 6.6: The region of mSUGRA parameter space consistent with the measurement of the cross-section of events with missing p_T greater than 500 GeV and with the endpoint measurements obtained in Section 5.2.3, for positive μ.

mass measurements. Furthermore, there is enough constraint on the system to start looking beyond mSUGRA, and to investigate the effects of ambiguity in the determination of the decay chain. This ambiguity has only rarely been considered before (see for example[66, 67]), and is investigated in Section 6.2.3.

6.2   Generalising mSUGRA

6.2.1   Background

It has been shown so far that one can sample from the mSUGRA parameter space using both kinematic endpoint data and a simple cross-section measurement. Endpoint data alone gives more than adequate precision within the framework of mSUGRA, provided one samples the mSUGRA parameter space and assumes that one has identified the particles in the decay chain correctly. The aim of this section is to generalise this analysis to include both ambiguity in the decay chain and more general assumptions about the mechanism of SUSY breaking. The effect of a jet energy scale error on the endpoint positions is also investigated, thus demonstrating how one would include correlated experimental effects in the analysis.

Figure 6.7: The region of mSUGRA parameter space consistent with the measurement of the cross-section of events with missing p_T greater than 500 GeV and with the endpoint measurements obtained in Section 5.2.3, for negative μ.

6.2.2 Effect of a jet energy scale error

As we saw in Section 2.2.5, ATLAS does not measure the energy of jets perfectly, but instead has some energy scale error. Given that most of the endpoints feature a quark jet, it is worth investigating the effect of the energy scale error on the positions of the endpoints, and the subsequent effect on the precision obtained in the mSUGRA parameter space.

Firstly, it is noted that for jets whose energy exceeds 70 GeV (the likely energy of the jet contributing to the endpoint plots given the relatively large mass difference between the squarks and the neutralinos), the energy scale error is expected to be of the order of 1 per cent[25]. This is much lower than the errors attributed to the endpoints that arise from mismeasurement, and hence the effect will not cause a discernible difference to the results. Nevertheless the following analysis serves as a practical example of how one can incorporate experimental effects into the analysis.

To accommodate the effect of an unknown shift s in the absolute jet energy scale, a parameter s is added to the set explored by the sampler. In other words, the sampler now wanders around the extended space $Q = P \otimes S$ defined as the product of the mSUGRA parameter space $P = \{\mathbf{p}\}$ with the set S of possible values of s. At each point $\mathbf{q} = (\mathbf{p}, s) \in Q$ the masses $\mathbf{m}(\mathbf{p})$ of the particles in the decay chain are worked out. A calculation of the "idealised" positions of

Figure 6.8: The region of mass space consistent with a measurement at 10% precision of the cross-section of events with missing p_T greater than 500 GeV combined with the endpoints measured in Section 5.2.3, obtained using a Markov chain sampler in mSUGRA space. The two bands arise from the fact that the data are a weighted average of the two initial data sets taken with positive and negative μ. The data demonstrate that there is a negative μ solution that is not substantially less likely than the positive μ solution, and in which the masses are slightly different.

the edges corresponding to these masses is performed (as before) but these are then *moved* by the amount predicted by the current hypothesis s for the absolute jet energy scale correction. The resulting modified edge positions $\mathbf{e}^{pred} = \mathbf{e}^{pred}(\mathbf{q}) = \mathbf{e}^{pred}(\mathbf{m}(\mathbf{p}), s)$, which now depend on s, are the values which are used in the new version of Equation (5.11).

Having extended P to the larger space Q, our goal is now to sample not from $p(\mathbf{p}|\mathbf{e}^{obs})$ but from $p(\mathbf{q}|\mathbf{e}^{obs})$. The latter is proportional to $p(\mathbf{e}^{obs}|\mathbf{q})p(\mathbf{q})$. The first term $p(\mathbf{e}^{obs}|\mathbf{q})$ may be calculated almost exactly as before in Equation (5.11) but with the new modified edge positions $\mathbf{e}^{pred}(\mathbf{m}(\mathbf{p}), s)$ described above. The last term $p(\mathbf{q})$ may be decomposed by independence into two parts: $p(\mathbf{p})p(s)$. The first of these, $p(\mathbf{p})$, is the mSUGRA-space prior which has already been seen,[5] while the other, $p(s)$, is the expected distribution of the final uncertainty in the absolute jet energy scale. Following[25], $p(s)$ is taken to be a Gaussian of width 1%.

In order to determine the particular amounts δ_i by which the i^{th} endpoint should be shifted for a given jet energy scale correction factor s, a toy Monte Carlo simulation is run at that point and for that edge.[6] This is done once with and once without the correction factor s multiplying the jet energies. The positions of the endpoints are compared in the two cases. Different endpoints are shifted by different fractions of the energy scale error s.

The results including uncertainty in the jet energy scale are shown in Figures 6.9 and 6.10 for positive and negative μ respectively and are comparable to those obtained previously (Figures 6.6 and 6.7) when uncertainty in the jet energy scale was not considered.

6.2.3   Chain ambiguity in mSUGRA

In order to investigate the effect of chain ambiguity on the mSUGRA parameter space, the edge data from Section 5.2.3 are here used in an mSUGRA fit *without* the assumption that the particles in the decay chain have been identified correctly. There are relatively few processes that can give the characteristic endpoints associated with the squark cascade decay, and it should be sufficient merely to extend the decay given in Equation 5.3 to include the possibility that any two neutralinos may be produced in the decay (provided of course that the one further down the chain is lighter than that higher up the chain) and that the slepton can be either left or right-handed. This gives twelve possible mass hierarchies, and each of these gives a series

---

[5]It must be remembered that, as in earlier sections, the likelihood $p(\mathbf{e}^{obs}|\mathbf{q})$ will be zero (given the model) at points where the masses of the particles in the chain do not obey the necessary mass hierarchy. It is computationally easier to place this veto into the prior $p(\mathbf{p})$ as before.

[6]Strictly speaking the toy Monte Carlo simulation is only needed for the *llq* edge and the *llq* threshold as the shifts in the edge positions for the *other* edges are linear in $\sqrt{s}$ and may be calculated analytically.
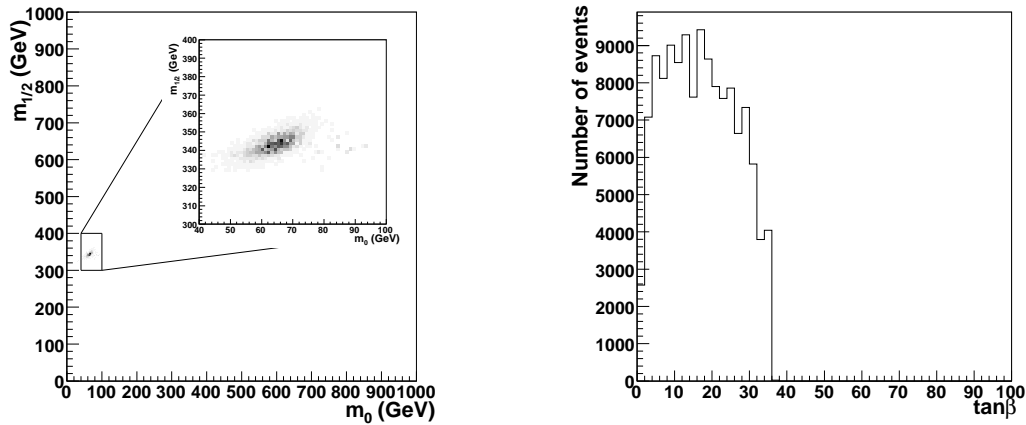
Figure 6.9: The region of mSUGRA parameter space consistent with the endpoint measurements obtained in Section 5.2.3, for positive μ, with a 1 per cent jet energy scale error included. The sampler was run using endpoint data only.
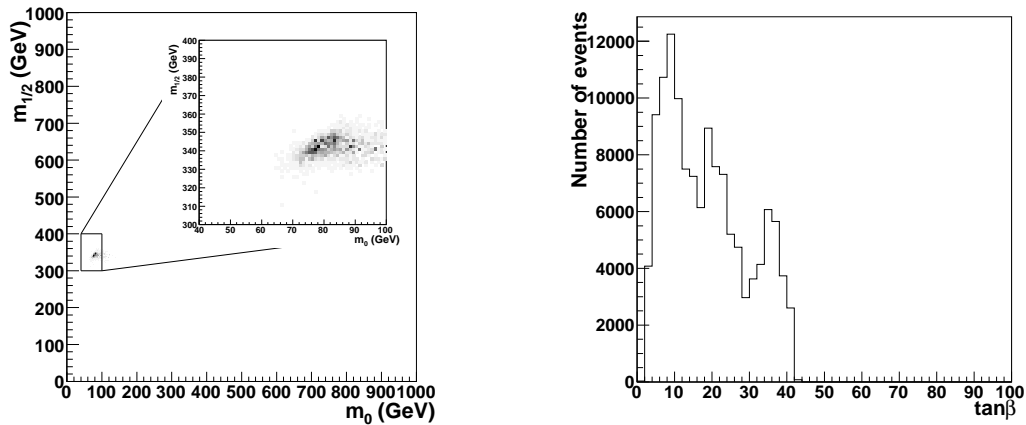
Figure 6.10: The region of mSUGRA parameter space consistent with the endpoint measurements obtained in Section 5.2.3, for negative μ, with a 1 per cent jet energy scale error included. The sampler was run using endpoint data only.

of possible endpoints in the mass spectra. This assumes that the only ambiguity arises from two body processes- in fact it may be possible for processes involving a three body decay to mimic the observed signal. This is considered in Chapter 7.

If we label the N_a different mass assignments with a tag a_i, the likelihood for the i-th observed edge at each point $\mathbf{p}$ in the mSUGRA parameter space P now becomes:

$$
\begin{aligned}
p(e_i^{obs}|\mathbf{p}) &= \sum_{j=1}^{N_a} p(e_i^{obs}|\mathbf{p}, a_j)p(a_j) \\
&= \sum_{j=1}^{N_a} p(e_i^{obs}|\mathbf{m}_{a_j}(\mathbf{p}))p(a_j)
\end{aligned}
\tag{6.3}
$$

where $p(a_i)$ is the prior for the mass assignments, and N_a gives the number of assignments open at that point in parameter space. If we assume that each of the assignments is equally likely, the prior $p(a_i)$ is simply $1/N_a$. The term $p(e_i^{obs}|\mathbf{m}_{a_i}(\mathbf{p}))$ is calculated using Equation (5.11) with the masses corresponding to the assignment a_i.

Equation (6.3) makes the conservative assumption that any observed edge could have come from any observed chain (i.e. not necessarily from the same chain as that generating a different observed edge). Furthermore (but less realistically) it assumes that there is no correlation between the chains generating each of the edges, whereas in many parts of parameter space it is highly likely that there is only one dominant chain. It is thus arguable that Equation (6.3) should be replaced by the stronger statement

$$
\begin{aligned}
p(\mathbf{e}^{obs}|\mathbf{p}) &= \sum_{j=1}^{N_a} p(\mathbf{e}^{obs}|\mathbf{p}, a_j)p(a_j) \\
&= \sum_{j=1}^{N_a} p(\mathbf{e}^{obs}|\mathbf{m}_{a_j}(\mathbf{p}))p(a_j)
\end{aligned}
\tag{6.4}
$$

which says that *all* the observed edges were the result of the same (albeit unknown and unidentified) chain of sparticles. I choose to present results using (6.3) rather than (6.4).

The results for positive μ are seen in Figure 6.11, whilst those for negative μ are in Figure 6.12. The precision is worse than that encountered previously, but not by much. It appears that there are two favoured regions in each plot, rather than the single region encountered previously. Further investigation reveals that one of these features a decay chain with a left handed slepton whilst the other features a decay chain with a right handed slepton.

The next step is to view the regions in the weak scale mass space that correspond to the

chosen mSUGRA points, and here there is a problem. Since we are now assuming that we do not know exactly which particles are in the decay chain, one can no longer take the points in the mSUGRA plane and claim that they give the masses of the lightest two neutralinos and the left handed slepton. Instead, it is merely possible to say that we have measured a neutralino-like object and a slepton-like object, but that more facts are needed to make a stronger statement.

It is possible, however, to use some other information to learn more about the particles in the decay chain. For a start, one can look at the width of the distribution for each mass ($m_{\tilde{\chi}_1^0}$, $m_{\tilde{\chi}_2^0}$, etc) that results from the mSUGRA points and use these widths as a qualitative guide. If the endpoints are really caused by a single mass hierarchy, the masses in this chain should generally fit the data better than other hierarchies, and this will manifest itself in a smaller spread of masses for the masses involved in the correct hierarchy. In the present case, the endpoints should all be caused by a decay chain featuring the lightest two neutralinos and the left handed slepton, so one expects these masses to have narrower distributions. This is indeed the case for the neutralinos, as seen in Figure 6.13, though the selectron results are less different. Note that these should not be confused with mass measurements; we are merely trying to get a hold on which decay chains are most likely to fit with the results.

There are other things that can be done. Having selected a region of the mSUGRA parameter space, one can look within that region at the branching ratios for the different possible mass hierarchies, after which it might be found that there are not enough events of a given type to be consistent with the observed endpoints. Therefore, although a decay chain featuring a $\tilde{\chi}_3^0$ and $\tilde{\chi}_2^0$ may fit a given endpoint slightly better than the correct chain, it might be impossible for that chain to produce an endpoint with the same number of events present as has been observed. This, in conjunction with the width of the mass distributions, might be enough to confirm that the decay chain must be of the correct type, but it would be foolish to make this assumption without first considering all the options.

Given that the region in mSUGRA space has not substantially increased in size, the cross-section information will not be added at this stage. Instead, I will investigate the effect of relaxing some of the assumptions of the mSUGRA model.
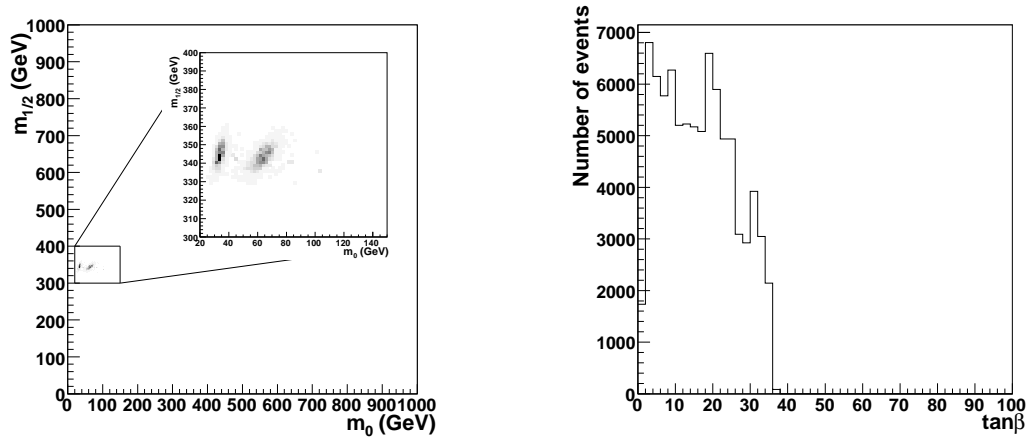
Figure 6.11: The region of mSUGRA parameter space consistent with the endpoint measurements of Section 5.2.3, without the assumption that the neutralinos and slepton in the squark decay chain have been correctly identified. For full details, see text. Results are shown for positive μ.
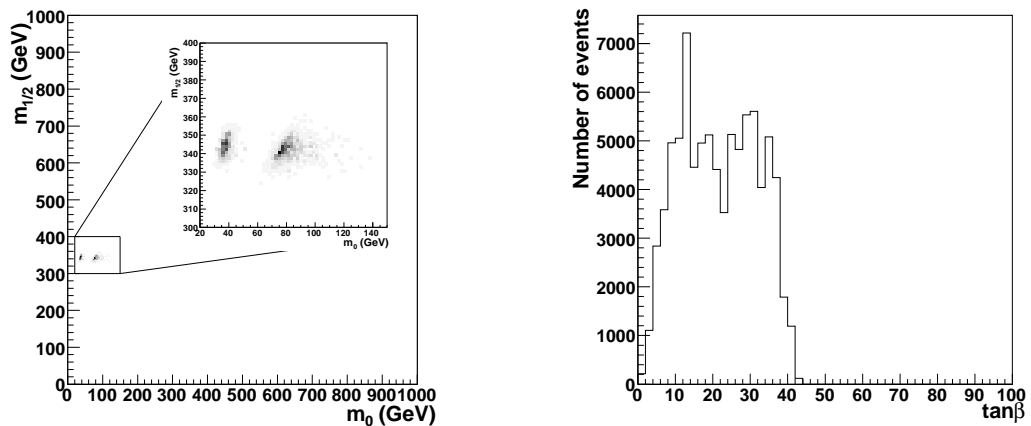
Figure 6.12: The region of mSUGRA parameter space consistent with the endpoint measurements of Section 5.2.3, without the assumption that the neutralinos and slepton in the squark decay chain have been correctly identified. For full details, see text. Results are shown for negative μ.
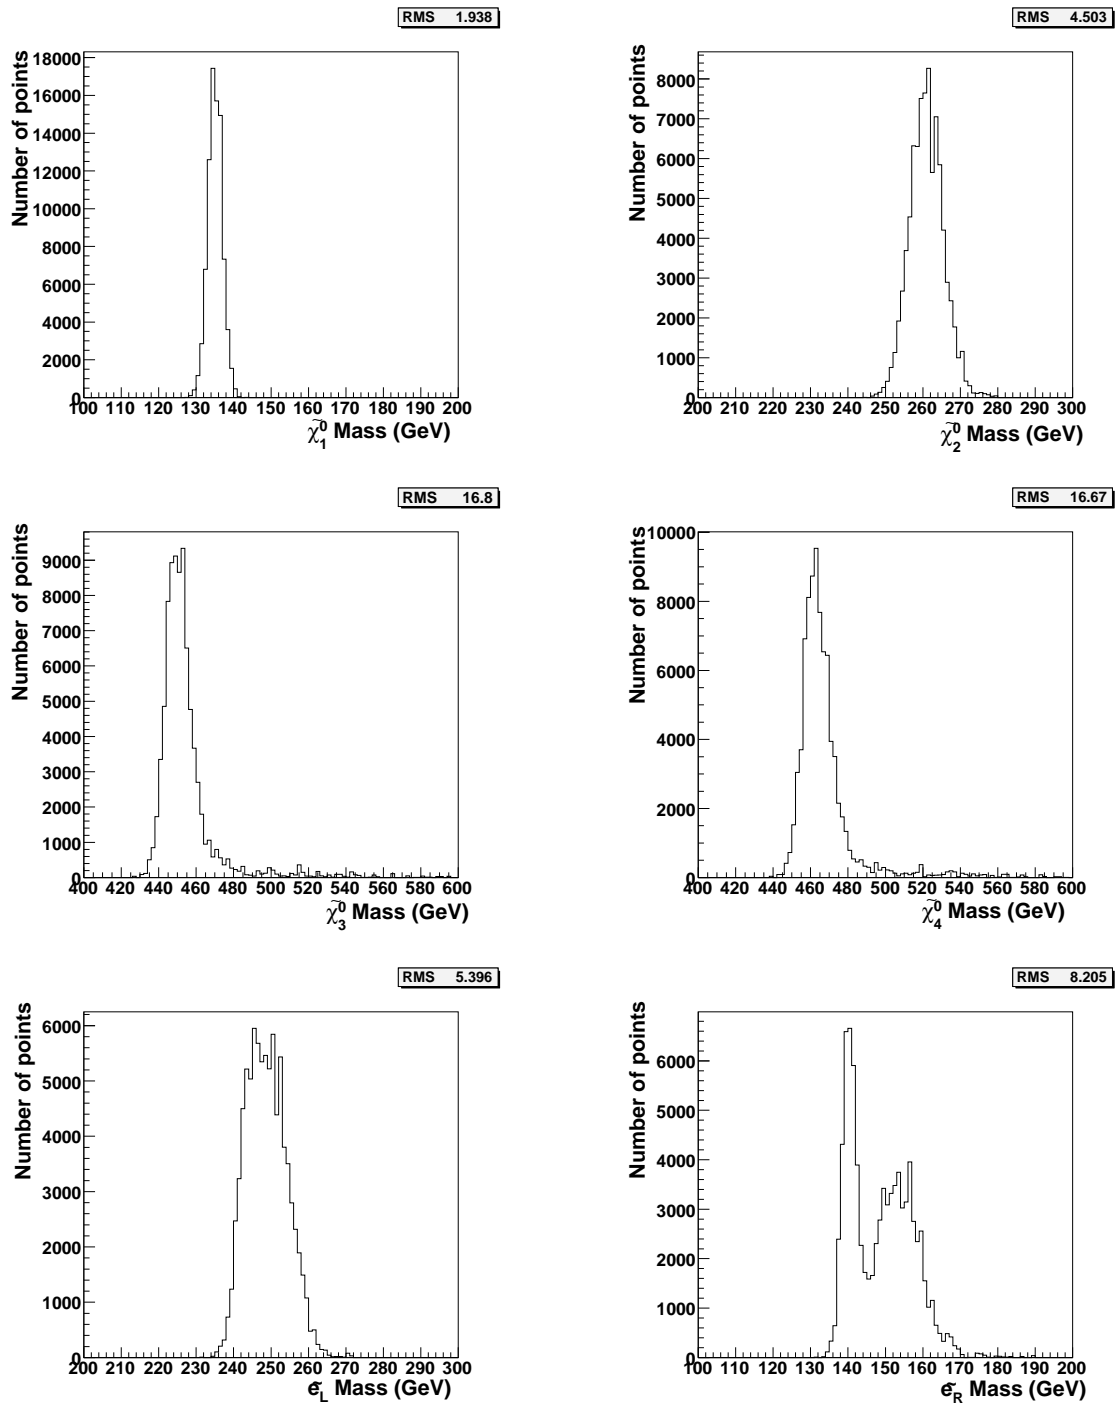
Figure 6.13: The mass distributions obtained from the mSUGRA decay chain, in which one has not assumed that the particles in the decay chain have been identified. These are not to be confused with mass measurements!

