Slepton Spin Determination and Simulation of the Transition Radiation Tracker at the ATLAS Experiment

Thesis submitted for the degree of
Ph.D. in Physics

by

Thomas Høvring Kittelmann

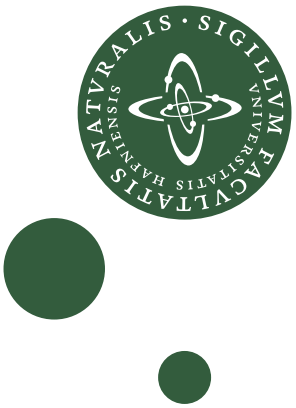

January 8, 2007

Experimental High Energy Physics Group
Niels Bohr Institute
University of Copenhagen

Abstract

At the end of 2007, the Large Hadron Collider, a proton-proton collider delivering collisions at unprecedented centre of mass energies and luminosity, will be put into service at CERN. Two multi-purpose experiments, ATLAS and CMS will collect data, suitable for uncovering new physics, believed to exists at scales just outside the reach of present experiments.

After a thorough review of the present theoretical status of particle physics, and an outline of the capabilities of the ATLAS experiment, the thesis will discuss the detector simulation of the ATLAS Transition Radiation Tracker (TRT) in great detail. The TRT is a sub-detector, which provides tracking information through a large number of thin drift tubes, in addition to yielding information suitable for electron identification. The constructed simulation algorithms are to a large extend based on first-principles physics concepts, and following their tuning and validation against data taken in a test beam at CERN in 2004, they can be expected to describe the complex setup of the full detector installed as part as ATLAS reliably.

The thesis presents a method for measuring the spin of new physics particles with leptonic quantum numbers, which might allow for a discrimination between models of new physics that otherwise have similar experimental signatures. The method is tested in the context of supersymmetry, and the potential for distinguishing such models from models with universal extra dimensions is explicitly addressed. Issues such as event selection, background treatment and efficiency correction are discussed in detail.

In loving memory of
Anna-Lise Kittelmann Pedersen
(1919 – 2005)

Contents

List of Figures

List of Tables

There is a theory which states that if ever anyone discovers exactly what the Universe is for and why it is here, it will instantly disappear and be replaced by something even more bizarre and inexplicable.

There is another which states that this has already happened.

Douglas Adams, in "The Restaurant at the End of the Universe"

1 Introduction

The field of particle physics has come a long way since J. J. Thomson's discovery of the electron in 1897. After the development of Quantum Mechanics and the unravelling of the physics of the atom in the first half of the twentieth century, the discovery of a plethora of particle states in the earliest particle accelerators during the 1950s lead to a series of theoretical endeavours, which culminated in the early 1970s with the formulation of the model now known as the Standard Model of Particle Physics (SM). In this model, the description of matter and all known forces of Nature except gravity, is incorporated in the framework of Quantum Field Theory (QFT), which itself combines and extends Special Relativity and Quantum Mechanics. The model requires three generations of matter particles, three types of force mediating bosons and one so-called Higgs boson to account for the particle masses.

With the single notorious exception of the Higgs boson, all of these particles have since been observed in direct accelerator based searches and their properties found to be in excellent agreement with the SM predictions. Despite its astounding success, however, a number of cosmological and theoretical issues remain unresolved by this model, and there are several indications that it will break down at energy scales just above the reach of present experiments. The exact nature of the expanded model which should replace the SM is unknown, but it might involve either supersymmetry (SUSY), extra spacial dimensions, or something even more exotic.

In less than a year from now, a monumental new circular collider will be put into service at CERN, which will make it possible to explore physics at this exciting scale. This machine, the Large Hadron Collider (LHC), will ultimately accelerate protons to an energy of 7 TeV, and bring them into head-on collisions at unprecedented luminosities of 10^{34} cm^{-2}s^{-1}, corresponding to 100 fb^{-1} of recorded data per year (with initial luminosities an order of magnitude lower). The accelerator is placed in the 27 km long tunnel formerly used by the LEP experiment (cf. figure 1.1), and the high energies are made possible by controlling the beam with strong 8 T magnetic fields, provided by superconducting dipole magnets around the ring. The high luminosities will be achieved by using closely packed proton bunches, interspersed by 25 ns, and a high average number of collisions at each crossing (~ 25).

Four major experiments will be collecting data at the LHC during its operation, as illustrated in figure 1.2. The largest of these, ATLAS and CMS, are designed as multi-purpose detectors, able to reliably and redundantly observe and analyse new physics based on a wide range of experimental signatures. The high event rates impose strong demands upon the technologies used in these experiments, and furthermore results in tremendous rates of data, that must be filtered and, if appropriate, stored for later analysis. The efficient handling of such data amounts would certainly not be possible if not for the explosive developments in processing power, network speed and storage capacity seen during the last decades. Even taking this progress into account, however, the handling of the filtered data output of $\mathcal{O}$(PB/year), is a non-trivial task and has required large development efforts into custom GRID technologies.

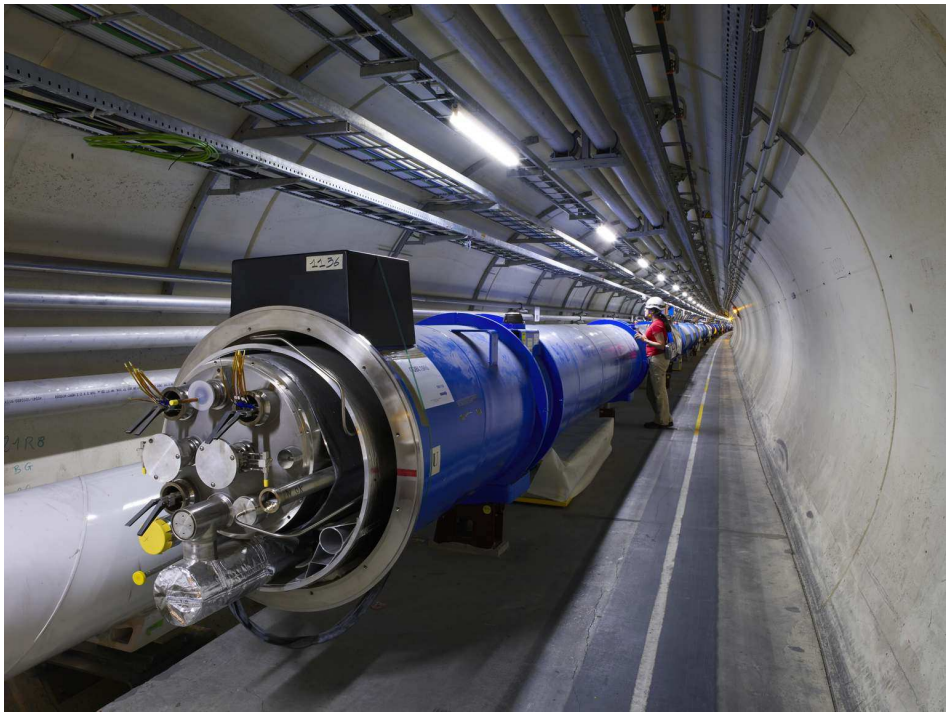

Figure 1.1: *Photograph from the LHC tunnel, where the elements of the collider ring are currently being assembled. Source: CERN Press Office.*

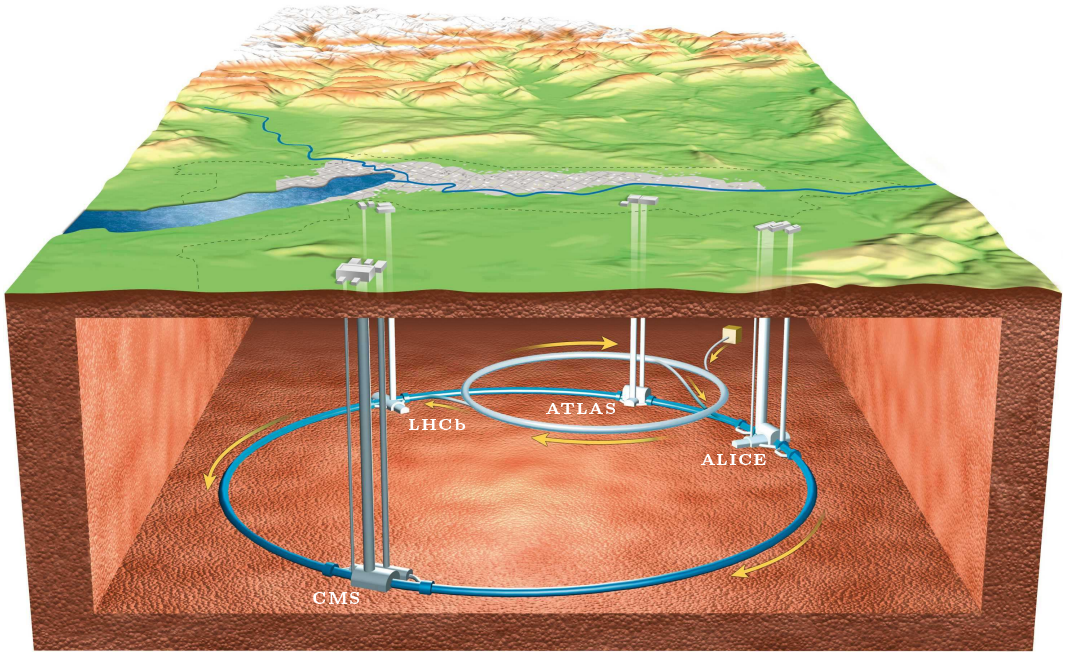

Figure 1.2: *Schematic overview of the LHC ring, showing the locations of the four main experiments ATLAS, LHCb, CMS, and ALICE which will collect data when it is running. Behind LHC the location of the city of Geneva is indicated. Source: CERN Press Office.*

The first part of the thesis is dedicated to a thorough review of the present theoretical status of particle physics. After a review of the SM and its framework, QFT, its present status and problems will be considered. Proposed extensions to the model, addressing these problems are discussed, along with their experimental consequences at the LHC. The second part consists of a relatively short overview of the ATLAS experiment, briefly outlining the capabilities of the detector and the utilised technologies.

The last parts of the thesis describes the Ph.D. work done by the author. Part three deals with the simulation of the response of the Transition Radiation Tracker (TRT), an ATLAS sub-detector providing tracking information through a large number of thin drift tubes, in addition to yielding information suitable for electron identification, based on the transition radiation emitted by ultra relativistic particles with gamma factors above ~ 1000. The discussion will focus on the simulation of the specific detector response and front-end electronics of the TRT, since those are the areas requiring specialised treatment in each sub-detector. In addition to the physics processes taking place in the active gas, a model is developed which takes into account the effects of noise in the electronics, as well as the fact that there are significant variation in performance between individual channels. Data from a test beam at CERN in 2004 are used along the way to tune and validate the performance of the implemented algorithms.

Finally, the fourth part of the thesis presents a method which can potentially be used to get a handle on the spin states of new physics particles with leptonic quantum numbers, and thus provide a way of distinguishing between models of new physics that might otherwise have identical experimental signatures. Issues such as event selection, background treatment and efficiency correction will described. Models investigated are SUSY and models with universal extra dimensions (UED). The method is tested in the scenario of new physics being SUSY, and, based on physics simulations and fast toy Monte Carlo studies, it is explicitly tested how well the hypothesis of UED could be ruled out in such a case.

Part I
Theory

The Standard Model of Particles Physics (SM) in a sense represents the pinnacle of the present understanding of nature at its most fundamental level. It unifies Quantum Mechanics with the Special Theory of Relativity while accomplishing the daunting and somewhat botanical exercise of describing the known forces and constituents of all known matter. Apart from failing to incorporate General Relativity there are, however, good reasons to believe that it is incomplete, and that the first hints of physics beyond it will be observed at the Large Hadron Collider (LHC) at CERN within the next couple of years.

The aim of this first part of the thesis is to summarise the SM along with its experimental and theoretical status and to discuss its possible TeV scale extensions and their phenomenological consequences at the LHC.

In order to understand the consequences and limitations imposed by the theoretical framework on the SM, the thesis will start out with a very cursory review of the general properties of Quantum Field Theories (QFT) in section 2. This is in no way meant to be a formal introduction of such, but is merely meant as a summary of their general features. Readers interested in more comprehensive and rigorous treatments are referred to the excellent literature on the subject such as [PS95], [HM84] and [Pet94] as well as references within the text.

Having thus set the stage, the SM itself is presented in section 3. This is again intended to be an overview, and detailed formulae are not given unless needed for clarification. In this vein section 4 proceeds with a discussion of the present experimental and theoretical status of the SM, and section 5 tries to summarise several of the more popular proposals for its extension at the TeV scale. Finally section 6 is devoted to a short overview of how these extensions would be observable at the LHC. As the analysis in the final part of this thesis is based on a Supersymmetric extension of the SM, slightly more emphasis will be put on this possibility in the two latter sections.

2 A Cursory Review of Properties of Quantum Field Theories

In a very general sense it has to be true that any particle physics theory must be able to describe the probability to find a certain final state, $|f\rangle$, given an initial state, $|i\rangle$. This probability can be expressed by the scattering amplitudes in the so-called S-matrix,

$$S_{fi} = \langle f| S |i\rangle \tag{2.1}$$

Unitarity of the S-matrix is a fundamental requirement of any sensible theory. Its dynamical part is usually separated out from the non-interacting part and four-momentum conservation into the T-matrix, defined by:

$$S_{fi} = \delta_{fi} + i(2\pi)^4 \delta^4(p_f - p_i) \langle f| T |i\rangle \tag{2.2}$$

When it is understood from the context what process is under consideration, the matrix element $\langle f| T |i\rangle$ is usually denoted by $\mathcal{M}$, and its calculation is the principal objective of any theory of particle physics.

The Standard Model is a Quantum Field Theory, which means that the basic entities described by the theory are not particles or classical fields but instead quantised fields. Whereas a classical theory of particles is quantised by turning the parameters of a particle such as $\mathbf{x}$ and $\mathbf{p}$ into operators, $\hat{\mathbf{x}}(t)$ and $\hat{\mathbf{p}}(t)$, satisfying commutation relations, a classical field, $\phi(\mathbf{x}, t)$, is instead quantised in coordinate space by turning the field itself into a position dependent

operator, $\hat{\phi}(\mathbf{x}, t)$.[1] Just as for the harmonic oscillator potential in standard Quantum Mechanics, one can then define creation and annihilation operators, in this case acting at each point $\hat{a}^\dagger(\mathbf{x}, t)$ and $\hat{a}(\mathbf{x}, t)$. The "particles" in QFT are thus manifested as excitations above the ground states (the vacua) of a number of field operators.

2.1 Lagrangian Formalism and Method of Feynman Diagrams

Actual work in the theory is performed within the Lagrangian formalism.[2] Classically one would write down the Lagrangian (density) for a system and then find the solutions by minimising the corresponding *action* by solving the so-called Euler-Lagrange differential equation. Analogously the actual content of a specific quantum field theory will be determined by the terms in the Lagrangian. Consider the following example dealing with a simple scalar field, ϕ:

$$\mathcal{L} = \frac{1}{2}(\partial_\mu \phi)^2 - \frac{1}{2}m^2\phi^2 - \frac{g}{4!}\phi^4 \tag{2.3}$$

The first term is shorthand for $\frac{1}{2}\partial_\mu\phi\partial^\mu\phi = (\partial\phi/\partial t)^2 - \nabla^2\phi$. Each of the terms has specific consequences for the field, and thus for the perceived properties of the particles appearing as its excitations. The first two terms alone would describe free non-interacting scalar particles of mass m, while the last term adds self-interactions between them.

The most widespread method of calculating the elements of the S matrix resulting from a given Lagrangian is via the so-called *Feynman diagrams*. The idea is that the terms of the Lagrangian can be translated into a set of fundamental building blocks for particle propagation and interactions, each associated with certain mathematical factors through *Feynman rules*, that can then be used to build up complete amplitudes. Figure 2.1 shows the Feynman rules associated with the Lagrangian in equation (2.3).[3] Figure 2.1.a shows the *propagator* for the particles of the ϕ field, which follows from the first two terms in equation (2.3). The last term gives rise to the vertex rule in figure 2.1.b. It represents a (self-) interaction term for the ϕ field, and is proportional to the *coupling constant* g, appearing in that term. Each factor of ϕ in the last term corresponds to exactly one of the four "legs" in the vertex. In a Lagrangian describing more than one field, interactions between different types of particles would thus be handled by different fields appearing in the same term. Furthermore, note how four-momentum conservation at the vertex is built into the rule by the delta function which contains a sum of the four-momenta carried by the particles in the four legs (with opposite signs for in- and out-going particles).

To actually calculate a given matrix element, $\mathcal{M}$, one must then use these building blocks to construct all possible topologically distinct Feynman diagrams that transforms the initial state into the final state. Some of the diagrams contributing to the matrix element for $\phi+\phi \to \phi+\phi$ scattering are shown in figure 2.2. The first two diagrams on the right-hand side are so-called *tree level* diagrams and have all propagator momenta fixed by the boundary conditions of the initial and final states. The remaining diagrams are *loop* diagrams, where one or more internal momenta are unconstrained. Consequently, these undetermined momenta needs to be integrated over in order to get the total contribution of those diagrams to $\mathcal{M}$. In a sense,

[1] This discussion blatantly ignores the (physically equivalent) Feynman path integral formalism, but is in any case merely an overview.

[2] Cf. [Gol80] for an in-depth review of its application in classical physics.

[3] In order to avoid the otherwise important discussion about the propagator poles, the rules are here shown in their form after a Wick rotation to Euclidean space.

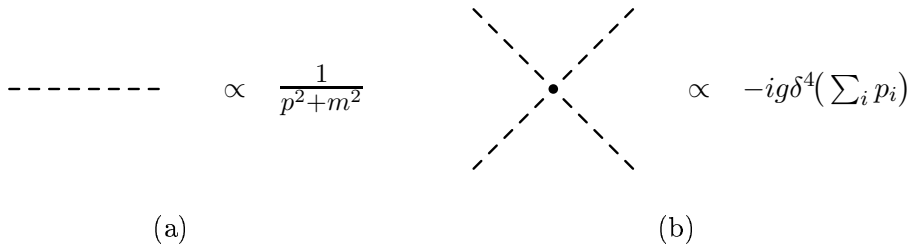

Figure 2.1: *Feynman rules in momentum space for the scalar ϕ^4 theory defined by the Lagrangian in equation (2.3).*

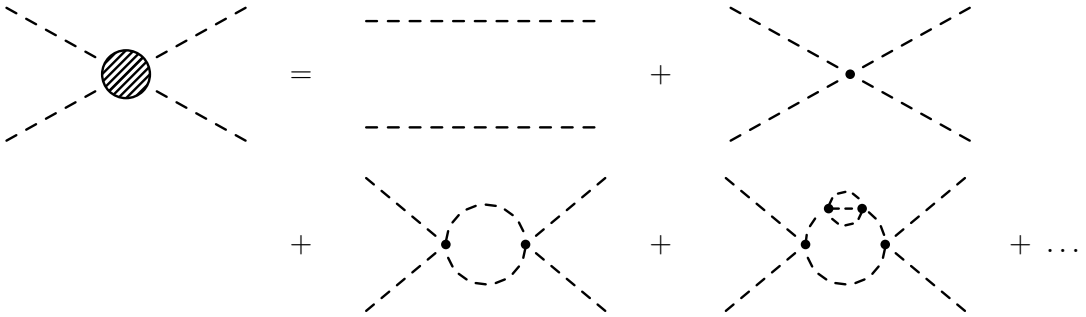

Figure 2.2: *The matrix element for $\phi + \phi \to \phi + \phi$ scattering in the scalar ϕ^4 theory defined by the Lagrangian in equation (2.3) expressed as a sum of Feynman diagrams, of which four of the simpler ones are shown. In the convention used here time is considered to run from left to right.*

the loop diagrams represent the contributions of QFT to the amplitude, while the tree-level diagrams alone would reproduce classical results.

In order to treat fermionic fields in a relativistically consistent manner one uses four component spinors, ψ, that are the objects of the so-called Dirac representation of the Lorentz group. As the Dirac representation is reducible, it is possible to form two 2 dimensional representations by considering appropriate subspaces with the eigenvectors of the two subgroups corresponding to the left and right-handed helicity parts of the fields respectively. This separation of the fields is explicit in the so-called *chiral* notation, where they are written as:

$$\psi = \begin{bmatrix} \psi_L \\ \psi_R \end{bmatrix} \tag{2.4}$$

In order to ensure the explicit Lorentz invariance of a theory, its Lagrangian should be based on Lorentz invariant components. Whereas the terms in the scalar Lagrangian in equation (2.3) are obvious Lorentz scalars, it is less obvious which combinations of the spinor ψ to consider. It turns out that $\bar{\psi}\psi$ is a Lorentz scalar while $\bar{\psi}\gamma^\mu\psi$ is a Lorentz vector, with $\bar{\psi} \equiv \psi^\dagger\gamma^0$.[1] The vector index, μ needs to be contracted with another vector in order to produce a scalar.

The explicitly Lorentz invariant Dirac Lagrangian describing a free fermion field of spin

[1] The four dimensional gamma matrices are an essential ingredient in the precise formulation of the theory, and are defined by the anti-commutator relation $\{\gamma^\mu, \gamma^\nu\} = 2g^{\mu\nu}$.

1/2 and mass m is:

$$\mathcal{L} = i\bar{\psi}\gamma^{\mu}\partial_{\mu}\psi - m\bar{\psi}\psi \tag{2.5}$$

Inserting equation (2.5) into the Euler-Lagrange equations yields two classes of solutions for ψ corresponding to the two spin-states of a particle. In general the left and right-handed components of ψ will mix under Lorentz boosts, unless the field is massless, $m \equiv 0$. This mass-induced mixing of chirality can explicitly be seen by the following valid rewriting of the mass term:

$$m\bar{\psi}\psi = m(\bar{\psi}_L\psi_R + \bar{\psi}_R\psi_L) \tag{2.6}$$

For future reference it should be noted that in addition to these so-called Dirac mass terms, it is also possible for neutral fermions to be their own anti-particles and have so-called Majorana mass terms. However, so far no observed fermion have been confirmed to possess this property (although there are some indications that neutrinos might, cf. section 4.2).

Care must be taken when quantising ψ due to its spinorial nature. This is most easily understood from the fact that standard creation operators would allow more than one identical particle in a given quantum state which would be in disagreement with Pauli's exclusion principle. Thus follows the requirement that the operation of a creation operator on a non-empty state will return it to vacuum and means that the usual commutation relationships for the field operators are replaced by anti-commutation relationships.

Fermions are usually drawn with arrows in Feynman diagrams: Arrows pointing forward in time signifies particles and arrows pointing backwards in time signifies anti-particles. This notation illustrates the fact that anti-particles can be thought of as particles running backwards in time, and also makes it clear which interactions are allowed by Lorentz invariance by requiring diagrams to always have unbroken fermion "lines" (e.g. a Lorentz variant boson-boson-fermion vertex would explicitly end a fermion line). Feynman rules can be derived for the fermionic fields just as in the scalar case, but due to the fermionic statistics they must be complemented by a few additional rules regarding the overall sign of diagrams: Each anti-fermion line passing all the way through a diagram and each fermion loop contributes an overall factor of -1 to a given diagram [Pet94, chap. 5.7].

While providing a picture of particle interactions that allows for an intuitive physical interpretation, the Achilles heel of the approach of Feynman diagrams is that, as should be clear from figure 2.2, it is always possible to construct an infinite number of different diagrams contributing to a given matrix element. The method thus only provides calculable results when it is possible to identify a manageable subset of the diagrams as providing the major part of the total amplitude. A necessary requirement for this to be possible is that the coupling constants of the theory, g in the example of equation (2.3), are sufficiently small that the contributions from diagrams with a high number of loops becomes negligible.[1] In this case, one can thus calculate amplitudes for transitions between free-field eigenstates using *perturbation theory*, keeping diagrams up to a certain order in g. A second reason for excluding certain diagrams follows from the form of the propagator factor in figure 2.1.a. While essentially any particle could appear in any amplitude in loops, particles with a mass, M, much higher than the typical momenta occurring in the process, p, will usually be suppressed by a factor of $\mathcal{O}\left(p^2/M^2\right)$.

[1] Often coupling constants are written in dimensionless normalised versions, e.g. $\alpha \equiv g^2/4\pi$ or $\alpha \equiv e^2/4\pi\epsilon_0$. This makes the relative strength of different couplings directly comparable, and the convergence of the perturbative expansion more apparent, as a diagram with n loops will be proportional to α^n.

Incidentally this also has the fortunate implication that it is possible to formulate theories with sensible results even though not all particles at higher mass scales are known.

For theories where the relevant coupling constants are not small, however, perturbation theory does not provide useful results and one must instead find the solutions of the theory by non-perturbative investigations. As will be discussed in section 3.1 this unfortunately turns out to be exactly the scenario encountered in the theory concerning strong interactions at low energies. The most widespread and generally applicable technique for non-perturbative investigations are the Lattice method where spacetime is discretised and solutions are found numerically.

2.2 Renormalisation

A deeper problem with quantum field theories surfaces when one tries to perform the integration over all possible unconstrained momenta in the loops of Feynman diagrams. As an example, if the incoming momenta are denoted p_1 and p_2, the contribution from the third diagram on the right side of figure 2.2 will be proportional to the integral:

$$\int \frac{d^4q}{(2\pi)^4} \frac{1}{(p_1+p_2-q)^2+m^2} \frac{1}{q^2+m^2} \tag{2.7}$$

Which is *divergent*! This is obviously a significant problem since the S matrix *must* be unitary in a sensible theory. A priori one could hope that the divergence is merely an artifact of the fact that only one of the diagrams in figure 2.2 is singled out for analysis, and that the summed amplitude of all of the possible diagrams would be without such problems. However, the problem is more fundamental than that. In a sense the problem arises because there are too many degrees of freedom involved when field operators, $\hat{\phi}(\mathbf{x},t)$, exist independently at each and every spacetime point. This is a short-distance problem which translates into there being too many available high-momenta states for the virtual loop particles, causing the tails of the momentum integrals to blow up. In order to recover the usefulness of quantum field theories one must thus somehow limit the available high-momentum phase-space, or equivalently impose a discretisation of the otherwise continuous spacetime with a minimal distance scale. To understand the physical significance of such a cutoff, it is useful to note that QFT can also be used to describe many-particle problems in condensed matter physics. In these applications there is a quite natural minimal distance where the continuum approximation is known to break down, namely the inter-atomic distance scale. With this in mind it is not hard to imagine the corresponding cutoffs in particle physics quantum field theories as being due to new and currently unknown effects. From this point of view, it is of course inevitable that a future Theory of Everything can not be a quantum field theory.

But contrary to the case in condensed matter physics, the nature of the QFT cutoffs in particle physics theories are unknown, and exactly how to implement them is consequently ambiguous. Their implementation is the mathematical equivalent of subtracting the infinities of the theory, and it turns out that this subtraction leaves a number of free parameters, which is not surprising as the difference between two infinite numbers is a priori undefined. In general the number of free parameters introduced in this manner is infinitely high, meaning that the theory looses predictive power - it can not be "repaired" and is said to be *non-renormalisable*. However, a small subclass of QFT theories, *renormalisable* theories, only requires a finite number of parameters to be fixed. Once these are fixed by experimental input the theory will be able to make consistent predictions. That this works is best confirmed by the theory

of electromagnetic interactions, Quantum Electrodynamics (QED), that will be discussed in more detail in section 2.3.1. The predictions of QED have in some cases been verified to 10 significant digits (cf. the anomalous magnetic moment of the electron [KN06] and [MT05]), which is arguably the highest precision in the verification of any physics theory to date.

A couple of very interesting results follow from this procedure. First of all it turns out that the quantum field theories that are renormalisable are the ones where only a few select types of terms are present in the Lagrangian. For scalar fields for instance, only terms proportional to ϕ^3 and ϕ^4 are allowed in addition to the kinematic terms. If QFT is applicable at all this means that, even before considering experimental input, the types of particles and the form of their interactions will be strongly constrained. Compared, for instance, with the potential, $V(r)$, in classical mechanics which a priori can contain any kind of term, this seems to be quite an astonishing fact. Additionally, the cutoff procedure turns out to change the view of masses and couplings considerably, which can be seen by looking in slightly more detail at how the renormalisation procedure might be applied. Assuming that a theory is renormalised by some cutoff parameter, Λ, which for simplicity will here be interpreted as the upper limit on momentum integrals, the amplitudes will a priori be a function of this scale in addition to the masses and coupling constants of the Lagrangian:

$$\mathcal{M} \equiv \mathcal{M}\left(m, g, \Lambda\right) \tag{2.8}$$

Remarkably, it turns out that this dependency on Λ can absorbed fully into the masses and couplings such that instead:

$$\mathcal{M} \equiv \mathcal{M}\left(m(\Lambda), g(\Lambda)\right) \tag{2.9}$$

The renormalisation conditions, i.e. the conditions used to fix the finite number of free renormalisation parameters, can then be taken as:

$$\lim_{\Lambda \to 0} m(\Lambda) = m_0 \tag{2.10}$$

$$\lim_{\Lambda \to 0} g(\Lambda) = g_0 \tag{2.11}$$

Where m_0 and g_0 are the usual long-distance laboratory values of the particles. If for instance the theory is QED this amounts to fixing the electron mass and charge to the values measured in traditional low-energy experiments. Fixing the parameters this way thus removes the direct dependency of $\mathcal{M}$ on Λ, and it is even possible to assign a physical interpretation to the picture of $(m, g) \to (m(\Lambda), g(\Lambda))$. In a sense it means that the charge and mass measured by a particle will depend on the energy scale at which it is measured. To get a physical view of this phenomenon, one can e.g. consider the electromagnetic charge of an electron. Due to quantum fluctuations an electron will actually never appear as a lonely particle, but will be surrounded by a "quantum cloud" of photons and electron-positron pairs as illustrated in figure 2.3. The net effect of the charges in this cloud is to shield the original "bare" value of the charge, just as if it had been immersed in a dielectric material. When probed at higher energies, i.e. with a hard photon able to penetrate this material, the charge will be shielded less and less. The net effect of the quantum fluctuations is thus to make the charge appear to grow when probed at higher and higher energy scales.

Finally, it should be noted that while renormalisability is thus an essential requirement for an entirely consistent theory, it can still be possible to work with theories that are not

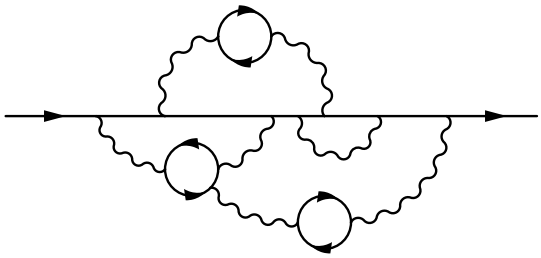

Figure 2.3: *Illustration of the quantum fluctuations surrounding an electron in QED.*

renormalisable, and section 5 will even introduce a few examples of such. The problem with such theories are that they can usually only provide reliable results at the scale where the parameters are tuned, and that the results will often be dependent on the chosen cutoff scale.

2.3 Symmetries

Symmetries have always played a special role in physics. Not only can they be a useful aid when deriving solutions to specific problems, but Nöther's theorem [EN18] additionally promises a one-to-one correspondence between the continuous symmetries and conservation laws of a physical system.

Symmetries of the S matrix involving transformations of $(\mathbf{x}, t)$ are aptly called *spacetime* symmetries, and any other symmetries *internal*. A priori one would believe that any kind of symmetries could be added to a theory, but in 1967 Coleman and Mandula [CM67] showed that there are severe restrictions on which symmetries can be allowed. The essence of the argument is that a theory required to possess too many symmetries, will become over-constrained and reduce to triviality ($S \equiv 1$). The largest allowed symmetry group that does not result in this collapse is a direct product between internal symmetries and the spacetime symmetries of the so-called Poincaré group. This specifically means that no mixing between internal and spacetime symmetries can occur. The Poincaré group is the combination of spacetime translations with Lorentz boosts and rotations, as well as three discrete symmetries, C, P and T. The associated conserved quantities promised by Nöther's theorem of translations and rotations are four-momentum and angular momentum respectively.[1] The continuous part of the Poincaré group is exactly the group containing all the continuous isometries of Minkowski space, which is not surprising given that potentials naturally depend upon relative distances. In order for a theory to be consistent with special relativity, it must be invariant under the continuous part of the Poincaré group. When this is satisfied, quantum field theories can be said to unify special relativity and Quantum Mechanics explicitly.

The three discrete spacetime symmetries that could be realised by a theory are the symmetry under Parity transformation (P), $(\mathbf{x}, t) \to (-\mathbf{x}, t)$, time-reversal ($T$), $(\mathbf{x}, t) \to (\mathbf{x}, -t)$ and charge conjugation (C), where upon particle and anti-particle states are exchanged. A very important result is that for any Quantum Field Theory containing a minimal energy state and satisfying Lorentz invariance and locality (and thus causality), the combination CPT must be a symmetry - even if the individual C, P and T are not [Lüd57]. This has

[1] There is of course also a conserved quantity associated with invariance under Lorentz boosts. In the non-relativistic limit its conservation implies that the centre of mass of a system moves linearly in time according to its momentum.

as a direct consequence that particles and their anti-particles are guaranteed to have equal masses and lifetimes and identical but opposite charges. The appearance of complex factors in a Lagrangian is usually associated with violations of T parity, and thereby of CP (but might not be if these phases can e.g. be absorbed through a field redefinition). This is perhaps not so surprising, if one considers the time evolution operator in standard quantum mechanics, $\exp(i\widehat{H}t)$. Here complex terms would lead to amplitude decay in one time direction and growth in the other.

In their derivation Coleman and Mandula only considered bosonic degrees of freedom, however, and in 1974 it was shown that the inclusion of fermionic degrees of freedom admitted one additional type of symmetry - that in a sense mixes internal and spacetime symmetries [HLS75]. This is the so-called *supersymmetry* between fermions and bosons. This interesting possibility is not realised in the Standard Model, but will be discussed in more detail in section 5.4.

One interesting aspect of symmetries of a Lagrangian is that they need not be respected by the ground state of the theory. If this is the case, the symmetry is said to be *spontaneously broken*, or "hidden". Furthermore, Goldstone's theorem [GSW62] states that every symmetry of a Lagrangian that is spontaneously broken will result in the appearance of massless scalar particles, *Goldstone bosons*, one for each dimension (generator) of the broken symmetry. If the symmetry of the Lagrangian was approximate rather than exact, these particles will be light instead of exactly massless and they are merely *pseudo Goldstone bosons*. An example of this mechanism will be treated explicitly in section 3.3.

2.3.1 Gauge Symmetries

A special class of internal symmetries occur when the mathematical description of a quantum field theory possesses unphysical degrees of freedom. These are the so-called *gauge symmetries*. As an example, consider the Lagrangian in equation (2.5) describing free non-interacting spin–1/2 fermions of mass m. This Lagrangian exhibits a $U(1)$ global gauge invariance in that it is invariant under $\psi \to \exp(-i\alpha)\,\psi$. In accordance with Nöther's theorem there is an associated conserved current, which can be interpreted as the electromagnetic charge. To derive it, one simply performs the infinitesimal transformation $\psi \to (1 - i\epsilon\alpha)\psi$ and uses the Euler-Lagrange equation to obtain $\partial_\mu j^\mu = \partial_\mu(\bar{\psi}\gamma^\mu\psi) = 0$ which is a continuity equation for the conserved charge $\int j^0 d^3x$. If ψ is taken to describe electrons, this can be seen to mean that the number of electrons minus the number of positrons is a conserved number – in other words, electromagnetic charge is conserved.

The Lagrangian in equation (2.5) thus exhibits $U(1)$ global gauge invariance. One might wish to promote this to be a stronger *local* symmetry. This would allow the $U(1)$ transformation to be different at each spacetime point, i.e. $\psi \to e^{-i\alpha(\mathbf{x},t)}\psi$. But while the mass term $m\bar{\psi}\psi$ is indeed invariant under this transformation, as $\bar{\psi} \to e^{i\alpha(\mathbf{x},t)}\bar{\psi}$, the derivative in the $i\bar{\psi}\gamma^\mu\partial_\mu\psi$ term explicitly violates this symmetry as $\partial_\mu\psi \to e^{-i\alpha(\mathbf{x},t)}\psi - ie^{-i\alpha(\mathbf{x},t)}\psi\partial_\mu\alpha(\mathbf{x},t)$. This means that the Lagrangian can only be invariant under a local $U(1)$ gauge transformation if the change in the derivative is somehow offset by something else. It turns out that this requires the derivative to be modified by exchanging ∂_μ with $D_\mu \equiv \partial_\mu - ieA_\mu$, where A_μ is a vector (spin 1) field transforming as $A_\mu \to A_\mu - e^{-1}\partial_\mu\alpha(\mathbf{x},t)$. To be consistent, the theory must then also include a kinematic term for this field. The only such term that does not itself violate the gauge symmetry is $F_{\mu\nu}F^{\mu\nu}$ where $F_{\mu\nu} \equiv \partial_\mu A_\nu - \partial_\nu A_\mu$. In particular this means

that there can be no mass term for the field A_μ. The complete Lagrangian is thus:

$$\begin{aligned}\mathcal{L} &= i\bar{\psi}\gamma^\mu D_\mu\psi - m\bar{\psi}\psi - \frac{1}{4}F_{\mu\nu}F^{\mu\nu} \\ &= i\bar{\psi}\gamma^\mu\partial_\mu\psi - m\bar{\psi}\psi + e\bar{\psi}\gamma^\mu A_\mu\psi - \frac{1}{4}F_{\mu\nu}F^{\mu\nu}\end{aligned} \tag{2.12}$$

This is *exactly* the Lagrangian of Quantum Electrodynamics! The first two terms are the usual kinematic terms for spin 1/2 electron and positrons, the third term describes the interaction with the photon field and the final term is the kinematic term for (massless) spin 1 photons. This is quite remarkable, and shows that the only way to construct a theory of spin 1/2 fermions obeying local $U(1)$ gauge invariance is by embedding the fermions in a quantum field theory of the electromagnetic interaction by adding to the theory a massless *gauge boson*, namely the photon. This connection between symmetries and the forms of the interactions is of course not a new idea. Indeed, Einstein derived the form of the interactions of General Relativity by requiring invariance under general coordinate transformations.

The example of $U(1)$ gauge symmetry can be readily generalised to invariance under essentially any interesting continuous symmetry.[1] Such groups are generally defined in terms of a number of structure constants, f^{abc}, that specify the group generators, t^a, through:

$$[t^a, t^b] \equiv if^{abc}t^c \tag{2.13}$$

As a concrete example, for $SU(2)$ the t^a's are proportional to the Pauli spin matrices, σ^i. Unless all $f^{abc} \equiv 0$, the group will not be Abelian, and this property will turn out to be the main difference with respect to the $U(1)$ case just considered for QED. Non-Abelian gauge theories are also called *Yang-Mills theories*. The general gauge transformation of ψ will look like:

$$\psi \to \exp\left(-i\alpha^a(\mathbf{x},t)t^a\right)\psi \tag{2.14}$$

where it is implicitly understood that the ψ's now contain some new index that the t^a's permute, i.e. the t^a's can be thought of as $N \times N$ matrices acting on N dimensional vectors of ψ's. Just as in the case of $U(1)$ transformations, the spacetime dependency of the gauge transformation will mean that the first term of equation (2.5) will not be gauge invariant unless the derivative is modified to:

$$D_\mu \equiv \partial_\mu - igA^a_\mu t^a \tag{2.15}$$

where the A^a_μ's are vector fields (one for each t^a) that has to be added to the theory, with appropriate transformations. The general field strength will be:

$$F^a_{\mu\nu} = \partial_\mu A^a_\nu - \partial_\nu A^a_\mu + gf^{abc}A^b_\mu A^c_\nu \tag{2.16}$$

And again, the most general form allowed by the gauge symmetry and renormalisation of the final Lagrangian will be given by:

$$\begin{aligned}\mathcal{L} &= i\bar{\psi}\gamma^\mu D_\mu\psi - m\bar{\psi}\psi - \frac{1}{4}F^a_{\mu\nu}F^{\mu\nu a} \\ &= i\bar{\psi}\gamma^\mu\partial_\mu\psi - m\bar{\psi}\psi + g\bar{\psi}\gamma^\mu A^a_\mu t^a\psi - \frac{1}{4}F^a_{\mu\nu}F^{\mu\nu a}\end{aligned} \tag{2.17}$$

[1] More precisely, to any Lie group. Lie groups can be thought of as groups whose multiplications are closed on continuous geometrical manifolds, with the most straight-forward example being rotations on spheres, represented by $SO(N)$.

Once again the first two terms represents the kinematics of the ψ field, and the third provides interactions between the ψ particles and the gauge vector bosons (with the same coupling strength for each of these interactions). The final term provides not only the kinetic term for the massless bosons, but due to the last term in equation (2.16), it also introduces third and fourth order interactions between the gauge bosons themselves – and they can thus be interpreted as carrying charge themselves. This phenomenon of charged gauge bosons turns out to be the essential difference between non-Abelian gauge theories and Abelian gauge theories such as the simple $U(1)$ example of QED.

In principle one could add one additional term involving $F^a_{\mu\nu}$ to equation (2.17) which would result in violations of P and T invariance (and thus also of CT and CP because of the CPT theorem). The fact that experimental evidence indicates the absence of such terms in the theory of the strong interactions is known as the *strong CP problem*, and is discussed further in section 4.4. A similar thing happens in $U(1)$ gauge theories like QED, but here the additional term turns out to have no effect in the theory (in the same way a constant term does not affect a quantity that only contributes to the dynamics of a theory through its derivative).