6.2.4   A non-universal SUGRA model

Background

The mSUGRA model assumes universality of the scalar and gaugino masses at the GUT scale, and also unifies the trilinear couplings at the GUT scale. Although this helps in reducing the SUSY breaking parameter set to a manageable level, reality may present a more complicated case. Hence, there is a very strong motivation for developing techniques that are either model independent or are at least able to tackle some more general SUSY models.

In this subsection, the effect of relaxing the assumption of universal GUT scale gaugino masses is examined, whilst still retaining the chain ambiguity and jet energy scale effects encountered in the Sections 6.2.2 and 6.2.3. It is important to realise that this is merely a first example of the use of the techniques developed here; one could just as easily relax more of the mSUGRA assumptions provided that one has made enough measurements to provide suitable constraints on the resulting model.

Kinematic edge constraints on non-universal SUGRA

The parameter set for the SUGRA model now becomes m_0, $\tan\beta$, A_0, $\text{sgn}(\mu)$, M_1, M_2 and M_3, where M_1, M_2 and M_3 are the GUT scale gaugino masses. A Metropolis sampler was used to sample from this parameter space (along with the jet energy scale error s), with the mass spectrum of each point found using `ISAJET 7.69`. Chain ambiguity was incorporated in the same way as described in Section 6.2.3. The results are seen in Figures 6.14 and 6.15: it should be noted that the previous m_0 vs $m_{1/2}$ plot has been superseded by three plots against the various GUT scale gaugino masses. The plots shown contain 800,000 points, after which the sampler was still clearly exploring new areas of the parameter space. Thus the plots are not final (i.e. they have not yet "converged") but they are sufficient to show that the endpoint data *alone* do not provide sufficient information to adequately constrain the non-universal SUGRA model, and so we have indeed reached a point where we need to consider additional measurements – such as the cross-section.

Kinematic edge data *and* cross-section constraints on non-universal SUGRA

A further Metropolis sampler was used to explore the parameter space of the non-universal SUGRA model using both the cross-section information *and* the edge data in the definition

of the probability weight for each point. The results for positive μ are seen in Figure 6.16, whilst those for negative μ are seen in Figure 6.17, and the difference from the plots described above is immediately apparent. The system is much more tightly constrained, and it has not wandered too far from the region corresponding to an mSUGRA model in which M_1, M_2 and M_3 are degenerate. One can convert this GUT scale region to a region in mass space as before (see Figure 6.18), though with the previous disclaimer that we have not yet identified which of the particles are involved in the decay chain but merely the range on the various masses that might be involved. Further work in the form of Monte Carlo studies targeted in the selected region at the GUT scale might possibly identify which masses are involved and hence improve the precision further, a study that is perfectly feasible given the relatively small extent of the region allowed by the data.

The results presented here are very encouraging, however, showing that even with only one extra observable one can afford to be more honest about the general lack of information regarding decay processes whilst still obtaining adequate precision within the framework of mSUGRA, and reasonable precision in a more general model.

6.3   Summary

Markov Chain sampling techniques have been used to combine kinematic endpoint measurements with a cross-section measurement in order to obtain precision SUSY mass measurements in simulated ATLAS data. Previous analyses have been extended to include ambiguity in the decay chain responsible for the endpoints, and a preliminary study has been made of a non-universal SUGRA model. Throughout it has been shown that the precision of mass measurements is greatly improved through the use of inclusive data, and the technique described offers a rigorous and general approach to the problem of constraining SUSY at the LHC. Reasonable precision has been obtained even with a fairly conservative estimate of the errors on the endpoints themselves.

The work described here is the first step toward what is hoped will be a powerful technique for future analysis. By collecting inclusive observables, one can start to look at more and more general models, with the final result limited only by the ability of physicists to come up with new pieces of information. At the very least, the Markov Chain approach is a powerful framework for combining information and exploring multi-dimensional parameter spaces in an efficient manner. The modification of `HERWIG` documented in appendix C is an essential part of this analysis, and is a powerful tool.
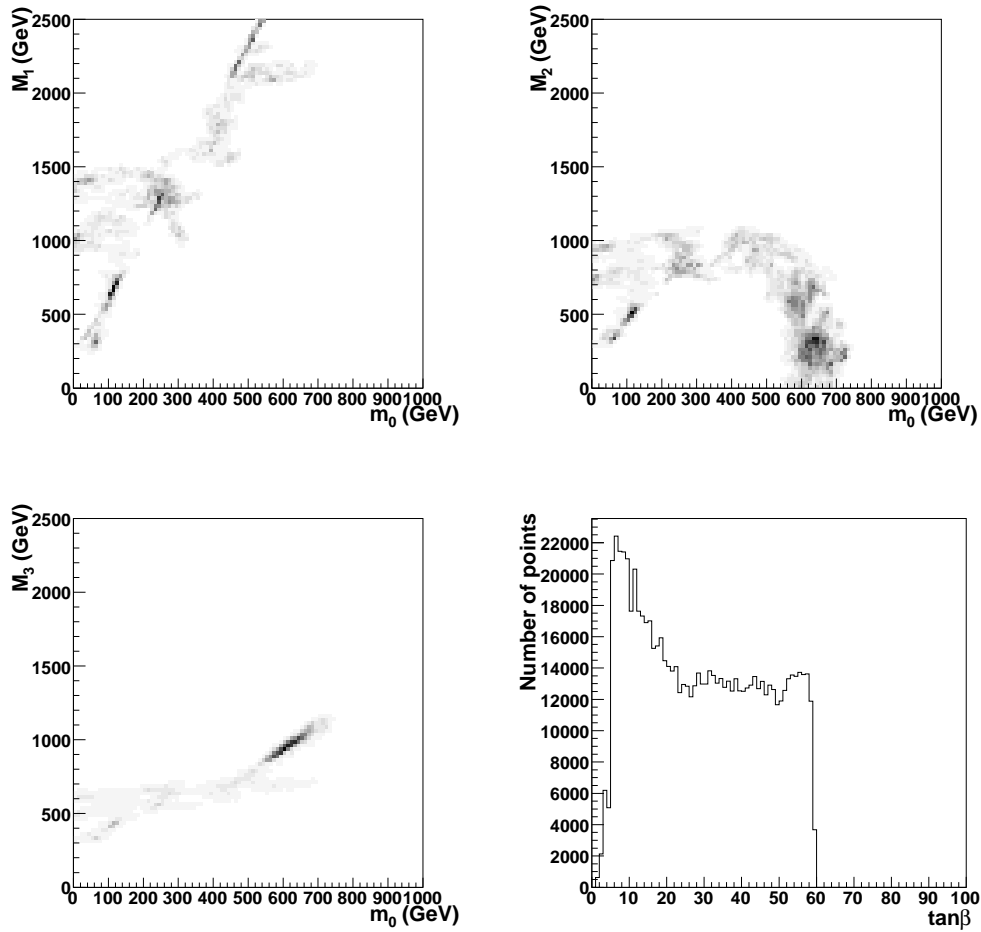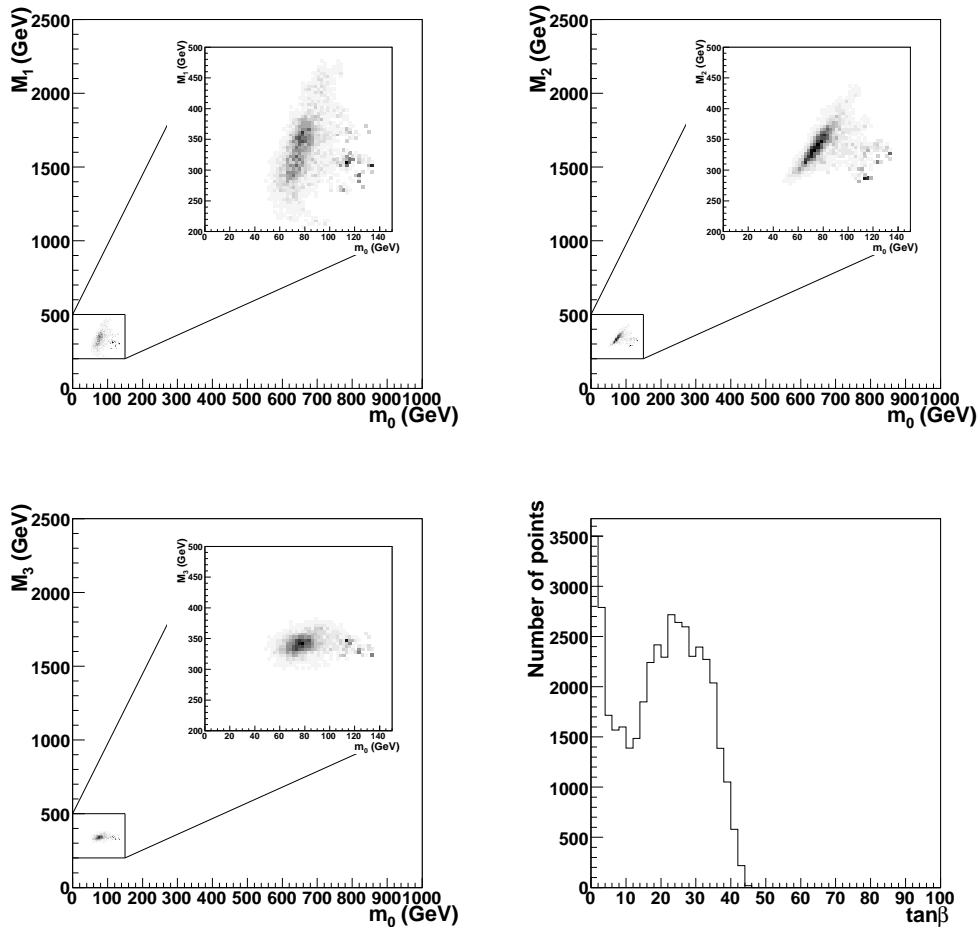
Figure 6.14: The region of the non-universal SUGRA parameter space consistent with the endpoint measurements of Section 5.2.3, with chain ambiguity included. Results are shown for positive μ. The Markov Chain has not converged.
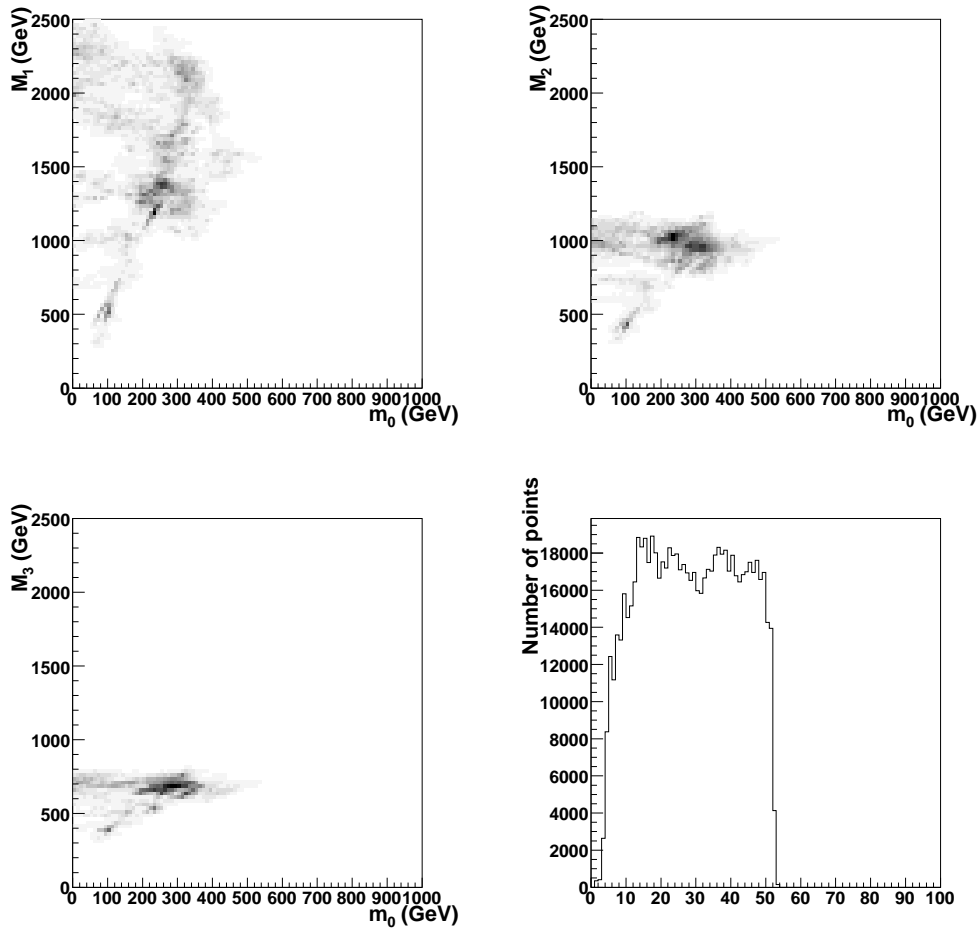
Figure 6.15: The region of the non-universal SUGRA parameter space consistent with the endpoint measurements of Section 5.2.3, with chain ambiguity included. Results are shown for negative μ, and the Markov Chain has not converged.
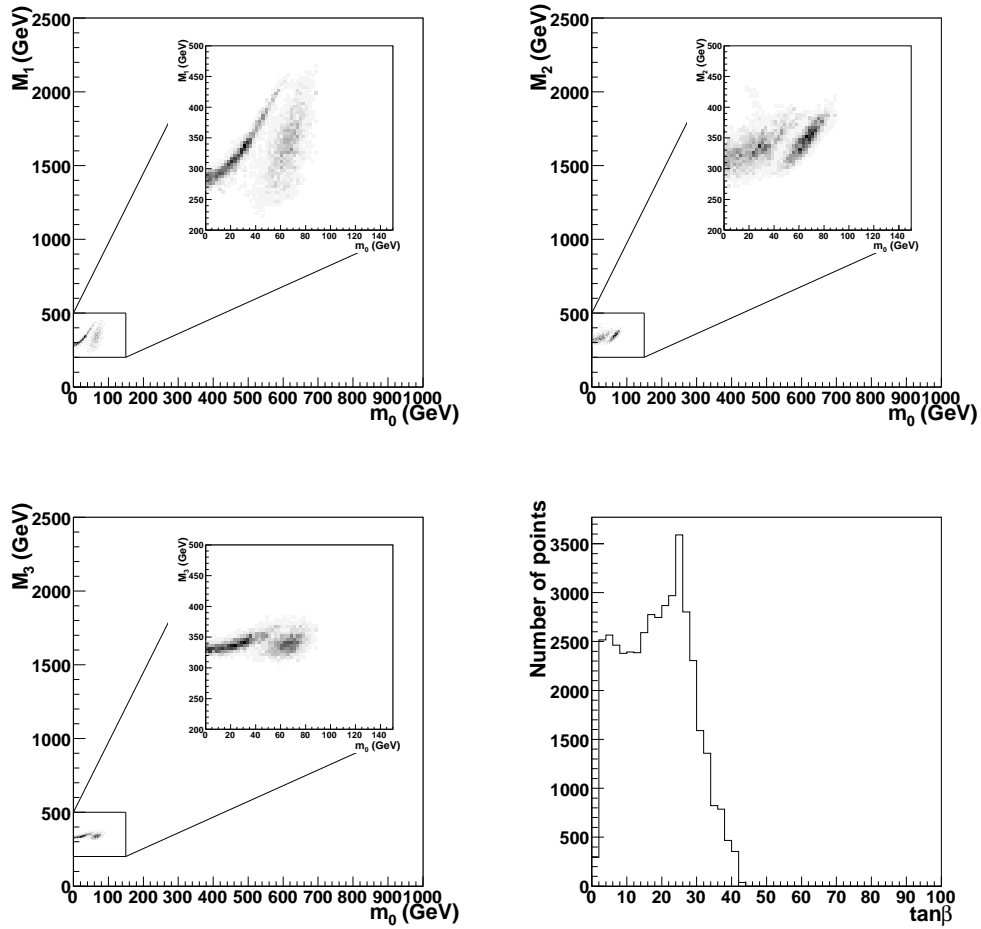
Figure 6.16: The region of the non-universal SUGRA parameter space consistent with the endpoint measurements of Section 5.2.3 and the cross-section measurement, with chain ambiguity included. Results are shown for positive μ. The Markov Chain has converged.

Figure 6.17: The region of the non-universal SUGRA parameter space consistent with the endpoint measurements of Section 5.2.3 and the cross-section measurement, with chain ambiguity included. Results are shown for negative μ. The Markov Chain has converged.
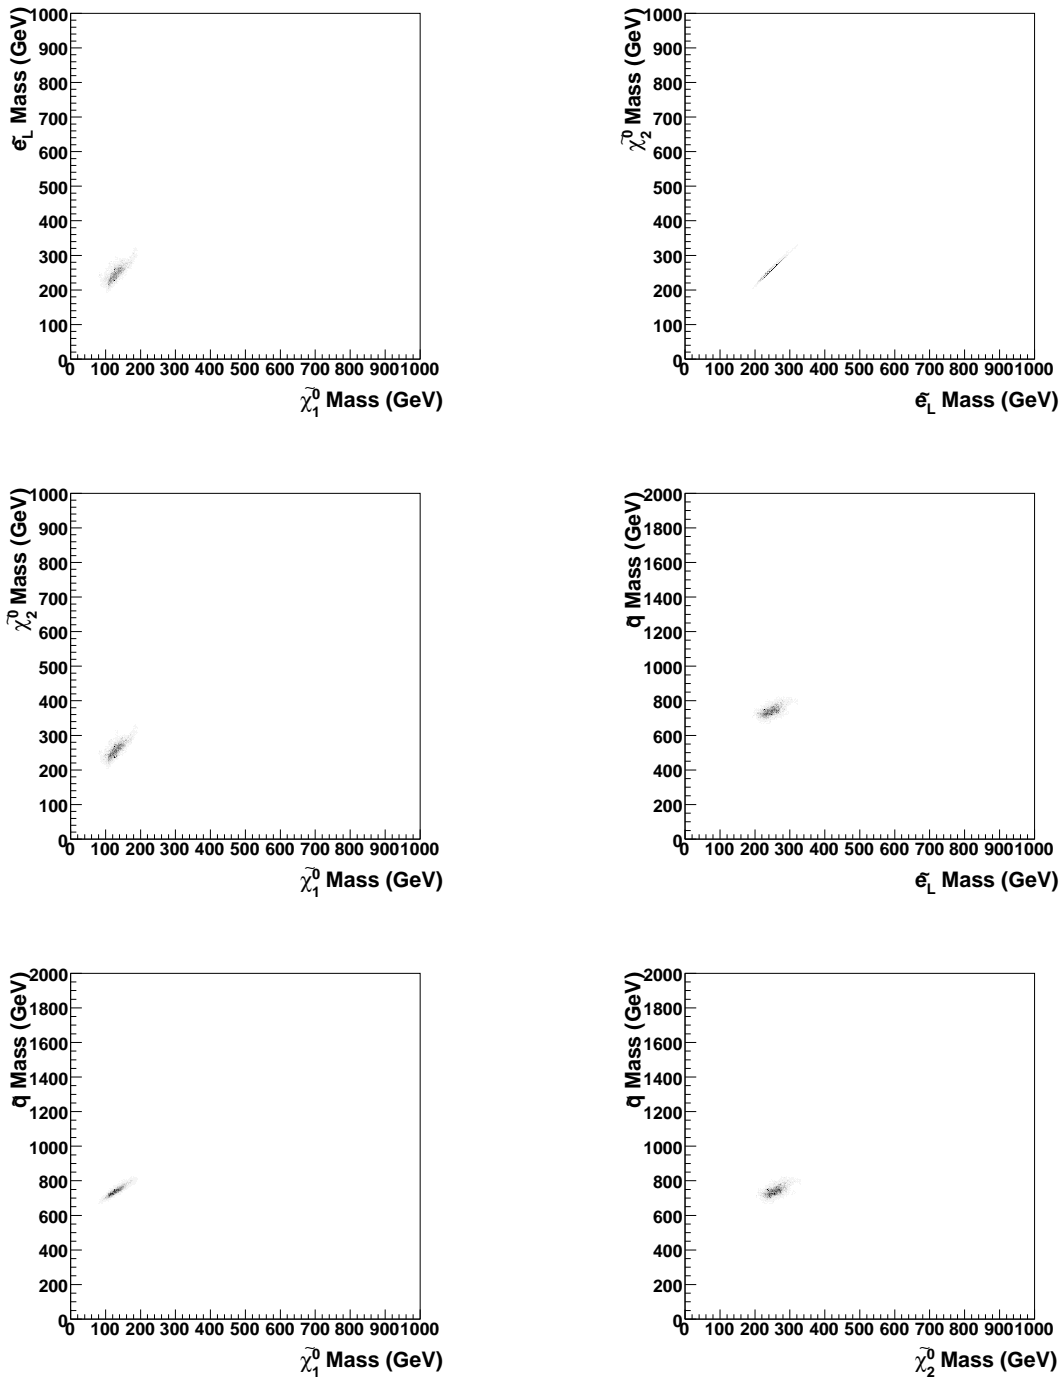
Figure 6.18: The region of mass space corresponding to the non-universal SUGRA parameter space region obtained in the text. The Markov Chain has converged.

As a final note, it is worth remarking that the technique is not limited solely to data obtained at the LHC. Any piece of relevant data is potentially useful, with obvious examples being cross-section limits for rare decay processes, and dark matter measurements that are currently already being used to set limits on theories (see, for example, references[72],[73] and[74]). This data may prove to be essential when one starts to explore models with greater numbers of parameters.

Chapter 7

NUHM models

Chapter 6 introduced a method for exploring SUSY parameter spaces with, in principle, an arbitrary number of parameters, whilst also demonstrating how to include effects such as errors on the jet energy scale and decay chain ambiguity. In the process, a short investigation of a model with non-universal gaugino masses was performed in order to demonstrate the utility of the method.

The purpose of this chapter is to briefly investigate another class of SUSY model, in which the GUT scale Higgs masses are non-universal (NUHM models). There has been much interest in these recently as viable extensions of mSUGRA since, although one can devise relatively strong theoretical motivations for gaugino mass universality at the GUT scale, there is no reason why the soft supersymmetry-breaking scalar masses of the electroweak Higgs multiplets should be universal. It is thus particularly important to consider models in which one breaks the degeneracy of these masses, and a set of benchmark points consistent with current measurements of the dark matter relic density, the $b \rightarrow s\gamma$ decay branching ratio and $g_\mu - 2$ was presented in reference[75]. Furthermore, it was observed in reference[76] that relatively rare phenomena in the mSUGRA parameter space become much more 'mainstream' in NUHM models, and hence they make important cases for study, given that they are not excluded experimentally.

For a Markov Chain enthusiast, NUHM models offer an interesting testing ground for new effects that are not commonly encountered in mSUGRA. For some model choices, decays featuring chains of successive two body decays are not present, and yet it is possible to observe decay chains involving a combination of two and three-body decays. Thus, it should still be possible to measure masses by the standard method of searching for kinematic endpoints, but

one first needs to derive expressions for their expected position. This introduces an extra layer of ambiguity when attempting to assign decay processes to the edges observed in invariant mass distributions, and requires an extension of the decay chain ambiguity work presented in Section 6.2.3. Cascade decays featuring the three-body decay mode could also occur in, for example, mSUGRA models [1], though they apparently have not been studied before.

The aims of this chapter are to investigate the collider phenomenology of an NUHM model, and to look at how the presence of the additional endpoints affects the decay chain ambiguity. Section 7.1 summarises the particular choice of NUHM model used in this chapter, reviewing both the mass spectrum and the relevant decay channels. The three-body endpoint positions for a general decay are derived in Section 7.2, and are applied to the NUHM model in Section 7.3. Section 7.4 discusses the prospects for more detailed analysis of the NUHM model, and considers how one should extend the decay chain ambiguity work discussed in the previous chapter.

7.1   Selection of NUHM model

The NUHM parameter space is related to that of mSUGRA by the addition of two extra parameters that express the non-universality of the two MSSM Higgs doublets. These can be specified at the GUT scale as the masses $m^2_{H_u}$ and $m^2_{H_d}$, or alternatively the conditions of electroweak symmetry breaking allow one to trade these for the weak scale parameters μ and m_A. In selecting a model for study, the benchmark model γ in reference[75] was found to be particularly interesting. The two body decay modes of the $\tilde{\chi}^0_2$ are not allowed, and hence one will not observe the characteristic two body endpoints seen in a variety of mSUGRA parameter space points but, rather, will have to develop other strategies for analysis. It is compatible with all current experimental constraints arising from, e.g. WMAP and limits on the branching ratios of rare decays, and is specified as follows:

$$m_0 = 328 \text{ GeV}, m_{1/2} = 247 \text{ GeV}$$

$$\tan\beta = 20, A_0 = 0, |\mu| = 325 \text{ and } m_A = 240$$

with a top quark mass of 178 GeV, and μ greater than zero.

---

[1] A suitable mSUGRA model for study would be obtained by taking the parameters of the NUHM benchmark point studied here and setting the GUT scale Higgs masses to the universal scalar mass m_0.
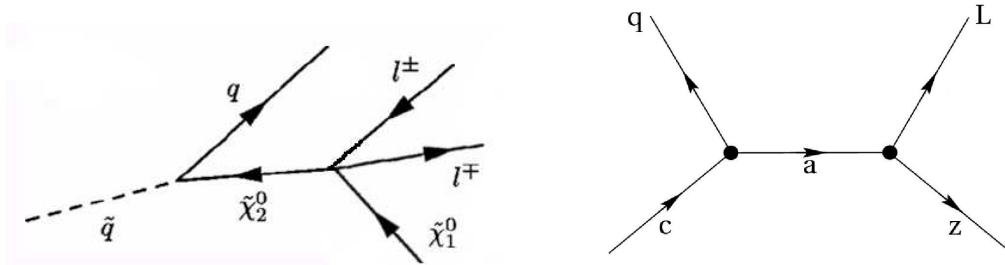
Figure 7.1: The cascade decay chain on the left, including a three-body decay, will give rise to kinematic endpoints. On the right, the same decay is shown but with the two leptons treated as a single object *L*, and with the sparticle symbols replaced by the notation used in the text.

`ISAJET 7.72`[77] has been used to generate the mass spectrum and decay information for the point, and the results are summarised in Tables 7.1, and Tables 7.2 and 7.3 respectively. In addition, `HERWIG 6.5`[65] was used to estimate a total SUSY production cross-section of 33 pb at this point. This differs from that given in reference[75] (where a sum of the $\tilde{g}\tilde{g}$, $\tilde{g}\tilde{q}$ and $\tilde{q}\tilde{q}$ cross-sections gives 55 pb), though is consistent with the fact that the Herwig calculation is only performed to leading order, whereas reference[75] quotes a next to leading order result.

The most relevant part of the decay table concerns the decay modes of the $\tilde{\chi}_2^0$, and Table 7.2 shows that one does not obtain two body decays to sleptons, but rather has three-body decays to quarks or leptons. Given that there are appreciable branching fractions for squark decays featuring $\tilde{\chi}_2^0$'s one will obtain decay chains of the form shown in Figure 7.1, and thus should be able to observe kinematic endpoints in lepton-jet invariant mass distributions using a similar method to that which has been previously documented for chains of successive two body decays. Each maximum will occur at a position given by a function of the three sparticle masses in the decay chain. Note that although the branching ratio for the decay $\tilde{\chi}_2^0 \rightarrow \tilde{\chi}_1^0 l^+ l^-$ is small, the large SUSY production cross-section will guarantee a reasonable sample of events (approximately 3000 events for an initial ATLAS sample of 30 fb^{-1}).

Particle	Mass/GeV
$\tilde{\chi}_1^0$	95
$\tilde{\chi}_2^0$	179
$\tilde{\chi}_3^0$	332
$\tilde{\chi}_4^0$	353
$\tilde{\chi}_1^\pm$	179
$\tilde{\chi}_2^\pm$	353
$\tilde{e}_L$	377
$\tilde{e}_R$	329
$\tilde{\nu}_e$	368
$\tilde{\tau}_1$	315
$\tilde{\tau}_2$	378
$\tilde{\nu}_\tau$	365
$\tilde{g}$	615
$\tilde{u}_L$	631
$\tilde{u}_R$	624
$\tilde{d}_L$	636
$\tilde{d}_R$	617
$\tilde{b}_1$	560
$\tilde{b}_2$	604
$\tilde{t}_1$	455
$\tilde{t}_2$	614
h^0	116
H^0	242
A^0	240
$H^\pm$	255

Table 7.1: The mass spectrum of the NUHM point defined in the text, as given by
`ISAJET 7.72`.

Decay Mode	BR
$\tilde{\chi}_2^0 \to \tilde{\chi}_1^0 q\bar{q}$	62%
$\tilde{\chi}_2^0 \to \tilde{\chi}_1^0 b\bar{b}$	19%
$\tilde{\chi}_2^0 \to \tilde{\chi}_1^0 l^+ l^-$	3.5%
$\tilde{\chi}_2^0 \to \tilde{\chi}_1^0 \tau^+ \tau^-$	2.7%
$\tilde{\chi}_2^0 \to \tilde{\chi}_1^0 \nu_l \bar{\nu}_l$	7.9%
$\tilde{\chi}_2^0 \to \tilde{\chi}_1^0 \nu_\tau \bar{\nu}_\tau$	3.9%
$\tilde{\chi}_3^0 \to \tilde{\chi}_1^\pm W$	62%
$\tilde{\chi}_3^0 \to \tilde{\chi}_1^0 Z$	14%
$\tilde{\chi}_3^0 \to \tilde{\chi}_2^0 Z$	21%
$\tilde{\chi}_3^0 \to \tilde{\chi}_1^0 h$	2.7%
$\tilde{\chi}_4^0 \to \tilde{\chi}_1^\pm W$	67%
$\tilde{\chi}_4^0 \to \tilde{\chi}_1^0 Z$	3.4%
$\tilde{\chi}_4^0 \to \tilde{\chi}_2^0 Z$	2.7%
$\tilde{\chi}_4^0 \to \tilde{\chi}_1^0 h$	9.1%
$\tilde{\chi}_4^0 \to \tilde{\chi}_1^0 H$	1.2%
$\tilde{\chi}_4^0 \to \tilde{\chi}_2^0 h$	16.5%
$\tilde{\chi}_1^\pm \to \tilde{\chi}_1^0 W^\pm$	99%
$\tilde{\chi}_2^\pm \to \tilde{\chi}_1^0 W^\pm$	9.7%
$\tilde{\chi}_2^\pm \to \tilde{\chi}_2^0 W^\pm$	39%
$\tilde{\chi}_2^\pm \to \tilde{\chi}_1^\pm Z$	30%
$\tilde{\chi}_2^\pm \to \tilde{\chi}_1^\pm h$	20%

Table 7.2: The dominant chargino and neutralino decay processes at the NUHM point defined in the text, as given by ISAJET 7.72, where q denotes a quark from the first two generations, and l is a lepton from the first two generations.