An important consequence of these self-interactions is that they change the running of the coupling constant qualitatively from the Abelian case. In QED the charge observed would grow with the energy scale as the "probe" is able to extend further and further inside the screening dielectric of virtual e^+e^- pairs. But in non-Abelian gauge theories where the gauge bosons themselves are charged, this effect instead becomes an "anti-screening", and the observed charge thus becomes smaller when probed at higher energies.

Another important feature of local gauge invariant theories is that it turns out that they are always renormalisable in four spacetime dimensions (cf. [Sal62] and [tH71b]). And although it is not strictly true that all renormalisable theories are local gauge theories (as exemplified by the ϕ^4 theory of equation (2.3)), renormalisable theories without local gauge invariance are rare. Furthermore, as will be described in the next section, it turns out that the interactions of the Standard Model Lagrangian all can be understood in terms of local gauge invariance. So although it is not really understood why the idea of local gauge invariance seems to be so important, there are every indication that it is indeed an essential part of Nature.

3 The Standard Model

Having thus described the framework of Quantum Field Theory that the Standard Model of Particle Physics lives in, the next step is then to describe the actual building blocks of the theory, i.e. the actual terms in the Lagrangian. It is fascinating to note how many constraints are already put on the theory simply by the framework. Not only are the symmetries of the theories severely constrained, but the form of the individual terms that can appear in the Lagrangian is limited to be within a select few types.

The Standard Model Lagrangian describes 12 fermion fields: 6 flavours of leptonic fields and 6 flavours of quark fields as shown in table 3.1. The Lagrangian for these spin–1/2 fermions is invariant under three different local gauge transformations and several gauge vector boson fields enters the Lagrangian accordingly. These bosonic fields are responsible for the interactions in the Standard Model, corresponding to the electromagnetic, weak and strong forces. The fields of the standard model are listed in table 3.1, and it is indicated which fermions feel which forces (i.e. interact with which gauge bosons): The quark fields feel them

Generation	Leptons		Quarks	
1	e	ν_e	u	d
	(0.511 MeV)		(~ 2 MeV)	(~ 5 MeV)
2	μ	ν_μ	c	s
	(106 MeV)		(1250 MeV)	(95 MeV)
3	τ	ν_τ	t	b
	(1777 MeV)		($172 \cdot 10^3$ MeV)	(4500 MeV)
Electromagnetic charge	-1	0	$+2/3$	$-1/3$
Weak interactions	yes	yes	yes	yes
Strong interactions	no	no	yes	yes

Table 3.1: *The fermion fields of the Standard Model including approximate observed masses. Detailed measurements are available in [PDG06]. The electromagnetic charges indicated are for particles. Anti-particles such as e^+, $\bar{u}$, etc. have charges of opposite sign. The neutrinos, ν_e, ν_μ and ν_τ are per construction massless in the Standard Model, but recent indirect measurements indicate otherwise (cf. section 4.2).*

all, the charged leptons only the electromagnetic and weak force and the elusive neutrinos only the weak force.

As indicated by the names "weak" and "strong", the coupling constants for the three interactions differs by several orders of magnitude, and the three forces thus each govern quite different phenomena. The strong force is described in section 3.1. It is mediated by the $SU(3)$ gauge boson, the aptly named gluon, and is responsible for binding quarks into hadrons such as the proton and neutron (and into nuclei through its residual strength). The electromagnetic force is mediated by a $U(1)$ gauge boson, the photon, and typically governs physics at atomic scales. Finally, the weak force, the subject of section 3.2, is really so weak that its effect is only rarely encountered in low-energy phenomena, and then only in processes disallowed by the other forces such as the flavour changing beta decays $n \to pe^-\bar{\nu}_e$. It turns out that its associated gauge bosons are required to be massive, which is in immediate conflict with gauge invariance. The solution to this problem is thought to be the Higgs mechanism, which will be discussed in section 3.3.

3.1 Quantum Chromo Dynamics

One of the challenges successfully overcome by particle physics in the last century was to make sense of the large number of hadronic states discovered in particle physics experiments in the years around 1950-1975. With a spectrum of literally hundreds of mesons (bosonic hadrons) and baryons (fermionic hadrons), the first step towards an understanding of the underlying structure came with the realisation by Gell-mann and Zweig in 1963 that the hadronic states in the spectrum could be classified by the introduction of fermionic *quarks* of spin–1/2 and electromagnetic charge of $2/3$ or $-1/3$. Mesons were bound states of a quark and an anti-quark, $q\bar{q}$, while baryons contained three quarks, qqq or anti-quarks, $\bar{q}\bar{q}\bar{q}$.

The quark theory was, however, initially not believed to be much more than a handy mathematical bookkeeping tool. It was not until around 1970 when Feynman and Bjorken interpreted results of high-energy electron-proton experiments by SLAC and MIT as showing a substructure of point-like "partons" inside protons, that quarks were generally acknowledged

Baryons

Particle	Mass [MeV]	Spin	Lifetime [s]	Decay modes	Quark content
p	938.3	$\frac{1}{2}$	stable ($> 10^{38}$)		uud
n	939.6	$\frac{1}{2}$	886	$pe^-\bar{\nu}_e$	udd
Λ^0	1116	$\frac{1}{2}$	$2.6 \cdot 10^{-10}$	$p\pi^-$, $n\pi^0$	uds
$\Lambda^0(1405)$	1406	$\frac{1}{2}$	$1.3 \cdot 10^{-23}$	$\Sigma\pi$	uds
$\Lambda^0(1520)$	1520	$\frac{3}{2}$	$4.2 \cdot 10^{-23}$	$\Sigma\pi$, nK	uds
$\Lambda^0(1830)$	1830	$\frac{5}{2}$	$0.7 \cdot 10^{-23}$	$\Sigma\pi$, nK	uds
$\Lambda^0(2350)$	2350	$\frac{9}{2}$	$4.4 \cdot 10^{-24}$		uds
Λ_c^+	2286	$\frac{1}{2}$	$2.0 \cdot 10^{-13}$		udc
Λ_b^0	5624	$\frac{1}{2}$	$1.2 \cdot 10^{-12}$		udb
Δ^{++}	1232	$\frac{3}{2}$	$0.6 \cdot 10^{-23}$	π^+p	uuu
Δ^+	1232	$\frac{3}{2}$	$0.6 \cdot 10^{-23}$	π^+n, π^0p	uud
Δ^0	1232	$\frac{3}{2}$	$0.6 \cdot 10^{-23}$	π^0n, π^-p	udd
Δ^-	1232	$\frac{3}{2}$	$0.6 \cdot 10^{-23}$	π^-n	ddd
Σ^+	1189	$\frac{1}{2}$	$0.8 \cdot 10^{-10}$	π^0p, π^+n	uus
Σ^0	1193	$\frac{1}{2}$	$0.7 \cdot 10^{-19}$	$\Lambda^0\gamma$	uds
Σ^-	1197	$\frac{1}{2}$	$1.5 \cdot 10^{-10}$	π^-n	dds
Ξ^0	1315	$\frac{1}{2}$	$2.9 \cdot 10^{-10}$	$\Lambda^0\pi^0$	uss
Ξ^-	1321	$\frac{1}{2}$	$1.6 \cdot 10^{-10}$	$\Lambda^0\pi^-$	sss
Ω^-	1672	$\frac{3}{2}$	$0.8 \cdot 10^{-10}$	Λ^0K^-, $\Xi^0\pi^-$	sss

Mesons

Particle	Mass [MeV]	Spin	Lifetime [s]	Decay modes	Quark content
π^+	139.6	0	$2.6 \cdot 10^{-8}$	$\mu^+\nu_\mu$	$u\bar{d}$
π^0	135.0	0	$0.8 \cdot 10^{-16}$	$\gamma\gamma$	$(u\bar{u} + d\bar{d})/\sqrt{2}$
K^+	493.7	0	$1.2 \cdot 10^{-8}$	$\mu^+\nu_\mu$, $\pi^+\pi^0$	$u\bar{s}$
K_S^0	497.6	0	$0.9 \cdot 10^{-10}$	$\pi^+\pi^-$, $\pi^0\pi^0$	$\approx (d\bar{s} + s\bar{d})/\sqrt{2}$
K_L^0	497.6	0	$0.5 \cdot 10^{-7}$	$\pi^\pm\ell^\mp\nu_\ell$, 3π	$\approx (d\bar{s} - s\bar{d})/\sqrt{2}$
η	547.8	0	$0.5 \cdot 10^{-20}$	$\gamma\gamma$, 3π	$\approx (u\bar{u} + d\bar{d} - 2s\bar{s})/\sqrt{6}$
η'	958	0	$3.2 \cdot 10^{-21}$	$2\pi\eta$, $\rho^0\gamma$	$\approx (u\bar{u} + d\bar{d} + s\bar{s})/\sqrt{3}$
ρ^0	776	1	$4.4 \cdot 10^{-24}$	$\pi^+\pi^-$	$u\bar{u}$, $d\bar{d}$
ρ^+	776	1	$4.4 \cdot 10^{-24}$	$\pi^+\pi^0$	$u\bar{d}$
ϕ	1019	1	$1.5 \cdot 10^{-22}$	K^+K^-, $K_L^0K_S^0$	$s\bar{s}$
D^+	1869	0	$1.0 \cdot 10^{-12}$	$K + X$	$c\bar{d}$
D^0	1865	0	$4.1 \cdot 10^{-13}$	$K + X$	$c\bar{u}$
D_s^+	1968	0	$0.5 \cdot 10^{-12}$	$K + X$	$c\bar{s}$
J/ψ	3097	1	$0.7 \cdot 10^{-20}$		$c\bar{c}$
B^0	5279	0	$1.5 \cdot 10^{-12}$	$D + X$	$d\bar{b}$
B^+	5279	0	$1.6 \cdot 10^{-12}$	$D + X$	$u\bar{b}$
B_s^0	5368	0	$1.5 \cdot 10^{-12}$	$D_s^- + X$	$s\bar{b}$
Υ	9460	1	$1.2 \cdot 10^{-20}$		$b\bar{b}$

Table 3.2: *A small excerpt of the many observed hadrons including key properties. The decay modes listed are only the dominant decay modes, where such exists. All the particles listed also have anti-particles of opposite charge. The exceptions are the particles such as the π^0 and Υ, where the listed quark content is symmetric under charge conjugation. So far, no proton decays have been observed, and the listed lifetime is an approximate lower bound. Apart from a few examples of excited Λ^0 states, all of the listed states represents the lightest hadrons of the given quark content. All data are taken from [PDG06].*

as physical objects in their own right.

But at this point, several glaring issues still demanded explanations. The first problem was the observation of the Δ^{++} particle of spin 3/2. In the quark picture this had to be interpreted as the lightest uuu baryon, i.e. a state consisting of three identical spin-half quarks with aligned spin combined in an $L = 0$ orbital state. This is, however, a picture of fermions in a symmetric state and the only way to avoid a contradiction with the spin-statistics theorem would be if the quarks carried a new internal quantum number, later dubbed *colour-charge* (as discussed shortly), under which the quark states should be anti-symmetric. This would again, in the simplest scenario, indicate $SU(3)$ as a global symmetry of the theory.

In a most satisfying way, promoting this to a local $SU(3)$ gauge symmetry would turn out to account for other unexplained features of the model. The first is the issue of *confinement*: Not only are the quarks bound into hadrons, but as no particles of fractional charge are ever observed in experiments, they must be bound in such a way that it is impossible for a quark to "break free" and be observed outside a hadron. This speaks of a very strongly coupled theory, but experiments also indicated that the coupling becomes smaller and smaller when probed at increasingly large energies, a property denoted *asymptotic freedom*. As discussed in section 2.3.1, this scaling of the coupling constant is exactly what is predicted by a non-Abelian gauge theory like $SU(3)$.

One could now proceed to writing down the actual $SU(3)$ gauge invariant Lagrangian for the quarks, but this would be somewhat redundant as this Lagrangian would essentially be similar to the one shown in equation (2.17). The only differences being that the terms involving ψ are repeated for each quark field in the theory and that the f^{abc}'s and t^a's must specifically represent $SU(3)$. As $SU(3)$ needs $3^2 - 1 = 8$ generators, the theory will contain 8 massless vector bosons, the carriers of the strong forces. These are ultimately responsible for binding the quarks into hadrons, and are thus are aptly named *gluons*.

Before one can derive and interpret the Feynman rules of the theory one must in principle proceed with a careful study of the representations of $SU(3)$. Fortunately it turns out that the relevant results can be summarised conveniently by borrowing imagery from colour-theory, and it is for this reason that the presently discussed gauge theory of quarks and gluons is given the name Quantum Chromo Dynamics, or QCD.[1] The idea is that each of the quarks carries a preserved colour-charge. This charge can have the value r (red), g (green) or b (blue). Anti-quarks on the other hand carries anti-colour, $\bar{r}$ (anti-red), $\bar{g}$ (anti-green) and $\bar{b}$ (anti-blue). Colours can be combined to $SU(3)$ colour-singlets, "white", by either combining a colour with its anti-colour ($r\bar{r}$, $g\bar{g}$ or $b\bar{b}$) or by using the combinations known from additive colour-mixing: rgb and $\bar{r}\bar{g}\bar{b}$.[2] Gluons on the other hand carries double charges consisting of one colour and one anti-colour, e.g. $r\bar{b}$, $r\bar{g}$, $g\bar{b}$, etc., and never in combinations that would make them colour-singlets.

The QCD vertices and associated power of the strong coupling constant, g, are shown in figure 3.1. The vertex in figure 3.1.a is similar to the interaction vertex in QED, while the gluon self-interactions in figure 3.1.b and 3.1.c are explicit features of the non-Abelian gauge structure.

To really understand low-energy QCD phenomena such as confinement, the vacuum struc-

[1] More in depth treatment of QCD and its development can be found in e.g. [ESW96], [PS95, chap. 14-17] and [Fri91].

[2] In principle the colour analogy would be more consistent if anti-red, anti-green and anti-blue were denoted cyan, magenta and yellow respectively. Using the nomenclature of anti-colours is, however, preferred simply because it is easier to remember.

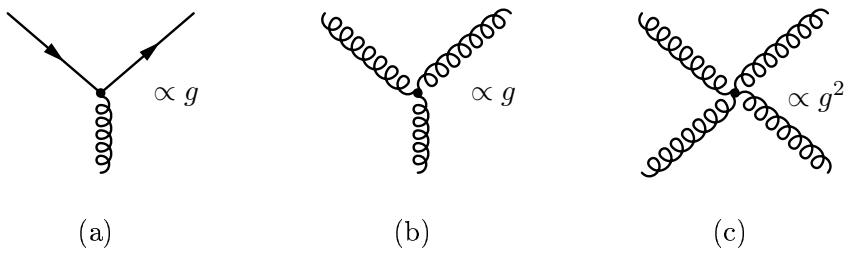

Figure 3.1: *QCD vertices. Curly lines indicates gluons while the fermion in figure (a) must be one of the quark fields. Note in particular that colour is conserved at each of the vertices.*

ture and the hadronic spectrum one would need to solve QCD explicitly without perturbative methods. Efforts in this direction has, however, so far proven fruitless. Despite this, momentous efforts has resulted in several simplified models that anyway manage to describe and provide insight into various interesting features of QCD. Few of these typically provide results of high accuracy, but with the growth of computational power in recent years the numerical lattice methods are becoming increasingly accurate and have been able to put the results of other models on firmer ground.

To get a feel about the nature of *confinement* in QCD, the so-called Lund String model is rather intuitive and has been successfully utilised for years in Monte Carlo simulations of QCD effects [SMS06]. Consider figure 3.2.a, which shows the well-known field lines between two electromagnetic charges. As the photon field is not self-interacting, these lines can be derived from a simple superposition of the individual potentials of the two charges. This means that the field lines will extend infinitely far from the charge configuration (as $V(r) \propto r^{-2}$ to be precise). In QCD the gluon field is on the other hand strongly self-interacting, and such a superposition is no longer possible. In fact, it turns out that the field lines will all be constricted to a narrow tube, or "string", stretched between the two colour-charges as illustrated in figure 3.2.b. The tube will have a width of $\mathcal{O}(1\text{ fm})$, comparable to typical hadron radii, and a constant energy density of $\mathcal{O}(1\text{ GeV/ fm})$. Consequently, if one tries to create a free quark by kicking it out of a hadron (e.g. by hitting it with a multi-GeV photon), the energy in the resulting field would quickly exceed the quark pair-creation, $g \to q\bar{q}$ threshold of $\mathcal{O}(5 - 10\text{ MeV})$, with the new quarks each matching up with the two primary quarks and thus "breaking" the string in 3.2.b into smaller strings. The end result of trying to kick a quark out of a hadron would thus simply be a fragmentation of the original system into a number of hadrons of lower energy. In very high energetic collisions involving hadrons one even speaks of entire *jets* of hadrons being emitted in the approximate direction of the initial struck quark.

By comparing the masses in table 3.1 and 3.2 it is clear that the masses of the hadrons, even those in the ground states of the "quark atoms", can not be described even approximately by adding up the masses of the constituent quarks. Without the luxury of the modern knowledge of the quark masses this would also follow simply by comparing the masses of e.g. the pions with the proton. It thus follows that the generation of hadron masses is mainly dynamical in nature. This is contrary to the situation in regular atoms where the dynamical mass effects are small, and one can therefore not expect a meson to consist of clearly distinct quarks exchanging the odd gluon. Rather it turns out, that the situation is infinitely more complex as illustrated by figure 3.3. The topmost picture illustrates that when probed at low energies, the proton appears simply as a fundamental particle of specific mass and charge. Probing a little deeper, the second picture shows 3 quarks bound by exchange of gluons. The momentum

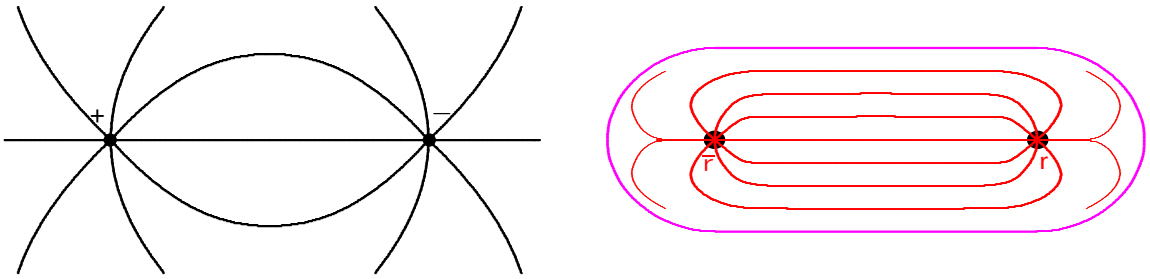

Figure 3.2: *Figure (a) illustrates how the electromagnetic field lines between a positive and a negative charge extends all the way to infinity. Figure (b) illustrates the corresponding field lines in QCD between a red and an anti-red charge. In the latter case the field lines are confined to a tube-shaped envelope of fixed width due to the self-interactions of the gluons. Source: [Sjö06].*

of a given quark will on average be approximately one third of the full proton momentum, but will vary due to the continual exchange of four-momentum. Probing at even higher energies, the third and final picture finally gives a realistic picture of the proton content. Due to gluon splitting and quark pair creation, the proton will at any given time actually consist of an entire *sea* of quarks and gluons in addition to the three *valence quarks*, each carrying a fraction of the total proton momentum.

This sea of QCD particles inside the hadrons is thus thought to be responsible for most of the hadron masses. This fact is, however, far from straight-forward, as should be clear by comparison to other bound systems such as nuclei. In these "classical" cases it is well known and easily understood that the mass of the combined bound state, the nucleus, here is *smaller* than the sum of the masses of the individual nucleons. So the phenomenon of hadron masses goes beyond a mere binding energy. To understand the phenomenon, one must realise that quark mass terms are related to a mixing of *chirality*, e.g. of right and left-handed states, as shown in equation (2.6). If one considers the QCD Lagrangian in the limit of vanishing quark masses this chiral symmetry should thus be explicit. But even so, it turns out that it is spontaneously broken by the QCD vacuum such that e.g. $\langle\bar{\psi}\psi\rangle$ obtains a non-zero value corresponding to approximately $\Lambda_{\mathrm{QCD}} = \mathcal{O}\,(200 - 300\ \mathrm{MeV})$. Without going into the details regarding the exact mechanism behind this breaking, it is nonetheless possible to get a feeling of it by considering a virtual $q\bar{q}$ pair appearing in the QCD vacuum in the frame where the emerging quarks are back to back. Unless the pair is created in an energetically unfavourable state of odd orbital angular momentum, angular momentum conservation will require the two quarks to have different handedness. In a sense it becomes energetically favourable for the vacuum to break the chiral symmetry.

With the non-zero expectation value of $\langle\bar{\psi}\psi\rangle$ in mind, one can now try to understand the hadron masses in table 3.2. Apart from the pions, kaons and the eta particle, the masses do seem roughly consistent with having values around 300 MeV per constituent quark in addition to the masses of the constituent quarks themselves (with higher values for exited states such as those of the Λ^0 shown in the table). Although the pion, kaon and eta masses seems to be too low in this picture, they actually *confirm* the idea of a spontaneously broken chiral symmetry. This can be understood in the context of Goldstone's theorem mentioned in section 2.3. In the approximation where the masses of the u and d quarks are assumed to be negligible, the chiral symmetry breaking will change the symmetry group from $SU_L(2) \times SU_R(2)$ to $SU(2)$, a loss of

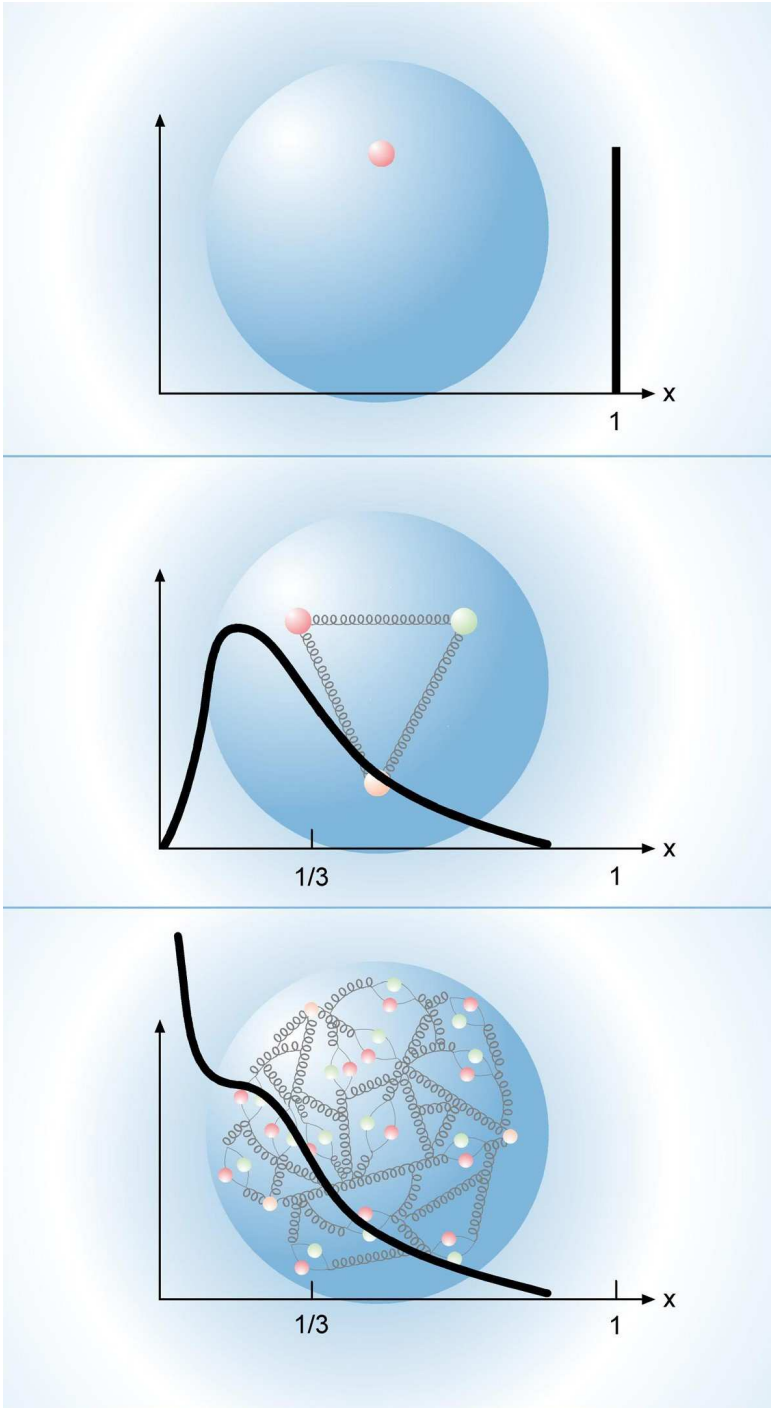

Figure 3.3: *Illustration of the proton content relevant at various scales as discussed in the text. Figure borrowed from* `http://desy.de/`.

$2^2 - 1 = 3$ generators. This accounts for the appearance of three massless Goldstone bosons. However, m_u and m_d are not exactly 0, so these bosons, identified with the three pions, will acquire a small mass. In a sense, the Goldstone mechanism can not fully protect them from acquiring a mass when the original Lagrangian symmetry was not exact. If also the s quark is considered light, the breaking of symmetry will instead be $SU_L(3) \times SU_R(3) \to SU(3)$, and there will be $3^2 - 1 = 8$ light scalar bosons. As the s quark is rather heavier than the u and d quarks, the 5 additional pseudo Goldstone bosons, the 4 kaons and the eta particle, will be slightly heavier again.

In addition to strong interactions, quarks also are also subject to the electromagnetic and weak forces that are the subject of the next two sections. These forces play a very limited role in the formation and structure of hadrons, but can be crucial for decay processes.

3.2 Weak Interactions

Looking at the various lifetimes in table 3.2, most of them seem to fall into two distinct groups, being either relatively long with values around $10^{-12} - 10^{-8}$ s, or extremely short around $10^{-24} - 10^{-20}$ s. The latter group is straight-forward to identify as strong decays. Not only are they the most numerous (this would be even more pronounced if the table had room for more of the many excited states such as those shown for the Λ^0), but the decays can be described at quark-level using QCD with a compatible interaction strength as $\hbar/100$ MeV $\approx 10^{-23}$ s. However, the decays of the longer lived hadrons with lifetimes around $10^{-12} - 10^{-8}$ s must be governed by a different fundamental interaction. Not only must such decays be associated with a much weaker coupling constant, but the decays would not be possible at all in QCD. This is essentially because the fundamental quark interaction vertex in figure 3.1.a is *flavour conserving*, i.e. $N(u) - N(\bar{u})$, $N(s) - N(\bar{s})$, $N(c) - N(\bar{c})$, etc. are all conserved by the strong interaction, but required to change under the long-lived decays in question. This also excludes QED as being the interaction responsible for these decays, as it too conserves flavour and is otherwise too strongly coupled to produce the required long lifetimes, as exemplified by the two decays in table 3.2 that actually *are* electromagnetic in nature, namely those of π^0 and Σ^0.

The only viable option left is thus the introduction of a third fundamental force, the *weak* force. Additional experiments and careful scrutiny revealed that not only did the new interaction seem to affect leptons in the same way as quarks – hence the leptons in the decay products in table 3.2, but it also violated P and C *maximally* (cf. [LY56] and [WAH$^+$57]). These effects are all apparent in the muonic decay, $\mu^- \to e^- \bar{\nu}_e \nu_\mu$. One finds that the decay rate of left-handed muons is $\Gamma(\mu_L^- \to e_L^- \bar{\nu}_e^R \nu_\mu^L) \approx 1/(2.2 \cdot 10^{-6}\text{ s})$, consistent with the rates of weak hadron decays. P and C invariance would predict that this rate should be exactly equal to $\Gamma(\mu_R^- \to e_R^- \bar{\nu}_e^L \nu_\mu^R)$ and $\Gamma(\mu_L^+ \to e_L^+ \nu_e^R \bar{\nu}_\mu^L)$ respectively, but experiments show these two rates to be vanishing![1]

It turns out that the data can be explained by the introduction of a local $SU(2)$ gauge symmetry acting, as shown in table 3.3, on doublets of the left-handed parts of the fermionic fields, while leaving the right-handed parts alone, i.e. as singlets of the transformation. A construction such as this would actually only be possible if all fermion masses were exactly vanishing, as their presence mixes left and right-handed fields and thus violates the newly

[1] Of course, any muon will decay eventually, as its mass term mixes its right and left-handed components. A more precise statement would be that the decay amplitude in any given frame of reference is vanishing for the right-handed part of the muon field.

Generation	Quarks	Leptons
1	$\begin{pmatrix} u \\ d \end{pmatrix}_L$ u_R, d_R	$\begin{pmatrix} \nu_e \\ e \end{pmatrix}_L$ e_R
2	$\begin{pmatrix} c \\ s \end{pmatrix}_L$ c_R, s_R	$\begin{pmatrix} \nu_\mu \\ \mu \end{pmatrix}_L$ μ_R
3	$\begin{pmatrix} t \\ b \end{pmatrix}_L$ t_R, b_R	$\begin{pmatrix} \nu_\tau \\ \tau \end{pmatrix}_L$ τ_R

Table 3.3: *The structure of the Standard Model fermions with respect to the weak force. The left-handed fields are placed in doublet representations of $SU(2)$ while the right-handed parts are placed in singlets.*

introduced gauge symmetry (cf. equation (2.6)). The solution to this problem requires the introduction of a so-called *Higgs boson*, and will be discussed in section 3.3. The omittance of the right-handed neutrino fields from table 3.3 is no coincidence. In the original version of the Standard Model the neutrinos are exactly massless, and as the neutrino carries neither electromagnetic nor colour-charge, a right-handed neutrino will be entirely non-interacting, and its presence in the model will provide no observable effects at all. In section 4.2, the recent indirect measurements of non-zero neutrino masses and their implications will be discussed further.

Now, instead of simply assigning a very small value to the coupling constant associated with the new gauge symmetry, it turns out that the correct procedure is instead to use a coupling constant of the same order as the electromagnetic coupling but to make the $2^2 - 1 = 3$ new gauge bosons very massive, thus limiting the effective range of the interaction. How these apparently gauge-violating and non-renormalisable mass terms can be accommodated in the theory will be the subject of section 3.3, but their magnitude can be very roughly estimated simply by comparing the electromagnetic and weak lifetime ratios to ratios of massless and massive propagators, assuming typical momentum transfers of $q \approx \mathcal{O}(1\text{ GeV})$:

$$\frac{\tau_{em}}{\tau_{weak}} \approx \frac{10^{-10}\text{ s}}{10^{-17}\text{ s}}, \quad \left(\frac{\tau_{em}}{\tau_{weak}}\right)^{\frac{1}{2}} \approx \left.\frac{(0^2+q^2)^{-1}}{(M^2+q^2)^{-1}}\right|_{q\approx 1\text{ GeV}} \Rightarrow M \approx 60\text{ GeV} \tag{3.1}$$

The precise formulation of the theory introduces two charged and one neutral massive gauge vector boson, and experiments show them to have masses of $m_{\mathrm{W}} \approx 80.4$ GeV and $m_{\mathrm{Z}} \approx 91.2$ GeV. Their interactions with the Standard Model fermions are shown in figure 3.4 along with that of the photon. As $SU(2)$ is a non-Abelian gauge group, the theory will also contain self-interactions of the gauge bosons as shown in figure 3.5 (because of the gauge boson masses, however, this does not lead to confinement as in QCD).

At tree-level, decays of muons and charged pions can thus be described by the Feynman diagrams shown in figure 3.6.[1] However, it would seem that e.g. the decay $K^- \to \mu^-\bar{\nu}_\mu$ is not accommodated in this theory. In principle it should proceed much like the pionic decay

[1] Incidentally, the fact that $\pi^- \to \mu^-\bar{\nu}_\mu$ is heavily preferred compared to the energetically favourable $\pi^- \to e^-\bar{\nu}_e$, is a clear experimental confirmation that the weak gauge bosons only couple to the left-handed fields. Angular momentum conservation and $m_\nu = 0$ requires the charged lepton to be created in a ℓ_R^- state, which is suppressed by the smallness of the lepton mass.

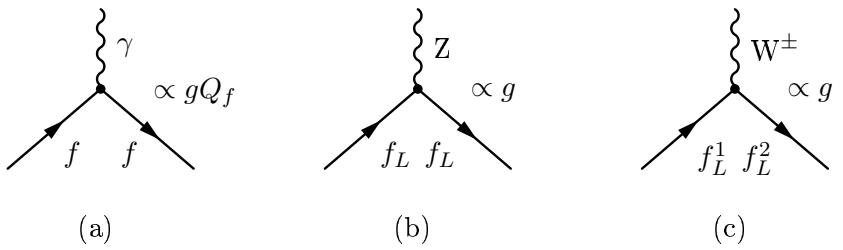

Figure 3.4: *Interaction vertices between fermions and electroweak gauge bosons. Photons couple to any fermions proportional to their electromagnetic charge, In a similar diagram the Z couples to the left-handed parts of the fermion fields. W's likewise couples only to left-handed fields, but connects the two fields in a weak doublet. This means that the pair of fermions fields, (f_L^1, f_L^2), in figure (c) can be any of $(e_L, \nu_{e,L})$, $(\mu_L, \nu_{\mu,L})$, $(\tau_L, \nu_{\tau,L})$, (u_L, d_L), (c_L, s_L) or (b_L, t_L).*

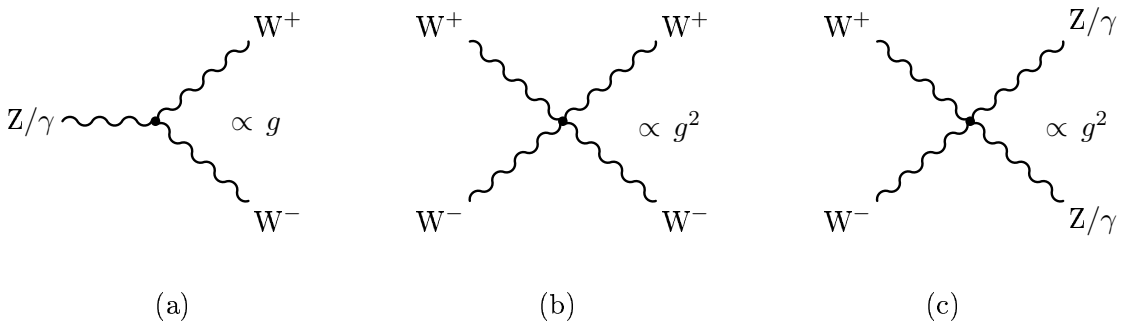

Figure 3.5: *Self-interaction vertices for electroweak gauge bosons.*

in figure 3.6.b, with an s quark instead of the d quark, but the vertex $s\bar{u}W^-$ would a priori not seem to appear in figure 3.4. The reason for this is that it was implicitly assumed when writing up the weak doublets in table 3.3, that the flavour eigenstates appearing there were the same as the mass eigenstates from table 3.1. However, there is a priori no reason why this has to be the case, and indeed experimental evidence such as the kaon decay show that it is not. Using i as a flavour index, the left-handed parts of the up and down-type quark fields can be written as u_L^i and d_L^i respectively, so that $u^1 = u$, $u^2 = c$, $d^3 = b$, etc. Denoting the states in the weak basis (those in table 3.3) with a prime and those in the mass-eigenstate basis unprimed, there must exist unitary basis changing matrices U_u and U_d satisfying

$$\begin{aligned} u_L^i &= U_u^{ij} u_L'^j \\ d_L^i &= U_d^{ij} d_L'^j \end{aligned} \tag{3.2}$$

Most of the terms in the Standard Model Lagrangian are invariant under this change of basis. For instance, the terms responsible for the left-handed part of the up-type quark couplings to photons, gluons and Z bosons are all proportional to:

$$\bar{u}_L'^i \gamma^\mu u_L'^i = \bar{u}_L^i \gamma^\mu (U_u^{\dagger\, ij} U_u^{jk}) u_L^k = \bar{u}_L^i \gamma^\mu \delta^{ik} u_L^k = \bar{u}_L^i \gamma^\mu u_L^i \tag{3.3}$$

However, the W interaction terms connects the up and down-type quarks and a similar can-

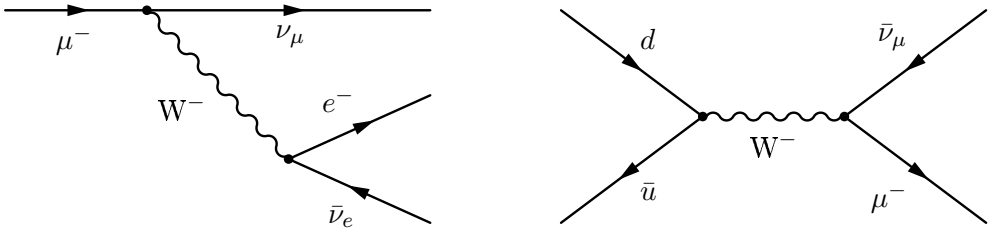

(a) (b)

Figure 3.6: *Examples of weak decays in the Standard Model at tree-level. Figure (a) shows the decay* $\mu^- \to e^- \bar{\nu}_e \nu_\mu$ *and figure (b) shows* $\pi^- \to \mu^- \bar{\nu}_\mu$.

cellation of the unitary matrices will not occur in those:

$$\begin{aligned} W_\mu^+ \bar{u}_L'^i \gamma^\mu d_L'^i + W_\mu^- \bar{d}_L'^i \gamma^\mu u_L'^i &= W_\mu^+ \bar{u}_L^i \gamma^\mu (U_u^{\dagger\, ik} U_d^{kj}) d_L^j + W_\mu^- \bar{d}_L^i \gamma^\mu (U_d^{\dagger\, ik} U_u^{kj}) u_L^j \\ &\equiv W_\mu^+ \bar{u}_L^i \gamma^\mu {V_{\text{CKM}}}^{ij} d_L^j + W_\mu^- \bar{d}_L^i \gamma^\mu {V_{\text{CKM}}}^{\dagger\, ij} u_L^j \end{aligned} \tag{3.4}$$

where the Cabbibo-Kobayashi-Maskawa mixing matrix has been introduced [Cab63], [KM73]:

$${V_{\text{CKM}}}^{ij} = U_u^{\dagger\, ik} U_d^{kj} \tag{3.5}$$

For a theory of N quark doublets, V_{CKM} will a priori contain $2N^2$ independent parameters. Unitarity constraints and redefinitions of field phases will reduce this number to $(N-1)^2$ of which $\frac{1}{2}(N-1)(N-2)$ will be complex phases. As mentioned in section 2.3, the presence of such complex phases in a Lagrangian is equivalent to violation of the combined symmetry CP, and the $\frac{1}{2}(3-1)(3-2) = 1$ such complex phase in the Standard Model is the only known source of such violations. In fact, even before the discovery of the third generation b and t quarks, they were predicted based on the observation of CP violation in the neutral kaon systems [CCFT64] consisting of only u, d and s quarks!

As to the actual contents of the CKM matrix, experimental data suggest that the unitary "rotation" it represents is rather small, and it is thus convenient to write it in terms of an expansion parameter, $\lambda \approx 0.23$, using the so-called Wolfenstein parametrisation [Wol83]:

$$V_{\text{CKM}} = \begin{pmatrix} V_{ud} & V_{us} & V_{ub} \\ V_{cd} & V_{cs} & V_{cb} \\ V_{td} & V_{ts} & V_{tb} \end{pmatrix} \approx \begin{pmatrix} 1 - \frac{\lambda^2}{2} & \lambda & A\lambda^3(\rho - i\eta) \\ -\lambda & 1 - \frac{\lambda^2}{2} & A\lambda^2 \\ A\lambda^3(1 - \rho - i\eta) & -A\lambda^2 & 1 \end{pmatrix} + \mathcal{O}\left(\lambda^4\right) \tag{3.6}$$

Where $A \approx 0.81$, $\rho \approx 0.20$ and $\eta \approx 0.33$.[1] Given that $V_{\text{CKM}} \neq 1$, the W interaction will now connect up and down-type quarks, but the interactions vertices will pick up CKM factors as illustrated in figure 3.7. The hierarchic structure of the CKM matrix evident in equation (3.6) means that transitions between generations will be suppressed by powers of λ compared to transitions within generations. This suppression is termed *Cabibbo suppression*, and is the reason that the charged kaon and pion lifetimes are similar despite huge differences in the available phase-space. The inter-generational quark mixing furthermore means that the flavour numbers, $N(u) - N(\bar{u})$, $N(s) - N(\bar{s})$, $N(c) - N(\bar{c})$, etc. that are otherwise conserved by the QED and QCD interactions are violated. However, the violation is such that generation agnostic *baryon number*, $\frac{1}{3}[N(q) - N(\bar{q})]$ is still conserved.

[1] The numbers are from the summer 2006 updates to [CKM05].

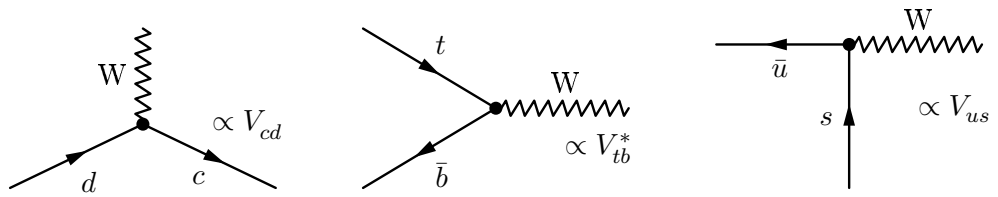

Figure 3.7: *The $u^i d^j W$ vertices pick up CKM factors when considering the quarks in the mass-eigenstate basis.*

One could try to use a similar chain of arguments to introduce CKM like factors in the leptonic sector of the Standard Model Lagrangian, but this does not happen with massless neutrinos, as such mass degeneracy would make it possible to choose the same diagonalising matrix for the neutrinos as for the charged leptons – making them cancel out even in the interactions with the W bosons.