Decay Mode	BR	Decay Mode	BR	Decay Mode	BR
$(\tilde{u}_L, \tilde{c}_L) \to \tilde{\chi}_2^0 q$	30%	$(\tilde{u}_R, \tilde{c}_R) \to \tilde{\chi}_1^0 q$	96%	$\tilde{g} \to \tilde{b}_1 b$	81%
$(\tilde{u}_L, \tilde{c}_L) \to \tilde{\chi}_4^0 q$	2%	$(\tilde{u}_R, \tilde{c}_R) \to \tilde{\chi}_2^0 q$	1%	$\tilde{g} \to \tilde{b}_2 b$	4%
$(\tilde{u}_L, \tilde{c}_L) \to \tilde{g} q$	1.5%	$(\tilde{u}_R, \tilde{c}_R) \to \tilde{g} q$	2.3%	$\tilde{g} \to \tilde{\chi}_1^\pm q \bar{q}$	6.8%
$(\tilde{u}_L, \tilde{c}_L) \to \tilde{\chi}_1^+ q$	63%			$\tilde{g} \to \tilde{\chi}_1^0 q \bar{q}$	2.2%
$(\tilde{u}_L, \tilde{c}_L) \to \tilde{\chi}_2^+ q$	2.5%			$\tilde{g} \to \tilde{\chi}_2^0 q \bar{q}$	3.4%
$(\tilde{d}_L, \tilde{s}_L) \to \tilde{\chi}_1^0 q$	2.1%	$(\tilde{d}_R, \tilde{s}_R) \to \tilde{\chi}_1^0 q$	98%		
$(\tilde{d}_L, \tilde{s}_L) \to \tilde{\chi}_2^0 q$	30%	$(\tilde{d}_R, \tilde{s}_R) \to \tilde{\chi}_2^0 q$	1%		
$(\tilde{d}_L, \tilde{s}_L) \to \tilde{\chi}_4^0 q$	2.7%				
$(\tilde{d}_L, \tilde{s}_L) \to \tilde{\chi}_1^- q$	56%				
$(\tilde{d}_L, \tilde{s}_L) \to \tilde{\chi}_2^- q$	8%				
$\tilde{b}_1 \to \tilde{\chi}_1^0 b$	3.6%	$\tilde{t}_1 \to \tilde{\chi}_1^0 t$	17%		
$\tilde{b}_1 \to \tilde{\chi}_2^0 b$	26%	$\tilde{t}_1 \to \tilde{\chi}_2^0 t$	13%		
$\tilde{b}_1 \to \tilde{\chi}_3^0 b$	2.2%	$\tilde{t}_1 \to \tilde{\chi}_1^+ b$	50%		
$\tilde{b}_1 \to \tilde{\chi}_4^0 b$	2.3%	$\tilde{t}_1 \to \tilde{\chi}_2^+ b$	20%		
$\tilde{b}_1 \to \tilde{\chi}_1^- t$	36%				
$\tilde{b}_1 \to \tilde{\chi}_2^- t$	26%				
$\tilde{b}_1 \to \tilde{t}_1 W$	3.8%	$\tilde{t}_2 \to \tilde{t}_1 h$	7%		
$\tilde{b}_2 \to \tilde{\chi}_1^0 b$	13%	$\tilde{t}_2 \to \tilde{\chi}_1^0 t$	1.8%		
$\tilde{b}_2 \to \tilde{\chi}_2^0 b$	2.4%	$\tilde{t}_2 \to \tilde{\chi}_2^0 t$	8.5%		
$\tilde{b}_2 \to \tilde{\chi}_3^0 b$	13%	$\tilde{t}_2 \to \tilde{\chi}_3^0 t$	9.5%		
$\tilde{b}_2 \to \tilde{\chi}_4^0 b$	14%	$\tilde{t}_2 \to \tilde{\chi}_4^0 t$	27%		
$\tilde{b}_2 \to \tilde{\chi}_1^- t$	3.2%	$\tilde{t}_2 \to \tilde{\chi}_1^+ b$	22%		
$\tilde{b}_2 \to \tilde{\chi}_2^- t$	46%	$\tilde{t}_2 \to \tilde{\chi}_2^+ b$	21%		
$\tilde{b}_2 \to \tilde{t}_1 W$	8.2%	$\tilde{t}_2 \to \tilde{t}_1 Z$	7%		
$\tilde{l}_L \to \tilde{\chi}_1^0 l$	12%	$\tilde{l}_R \to \tilde{\chi}_1^0 l$	99%	$\tilde{\nu}_l \to \tilde{\chi}_1^0 \nu_l$	17%
$\tilde{l}_L \to \tilde{\chi}_2^0 l$	33%			$\tilde{\nu}_l \to \tilde{\chi}_2^0 \nu_l$	24%
$\tilde{l}_L \to \tilde{\chi}_1^- \nu_e$	54%			$\tilde{\nu}_l \to \tilde{\chi}_1^+ l$	59%
$\tilde{\tau}_1 \to \tilde{\chi}_1^0 \tau$	81%	$\tilde{\tau}_2 \to \tilde{\chi}_1^0 \tau$	16%	$\tilde{\nu}_\tau \to \tilde{\chi}_1^0 \nu_\tau$	17%
$\tilde{\tau}_1 \to \tilde{\chi}_2^0 \tau$	6.9%	$\tilde{\tau}_2 \to \tilde{\chi}_2^0 \tau$	32%	$\tilde{\nu}_\tau \to \tilde{\chi}_2^0 \nu_\tau$	24%
$\tilde{\tau}_1 \to \tilde{\chi}_1^- \nu_\tau$	12%	$\tilde{\tau}_2 \to \tilde{\chi}_1^- \nu_\tau$	50%	$\tilde{\nu}_\tau \to \tilde{\chi}_1^+ \tau$	60%

Table 7.3: The dominant sfermion decay processes at the NUHM point defined in the text, as given by `ISAJET 7.72`, where q denotes a quark from the first two generations, and l is a lepton from the first two generations.

7.2   Kinematic endpoint derivation

7.2.1   Introduction

In this section, the decay $\tilde{q} \to q\tilde{\chi}_2^0$ followed by $\tilde{\chi}_2^0 \to l^+l^-\tilde{\chi}_1^0$ (shown in Figure 7.1) is studied. As with the two body case, such decays are fairly easy to select given that one can look for events with opposite-sign-same-flavour (OSSF) leptons, combined with the missing energy from the undetected neutralinos [2]. By taking different combinations of the visible decay products, one can form various invariant masses; m_{ll}, m_{llq}, m_{lq}^{max} and m_{lq}^{min}, where m_{lq}^{max} is the higher of the two m_{lq} invariant masses that can be formed in the event, and m_{lq}^{min} is the lower. These will have maxima resulting from kinematic limits, whose position is given by a function of $m_{\tilde{q}}$, $m_{\tilde{\chi}_2^0}$ and $m_{\tilde{\chi}_1^0}$, and these are derived for each case below.

In the following derivations, bold type is used for three momenta, and four vector quantities are denoted by explicitly showing Lorentz indices. Squared masses are represented by a non-bold character (e.g. $c = m_{\tilde{q}}^2$).

7.2.2 m_{ll} endpoint

The endpoint of the m_{ll} distribution results from the three-body decay of the $\tilde{\chi}_2^0$, and is given trivially by the mass difference between the $\tilde{\chi}_2^0$ and the $\tilde{\chi}_1^0$:

$$(m_{ll}^2)^{\text{max}} = (m_{\tilde{\chi}_2^0} - m_{\tilde{\chi}_1^0})^2. \tag{7.1}$$

7.2.3 m_{llq} endpoint and threshold

The calculation of the m_{llq} endpoint presented here follows the method given in appendix E of[66] by treating the decay as shown in Figure 7.1, where the two leptons are combined into a single object L, with an invariant mass given by m_{ll}. It is known from the dilepton invariant mass that $m_{ll} \equiv m_L$ must lie within a specific range:

$$m_L = \lambda(m_a - m_z), \lambda \in [0, 1] \tag{7.2}$$

---

[2]Note that this does not preclude the possibility of selecting the usual two body cascade process, and ways to resolve this ambiguity are considered in Section 7.4.

If one looks at the decay of Figure 7.1 in the rest frame of a, one can conserve four momentum to obtain the following expressions for the three momenta of q and L:

$$\mathbf{L}^2 = \mathbf{z}^2 = [m_L^2, m_a^2, m_z^2] \tag{7.3}$$

$$\mathbf{q}^2 = \mathbf{c}^2 = [0, m_a^2, m_c^2] \tag{7.4}$$

where

$$[x, y, z] \equiv \frac{x^2 + y^2 + z^2 - 2(xy + xz + yz)}{4y} \tag{7.5}$$

and the quark has been treated as massless.

In general, the invariant mass of q and L is given by

$$m_{qL}^2 = g_{\mu\nu}(L^\mu + q^\mu)(L^\nu + q^\nu) \tag{7.6}$$

$$= m_L^2 + 2|\mathbf{q}|(E_L - |\mathbf{L}|\cos\theta) \tag{7.7}$$

in the case where q is massless, and where θ is the angle between L and q in the rest frame of a. The maximum will occur when $\cos\theta$ is equal to -1, and hence L and q are back to back in the a rest frame. Combining this with the knowledge of $\mathbf{L}$ and $\mathbf{q}$ from Equations 7.3 and 7.4, one obtains the expression for the endpoint of the m_{llq} distribution in terms of m_L:

$$(m_{llq})^2 = L_m + \frac{(c - a)}{2a}\left[L_m - (z - a) + \sqrt{((a + z) - L_m)^2 - 4az}\right] \tag{7.8}$$

where $L_m = m_L^2$, $c = m_c^2$, $a = m_a^2$ and $z = m_z^2$. L_m can take any value in the range specified by Equation 7.2, and one now needs to maximise Equation 7.8 by considering separately the cases where $\lambda = 0$, $0 < \lambda < 1$ and $\lambda = 1$. After doing this, two possible expressions for the maximum of the m_{llq} distribution are obtained:

$$(m_{llq}^2)^{\text{max}} = \begin{cases} (m_{\tilde{q}} - m_{\tilde{\chi}_1^0})^2 & \text{if } m_{\tilde{\chi}_2^0}^2 > m_{\tilde{q}}m_{\tilde{\chi}_1^0} \\ (m_{\tilde{q}}^2 - m_{\tilde{\chi}_2^0}^2)(m_{\tilde{\chi}_2^0}^2 - m_{\tilde{\chi}_1^0}^2)/m_{\tilde{\chi}_2^0}^2 & \text{otherwise.} \end{cases} \tag{7.9}$$

In addition to finding an edge in the m_{llq} distribution, one can observe a threshold. Equation 7.7
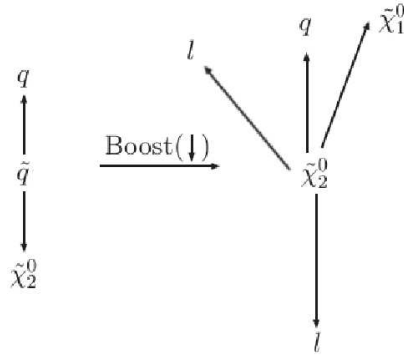
Figure 7.2: The two rest frames involved in the squark cascade decay.

has a minimum when $\cos\theta$ is equal to 1, in which case one obtains a minimum value of m_{llq}:

$$(m_{llq}^{min})^2 = L_m + \frac{(c-a)}{2a}\left[L_m - (z-a) - \sqrt{((a+z)-L_m)^2 - 4az}\right] \tag{7.10}$$

If L_m lies at the lower end of its allowed range, then $L_m = m_{llq}^{min} = 0$. However, one can raise the minimum value of m_{llq} by looking at the subset of events for which L_m is greater than some arbitrary cut value. This will then give an observable threshold in the m_{llq} distribution.

7.2.4 m_{lq}^{max} and m_{lq}^{min} endpoints

In the case of the m_{llq} endpoint, it has been shown that there are in fact two expressions, each of which applies in a specific region of mass space. In anticipation of this, a general method was used to perform the calculation in order to avoid missing one of the solutions.

The m_{lq}^{max} endpoint is conceptually much easier, however, as one only has to maximise the invariant mass that can be made from one of the leptons. The two sequential decays are shown in Figure 7.2, which shows the effect of a boost from the $\tilde{q}$ rest frame to the $\tilde{\chi}_2^0$ rest frame. Any maximum in the m_{lq} invariant mass must arise from having the relevant lepton (the 'interesting lepton') back to back with the quark in the $\tilde{\chi}_2^0$ rest frame. One can thus consider three extreme cases for the configuration of the leptons and $\tilde{\chi}_1^0$ in the $\tilde{\chi}_2^0$ rest frame:

1. The $\tilde{\chi}_1^0$ is produced at rest, and the two leptons are back to back with one of them anti-parallel to the quark.

2. One of the leptons is produced at rest, and so the $\tilde{\chi}_1^0$ is produced back to back with the

other lepton, with the interesting lepton being anti-parallel to the quark.

3. None of the particles from the three-body decay is produced at rest, in which case
   we will get the highest invariant mass by having the interesting lepton emerging anti-
   parallel to the quark, and the other two particles travelling in the same direction as the
   quark.

Obtaining the m_{lq} endpoint is simply a case of working out which of these gives the highest
invariant mass. A short calculation gives:

$$(m_{lq}^2)_{\text{max}}^{\text{max}} = \frac{(m_{\tilde{q}}^2 - m_{\tilde{\chi}_2^0}^2)(m_{\tilde{\chi}_2^0}^2 - m_{\tilde{\chi}_1^0}^2)}{m_{\tilde{\chi}_2^0}^2} \tag{7.11}$$

The m_{lq}^{min} endpoint is harder to obtain than the m_{lq}^{max} endpoint, given that one wants the max-
imum value of the smallest m_{lq} invariant mass it is possible to form in an event. A similar
approach to that used in the previous subsection is applied, whereby the decay configuration
that will give the maximum is visualised before working out the endpoint.

In this case, it has been proved that there is a local m_{lq}^{min} maximum when the two leptons are
produced parallel in the $\tilde{\chi}_2^0$ rest frame, travelling anti-parallel to the $\tilde{\chi}_1^0$ (which therefore travels
parallel to the quark). Note that this does not exclude the possibility of other local maxima,
but numerical simulation has not revealed any. Therefore, this configuration is assumed to
give the global maximum of m_{lq}^{min} [3].

A short calculation gives:

$$(m_{lq}^2)_{\text{min}}^{\text{max}} = \frac{(m_{\tilde{q}}^2 - m_{\tilde{\chi}_2^0}^2)(m_{\tilde{\chi}_2^0}^2 - m_{\tilde{\chi}_1^0}^2)}{2m_{\tilde{\chi}_2^0}^2} \tag{7.12}$$

7.2.5   Summary

Having obtained endpoints for the m_{ll}, m_{llq}, m_{lq}^{max} and m_{lq}^{min} distributions, it can be seen that
the expressions are very similar. The ratio of the m_{lq}^{max} and m_{lq}^{min} endpoint positions will al-
ways be $\sqrt{2}$, and, in a particular mass region, the m_{llq} endpoint is coincident with the m_{lq}^{max}
endpoint. This ultimately means that there may not be enough information to precisely deter-

---

[3]Other local maxima, were any to exist, would occur in configurations in which m_{lq}^{min} is equal to m_{lq}^{max} but in
which the moduli of the three momenta of the leptons are unequal (but coplanar with the quark) in the rest frame
of the heavier neutralino.

mine the mass differences involved in the decay chain, a point that will be discussed further in Section 7.4.

7.3   Observation of three-body endpoints in NUHM model

Having derived the endpoints for the process depicted in Figure 7.1, it is interesting to look at a concrete physics example by performing a Monte Carlo study of the NUHM model described in Section 7.1.

7.3.1   Event generation and simulation

The mass spectrum and decay table of the NUHM point were taken from `ISAJET 7.72` using the `ISAWIG` interface. 3,300,000 signal events were subsequently generated using `HERWIG 6.5` (corresponding to an integrated luminosity of $100 \, \text{fb}^{-1}$). This implements three-body decays of SUSY particles with spin correlations, with the decays of interest here being:

1. $\tilde{\chi}_2^0 \to Z \tilde{\chi}_1^0 \to l^+ l^- \tilde{\chi}_1^0$

2. $\tilde{\chi}_2^0 \to l \tilde{l} \to l^+ l^- \tilde{\chi}_1^0$

where the Z and the slepton are off-shell. These generated events are then passed through the `ATLFAST` detector simulation (from the same release), whose jet cone algorithm used a cone with $\Delta R = 0.4$. Electrons, muons and jets were subject to a minimum p_T cut of 5, 6 and 10 GeV respectively.

7.3.2   Selection cuts

In order to observe clear endpoints from the cascade decay process, it is necessary to first isolate a clean sample of squark decay events. As in the two body case, one can select events with two OSSF leptons and a large amount of missing energy, and can also exploit the fact that one expects hard jets in SUSY events. Hence, all plots that follow are subject to the following cuts:

- $E_T^{\text{miss}} > 300$ GeV;

- exactly two opposite-sign leptons with $p_T > 5$ GeV for electrons and $p_T > 5$ GeV for muons, with $|\eta| < 2.5$;

- at least two jets with $p_T > 150$ GeV;

There are of course many SM processes that contribute to the dilepton background in any SUSY analysis, though it has been shown that these are highly suppressed once an OSSF cut is used in conjunction with cuts on lepton and jet p_T, and on missing energy. Although there is in principle still a tail of SM events that can contribute, it has found to be negligible in the past (see, for example, reference[78]), and a full study of this background is considered to be beyond the scope of this thesis. In addition to the SM background, it has already also noted that the OSSF lepton signature can be produced by SUSY processes other than the decay of the $\tilde{\chi}_2^0$. and that this can often be subtracted by producing "flavour subtracted plots" in which one plots the combination $e^+e^- + \mu^+\mu^- - e^+\mu^- - e^-\mu^+$. Each of the plots presented below has been flavour subtracted and some are subject to additional cuts which are stated in the text.

A 610 GeV "typical" squark mass is chosen here for the calculation of endpoints.

7.3.3   Invariant mass plots

m_{ll} **plot**

The m_{ll} distribution is shown in Figure 7.3. 4566 events survive after cuts and background subtraction, though the effect of the trigger which may cut more harshly on lepton p_T has not been considered. Using the mass spectrum given in Section 7.1, one expects to find an endpoint at approximately 80 GeV, and this is clearly visible.

The shape of the distribution is very different from the triangular shape normally encountered in the case of successive two body decays (resulting from the phase space for that process), and this might prove important when attempting to distinguish three-body from two body decays. This is considered further in Section 7.4.

m_{llq} **plots**

As soon as one starts to form invariant masses involving quarks, it is important to consider how to select the correct quark from the cascade decay. As seen earlier, a reasonable assumption is that the two hardest jets in the event will come from squark decay on either side of the event,
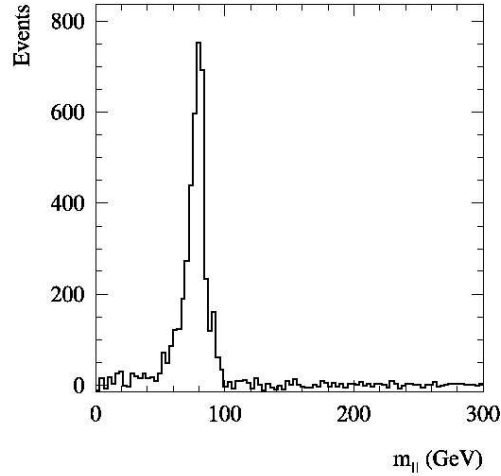
Figure 7.3: The m_{ll} distribution for the NUHM model defined in the text, with flavour subtraction.
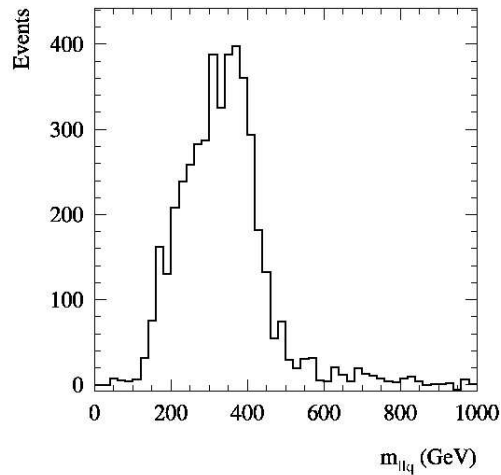
Figure 7.4: The flavour subtracted m_{llq} distribution for the NUHM model defined in the text, constructed by taking the lowest m_{llq} invariant mass that can be formed from the two hardest jets in the event.

and if one takes the lowest of the two m_{llq} invariant masses formed from the two hardest jets in the event, this should lie below the m_{llq} endpoint. The distribution of this m_{llq} is shown in Figure 7.4, and there is a visible endpoint consistent with the predicted value of approximately 490 GeV. The plot contains the same number of events as the m_{ll} plot, as the cuts are the same.

In order to obtain a further constraint on the physical model underlying the data, a threshold
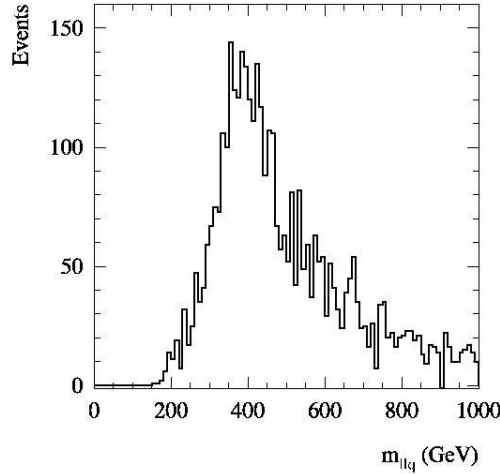
Figure 7.5: The flavour subtracted m_{llq} threshold plot, constructed using the highest of the m_{llq} invariant masses that can be formed from the two hardest jets in each selected event.

is constructed in the m_{llq} distribution. Here, the choice is made to follow the convention used in the study of successive two body decays, and thus to examine the subset of events for which $m_{ll} > m_{ll}^{max}/\sqrt{2}$. Substituting $m_{ll}^{max}/\sqrt{2}$ in place of L_m in Equation 7.10 gives the following threshold:

$$(m_{llq}^{min})^2 = \frac{(\sqrt{a} - \sqrt{z})^2}{2} + \frac{(c - a)}{4a}\left(3a - z - 2\sqrt{az} - \sqrt{a^2 + z^2 + 4\sqrt{az}(a + z) - 10az}\right) \quad (7.13)$$

where $c = m_{\tilde{q}}^2$, $a = m_{\tilde{\chi}_2^0}^2$ and $z = m_{\tilde{\chi}_1^0}^2$. Traditionally (i.e. in chains with successive *two-body* decays) this additional constraint requires, somewhat arbitrarily, that the angular separation of the two leptons in the rest frame of the slepton be greater than a right angle. For the *three-body* neutralino decay considered in this chapter, that geometrical interpretation is lost, but this is of no serious consequence.

A plot of the m_{llq} distribution is given in Figure 7.5, where, because a threshold is being searched for, the highest of the two m_{llq} invariant masses formed with the two hardest jets in the event is used to make the plot. 4172 events are contained in the plot. A threshold structure of some form is clearly observed, though it is difficult to ascertain the precise position, as the shape of the edge is not yet a well understood function of the sparticle masses and cut-induced 'detector' effects. To use the constraint from this edge to the full, it may be necessary to repeat the analysis of reference[79] in the context of a three-body final decay. The predicted value is approximately 240 GeV.

m_{lq} plots

The m_{lq}^{max} distribution is plotted by forming m_{llq} invariant masses with the two hardest jets in the event. The jet from the lowest mass combination (which is the best guess for the quark emitted in the squark cascade decay) is then used to form the m_{lq} invariant mass with each of the leptons in the events. The maximum of these is plotted in the m_{lq}^{max} plot (shown in Figure 7.6), where there is an additional cut that the dilepton invariant mass in each selected event must be less than the dilepton endpoint. 4161 events are in the plot. There is an endpoint predicted at about 490 GeV (we are in the mass region where it should appear at the same position as the m_{llq} endpoint) and this is consistent with the plot, though it is difficult to identify the endpoint due the fact that the shape is easily confused with the tail. If detector effects are removed and the shape of the distribution is observed in a clean environment in which there is only one squark mass and one only considers the effect that phase space has on the shape of the distribution near the endpoint (Figures 7.7 and 7.8) then it is seen that the endpoint is only approached quadratically. Although a full analysis of the tail would probably require full simulation (and thus a separate study), an attempt has been made to determine how much is caused by detector smearing, and how much is caused by background SUSY processes. Figures 7.9 to 7.11 examine the m_{eq}^{max} distribution, showing the Monte Carlo truth plot, the plot obtained by selecting events on the basis of truth but with the particles reconstructed by the `ATLFAST` detector simulation, and finally a plot which contains only SUSY background processes. The tail does not predominantly arise from detector smearing (which will nevertheless smear the endpoint), but has instead a large contribution from the SUSY background. There is also a combinatoric background related to the wrong choice of squark, but this is harder to isolate.

The m_{lq}^{min} plot is constructed in a similar fashion to the m_{lq}^{max} plot, with the exception that the lowest of the two possible m_{lq} combinations is plotted for each event. The result is shown in Figure 7.12, where there is an additional cut; one of the m_{llq} invariant masses formed from the two hardest jets in the event must lie below the approximate observed position of the m_{llq} endpoint, whilst the other must lie above, leaving 1664 events in the plot. This removes much of the tail due to incorrect squark choice, and leaves a very clean endpoint at the predicted value of approximately 350 GeV [4].

---

[4]Note that this extra cut is not possible in the case of the m_{lq}^{max} endpoint, as the m_{lq}^{max} distribution is highly correlated to the m_{llq} distribution (the events at the endpoint are the same in both cases). Hence, performing this cut on the m_{lq}^{max} distribution artificially removes any events beyond the endpoint.
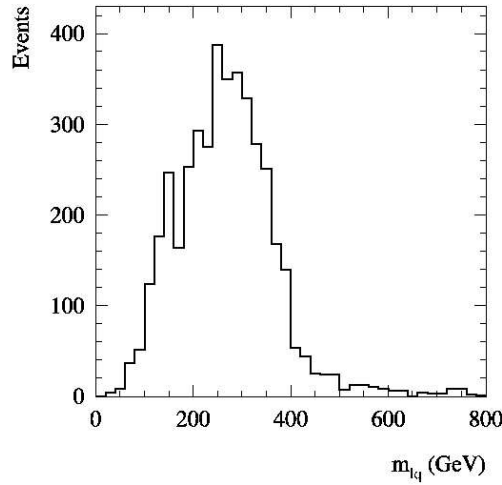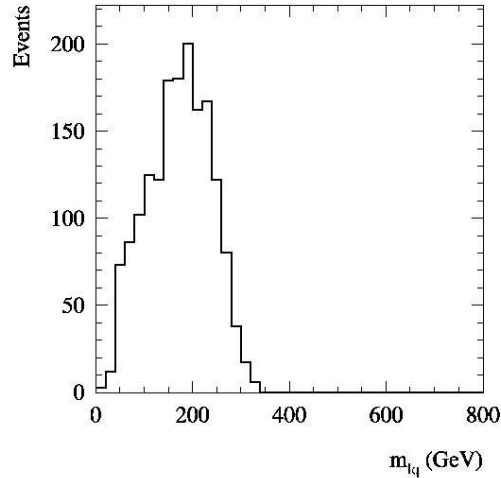
Figure 7.6: The flavour subtracted m_{lq}^{max} distribution for the NUHM model defined in the text, constructed by taking the jet (from the two hardest jets in the event) that gives the lowest m_{llq} invariant mass and forming the highest invariant mass that one can make with the two leptons in the event.
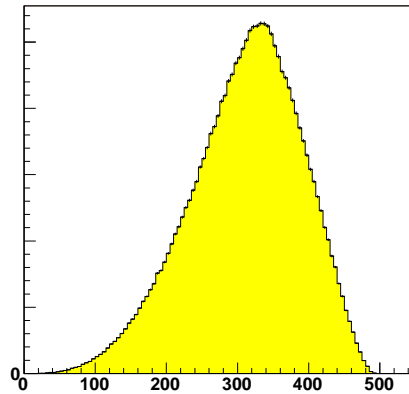
Figure 7.7: A calculation of the shape of an example m_{lq}^{max} distribution (for particle masses of $m_{\tilde{\chi}_1^0} = 95$ GeV, $m_{\tilde{\chi}_2^0} = 179$ GeV and $m_{\tilde{q}} = 610$ GeV similar to those in the NUHM model defined in the text) using the approximation in which all particles are taken to be scalars, i.e. phase-space only. The calculation assumes 100% acceptance and does not model any detector effects. m_{lq}^{max} values are in GeV.
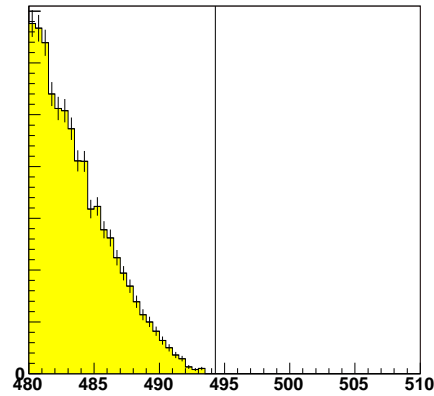
Figure 7.8: A zoomed view of the m_{lq}^{max} phase space distribution (left) in the vicinity of the end point $(m_{lq})_{max}^{max}$, the position of which is given by the vertical line at 494.24 GeV. m_{lq}^{max} values are in GeV.
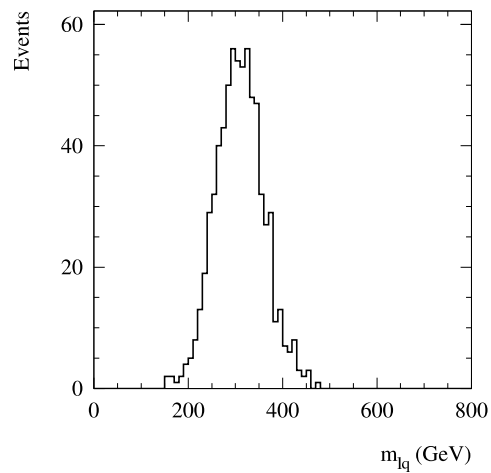
Figure 7.9: The truth distribution for the m_{eq}^{max} invariant mass, taken from the Herwig event record as recorded in `ATLFAST`. The plot exhibits a clean edge with no tail.
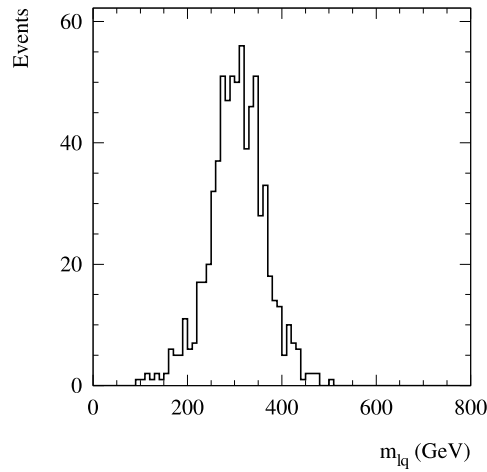
Figure 7.10: The m_{eq}^{max} distribution obtained by selecting events on the basis of Monte Carlo truth information, but with the electrons and jets reconstructed by `ATLFAST`. The plot has a slightly higher endpoint than the truth distribution, but no significant tail.
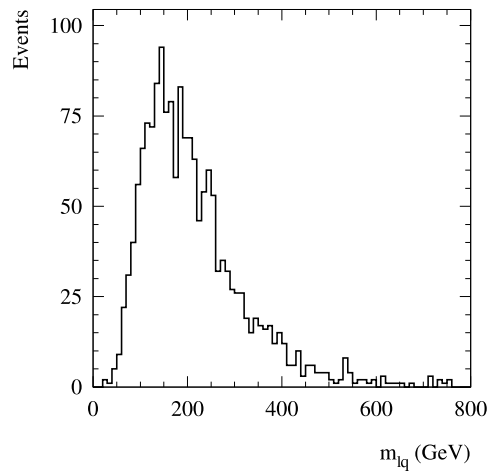
Figure 7.11: The m_{eq}^{max} distribution obtained using events with no decay of the form $\tilde{\chi}_2^0 \rightarrow \tilde{\chi}_1^0 e^+ e^-$, representing the contribution to the plot from SUSY background processes.

Figure 7.12: The flavour subtracted m_{lq}^{min} distribution for the NUHM model defined in the text, constructed by taking the jet (from the two hardest jets in the event) that gives the lowest m_{llq} invariant mass and forming the lowest invariant mass that one can make with the two leptons in the event.

7.4 Discussion

Having collected a series of endpoint plots and derived the edge positions in terms of the sparticle masses involved in the decay chain, one ought in principle to be able to reconstruct the masses in the chain. This is not as simple as it first appears, however, for the following reasons:

1. The NUHM point studied here is in the mass region where the m_{lq}^{max} endpoint is in the same position as the m_{llq} endpoint, and it is also known that the m_{lq}^{min} edge is related to these other two merely through multiplication by a constant factor. Discounting the poorly measured threshold, there are thus only two equations in three unknowns (the m_{ll} endpoint can be used as the second equation), and there is not enough information to constrain the masses.

2. The observation of endpoints does not reveal anything about the decay process that produces them, and the shapes of the lepton-jet invariant mass distributions in this chapter are not dissimilar to those produced by chains of successive two body decays. Hence, one needs to consider how one would in principle distinguish between the two cases to be certain that the correct masses are being reconstructed.

The adequate solution of these problems is unfortunately too involved to be considered for inclusion in this thesis, and will have to wait for a subsequent study. However, the following discussion will suggest possible answers to both.

7.4.1   Mass reconstruction

In order to reconstruct the sparticle masses in the three-body process, one will need to supply extra constraints. One possible method involves going one step higher in the decay chain, and searching for decays of the form:

$$\tilde{g} \to \tilde{q}q \to \tilde{\chi}_2^0 qq \to \tilde{\chi}_1^0 qqll \tag{7.14}$$

This will give more endpoints, and hence will provide direct constraints on the masses of the system. This is similar to an approach used for two body decays[80], and would have the advantage of providing a measurement of the gluino mass. The problem here is that the NUHM model has a relatively light gluino, which is lighter than the squarks of the first two generations. Hence, the method will only be applicable to decay chains involving stop and sbottom squarks, and we see that whilst this approach may be helpful in some cases it is certainly not generally applicable to all regions of parameter space.

Another approach is to use other observables to constrain sparticle masses. The work presented in the previous chapter demonstrated that one can obtain a dramatic improvement in mass measurements by combining exclusive data (such as endpoint information) with inclusive observables (such as the cross-section of events passing a missing p_T cut, to give one example). This analysis could be repeated here, and with enough inclusive observables one could obtain good measurements of the sparticle masses involved in our cascade process. This has the advantage of being generally applicable regardless of the mass spectrum, though it relies on specifying a particular SUSY model in which to perform the analysis. Given that the distinguishing feature of NUHM models is a more complicated Higgs sector (with non-degenerate Higgs mass parameters), these extra constraints must be Higgs related.

7.4.2   Decay chain ambiguity

It has already been remarked that it is not trivial to distinguish three-body endpoints from two body endpoints, and hence the observation of endpoints alone is not enough to be able to reconstruct masses. This adds an extra layer of ambiguity to that investigated in Section 6.2.3,

and there follow some suggestions for how this might be resolved.

Previous sections have illustrated one characteristic feature of three-body decay chains; the ratio of the m_{llq} and m_{lq}^{min} endpoint positions is always $\sqrt{2}$. Provided one can obtain precise measurements of these quantities, one would have a clue that three-body processes were being observed, although this could easily be faked by two body decays that conspired to produce endpoints in similar positions.

For an extra clue, consider that although the shapes of the lepton-jet distributions in three-body decays are not dissimilar to those encountered in two body processes, this is not true of the m_{ll} distribution, and hence there is potentially some information contained in the shape of the dilepton distribution. In the two-body case, one obtains a triangular shape that is identical to the phase space distribution. In contrast, the three-body distribution in Figure 7.3 is heavily peaked toward the endpoint. Unfortunately, this is unlikely to be true over the whole of parameter space; the three-body decay proceeds via an off-shell Z or slepton, and the m_{ll} distribution is peaked toward the endpoint when the endpoint (which is the same as the mass difference between the $\tilde{\chi}_2^0$ and the $\tilde{\chi}_1^0$) approaches, for example, the Z mass, such that the shape will depend heavily on the SUSY point. Furthermore, a previous study investigated the effect on the m_{ll} shape of incorporating the matrix element for the three-body decay process in addition to the phase space, for different points in mSUGRA parameter space, and although for some points it dramatically altered the pure phase space result there were other points where the three-body and phase space shapes were virtually indistinguishable[81].