Specifically this means that a "lepton flavour number" is still conserved for each generation in the Standard Model. These conserved *lepton numbers* are:

$$\begin{aligned} L_e &\equiv N(e^-) + N(\nu_e) - N(e^+) - N(\bar{\nu}_e) \\ L_\mu &\equiv N(\mu^-) + N(\nu_\mu) - N(\mu^+) - N(\bar{\nu}_\mu) \\ L_\tau &\equiv N(\tau^-) + N(\nu_\tau) - N(\tau^+) - N(\bar{\nu}_\tau) \end{aligned} \tag{3.7}$$

Recent indirect measurements have, however, proven the existence of very small but nonzero neutrino masses as will be discussed in section 4.2. Analogously to the quark sector, these observations would thus imply small violations of the generational lepton numbers in equation (3.7), but with $L_e + L_\mu + L_\tau$ still being conserved.

3.3 The Higgs Mechanism and Electroweak Unification

While the special structure of the gauge theory of weak interactions accounts for many experimentally observed features of hadronic and leptonic interactions, the discussion in the previous section ignored the fact that mass terms for both the fundamental fermions as well as for the weak gauge bosons themselves break the gauge symmetry explicitly in the Lagrangian, making the theory non-renormalisable in the process.[1]

The solution to this problem is to keep the gauge symmetry intact in the Lagrangian and instead let the ground state or vacuum break the asymmetry, the phenomenon known as *spontaneous symmetry breaking* mentioned in section 2.3. Such a theory will be renormalisable [tH71a] despite of the appearance of "fake" mass terms. As an example, consider the mass term $m\bar{\psi}\psi$, which a priori is forbidden in the theory. It is, however, not problematic for the theory to include a scalar field, ϕ, that couples to ψ through the interaction term $\lambda_\psi \phi \bar{\psi}\psi$ for some coupling λ_ψ. If the vacuum then breaks the symmetries of the Lagrangian in such a way that ϕ achieves a non-zero vacuum expectation value, v, it can be expanded about the new

[1] The fact that gauge boson mass terms break local gauge invariance was already implicitly mentioned in section 2.3.1. That fermion mass terms likewise break it follows from equation (2.6) and the fact that the left and right-handed parts of the fields transform in different representations under the weak gauge transformation.

minimum $\phi(\mathbf{x},t) = v + h(\mathbf{x},t)$ (note that the requirement of Lorentz invariance implies that such a scenario is only possible for a scalar field). The fermionic coupling of ϕ would then give rise to a mass term for ψ, as well as a coupling between ψ and the *Higgs field*, h:

$$-\lambda_\psi \phi \bar{\psi}\psi \rightarrow -\lambda_\psi v \bar{\psi}\psi - \lambda_\psi h \bar{\psi}\psi \tag{3.8}$$

In order for the spontaneous symmetry breaking to additionally provide non-zero masses for gauge bosons, ϕ must transform as a non-singlet under the associated local gauge symmetry. This mechanism is most easily illustrated using an example of a $U(1)$ local gauge symmetry. Consider the QED Lagrangian in equation (2.12), extended with a complex scalar field, $\phi \equiv \phi_1 + i\phi_2$, also transforming under the local gauge invariance as $\phi \rightarrow e^{-i\alpha(\mathbf{x},t)}\phi$:

$$\mathcal{L} = i\bar{\psi}\gamma^\mu D_\mu \psi - \lambda_\psi \phi \bar{\psi}\psi - \frac{1}{4}F_{\mu\nu}F^{\mu\nu} + |D_\mu \phi|^2 - V(\phi) \tag{3.9}$$

Here the original ψ mass term of equation (2.12) has been substituted with an interaction with ϕ, although this is not strictly necessary in the QED. $V(\phi)$ is the *Higgs potential* which can be written in the most general form that is renormalisable and respects the $U(1)$ symmetry:

$$V(\phi) = \mu^2 \phi\phi^* + \lambda(\phi\phi^*)^2 \tag{3.10}$$

λ must always be positive if the theory is to have a stable vacuum at all. If μ^2 is also positive, the vacuum ground state will be $\phi = 0$ as seen in figure 3.8.a, which is invariant under the gauge transformation. If, however, $\mu^2 < 0$, the solution at $\phi = 0$ becomes unstable and a multitude of degenerate ground states appear at $|\phi| = \sqrt{-\mu^2/2\lambda} \equiv v/\sqrt{(2)}$, as illustrated by the dashed line in figure 3.8.b. No matter which of these states are chosen by Nature as the vacuum state, it will no longer respect the $U(1)$ gauge symmetry. For simplicity, and without loss of generalisation, this vacuum can be chosen to lie along the ϕ_1 axis, and the components of ϕ are naturally written in terms of radial and tangential deviations from this vacuum using two real scalar fields h and χ:

$$\sqrt{2}\phi(\mathbf{x},t) = \big[v + h(\mathbf{x},t)\big] \exp\left(i\,\frac{\chi(\mathbf{x},t)}{v}\right) \approx v + h(\mathbf{x},t) + i\chi(\mathbf{x},t) \tag{3.11}$$

Where the last approximation holds when h and χ represents small fluctuations. In these variables, the potential becomes:

$$V(\phi) = -\frac{\lambda}{4}v^4 + \lambda v^2 h^2 + \lambda v h^3 + \frac{\lambda}{4}h^4 \tag{3.12}$$

The first of these terms is just a constant and will not affect the dynamics of the model. The second is a mass term for the scalar *Higgs particles* represented by the field h, and the remaining terms represents cubic and quartic self interactions of this field. No mass term for χ appears, and χ thus represents massless scalar particles, which is not so surprising as it represents excitations among states of equal energy. This is actually an explicit example of the predictions of Goldstone's theorem: $U(1)$ is a continuous symmetry of the Lagrangian that is spontaneously broken by the vacuum, and this must result in the appearance of massless Goldstone bosons. In fact, with the present example in mind, the general validity of Goldstone's theorem becomes intuitively clear: If a Lagrangian is invariant under some continuous symmetry group, a spontaneously broken vacuum state can always be connected to a continuum of other states of equal energy by the generators of this group. It will thus always

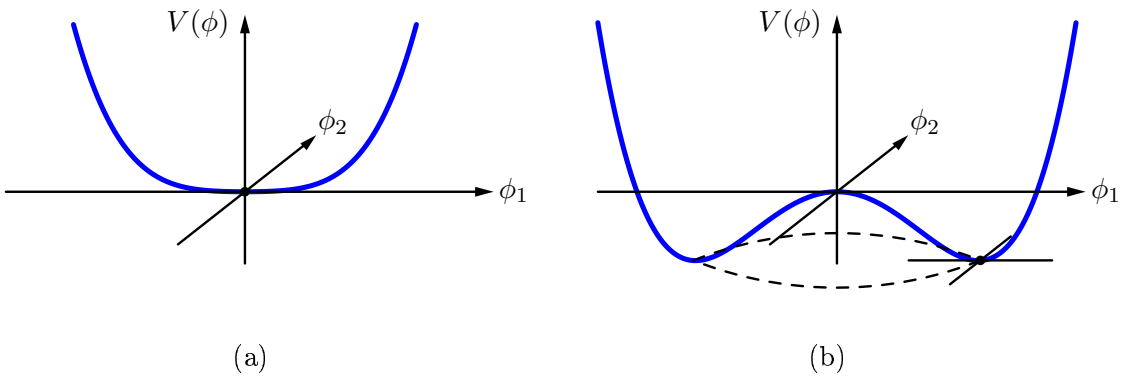

Figure 3.8: *The Higgs potential of equation (3.10). Figure (a) and (b) shows the potential for* $\mu^2 > 0$ *and* $\mu^2 < 0$ *respectively.*

be possible to obtain energetically costless, i.e. massless, fluctuations away from the vacuum state.

One can now proceed to apply the variable change of equation (3.11) to the rest of the Lagrangian in equation (3.9) also. This results in mass terms for both the ψ and gauge vector boson fields, as well as interactions between each of these and h and χ. Initially this does *not* look like a feasible mechanism for providing mass, however, as it invariably predicts the existence of a massless scalar particle, χ, in stark contradiction with experiments. This problem would indeed be insurmountable if the $U(1)$ gauge symmetry had merely been global instead of local. But it turns out that the interactions between the ϕ and vector boson fields due to the derivative, D_μ in equation (3.9), are precisely such that the, χ field will never give rise to physical particles and that the degree of freedom associated with it instead manifests itself in longitudinal polarisation states of the vector boson that became available when it acquired mass. In a sense, the vector boson has "eaten" the Goldstone boson and acquired mass in the process. That this is the result is definitely not obvious from merely looking at the Lagrangian. It can, however, be made immediately clear by using the gauge invariance to eliminate χ from equation (3.11), by

$$\phi(\mathbf{x},t) \to e^{-i\alpha(\mathbf{x},t)}\phi(\mathbf{x},t) \quad \text{with} \quad \alpha(\mathbf{x},t) = \frac{\chi(\mathbf{x},t)}{v} \tag{3.13}$$

In this choice of gauge, denoted *the unitary gauge*, the χ field does not even exists. Any other choice of gauge will result in exactly similar physical amplitudes, but with contributions from an unphysical *ghost field*, χ.

The $U(1)$ example of the Higgs mechanism discussed so far, essentially captures all of its qualitative features. The rest of this section will consist of a brief review of how the Higgs mechanism is applied to the weak and electromagnetic sectors of the Standard Model [Wei67]. ϕ is here taken to be a tuple of complex scalar fields transforming as a doublet under the weak $SU(2)$ gauge transformation (transforming as usual under the $U(1)$ electromagnetic gauge transformation):

$$\phi = \begin{pmatrix} \phi_1 + i\phi_2 \\ \phi_3 + i\phi_4 \end{pmatrix} \tag{3.14}$$

with real ϕ_i's. This Higgs potential will the keep same algebraic form as in equation (3.10), but with $\phi\phi^*$ replaced by $\phi^\dagger\phi = \phi_1^2 + \phi_2^2 + \phi_3^2 + \phi_4^2$. The one-dimensional circle of degenerate ground states indicated by the dashed line in figure 3.8.b is consequently replaced by a three-dimensional four-sphere, resulting in a physical Higgs boson of mass $m_h^2 = -2\lambda v^2$ as well as 3 Goldstone bosons that end up being "eaten" by three of the four gauge bosons that thus acquires masses. This leaves three massive gauge bosons, W^+, W^- and Z and one massless, the photon (γ). When the detailed calculations are carried out, it turns out that the physical fields associated with the Z and γ particles do not directly correspond to vector boson fields associated with the $SU(2)$ and $U(1)$ generators, respectively. Rather, they are the orthogonal mixtures of the vector boson field associated with $U(1)$ and the neutral vector boson field associated with $SU(2)$ that defines the mass-eigenstates of these neutral vector bosons. One important consequence of this mixing is that the Z mass will not be precisely equal to that of the W bosons. If θ_W denotes the angle between the two bases of the neutral vector bosons, the mass relation turns out to be:

$$m_{\mathrm{W}} = m_{\mathrm{Z}} \cos\theta_{\mathrm{W}} \tag{3.15}$$

As θ_{W} can be extracted from experimental measurements of couplings to be $\sin^2\theta_{\mathrm{W}} \approx 0.23 \Rightarrow \cos\theta_{\mathrm{W}} \approx 0.87$, this provides a very convincing verification of the model, as the observed values of $m_{\mathrm{W}} \approx 80.4$ GeV and $m_{\mathrm{Z}} \approx 91.2$ GeV are indeed in very good agreement with this. Furthermore, these masses also indicate a value of the vacuum expectation value of the Higgs field of $v \approx 246$ GeV.

It is an explicit part of the Standard Model that the electromagnetic and weak forces are associated with the same coupling constant, g, and one therefore usually says that the two forces are *unified* in the electroweak force. However, it is interesting to note that the gauge group $SU(2) \times U(1)$ is reducible, so in that respect the two forces still require some separate algebraic treatment. Whether or not the two forces are truly unified is thus a matter of opinion, although it is certainly a fact that they mix.

The fermion-interaction terms in the $SU(2) \times U(1)$ electroweak theory must be somewhat more complex than the $\lambda_\psi \phi \bar{\psi}\psi$ used in the simple example above. The right-handed fields transform as $SU(2)$ singlets while the Higgs and left-handed fields transform as $SU(2)$ doublets. Assuming that e.g. $(0, 1)$ in equation (3.14) represents the direction of the vacuum expectation value (i.e. along the ϕ_3 direction), an example of gauge invariant interaction terms, so-called *Yukawa couplings*, leading to down-type quark masses would be:

$$-\lambda_d(\bar{u}_L, \bar{d}_L)\phi\, d_R + h.c. \quad \xrightarrow{\text{unitary gauge}} \quad -\frac{\lambda_d}{\sqrt{2}} v\bar{d}d - \frac{\lambda_d}{\sqrt{2}} h\bar{d}d \tag{3.16}$$

The first resulting term in this example is thus a mass term for the d-quark with $m_d = \lambda_d v/\sqrt{2}$ and the second represents its coupling to the Higgs field, h. An important result that follows from this is that *the coupling of a fermion to the Higgs field is proportional to its mass.* To provide masses to up-type particles also, additional terms similar to the ones in relation (3.16) are needed, but with ϕ replaced by $\phi^c \equiv i\sigma^2\phi^*$. This necessity is clear in the unitary gauge where the first entry in $\phi = (0, v + h)$ would make such a term vanish, unless ϕ is replaced by $\phi^c = -(v + h, 0)$.

Even though the electroweak theory thus seems to be in very good agreement with experimental data (as discussed in greater detail in section 4), the main problem of the theory is simple: The Higgs particle associated with the h field has so far not been seen in experiments.

The experimental limits does, however, not rule out the model and one of the main objectives of the LHC is indeed to try to verify this final missing piece of the Standard Model.

It is clear, however, that *something* is needed to explain the gauge boson masses, as the theory is a priori not renormalisable without. One consequence of this non-renormalisability in a theory without Higgs bosons would be that the cross section for WW scattering, $W^+W^- \to W^+W^-$ would grow with energy and thus violate unitarity. Negative interference from the amplitudes with intermediate Higgs bosons such as $W^+W^- \to h \to W^+W^-$ solves this problem, but if the Higgs boson mass is not higher than around 800 GeV, which thus represents a very stringent cut on the value of this parameter.

The particular structure of the Higgs mechanism explained here represents the most simple choice, but there is no particular reason why it would need to be realised. One straight-forward alternative would be to introduce two distinct Higgs doublets, ϕ_1 and ϕ_2, each acquiring different vacuum expectation values, v_1 and v_2, respectively, when the theory is spontaneously broken. The vacuum must still be degenerate with three degrees of freedom: Fewer would not be enough to provide masses for all gauge bosons and more would imply observable massless Goldstone bosons. This means that the theory will contain a total of $8-3=5$ massive scalars. These will consist of two CP even neutral scalars h^0 and H^0, and one which is CP odd and neutral, A^0, in addition to two charged scalars H^+ and H^-. The couplings of these scalars to fermions will no longer simply depend on the fermion mass, but will in general also depend upon the ratio of the two vacuum expectation values, $\tan\beta \equiv v_1/v_2$, the mixing angle between the h^0 and H^0 state and, in case of the charged scalar coupling to quarks, CKM factors. The details of the theory will furthermore depend on the exact couplings between fermions and the Higgs doublets responsible for the fermion masses. In principle there would be many combinations pairing one of ϕ_1, ϕ_2, ϕ_1^c and ϕ_2^c with either up-type quarks, down-type quarks and charged leptons respectively.

A final note might be in order regarding the potential $V(\phi)$ inserted in equation (3.9). Some people might find it troubling that terms proportional to ϕ^2 and ϕ^4 that in previous sections were interpreted as mass and interaction terms, are now interpreted as describing a potential that needs to be minimised in the classical sense. However, this viewpoint is equivalent and is merely assumed out of convenience. One might just as well think of the Higgs field as describing fundamental tachyons of mass $i\,|\mu|$ that through its ϕ^4 self-interaction term achieves a vacuum expectation value different from 0, with the h field representing a condensate of non-tachyonic mass. In this picture, the Higgs mechanism has conceptual similarities to how the chiral symmetry breaking in QCD provides hadron masses. The difference is mainly that whereas the ground states away from zero are put explicitly into the Higgs potential while they are being dynamically built up in QCD.

4 Experimental and Theoretical Status of the Standard Model

Having thus described the Standard Model of particle physics in the previous sections, it is now time to look at its present experimental status. As in previous sections, the aim is merely to provide a brief overview, leaving the details to the references.

As discussed in section 3.1, there are many mathematical difficulties in making high precision predictions when dealing with QCD – especially at low energies where the theory is non-perturbative. The most precise tests of the SM are thus usually performed in the electroweak sector, and the many experimental tests of QCD mentioned in section 3.1 will not be

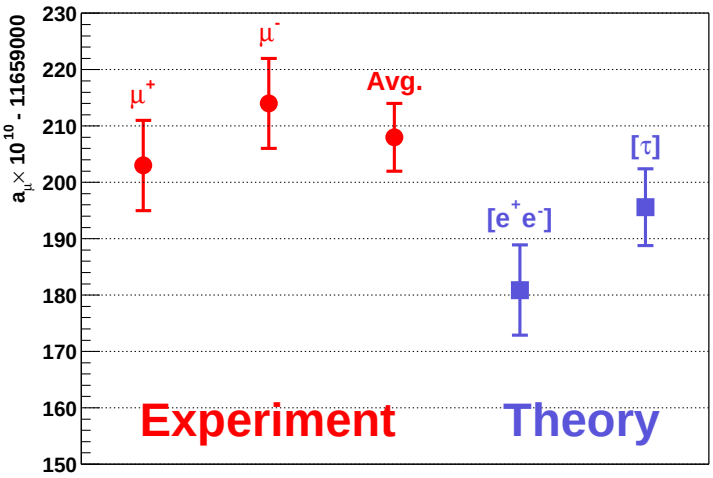

Figure 4.1: *Comparison between SM predictions and measurements of the muon anomalous magnetic moment as discussed in the text. Source [GM204].*

discussed further here. The perhaps most precise verification of the electromagnetic part of the electroweak sector has already been mentioned, namely the calculation and measurements of the anomalous magnetic moment of the electron. Here high order loop calculations and extremely sensitive experimental setups have shown to agree to 10 significant digits ([KN06] and [MT05]). Due to its higher mass, the corresponding anomalous magnetic moment of the muon (the famous "g-2" parameter) is more sensitive to physics at higher mass scales, and can not be calculated entirely within QED, but receives significant loop contributions due to the more massive particles of the SM. This also means that this quantity in principle should be sensitive to effects from new physics beyond the SM, and the latest measurements of it in a spin-precession experiment [GM204] shown in figure 4.1 have received a lot of attention. Unfortunately it turns out that the theoretical uncertainties are too large to be able to conclude anything. The reason is that coloured particles appears in the loop diagrams, which means that parts of the relevant amplitudes will contain non-perturbative quantities. It is, however, in principle possible to relate these incalculable parts of the formulae to other processes such as $e^+e^- \to q\bar{q}$ cross sections and hadronic decays of taus, and use experimental data from observations of these to correct the theoretical predictions for the muon anomalous magnetic moment. Unfortunately it appears that these two methods are not compatible and the final result is a 2.7σ deviation from the SM if the $e^+e^- \to q\bar{q}$ data are used but only 1.4σ if the tau decay data are used. It is thus hard to conclude anything else than agreement with the SM based on the present status.

In the early eighties the knowledge of electroweak model was based upon a series of indirect measurements, such as the weak decays of hadrons and neutral current scattering [GAR73]. The direct observations of the W [UA183a][UA283b] and Z [UA183b][UA283a] bosons at CERN around 1983 thus signified a major milestone and put the theory on solid ground. But the to date most impressive verification of the electroweak sector of the Standard Model is based upon data from e^+e^- annihilation data from high-energy experiments where the degrees of freedom associated with the massive, $\mathcal{O}(80-90\ \text{GeV})$, weak gauge bosons become active. Recently a combined analysis of the relevant data was published ([LEP05] and updates from summer 2006), mainly focusing on data from the LEP and SLC experiments. A particularly

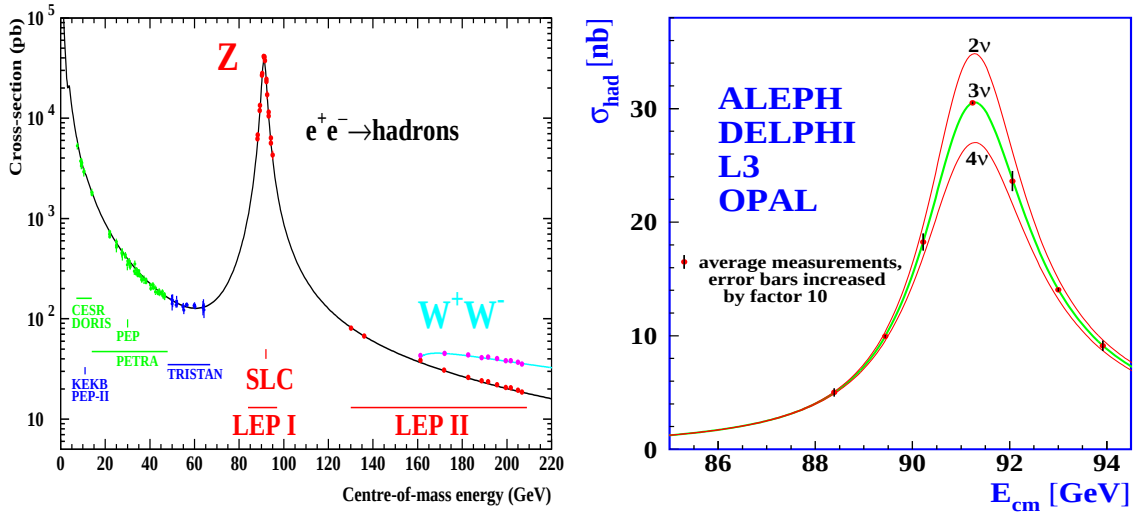

(a) (b)

Figure 4.2: *Cross section for $e^+e^- \to \gamma \to q\bar{q}$ as a function of centre of mass energy. Figure (a) shows the cross section in the entire range of its measurement and includes data from several different experiments. Figure (b) is a close-up of the Z resonance showing data from the four LEP experiments versus different theoretical predictions depending upon the number of light neutrinos. Source [LEP05].*

nice measurement that includes effects from all four electroweak gauge bosons is that of the inclusive cross section for e^+e^- into hadrons, σ_{had} as a function of the centre of mass energy shown in figure 4.2.a. At energies below $\mathcal{O}(50\text{ GeV})$ the contribution from the massive gauge bosons is negligible, and the dominant process is $e^+e^- \to \gamma \to q\bar{q}$ whose cross section falls off quadratically with centre of mass energy as predicted by the theory. At energies around the mass of the Z, $m_{\text{Z}} \approx 91.2$ GeV the process $e^+e^- \to \text{Z} \to q\bar{q}$ is highly resonant, and the cross section rises dramatically as this process takes over (note that the curve shown appropriately takes interference between the γ and Z channels into account). After the resonance, the cross section continues to fall off quadratically with centre of mass energy, but is higher than an extrapolation from the original photon line would have predicted as the Z channel is now comparable to the γ channel. At a threshold energy of $2m_{\text{W}} \approx 161$ GeV the process $e^+e^- \to \text{Z} \to \text{W}^+\text{W}^-$ becomes kinematically accessible. The exact turn-on shape represents a direct test of the gauge boson self-interactions shown in figure 3.5.

Assuming universality in the couplings between leptons and the Z boson, the exact shape and height of the Z resonance can furthermore be used to test the picture of 3 fermionic generations provided by the SM. Figure 4.2.b shows a close-up of the data at the resonance overlaid with three different theoretical curves depending upon the number of neutrinos in the model: The more different neutrino types, the more of the Z boson branching ration will be "eaten" by neutrino decays, and less of it will be available for the observed σ_{had}. The data are in remarkable agreement with the 3 Standard Model neutrino types, and seriously constrains attempts to expand it with a fourth generation. A fourth neutrino would have to be quite different from its three known cousins - it would either have to be heavier than $m_{\text{Z}}/2$ or have a very small coupling to the Z boson.

Many other tests of the electroweak sector have been carried out at e^+e^- colliders such as LEP. For instance, it is not surprising that the chiral nature of the weak couplings predicts

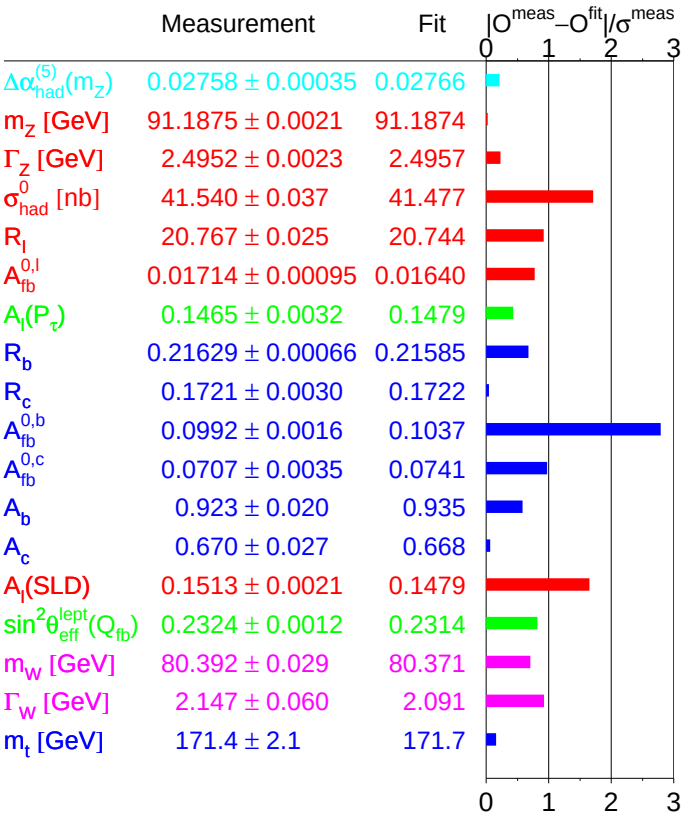

Figure 4.3: *Absolute pulls of various electroweak observables from e^+e^- interactions. Pull is here defined as the difference between the direct and indirect measurement of an observable in units of its experimental error. Source: Summer 2006 update of [LEP05].*

asymmetries in weak decay angle distributions. To see how well the Standard Model agrees with all of the data, and if a consistent tuning of parameters is possible, one can compare the directly measured values of the various observables with that inferred indirectly through their theoretical relationships to other observables. The result of such a fit is shown in figure 4.3. Even though a single parameter has deviations almost at the level of 3σ, the global χ^2/N_{dof} is 18.3/13 corresponding to a significance level of 15%, which means that all deviations from the SM are consistent with being statistical fluctuations.

For many years the experimental knowledge of the CKM matrix in equation (3.6) was rather limited, mainly because its third column and row is only directly accessible in experiments with the heavy third generation quarks. This meant that knowledge about e.g. the CP violating complex parameters were only accessible as extremely small effects in the decays of the neutral kaons [CCFT64]. However, with the advent of the experiments BaBar [HQ98] [BBR02] and Belle [BEL98] at asymmetric B-factories around the turn of the century, this situation changed. Here, e^+e^- colliders produced $B\overline{B}$ pairs through the $b\bar{b}$ resonance $\Upsilon(4S)$, at a relative velocity compared to the laboratory frame. This allowed precise vertex measurements of the decays to be turned into precise decay time measurements, and with the multitude of possible B decays, many independent entries of the CKM matrix became accessible.

One effort to combine the multitude of data regarding the CKM matrix is presented in [CKM05]. Figure 4.4 shows the present constrains on the parameters ρ and η. The different

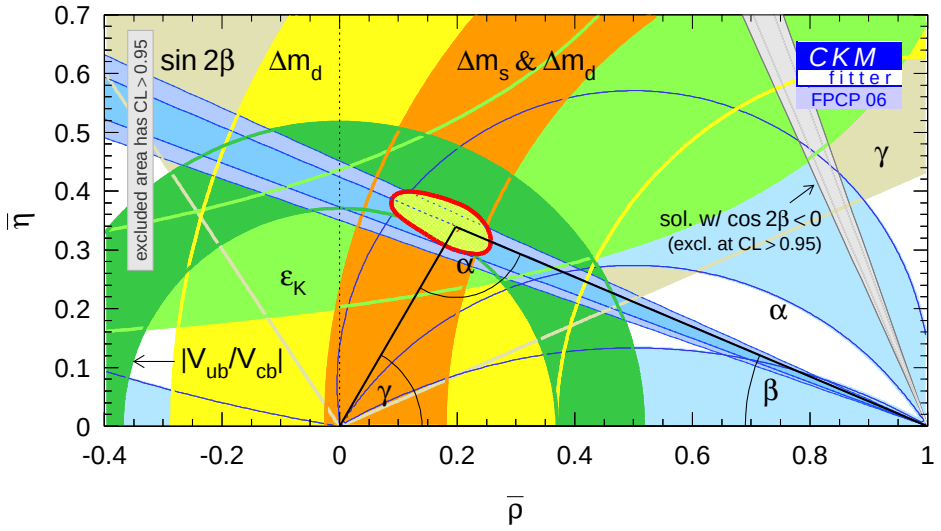

Figure 4.4: *Experimental constrains on* $(\bar{\rho}, \bar{\eta}) = [1 - \frac{1}{2}\lambda^2]^{-1}(\rho, \eta) + \mathcal{O}\left(\lambda^4\right)$ *inferred from a variety of measured observables. Source: Summer 2006 update of [CKM05].*

allowed regions arise from measurements of a variety of different quantities in an assortment of different analyses under the assumption that these are correctly related by the 3-by-3 unitary CKM matrix. If the various bounds had not been found to all be consistent with a single point, this would have meant either that the 3 generation CKM matrix would not be unitary (e.g. that it would really be a sub-matrix of a larger unitary matrix), or that the theoretical assumptions entering into one or more of the individual measurements would be invalid due to the presence of new physics such as new particles in loop diagrams.

4.1 Status of Higgs Boson Searches

Despite the impressive results of the electroweak model, its final verification is still pending, however, as experiments have so far never been able to verify a direct production of the Higgs boson, which is an unavoidable prediction of the model as discussed in section 3.3. The strongest direct limits on its mass is presently based on the final data samples from LEP at centre of mass energies up to around 209 GeV, where the Higgs boson could have been produced in so-called Higgs strahlung events, $e^+e^- \to \mathrm{Z} \to \mathrm{Z}h$. Despite tantalising reports of 3σ evidence for a Higgs boson at masses around 114 GeV in the dramatic final period of the running of the experiment [AL00], this signal did not persist after more careful studies and combinations of data from all the LEP experiments. The final result was instead a direct exclusion of the Higgs boson at masses lower than 114.4 GeV (95% C.L.) [LEP03].

It is, however, still possible to constrain the mass of the Higgs boson through its indirect relationship to other observables. Figure 4.5.a shows the preferred value at the minimum of the reduced χ^2 of the indirect fit as well as the directly excluded region. The indirect data actually have the peculiar feature of favouring a mass below the direct exclusion limit, although not at a significant level. It should be stressed that figure 4.5.a does not give a proof of the existence of the Higgs boson, as only a direct observation could do that. Instead, it shows that the direct lower limit on its mass does not exclude all of the values preferred by indirect data, meaning that the Higgs mechanism could still explain all observations. Furthermore, it gives valuable

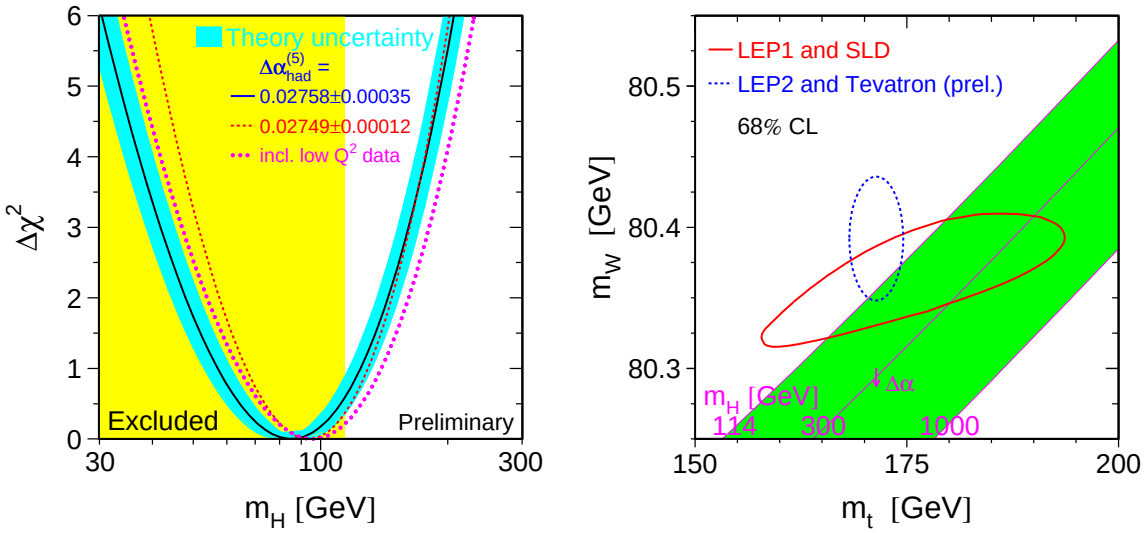

(a) (b)

Figure 4.5: *Current experimental knowledge about the Standard Model Higgs boson. Figure (a) shows the directly excluded region corresponding to $m_h > 114$ GeV in yellow along with the indirect result from a combined fit to electroweak observables. Figure (b) shows explicitly the correlation between the observed values of the W and top masses and the inferred Higgs mass. Source: Summer 2006 update of [LEP05].*

information about the range of possible values of the Higgs mass, which is $m_h < 166$ GeV or $m_h < 199$ GeV (both at 95% C.L.) depending on whether the direct limit is used to push the indirect limit upwards or not.

Of particular interest among observables showing a relationship with the mass of the Higgs boson are the masses of the W boson and top quark. At lowest order, the W mass is indeed given by equation (3.15), but it will in fact receive small loop corrections that are mainly dependent on the masses of the Higgs boson and top quark. Not only does this again mean that the Higgs mass can be predicted indirectly as shown in figure 4.5.b, but it will provide a very transparent test of the Higgs mechanism if and when the Higgs boson is discovered.

While the SM is not in any statistically significant way ruled out by the allowed region in figure 4.5.b, the region is actually in very good agreement with the supersymmetric extensions of the SM. In these theories, that will be discussed in more detail in section 5.4, the new symmetry explicitly forbids the Higgs couplings of the ϕ^c type discussed in section 3.3, so in these models it is always necessary to have (at least) two Higgs doublets in order to provide masses to both up and down-type fermions. Furthermore, these theories usually predict that four of the resulting Higgs bosons, H^0, A^0 and $H^\pm$, will be heavy compared to h^0, which from an experimental point of view will look very similar to a SM Higgs boson. The couplings of the h^0 will, however, not always be exactly similar to its Standard Model counterpart, but might be smaller. In fact, when $\tan\beta$ is larger than about 5, the limit on the Higgs boson mass in these models will be reduced from $m_h > 114$ GeV to about $m_h > 90$ GeV [Djo05].

In any case, if the Higgs boson has Standard Model like couplings, LHC should be able to

observe it over the entire allowed range up to the unitarity bound $m_h \approx 800$ GeV and even a bit further, with just a few years of data taken at full luminosities [ATL99b] [LP99].

4.2 Neutrino Masses

In the Standard Model neutrinos are by construction massless, so when the first experimental evidence proved otherwise [SK98], it was essentially a proof of physics beyond the Standard Model. However, as the Standard Model is readily extended to include non-zero neutrino masses also, the discussion of neutrino masses will be taken in this section.

If the neutrinos do not have degenerate masses, the leptonic sector of the Standard Model will need to be extended with matrices for basis change between mass and flavour eigenstates similar to those for the quark sector in equation (3.2). In the leptonic sector this matrix is usually denoted the Maki-Nakagawa-Sakata (MNS) matrix [MNS62], and is given by

$$V_{\mathrm{MNS}}{}^{ij} = U_\ell^{\dagger\, ik} U_\nu^{kj} \tag{4.1}$$

Where U_ℓ and U_ν denotes the transformation between mass and flavour eigenstates for the charged leptons and neutrinos, respectively. The introduction of the matrix requires 4 new parameters: 3 real mixing angles and one CP violating complex phase, in addition to the neutrino masses. To show how this will result in experimentally observable neutrino flavour oscillations, the present discussion will for simplicity be limited to 2 generations, with straightforward generalisations to the 3 generation case. In this case the basis changing matrix will only contain one real parameter and no complex phases. Thus, denoting the two mass eigenstates corresponding to masses m_1 and m_2 with ν_1 and ν_2 respectively, and denoting the two flavour eigenstates with ν_α and ν_β (these could for instance represent $\nu_{e,L}$ or $\nu_{\mu,L}$), their relationships can be expressed by the real mixing angle ϕ:

$$\begin{bmatrix} |\nu_1\rangle \\ |\nu_2\rangle \end{bmatrix} = \begin{bmatrix} \cos\theta & \sin\theta \\ -\sin\theta & \cos\theta \end{bmatrix} \begin{bmatrix} |\nu_\alpha\rangle \\ |\nu_\alpha\rangle \end{bmatrix} \tag{4.2}$$

If at time $t = 0$ a neutrino is created in the flavour eigenstate, its time evolution will be given by

$$e^{i\widehat{H}t}|\nu_\alpha\rangle = e^{iE_1 t}\cos\theta|\nu_1\rangle - e^{iE_2 t}\sin\theta|\nu_2\rangle \tag{4.3}$$

where $E_i^2 = p^2 + m_i^2$. It follows that the probability for an experiment at a later time to observe a ν_β flavour state is:

$$\begin{aligned} P_{\nu_\alpha\to\nu_\beta}(t) &= \left|\langle\nu_\beta|e^{i\widehat{H}t}|\nu_\alpha\rangle\right| = \sin^2 2\theta \sin^2\left[\frac{E_2 - E_1}{2}\right] \\ &\approx \sin^2 2\theta \sin^2\left[\frac{(m_2^2 - m_1^2)t}{4p}\right] \\ &\approx \sin^2 2\theta \sin^2\left[1.27\frac{(m_2^2 - m_1^2)[\mathrm{eV}^2]L[\mathrm{km}]}{E[\mathrm{GeV}]}\right] \end{aligned} \tag{4.4}$$

Where it has been used that neutrinos are always ultra relativistic in all experimental circumstances and where L is the distance travelled by the neutrino. Even though the neutrino masses are small enough to be outside direct experimental reach, equation (4.4) allows one to

probe very small mass differences by observations over very large differences. As an example, for neutrinos of $E \approx 2$ MeV, oscillation lengths of 1000 km will correspond to mass squared differences of 10^{-5} eV^2.

Numerous experiments in the past decade have been conducted studying oscillations in neutrinos originating from the sun, from cosmic rays (so-called atmospheric neutrinos), from nuclear reactors or from particle accelerators, and with different target materials, thus measuring transitions between different neutrino species for various values of E and L. The analyses usually need detailed Monte Carlo simulations, taking into account issues such as the precise geological structure of the earth, energy spectra of the neutrinos as well as the relevant interaction cross sections of the target material. Details can be found in the recent review [MSTV04] and references therein, but a lot of the present data is nicely summarised in figure 4.6, which shows that the different experiments are consistent with two different solutions. The first corresponds to a squared mass splitting around $8 \cdot 10^{-5}$ eV^2 and a mixing angle around 30 degrees and the second corresponds to a squared mass splitting around $3 \cdot 10^{-3}$ eV^2 and a mixing angle around 45 degrees. The third mixing angle has not been directly measured and has an experimental upper limit around 10 degrees. This also means that the CP violating phase is so far unmeasured.

All in all, the data gives a picture of the neutrino mass spectrum as shown in figure 4.7. As only the absolute relative mass splittings are constrained, two hierarchies are possible. Most importantly, however, is that there is so far no experimental measurement of the absolute scale of the masses. The strongest direct bounds are given on the electron neutrino mass (which according to figure 4.7 is a combination of m_1 and m_2), due to analyses of the energy spectrum of electrons from beta decays of Tritium near the kinematic endpoint. The results from the two most recent and most accurate experiments, Troitsk and Mainz are upper limits of $m_\nu < 2.05$ eV and $m_\nu < 2.2$ eV respectively [Lob03]. The next generation experiment, KATRIN, aims to improve this limit to around 0.2 eV (90% C.L.), with possibility of a 5σ discovery down to 0.35 eV [KAT04].

The most popular way to include non-zero neutrino masses in a manner which naturally accounts for their very small values is by assuming that the right-handed neutrino (which has no interactions, and is thus "sterile"), in addition to the typical Dirac mass, m_D, shared with the left-handed field, has a large Majorana mass term M. If furthermore the left-handed field itself is without such a Majorana mass term (or if it at least negligible compared to the small observed neutrino masses), the particles in the mass eigenstate will be a mainly sterile ν_R of mass $\sim M$, and one mainly weakly interacting left-handed neutrino ν_L of mass $(m_D/M)m_D$.

Even if the Dirac mass, m_D, is not smaller than the typical masses of the quarks and charged leptons, the small resulting value of the mass of the left-handed neutrino is thus explained by the huge value of the mass of the right-handed field. This mechanism is consequently denoted the *seesaw mechanism* [Yan80], and is popular because large M are easily accommodated in some extensions of the SM (cf. section 4.4). Another way to accommodate the observed masses in the SM, which would not necessitate right-handed neutrinos would be to simply add a small Majorana mass term for the left-handed neutrinos. In this case there would be no apparent mechanism to explain its small value. One consequence of such a Majorana mass is that a rare version of the double beta decays, $2n \to 2p2e^-2\bar{\nu}_e$, would be possible in which no neutrinos appear in the final state. These would be characterised experimentally by the absence of the missing momentum associated with neutrinos, i.e. a peak in the summed energy spectrum of the electrons. In addition to dedicated experiments (cf. [NEMO05] and [P$^+$06]) other beta decay experiments such as KATRIN will be looking for these as well. A

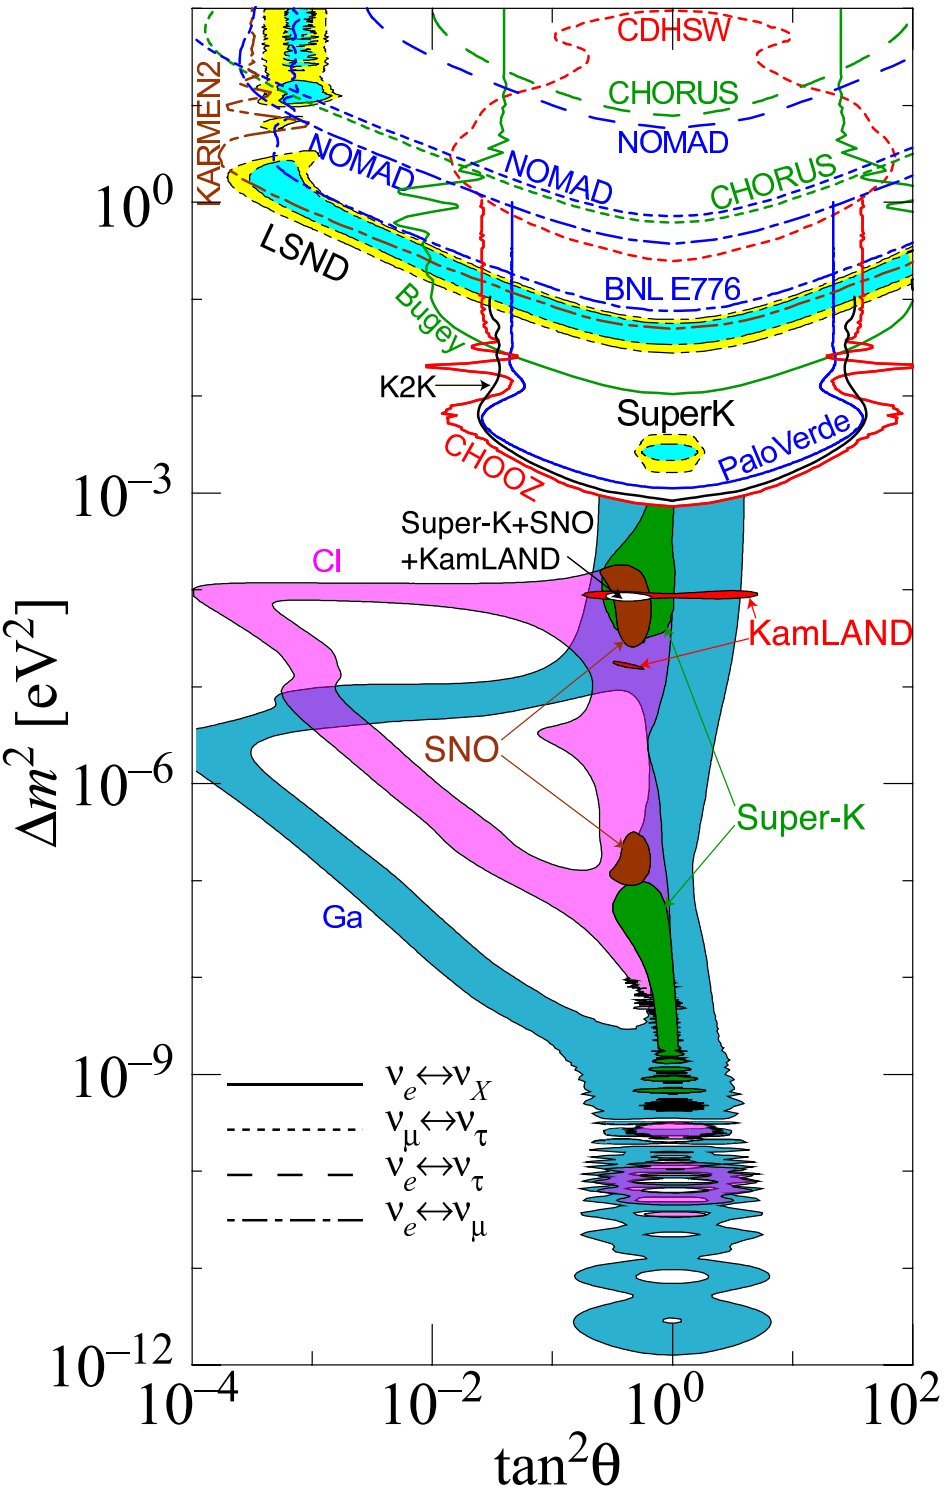

Figure 4.6: *Experimentally allowed regions of neutrino mixing angles and squared mass splittings (all at 90% C.L.). All data are consistent with two solutions. The figure is borrowed from* `http://hitoshi.berkeley.edu/neutrino`, *from where detailed references are available.*

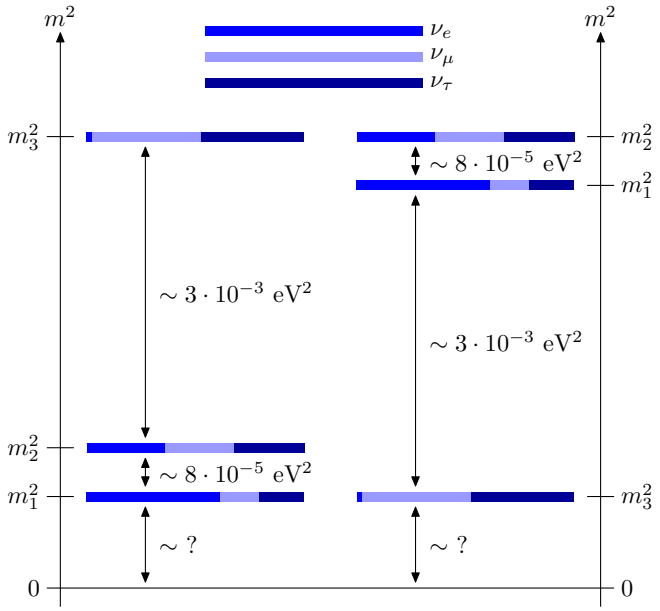

Figure 4.7: *Possible neutrino mass spectra.*

seesaw mass would be harder to confirm experimentally (but not impossible [BM06]).

4.3 Cosmological Motivations for Extensions of the Standard Model

Using astrophysical and cosmological observations as input to particle physics models is always a little dangerous, as the conclusions that can be drawn from such observations are invariably model dependent and indirect. However, they provide access to energy and distance scales that are clearly out of reach for laboratory experiments in addition to providing unique knowledge regarding the birth of the universe (i.e. its boundary conditions). Additionally, it would seem that an increasing number of unresolved puzzles in astrophysics requires progress in particle physics to solve.

The cosmological standard model of inflationary big bang originally motivated by observations such as Hubble expansion and the apparent flatness of the universe ($\Omega = 1$) have in the last decade truly entered a precision era with highly sensitive measurements of the cosmic microwave background (cf. figure 4.8.a) as well as observation of highly luminous supernovae standard candles (cf. figure 4.8.b). Combined with observations of galaxy rotation curves and insights into the large scale structure of the universe, a consistent, precise and somewhat surprising picture of the cosmological composition of the universe has emerged (cf. e.g. [SDSS04] and [Pad03] for reviews). Most impressively it manages to describe a large variety of phenomena from very few cosmological parameters. To high precision it is consistent with $\Omega = 1$, with contributions from regular baryonic matter, neutrinos, cold dark matter and the enigmatic dark energy broken down roughly as:

$$\begin{aligned} \Omega_\Lambda &\approx 73\% \\ \Omega_{\mathrm{CDM}} &\approx 23\% \\ \Omega_{\mathrm{B}} &\approx 4\% \end{aligned} \tag{4.5}$$

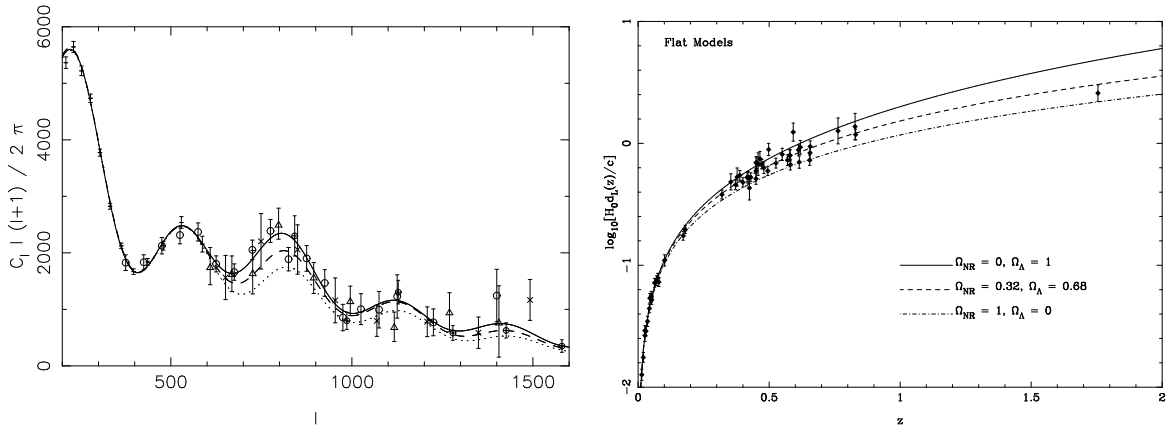

(a) (b)

Figure 4.8: *Two of the recent achievements of observational cosmology. Figure (a) shows the higher order angular power spectrum of the cosmic microwave background, including the predictions from a couple of competing theories (cf. [SMS05] and references therein). Figure (b) shows how observations of very distant supernovae standard candles indicates an accelerating universe, clearly inconsistent with* $\Omega_\Lambda = 0$ *(cf. [Pad03] and references therein).*

Thus, the main contribution to Ω is thus due to "dark energy", i.e. a negative pressure exerted by the vacuum counteracting gravity at large distances (leading to the observed accelerated expansion of the universe). This dark energy can be parametrised by the introduction of a cosmological constant into Einsteins equations of general relativity, but is completely unexplained by fundamental particle physics theories. In fact, the energy density it represents is typically 120 orders of magnitude smaller than the value predicted by such theories! An additional problem with the observation is that a universe with non-zero cosmological constant would start by being heavily matter dominated $\Omega_\Lambda \approx 0$, but with its expansion the matter density would fall, while the density due to the cosmological constant remain the same. This means that at large times $\Omega_\Lambda \approx 1$. The observation of $\Omega_\Lambda \approx 0.7$ thus indicates that the current time is very special, which is somewhat unsettling.

The second largest contribution is from cold dark matter, which is matter consisting of non-relativistic (i.e. heavy) particles that at most interact weakly with normal matter. Presently, all of the candidates for the content of dark matter that are based in the Standard Model, such as massive neutrinos or non-luminous astrophysical objects such as black holes, are ruled out in the sense that they can not contribute significantly to the observed value of Ω_{CDM}. The most plausible candidate for cold dark matter would thus seem to be some sort of weakly interacting massive particle (a WIMP), such as e.g. a new neutrino with mass larger than 50 GeV. The Standard Model does not provide any such WIMP candidate, but if one were to be discovered at e.g. the LHC it would be a tremendous achievement for particle and astrophysics.

Although the last few percent of Ω, represented by Ω_{B}, would seem to consist of matter that is contained in the Standard Model (appearing as planets, stars and galaxies), its value is also a bit of the mystery. The inflationary big bang model essentially predicts that equal amounts of baryonic and anti-baryonic matter should exist at the end of inflation. During the subsequent expansion and cooling of the early universe, the baryonic and anti-baryonic matter should annihilate symmetrically leaving essentially $\Omega_{\mathrm{B}} \approx 0$ (usually referred to as the problem of *baryogenesis*). This scenario is only avoided if matter and anti-matter asymmetries appears

in such a way that the following criteria, put forward by A. D. Sakharov in 1967 [Sak67], are satisfied in the early universe:

Baryon number violation An obvious prerequisite for going from a baryon-antibaryon symmetric state to a state with predominantly baryonic matter.

Violation of C and CP symmetries If C or CP invariance holds, then the rate of processes producing net excesses of baryons will be equal to the rate of processes producing net excesses of antibaryons.

Departure from thermal equilibrium Departure from thermal equilibrium is needed as the baryons and anti-baryons have equal masses (due to CPT invariance), and thermal equilibrium therefore enforces matter-antimatter symmetry even in the face of fundamental asymmetries.

The first and third criteria are usually easily satisfied in theories of inflation (with the baryon number violation in the SM due to vacuum effects at ultra high energies), and C is already known to be maximally broken in the weak interactions. Unfortunately the CP violation in the Standard Model is too small to play a role for baryogenesis, again indicating the inadequacy of the Standard Model in the description of physics at the energy scales outside direct experimental reach.

A seemingly unrelated puzzle is the observations of cosmic ray particles of ultra high energies above $\mathcal{O}\left(10^{11}\text{ GeV}\right)$ [HI93]. Due to the interactions between ultra high-energy protons and the 2.7 K photons in the cosmic background radiation, these particles must originate within 10^8 ly, and thus be of non-cosmological origin. There is presently no known mechanism in astrophysics based on the Standard Model that could explain these observations.

In addition to these observations, it seems obvious that phenomena such as inflation and the singularity in black holes ultimately requires clarifications from particle physics.

4.4 Theoretical Motivations for Extensions of the Standard Model

Despite the tremendous success of the Standard Model, it has several inherent theoretical issues that shows that it can not be a final theory. This is immediately apparent as it not only, like any quantum field theory, works in the context of an underlying cutoff at high energies and small distance scales, but it also completely fails to provide a fundamental description of gravity. The reason for this is that quantum field theories reproducing General Relativity at large scales invariably are not renormalisable (cf. e.g. [D+74]). A direct consequence of this is that any quantum field theoretical description of Nature must break down at the scale where quantum effects of gravity can no longer be neglected. This so-called *Planck scale*, Λ_{Planck}, is the scale where the gravitational energy of two particles of masses Λ_{Planck} placed a Compton wavelength apart is equal to their self energy, i.e.:

$$\left|V_g(m_1 = m_2 = \Lambda_{\text{Planck}}, r = \Lambda_{\text{Planck}}^{-1})\right| \equiv \Lambda_{\text{Planck}} \tag{4.6}$$

This results in

$$\Lambda_{\text{Planck}} = \frac{1}{\sqrt{G_N}} \approx 1.22 \cdot 10^{19}\text{ GeV} \tag{4.7}$$

Where G_N is the gravitational constant. There are, however, several theoretical indications that the Standard Model will not be adequate for a description of physics all the way from the electroweak scale, $\mathcal{O}(100\text{ GeV})$, to the Planck scale.

In fact, the so-called *hierarchy problem* strongly indicates the presence of new physics just around the corner, at $\mathcal{O}(1\text{ TeV})$. To understand this remember that, as discussed in section 2.2, the mass, m_0, of a particle defined at the cutoff scale, Λ, of a theory and the mass observed at low energies, m, will not be identical due to the quantum corrections of the theory. For fermions, the difference between the two is of the same order of magnitude as the mass itself [Daw06]:

$$\delta m_f = \mathcal{O}\left(m_f \log \frac{\Lambda}{m_f}\right) \tag{4.8}$$

Which, once the consequences of renormalisation are accepted, is unproblematic. For a scalar particle like the Higgs boson, however, the corrections will be of the order of the cutoff scale instead:

$$\delta m_h^2 \equiv m_h^2 - m_{h0}^2 = \mathcal{O}\left(\Lambda^2\right) \tag{4.9}$$

This means, if $\Lambda = \Lambda_{\text{Planck}}$, that the value of m_{h0} must be incredibly fine-tuned if m_h is to end up around $100 - 200$ GeV as required by the model (cf. section 4.1). As an example, if m_{h0} is tuned to provide m_h in the required range, then a redefinition of $m_{h0} \to m_{h0}[1 + 10^{-12}]$ would raise m_h to $\mathcal{O}\left(10^9\text{ GeV}\right)$. The only natural mass value for a scalar particle such as the Higgs boson is thus $m_h = \mathcal{O}(\Lambda_{\text{Planck}})$. The obvious way to avoid this fine-tuning problem would be if some new physics appeared at a scale, Λ, not much higher than m_h. This new physics might thus be open for direct experimental investigations at the next generation of accelerator based experiments starting with the LHC, and a dedicated discussion of the available models will take place in sections 5 and 6.

A similar fine-tuning problem actually appears in the strong interactions. As mentioned in section 2.3.1, it is possible to include CP violating terms in non-Abelian gauge symmetries such as QCD. These will effectively correspond to the introduction of a CP violating phase, θ. At the classical level, a chiral rotation of the quark fields can completely remove this phase, but the QCD vacuum breaks this classical result and obtains a phase, θ_{QCD}. The effective CP violating phase in the Lagrangian will be the difference between this phase and the phase of the quark mass matrix, $\overline{\mathbf{M}}$ (after the masses are provided by the electroweak symmetry breaking):

$$\theta_{\text{eff}} = \theta_{\text{QCD}} - \arg\det\overline{\mathbf{M}} \tag{4.10}$$

As the two terms on the right side have different origin, it is natural to expect that $\theta_{\text{eff}} = \mathcal{O}(1)$, implying CP violation in the strong interactions. The most stringent bound of this comes by comparing the limit on the electric dipole moment of the neutron (cf. [PDG06]), with the theoretical predictions given θ_{eff} (cf. [Bal79] and [C+79]), and is:

$$\theta_{\text{eff}} < 1 \cdot 10^{-10} \tag{4.11}$$

This apparent unnatural elimination of the effective CP violating phase in QCD is known as the *strong CP problem*. The reader is referred to the excellent overview of this problem, its proposed solutions and their status given in [Pec98]. Of particular note is one of the first

proposed solutions to this mystery, which was put forward already in 1977 [PQ77]. In this model the idea is to connect θ_{eff} to a broken global $U(1)$ symmetry and dynamically ensure extremely small values of θ_{eff} in a manner similar to the Goldstone mechanism. A side effect of this would be the presence of a new type of very light boson, the *axion*, which would be very light and have extremely small strong and weak interactions. It should, however, couple to photons in the presence of very strong magnetic fields, and as such they can be searched for experimentally. In fact, one such experiment, CAST, is presently using a decommissioned LHC magnet to search for axions produced in the core of the sun [CAST05]. What makes this even more interesting is the fact that the axion could also be an important component of dark matter.

In addition to these crucial fine-tuning problems, there are a series of intriguing questions left open by the Standard Model which, while not directly devastating, one would hope would be answered in a more fundamental theory. First of all it is unsatisfying that there are so many free parameters left in the theory: Gauge couplings, Higgs couplings (i.e. masses), the CKM parameters and the parameters associated with the electroweak symmetry breaking. Furthermore, the structure of their values is puzzling: Why is the top quark the only fermion with a Higgs coupling of $\mathcal{O}(1)$? Why are all the other fermions so light? And why do they still have such a large spread between e.g. the b quark and the electron. Why is the CKM matrix so diagonal? The fact that it is means that the "rotation" between the weak eigenstates of the up-type quarks and the states where their Higgs couplings are diagonal is very similar to the corresponding rotation of the down-type quarks. Furthermore, the form of equation (3.6) means that the overlap between quark states is small for the first and second generation, even smaller between the second and third generation and very small between the first and third generation. This hierarchical structure means that CP violating effects in the SM are small, and one could wonder whether there is a reason for this hierarchy? On that note, one could wonder why the fermions seems to be grouped into generations,[1] why there are exactly three (the minimum number required for the CKM mechanism to include CP violating effects), and why the masses between the generations also seems to follow a hierarchy? It is also puzzling why the one gauge group, $SU(2)$ which seems to be chiral in nature is the one that is spontaneously broken? In the light of Coleman and Mandula's theorem, it is also surprising that an internal gauge symmetry should know at all about the chirality of fermions, which is a spacetime property. Finally the gauge structure of the couplings, $SU(3) \times SU(2) \times U(1)$, seems like a rather arbitrary choice.

One esthetically pleasing idea which was proposed already in the seventies (cf. [GG74] and [PS73]), that could potentially lead to an improved understanding of a few of the questions above is based on the observation that the running coupling constants of the three known gauge groups seem to converge towards a common value at scales around $\mathcal{O}\left(10^{15}\text{ GeV}\right)$. This opens for the intriguing possibility that the fundamental gauge group is not $SU(3) \times SU(2) \times U(1)$, but rather some larger encompassing *irreducible* group, G, as such a group is only associated with a single coupling constant. The reason the low-energy gauge structure of the theory looks like $SU(3) \times SU(2) \times U(1)$ instead of G would have to be some sort of spontaneous symmetry breaking taking place at $\mathcal{O}\left(10^{15}\text{ GeV}\right)$. With the notable exception of gravity, such a theory would thus provide a unification of all of the fundamental interactions and is usually referred

[1] In fact, the different fields in a complete generation have *exactly* the quantum numbers needed to cancel out the potentially dangerous *gauge anomaly* in the theory. This fact certainly does not seem coincidental, and thus further emphasises the fact that the generational structure points towards the existence of some deeper underlying theory.

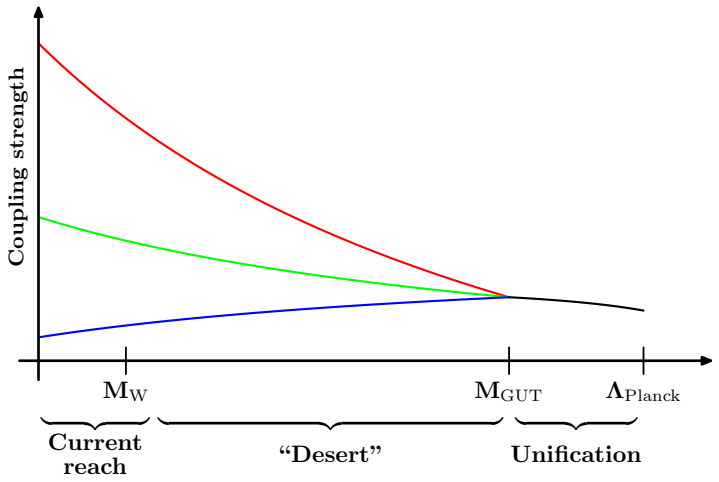

Figure 4.9: *Illustration of unification of the gauge couplings. The annotation "desert" indicates the assumption inherent in the extrapolation outside the region currently reached experimentally, namely that no new physics between the weak and the GUT scale will ruin the unification.*

to as a "Grand Unification Theory", with the corresponding energy scale denoted M_{GUT}. The unification of the gauge couplings is illustrated in figure 4.9.

It is safe to say that no entirely self consistent and satisfactory GUT theory is known presently, but two longstanding candidates for a unifying gauge group are $SO(10)$ and $SU(5)$. In $SU(5)$ theories, the fermions in a each generation are combined into two different representations (in the same way that the fermion fields in each generation are combined in two doublets and several singlets under the weak $SU(2)$), and in $SO(10)$ theories all the fermions in one generation are combined into one single representation, which notably needs to include a right-handed neutrino, which would naturally get a very large Majorana mass of $\mathcal{O}(M_{\text{GUT}})$, needed for the seesaw mechanism as discussed in section 4.2. This is a strong point in favour of such a GUT. The combined representation of the fermions ultimately means that the Higgs couplings of the fields becomes related, thus reducing the number of free parameters in the theory and predicting mass relations such as m_e/m_u or m_τ/m_b, but unfortunately not in very good agreement with the observed values. Furthermore, some of the gauge bosons of the new group will invariably connect the leptons and quarks within each generation, leading to baryon and lepton number violations. However, below the GUT scale the spontaneous symmetry breaking will have supplied masses to these gauge bosons around M_{GUT}, meaning that such couplings will be very suppressed at low energies. But the stringent bounds on the proton lifetime nevertheless impose tight limits on proposed GUT theories, and have more or less ruled out those based on $SU(5)$. Finally, it should be noted that the running coupling constants do not exactly meet in one point as otherwise suggested by figure 4.9. This might not be a problem as the hierarchy problem indicate the presence of new physics at the TeV scale which might affect the precise running (in fact, a feature of the supersymmetric theories to be discussed in section 5.4 is that they improves this unification considerably [GKL91]). Another possibility would be for the spontaneous symmetry breaking at the GUT scale to happen in two steps.

A final interesting point is that GUT theories are known to predict the existence of magnetic monopoles, making the experimental search for such objects at e.g. the LHC particularly interesting.

5 Solutions to the Hierarchy Problem – Models Beyond the Standard Model

In this section some of the most promising new models of physics beyond the Standard Model that address the hierarchy problem are reviewed. Both because it is one of the biggest problems of the Standard Model, but in particular also because its nature almost guarantees that its solution will require new physics at the TeV scale, with potential impact on the next collider experiments starting with the LHC. After a review of the features that such models are generally expected to possess in section 5.1, particular models are reviewed in sections 5.2, 5.3 and 5.4.

5.1 General Features of Models

As discussed in section 4.4, the nature of the hierarchy problem is that radiative corrections tends to drive the mass of the Higgs boson towards the cutoff scale, $\Lambda_{\text{Planck}} \approx 10^{19}$ GeV, unless prevented by *extreme* fine-tuning. As this is unnatural, new physics must enter the scene and prevent this at much lower scales. In fact, if the tuning should remain as small as possible, the new physics should enter at scales very close to m_h which is expected to be around 100 GeV to 200 GeV. It follows from equation (4.9) that a relative change in m_{h0} will result in a relative change in m_h that is inflated by (at least) a factor of M^2_{new}/m^2_h, where M_{new} is the mass scale of the new physics. So if one for instance requires that a 1% change in m_{h0} should not change the mass of the Higgs boson by more than 10%, the limit on the new physics becomes:

$$10\% < 1\% \left(\frac{M_{\text{new}}}{m_h}\right)^2 \quad \Rightarrow \quad M_{\text{new}} < \mathcal{O}\left(300 - 600 \text{ GeV}\right) \tag{5.1}$$

The main problem facing models at such low scales is that new physics close to the electroweak scale should affect the values of electroweak observables that are already well constrained to within $\mathcal{O}\left(0.1 - 1\%\right)$ by precision data as discussed in section 4. New particles in loops will be quadratically suppressed by their masses, implying large lower limits on the new scale:

$$\frac{M^2_{\text{weak}}}{M^2_{\text{new}}} < \mathcal{O}\left(0.1 - 1\%\right) \quad \Rightarrow \quad M_{\text{new}} > \mathcal{O}\left(1 - 3 \text{ TeV}\right) \tag{5.2}$$

Equations (5.1) and (5.2) taken together obviously represents a problem for models of new physics trying to address the hierarchy problem. Viable models of new physics falls into two categories: Those that try to circumvent the upper limit in equation (5.1) by changing the implicit assumptions leading to the hierarchy problem, and those that tries to circumvent the lower limit in equation (5.2) by inclusion of mechanisms that could suppress the indirect effect of new physics further than $m^2_{\text{weak}}/M^2_{\text{new}}$. One way to achieve this is by associating a global symmetry or, equivalently, a new charge, with the particles at the new scale, which would severely limit the interactions between SM and new particles (in the same way that conservation of lepton number forbids an $e\mu\gamma$ vertex). If all SM particles are neutral with respect to the new charge and all new particles are not, then all vertices with two SM particles and one new particles would be forbidden. These effects can lower, if not practically eliminate, the indirect limits in equation (5.2).

An appealing bonus feature of models with such a new global charge is that they can provide a natural candidate for dark matter, as the lightest of the new particles, being charged

under the new symmetry, is naturally stable. This again implies, that it must be electromagnetic and colour-neutral, or be a lot heavier than the TeV scale. If not, it should have shown up in searches for heavy isotopes or create cosmological problems by changing nucleosynthesis, or making neutron stars inherently collapse into black holes (cf. [ACD01, chap. 4] and references therein). Based on annihilation cross sections between the new particles it will usually be possible to calculate the relic density of these stable particles after the Big Bang freeze out (that is, after the temperature drops below the mass scale of the new particles), and the resulting prediction for Ω_{CDM} can be compared to the precise limits associated with equation (4.5). This can then be used as indication of which parameter values of the theory are sensible (if any). It should be noted, however, that such arguments would be obviously void if the dark matter candidate is not *incredibly* stable, with lifetimes larger than the present age of the universe.

Investigating the radiative corrections to the Higgs boson mass in detail, it turns out that one-loop diagrams involving top quarks, weak gauge bosons and the Higgs boson itself provide the main contribution to the radiative corrections. This is not surprising of course, as these are the fields with the strongest couplings to the Higgs field. If loops involving new particles are to cancel these contributions in a way that does not involve tuning of the new couplings, it is natural to expect that the new particles should be related to the SM particles through some symmetry principles. This means that it is not unreasonable to expect particles at the scale M_{new} with quantum numbers and couplings similar to those of the top quark and weak gauge boson and Higgs field itself. Other SM particles of course also provide radiative corrections to the Higgs boson mass, but as their effect is smaller the upper limit on their appearance could be significantly higher than that indicated in equation (5.1).

It should also be noted that the limits on violations of baryon and lepton number and CP and flavour invariance are much typically much stronger than 0.1%, imposing additional limits on the models of new physics.

A final note regarding the hierarchy problem is that it could also be "solved" by anthropic arguments, i.e. arguments stating that fine-tuned parameters have the values they need to have exactly because they lead to a universe wherein intelligent life could evolve and wonder about hierarchy problems. Such arguments are usually given in context of theoretical models that would allow Nature to "throw the dice" again and again (causally disconnected regions, multiple universes, multiple vacua, etc.), but which are often associated with low predictive power. While anthropic arguments are, per construction, very hard to disprove, they are in the authors opinion uninteresting for precisely the same reason.

5.2 Higgs Compositeness (Technicolour and Little Higgs models)

To understand the idea in Technicolour models one should start by realising why it is a problem that $m_h \ll \Lambda_{\mathrm{Planck}}$, but not a problem that $\Lambda_{\mathrm{QCD}} \ll \Lambda_{\mathrm{Planck}}$. Λ_{QCD} does not itself appear in the Lagrangian, but is generated dynamically with a value depending on the strong coupling constant. If written in its dimensionless normalised form, $\alpha_s \equiv g^2/4\pi$, then its value at a given scale, Λ, is related to Λ_{QCD} by the relation [Str06]:

$$\Lambda_{\mathrm{QCD}} \approx \Lambda \exp\left(-\frac{2\pi}{b\alpha_s(\Lambda)}\right) \tag{5.3}$$

Where b is a constant given by the details of the radiative diagrams of QCD, $b \approx 5$. It is thus quite easy for natural values of the coupling constant to induce a large hierarchy

between Λ_{QCD} and Λ. $g(\Lambda) = 1.0$ and $g(\Lambda) = 0.5$ will for instance result in $\Lambda_{\mathrm{QCD}}/\Lambda \approx 10^{-7}$ and $\Lambda_{\mathrm{QCD}}/\Lambda \approx 10^{-28}$, respectively. The idea in so-called *technicolour* models [FS81] is to use dynamics similar to that of QCD to generate masses for the weak bosons, which would thus naturally be much lighter than $\Lambda_{\mathrm{Planck}}$. In fact, if the Higgs boson did not break the electroweak symmetry, the QCD vacuum would have done it on its own. The pions would have been massless, and would have been eaten by the W and Z bosons, that would have ended up with masses, but unfortunately only in the tens of MeV range. The σ meson, a scalar state consisting of u and d quarks, would have been the Higgs particle in a sense. Inspired by this, the solution proposed in technicolour models is thus be to introduce a copy of QCD at a higher mass scale, Λ_{TC}, that could give rise to a Higgs particle and W and Z masses at the right weak scale, $\mathcal{O}(100\ \mathrm{GeV})$. This requires the introduction of new fermions, techniquarks, and a new strongly interacting gauge group associated with a new technicolour quantum number and force mediating technigluons. Just as in QCD one would expect the bound states of the techniquarks to produce an entire spectrum of technihadrons, and the lightest such would be a dark matter candidate unless technibaryon number is broken.

The main problem facing technicolour models is that Λ_{TC} can not be too high if the generated values of m_{Z}, m_{W} and m_h should be consistent with data. The new particles around the low value of Λ_{TC} in turn leads to conflicts with the indirect precision electroweak data. Proposed solutions to this problem alters the dynamics of the new gauge group to be less QCD like, allowing most of the new particles to appear at higher scales (for a recent example see [Pia06] and references therein).

A recent and not entirely unrelated idea is used in the *Little Higgs* models [STS05]. Here the Higgs boson mass is kept lower than the rest of the new physics by making it a pseudo Goldstone boson of some new broken symmetry. This is a priori not easy to do as the interactions of Goldstone bosons are limited (they are after all dynamically kept from acquiring masses through radiative corrections), and requires non trivial breaking of more than one symmetry at once if they should be able to play the part of the Higgs boson.

5.3 Extra Dimensions

A different type of solutions of the hierarchy problem is based on the observation that the high value of the Planck scale in equation (4.7) is simply an artifact of considering the problem in too few dimensions [ADD98]. Although direct experience as well as all observations so far seems to agree with a $3+1 = 4$ dimensional spacetime, there is no fundamental problems with the addition of additional spacial dimensions, as long as they are added in a way that does not contradict previous data and experience. One way to achieve this would be if the extra dimensions are *compactified* at a microscopic scale, i.e. if rather than being unbounded like $\mathbb{R}^n$ they would have the topological structure of a closed bounded surface of a small length scale R. While the strong and electroweak interactions have been tested down to length scales of approximately $(100\ \mathrm{GeV})^{-1} \sim 10^{-15}$ mm, the corresponding limits on gravitational interactions are much weaker, $\mathcal{O}(0.1\ \mathrm{mm})$.[1] If there is a mechanism that constrains the SM fields to the usual 4 dimensional spacetime, but allows gravity (i.e. gravitons) to propagate in all dimensions, these extra dimensions could be large, i.e. having length scales up to $R \approx 0.1$ mm. At scales $r \ll R$ it follows from Gauss' Law that the gravitational potential of two

[1] One recent precise measurement is [H+04], and an overview of the various existing and planned measurements is given in [Gun05].

test masses, m_1 and m_2 should go as:

$$V(r) \approx -G_{N(4+d)} \frac{m_1 m_2}{r^{1+d}}, \quad r \ll R \tag{5.4}$$

Where d is the number of extra dimensions. However, it is only the field lines in the macroscopic spacetime that will extend to distances $r \gg R$, and the effective strength of gravity will thus appear smaller:

$$V(r) \approx -k \frac{G_{N(4+d)}}{R^d} \frac{m_1 m_2}{r} \equiv -G_N \frac{m_1 m_2}{r}, \quad r \gg R \tag{5.5}$$

Where k is a constant, $\mathcal{O}(1)$ for low d, depending on the exact structure of the extra dimensions[1] and where it has been assumed that the extra dimensions are factorisable, i.e. their structure is the same at every point in the 4 spacetime (the surface of a cone would be an example of a surface where the extent of one dimension is not independent of the other).

Equations (5.4) and (5.5) taken together shows that the huge value of the Planck scale in 4 dimensions is really just an effective large distance scale. Using the definition of the Planck scale from equation (4.6), the true Planck scale is found to be (ignoring k):

$$\Lambda_{\text{Planck}}^{\text{true}} \sim {\Lambda_{\text{Planck}}}^{2/(2+d)} R^{-d/(2+d)} \tag{5.6}$$

This relationship implies that, for a fixed $d \geq 1$, the hierarchy $\Lambda_{\text{Planck}} \gg m_h$ can be traded away with a different hierarchy by instead raising the scale of the extra dimensions to $R \gg m_h^{-1}$. The point is of course that the latter hierarchy is not associated with a fine-tuning problem. The value of R that will result in a $\Lambda_{\text{Planck}}^{\text{true}}$ within the bound of equation (5.1) will depend on the number of extra dimensions. Requiring $\Lambda_{\text{Planck}}^{\text{true}} \approx 1$ TeV, the value of R will be given roughly by $10^{32/d-19}$ m. This means that $R_{d=1} \approx 10^{13}$ m is clearly wrong, $R_{d=2} \approx 1$ mm is around the present limit (~ 0.1 mm) of direct Cavendish like measurements looking for deviations from the $\sim r^{-1}$ behaviour of the gravitational potential. $R_{d=3} \approx 1$ nm, $R_{d=4} \approx 0.01$ nm, etc. are all well below the bounds. Stringent bounds on these models can be inferred from their hypothetical effect upon astrophysical phenomena such as neutron stars and supernovae, and these turns out to disfavour $d \leq 3$ (cf. [Kan04] and references therein).

When particles such as gravitons move in all $4+d$ dimensions, the net effect on their kinematics from a 4 dimensional viewpoint will be that the momentum in the d extra dimensions will appear as mass:

$$\begin{aligned} E^2 &= |\vec{p}|^2 + m^2 \\ &= p_x^2 + p_y^2 + p_z^2 + p_5^2 + p_6^2 + \ldots + p_{4+d}^2 + m^2 \\ &= p_x^2 + p_y^2 + p_z^2 + M^2, \quad M \equiv \sqrt{p_5^2 + p_6^2 + \ldots + p_{4+d}^2 + m^2} \end{aligned} \tag{5.7}$$

For gravitons, $m = 0$ of course, but they would nonetheless appear massive in 4 dimensions. As the extra dimensions are compact, their possible momentum states will be discretised. If the extra dimensions are compactified toroidally (i.e. as a direct product of one-spheres $S^1 \times \ldots \times S^1$) with universal radii R, their allowed masses will be:

$$\frac{\sqrt{n_5^2 + n_6^2 \ldots + n_d^2}}{R}, \quad n_i = 0, 1, 2, \ldots \tag{5.8}$$

[1] Pedagogical and precise calculations in the case of toroidal compactification are given in [Kje04].

With R^{-1} ranging from $\sim 10^{-4}$ eV for $d = 2$ to ~ 100 MeV for $d = 7$. Effectively, due to the large value of R, there will almost be a continuum of mass states available for a graviton from a 4 dimensional viewpoint. The resulting tower of available states is denoted *Kaluza Klein* states (KK) and their high density will enhance the graviton production cross section at the TeV scale and thus make up for the small value of Λ_{Planck}. From a $4+d$ dimensional viewpoint on the other hand, there will simply be a high cross section for production of massless gravitons at the TeV scales, as this is really the true scale of quantum gravity, $\Lambda^{\text{true}}_{\text{Planck}} \sim$ TeV. If there is any truth to the scenario of large extra dimensions, it thus opens up the very exciting scenario of being able to probe quantum gravity directly at the LHC! Some people even speculate that this might result in the creation of mini black holes inside the detector [Kan04].

Many variations exists on models of large extra dimensions. In some models massive gravitons can leave the 4 dimensional sub-manifold ("brane"), with the apparent violation of momentum conservation being balanced by the brane itself (just as a person jumping of the surface of the earth does not violate momentum conservation). Or perhaps the mechanism that keeps SM particles on the 4 dimensional sub-manifold is similar to quark confinement, only letting particles without SM charges propagate into the $4 + d$ dimensions. In this case a SM particle trying to leave the brane could, if it had enough energy, result in pair creation and resulting neutral composite states moving out in the $4 + d$ dimensions (in the same way a highly energetic quark trying to leave hadronic confinement results in colour-neutral pions). As the extra dimensions are compactified, such neutral composites would thus cross the 4 dimensional sub-manifold of the SM periodically, possibly resulting in interactions each time. Imagine reconstructing a long line of vertices all $\sim R$ apart! Models have even been constructed where the extra dimensions are not even compactified [RS99a] or large [RS99b], while still resolving the hierarchy problem. Such scenarios obviously requires other "tricks".

Having accepted the idea of extra dimensions, one could of course also imagine having extra dimensions where some or all Standard Model fields could propagate, *universal extra dimensions* (UED) [ACD01]. These models might not require the introduction of a mechanism to keep the SM fields in the 4 non-compactified dimensions, but they will instead have to be a lot smaller, $R <\sim m^{-1}_{\text{weak}} \approx 10^{-18} m$, in order to escape electroweak constraints.

Momentum in the extra dimensions will again result in effective mass contributions in 4 dimensions according to equation (5.7). If for instance $d = 1$ and the SM field has mass m the resulting Kaluza-Klein states should to a very good approximation be towers of almost evenly spaced masses: $M_n = \sqrt{n^2/R^2 + m^2} \approx n/R + \frac{1}{2}(mR)^2$, which are almost degenerate for different SM particles (although the degeneracy is lifted somewhat due to radiative corrections to the masses). Additionally, momentum conservation in the extra dimensions should a priori amount to conservation of the KK excitation number, n.[1] However, the technical details of allowing chiral fermions to extend into the extra dimensions means that the compactification scheme can not use a simple manifold such as e.g. S^1, but need to be a more complex orbifold such as $S^1/\mathbb{Z}_2$. The effect of this is that n will not actually be conserved, but that the KK parity, $(-1)^n$ will. This means, unless KK parity is somehow broken, that the lightest $n = 1$ state will be stable. Not surprisingly this usually turns out to be a KK excitation of the photon, γ^*. As it is neutral, this is actually a dark matter candidate.

The conservation of KK parity also implies that the lightest KK states can only affect

[1] A similar thing does not happen in the case of Large extra dimensions due to the localisation of the 4 dimensional branes within the new dimensions (thus breaking translational invariance). As already mentioned, the brane itself can absorb momentum in models of large extra dimensions, resulting in gravitons "vanishing" from the 4 dimensional world.