To summarise, physicists may be fortunate enough to find that Nature presents a point at which they can distinguish three-body decays from two body processes simply by looking at the endpoint shapes and positions, but this is certainly not guaranteed. For this reason, a better approach to the problem of mass reconstruction is to use the Markov Chain techniques presented in Chapter 6, where no assumption is made about the processes causing the observed endpoints. This allows the selection of a region of the parameter space consistent with the data which can be used as a basis for further investigation, in a way directly analogous to that used in Section 6.2.3.

7.5   Summary

The expressions for the position of the kinematic endpoints arising in cascade decays featuring a two body decay followed by a three-body decay have been derived, and have been applied to

the lepton-jet invariant mass distributions given by a squark decay process. This has involved the first analysis of an NUHM model as it would appear within the ATLAS detector, and, using standard cuts, it has been possible to observe endpoints in simulated data that are consistent with the predicted positions. It is thus concluded that the technique is a viable extension of the current method used in chains of two body decays.

The existence of models similar to that studied here is of great relevance to the work presented in Chapter 6. The mass reconstruction is hampered by both a lack of constraint from the endpoint equations themselves, and the problem of distinguishing three-body from two body decay processes. In the case of the NUHM benchmark point γ, one would be able to identify the decays as three-body decays by using the shape of the dilepton invariant mass distribution, though this is not expected to be generally true over the whole of the parameter space.

Chapter 8

Summary

The Standard Model of particle physics is an immensely successful theory of Nature, but it has obvious deficiencies. There are strong hints that new physics will appear at the TeV scale.

This thesis has looked at both the machinery required to find this new physics, and the analytical method by which one of the proposed new theories might be found. The Large Hadron Collider is an excellent discovery machine, and represents a new frontier in high energy physics. The discovery of new particles requires the accurate identification of decay products such as leptons and jets, as well as precise measurements of particle momenta. These are both greatly enhanced by having precise measurements of particle tracks, and this thesis has contributed to these efforts by developing monitoring software for the semi-conductor tracker of the ATLAS detector. This has been tested against existing calibration data, and has already made measurements of the noise occupancy under the various conditions imposed during the SR1 cosmic commissioning tests at CERN. It was found that all SCT modules are operating within the design specification, once a few problem strips with noise occupancy greater than 500×10^{-5} have been removed. There is no evidence that the SCT picks up noise from the TRT readout cycle or that the trigger rate affects the noise occupancy. In addition, experiments with two different grounding schemes revealed no effect on the noise occupancy. The monitoring software is in continuous development, and suggestions for the improvement of the package have been provided.

If supersymmetry really does present an elegant solution to the problems of the Standard Model, then sparticles ought to be observed at the LHC. It is, however, non-trivial to take ATLAS data and reconstruct the parameters of the supersymmetric theory, or even to measure the masses of sparticles.

This thesis has tried to cultivate a general approach to SUSY parameter measurement, and has succeeded in advancing the current approach which involves studying benchmark points in models with a reduced set of parameters. The current technique has been illustrated with two model points from the mSUGRA parameter space, involving two distinct phenomenologies. It was found that there are unavoidable ambiguities in the assignment of cascade decay chains to kinematic endpoints, and furthermore that these endpoints are mostly sensitive to sparticle mass differences rather than the absolute masses.

To address these problems, a Markov Chain Monte Carlo sampling technique has been used to combine the endpoint data with a cross-section measurement. Not only does this solve the problem of setting the mass scale (because the cross-section is correlated to the mass scale), but it also allows decay chain ambiguity to be included whilst allowing the analysis to be performed in models with arbitrary numbers of parameters. This approach is in principle completely general, though it relies on having an adequate level of constraint on the parameter space through a suitable series of measurements. The work presented here may be considered a proof of principle, and it still remains to develop a series of variables, and to implement these in the presence of Standard Model background.

The final chapter of this thesis examined a SUSY model with non-universal Higgs masses, and found that the decay chain ambiguity must be extended to include the possibility of cascade chains featuring three body decays of the $\tilde{\chi}_2^0$. The kinematic endpoint positions for this decay have been calculated, and have been tested on simulated data for a sample model.

The work presented here can obviously be extended and refined further, but it has already demonstrated a powerful technique that may be used on ATLAS (and other) data over the next 10 years in order to learn as much about SUSY as possible.

Appendix A

Markov Chain Sampling

There follows a brief review of the relevant techniques involved in the Markov chain methods used in our analysis. For a more comprehensive explanation, see reference[82]).

A.1 Markov Chains

Let X_i be a (possibly infinite) discrete sequence of random variables. $X_1, X_2,...$ is said to have the *Markov property* if:

$$P(X_{i+1} = x_{i+1}|X_i = x_i, X_{i-1} = x_{i-1}, ..., X_1 = x_1) = P(X_{i+1} = x_{i+1}|X_i = x_i) \qquad (A.1)$$

for every sequence $x_1, ..., x_i, x_{i+1}$ and for every $i \geq 1$. A sequence of random variables with the Markov property is called a *Markov Chain*.

Suppose i is a discrete step in a time variable. The Markov property is then equivalent to stating that, given a present element of the sequence X_i, the conditional probability of the next element in the sequence is dependent only on the present. Thus, at each time i the future of the process is conditionally independent of the past given the present.

A.2 Sampling and probability distributions

Suppose we wish to determine a probability distribution $P(\mathbf{x})$; for example, the posterior probability of a model's parameters given some data. It is assumed in general that $\mathbf{x}$ is an N-

dimensional vector and that $P(\mathbf{x})$ can be evaluated only to within a normalisation constant Z; i.e. we can evaluate the function $P^*(\mathbf{x})$ where:

$$P(\mathbf{x}) = \frac{P^*(\mathbf{x})}{Z} \tag{A.2}$$

Although $P(\mathbf{x})$ cannot be obtained analytically, we can in theory solve the problem by sampling from $P(\mathbf{x})$ and plotting the results. Two immediate problems present themselves; the first is that Z is in general unknown. The second, which holds true even if we know Z, is that it is not obvious how to sample from $P(\mathbf{x})$ efficiently without visiting every position $\mathbf{x}$. We would like a way to visit places in $\mathbf{x}$-space where $P(\mathbf{x})$ is large in preference to places where $P(\mathbf{x})$ is small, thus giving a description of the probability distribution with a minimum of computational effort.

A.3   The Metropolis-Hastings Algorithm

The above problem can be solved through the use of Markov Chain Monte Carlo methods, one example of which is the Metropolis-Hastings algorithm. This makes use of a proposal density Q which depends on the current state of a system, which we label $\mathbf{x}^{(t)}$. (This state is really a point in a Markov Chain, and may be, for example, a particular choice of the parameters in the model whose probability distribution we are trying to sample). The density $Q(\mathbf{x}'; \mathbf{x}^{(t)})$ (where $\mathbf{x}'$ is a tentative new state, or the next point in the Markov chain) can be any fixed density from which it is possible to draw samples; it is not necessary for $Q(\mathbf{x}'; \mathbf{x}^{(t)})$ to resemble $P(\mathbf{x})$ for the algorithm to be useful, and it is common to choose a simple distribution such as a Gaussian.

Assuming that it is possible to evaluate $P^*(\mathbf{x})$ for any $\mathbf{x}$ as above, the first step in the Metropolis-Hastings algorithm is to generate a new state $\mathbf{x}'$ from the proposal density $Q(\mathbf{x}'; \mathbf{x}^{(t)})$. The decision on whether to accept the new state is made by computing the quantity:

$$a = \frac{P^*(\mathbf{x}')Q(\mathbf{x}^{(t)}; \mathbf{x}')}{P^*(\mathbf{x}^{(t)})Q(\mathbf{x}'; \mathbf{x}^{(t)})} \tag{A.3}$$

If $a \geq 1$ the new state is accepted, otherwise the new state is accepted with probability a. It is noted that if Q is a simple symmetric density, the ratio of the Q functions in Equation (A.3)

is unity, in which case the Metropolis-Hastings algorithm reduces to the Metropolis method, involving a simple comparison of the target density at the two points in the Markov Chain.

If Q is chosen such that $Q(\mathbf{x}'; \mathbf{x}) > 0$ for all $\mathbf{x}, \mathbf{x}'$, the probability distribution of $\mathbf{x}^{(t)}$ tends to $P(\mathbf{x}) = P^*(\mathbf{x})/Z$ as $t \to \infty$. Thus, by choosing points via the Metropolis algorithm and then plotting them, we have achieved our goal of obtaining a description of $P(\mathbf{x})$ in an efficient manner.

A.4   The Metropolis Algorithm as a Markov Chain Monte Carlo method

Note that the presence of the caveat $t \to \infty$ implies that there is an issue of convergence in the application of the Metropolis algorithm, and this is to be expected from the Markov Chain nature of the method. Each element in the sequence $\mathbf{x}^{(t)}$ has a probability distribution that is dependent on the previous value $\mathbf{x}^{(t-1)}$ and hence, since successive samples are correlated with each other, the Markov Chain must be run for a certain length of time in order to generate samples that are effectively independent. The exact details of convergence depend on the particular $P(\mathbf{x})$ being sampled, and on the details of Q, and hence there is some element of tuning involved in getting the algorithm to run successfully.

As a final note, it is worth stating that the purpose of definition (A.3) is to ensure that the Markov Chain used in the Metropolis method is *reversible*. By this it is meant that the chain satisfies the *principle of detailed balance*, i.e. that the probability $T(\mathbf{x}_a; \mathbf{x}_b)$ of the chain making a transition from a state $\mathbf{x}_a$ to a state $\mathbf{x}_b$ is related to the reverse transition via $T(\mathbf{x}_a; \mathbf{x}_b)P(\mathbf{x}_a) = T(\mathbf{x}_b; \mathbf{x}_a)P(\mathbf{x}_b)$. This property is necessary if we require the distribution of samples from the chain to converge to $P(\mathbf{x})$.

A.5   Convergence of Markov Chains

Assessing whether a Markov Chain has converged often involves intuition. For example, if one has a firm knowledge of the likely shape of the probability distribution being sampled (limited to the knowledge that there is only one minimum for example), one can watch the progression of the size of the region that has been sampled and obtain a reasonably good estimate of when the system has converged. This is not particularly scientific, however, and

it is useful to employ a quantitative test of convergence whilst bearing in mind that no test is completely definitive. I have chosen the relatively simple technique of observing the variation in the mean of the parameter whose convergence one wishes to test, as demonstrated below for the mSUGRA data presented in Chapter 4.

Figures 6.2 and 6.3 display the region of mSUGRA parameter space consistent with the cross-section of events with missing p_T greater than 500 GeV. Three parameters were varied in the study (m_0, $m_{1/2}$ and $\tan\beta$), and the mean of each of these parameters can be plotted against 'time'[1]. Convergence can be said to occur when the mean tends toward a constant value, though one has to assume that there is no hidden region of high probability that the sampler failed to find. If a narrow region of high probability is located far from the region that is currently being explored, the probability of the sampler reaching that region is very small. Although it will eventually be found it is not clear how to ensure that this has occurred at any given point in time. Strategies to address this problem include the use of some prior knowledge about the probability surface being explored, and the use of proposal functions that tend to increase the probability of occasionally moving through a greater distance in parameter space.

Results are shown in Figures A.1, A.2 and A.3. The sampler was started at a point far from the true point, and in each case one can observe a rapid move toward a region of higher probability, followed by a gradual convergence. The final mean value does not necessarily correspond to the correct value; in the case of positive μ this is because the endpoints are mismeasured whilst in the negative μ case there is in fact no correct mean parameter value as the model being investigated has positive μ. The same method was used to assess the convergence of all other data sets in the main body of the text.

It is noted that the case for the convergence of $\tan\beta$ is less compelling than that for the other parameters, and it is suggested that this is a reflection of the fact that the constraints used in the analysis have little dependence on $\tan\beta$. A flat distribution in $\tan\beta$ should ultimately arise from the sampling, and it takes a relatively long time to establish given the poor level of constraint. Essentially, this illustrates that a knowledge of the underlying physics is often as useful as quantitative tests in establishing the behaviour of a system near convergence.

SUSY models in general present quite a challenge to sampling algorithms, as the probability surfaces can in principle have lots of discrete minima (see, for example, reference[83]). Therefore, any convergence test must be coupled with the knowledge that the system has successfully located all minima in the problem. A crude way to do this is to start the sampler

---

[1]the time variable actually labels successive points in the Markov Chain.
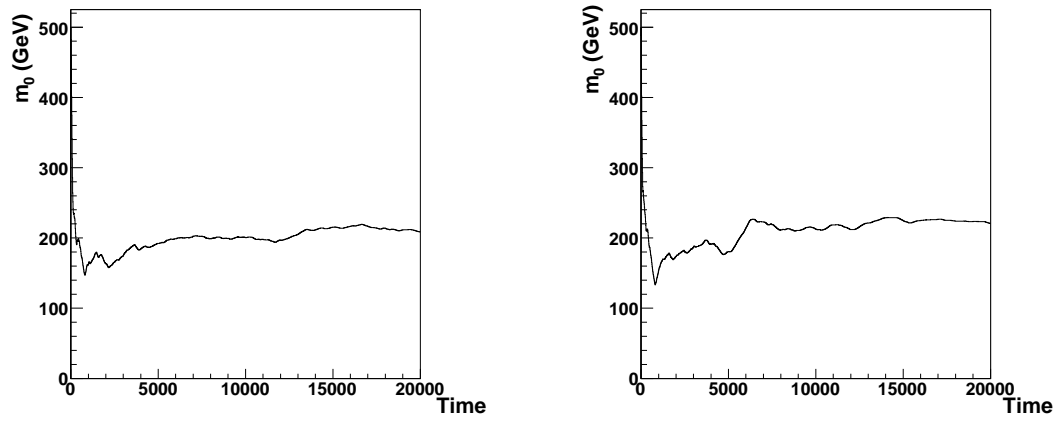
Figure A.1: The variation of the mean m_0 value with time for the mSUGRA data presented in Figures 6.2 and 6.3. Positive μ data is shown on the left, whilst negative μ data is on the right.
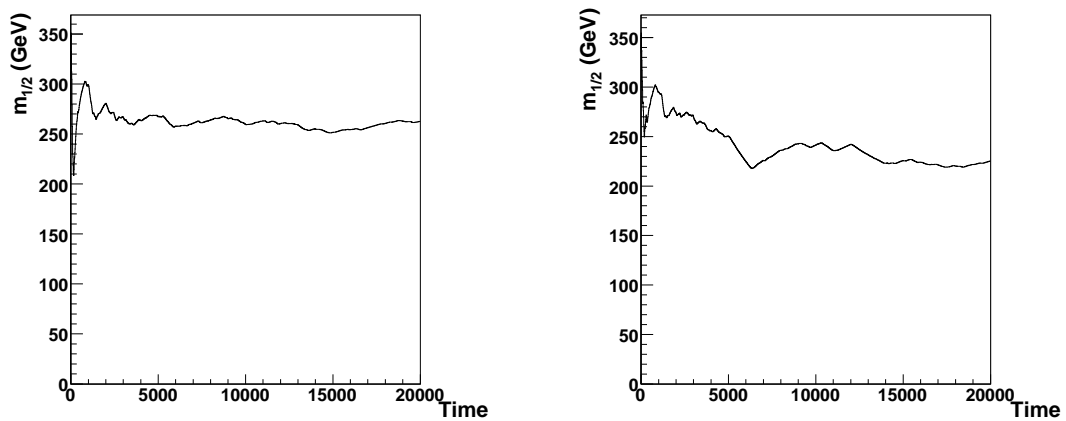
Figure A.2: The variation of the mean $m_{1/2}$ value with time for the mSUGRA data presented in Figures 6.2 and 6.3. Positive μ data is shown on the left, whilst negative μ data is on the right.
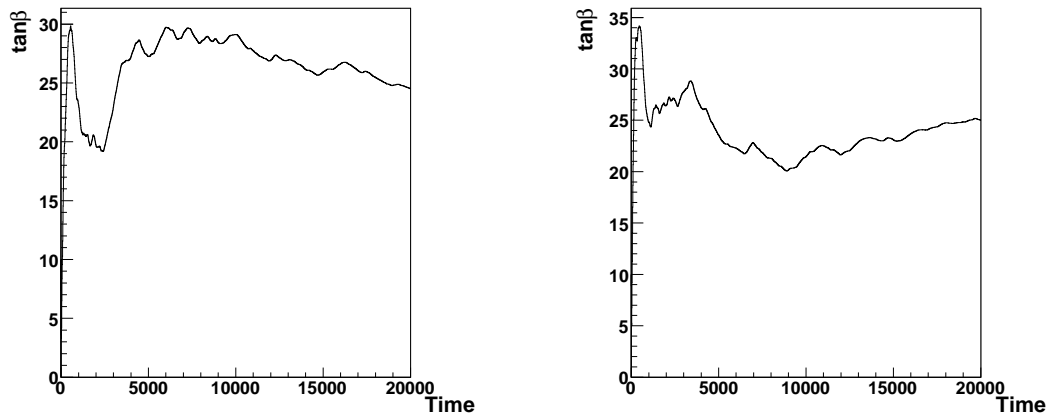
Figure A.3: The variation of the mean tanβ value with time for the mSUGRA data presented in Figures 6.2 and 6.3. Positive μ data is shown on the left, whilst negative μ data is on the right.

running in lots of different places, or to have some prior physical intuition regarding the parameter space. A general, rigorous and elegant solution to this problem is impossible, and the development of an adequate coping strategy is at the frontier of current research efforts.