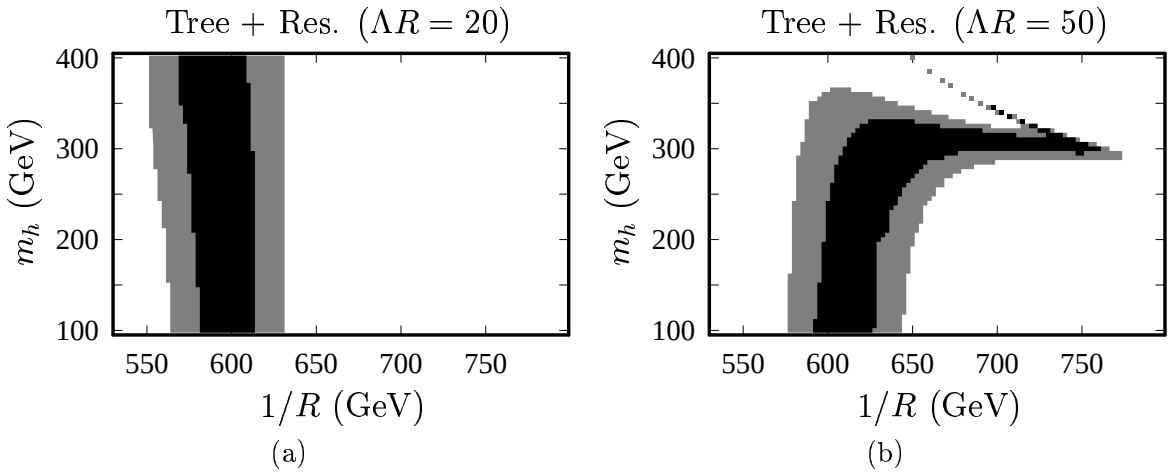

Figure 5.1: *Bounds corresponding to regions where a model with one extra universal dimension is consistent with the observed dark matter density at the 1σ and 2σ level, respectively. The results in figures (a) and (b) shows the result for a cutoff that is 20 and 50 times greater than $1/R$ respectively. Source: [KMSS06].*

SM processes if they appear in pairs, and the indirect lower bound on $1/R$ from electroweak observables is thus not very strict. For $d = 1$, $1/R > 300$ GeV and for $d = 2$, $1/R > 400-800$ GeV [ACD01]. The reason for the big uncertainty in the $d = 2$ bound, and the reason for the absence of bounds for $d \geq 3$ is that gauge theories are actually not renormalisable in more than 4 spacetime dimensions. It turns out that UED theories anyway provides relatively accurate results for $d = 1$ and somewhat accurate results for $d = 2$, with the results depending on a effective cutoff, Λ. In figure 5.1 predicted annihilation cross section of KK states are translated into values of Ω_{CDM} for different values of m_h, R and Λ. The grey regions indicate compatibility with the observed value (from equation (4.5)). While the difference between figures 5.1.a and 5.1.b indicate a substantial theoretical uncertainty, it is clear that if a $d = 1$ model of universal extra dimensions is to explain the cosmologically observed dark matter, it will be associated with a mass spectrum where at least the $n = 1$ KK excitations should be within reach of the LHC.

The small extra dimension in UED does not in itself solve the hierarchy problem as such, although the effective cutoff, Λ, at tens of TeV's, reduces it somewhat (effectively delegating the problem to a more fundamental model). More direct solutions might of course be incorporated. It is for instance straight-forward to simply add also large extra dimensions to an UED theory and let them solve the problem in the usual manner. It has also been suggested that if the third quark doublet alone was allowed to extend into another dimensions, that a Higgs field would naturally appear as a condensate of these (using the non-renormalizability of the extra dimensional theory as an argument for the strong interactions needed for such a condensate) [A+00]. This would naturally give a large Higgs coupling to the t (and possibly b) quark, and the other couplings of this extra dimensional Higgs would become effectively small in the same way that extra dimensions can account for small graviton couplings.

5.4 Supersymmetry

A final, and arguably the most elegant, solution of the hierarchy problem is provided by supersymmetric (SUSY) theories (for further reading the reader is referred to the excellent review [Mar97]). As described in section 2, the Poincaré group does not quite exhaust the allowed spacetime symmetries in theories describing fermions. The final allowed possibility in a non-trivial theory would be the so-called *supersymmetries*, i.e. symmetries between the fermionic and bosonic degrees of freedom in a theory [HLS75]. If Q is a generator for such a symmetry, then

$$Q\,|\text{boson}\rangle = |\text{fermion}\rangle \quad \text{and} \quad Q\,|\text{fermion}\rangle = |\text{boson}\rangle \tag{5.9}$$

Which makes it immediately clear that the operator itself must transform as spinor of spin $1/2$, thus emphasising its spacetime origin as well as implying that $Q^\dagger \neq Q$. It further follows from equation (5.9) that $Q^\dagger$ will also generate transitions between fermions and bosons.[1] A minimal supersymmetry group is thus defined by the set of generators $\{Q, Q^\dagger\}$, and a theory with such a set of symmetry generators is denoted a $N = 1$ supersymmetric theory. It is also possible to construct $N \geq 2$ supersymmetries, with a larger set of generators, $\{Q_1, Q_1^\dagger, \ldots, Q_N, Q_N^\dagger\}$, but the discussion here will for simplicity avoid those.

5.4.1 Supermultiplet Assignments

The fields of a SUSY theory must be ordered into irreducible representations (analogously to how the SM fields are ordered into irreducible representations of $SU(2)$ in table 3.3). These are denoted *supermultiplets* and are per definition closed under the supersymmetry generators. This means that each supermultiplet must thus contain both fermionic and bosonic fields per equation (5.9), and it can in fact be shown that the number of bosonic and fermionic degrees of freedom in a supermultiplet must be equal. A specific consequence of this is that no field can be placed in a singlet supermultiplet, meaning that each field will be placed in a multiplet with other fields, its so-called *superpartners*. Careful analysis shows that there essentially only exists two types of supermultiplets can exist in a renormalisable theory. The first type are the *chiral supermultiplets*, where the two spin states of one chiral part of a fermion (left or right-handed) are paired with two scalar degrees of freedom, for instance as a complex scalar field. The second allowed type is the *gauge supermultiplets*, where such two fermionic degrees of freedom are instead paired with the two degrees of freedom of a massless gauge field (the supermultiplets are thus defined in terms of the fields before electroweak symmetry breaking).

The operators Q and $Q^\dagger$ are neutral under the gauge interactions, which implies that the particles of a given supermultiplet will share gauge quantum numbers with their superpartners, i.e. they will share the same electromagnetic and colour-charge, and they will all be in similar (singlet or doublet) states under $SU(2)$. Furthermore, Q and $Q^\dagger$ will commute with the generators of the Poincaré group, meaning that they will all share the same mass. This last requirement of course leads to predictions in clear contradiction with observations (such as a boson with the same quantum numbers and mass as the electron). If SUSY is at all a symmetry of nature, the only unavoidable conclusion is thus that it is spontaneously broken in a way that provides large effective masses to the unobserved superpartners at scales above the present direct experimental reach.

[1] For instance, using $Q^\dagger$ on the first equation yields $Q^\dagger\,|\text{fermion}\rangle = Q^\dagger Q\,|\text{boson}\rangle$. It is clear that the right-hand side transform as a boson.

It is now easy to see how the supersymmetry solves the hierarchy problem in a natural way. An unbroken supersymmetry would equate the mass of a scalar such as the Higgs particle to the mass of a fermion. As fermion masses only acquire moderate mass corrections comparable to the mass itself (cf. equation (4.8)), the scalar particles inherits this property. The actual mechanism that ensures this is a remarkable conspiracy between the spin-statistics theorem and the equal number of fermionic and bosonic degrees of freedom of a given mass scale ensured by SUSY. A scalar couples directly to both fermions and bosons and, as remarked in section (2.1), the spin-statistics theorem implies a sign difference between fermionic and bosonic loops, thus resulting in a cancellation of the radiative corrections of scalar masses. Of course, as SUSY is broken, this cancellation will be imperfect, with radiative corrections to the scalar masses proportional to the square of the mass splitting reappearing. If for instance the mass splitting between the top and its superpartner is taken to be δm_t, then the resulting contribution to the radiative Higgs boson mass correction will be $\delta m_h^2 \sim \delta m_t^2 \log \frac{\Lambda}{m_t}$, with similar results for the W and Z and their superpartners as well as for the Higgs itself and its superpartner.[1] This means that the superpartners of these particles should certainly have a mass of at $1 - 2$ TeV, if they should solve not reintroduce the hierarchy. Other fermions have smaller Higgs couplings, thus suppressing the contribution with $\sim m_f/m_t$, and they could thus be slightly heavier. Even the gluon and photon that does not couple to the Higgs boson directly would be bounded by this limit, as they couple to the Higgs in diagrams in higher order diagrams, and directly to the superpartner of the top which is also a scalar with potentially troublesome radiative mass correction. In conclusion, to avoid the hierarchy problem all new sparticles should have masses at the TeV scale, with slightly stricter bounds on the superpartners of the top quark, W, Z and Higgs boson which should all be within reach of the LHC.

5.4.2 Sparticle Spectrum

After all of these realisations, the next step in the building of a supersymmetric theory is then find out which particles should go into which supermultiplets. From an economical point of view one might wish to pair up as many of the boson and fermion fields of the SM model with each other as possible, and thus minimise the number of required new fields. However, due to the requirements mentioned in the previous section such pairings are not possible at all within the SM. For instance, the right-handed electron field, e_R, must be combined with a complex scalar field carrying unit electromagnetic charge and electron lepton number. Such a scalar field simply does not exists among the other SM fields, and thus must be a new field to be added to the theory. A similar line of reasoning can be used for each of the other SM fields, and in the end it is clear that several new scalar and fermion fields have to be added to the theory. Out of convenience the notation and names for these new "SUSY particles" have all been chosen to serve as handy mnemonics for their properties by referring to their SM superpartners. Thus, the notation for the new fields are all given by adding a tilde ("$\sim$") to the symbols of the SM superpartners, the SUSY particles or *sparticles*, and their names are given by that of their SM superpartners, appending a "s-" for new scalars and prepending "-ino" for

[1] It is implicit in the discussion here that SUSY is *softly* broken, i.e. broken in a way that removes the mass degeneracy, but does not introduce different gauge couplings for the superpartners. If this were the case the radiative mass corrections would only be reduced to the difference between couplings multiplied with $\Lambda^2_{\text{Planck}}$, which would generally not be very small at all. Fortunately, this soft breaking arises naturally in many scenarios of supersymmetry breaking.

new fermions. This adds 12 new scalar fields (squarks) as superpartners of the quark fields, $\tilde{u}_L$, $\tilde{d}_L$, $\tilde{c}_L$, $\tilde{s}_L$, $\tilde{t}_L$, $\tilde{b}_L$, $\tilde{u}_R$, $\tilde{d}_R$, $\tilde{c}_R$, $\tilde{s}_R$, $\tilde{t}_R$ and $\tilde{b}_R$ and superpartners (sleptons) of the leptonic fields $\tilde{e}_L$, $\tilde{\mu}_L$, $\tilde{\tau}_L$, $\tilde{e}_R$, $\tilde{\mu}_R$, $\tilde{\tau}_R$, $\tilde{\nu}_{e,L}$, $\tilde{\nu}_{\mu,L}$ and $\tilde{\nu}_{\tau,L}$ (the discussion here is, for simplicity, ignoring the non-zero neutrino masses). Note that the designations "L" and "R" does not imply any form of chirality of the SUSY scalars themselves. They are, after all, scalars and such a concept would thus be meaningless. Rather they designate the chirality of the SM superpartner, and of course indicate whether the scalar itself is in a singlet or a doublet state with respect to the $SU(2)$ transformation. This thesis will use the convention of marking all SUSY particles with a double line in a Feynman diagram, and only using arrows on propagators that have spin half. Furthermore, neutralinos, charginos and gluinos will be marked with full lines and scalars with dashed lines.

The superpartners of the gauge and Higgs fields will all be spin 1/2 fermions and such field are generally called *gauginos* and *Higgsinos* respectively. Denoting the SM $SU(2)$ and $U(1)$ gauge bosons with $W^i = (\mathrm{W}^0, \mathrm{W}^\pm)$ and B respectively (i.e. the Z and γ fields are mixtures of W^0 and B), the corresponding superpartners will be the *winos*, $\widetilde{W}^0$ and $\widetilde{W}^\pm$, and the *bino*, $\widetilde{B}$. Likewise, the gluon gets a superpartner, the gluino, $\tilde{g}$. Regarding the Higgs sector, it was described in section 3.3 how up-type and down-type fermions masses can only be simultaneously accommodated in a model with a single Higgs doublet, ϕ, if Yukawa couplings involving both ϕ and ϕ^c exists. However, the couplings involving ϕ^c are forbidden by the structure of supersymmetry, and any supersymmetric theory must thus have a second Higgs doublet.[1] One of these doublets will provide masses to all down-type SM fermions and one will provide masses to up-type SM fermions. As described in section 3.3, this thus leads to 5 massive Higgs states, h^0, H^0, A^0, H^+ and H^-, and consequently also to 5 Higgsino fermions: $\widetilde{h}^0$, $\widetilde{H}^0$, $\widetilde{A}^0$, $\widetilde{H}^+$ and $\widetilde{H}^-$.

The prediction of new fields in a generic $N = 1$ SUSY model is thus quite specific. The only allowed variation on the scheme mentioned above would be the introduction of additional supermultiplets, whose particles were all too massive or weakly coupled to have been observed so far. The so-called *Minimal Supersymmetric Standard Model* (MSSM), makes the assumption that only the minimally required supermultiplets, i.e. those including SM particles, exists (it usually also assumes R-parity conservation which will be introduced later). Remarkably, the unbroken MSSM thus introduces no new parameters, and its particle content and their interactions are entirely given by the supersymmetry! This amazing predictability is, however, lost since the supersymmetry must be broken. The MSSM does not a priori try to implement a specific mechanism for the spontaneous supersymmetry breaking, but rather explicitly includes mass terms for all sparticles. However, as all sparticles (with the exception of the gluino) share all conserved quantum numbers with at least one other SUSY particles, the flavour states will in general be mixed in the mass eigenstates, thus requiring mixing angles and CP violating phases in addition to just mass terms. All in all the MSSM, which is remarkably constrained in its unbroken form, requires the introduction of a total of 105 parameters when broken generically. Of course, if one knew the mechanism behind the spontaneous breaking of supersymmetry, one would also know how these parameters are related to each other and the number of free parameters should hopefully be reduced drastically. For practical reason, most studies of supersymmetry in the broken regime are thus performed in the context of one or more specific model of supersymmetry breaking. In section 5.4.4 one such model will be

[1] In a supersymmetric theory the second Higgs doublet is also needed in order to cancel a disastrous contribution of the first to the gauge anomaly.

introduced, which has been widely popular and is often used as a benchmark scenario, even though it is not particularly well motivated.

However, even before choosing a specific model for how the supersymmetry is broken, the free parameters of the MSSM are severely constrained by experimental data. This is because the majority of the 105 new parameters of the MSSM are associated with CP and flavour violations, which in many cases have very strict bounds. To put it more precisely, the 6 sleptons of unit charge can a priori mix freely and the 6 mass eigenstates are found by diagonalising a 6×6 mass matrix. However, the indirect limits on flavour and CP violating effects beyond the Standard Model (such as in $\mathcal{B}(\mu \to e\gamma)$ and neutral kaon mixing) implies that mixing angles between generations must be very small and that the first and second generation scalars with a given chirality superpartner must be nearly degenerate in mass, so that e.g. $m_{\tilde{e}_L} \approx m_{\tilde{\mu}_L}$ and $m_{\tilde{e}_R} \approx m_{\tilde{\mu}_R}$ [DKS90]. Similar degeneracy can on the other hand not exists for the third generation scalars, as large Higgs couplings results in large radiative mass corrections. In reality the 105 free parameters of the MSSM are thus reduced to roughly 20 unconstrained parameters.

The result is that the following slepton spectrum arises at the electroweak scale: $\tilde{e}_L$ and $\tilde{\mu}_L$ will be nearly degenerate in mass but will not mix. The same is true for $\tilde{e}_R$ and $\tilde{\mu}_R$. The two tau states, $\tilde{\tau}_R$ and $\tilde{\tau}_R$, on the other hand will mix and end up in two non-degenerate mass eigenstates: $\tilde{\tau}_1$ and $\tilde{\tau}_2$ (with the lower index always indicating the lightest sparticle). The same will be true in the squark sector where $(\tilde{u}_L, \tilde{c}_L)$, $(\tilde{d}_L, \tilde{s}_L)$, $(\tilde{d}_R, \tilde{s}_R)$ and $(\tilde{u}_R, \tilde{c}_R)$ will all be non-mixing pairs of mass degenerate states, and $(\tilde{t}_1, \tilde{t}_2)$ and $(\tilde{b}_1, \tilde{t}_b)$ will be non-degenerate mass eigenstates of the scalar top and bottom states, respectively.

The gauginos and Higgsinos will likewise mix. The two CP even Higgsinos, the bino and the neutral wino will combine to form four *neutralinos*, $\tilde{\chi}_1^0$, $\tilde{\chi}_2^0$, $\tilde{\chi}_3^0$ and $\tilde{\chi}_4^0$. Likewise, charged winos and Higgsinos will mix and form four *charginos*, $\tilde{\chi}_1^\pm$ and $\tilde{\chi}_2^\pm$.

In general radiative corrections tends to reduces the overall qualitative low-energy differences between different SUSY breaking scenarios somewhat, as e.g. some masses are always driven either upwards or downwards at low energies by radiative corrections. It is for instance natural to expect that the squarks and gluino will tend to be heavier than the sleptons and gauginos. It is also typical for the radiative corrections to result in a tachyonic Higgs mass at the electroweak scale (but not for any of the other scalars), which means that supersymmetric models provide a natural explanation of electroweak symmetry breaking. In addition it is rather interesting that the gauge unification in figure 4.9 is not exact in the sense that the three lines do not actually meet in one point. However, the presence of SUSY particles changes the radiative corrections to the couplings and truly makes them meet in one point [GKL91].

5.4.3 Interactions and R-Parity

As the particles in a given supermultiplet carries the same quantum numbers and are in similar gauge representations, all SM interaction vertices can be translated into valid interaction vertices involving SUSY particles by substituting one or more SM particles with their superpartners. The supersymmetry furthermore guarantees that the coupling constant associated with the vertex should be unchanged. However, due to the spin difference between SM particles and their superpartners, Lorentz invariance dictates that only the new vertices that results when an even number of SM particles are exchanged with superpartners are allowed. This situation alone would limits the potential indirect effect of SUSY on precision electroweak

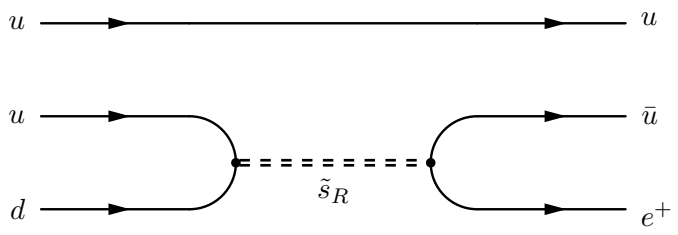

Figure 5.2: *The presence of baryon and lepton number violating vertices in supersymmetric models potentially leads to proton decays.*

observables.

However, the argument can also be reversed. If it is possible to arrive at vertices involving SUSY particles by a "supersymmetrisation" of SM vertices, it should of course also be possible to arrive at the SM vertices by a "desupersymmetrisation" of vertices involving SUSY particles. However, once again only the vertices that are reached by exchanging two SUSY particles with SM particles should be considered. This leaves the possibility of the existence of vertices with an odd number of SUSY particles that have no analogy in the SM. This is really just a symmetry argument: An unbroken supersymmetry should not give special preference to SM particles over their superpartners, and if there are vertices with one SM particle and two SUSY particles, there should also be a possibility of vertices with one SUSY particle and two SM particles. The problem with the new vertices are that they invariably violate either lepton or baryon number conservation, usually with disastrous predictions such as rapid proton decay as shown in figure 5.2. In order to remain in agreement with experiments, the coupling constants of the new vertices must be extremely small. Such smallness would be most natural if it was a consequence of a symmetry principle, and thus the conservation of a multiplicative quantum number, R-parity, is usually postulated:

$$R = (-1)^{3B+L+2s} \tag{5.10}$$

Where $B = \frac{1}{3}[N(q) - N(\bar{q})]$ is the baryon number, $L = L_e + L_\mu + L_\tau$ is the generation agnostic lepton number (cf. equation (3.7)) and s is the spin quantum number. In Lorentz invariant theories, the factor of $(-1)^{2s}$ is actually superfluous, as its product in an interaction vertex will always be one, $(-1)^{2s_1}(-1)^{2s_2}(-1)^{2s_3} = (-1)^{2(s_1+s_2+s_3)} = (-1)^{2\cdot 0} = 1$. However, its presence has the convenient effect that all SM particles will have $R = +1$ and all new SUSY particles, i.e. all those marked with a tilde ("$\sim$"), will have $R = -1$. This makes it immediately obvious that every vertex must have an even number of sparticles, thus avoiding all B and L violating vertices. A valid Feynman diagram will thus always have an unbroken sparticle line, which is the reason for the convention in the present thesis of always marking sparticles with double lines.

A direct consequence of R-parity is that sparticles can only be produced directly in pairs in collider experiments, thus lowering the mass reach in direct searches. Furthermore, the lightest supersymmetric particle (LSP) will become stable, and thus a dark matter candidate, if it is also uncharged. In fact, taken at face value R-parity thus excludes models where the LSP is charged.

It should be noted that the absence of R-parity violating vertices like $ud\tilde{s}_R$ is not more unnatural than the absence of e.g. the $e\mu\gamma$ vertex in the SM. However, the conservation of R-parity is not particularly natural either, and it is of course quite possible that R parity could

be broken, albeit in a very small way so as to not come into conflict with indirect experimental limits such as those from the proton decay. The most important consequence of this would be that the LSP would no longer be stable. As mentioned, R-parity is usually taken as an assumption in the MSSM.

5.4.4 The Constrained MSSM – A Scenario of SUSY Breaking

It is beyond the scope of the present thesis to describe the details of the various models of SUSY breaking that have been proposed, and the reader is once more referred to the literature such as [Mar97]. To avoid difficulties with low-scale phenomenology more or less all of the proposed models introduce a new sector of supermultiplets wherein the spontaneous SUSY breaking actually occurs. The new sector is typically only coupled radiatively to MSSM particles through flavour blind interactions mediated by such diverse things as gravitons, a gaugino compensate, regular SM gauge bosons or even extra dimensional states [RS99c]. As always when a symmetry is spontaneously broken a massless Goldstone boson appears. However, due to the special nature of supersymmetry, the massless particle that appears when it is spontaneously broken is actually a massless spin 1/2 fermion of odd R-parity, the *goldstino*. Quite analogously to the Higgs mechanism in electroweak symmetry breaking, the goldstino will be eaten by a hungry gravitino, which is a massless spin 3/2 particle thus supplying it with mass and spin $\pm 1/2$ longitudinal polarisations (which can thus couple a lot stronger than the usual gravitational strength). In the theories where the SUSY breaking is mediated by gauge bosons the massive gravitino could be the LSP and thus a dark matter candidate, and its goldstino part could have an interaction strength sufficient for it to be important inside collider experiments. Of course, theories involving gravitons and gravitinos suffer to some degree from problems of renormalisability.

The analysis presented in the last part of this thesis will work inside the assumptions of a gravity mediated spontaneous supersymmetry breaking mechanism, which works under the assumption of unified particle masses at M_{GUT}. This so-called *Constrained Minimal Supersymmetric Standard Model* (CMSSM)[1] is perhaps not particularly well motivated theoretically (it has even been called a "fairly arbitrary parameter reduction scheme" [Str06]), but it has historically been *the* benchmark scenario for proposed SUSY analyses at collider experiments. As the actual mechanism behind SUSY breaking is unknown anyway, the CMSSM was chosen in the present thesis also out of convenience.

The CMSSM assumes that, at M_{GUT}, all sleptons and squarks have unified SUSY breaking masses of m_0 (which also contributes to Higgs masses), and all gauginos have unified SUSY breaking masses of $m_{1/2}$. The masses of the Higgsinos will be given by μ from the Higgs potential. Furthermore, all cubic SUSY breaking interactions are assumed to have a common interaction strength, A_0. Some of the free parameters of the theory can be related to the values of parameters such as m_{Z}, m_t and $\tan\beta$ at the weak scale and the freedom of the CMSSM is reduced to four free parameters and a sign [KKRW94]:

$$m_0,\ m_{1/2},\ A_0,\ \tan\beta,\ \mathrm{sign}\mu \tag{5.11}$$

For a given value of these parameters, careful computations of radiative corrections must be performed in order to get the resulting particle spectrum at the weak scale. A typical situation

[1] It is also known as the minimal supergravity model (MSUGRA), although that name is becoming less popular, as it has become apparent that its connection to supergravity models is perhaps not really that clear cut.

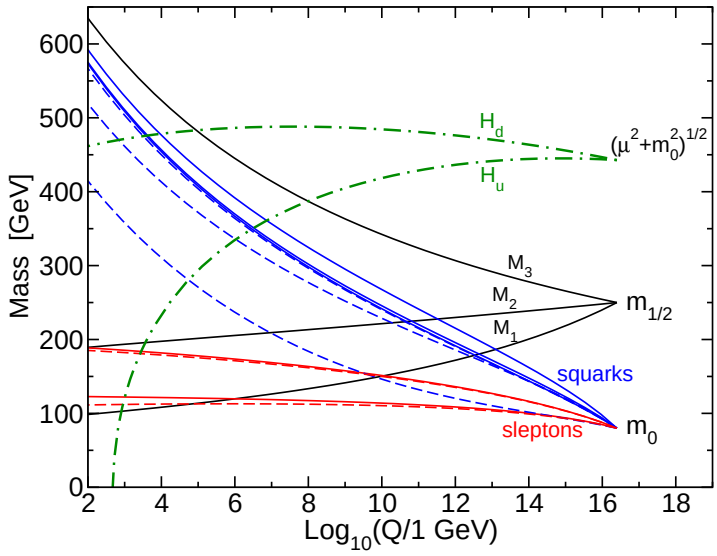

Figure 5.3: *Typical CMSSM spectrum of SUSY masses at different energy scales as discussed in the text. The parameters used for the plot were* $M_{GUT} = 2.5 \cdot 10^{16}$ *GeV,* $m_0 = 80$ *GeV,* $m_{1/2} = 250$ *GeV,* $A_0 = -500$ *GeV,* $\tan\beta = 10$*, and* $\mathrm{sign}\mu = +1$*. Source: [Mar97].*

is illustrated in figure 5.3, where the dashed squark and slepton lines corresponds to the third generation particles.

The radiative corrections of the slepton masses will in general be relatively small and will split the 6 slepton mass eigenstates into two groups, with the lightest consisting of the states that are primarily superpartners of right-handed fields ($\tilde{e}_R$, $\tilde{\mu}_R$ and $\tilde{\tau}_1$). The scalar muon and electron fields are almost mass degenerates, while the tau states are slightly lighter.

As they carry colour-charge, the radiative contribution to squark masses are much larger. The states are all driven to approximately the same masses, apart from the lightest top state, $\tilde{t}_1$, that becomes a lot lighter than the rest due to their very strong Higgs couplings. The same large coupling is also responsible for driving one of the Higgs masses to tachyonic mass, thus inducing electroweak symmetry breaking (the two Higgs doublets are labelled H_d and H_u in the figure).

Finally, the masses of the gauginos, the bino (M_1), the winos (M_2) and the gluino (M_3) are also shown (approximate in the first two cases as they are not mass eigenstates). The gluino will naturally obtain a mass similar to those of the squarks, while the other gauginos will be lighter.

It is possible to derive approximate formulae for how the sparticle masses at the weak scale are related to the parameters in equation (5.11). For instance the approximate masses of the $\tilde{\chi}_1^0$, $\tilde{\chi}_2^0$, $\tilde{\chi}_1^\pm$, $\tilde{g}$ and scalars of the first and second generations are given approximately by (cf. [Mar97] and [CST03] and references therein):

$$\begin{aligned}
&m_{\tilde{\chi}_2^0} \approx m_{\tilde{\chi}_1^\pm} \approx 2\, m_{\tilde{\chi}_1^0} \approx 0.9\, m_{1/2}, \quad m_{\tilde{g}} \approx (3.1 \pm 0.5)\, m_{1/2} \\
&m^2_{\tilde{e}_L} \approx m^2_{\tilde{\mu}_L} \approx m_0^2 + 0.53\, m^2_{\frac{1}{2}}, \quad m^2_{\tilde{e}_R} \approx m^2_{\tilde{\mu}_R} \approx m_0^2 + 0.15\, m^2_{\frac{1}{2}} \\
&m^2_{\tilde{u}_L} \approx m^2_{\tilde{c}_L} \approx m^2_{\tilde{d}_L} \approx m^2_{\tilde{s}_L} \approx m_0^2 + (6.0 \pm 1.0)\, m^2_{\frac{1}{2}} \\
&m^2_{\tilde{u}_R} \approx m^2_{\tilde{c}_R} \approx m^2_{\tilde{d}_R} \approx m^2_{\tilde{s}_R} \approx m_0^2 + (5.5 \pm 1.0)\, m^2_{\frac{1}{2}}
\end{aligned} \tag{5.12}$$

Due to the electroweak symmetry breaking, the masses for sleptons and squarks receives corrections typically less than $\mathcal{O}\left(m_{\mathrm{Z}}^2\right)$, thus introducing a "hyperfine splitting" between e.g. $m_{\tilde{u}_L}$ and $m_{\tilde{c}_L}$. However, in order to get precise values of all sparticle masses and mixing parameters (and thus also branching ratios), detailed computations must be performed. Several programs have been developed for this purpose, based on slightly different approaches, but generally agreeing at the $\mathcal{O}\left(10^{-2}\right)$ level or better [BFKP06]. The weapon of choice in the present thesis is the ISASUGRA program from ISAJET version 7.72 [PPBT03]. A top mass of $m_t = 173$ GeV is assumed in all instances.

It is interesting to note that indirect measurements such as the values of the muon anomalous magnetic moment in figure 4.1 and the (cold) dark matter density in equation (4.5) already put tight constraints on the values of the CMSSM parameters defined in equation (5.11). Starting from a sparticle spectrum, the program MicrOMEGAs version 1.3.6 [BBPS06] [BBPS02] is used to calculate predictions of the observables in question.

In CMSSM the lightest neutralino is the LSP and its primary component is the bino, which has small couplings. This means that the necessary annihilation cross sections at the early universe will generally be too small and the relic density far too large. However, in certain regions of CMSSM parameter space, the annihilation cross sections are enhanced due to resonance effects, and those will be the regions that are compatible with the observed dark matter density. The cosmological observations actually determine the parameter Ω_{CDM} at the 10% level, which translates directly into rather narrow allow regions of CMSSM parameter space. However, due to the *many* model dependent assumptions entering into such a translation, and since it would be surprising if the CMSSM was actually valid at such a precise level, the dark matter bounds shown in the present thesis will be less severe than the actual uncertainty on Ω_{CDM} would imply. The upper bound is thus raised by approximately a factor of two to $\Omega_{\mathrm{CDM}}h^2 < 0.25$ (h is the dimensionless Hubble parameter, $h^2 \approx 0.5$), and the lower bound is removed entirely. It is after all possible that the dark matter density could get contributions from other sources. Also, remember that even an incredibly small violation of R-parity conservation could turn the LSP unstable and thus remove the upper bound entirely.

The predicted dark matter density is shown in $(m_0, m_{1/2})$ planes for four different values of $(\tan\beta, A_0, \mathrm{sign}\mu)$ in figure 5.4. The grey areas in the upper left corners are excluded as they corresponds to mass spectra where the LSP is the charged $\tilde{\tau}_1$, and the grey area at very low values of $m_{1/2}$ are excluded as they corresponds to scenarios without electroweak symmetry breaking (i.e. where the H_u line in figure 5.3 is not driven towards negative values). Generally the allowed regions can be described as two long regions extending along the excluded edges and a central larger region. At the time of freeze out, most of the SUSY particles will either be the LSP or those very close to it in mass. This mean that the SUSY models producing a suitably low relic density will be those with enhanced cross section between either two LSPs or between the LSP and those particles very close to it in mass. The three allowed regions in figure 5.4 have enhanced annihilation cross sections for different reasons. In the horizontal area stretching towards high m_0 the mass of the Higgs boson is very close to twice the LSP mass, thus resulting in resonant enhancement of $\tilde{\chi}_1^0\tilde{\chi}_1^0 \to h \to b\bar{b}, \tau^+\tau^-$. In the vertical region, close to the excluded region, where the $\tilde{\tau}_1$ would be the lightest SUSY particle, the resonant decay $\tilde{\chi}_1^0\tilde{\tau}_1 \to \tau \to \tau\gamma$ is enhanced, and in the large central region, where both m_0 and $m_{1/2}$ are low, the sleptons will generally be very light, resulting in $\tilde{\chi}_1^0\tilde{\chi}_1^0$ annihilation into leptons through the exchange of a slepton.

Regarding the contribution of sparticles to the muon anomalous magnetic moment, it is estimated in [BBPS06] that it should be in the interval $[5.1 \cdot 10^{-10}, 64.1 \cdot 10^{-10}]$ (at the

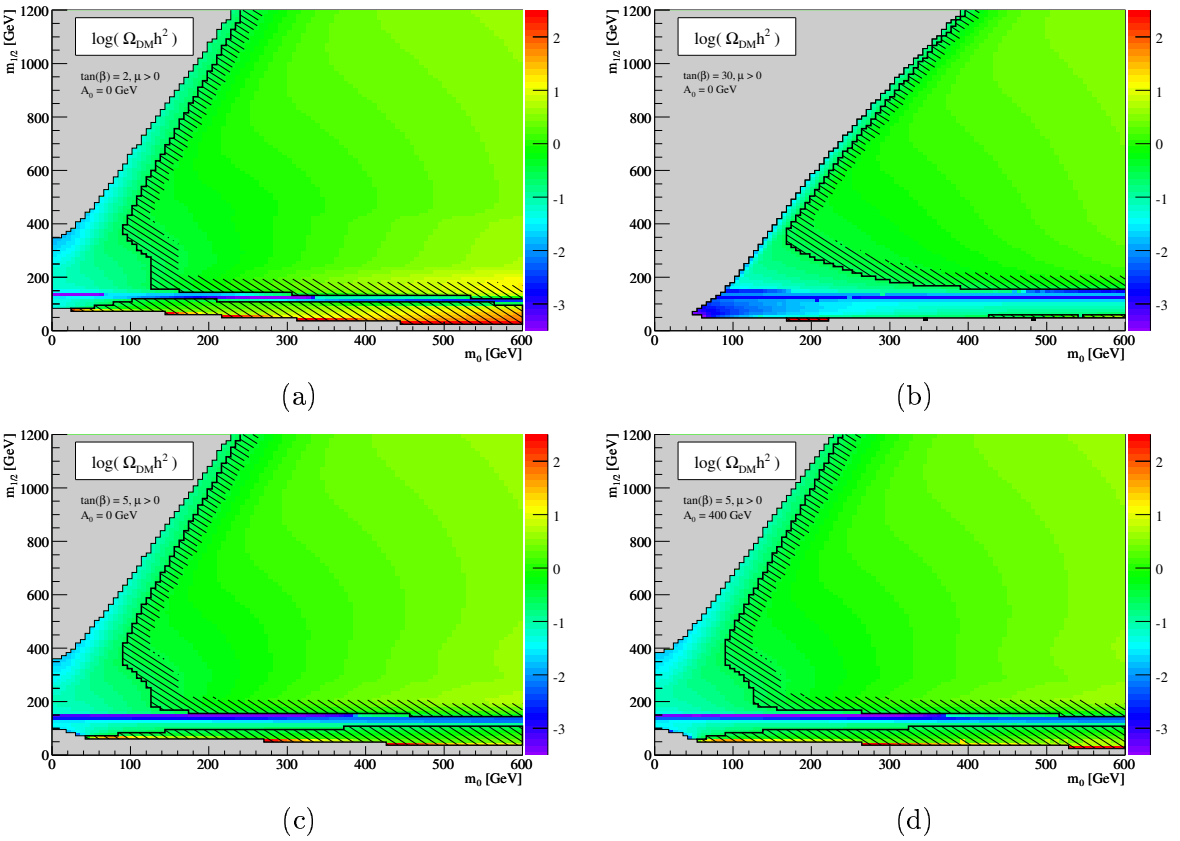

Figure 5.4: *Values of $\Omega_{DM}h^2$ resulting from different values of CMSSM parameters. The excluded region corresponds to $\Omega_{DM}h^2 > 0.25$.*

3σ confidence level). It turns out that high values of m_0 and $m_{1/2}$ will lead to too low contributions, which means that the muon anomalous magnetic moment data prefers relatively light supersymmetry unless $\tan\beta$ is high. Again, to avoid putting too much emphasis on the detailed predictions of the CMSSM the excluded region will be defined as the region where the contribution is less than $0.5 \cdot 5.1 \cdot 10^{-10}$. The predicted values and the excluded regions are again shown in the $(m_0, m_{1/2})$ planes for four different values of $(\tan\beta, A_0, \mathrm{sign}\mu)$ in figure 5.5. Note that the text "g-2" in the figure texts is really short-hand for "the contribution of sparticles to the muon anomalous magnetic moment".

The final two constraints that will be shown here are those due to the Higgs boson mass and the direct LEP searches for sleptons and charginos. These are shown simultaneously in figure 5.6 along with the bounds from the dark matter and contribution to the muon anomalous magnetic moment already discussed. In general the LEP bounds more or less corresponds to $m_{1/2} > 175$ GeV (C.L. 95%) (due to the chargino searches), with a slightly higher value when $m_0 < 60$ GeV (due to the slepton searches).[1] The predicted values of the mass of the lightest Higgs boson are always rather light, which is a general feature of supersymmetric models. It would appear that a value of $\tan\beta = 2$ is excluded. Furthermore, higher values of $m_{1/2}$ are

[1] The LEP bounds implemented in MicrOMEGAs appear to differ by a few GeV's with regard to the final LEP limits reported in [LEP04] and [LEP01]. Such small deviations are, however, irrelevant for the present purposes.

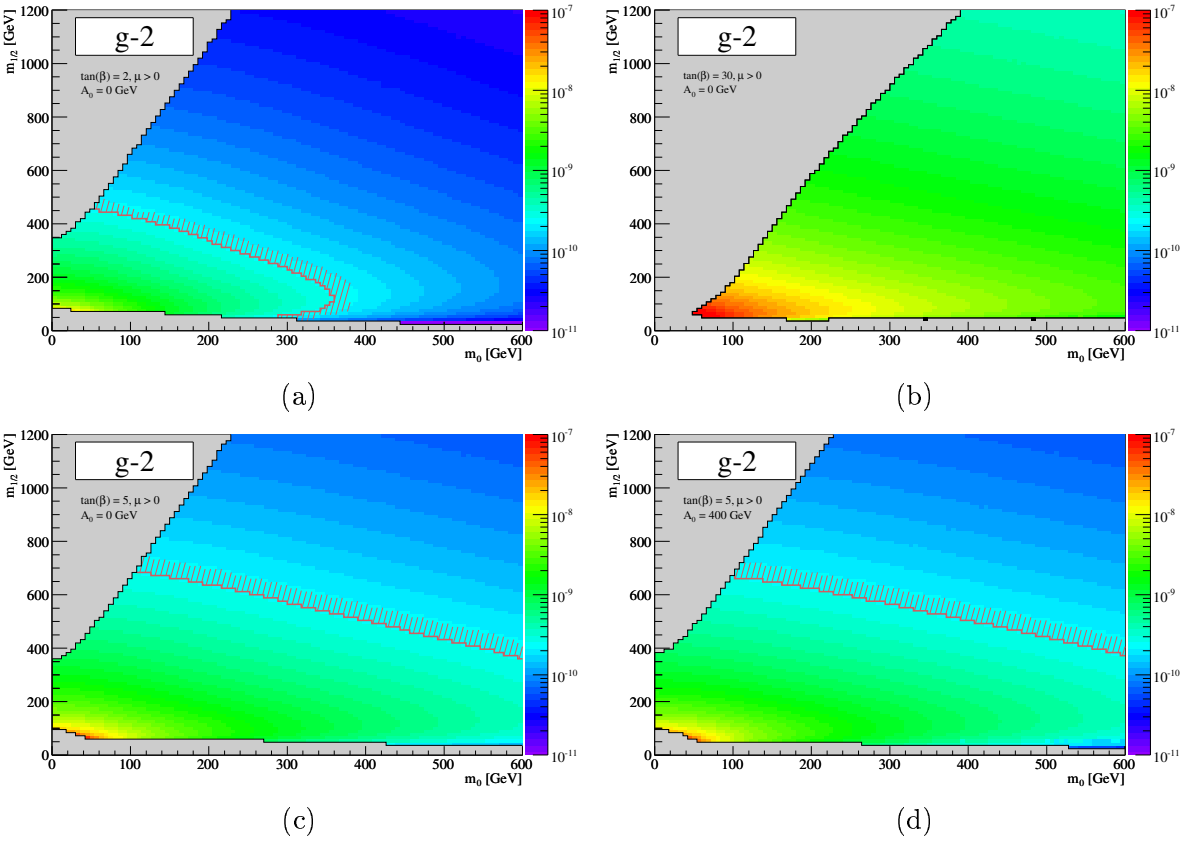

Figure 5.5: *Contribution to the anomalous magnetic moment of the muon due to the presence of SUSY particles given different values of CMSSM parameters. The excluded region corresponds to contributions smaller than* $0.5 \cdot 5.1 \cdot 10^{-10}$.

preferred but less strongly when $\tan\beta$ is high, $\mu > 0$ or $A_0 = 0$ GeV.

6 Phenomenological Expectations

Having thus reviewed several various proposed scenarios of physics that might exist at the TeV scale, this section will present a brief discussion of the experimental signatures these might lead to at the multipurpose experiments ATLAS and CMS at the LHC.

6.1 Physics at a 14 TeV Proton-Proton Collider

At the LHC two beams of 7 TeV protons will collide head-on. Unfortunately, due to the composite structure of the proton, only a small fraction of the 14 TeV that are available in the centre of mass frame will be available for the direct production of the hypothetical heavy particle states that are thought to exists at the TeV scale. As remarked in section 3.1 (and illustrated in the third picture in figure 3.3), the proton structure is rather complex, consisting of a multitude of partons, i.e. gluons, quarks and anti-quarks of various flavours, each usually carrying but a small fraction of the total proton momentum. For the same reason that quantities such as charge and mass are dependent on the scale at which they are probed, the details of how the total momentum is distributed across different partons will also depend

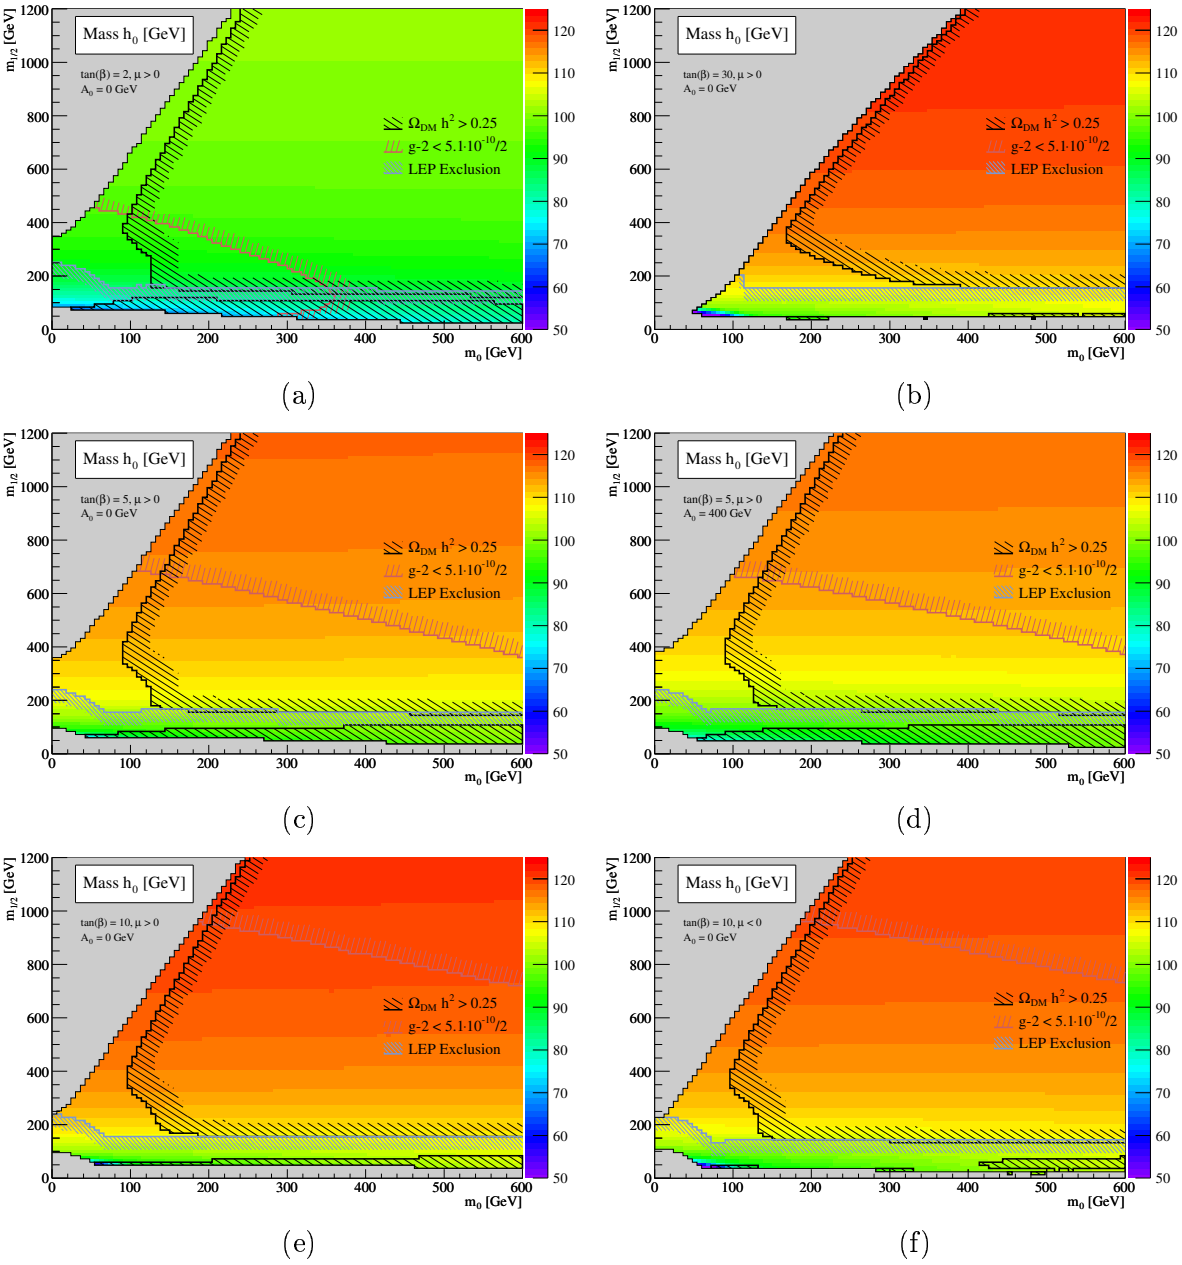

Figure 5.6: *Mass of the lightest Higgs boson as a function of different values of CMSSM parameters. Superimposed are other limits as discussed in the text.*

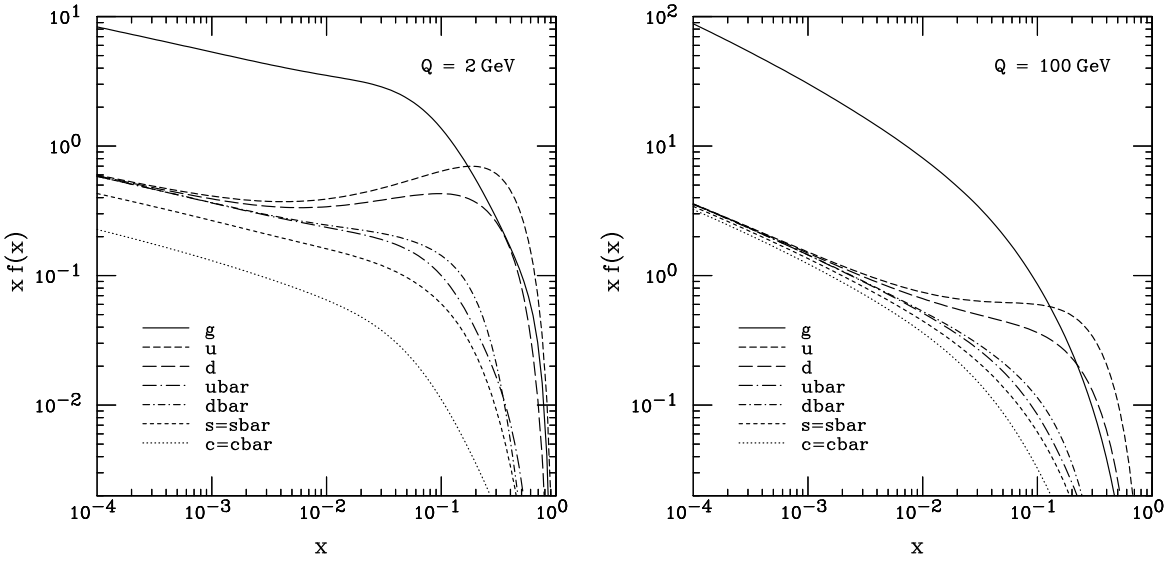

Figure 6.1: *Proton parton distribution functions at $Q = 2$ GeV and $Q = 100$ GeV respectively. Note that the vertical axis shows $xf(x, Q)$ and not $f(x, Q)$. Thus all $f(x, Q)$ are extremely peaked at low x. Source: [P+02].*

upon the energy scale, Q. These are usually parametrised into a parton distribution function (PDF) for each parton type. Thus, if p_{tot} is the total proton momentum, $f_g(x, Q)dx$ defines the fraction of partons that are gluons with momenta between xp_{tot} and $(x + dx)p_{\text{tot}}$, as observed at the energy scale Q. Conservation of energy of course ensures that:

$$\int_0^1 dx\, x\big[f_g(x,Q) + f_u(x,Q) + f_d g(x,Q) + f_{\bar{u}}(x,Q) + f_{\bar{d}}(x,Q) + \dots\big] \equiv 1 \qquad (6.1)$$

Uncertainties in parton distributions can often have a significant impact on physics analyses at a hadron collider. Figure 6.1 shows parton distributions at $Q = 2$ GeV and $Q = 100$ GeV. It is clear that all distributions are very peaked at low x, which underlines the fact that two partons sampled at random from two protons almost always will have a combined centre of mass energy much lower than that maximally available. The very highest x values are populated by both gluons and valence quarks, with gluons dominating at x less than ~ 0.1. This means that the potential gain from running an anti-proton proton collider (i.e. the high cross sections for resonant production through $q\bar{q}$ annihilation) are not necessarily sufficient to outweigh the extra technological challenges involved in using an anti-proton beam. The fact that the parton energies are distributed with small tails towards higher values furthermore means that the direct discovery potential of a hadron collider will always be limited both by beam energy and luminosity, whereas an e^+e^- collider like LEP or the proposed International Linear Collider, will be limited almost solely by beam energy once sufficient luminosities are recorded. To achieve as high as possible instantaneous luminosities, the LHC is going to run with very short bunch to bunch distances of 25 ns and with multiple (~ 25) collisions every bunch crossing.

It is also clear from the falling PDFs towards higher x that the cross section for the production of heavy particles (W, Z, top quarks, Higgs boson, SUSY particles, KK excitations,

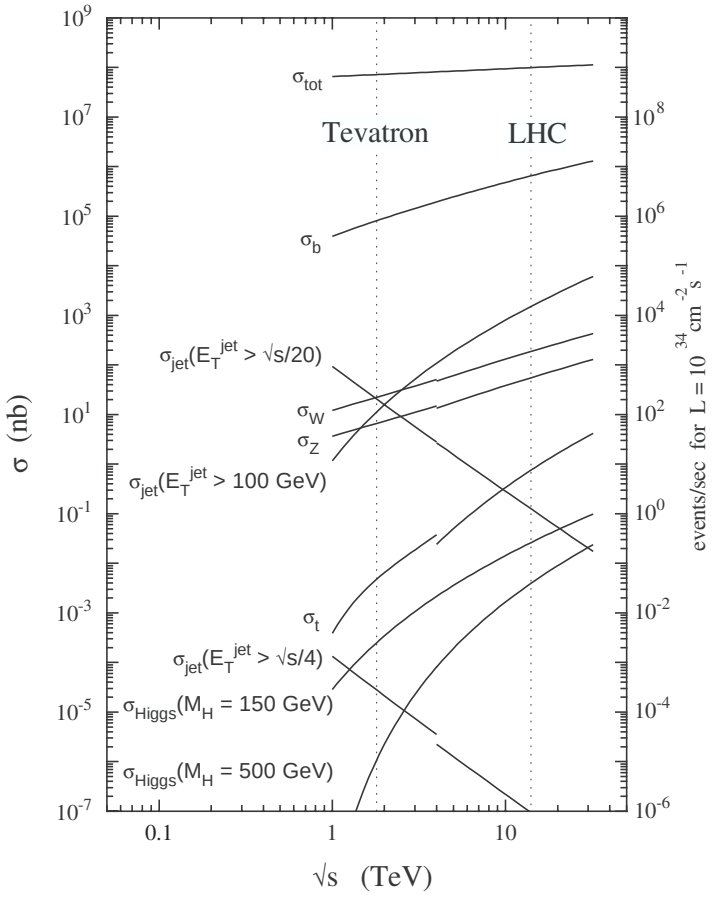

Figure 6.2: *Cross section for a variety of SM processes in proton-proton collisions. Source: [ATL03].*

etc.) will be but a fraction of the total proton-proton interaction cross section as can be seen in figure 6.2. Furthermore, even events where some interesting hard process takes place will be cluttered by the particles from the rest of the protons. Most particles from this "underlying event" will be emitted at low angles with respect to the beam (high rapidity), and it is inevitable that many will escape detection as even the best detector must have a blind spot, where the beam pipe enters it. This means that there will be no way to reliably determine the overall momentum going into the hard collision, apart from the knowledge that the momentum in the plane transverse to the beam should be almost zero, $p_t \approx 0$. With one component of the momentum thus unknown, resonances decaying into undetectable particles such as neutrinos, or perhaps a weakly interacting dark matter candidate, can not be reconstructed as mass peaks. Instead their masses will have to be inferred from threshold effects. To make matters worse, the initial state and the remnants are all coloured, and so will usually have non-trivial colour-connections to strongly produced coloured final states (in addition to the general problem of relating the hard process at parton level to the hadrons observed in the detector).

For all of these reasons, hadron colliders are a much more "dirty" environment than e^+e^- colliders. The reason that they are still very useful, is of course that it is feasible to operate proton beams at much higher energies than electron beams in a circular collider, as they are not limited by the constant energy loss through synchrotron radiation to the same degree. In fact the power lost due to synchrotron radiation in circular colliders is proportional to

γ^4, meaning that it will be extremely difficult for an e^+e^- collider to go beyond LEP beam energies of $\sim$ 100 GeV, while it is almost a non-issue for protons, as $(m_e/m_p)^4 \approx 10^{-13}$. Synchrotron radiation will be less of a problem in a linear collider, which on the other hand require much higher acceleration gradients. Acceleration units with higher gradients are thus the focus of the development of future generations of electron colliders.

6.2 Production of New Physics States

The distribution of partons inside the proton thus suppresses the production cross section of heavy particles. Furthermore, as discussed in section 5.1, the demand that new physics should not disturb precision electroweak data means that it either must be very weakly coupled, associated with a new global symmetry or have masses well above a TeV. This means that new physics will have production cross sections typically around picobarns. Furthermore, new particles that carries colour-charge should be produced much more copiously than those without.

The rest of this section will look at sparticles production within the CMSSM (as this will be relevant for the analysis in the last part of the thesis). Typical tree-level Feynman diagrams for the production of sparticles are given in figure 6.3. Figures 6.3.a to 6.3.f shows the dominant production of coloured gluinos and squarks. From the form of the parton distributions in figure 6.1, one would expect the diagrams in figures 6.3.e and 6.3.f to be most dominant, followed by figures 6.3.a, 6.3.b and 6.3.d. Figures 6.3.g to 6.3.j shows the more elusive direct productions of neutralinos, charginos and sleptons. These have smaller couplings and requires $\bar{q}$ in the initial states, but in the CMSSM the lightest of these particles are typically lighter than the lightest squarks, which thus partly compensates for these effects. It is reassuring that even if the coloured sparticles turns out to be too heavy to be produced at the LHC, it might still be possible to have Drell-Yan production of other sparticles through the diagrams 6.3.g, 6.3.i and 6.3.j.

The resulting total cross section to visible states (i.e. those that do not directly produce only sneutrinos or neutralino LSPs) is shown in figure 6.4. The biggest dependency seems to be on $m_{1/2}$, which can be explained by the presence of a gluino in the final states of figures 6.3.e and 6.3.f ($m_{\tilde{g}} \approx 3m_{1/2}$, cf. equation (5.12)). The production cross sections were evaluated using the HERWIG event generator version 6.510 [M$^+$92][C$^+$01][C$^+$02][MOR$^+$02].

The two directly produced sparticles per SUSY event are thus most likely to be squarks and gluinos, as shown explicitly in figure 6.5, which was made using a slightly modified version of HERWIG in which only the initial hard process of each event was generated. The exceptions are for very low $m_{1/2}$, where charginos are produced copiously and at very high $m_{1/2}$ and moderate m_0, where the coloured sparticles are very heavy and the sleptons are light (remember that the excluded regions in the upper left corners correspond to the case where the $\tilde{\tau}_1$ is the LSP). The very low regions of large chargino production are already excluded by LEP, however.

6.3 Signatures of New Physics

As should be clear from the discussion in section 5, there are many different scenarios for what new physics could appear at the TeV scale, and thus in principle many different potential signatures to be aware of in the experimental searches at the LHC. However, there are in many recurring phenomenological features between the different models, which is of course

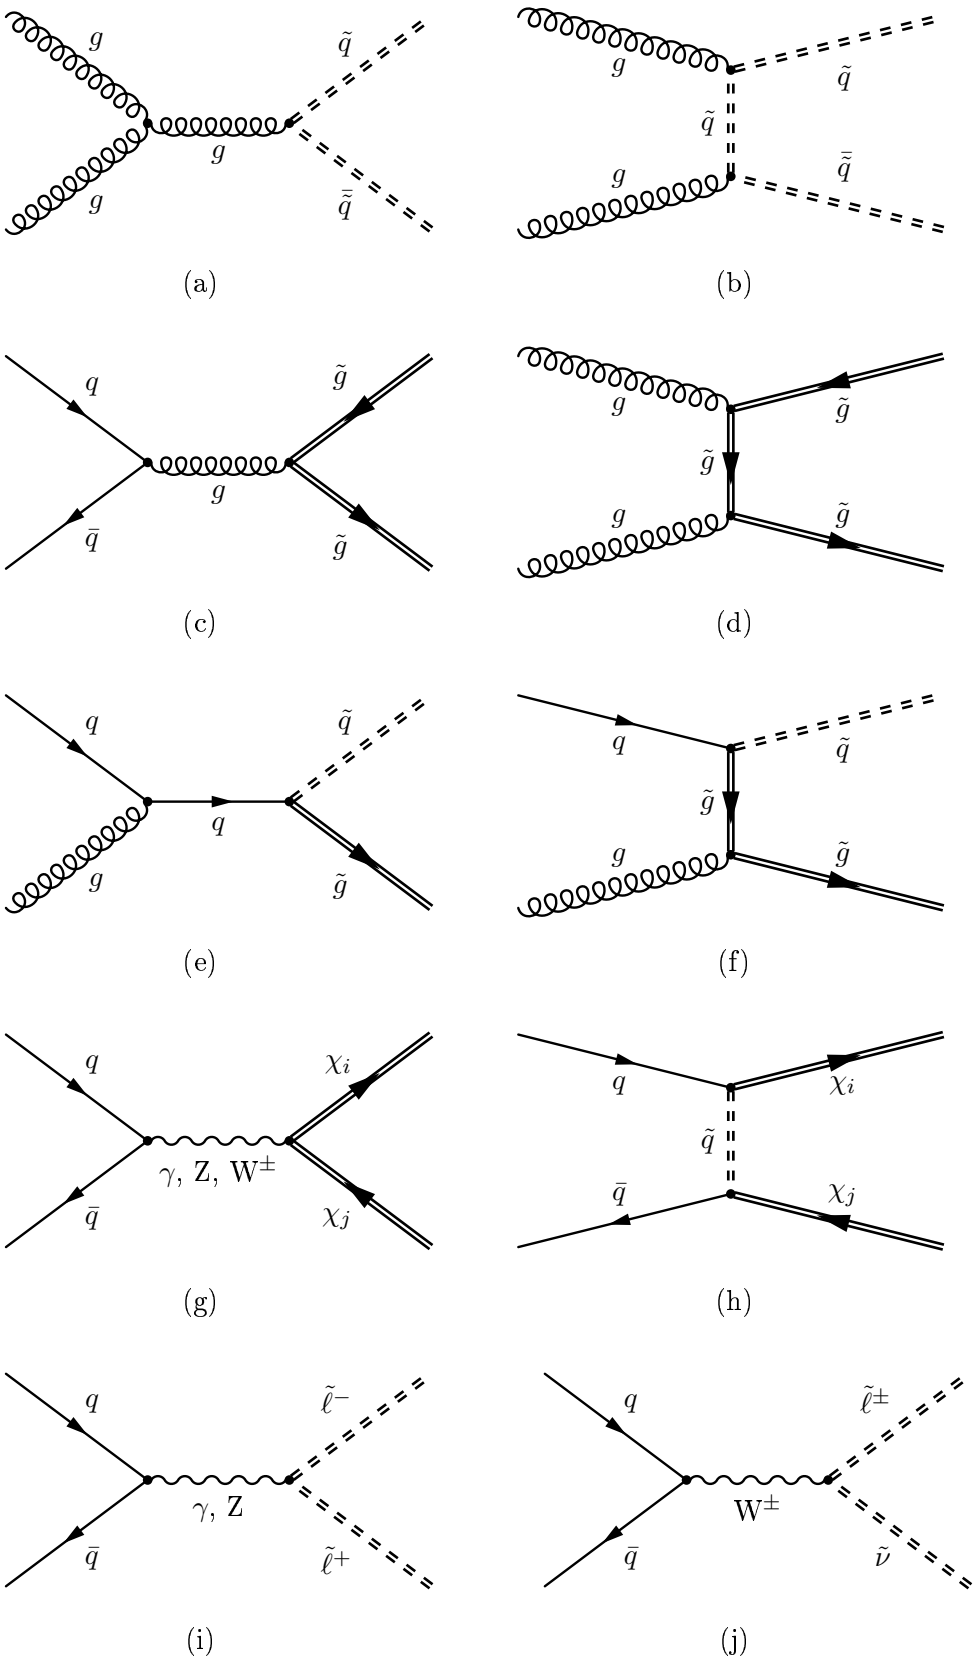

Figure 6.3: *Representative lowest order Feynman diagrams responsible for production of (MSSM) SUSY events at the LHC.*

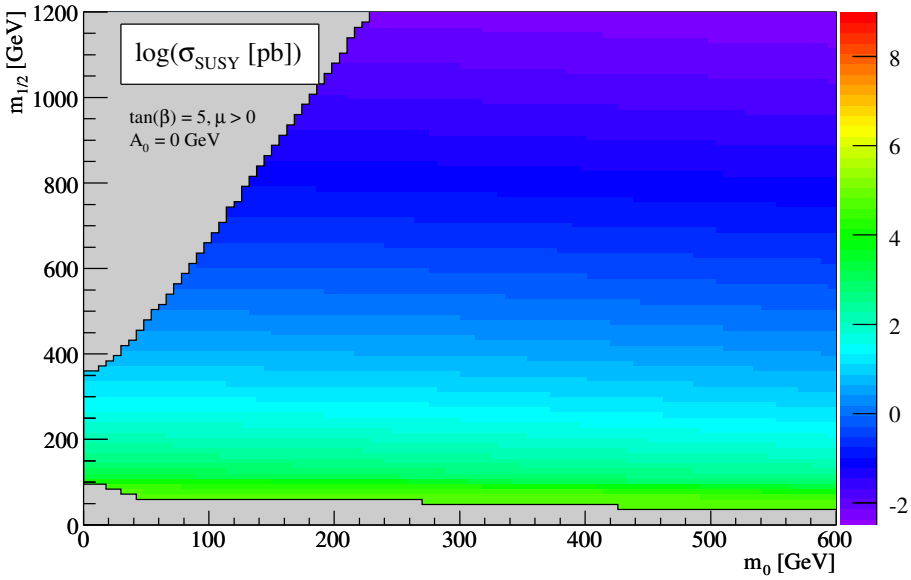

Figure 6.4: *Visible cross section for production of SUSY particles in proton-proton collisions.*

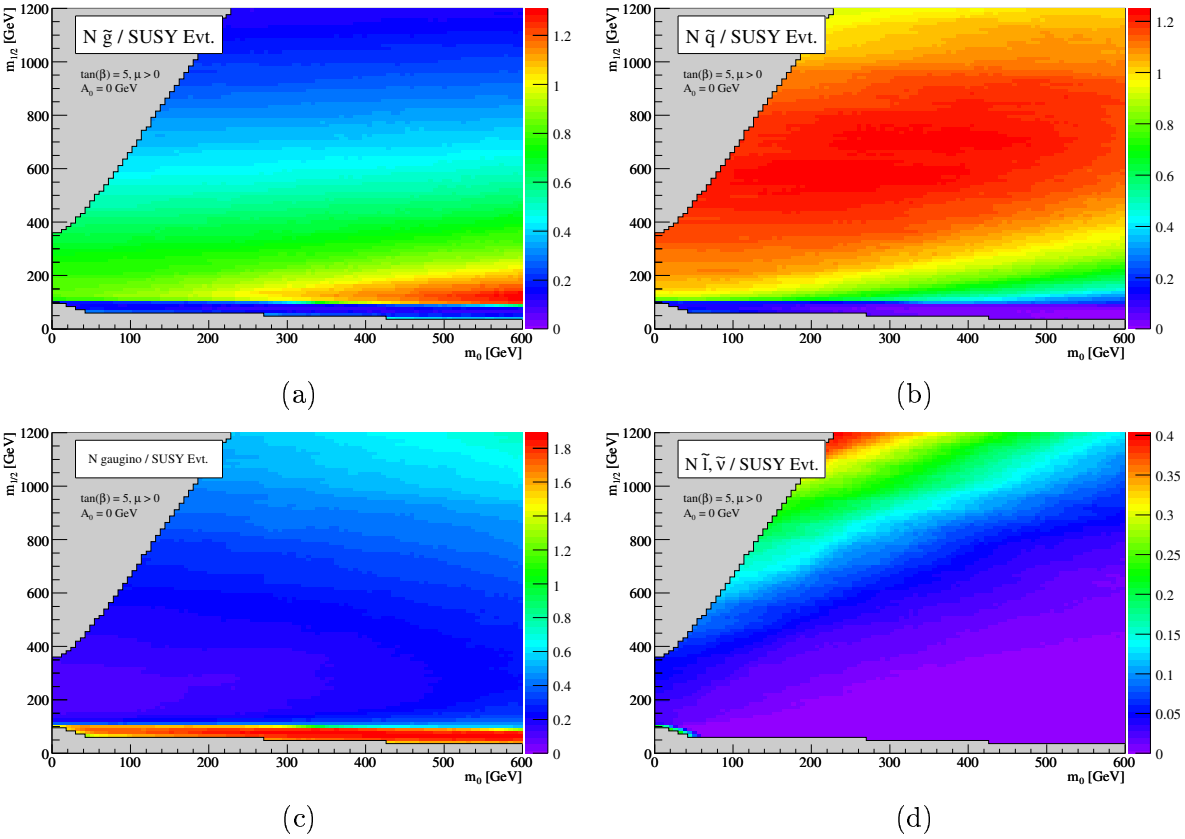

Figure 6.5: *Directly produced sparticles, here shown for $A_0 = 0$ GeV, $\mu > 0$, and $\tan\beta = 5$. Note that "gaugino" in the figure text actually refers to all $\tilde{\chi}_i^0$ and $\tilde{\chi}_j^{\pm}$ states (and thus also the Higgsino component). Note that the scale of the four plots are not similar.*

not really surprising given that they try to solve the same problems and are constrained by the same direct and indirect observations.

Several theories are supplying a weakly interacting long lived particle that is a potential candidate for being the content of the cold dark matter. Events where such are produced will in general be characterised by the associated missing momentum. As explained in the previous section, the total energy and momentum component along the beam axis going into the hard process is undetermined, which means that only the missing transverse momentum (or, missing transverse energy) can be used as a signature. In order to be able to do so reliably, it is essential that the amount of fake missing momentum is kept under control. This requires a quite hermetic detector, covering as much of the 4π of solid angle as possible – and with high and uniform reconstruction efficiencies. Even theories without a dark matter candidate might of course produce missing momentum through decays to neutrinos, merely long lived weakly interacting massive particles or perhaps something more exotic like KK excitations of gravitons that cause momentum imbalance when leaving the 4 dimensional brane.

When the WIMP candidate furthermore owes its long lifetime to a new global symmetry (such as KK parity or R parity), decays of the new physics particles will always result in at least one daughter that also carries the new global charge. Events with new physics will thus always consist of a pair of new particles being produced directly and, unless these are amongst the lightest new particles, a whole "cascade" of decays will subsequently take place, until the lightest new particle – the dark matter candidate – is reached. Each of the decays in the cascade will likely produce rather energetic SM particles. In such scenarios new physics events will be characterised by a large amount of missing transverse momentum combined with many highly energetic physics objects (jets, leptons, etc.). This is exactly the scenario expected in UED and SUSY. Of course, the new global symmetry might be violated in a very small way, and the lightest new particle might thus decay itself, producing yet another group of decay products. Depending on its actual lifetime this decay might take place either essentially as part of the primary vertex, as a secondary vertex, somewhere inside the detector or outside. Every single new physics event might thus be accompanied by two such decays, which could be really rather remarkable depending on the available decay modes (imagine 4 Z bosons accompanying each new physics event). Of course, the new global charge is not strictly conserved and the lifetime of the lightest new particle is significantly less than the age of the universe, then there is no longer a cosmological requirement for it to be neutral. If it carries colour-charge, it could, from a collider experimental viewpoint, result in the production of very heavy hadronic states with one parton carrying all the mass. Such an "R–Hadron" might undergo several reactions on its way out of the detector, and each time the regular quarks might be exchanged. The experimental signature would thus be a very heavy hadronic state working its way out of the detector while changing its charge several times on the way [Kra04].

If on the other hand the new global charge is indeed conserved, then all new particles will cascade down to the dark matter candidate that will escape undetected. Due to the unknown total momentum component this further has the immediate implication that one will never be able to measure the momentum of all decay products of any of the new particles, which means that they can not be detected directly as a traditional mass resonance, and one will instead have to rely on observation of threshold effects and total production cross-sections in order to perform observations and mass measurements. On the other hand, the fact that all produced new particles will result in cascade decays with particles emitted at each step could lead to some quite dramatic events, rather different than those of the QCD background. Thus, in addition to looking for large amounts of missing transverse momentum, one should also search

for events with several hard jets, leptons or photons. It is also natural for many models to predict the production of a large number of particles with top quantum numbers, due to its large Higgs coupling, and it will thus be important to look for events with b jets also.

All in all, it should be relatively straight-forward to discover most of the scenarios of new physics discussed in this thesis. But it is important to note that most of these models have deliberately been constructed as minimal models, where as few new parameters and particles as possible are introduced on top of the Standard Model. It might be that new physics exists at the TeV scale, but as a non-minimal model with different experimental signatures. It is thus important to develop model-independent analyses, which in a systematic way simply searches for deviations from the Standard Model in a very large number of different events.

The actual identification of observed new physics would require measurements of a larger number of coupling constants as well as the spin of several of the new particles. For models predicting towers of new particles (i.e. the KK towers in UED), it would furthermore be very fortunate if not only the first, but also the second level of excitations could be seen. These will obviously reside at higher mass scales so it is far from certain that this could be achieved at the LHC. In order to verify SUSY, the observation of scalar particles carrying the quark and lepton quantum numbers would be rather convincing, and if the couplings of these particles furthermore were identical to those of the SM fermions with similar quantum numbers, it would certainly be hard to argue that the new physics was not SUSY. However, even in such a fortunate scenario the nature of SUSY breaking would probably not be clear. Finally, a Higgs boson candidate might be observed early on due to its strongest couplings. But in order to verify that it is indeed a Higgs boson, and of what type, it would be necessary to measure its couplings to a large number of different particles, including some with less fortunate couplings.

In summary, if the logic analysis of the problems of the SM is not flawed, there should exist new physics at the TeV scale, and it seems rather unlikely that the LHC should not be able to discover it with great certainty. However, getting proof of physics beyond the Standard Model is one thing. Identifying its nature is another, and will require precise measurements of particle spins and branching ratios in more than just one or two channels. It is hard to believe that such a task could be completed in a fully satisfactory manner at a hadron collider, thus underlining the case of a future linear e^+e^- collider. However, there is no reason for the LHC not to go as far as possible also in this regard, and the final part of this thesis will present a feasibility study for a method to discriminate between e.g. supersymmetry and universal extra dimensions by investigating spin effects of the particle that could either be a slepton or a lepton Kaluza Klein excitation.

Part II
The ATLAS Experiment

7 The ATLAS Experiment

ATLAS is one of two multi-purpose detectors that will collect data at the LHC, and it is currently in the final stage of being assembled in its huge underground cavern. The design of the detector has been guided by a desire to be able to discover new physics signals, including Higgs bosons, over as broad a mass range as possible and with a large variety of experimental signatures, in order to provide several and redundant means of discovery.

The experiment, shown in figure 7.1, consists of an inner tracking system, providing tracking in a 2 T magnetic field over a large volume. In addition to providing directions and momenta of charged tracks, this central system provides both precise vertex reconstruction and electron identification, primarily based on transition radiation. Electromagnetic and hadronic calorimeters follows the central tracker, measuring the energies of all particles with electromagnetic or strong interactions. The combined calorimeter system is almost hermetic and provides excellent constraints on the overall missing transverse energy as well as accurate reconstruction of jets and identification of electrons and photons. Outside this hermetic system is located a muon spectrometer, whose immense size defines the physical dimensions of the experiment. This system consists of layers of detectors interspersed in a strong magnetic field of up to 4 T delivered by a system of low-density air-core toroids. This ensures that momentum measurements of muons can be performed up to very high momenta and independently of the backgrounds in the central detector that might be substantial at full luminosity. A sophisticated multi-level trigger system selects interesting events out of the high rate of QCD background processes, with the first level trigger being based on the calorimeter and muon system only. Taken together, this means that a variety of physics signatures can be reliably extracted with high precision:[1]

- Reconstruction of charged tracks with transverse momentum measurements at $\eta < 2.5$.
- Reconstruction of electromagnetic showers from photons and electrons at $\eta < 3.2$, with fine calorimeter segmentation, tracking and particle identification based on transition radiation allowing for separation between electrons, photons, π^0s and other particles at $\eta < 2.5$.
- Identification of muons and measurements of their transverse momentum over a very large momentum range and in presence of high backgrounds ($\eta < 2.7$).
- Reconstruction of jets and hadronic particles with energy measurements at $\eta < 4.9$, and tagging of b-jets at $\eta < 2.5$.
- Identification of hadronic tau decays at $\eta < 2.5$.
- Reconstruction of transverse missing energy with complete ϕ acceptance for $\eta < 4.9$.

In addition to the direct search for new physics, the experiment will of course also take advantage of the fact that an unprecedented large number of established but interesting SM particles such as W bosons, top quarks and b-hadrons will be produced at the LHC, and programmes exists to study the properties of these. In fact, one of the first and important

[1] The ATLAS coordinate system follows the standard pattern of having the z–axis along the beam axis, the origin at the central interaction point and azimuthal and polar angles denoted ϕ and θ respectively. η is the pseudo-rapidity defined as $\eta \equiv -\log[\tan\theta/2]$.

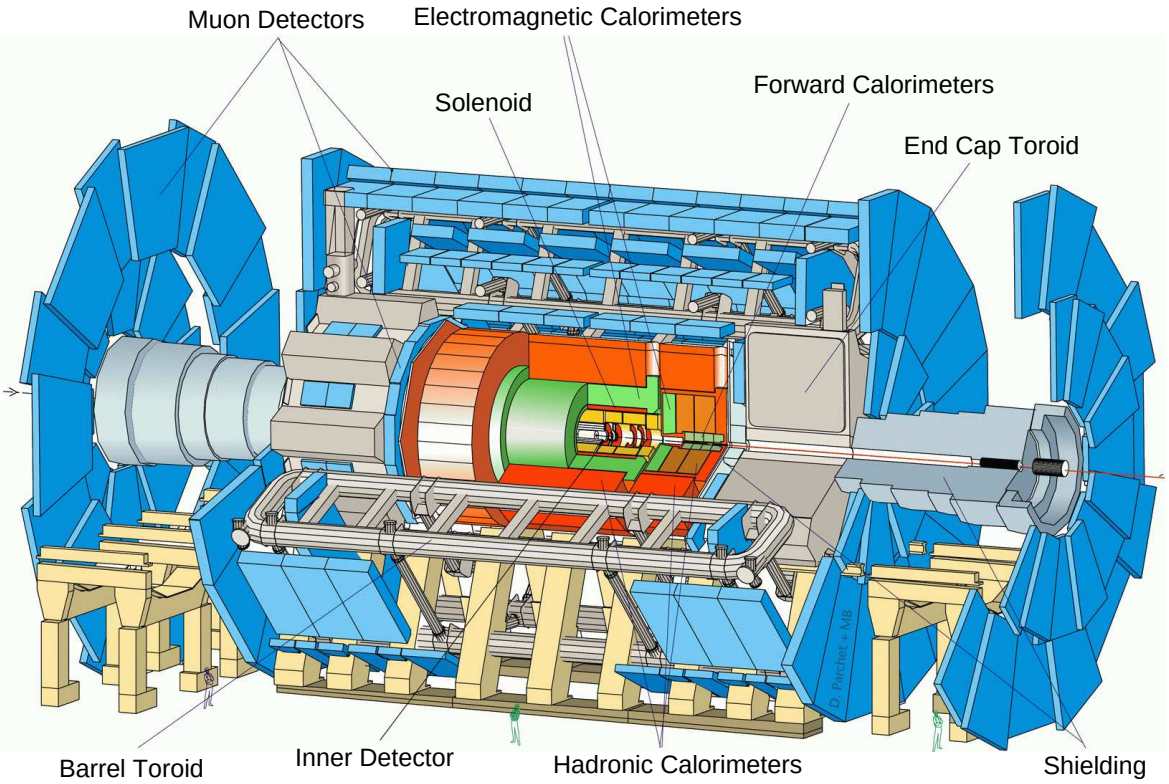

Figure 7.1: *The ATLAS Experiment. Source: [ATL99a].*

milestones that must be achieved soon after the experiment starts collecting data will be the "re-discovery" of SM states, which will be used to verify that the detector is working as intended, that the integration of the various sub-systems have been properly carried out, and the resulting overall detector performance properly understood.

The ATLAS detector has been described in great detail in several technical design reports,[1] and the reader is referred to these for more detailed descriptions. The rest of this section will be dedicated to a brief overview of each of the key sub-systems of the experiment. Special emphasis will be put on the Transition Radiation Tracker (TRT) as it is the subject of the third part of the thesis.

7.1 Magnet System

The magnet system in ATLAS is composed of a total of 25 individual superconducting components, 24 of which are making up 3 air-core toroids for the muon spectrometer and the last being a large solenoid around the inner tracker. Superconducting technology is the optimal choice given the energy requirements and excess heat produced by regular magnets. This means that the magnet systems need to be accompanied by large cryogenic systems to maintain superconductivity. Furthermore, the demands on the mechanical support structure for, in particular, the toroid magnets due to the combined effects of their weights and magnetic forces are severe. The barrel toroid system consists of eight large flat oblong coils of dimen-

[1] The physics performance TDR [ATL99a][ATL99b] and references therein is a good starting point. References to other sources are given in the text and figure captions.

sions 5.4 m $\times$ 25.3 m assembled by a number of supporting elements in a toroidal structure with an inner radius of 4.7 m. Similarly, each of the two end cap toroids consists of 8 coils of dimensions 4.5 m $\times$ 5 m arranged in a toroid with an inner radius of 1.65 m, but housed in one large cryostat in each end cap. The end cap toroids are located inside of the barrel toroid in a configuration, which is offset by a rotation of 22.5° around the beam axis compared to the barrel toroids, and with the coils extending radially slightly into the volume between the barrel coils in a manner resembling large gears, thus ensuring full radial field coverage. Despite the large volume provided with a magnetic field by the toroid magnets, the overall layout is very airy with a low average density, thus keeping the impact on tracking resolution due to multiple scattering low. The peak fields in the toroid systems is 4 T and the bending power is between 2 Tm and 8 Tm, depending on the direction of the traversing muon. Accurate field maps as well as sophisticated alignment systems and tracking algorithms will obviously be essential for reliable track reconstruction and momentum measurements.

The central solenoid provides a field of 2.09 T at its centre. The limit on this strength comes mainly from the fact that a high field strength makes particles with low values of transverse momentum curl up before they reach the detector elements. As it is, momenta of 16 MeV and 335 MeV are required to reach the first and last layers of the inner tracker respectively, for particles emitted at $\eta = 0$. Furthermore, a higher field strength will obviously also be associated with higher demands on the construction of the tracker and front-end electronics within. The axial extend of the central solenoid is slightly less than that of the inner trackers, which means that the axial part of the field will fall to around 1 T at the end of the tracker, where it will also pick up a substantial radial component.[1] Therefore advanced tracking algorithms and precise field maps will also be of paramount importance in the inner tracker.

In order to perform momentum measurements at as high energy scales as possible for a given field strength and solenoidal size, the radial extend of the inner tracking volume is maximised by placing the electromagnetic calorimetry outside the central solenoid. This means that the solenoid is constructed with as low amount of material as possible, in order not to degrade the performance of the electromagnetic calorimeter more than necessary. One specific consequence of this is that the electromagnetic calorimeter and the central solenoid are housed inside one common insulating vacuum vessel, thus eliminating two vacuum walls in front of the calorimeter.

Finally, it should be noted that the energy stored in the ATLAS magnet systems, and especially in the barrel toroid, will be an enormous 1.3 GJ, equalling more than one eight of the energy stored in the entire LHC. Fail-safe quench protection systems are thus constructed, which allows the stored energy to be quickly and safely deposited in resistor arrays outside the detector.