Appendix B

Event Shape Variables

Events in the ATLAS detector may simply involve standard model particles and interactions, or may involve new physics. The shape of an event is related to the underlying physics causing the event, and it is often possible to obtain at least some level of discrimination between different underlying theories on the basis of event shape information. Several event shape variables have been defined in the past, and a brief list is provided here (for more details, see reference[84]). All of the these are *inclusive* variables.

B.1 Sphericity

Sphericity is constructed from the normalised tensor:

$$S_{ab} = \frac{\sum_i p_a^i p_b^i}{\sum_i |\mathbf{p^i}|^2} \tag{B.1}$$

where the sum index i runs over the particles in the event, a and b run from 1 to 3 (i.e. over the components of the three momenta of the particles), and $\mathbf{p^i}$ is the three momentum of the ith (visible) particle in the event. The three eigenvalues $\lambda_{1,2,3}$ of S_{ab} satisfy $\lambda_1 \geq \lambda_2 \geq \lambda_3 \geq 0$ and $\lambda_1 + \lambda_2 + \lambda_3 = 1$ and, since only two of them are independent, the sphericity can be defined as $S = \frac{3}{2}(\lambda_2 + \lambda_3)$. S is a measure of the summed square of transverse momenta with respect to the event axis and is 1 for an ideal spherical event and 0 for a linear event.

B.2   Thrust

Thrust is defined by the quantity:

$$T = \max \frac{\sum_i |\mathbf{p^i} \cdot \hat{\mathbf{n}}_{\mathbf{T}}|}{\sum_i |\mathbf{p^i}|} \tag{B.2}$$

where $\hat{\mathbf{n}}_{\mathbf{T}}$ is a unit vector called the 'thrust axis' whose direction is chosen to maximise T. T is in the range $\frac{1}{2} \leq T \leq 1$, with the lower end of this range corresponding to an isotropic event, and the upper end of the range representing a linear event.

B.3   Oblateness

Oblateness is a measure of the distribution of energy around the thrust axis. Starting from the thrust axis $\hat{\mathbf{n}}_{\mathbf{T}}$ defined above, one can firstly construct the 'thrust major' T_M:

$$T_M = \max_{\hat{\mathbf{n}}_{\mathbf{M}}} \frac{\sum_i |\mathbf{p^i} \cdot \hat{\mathbf{n}}_{\mathbf{M}}|}{\sum_i |\mathbf{p^i}|}, \hat{\mathbf{n}}_{\mathbf{M}} \cdot \hat{\mathbf{n}}_{\mathbf{T}} = 0 \tag{B.3}$$

where the maximisation is performed over all directions of the unit vector $\hat{\mathbf{n}}_{\mathbf{M}}$ such that $\hat{\mathbf{n}}_{\mathbf{M}} \cdot \hat{\mathbf{n}}_{\mathbf{T}} = 0$. The 'thrust minor' T_m is given by:

$$T_m = \frac{\sum_i |\mathbf{p^i} \cdot \hat{\mathbf{n}}_{\mathbf{m}}|}{\sum_i |\mathbf{p^i}|}, \hat{\mathbf{n}}_{\mathbf{m}} = \hat{\mathbf{n}}_{\mathbf{T}} \times \hat{\mathbf{n}}_{\mathbf{M}} \tag{B.4}$$

The oblateness is constructed by taking the difference between the thrust major and the thrust minor:

$$O = T_M - T_m \tag{B.5}$$

Appendix C

Development of Parallel Herwig Code

C.1   Introduction

This appendix summarises the most relevant points regarding the development of a parallel version of the `HERWIG` Monte Carlo generator, produced by adapting the existing FORTRAN code. The modifications use the Message Passing Interface (MPI) standard[71], and the code was developed and run on the Franklin machine of the Cambridge-Cranfield High Performance Computing Facility (HPCF)[85]. This is a 972 processor Sun Galaxy-class supercomputer, containing 9 nodes of 100 CPUs located in Cambridge, and one node of 72 processors located in Cranfield. Each node is a shared memory system, allowing each CPU to access all of the available memory on the node. Jobs may run on no more than 96 processors at any one time.

The actual analysis detailed in Chapter 6 required the running of the software chain `ISAJET` → `HERWIG` → `ATLFAST` at each point in parameter space, as well as code that picked points in the Markov chain and made the decision regarding whether to jump. This was accomplished by having a script that looped over points and called each program in turn. The main `HERWIG` program calls `ATLFAST` at the end of the main event loop, and hence when `HERWIG` is parallelised (and thus runs over many CPUs), `ATLFAST` is run separately on each node, and thus each instance of `ATLFAST` therefore sees only a subset of the total number of events. This speeds up the run time, even though no part of the `ATLFAST` code was explicitly altered.

It is worth noting that the development of the MPI version of `HERWIG` required the building on the HPCF of the `CERNLIB` libraries on which it depends. This required some modification of the `CERNLIB` code, owing to the fact that the preprocessor statements cause problems for the Sun FORTRAN compiler.

The following notes both introduce the MPI syntax and explain, in some detail, the modifications that were made to the various `HERWIG` subroutines. This information has been provided here since it is not documented elsewhere, and should be considered as reference material. For the original `HERWIG` manual, see reference[65].

C.2   MPI Herwig

C.2.1   Introduction to MPI

Programs are adapted for the HPCF by splitting DO loops up into blocks, and sending each block to a separate processor (also referred to here as a 'node', a term that should not be confused with the boards of 100 CPUs referred to above). Each processor has a label `myid` with the `myid = 0` node being the 'master' node, and the others termed 'slaves'. All file-handling, etc, should be performed by only one node, usually the master node. The various variables required for MPI operation are setup in a standard include file `mpi_comm.h`, which must be included in all routines that require use of MPI. Furthermore, the main program must include the following configuration lines:

```
  CALL MPI_INIT(ierr)
C Next line gets number of processors--nproc
  CALL MPI_COMM_SIZE(MPI_COMM_WORLD,nproc,ierr)
C Next line gets number of this process -- myid
  CALL MPI_COMM_RANK(MPI_COMM_WORLD,myid,ierr)
```

and the program must be ended with:

```
  CALL MPI_FINALIZE(ierr)
```

Splitting up DO loops

After initialising the MPI with the configuration lines given above, one can then define the block for each node by dividing the number of elements in the DO loop by the number of processors:

```
c the size of the block(nr) for each processor
```

```
pnk = (number of elements in do loop)/nproc
nr = ceiling(pnk)
```

Each DO loop can then be split into blocks thus:

```
do 100 N=nr*myid+1, nr*myid+nr
.
.
.
end do
```

Note that the loops are split into integer blocks, and there may be a remainder if the number of elements in the DO loop is not divisible by the number of processors. In DO loops where it is vital to evaluate all the elements, care must therefore be taken to ensure that the process with the highest value of myid evaluates all remaining elements and not just an integer block that leaves some unevaluated (this is the case for the some of the Herwig loops, and details are given below).

Message passing

After a given DO loop has been completed, the results from each node must be sent back to the master node. This is accomplished using the following syntax for message passing:

```
tag30=30
if(myid.eq.0)then
do source=1,nproc-1
      call mpi_recv(MARTINEFFERR, 1, mpi_integer,
   $            source,tag30,MPI_COMM_WORLD, status, ierr)
      if(MARTINEFFERR.ne.0)evfail=1
end do
else
      call mpi_isend(MARTINEFFERR,1,mpi_integer,
   $          0,tag30,MPI_COMM_WORLD,request,ierr)
      call mpi_wait(request,status,ierr)
endif
CALL MPI_BARRIER(MPI_COMM_WORLD,ierr)
```

Here, the master node (with `myid` equal to zero) uses the `mpi_recv` command to receive a variable `MARTINEFFERR` of type integer from each slave node. This is used to check that each node has run without a problem, and the `evfail` flag is used to veto the point if not. One must specify the number of elements being passed (in this case 1), and the 'tag' for the process (this is defined differently for each instance of the message commands in the program- I have used tag30 here, set to 30. In principle, any integer can be used.). The remaining three arguments of the function have been defined in the include statement already described, and essentially are never changed by the user (there is a slight exception in that the syntax changes for arrays, but this will be covered below).

The slave nodes use the `mpi_isend` command to send the variable back to the master process. The `mpi_wait` command ensures that the messages are all received before continuing. A call to `MPI_BARRIER` prevents the program from continuing until all messages have been received.

C.2.2   Use of MPI in Herwig

There are three distinct sections of the Herwig code that are suitable for parallelisation:

1. The main event loop.

2. The initial weight search

3. The SUSY weight calculations

The last two steps occur during the initialisation of Herwig. There follows a summary of how each of these three steps has been implemented.

The main event loop

The main event loop is the easiest to parallelise in the program, as splitting it up simply corresponds to generating a subset of events on each node (with different seeds on each node). The results of any analysis (i.e. cross-sections or cut based analysis) can be performed on each node and then sent back to the master process as required. No information needs to then be sent back from the master node to the slave nodes to ensure successful operation of the program, in contrast to the other MPI steps in the program.

In the regular Herwig program, the main event loop runs from 1 to `MAXEV` and calls each of the Herwig subroutines in turn. In the MPI program, this is split up as above, so that a block of events is sent to each processor:

```
c the size of the block(nr) for each processor
  pnk = (number of elements in do loop)/nproc
  nr = ceiling(pnk)
  do 100 N=nr*myid+1, nr*myid+nr
  .
  .
  .
  end do
```

In the analysis of Chapter 6, the program was required to calculate the number of events passing a missing p_T cut. This was worked out in the analysis routine called in the main event loop in the normal way. Therefore, after the loop has been run on each node, all that remains is to pass the number of events passing the cuts back to the master node and add the numbers up. (A tally of the events calculated on each node was also kept, to ensure that the fraction of events passing the cuts was calculated properly). In addition, the mean cross-section and the error in the cross-section from each node were returned in order to improve the accuracy of the final results. The relevant code is given below.

```
  tag=1
  tag2=2
  tag3=3
  tag4=4

  if(myid.eq. 0) then
      sum_tally = TALLY
      sum_total = TOTAL
      mean_cross_section=cross_section
      total_error=sq_error
      do 30 source = 1, nproc-1
          call mpi_recv(TALLY, 1, mpi_real, source,
     $          tag,MPI_COMM_WORLD, status, ierr)
```

```
         call mpi_recv(TOTAL, 1, mpi_real, source,
    $           tag2,MPI_COMM_WORLD,status, ierr)
         call mpi_recv(cross_section,1,mpi_double_precision,
    $           source,tag3,MPI_COMM_WORLD,status, ierr)
         call mpi_recv(sq_error,1,mpi_double_precision,
    $           source,tag4,MPI_COMM_WORLD,status, ierr)
         sum_tally = sum_tally + TALLY
         sum_total = sum_total + TOTAL
         mean_cross_section= mean_cross_section + cross_section
         total_error=total_error+sq_error
 30   continue
  else
      call mpi_isend(TALLY, 1, mpi_real, 0,
    $      tag,MPI_COMM_WORLD, requests2(1), ierr)
      call mpi_isend(TOTAL, 1, mpi_real, 0,
    $      tag2,MPI_COMM_WORLD, requests2(2), ierr)
      call mpi_isend(cross_section,1,mpi_double_precision,
    $      0,tag3,MPI_COMM_WORLD,requests2(3), ierr)
      call mpi_isend(sq_error,1,mpi_double_precision,
    $      0,tag4,MPI_COMM_WORLD,requests2(4), ierr)
      call mpi_waitall(4,requests2,status,ierr)
  endif

  if (myid.eq.0) then
      mean_cross_section=mean_cross_section/nproc
      total_error=sqrt(total_error)
  endif
```

Note the use of the `mpi_waitall` command, which requires that the `request` variable encountered previously be replaced with an array; each message uses one element of the array, and then the whole array is used in the argument of the wait all command in order to make sure that each separate message passing instance is completed before the program continues.

The initial weight search

The initial weight search in Herwig is performed just before the main event loop and sets two variables- the maximum weight `WGTMAX` and the average weight `AVWGT`. The number of points in the search is controlled by the variable `IBSH` which must therefore be divided by nproc on each node to give a subset of the total weight search. Each node then sends its own values of `WGTMAX` and `AVWGT` back to the master process which takes the highest `WGTMAX`, and averages the `AVWGT` values. These results must then be sent back to the slave processes, as they are crucial to the subsequent running of the main event loop. The code is given below:

```fortran
      temp=10000/nproc
      IBSH=ceiling(temp)

      tag5=5
      tag7=7
      if(myid.eq.0)then
          WGTMAX=0
          CALL HWEINI
          maximum_wgt=WGTMAX
          avg_wgt_early=AVWGT
          do 40 source = 1, nproc-1
              call mpi_recv(WGTMAX, 1, mpi_double_precision, source,
     $          tag5,MPI_COMM_WORLD, status, ierr)
              if(WGTMAX.gt.maximum_wgt)maximum_wgt=WGTMAX
              call mpi_recv(AVWGT, 1, mpi_double_precision, source,
     $          tag7,MPI_COMM_WORLD, status, ierr)
              avg_wgt_early=avg_wgt_early+AVWGT
 40       continue
          WGTMAX=maximum_wgt
          AVWGT=avg_wgt_early/nproc
      else
          WGTMAX=0
          CALL HWEINI
          call mpi_isend(WGTMAX, 1, mpi_double_precision, 0,
     $        tag5,MPI_COMM_WORLD, requests(1), ierr)
          call mpi_isend(AVWGT, 1, mpi_double_precision, 0,
```

```
   $         tag7,MPI_COMM_WORLD, requests(2), ierr)
      call mpi_waitall(2,requests,status,ierr)
  end if


  tag6=6
  tag8=8
  if(myid.eq.0)then
      do 50 source = 1, nproc-1
          call mpi_isend(WGTMAX, 1, mpi_double_precision, source,
   $          tag6,MPI_COMM_WORLD, requests(1), ierr)
          call mpi_isend(AVWGT, 1, mpi_double_precision, source,
   $          tag8,MPI_COMM_WORLD, requests(2), ierr)
          call mpi_waitall(2,requests,status,ierr)
 50     continue
  else
      call mpi_recv(WGTMAX, 1, mpi_double_precision, 0,
   $          tag6,MPI_COMM_WORLD, status, ierr)
      call mpi_recv(AVWGT, 1, mpi_double_precision, 0,
   $          tag8,MPI_COMM_WORLD, status, ierr)
  end if
```

The SUSY weight calculations

Herwig needs to go through each of the 3 and 4 body SUSY decay modes and calculate a weight for each. The exact number of modes varies depending on which point in the mSUGRA parameter space is chosen, but there is no reason why different weights cannot be worked out on different processors. Hence, the decay weight calculations were 'parallelised' by simply getting different nodes to work out different weights before the master process collects them all and sends the full table of weights back to the slaves.

The weight calculations are implemented in the HERWIG routine HWISP3. Care must be taken to ensure that the entire DO loop is evaluated, as otherwise the weights in the end of the loop are not calculated and Herwig will definitely crash. A further complication arises from the fact that the number of processors can often be more than the number of elements in the loop, so the code needs to check the number of processors and run one of two sets of code.

Consider the three body modes first. The variable N3MODE stores the number of these

modes, such that the DO loop involving the weight calculation runs from 1 to N3MODE. The first thing to do is to work out the size of the block sent to each node:

```
if(nproc.lt.N3MODE)then
      decay_pnk=N3MODE/nproc
      decay_nr=ceiling(decay_pnk)
else
      decay_nr=1
end if
```

This code ensures that when the number of processors is greater than the number of modes, there is only one mode per node. The original Herwig source code subsequently has a DO loop that defines some variables and calls a routine to work out the weight for each mode. This must be called in subtly different ways depending on both the number of processors and, in the case where this number is less than the number of modes, the value of myid of a particular node. Take first the case where the number of processors is greater than the number of modes:

```
if(nproc.gt.N3MODE)then
if(myid.lt.N3MODE)then
DO 3000 I=myid+1,myid+1
IF(RSPIN(IDK(ID3PRT(I))).EQ.ZERO) THEN
      RHOIN(1,1) = ONE
      RHOIN(1,2) = ZERO
      RHOIN(2,1) = ZERO
      RHOIN(2,2) = ZERO
ELSE
      RHOIN(1,1) = HALF
      RHOIN(1,2) = ZERO
      RHOIN(2,1) = ZERO
      RHOIN(2,2) = HALF
ENDIF
PHEP(5,1) = RMASS(IDK(ID3PRT(I)))
PHEP(4,1) = SQRT(100.0D0**2+PHEP(5,1)**2)
PHEP(1,1) = 100.0D0
PHEP(2,1) = 0.0D0
PHEP(3,1) = 0.0D0
```

```
   IF(myid.gt.N3MODE)WRITE(6,*) "MYID OVER LIMIT!",myid
3000 CALL HWD3ME(1,0,I,RHOIN,1)
   end if
   (IF.NOT.SUSYIN) RETURN
   CALL MPI_BARRIER(MPI_COMM_WORLD,ierr)
```

Thus, the nodes with `myid` less than `N3MODE` each work out one weight, whilst those with `myid` greater than or equal to `N3MODE` remain idle. (An MPI barrier is required to ensure that the processes all meet up at some point).

Now consider the case where the number of processors is less than the number of modes (this does occur with some regularity):

```
   else
   if(myid.lt.nproc-1)then
   DO 3001 I=decay_nr*myid+1,decay_nr*myid+decay_nr
   IF(RSPIN(IDK(ID3PRT(I))).EQ.ZERO) THEN
   .
   .
   .
3001 CALL HWD3ME(1,0,I,RHOIN,1)
   else
   if(myid.eq.nproc-1)then
   DO 3002 I=decay_nr*myid+1,N3MODE
   IF(RSPIN(IDK(ID3PRT(I))).EQ.ZERO) THEN
   .
   .
   .
3002 CALL HWD3ME(1,0,I,RHOIN,1)
   end if
   end if
   end if
   IF(.NOT.SUSYIN) RETURN
   CALL MPI_BARRIER(MPI_COMM_WORLD,ierr)
```

This code executes the same loop as before but makes sure that the process with the highest `myid` works out all the weights up to `N3MODE`. An MPI barrier is used again to make sure that

none of the nodes goes on ahead of the others before the weights are calculated.

Having now obtained values for all of the weights, they must be sent back to the master process and then returned to the slaves. The weights are stored in an array called WT3MAX which is initialised in the Herwig include file to contain 500 elements. In any given run of Herwig, however, there are only N3MODE weights, and so the whole array does not to be sent back and forth, only the full elements. Furthermore, when sending weights back to the master process, one only needs to send those that have been calculated on each node which will only be a few elements of the total array, with the exact number depending on the number of modes and the number of processors.

The following code sends the array back to the master process. The master process has an array called MASTER_WT3MAX which it uses to collect all of the weights from the different processes, and an array MOBILE_WT3MAX which has less elements than WT3MAX is used to pass things back and forth:

```
  tag10=10
  if(myid.eq.0)then
      do I=1, decay_nr
          MASTER_WT3MAX(I)=WT3MAX(I)
      end do
  end if


C-----Now need to do one of two things depending on whether nproc>N3MODE

  if(nproc.le.N3MODE)then
  if(myid.ne.0)then
      if(myid.le.nproc-1)then
      do I=1,N3MODE
          MOBILE_WT3MAX(I)=WT3MAX(I)
      end do
      else
      if(myid.eq.nproc-1)then
          do I=1,N3MODE
              MOBILE_WT3MAX(I)=WT3MAX(I)
          end do
```

```
         end if
         end if
         call mpi_isend(MOBILE_WT3MAX,150,MPI_DOUBLE_PRECISION,
     $        0,tag10,MPI_COMM_WORLD,request,ierr)
         call mpi_wait(request,status,ierr)
   else
         do 911 source = 1, nproc-1
         call mpi_recv(MOBILE_WT3MAX,150, mpi_double_precision,
     $        source,tag10,MPI_COMM_WORLD, status, ierr)
         if(source.lt.nproc-1)then
             do I=decay_nr*source+1, decay_nr*source+decay_nr
                MASTER_WT3MAX(I)=MOBILE_WT3MAX(I)
             end do
         else
             do I=decay_nr*source+1, N3MODE
                MASTER_WT3MAX(I)=MOBILE_WT3MAX(I)
             end do
         end if
911    continue
   end if
   else
         if(myid.ne.0.and.myid.lt.N3MODE)then
             wt3=WT3MAX(myid+1)
             call mpi_isend(wt3,1,MPI_DOUBLE_PRECISION,
     $        0,tag10,MPI_COMM_WORLD,request,ierr)
             call mpi_wait(request,status,ierr)
         else
             if(myid.eq.0)then
             do 511 source = 1, N3MODE-1
                call mpi_recv(wt3,1, mpi_double_precision,
     $            source,tag10,MPI_COMM_WORLD, status, ierr)
                MASTER_WT3MAX(source+1)=wt3
511    continue
         end if
   end if
   end if
```

This code remains slightly wasteful (150 numbers are sent back in some cases rather than just the full elements of the array), and hence it could be improved further. The code above has got as far as receiving all the 3 body weights on the master process and storing them in an array (`MASTER_WT3MAX`). This must now be sent back to the slaves in order to complete the process:

```
C Now send WT3MAX back to the slave processes
   tag11=11
   if(myid.eq.0)then
        do I=1,N3MODE
            MOBILE_WT3MAX(I)=MASTER_WT3MAX(I)
            WT3MAX(I)=MASTER_WT3MAX(I)
        end do


        do 512 source = 1,nproc-1
            call mpi_isend(MOBILE_WT3MAX,150,MPI_DOUBLE_PRECISION,
     $         source,tag11,MPI_COMM_WORLD, request, ierr)
            call mpi_wait(request,status,ierr)
 512    continue
   else
        call mpi_recv(MOBILE_WT3MAX,150, mpi_double_precision,0,
     $      tag11,MPI_COMM_WORLD, status, ierr)
        do I=1,N3MODE
            WT3MAX(I)=MOBILE_WT3MAX(I)
        end do
   end if
```

This concludes the MPI code for the 3 body modes. The implementation for the 4 body modes is very similar except that there are `NBMODE` modes, and these are stored in an array called `WTBMAX`.

Colophon

This thesis was made in LaTeX 2_ε using the "hepthesis" class[86].

Bibliography

[1] Michael E. Peskin and Daniel V. Schroeder. *An Introduction to Quantum Field Theory*. Addison-Wesley, Reading, MA, 1995.

[2] Steven Weinberg. A model of leptons. *Phys. Rev. Lett.*, 19:1264–1266, 1967.

[3] S. L. Glashow, J. Iliopoulos, and L. Maiani. Weak interactions with lepton - hadron symmetry. *Phys. Rev.*, D2:1285–1292, 1970.

[4] Salam. A. In N. Svartholm, editor, *Elementary Particle Theory*, page 367. Almquist and Wiksells, Stockholm, 1969.

[5] Peter W. Higgs. Broken symmetries and the masses of gauge bosons. *Phys. Rev. Lett.*, 13:508–509, 1964.

[6] Benjamin W. Lee, C. Quigg, and H. B. Thacker. The strength of weak interactions at very high-energies and the Higgs boson mass. *Phys. Rev. Lett.*, 38:883, 1977.

[7] N. Cabibbo, L. Maiani, G. Parisi, and R. Petronzio. Bounds on the fermions and Higgs boson masses in grand unified theories. *Nucl. Phys.*, B158:295, 1979.

[8] Thomas Hambye and Kurt Riesselmann. Matching conditions and Higgs mass upper bounds revisited. *Phys. Rev.*, D55:7255–7262, 1997.

[9] Gino Isidori, Giovanni Ridolfi, and Alessandro Strumia. On the metastability of the standard model vacuum. *Nucl. Phys.*, B609:387–409, 2001.

[10] **LEP** Electroweak Working Group. A combination of preliminary electroweak measurements and constraints on the standard model. `hep-ex/0511027`, 2005.

[11] R. Barate et al. Search for the standard model Higgs boson at LEP. *Phys. Lett.*, B565:61–75, 2003.

[12] S. Eidelman et al. Supersymmetry. In *Review of particle physics*, volume B592. 2004.

[13] Stephen P. Martin. A supersymmetry primer. `hep-ph/9709356`, 1997.

[14] Steven Weinberg. *Supersymmetry*, volume 3 of *The Quantum Theory of Fields*. Cambridge University Press, Cambridge, England, 2005.

[15] Sidney R. Coleman and J. Mandula. All possible symmetries of the S matrix. *Phys. Rev.*, 159:1251–1256, 1967.

[16] R. Haag, J. Lopuszanski, and M. Sohnius. All possible generators of supersymmetries of the S matrix. *Nucl. Phys.*, B88:257, 1975.

[17] G. L. Kane, Chris Kolda, and James D. Wells. Calculable upper limit on the mass of the lightest higgs boson in any perturbatively valid supersymmetric theory. *Physical Review Letters*, 70:2686, 1993.

[18] J. R. Espinosa and M. Quiros. Upper bounds on the lightest higgs boson mass in general supersymmetric standard models. *Physics Letters B*, 302:51, 1993.

[19] Stuart Raby. Proton decay. `hep-ph/0211024`, 2002.

[20] D. N. Spergel et al. First year Wilkinson Microwave Anisotropy Probe (WMAP) observations: Determination of cosmological parameters. *Astrophys. J. Suppl.*, 148:175, 2003.

[21] Marco Battaglia et al. Proposed post-LEP benchmarks for supersymmetry. *Eur. Phys. J.*, C22:535–561, 2001.

[22] M. Battaglia et al. Updated post-WMAP benchmarks for supersymmetry. *Eur. Phys. J.*, C33:273–296, 2004.

[23] J Ellis, K A Olive, Y Santoso, and V C Spanos. Supersymmetric dark matter in light of WMAP. *Physics Letters B*, 565:176, 2003.

[24] The Tevatron web page. `http://www-bdnew.fnal.gov/tevatron/`.

[25] **ATLAS** Collaboration. *ATLAS detector and physics performance Technical Design Report*. CERN/LHCC 99-15, 1999.

[26] R. Bonciani, S. Catani, M. L. Mangano, and P. Nason. NLL resummation of the heavy-quark hadroproduction cross-section. *Nuclear Physics B*, 529:424, 1998.

[27] Szymon Gadomski. ATLAS inner detector performance, 2000.

[28] **ATLAS** Collaboration. *ATLAS Inner Detector Technical Design Report.* CERN/LHCC/97-16 and CERN/LHCC/97-17, 1997.

[29] **ATLAS** Collaboration. *ATLAS Pixel Detector Technical Design Report.* CERN/LHCC/98-13, 1998.

[30] C. Cemme. The ATLAS pixel detector. *Nucl. Instr. and Meth.*, A501:87–92, 2003.

[31] A. Carter et al. `http://atlas.web.cern.ch/Atlas/GROUPS/INNER_DETECTOR/` `SCT/module/SCTbarrelmod.html`.

[32] V. A. Mitsou and for the ATLAS TRT collaboration. The atlas transition radiation tracker, 2003.

[33] S Baranov, M Bosman, I Dawson, V Hedberg, A Nisati, and M Shupe. Estimation of radiation background, impact on detectors, activation and shielding optimization in ATLAS, 2005.

[34] S. Eidelman et al. Review of particle physics. *Phys. Lett.*, B592:1, 2004.

[35] **ATLAS** Collaboration. *Calorimeter performance technical design report*. CERN/LHCC 96-40, 1996.

[36] **ATLAS** Collaboration. *ATLAS Muon Spectrometer Technical Design Report.* CERN/LHCC 97-22, 1997.

[37] The barrel modules of the ATLAS SemiConductor Tracker. *Collaboration note*, ATL-INDET-PUB-2006-005.

[38] Lee Martin Drage. *The Evaluation of Silicon Microstrip Detectors for the ATLAS Semiconductor Tracker and Supersymmetry Studies at the Large Hadron Collider*. PhD thesis, University of Cambridge, 1999.

[39] Robinson D. et al. Silicon microstrip detectors for the ATLAS SCT. *Nucl. Instr. Methods A.*, 485:84, 2002.

[40] Silicon detector PRR papers. `http://edmsoraweb.cern.ch`, 2000.

[41] A. Carter et al. SCT barrel module: Requirements and specifications. SCT-BM-FDR.

[42] F. Campabadal et al. Design and performance of the ABCD3TA ASIC for readout of silicon strip detectors in the ATLAS semiconductor tracker. *Nucl. Instrum. Meth.*, A552:292–328, 2005.

[43] 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.*, 32:530–534, 1998.

[44] M. Dentan et al. RD29 final status report. CERN-LHCC-98-037, 1998.

[45] Alan Barr. *Studies of supersymmetry models for the ATLAS experiment at the Large Hadron Collider*. PhD thesis, University of Cambridge, 2002. CERN-THESIS-2004-002.

[46] M. J. Palmer. *Studies of Extra-Dimensional Models with the ATLAS Detector*. PhD thesis, University of Cambridge, 2005.

[47] Gareth Moorhead Lars Eklund, John Hill and Peter Phillips. SCTDAQ. `http://map.web.cern.ch/sct-testdaq/sctdaq/sctdaq.html`.

[48] A. J Barr, B. J. Gallop, C. G. Lester, M. J. Palmer, D. Robinson, and A. Tricoli. The SctRodDaq website. `www.hep.phy.cam.ac.uk/~atlasdaq/`.

[49] P. W. Phillips. System performance of ATLAS SCT detector modules. Prepared for 8th Workshop on Electronics for LHC Experiments, Colmar, France, 9-13 Sep 2002, on behalf of the ATLAS SCT collaboration.

[50] B. J. Gallop. *ATLAS SCT Readout and Barrel Macro-assembly Testing and Design of MAPS Test Structures*. PhD thesis, University of Birmingham, 2005.

[51] P. W. Phillips. Functional testing of the ATLAS SCT barrels. Prepared for the 7th International Conference on Large Scale Applications and Radiation Hardness of Semiconductor Detectors, Florence, Italy, 5-7 October 2005, on behalf of the ATLAS SCT collaboration.

[52] Alan Barr. Calibrating the ATLAS semiconductor tracker front end electronics. *IEEE NSS Conference Records*, ATL-INDET-CONF-2006-001:1192–1195, 2004.

[53] T. Dubbs et al. Noise determination in silicon micro strips. *IEEE Trans. Nucl. Sci.*, 43:1119–1122, 1996.

[54] The SCT Monitoring wiki page. `https://twiki.cern.ch/twiki/bin/view/Atlas/SCTMonitoring`.

[55] COOL - LCG conditions database project. `http://lcgapp.cern.ch/project/CondDB/`.

[56] ATLAS tracker takes major strides. *CERN Courier*, June, 2006.

[57] Heinz Pernegger. ID barrel testing in SR1. Talk given at the International Workshop on Alignment and Commissioning of the ATLAS Inner Detector, Ringberg Castle, Germany, June 12-16 2006.

[58] Reisaburo Tanaka. SCT noise occupancy. Talk given at the ATLAS ID Commissioning meeting, 6th July 2006 (`http://indico.cern.ch/conferenceDisplay.py?confId=a0625`).

[59] Vasiliki A. Mitsou. Production performance of the ATLAS semiconductor tracker readout system. *IEEE Trans. Nucl. Sci.*, 53:729–734, 2006.

[60] A. J. Barr. Measuring slepton spin at the LHC. *JHEP*, 02:042, 2006.

[61] D. R. Tovey. Measuring the SUSY mass scale at the LHC. *Phys. Lett.*, B498:1–10, 2001.

[62] Frank E. Paige, Serban D. Protopescu, Howard Baer, and Xerxes Tata. ISAJET 7.69: A Monte Carlo event generator for p p, anti-p p, and e+ e- reactions. `hep-ph/0312045`, 2003.

[63] G. Corcella et al. HERWIG 6: An event generator for hadron emission reactions with interfering gluons (including supersymmetric processes). *JHEP*, 01:010, 2001.

[64] G. Corcella et al. HERWIG 6.5 release note. `hep-ph/0210213`, 2002.

[65] Stefano Moretti, Kosuke Odagiri, Peter Richardson, Michael H. Seymour, and Bryan R. Webber. Implementation of supersymmetric processes in the HERWIG event generator. *JHEP*, 04:028, 2002.

[66] Christopher Gorham Lester. *Model Independant sparticle mass measurements at ATLAS*. PhD thesis, University of Cambridge, 2001.

[67] B. K. Gjelsten, D. J. Miller, and P. Osland. Measurement of SUSY masses via cascade decays for SPS 1a. *JHEP*, 12:003, 2004.

[68] Gianluca Comune. Co-annihilation model parameter measurements. Talk given to the CERN-ATLAS Supersymmetry Working Group, 16 February 2005 (`http://agenda.cern.ch/fullAgenda.php?ida=a05518`).

[69] G. Belanger, F. Boudjema, A. Pukhov, and A. Semenov. micrOMEGAs: A program for calculating the relic density in the MSSM. *Comput. Phys. Commun.*, 149:103–120, 2002.

[70] W. de Boer, C. Sander, V. Zhukov, A. V. Gladyshev, and D. I. Kazakov. The supersymmetric interpretation of the egret excess of diffuse galactic gamma rays. *Physics Letters B*, 636:13, 2006.

[71] Message Passing Interface Forum. MPI: A Message-Passing Interface standard. http://www.mpi-forum.org/docs/mpi-11-html/mpi-report.html.

[72] Howard Baer and Csaba Balazs. χ^2 analysis of the minimal supergravity model including WMAP, $g_\mu - 2$ and $b \to s\gamma$ constraints. *JCAP*, 0305:006, 2003.

[73] Edward A Baltz and Paolo Gondolo. Markov Chain Monte Carlo exploration of minimal supergravity with implications for dark matter. *JHEP*, 0410:052, 2004.

[74] B. C. Allanach and C. G. Lester. Multi-dimensional mSUGRA likelihood maps. *Phys. Rev.*, D73:015013, 2006.

[75] A. De Roeck et al. Supersymmetric benchmarks with non-universal scalar masses or gravitino dark matter. `hep-ph/0508198`, 2005.

[76] Howard Baer, Azar Mustafayev, Stefano Profumo, Alexander Belyaev, and Xerxes Tata. Direct, indirect and collider detection of neutralino dark matter in SUSY models with non-universal higgs masses. *JHEP*, 07:065, 2005.

[77] Frank E. Paige and Serban D. Protopopescu. *ISAJET 7.72: A Monte Carlo event generator for pp, $\bar{p}p$ and e^+e^- reactions.* `http://www.phy.bnl.gov/~isajet/`, 2005.

[78] I. Hinchliffe, F. E. Paige, M. D. Shapiro, J. Soderqvist, and W. Yao. Precision SUSY measurements at LHC. *Phys. Rev.*, D55:5520–5540, 1997.

[79] Christopher G. Lester. Constrained invariant mass distributions in cascade decays: The shape of the 'm(qll)-threshold' and similar distributions. `hep-ph/0603171`, 2006.

[80] B. K. Gjelsten, D. J. Miller, and P. Osland. Measurement of the gluino mass via cascade decays for SPS 1a. *JHEP*, 06:015, 2005.

[81] S. Chouridou, R. Ströhmer, and T . Trefzger. Study of the three body matrix element of the neutralino decay $\tilde{\chi}_2^0 \rightarrow \tilde{\chi}_1^0 l^+ l^-$. *ATLAS Physics note*, pages ATL–COM–PHYS–2002–032, 2003.

[82] David J. C. MacKay. *Information Theory, Inference and Learning Algorithms*. Cambridge University Press, 2003.

[83] Nima Arkani-Hamed, Gordon L. Kane, Jesse Thaler, and Lian-Tao Wang. Supersymmetry and the LHC inverse problem. `hep-ph/0512190`, 2005.

[84] Barger V. D. and Phillips R. J. N. *Collider Physics*. Addison-Wesley, USA, 1987.

[85] The Cambridge-Cranfield High Performance Computing Facility. `http://www.hpcf.cam.ac.uk/`.

[86] Andy Buckley. The hepthesis LaTeX class.

List of Figures

List of Tables