7.2 Inner Tracking System

The inner tracking system in ATLAS, illustrated in the $r - z$ plane in figure 7.2, combines high granularity silicon detectors at inner radii with a large number of gaseous proportional straw detectors at large radii and is limited in radial extent to around 1.1 m by the cryostat of the central solenoid magnet. The system is designed to provide reliable and precise tracking,

[1] It is of course important to note that the radial component will vanish at $r = 0$ for symmetry reasons, and that the combined fields from all of the ATLAS magnet therefore will be parallel with the beam directions everywhere, thus not influencing the beam in its path around the LHC.

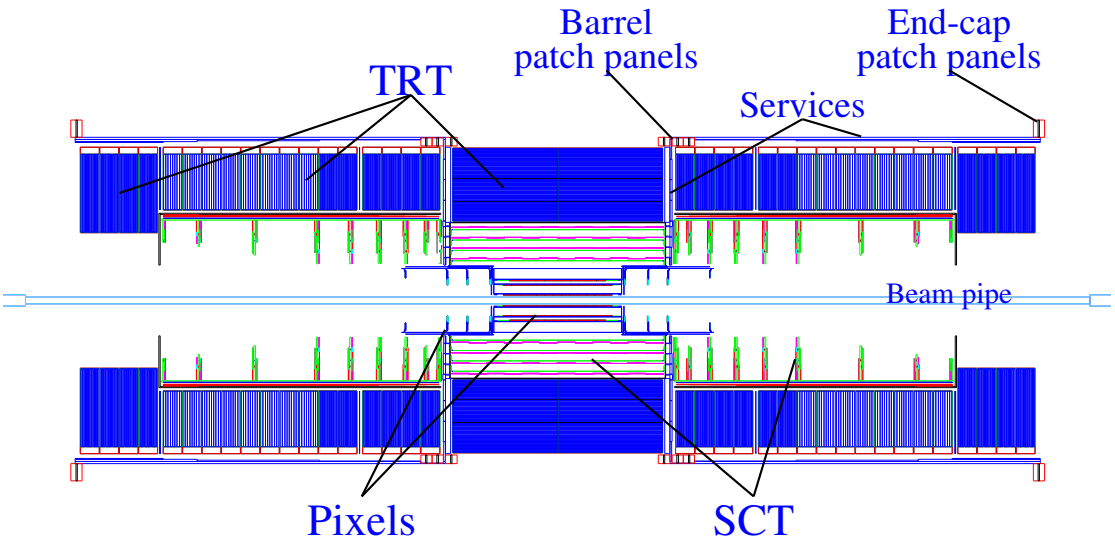

Figure 7.2: *View of the inner tracker in the $r-z$ plane. Source: [ATL97a].*

vertex reconstruction and momentum measurements of charged particles in the high intensity environment of the LHC. Additionally, it also provides electron identification capabilities, supplementary to those of the electromagnetic calorimeter. The tracking at large radii also allows for reconstruction of late decays of particles such as $K_S^0 \rightarrow \pi^+\pi^-$.

The innermost silicon based part of the detector utilises two different technologies. The sensitive part of the detector placed closest to the interaction region consists of (2×6) cm^2 silicon modules with a combined area of 2 m^2, each instrumented with $\sim 6\cdot 10^4$ small rectangular pixels providing hit resolutions of about 12×70 μm^2 [ATL05]. The modules are mounted in three cylindrical layers in the barrel, with modules overlapping slightly for complete azimuthal coverage as can be seen in figure 7.3.a. In the end cap the modules are assembled in cylindrical discs in the $r-\phi$ plane, with varying inner radii so as to provide 3 hit measurements for all particles at η up to 2.5 for all particles originating in the interaction region at $|z| < 11.2$ cm. The pixel system is expected to be able to measure the impact parameter of high-energy muons with a resolution of 10 μm.

The two-dimensional segmentation of the silicon in the pixel system not only provides tracking space points with precise three-dimensional information, but also reduces the channel occupancy in environments with high particle fluxes as well as making the overall performance more fault-tolerant in the presence of radiation damage of individual channels. The downside is the extra complications and expenses involved in the connection and readout of the large number of channels ($80 \cdot 10^6$ for the entire pixel detector). For this reason, pixel based technology is supplanted by silicon strip detectors with 80 μm pitch at larger radii, where the problems of occupancy and radiation are less severe. The third component of the measured space points is provided by a small stereo angle of 40 mrad between pairs of silicon strip detectors assembled back to back. The active area of the ATLAS silicon strip system will amount to 60 m^2 – more than an order of magnitude more than previous generations of silicon strip detectors. As for the pixel system, the silicon strip modules are mounted in cylinders in the barrel and in discs covering pseudo-rapidities up to 2.5 in the end cap.

The silicon strip system contains a total of $6 \cdot 10^6$ channels which, although significantly lower than the corresponding number in the pixel system, means that both systems face

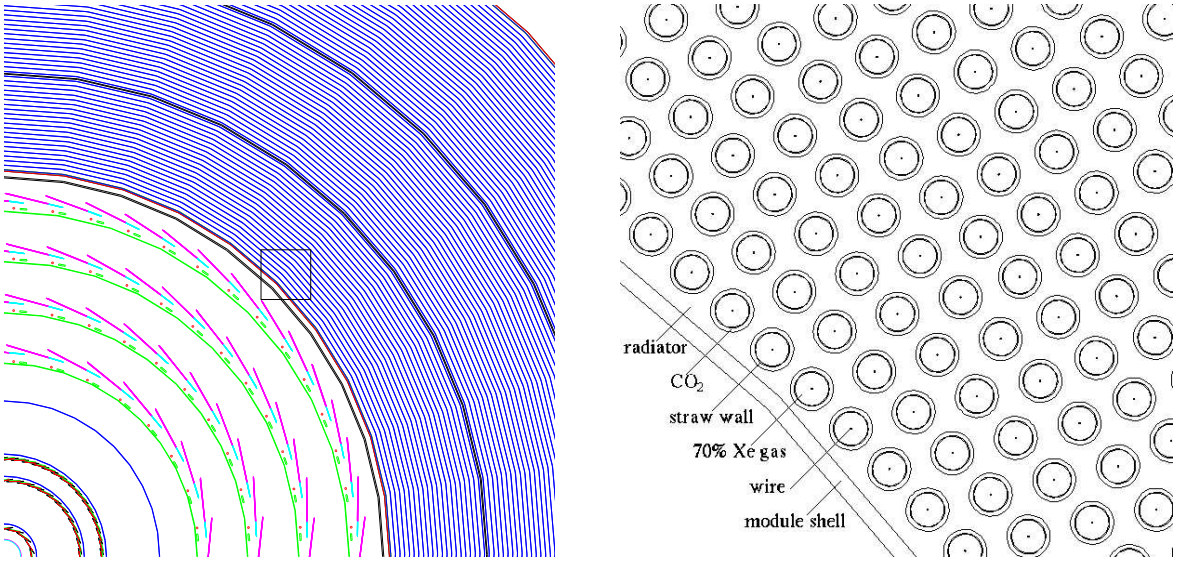

(a) (b)

Figure 7.3: *Figure (a) shows a cross-sectional view of the inner tracker in one quadrant of the $r-\phi$ plane at $\eta = 0$. Particles leaving the interaction point crosses 3 layers of silicon pixels, 4 layers of silicon strip modules (with two strip measurements in each layer), and 73 layers of straw tubes (providing about 36 hit measurements). Figure (b) provides a close-up view of the TRT straws corresponding to the square inlet in figure (a). Note that the view of the straw layers in figure (a) is somewhat simplified. Source: [ATL97a].*

complex issues of connecting, powering and reading out the large number of channels, without producing too much heat. Both systems thus have radiation hard front-end electronics mounted directly alongside the silicon sensors, and are operated under extensive cooling.

In addition to the ~ 11 silicon based hit measurements of each track at $\eta < 2.5$, about $3 \cdot 10^5$ straw tubes of 4 mm diameter in the TRT provides an average of ~ 36 hit measurements in the form of drift-cylinders with individual resolutions of around 140 μm. The straws are constructed from 72 μm thick kapton tubes with diameters of 4 mm and the 31 μm thick wires are made of gold-plated tungsten. Furthermore, the volumes between the straws are filled with materials of rapidly varying dielectric constant, resulting in the emission of transition radiation, when these "radiator" materials are passed by ultra relativistic charged particles. If only a single transition took place, the intensity of the emitted radiation would be proportional to the γ factor of the passing particle, but when a high number of transitions follow each other closely, interference effects change the emitted spectrum qualitatively in such a way that it achieves a threshold suitable for particle identification.[1] For the multi-layer setups used in the TRT the effect is such that only particles with gamma factors higher than $\mathcal{O}\left(10^3\right)$ emit transition radiation, and does so in the form of X-ray photons with energies of $\mathcal{O}\left(5\text{keV}\right)$. Furthermore, the transition radiation will essentially be emitted in the direction of flight, and thus follow the same path as that of the passing particle and traverse the same straw volumes. The active gas in the TRT consists of 70% Xenon, capable of absorbing X-ray photons, and the amount of energy deposited in the active straw volumes will consequently be

[1] The theory of transition radiation is described in great detail in general [AYM75] and in the context of the TRT [Ege98][Dri04] elsewhere.

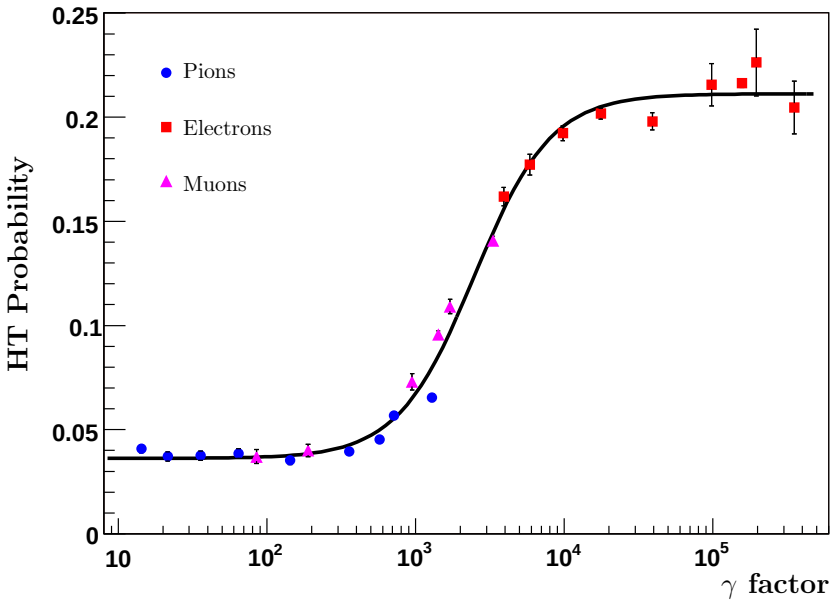

Figure 7.4: *Probability to exceed the high threshold in a single straw as a function of the gamma factor, as measured in the 2004 ATLAS combined test beam. Data points are plotted for electrons, muons and pions and represents energies between 2 GeV and 180 GeV (electrons and pions) and between 9 GeV and 350 GeV (muons). Source: [MP06].*

substantially higher for particles with a high gamma factor, which essentially means electrons at the energy scales relevant at the LHC. In order to capitalise on this, the TRT front-end electronics implements two different thresholds on the amplitude of the registered signal: A *low threshold* (LT) determining the signal arrival time used for tracking, and a $\sim$ 20 times larger *high threshold* (HT), used to record the large energy deposits from transition radiation (the wire signals and front-end electronics are discussed in further detail in section 9). Figure 7.4 shows the probability, measured in a test beam in 2004, for a particle passing a single straw to cause the HT to be exceeded as a function of its gamma factor. The non-vanishing value of the plateau at low values of gamma is due to dE/dx tails. Straws without a HT additionally contains information about the identity of the passing particles through the time over threshold of the LT which is essentially a measurement of dE/dx, although the discriminating power of this is dampened somewhat due to the signal shaping in the electronics. In total, the TRT is expected to provide typical levels of $e-\pi$ discrimination at a level between 2σ and 2.5σ [MP06]. Furthermore, the silicon based trackers can also contribute with dE/dx measurements to some degree, and the electromagnetic calorimeter will obviously also provide electron identification through measurements of shower shapes and energy measurements for E/p.

Although approximately similar in performance, the layout and choice of radiator materials differs somewhat between the TRT barrel and end cap sections. In the barrel section, 144 cm long straws are arranged axially and symmetrically around $z = 0$ in 73 radial layers, providing track measurements in the $r - \phi$ plane, which is the most important for momentum determination. As no stereo angle is used the z coordinate is undetermined by the TRT alone and tracking will thus usually be seeded from the silicon based detectors.

The straw layers extend from $r \approx 56$ cm to $r \approx 107$ cm, and straws within a layer are located with inter-wire distances of ~ 6.8 mm, and neighbouring layers have a relative shift of

half an inter-wire distance. The two halves of the 31 μm diameter wires in the barrel straws are electrically separated at $z = 0$, in the sense that they actually consists of two separate wires fused together by a tiny glass pearl, meaning that each of the 53 000 barrel straws are associated with two separated readout channels. The barrel wires are furthermore held in place at $z = 0$ by a plastic "twister", minimising the wire-sag considerably, given that the gravitational sag of a horizontal wire is proportional to the fourth power of the unsupported wire length [Tok90]. The wire signals are sampled and discriminated by electronics mounted at each end of the barrel, where also wire tension plates are located and high voltage is provided. In order to avoid occupancy issues, the inner 9 wire layers are constructed differently, with insulating pearls at $z = \pm 40$ cm. This leaves the central section of those straw tubes inactive, but the most important section of the inner layers are in any case near the end of the barrel, since these are required in order to ensure a high number of active straws in the transition region between the barrel and end cap.

To simplify the construction and to improve mechanical stability, the barrel straws are assembled in 96 modules that are placed in three coaxial layers with a 32-fold symmetry in ϕ around the beam as illustrated in figure 7.5.a, with modules in the three layers each containing 329, 520 and 793 straws, respectively. In order to ensure hermiticity in ϕ, the module shapes are constructed such that no module wall is parallel to the radial direction. Instead the module walls not pointing along ϕ are tilted with respect to the radial direction by an angle of 27.5°. Inside the 300 μm thick carbon composite module shells, the tubes are packed inside a large numbers of 3 mm thick fibre sheets, all with holes cut out for the straws to pass through, and are kept in place at five points along the axis of the modules by internal alignment sheets made of 100 μm thick kapton, as shown in figure 7.5.b. Figure 7.6 shows a close-up of the fibre material. It is clear that a particle passing through the fibre will undergo numerous transitions in and out of the fibre material, which thus acts as the radiator for the barrel part of the TRT.

The detector layout in the TRT end caps, where 246 000 of the 298 000 straws are located, is more straight-forward as the required azimuthal homogeneity is achieved by arranging the straws like spokes on a wheel. Each such arrangement of straws is denoted a "straw plane" and, analogous to the straw arrangement in the barrel, neighbouring straw planes are rotated by $\Delta\phi/2$ with respect to each other, where $\Delta\phi$ represents the ϕ difference between two straws. As in the barrel, no stereo angle is implemented, and the measurements thus occurs in the $\phi - z$ plane, with the ϕ component providing the main momentum resolution due to the direction of the solenoidal field as well as the direction of the particles origination from the interaction point. Radiator layers are placed between the straw planes and each layer typically consists of $16 \sim 15$ μm thick polypropylene foils placed with a regular spacing of ~ 200 μm. This is different from the situation in the barrel, where the fibre material gives irregular spacings and the interference effects, giving rise to transition radiation photons, are thus somewhat stronger in the end cap.

For mechanical stability and to maximise the modularity in the construction, 8 or 16 straw planes and the radiator planes between them are assembled inside one *wheel*, consisting of inner and outer support shells and have front-end electronics mounted at the outer shell. Each TRT end cap consists of 18 wheels, out of which four are staged for economical reasons and will not be part of the initial layout of the detector (the staged wheels can be identified in figure 7.2 as those that are furthest from the interaction point and with smaller inner radii). The 14 wheels in the initial version of the detector provide measurements down to $\eta \approx 2.1$, and contain 768 straws in each internal straw plane.

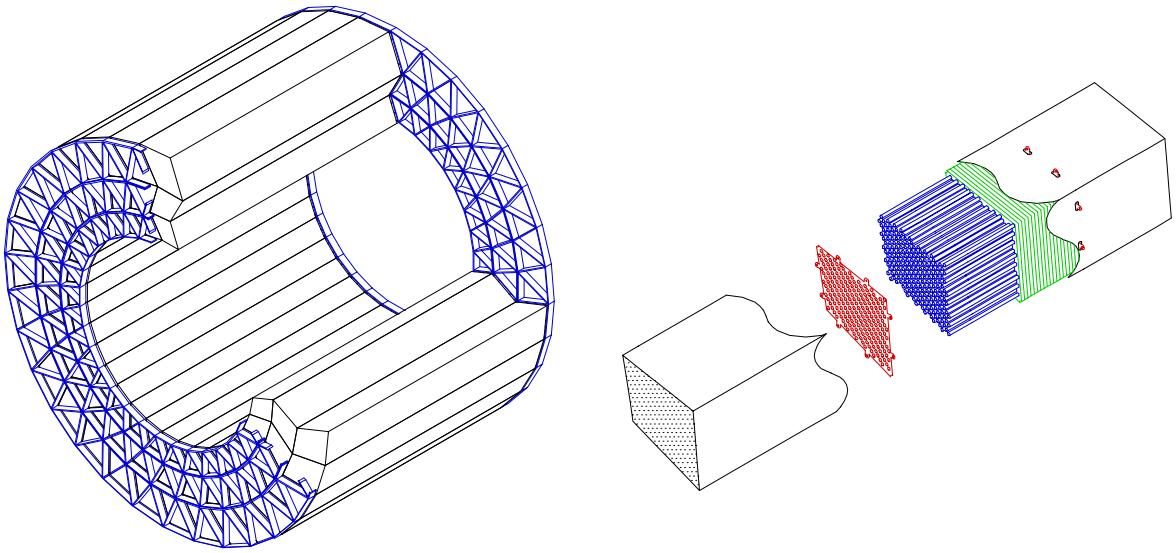

(a) (b)

Figure 7.5: *Figure (a) shows the overall layout of the TRT barrel modules as well as the support frames. Figure (b) illustrates the contents of a given module: Straw tubes (blue) are placed inside stacked fibre sheets (green) and are supported at five points by kapton sheets (red). Source: [ATL97b].*

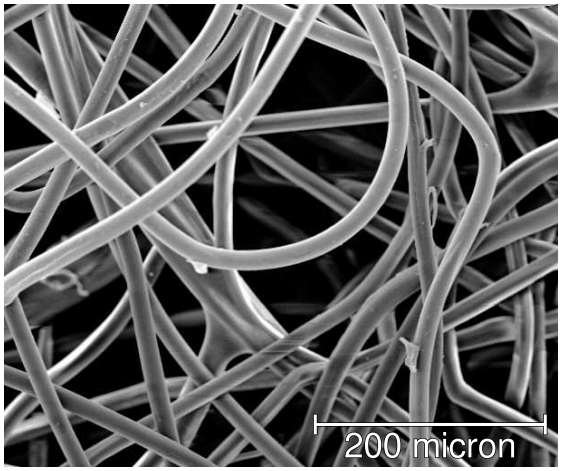

Figure 7.6: *Close-up of the TRT barrel radiator fibre material. Source: [ATL97b].*

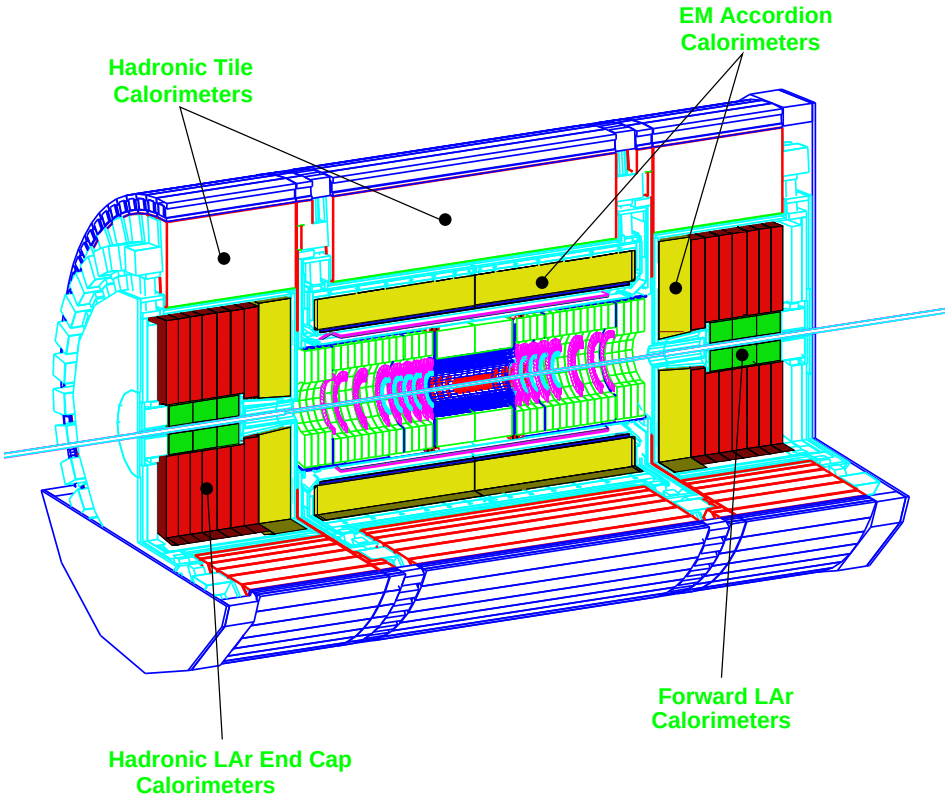

Figure 7.7: *Layout of the calorimeter systems. Source: [ATL99a].*

7.3 Calorimeters

The ATLAS calorimeter system, shown in figure 7.7, follows the standard pattern of having an electromagnetic calorimeter placed in front of a thick hadronic calorimeter. The former is responsible for providing measurements of electromagnetic showers due to electrons and photons, while the latter encapsulates the irregular hadronic showers. This combination is important for reconstruction of jets, for providing measurements of the missing transverse energy, and for reliably containing all particles apart from muons (and neutrinos), which is essential for the muon detectors located outside.

The total amount of material in front of the electromagnetic calorimeter amounts to approximately $2-4$ radiation lengths depending on eta (with a local peak of 7 interaction lengths in the gap region $1.4 < \eta < 1.5$), but most of this is located close to the front of the calorimeter in the solenoidal coil, cryostats or services with less impact on performance. A particle reaching the outer radius of the silicon based trackers will only have traversed ~ 0.2 radiation lengths of material, and at the outer radius of the TRT it will have encountered an additional ~ 0.3 radiation lengths, with the main source of material being the services for the silicon detectors.

The electromagnetic calorimeter has a thickness of half a metre and more than 24 radiation lengths. It covers $\eta < 3.2$ and is based on layers of liquid argon sandwiched between layers of lead absorber. This makes the detector inherently radiation tolerant as the liquid argon can be continuously replaced during the detector lifetime. In order to achieve complete azimuthal coverage, the layers are positioned in an accordion structure as illustrated in figure

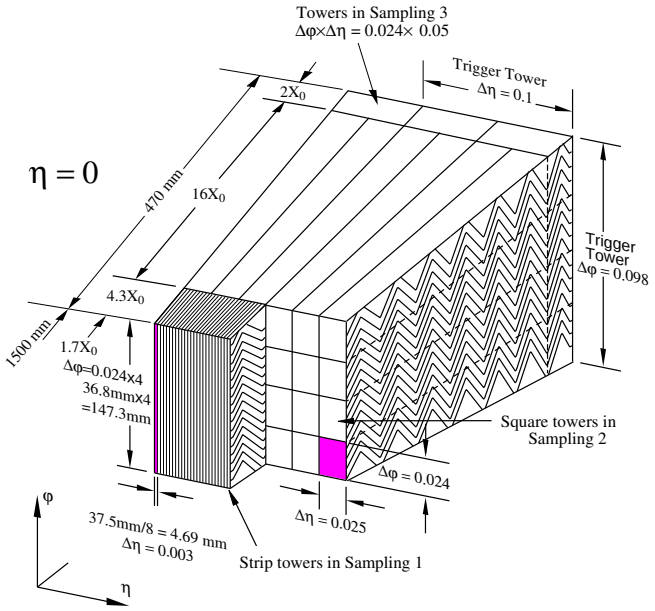

Figure 7.8: *Illustration of the segmentation and accordion geometry of the electromagnetic calorimeter. Source: [ATL96b].*

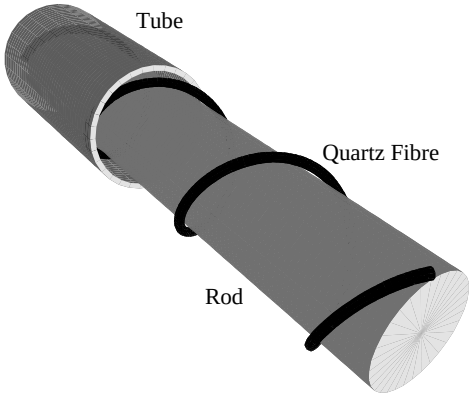

Figure 7.9: *Illustration of a narrow gap rod in the forward calorimeter. Source: [ATL94].*

7.8, which also shows the readout segmentation of the detector. A total of 1024 1.8 mm thick lead absorber sheets are folded in the accordion shape and stacked in ϕ. Each of these are mechanically reinforced by having two thin (0.2 mm) steel sheets glued on each side. The 3.9 mm thick gaps between the absorber sheets are filled with the liquid argon, and the signals deposited herein are read out by 300 μm thick copper plated sheets of kapton with a geometry mimicking that of the lead absorber sheets, and placed directly between them (held in place by a light honeycomb spacer). The signals from four double-gaps are collected to give a ϕ segmentation of $2\pi/256 \approx 0.025$, while the (pointing) readout segmentation in both η and the longitudinal direction is handled by etching appropriate divisions into the copper layers. The specific segmentation varies along η, but the general image conveyed by figure 7.8 holds throughout the electromagnetic calorimeter: Three longitudinal samplings with the η segmentation decreasing outwards. The first sampling contains a very fine segmentation of $\Delta\eta \approx 0.003$, and enhances π^0/γ separation and e/γ particle identification, while the second sampling absorbs most of the energy of electromagnetic showers. For the purpose of the first level trigger a less fine-grained 0.1×0.1 segmentation in $\eta \times \phi$ is read out with no longitudinal segmentation.

In order to compensate for the energy loss in the material in front of the calorimeter, a presampler consisting of a 10 mm thick liquid argon detector is placed in front of the electromagnetic calorimeter at $\eta < 1.8$. In total the expected energy resolution for the electromagnetic calorimeter is given by [AEBC06]:

$$\frac{\sigma_E}{E} \approx \frac{10\%}{\sqrt{E/\mathrm{GeV}}} \oplus 0.2\% \tag{7.1}$$

Where the first term is the sampling term and the second is due to detector imperfections and miscalibrations.

The hadronic calorimeters in ATLAS are based on three different technologies. At large radii where the radiation levels are relatively low, the so-called tile calorimeter is located which is a sampling calorimeter based on alternating layers of "tiles" of iron and plastic scintillator that are readout by wave-length shifting fibres attached to photo-multiplier tubes. The tile calorimeter consists of a barrel part as well as two "extended barrel" wheels, covering the outer regions of the end cap. Each of these three sections is assembled from a number of wedges each covering $2\pi/64$ of ϕ.

At lower radii behind the electromagnetic end caps, the radiation levels are higher and the calorimetry is based on liquid argon sandwiched between copper plates and covers the η region down to 3.2. Finally, the forward calorimeter located in the region at $3.1 < \eta < 4.9$ must be able to function in an environment of extreme particle fluxes and energies, but is vital for a reliable determination of the missing transverse energy. The radiation hardness is once again achieved by using liquid argon as the active material, and the very fast response time needed to avoid too high occupancy and buildup of positive ions due to pile-up is achieved by letting the active volumes between the absorber be very small. This is difficult to achieve with the absorbers in a parallel-plate layout, so instead stainless steel tubes are placed in holes in the absorber and are filled with rods also constructed from the absorber material. The outer radii of the rods are ~ 250 μm smaller than the inner radii of the tubes and the rods are kept centred by quartz fibres as shown in figure 7.9. This provides the desired narrow gap for the liquid argon. In order to limit the transverse spread of the showers the absorber is, at the highest pseudo-rapidities, based on tungsten, whose interaction length is about half that of iron. At lower rapidities copper is used.

In total the hadronic calorimeter provides around 11 interaction lengths of material, and the expected energy resolution is [ATL96a]:

$$\frac{\sigma_E}{E} = \begin{cases} \frac{50\%}{\sqrt{E/\mathrm{GeV}}} \oplus 3\%, & \eta < 3 \\ \frac{100\%}{\sqrt{E/\mathrm{GeV}}} \oplus 10\%, & 3 < \eta < 4.9 \end{cases} \tag{7.2}$$

7.4 Muon Spectrometer

The detector elements of the muon spectrometer are located outside the calorimeters, within and outside the toroid magnet volumes. The 11 interaction lengths of material before the muon system ensures that the hadronic punch-through will be sufficiently low for track segments in the muon detector elements to be reliably identified as muons.

A total of four different technologies are implemented throughout the muon system and are implemented in three cylindrical layers ("stations") in the barrel and in vertical concentric structures in the end cap as shown in figure 7.10. Several hit measurements are performed at each station, providing good resolution and pattern recognition capabilities, as well as reducing problems with noise through coincidence requirements. In the barrel part, detector layers are placed at inner and outer radii of the toroid in addition to one layer placed inside, providing measurements of track sagitta for momentum resolution. In the forward region only measurements outside the field region are available, and the momenta are thus provided by the track angles in the layers.

The most pervasive detector type in the muon system is the so-called monitored drift tubes (MDTs), consisting of several layers of aluminium tubes with diameters of 3 cm containing central sense-wires and an argon based active gas mixture, providing local single-wire resolutions of ~ 80 μm in the form of drift-cylinders. The global positions and deformations of the

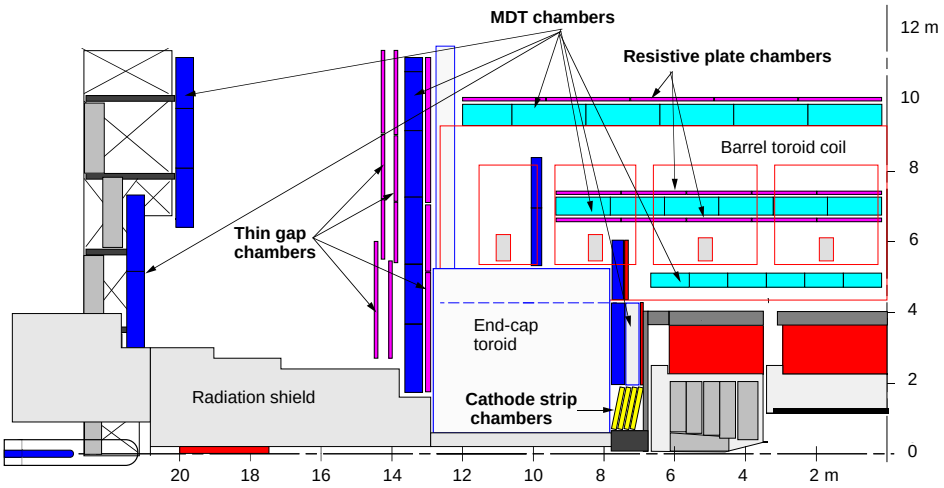

Figure 7.10: *View of one quadrant of the muon systems in the $r - z$ plane.*

tubes in the huge structure are constantly monitored through sophisticated laser alignment systems mounted directly on the structures, allowing the local hit information to be translated into the global reference system with negligible impact on the resolution. As the longest drift tubes are ~ 6 m in length, wire sagging can be considerable and this must also be taken into account, with some tubes even bent to follow the slope of the wire within.

In the region nearest to the beam pipe at $2 < \eta < 2.7$, the high particle fluxes makes the operation and performance of large radius drift tubes problematic due to issues of occupancy and the associated buildup of positive ions. Consequently detectors based on cathode strip chambers (CSC) replaces the drift tubes in this region, and delivers resolutions of ~ 60 μm. The CSCs deal with the occupancy issues by having a higher granularity, smaller gas gaps, and have been specifically designed so as to be highly resilient to the high levels of radiation in the forward part of the detector.

In addition to providing precise track measurements, the muon system is also, along with the calorimeters, responsible for providing data to the first level trigger system. This means that data with a time resolution much less than the 25 ns bunch crossing period is needed, and with a relatively low latency. The maximal drift time of the drift tubes is ~ 500 ns, making them unsuitable for that purpose. Therefore other faster and less precise detectors are interspersed between the MDTs, with the additional benefit of providing the second coordinate measurements inherently unobtainable in parallel drift tubes. In both the barrel and the end cap the measurements with fast response times are based upon detectors with very narrow volumes of gas, thus ensuring the necessary short drift times. In the barrel, resistive place chambers with parallel Bakelite plates separated by very small gas gaps are used, and in the end cap the chosen technology are so-called thin gas chambers, where sense wires are placed inside the narrow gap.

7.5 Trigger and Data Acquisition System

As illustrated in figure 6.2, the LHC production cross sections of interesting physics such as Higgs or SUSY events is not only expected to be rather low, but is expected to be at least seven orders of magnitude lower than that of minimum bias QCD events. This has two

consequences: First of all the instantaneous luminosity at the LHC needs to be very high, which is achieved by short 25 ns bunch separations and multiple (~ 25) interactions at each bunch crossing. Secondly, it is essential to be able to filter out minimum bias events with a rejection factor of at least 10^7, while keeping $\mathcal{O}(1)$ efficiency on the selection of interesting physics. Furthermore, at an event rate of 40 MHz and expected event sizes of ~ 2 MB, the raw data rate will be an astounding 80 TB/s, which is clearly outside current technological reach both concerning network speeds and available storage. Consequently the rejection needs to happen online and within fractions of a second, with large parts of the rejection taking place in specialised hardware and based on a subset of the recorded data.

The online filtering in ATLAS is carried out by a three level trigger system illustrated in figure 7.11, with the filtering at each level being based on increasing amounts of information and more refined selection criteria. The first level trigger decision is based on reduced granularity data from the muon and calorimeter systems and must be taken with a latency of no more than 2.5 μs. This means that full event data corresponding to ~ 100 bunch crossings are stored in pipelines in the front-end electronics on the detector. Upon a level one accept, data from the relevant bunch crossing are read out by read out drivers (RODs) into a number of read out buffers (ROBs) outside the detector. The first level trigger decision is taken in specialised hardware and is based on programmable trigger menus looking for high p_t objects such as muons, isolated electrons/photons, hadronic taus, jets, missing transverse energy, or total transverse energy. The first level trigger provides a rejection factor of around 500 and outputs an event rate of 75 kHz.

Instead of reading the entire event from the ROBs, the second level trigger uses the output from the first level trigger to identify regions of interests, i.e. the regions of the detector where hard physics objects were observed, and selectively only pulls the data from the ROBs where the relevant event fragments are located. The addition of information from the inner tracker and higher granularity data from the muon and calorimeter systems combined with a larger possible latency of $\mathcal{O}$(ms), allows the level two trigger, based on specialised software, to provide an additional rejection factor of around 100. In addition to straight-forward refinements of the first level trigger decisions, the second level trigger also has the possibility of using topological criteria such as requiring the invariant mass of low-energy muon pairs to be consistent with b and c hadrons. As a result of level two trigger decisions, ~ 1 kHz of complete events are read out from the ROBs and transferred to event filter farms. Here the full events are reconstructed using special optimised versions of the offline reconstruction and analysis software, with detailed calibration and alignment constants applied. The third level trigger provides an additional rejection factor of about 10, meaning that a rate of 100 events per second are passed on for permanent storage.

The precise content of the trigger menus at each of the three levels is presently the object of intense studies of the optimal trade off between the redundant and efficient selection of interesting physics events and the required reduction in background rates. Such studies will of course continue when the detector starts taking data, and are foreseen to result in gradual updates as the knowledge of detector and physics processes increases or the focus of particular physics goals shifts. The precise trigger criteria are thus not known at the present time, but when based on a combination of two or more signatures current baseline requirements involve p_t thresholds around 6 GeV for muons, 15 GeV for electrons, 20 GeV for photons, 10 GeV[1] for missing transverse energy, $50 - 120$ GeV for jets and 10 GeV for taus decaying into hadrons.

[1] In the most optimistic scenarios. The actual threshold might be considerably higher.

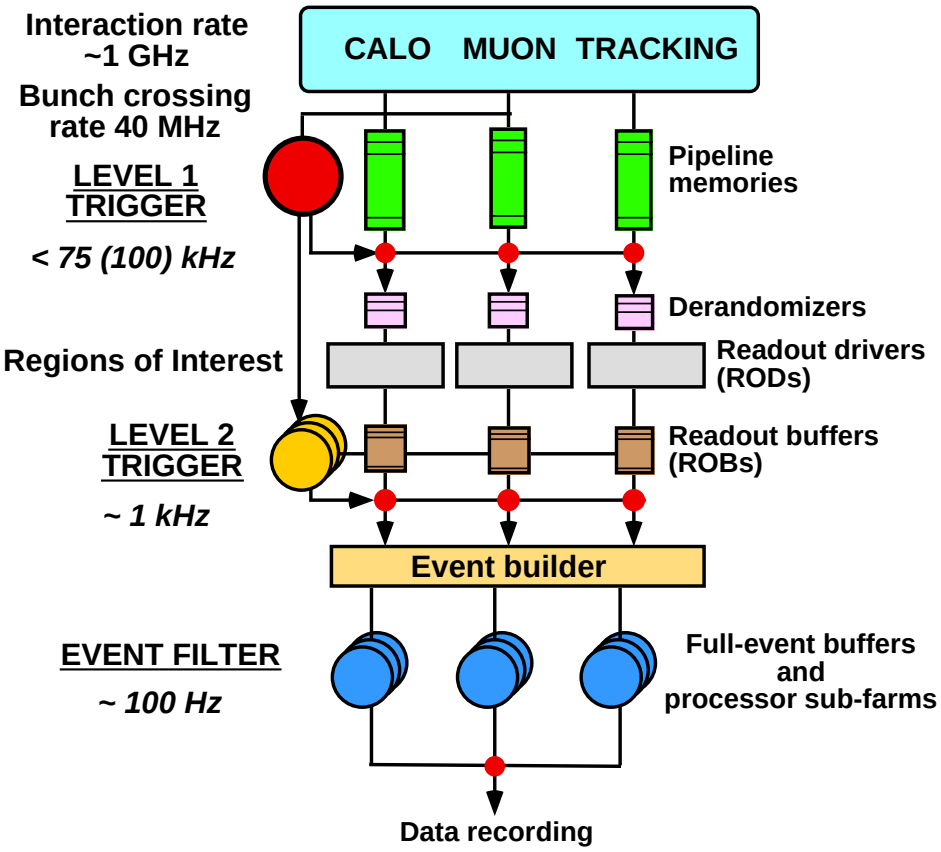

Figure 7.11: *Illustration of the ATLAS trigger system. Source: [ATL99a].*

When a single signature is the basis for a trigger decision, the requirements will necessarily be somewhat higher.

Even with the considerable reduction in the data stream applied, data will still be output at rates of 720 GB/hour and 2 PB/year (assuming a standard accelerator year of 10^7 s). It is far from trivial how to store and subsequently access these amounts of data, and the solution taken by the ATLAS collaboration is to utilise GRID technology to transparently spread the load over a large number of data centres, while using an analysis framework which allows one to only pull the data actually relevant for a given analysis from the remote storage centres. The latter part is achieved by starting analysis from heavily distilled event data and dynamically pulling additional data for interesting events as needed.

Part III
Simulation of the Transition Radiation Tracker

One of the major challenges encountered in both design and data analysis phases of an experiment like ATLAS is to understand how observed physics objects relates to underlying physics. Due to the level of complexity inherent in this process, it inevitably relies on the availability of realistic and trustworthy Monte Carlo simulations.

For physics analysis studies the first step is the phenomenological simulations of the collisions of beam particles, and the resulting (semi-) stable particles emerging from the interaction region. This part of the simulation, which, apart from beam content, is not specific to a given experiment, will be driven by various hypotheses regarding both parameters of established particles as well as assumptions about new physics, and should be based upon state of the art theoretical methods. At hadron colliders the uncertainties in this simulation step will always include those of QCD – mainly in the parton distribution functions and in the hadronisation.

The next step is to simulate the passage of these generated particles through the detector and the resulting detector response. The former requires a precise geometrical description of the detector material and magnetic field maps. The resulting particle trajectories, including effects of scatterings, energy losses and production of secondaries, is handled by specialised software tool kits, with the choice in ATLAS being GEANT4 [G403] (the sole exception to this is the production of transition radiation in the TRT, which is implemented as a custom process [Dam04]). All particle steps through the geometrical volumes corresponding to active detector elements (such as silicon modules or gas inside drift tubes) are recorded. Based on these "simulation hits", sub-system specific code subsequently takes over in order to model the detailed detector response and the resulting digitised output of the front-end electronics. This final "digitisation" step usually involves both simulation of the physics effects, such as the ionisation created by the particle traversing the active detector element, as well as of the signal processing in the electronics including signal shaping and the final discrimination.

Section 8 will give a short introduction to the way geometry in ATLAS is implemented in the GeoModel framework [BT04][SAB+04][Bou03], with focus on the TRT. Section 9 will then describe the digitisation of the TRT with focus on recent developments. The specific issues of noise and detector non-uniformity are delegated to sections 10 and 11, with the former describing the findings of a study in a test beam in 2004 and the latter describing the actual model implemented based on these. Finally, section 12 will, based on results from the digitisation, discuss how to get the optimal track resolution of the TRT in face of non-Gaussian effects, and will furthermore compare the output of the digitisation with data from the 2004 test beam.

8 Description of Detector Geometry in Software

Although simulations of the ATLAS detector were originally carried out in a `FORTRAN` framework and based on GEANT3,[1] the geometry description was reimplemented in `C++` around 2003 in order to better take advantage of the benefits of a modern object oriented language, and to allow a better integration with the rest of the ATLAS software.

Previously the geometry had been implemented separately in both simulation and reconstruction code. While perfectly adequate for early performance studies, such a strategy is, however, relatively inflexible as it requires synchronisation of geometry in several separate implementations. To solve this, the geometry description is separated from both simulation and reconstruction code and implemented in one central place in the GeoModel framework.

[1] This original framework is described in detail in [ATL99a, chap. 2].

Each sub-detector implements both raw and readout geometries: The former is represented in memory as nodes of material, shapes and transformations arranged in an ordered tree data structure, and includes the possibility of efficient parameterisations of large groups of identical volumes such as the TRT straws. For the purpose of the readout geometry, additional information might be stored as required in custom data structures. This might for instance be information related to how the TRT straws in the barrel are grouped into layers inside each barrel module, along with information about which layers follow each other in ϕ and r. The interface of the readout geometry can then be implemented as desired by the subsystems in order to facilitate optimal reconstruction performance and simplicity of reconstruction code. Behind the readout interface the internal algorithms access the raw geometry tree for information about the position of the active volumes, thus ensuring automatic synchronisation between the two geometries. Another advantage of the GeoModel approach is that (mis-) alignments of the detector elements can be introduced in a way, which affects both simulation and reconstruction by altering the values of the transformations stored in the relevant transformation nodes in the raw geometry tree.

The construction of the in-memory data structures for the geometry are driven by primary numbers in a versioned database. This means that a large number of geometries can be supported by a given software release in a transparent manner. Furthermore, these database versions are automatically embedded in the output of e.g. simulation jobs, and subsequent digitisation and reconstruction jobs running on that data will, in essence, automatically be configured to run with the same geometry, thus minimising issues of misconfiguration.

This support of multiple geometry versions is particularly important, since it means that data simulated with an older layout can be reconstructed with newer releases, while the geometry is allowed to evolve due to updated detector information. It also makes comparisions of the physics performance of different versions straight-forward, which is essential in order to ensure that any evolution is understood and based on actual improvements rather than inadvertent software bugs. Additionally it allows the code to support simulation of setups such as test beam or cosmic runs, where only parts of the full detector are present. As an example, figure 8.1 shows the six TRT barrel modules used in the 2004 test beam, including outer and inner support structures, but with the module shells as well as the radiator materials hidden in order to expose the straws and cooling tubes.

Figure 8.1 was created by the HEPVis visualisation tool and based on the in-memory translation of the raw geometry tree to one understandable by the Open Inventor 3D visualisation toolkit. In the same way the raw geometry tree is dynamically translated to a GEANT4 geometry tree during simulation jobs. While the straight-forward 3D visualisation of the geometry can provide the basis of an event display, it is also an essential tool for debugging the implemented geometry. First of all to investigate geometry clashes, but also to double check that the readout and raw geometries indeed are synchronised properly. The preferred way to do this is to make simulations of so-called geantinos, i.e. unphysical test particles without any interactions at all, and compare their expected straight trajectories with the global positions of the simulation hits as given by the readout geometry interface. An example of such a test is shown in figure 8.2, where one of the simulation hits, which should be created at the intersection between the track and straw, is obviously relocated to the wrong straw due to a problem with the straw identifier decodings.

The actual raw geometry implemented for the TRT is very close to the real layout as described in section 7.2. One difference is that the dead material in the form of support structures, electronics, cabling and services for the silicon detectors are approximated by

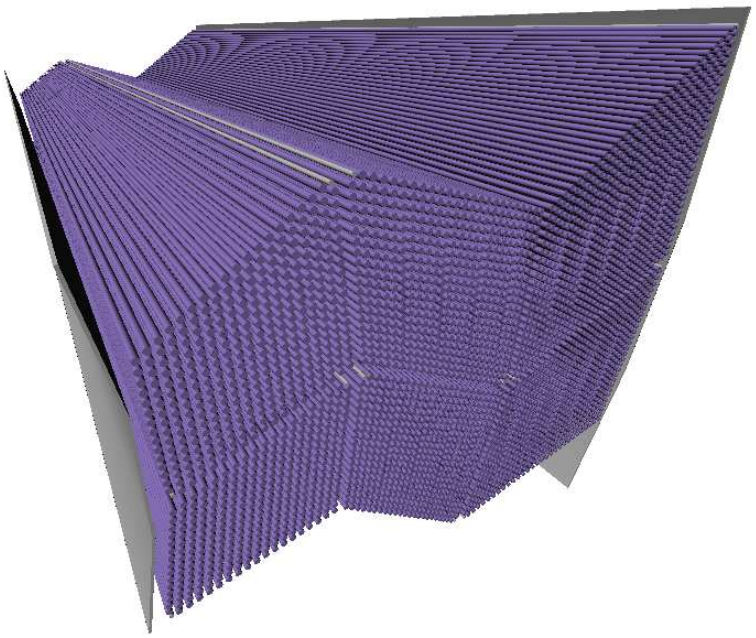

Figure 8.1: *A HEPVis visualisation of the GeoModel geometry used for the TRT in the 2004 test beam setup.*

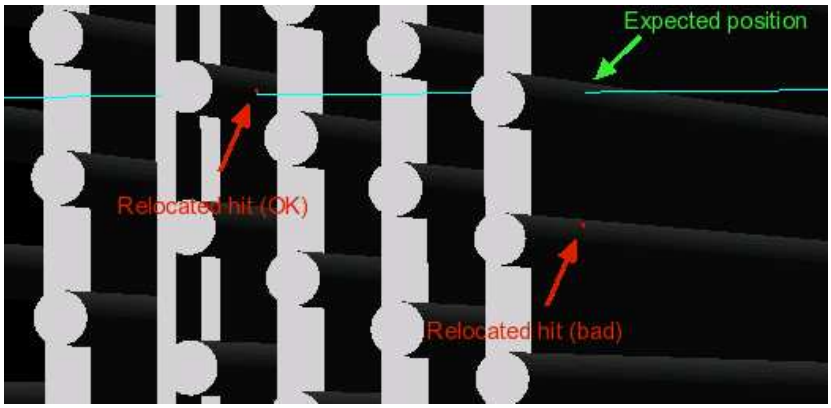

Figure 8.2: *A HEPVis visualisation of the TRT end cap geometry, showing the radiator planes and straws. The figure illustrates how the consistency between raw and readout geometries can be tested with geantino simulations as discussed in the text.*

a number of discs consisting of an averaged material. Likewise, the radiator materials are approximated by continuous volumes (seen as the planes between the straws in the end cap in figure 8.2). The detailed microscopic structure, important for the creation of transition radiation, is instead taken into account inside the custom transition radiation process which is active during the GEANT4 simulation.

9 Digitisation: Simulation of Detector Response and Front-End Electronics

While the geometrical description and simulation of energetic particles passing through the detector is generally handled in the same way for all subsystems, the digitisation, i.e. the simulation of the response in the active detector elements and the signal treatment in the front-end electronics, is handled by specialised code adapted to the needs of each subsystem.

In the case of the TRT the code needs to simulate the processes in the active gas such as ionisation, electron drift and multiplication in addition to the amplification, signal shaping and analog to digital conversion in the front-end electronics. Since the tracking functionality of the drift tubes is based on a precise relationship between drift times and track-wire distance at closest approach, issues of timing are important at every step. Section 9.1 will describe the parts of the simulation that take care of the processes in the active gas, while section 9.2 describes the processing in the front-end electronics.[1]

In addition to simulating the actual physics processes in the active gas and the processing in the front-end electronics, the digitisation code also passes appropriate Monte Carlo truth information along, and handles the "pileup" of minimum bias QCD background events due to the multiple collisions in each bunch crossing. This effect, which also includes particles created in neighbouring bunch crossings, is handled by having the framework merge simulation hits from different events before running the digitisation code. Furthermore, the code is based on versioned primary numbers, thus ensuring full backwards compatibility as well as allowing for a flexible integration of new features. Finally, it should be noted that the digitisation code needs to be extremely efficient, since each event at full luminosity involves $\mathcal{O}\left(10^5\right)$ straw crossings and the deposition of $\mathcal{O}\left(10^6\right)$ primary ionisation clusters. To achieve the necessary performance, caching techniques and lookup tables are used heavily throughout the code.

9.1 Physics of Ionisation, Drift and Gas Amplification

The physical processes in the TRT straws are essentially similar to those taking place in any system of drift tubes (see e.g. [Leo93]), but the main difference is that the straw radii of $r = 2$ mm is particularly small, meaning that the number of primary ionisations are also small and the random nature of their number and distribution along the track path will thus influence the performance of the detector heavily. In the TRT gas mixture of 70% Xe, 27% CO_2 and 3% O_2 at a pressure of 1 atm., the free path length between ionisations for a 80 GeV pion is around 165 μm. For a particle crossing the straw at perpendicular incidence to the wire this implies average numbers of primary ionisations along the track ranging from zero to

[1] The TRT digitisation code has a long history, with many contributers along the way, with the earliest `FORTRAN` versions being written and maintained by P. Nevski and F. Luehring, and the first `C++` versions due to D. Barberis and K. Assamagan.

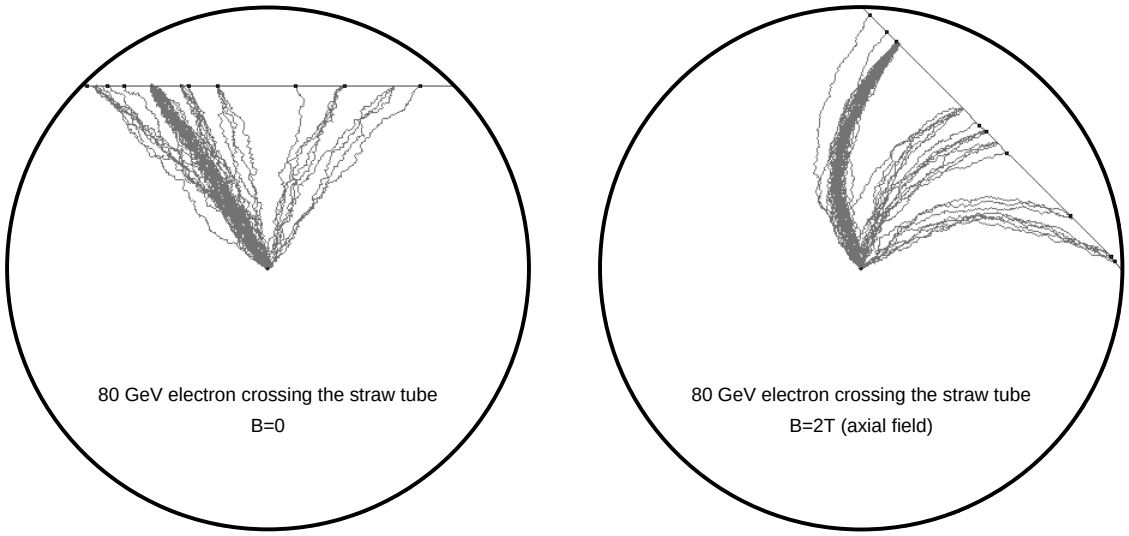

(a) (b)

Figure 9.1: *Examples of ionisation and electron drift with and without an external magnetic field, as discussed in the text. Source: [Cwe06].*

24 as the distance of closest approach between the wire and the track goes from $r = 2$ mm to $r = 0$ mm.

Figure 9.1 shows two examples where ~ 11 primary ionisation clusters along the track gives rise to electrons drifting towards the wire, driven by the field due to the potential difference of 1530 V between the tube (the cathode) and the sense wire (the anode). The distance at closest approach measurement, which is the one important for tracking, is given by the measurement of the time the low threshold is exceeded, which constitutes an estimate of the "leading edge" of the signal arriving at the wire. As the electrons essentially drift at a constant velocity around $v_{\text{drift}} = 50$ μm/ns due to the combined effects of the field and the constant collisions with the gas particles, this so-called $r - t$ relationship should ideally be given approximately by $r = v_{\text{drift}} t$. However, this will only be the case when a large amount of ionisation energy is deposited near the very central part of the track segment crossing the drift tube. Given the relatively low number of initial ionisation clusters, their random distribution is the major limiting factor on the overall tracking resolution and is likewise the major source of non-Gaussian resolution effects, which will be studied further in section 12, where it will also be discussed how to construct a refined $r - t$ relationship that takes such non-linear effects into account.

The digitisation code relies on the so-called photo absorption ionisation (PAI) model [AC80], in which ionisation cross sections of charged particles of given gamma factors and in a given gas are derived numerically from tabulated photon absorption cross sections for the gas. This is done by using the fact that both cross sections can be related to the (complex) dielectric constant of the media, which is a function of energy. Even without going into details, it is of course clear that absorption cross sections of real photons should be related to ionisation due to charged particles, since such ionisation processes are really the result of absorption of the (virtual) photons "surrounding" the energetic charged particle. The current object oriented `C++` implementation of the PAI model is based upon an earlier `FORTRAN` implementation by P. Nevski, and constructs the photo absorption cross section for the TRT

gas as a weighted average of tabulated absorption cross sections for individual atoms. Figure 9.2.a shows the resulting photo absorption cross section, where the energy levels around 5 keV, important for absorption of transition radiation photons, are clearly seen. Figure 9.2.b shows the resulting distribution of the energies deposited in the primary ionisation clusters. Irrespective of the details of how the energy is deposited during a particular ionisation process, the effective result should ideally be a number of freed electrons given by the ratio between the total deposited energy and the ionisation threshold of the gas: A single electron freed with enough surplus energy will itself be the cause of other ionisations, and an electron freed from an inner shell implies subsequent ionising radiation, when the level is repopulated from outer shells. In practise, however, a fraction of the energy will be deposited in other energy modes such as kinetic energy and excitation of molecular bindings. Thus, while the ionisation potential for the TRT gas mixture is around 12 eV, the average energy used per resulting free electron is $W = 25.3$ eV (based on [Cwe06, table 4-I]), and the number of freed electrons in a given cluster will be given by the formula:

$$N_e = \text{floor}\left[\frac{E_{\text{deposited}}}{25.3 \text{ eV}} + 1\right] \tag{9.1}$$

Where floor(x) is the function returning the largest integer less than or equal to x. It is clear that a lower value of W gives a better performance of the detector. For this reason, a noble gas such as Xenon is ideal as the main gas component, since such a gas has no vibrational or rotational modes and no molecular bindings, thus ensuring a relatively low value of W (and, of course, good detector ageing properties since it is inert). However, in a pure Xenon gas mixture, ultraviolet photons originating from ionisation clusters can start secondary ionisation clusters at sizeable distances from the original clusters. This effect would ruin the necessary $r-t$ relationships, and is avoided by the addition of CO_2 to the gas mixture. The additional energetic modes available in these molecules effectively absorbs the ultraviolet photons. The digitisation code thus relies on the approximation of all free electron emerging from the localised region of the initial ionisation cluster.

For the special case of Transition Radiation (TR) photons ending inside the active gas volume, they are assumed to have been absorbed fully in a single Xenon atom, and their entire energy are thus implemented in the code as being deposited in a single highly energetic ionisation cluster, with several hundreds of resulting freed electrons. Furthermore, as the model used to produce the TR photons during the GEANT4 simulation step assumes an ideal geometry with regular spaced radiators, the amount of transition radiation will tend to be overestimated somewhat, and particularly in the barrel where the radiator material consists of manifestly irregular fibres. On the other hand, the model does not account for the fact that transition radiation can also be produced to some degree when the straw walls are passed. The net result is that the amount of simulated transition radiation is approximately correct in the end caps, while it is 25% too high in the barrel. To correct for this, the digitisation code ignores 20% of the simulated TR photons in the barrel.

Electrons freed by the ionisation will be accelerated towards the wire, and will undergo elastic and inelastic collisions with the gas molecules along the way. At large radii compared to the wire radius, $r_{\text{wire}} = 15.5$ μm, where the field strength is relatively low, the frequent (elastic) collisions will keep the energy of the drift electron below the ionisation threshold of the gas, resulting in the effectively constant drift velocity towards the wire. Close to the wire, however, the field is so strong that the electrons will gain enough energy between collisions that they themselves will cause new ionisations, with the result being a cascade of electrons

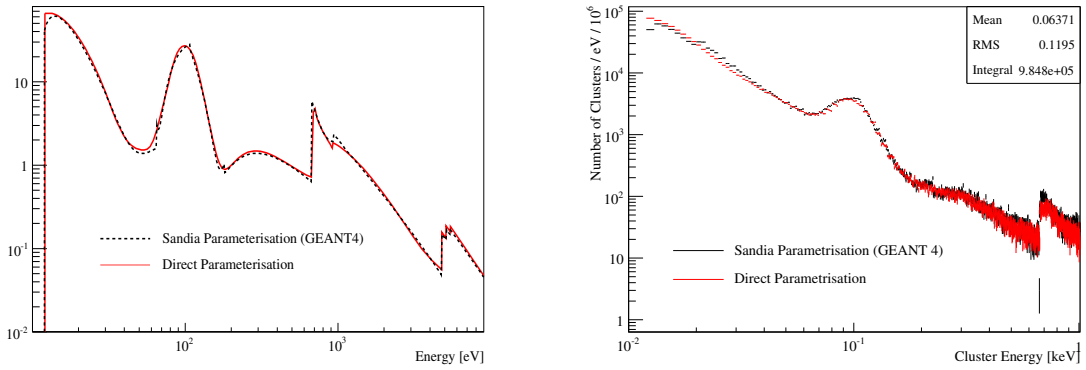

(a) (b)

Figure 9.2: *Photo absorption cross section for the TRT gas (a), and the resulting distribution of energies of the primary ionisation clusters (b). The black curves are based on the Sandia parameterisation of the cross section, which is available in GEANT4, while the red curves are based on more detailed and direct piecewise linear parameterisations due to I. Gavrilenko. While the difference between the two parameterisations appears small, it does give a 7% difference in the mean free path length, which is significant.*

reaching the wire and inducing a current strong enough to be measurable in the front-end electronics. Figure 9.3.a shows the average number of electrons reaching the wire for each free electron in the original cluster. At radii less than $r \approx 150$ μm, the field strength is high enough that gas multiplications occurs, and the threshold at $r \approx 150$ μm is clearly seen in the figure. To first approximation this implies that all clusters located at $r > 150$ μm will result in cascades of similar magnitudes, while electrons in clusters starting closer to the wire will have less distance over which to develop the cascade, resulting in a smaller gain factor that vanishes entirely at the wire.

It is important to note that the picture of electrons drifting towards the wire, cascading, and inducing a signal when reaching the wire is only effectively true. In reality, the movement of electrons and ions inside the straw volume essentially immediately induce currents between the anode and cathode due to capacitive couplings. However, only the last steps of the cascade developments contains enough charged particles for the effect to be significant. In fact, the main part of the current will be induced by the slower movement of the ions, rather than of the electrons [Leo93, chap. 6.5]. Figure 9.4 shows the signal following from a cascade in a TRT straw. The fast electron component only constitutes 3-5% of the integrated current, but is the one important for the $r - t$ relationship since that is given by the leading edge. The ion tail stretches over hundreds of nanoseconds, corresponding to tens of bunch crossings. If uncompensated, the net effect in a high occupancy environment would be unreasonable baseline fluctuations in subsequent events. To avoid this issue, the front-end electronics implements an active baseline restoration which effectively eliminates the ion tail, reducing its contribution to around 1.5% of the fast electron component. The correct way to implement the signal from a given cascade in the digitisation code is thus to register a cascade at the time a given drift electron reaches the wire, and to provide it with an appropriate shape which will be discussed in section 9.2.

An additional potential problem caused by drifting ions is that they might liberate electrons when reaching the cathode. In order to suppress such sources of free electrons, the

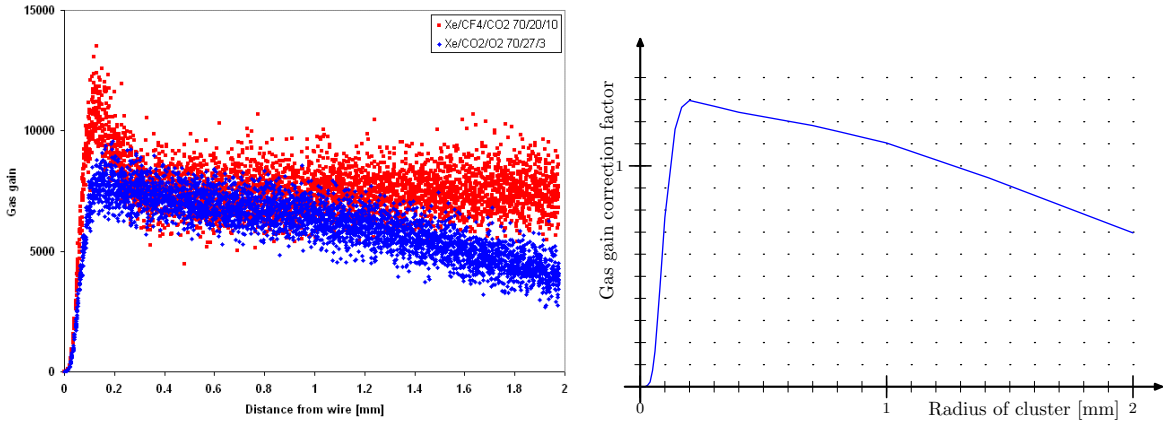

(a) (b)

Figure 9.3: *Figure (a) shows the average gas gain in the TRT gas, using the final gas mixture (blue) as well as a previously considered gas type (red). Each point represent the average gain per electron found by considering initial clusters containing about 250 electrons. Figure (b) shows the normalised gain correction factor used in the code, which is based upon a parameterisation of the curve in figure (a). Figure (a) was produced by P. Cwetanski, based on Garfield simulations.*

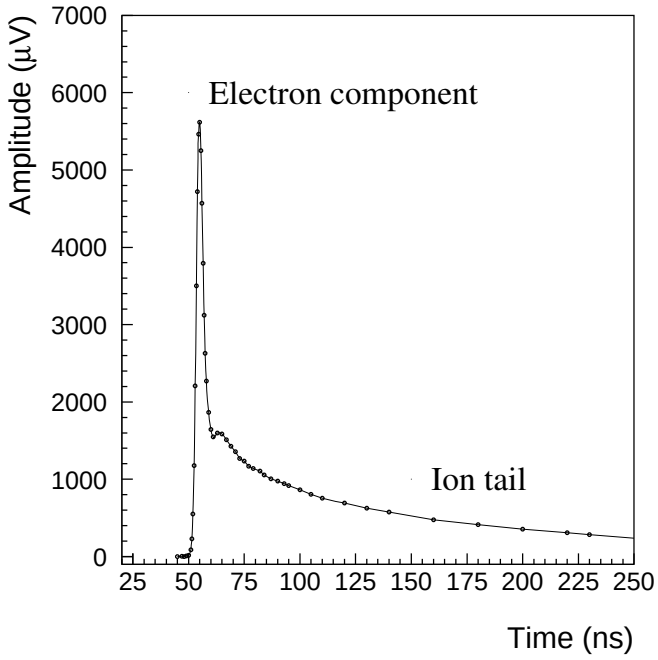

Figure 9.4: *Signal following from a cascade in a TRT straw, showing both the fast electron component and the long ion tail. Source: [ATL97b].*

small component of O_2 with high electron affinity is added to the gas mixture. The effect is that $\mathcal{O}(60\%)$ of drift electrons are reabsorbed by the gas by attaching themselves to the gas constituents, forming negative (molecular) ions. Figure 9.3.a includes the effect of this reabsorption and shows that, even though cascades only starts at $r \approx 150$ μm, the gain at $r > 150$ μm is not constant. This is because the cross section for reattachment depends on the field strength. Figure 9.3.a is based upon gas simulations performed with the dedicated programs Garfield [Vee84] and Magboltz [Bia99] by P. Cwetanski [Cwe06]. These simulations are very detailed and provides reliable results at the $\mathcal{O}(1\%)$ level. However, with $\mathcal{O}(10^5)$ straw crossings in each ATLAS event, the computational requirements would be infeasible if such detailed simulations were performed in large scale productions of simulated events. Therefore, the digitisation code is to a large extent based upon fast parameterisations of the results of the more detailed simulations.

In the digitisation code, cascade and reattachment effects are implemented as a flat reattachment probability of 0.6 for each free electron in a ionisation cluster. The surviving electrons each then contribute a signal on the wire which is given by $[25.3 \text{ eV}/(1-0.6)] \cdot x \cdot g(r)$ where x is a random exponential fluctuation factor with $\langle x \rangle = 1$ ($dP/dx = \exp[-x]$, $x > 0$), accounting for the fluctuations in the cascade developments and number of initial electrons, and $g(r)$ is a gas gain factor shown in figure 9.3.b, which is based on a direct parameterisation of figure 9.3.a.[1]

As should be clear from figure 9.3.b, $g(r)$ is normalised such that its average height considered over r is one. The reason for this approach, which ignores the absolute scale of figure 9.3.a but instead include the factor of 25.3 eV$/(1-0.6)$, is that the low threshold is implemented in the real life front-end electronics as a discrimination against a given electron count (current strength), while the interesting parameter from a physics point of view is the amount of energy deposited in the active gas by ionisation processes. The conversion between the two quantities will depend on details such as gas gain, reattachment probability and signal amplification in the electronics, and as such will be difficult to estimate accurately. Therefore each step of the digitisation will be normalised in such a way as to, on average, preserve the originally deposited energy all up to the point of discrimination against thresholds. In other words, a 300 eV threshold in simulation is defined such that it is, on average, passed by a single ionisation cluster of 300 eV or more. Ultimately, the threshold values will have to be tuned against real data, which will be done for the 2004 test beam data in section 12.2.

Having simulated the signal registered on the wire due to the individual free electrons from the primary ionisation clusters, it is of course also essential to model the registration time of the signals at the wire correctly, since these are essential for the tracking performance of the detector. As the effective drift velocity of the electrons towards the wire is around 50 μm/ns, while the final tracking resolutions are $\mathcal{O}(150$ μm), it is clear that the timing must preferably be modelled with a precision of at least $\mathcal{O}$(ns).

As a particle with $\beta \approx 1$ uses about 0.01 ns to cross the straw, electrons from all ionisation clusters created from a particular passing particle can be safely assumed to start their drifting at the same time. Of course, as it takes several nanoseconds for particles from the interaction region to reach the drift tubes, the precise time of flight of the particle is important, and is

[1] This is a somewhat ad hoc model, and it might be interesting to investigate whether improvements would result from a more fundamental modelling: The reattachment probability and gas gain should depend on the radius, and so should the spread in the gas gain. Finally, the number of free electrons resulting in each primary ionisation cluster, whose average, N, is given by equation (9.1), should fluctuate with $\sigma^2 = f \cdot N$, where $f = 0.19$ is the Fano factor of the TRT gas [Cwe06, table 4-II].

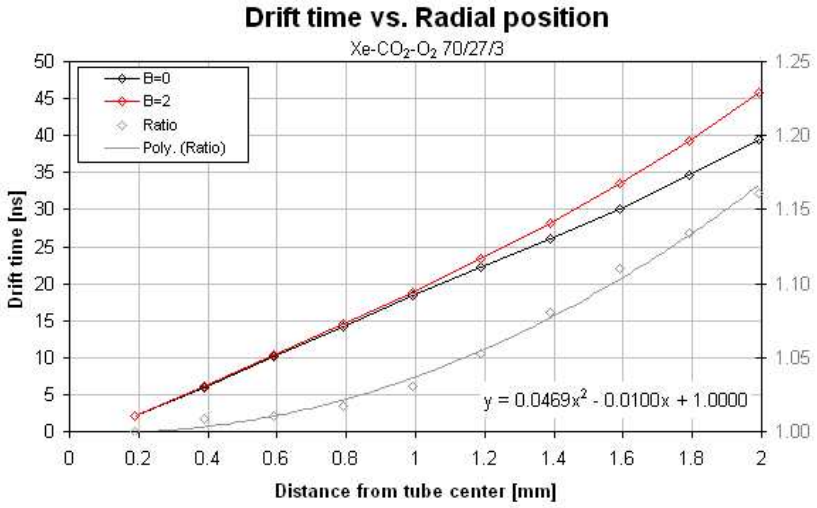

Figure 9.5: *Drift velocity as a function of the wire-to-cluster distance. At low radii the function is fixed by the requirement* $t(r_{wire}) = 0$. *Figure due to P. Cwetanski.*

encoded in the simulation hits by GEANT4. The maximal drift times are up to around 45 ns and output digits contain threshold information covering 75 ns, so it is important, for simulations as well as for the real detector, that the individual output channels are implemented with timing offsets such that the interval of possible arrival times is placed roughly in the middle of the output interval for all straws. The chosen implementation sets the timing offset for a particular channel to be equal to the time required for a particle with $\beta = 1$ to reach the instrumented end of the corresponding straw, thus giving the earliest possible time a signal could be registered, and subtracting 15 ns from this number.

For the drift times in the gas itself, the simulation is again based upon the results of Garfield studies. It is clear from figure 9.1 that the drift time, in addition to the dependence on the distance from the wire to the ionisation cluster, where the electron starts, also depends upon the external magnetic field and will contain small random contributions due to the stochastic nature of the collisions with the drift gas constituents. Figure 9.5 shows average drift times as a function of start radius for the case of no external magnetic field as well as for an external axial magnetic field of 2 T. The figure also shows the relative magnitude of the two curves, which grows from being at unity at low radii to a 15% effect at $r = 2$ mm, corresponding to a difference of about 6 ns. In the barrel, the field is to a good approximation parallel to the wires and of constant strength, $B = 2$ T, thus resembling the situation in figure 9.1.b. In the end cap, however, the field is mostly orthogonal to the wires and depending on the direction of the passing tracks, the situation will be somewhere between the two extreme cases illustrated in figure 9.6. In order to simulate this, the drift times of a given cluster will be approximated as being in an axial field with a strength given by $2\ \mathrm{T} \times \sin\theta$, where θ is the angle between the field direction and the direction from the cluster towards the wire.[1]

While the curves in figure 9.5 represent the average drift times, they do not account for drift

[1] A small gain in realism might be achievable by using numbers based directly on Garfield simulations of the specific cluster position and field orientation, instead of simply basing them on the numbers for the setup with an axial field. It might also be interesting to drop the assumption of a perfect solenoidal field of 2 T. The present approach is, however, computationally efficient and is already quite realistic.

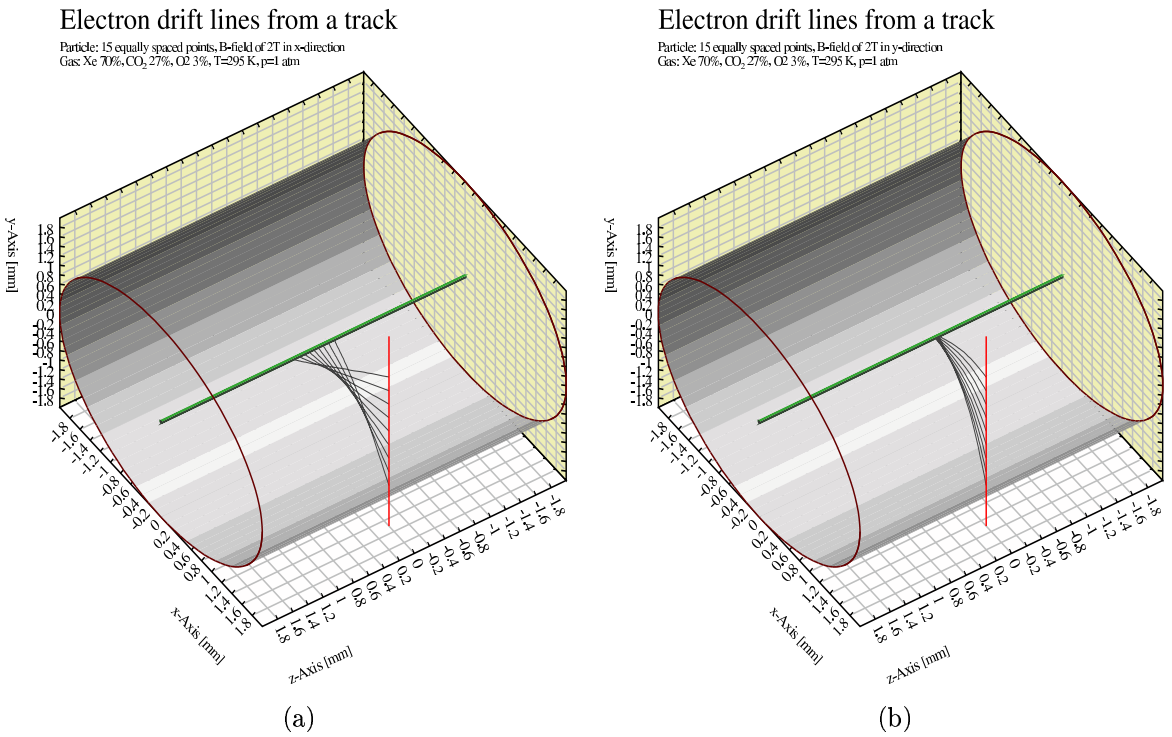

Figure 9.6: *Average electron drift paths in a straw with a 2 T magnetic field orthogonal to the wire and either orthogonal (a) or parallel (b) to the direction of the passing particle. Source: [Cwe06].*

time fluctuations. As these are the result of a very large number of collisions, they should be well described by a Gaussian distribution around the mean, which is also how they are implemented in the code. Again, detailed Garfield simulations are used to find the magnitude of the fluctuations, and figure 9.7.a shows the resulting spread in drift times – both relative to the average total drift time and its absolute value in nanoseconds for the case of no external magnetic field. Electrons starting their drift close to the wire will arrive at the wire before significant fluctuations are build up, but at larger radii they can be substantial. Figure 9.7.b shows the resulting drift times for a number of test electrons.

Having reached the wire, the signal will still have to propagate along the straw before being registered at the electronics. This happens with a speed of $0.75c$, and is complicated by the fact that half of the signal will start the propagation in the direction away from the electronics before it is reflected at the far end of the straw. For long straws such as those in the barrel layers 10-73 the difference between the arrival times of the two parts of the signal can be as large as 6 ns for particles passing the straw near the electronics. If the straw is crossed at the other end, the two halves of the signal will essentially reach the electronics simultaneously, but will take 3 ns to do so. These important effects are straight-forward to implement in the digitisation code, and might be worth to consider in advanced track reconstruction algorithms also (cf. section 12).

9.2 Front-End Electronics and Digitisation Scheme

The signal arriving at the instrumented end of a drift tube is amplified and shaped before being discriminated against the two threshold levels. These steps are all applied in a so-called

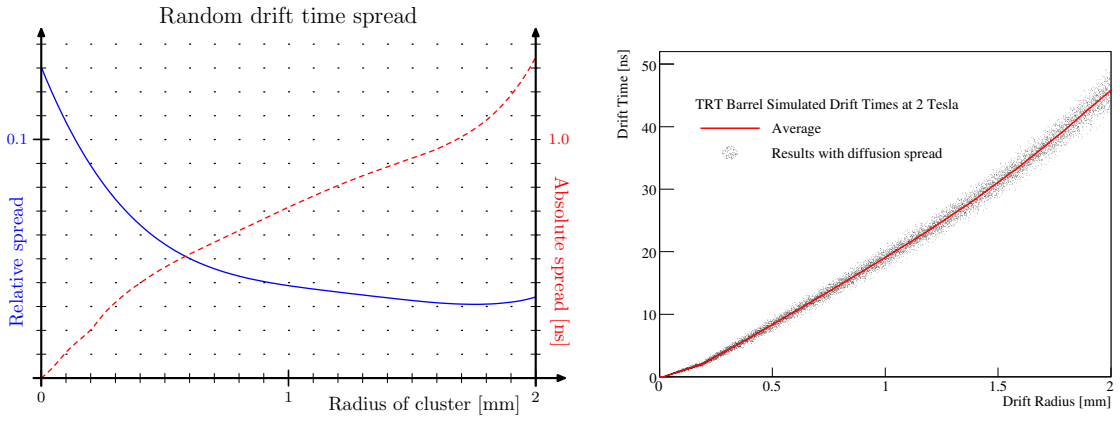

Figure 9.7: *Figure (a) shows the relative (solid blue) and absolute (dashed red) spread in drift times due to the random nature of the collisions during the electron drift (for the case of no external field). Figure (b) shows the resulting drift times for a number of drift electrons as a function of initial distance from the straw centre (with an external axial magnetic field of 2 T).*

Amplifier/Shaper/Discriminator with Baseline Restoration (ASDBLR) chip [B+96], while the links to timing, trigger and read out systems are handled by a separate *Drift Time Measuring Read Out Chip* (DTMROC) [A+01].[1] The actual amplification will only modify the signal with a constant factor which, as per the previous discussion about signal normalisation, should not be implemented in the digitisation code. Since the buffering and transferring of binary data can to a very good approximation be considered flawless, the digitisation code thus only needs to implement signal shaping and discrimination in order to simulate the processing in the front-end electronics.

Signal shaping essentially constitutes a convolution with a shaping function, which is slightly different for the low and high thresholds. The shaping functions, based on charge injection measurements by A. Romaniouk, are shown in figure 9.8. Included in the plot is a separate shaping function, which is used for the shaping of noise pulses as will be discussed in section 11. Note that the shaping functions "overshoots" and give small and negative contributions after the initial peak. This is an effect of the baseline restoration, which is applied in the electronics and which is necessary to ensure that the signal returns quickly to the baseline, instead of having tails at high times due to the slow drift of the positive ions. This is crucial in order to keep the drift tubes operational in the high occupancy environment of the LHC. The observation that the shaping functions for LT and noise pulses are not exactly similar can be taken as an indication of a small scale dependency of the shaping function on the strength of the incoming signal (the difference between e.g. the shaping of the LT and the HT signals is, on the other hand, a deliberate effect, implemented in the electronics). The shaping functions are all normalised such that their heights rather than their areas, are at unity, consistent with the normalisation throughout the digitisation chain. The convolution between the signal and the shaping functions are performed numerically using a binning of 3.125 ns/4 $\approx$ 0.78 ns. The artificial effects introduced by this binning should be minimal as 50 μm/ns $\times$ 0.78 ns/$\sqrt{12} \approx$ 11 μm.

Ideally, the discrimination of a given time dependent signal, $S(t)$, against a threshold, S_T,

[1] A concise and thorough description of the front-end electronics is given in [ATL97b, chap. 12.8].

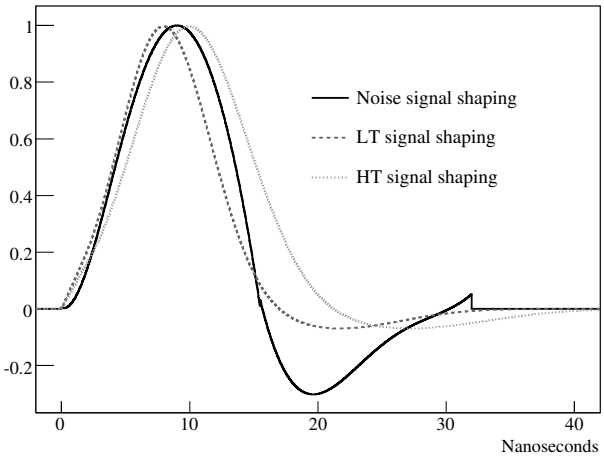

Figure 9.8: *Signal shaping functions used to model the shaping in the ASDBLR for signal passed onto both the low and high threshold discriminator as well as for noise in the low threshold discriminator.*

Bit Pattern	HT	LT Bit Pattern	Typical cause
0 00000000 0 00001111 0 11111000	off	000000000000111111111000	Passing π
0 00000000 1 11111110 0 00000000	on	000000001111111000000000	Passing e
0 00111110 1 00000011 1 11111110	on	001111100000001111111110	π followed by e

Table 9.1: *Three examples of TRT digits and their typical physics causes.*

should yield a time dependent boolean answer given by $\Theta\left(S(t) - S_T\right)$, where Θ is the step function. However, the actual discriminating electronics need finite amounts of time to change their state, when the signal cross the threshold. This is accounted for in the digitisation code by including a small settling time (0.78 ns) – the additional time the discriminator will stay on after the signal drops below threshold – and a period of dead time (6.25 ns) after this settling, in which the threshold will never register as exceeded, regardless of the height of the signal.

Effects of noise and channel-to-channel variations are implemented by adding a noise component to $S(t)$ and letting S_T vary from channel to channel with, to a lesser extend, event-to-event fluctuations. The modelling of this will be discussed further in section 11. Channel-to-channel variations are inevitable since the threshold setting can only be tuned for each of the ASDBLR chips, which are connected to groups of 8 channels.

The final step in the digitisation code is the actual encoding of the discriminator output into 27 bit long digits, containing discrimination information covering 75 ns, or 3 bunch crossings. The time structure of the low threshold is vital, while the useful information regarding the high threshold is mainly whether it was exceeded at all during the passage of a particular particle. Therefore, 24 and 3 bits are used for the low and high threshold discrimination results, respectively, as illustrated in figure 9.9. This means that the low threshold is binned into 3.125 ns bins, and with bin edges aligned with the 0.78 ns internal binning. The resulting impact on the track measurements is 50 μm/ns $\times$ 3.125 ns/$\sqrt{12} \approx 45$ μm. Table 9.1 shows a few examples of the digits that might result from various passing particles.

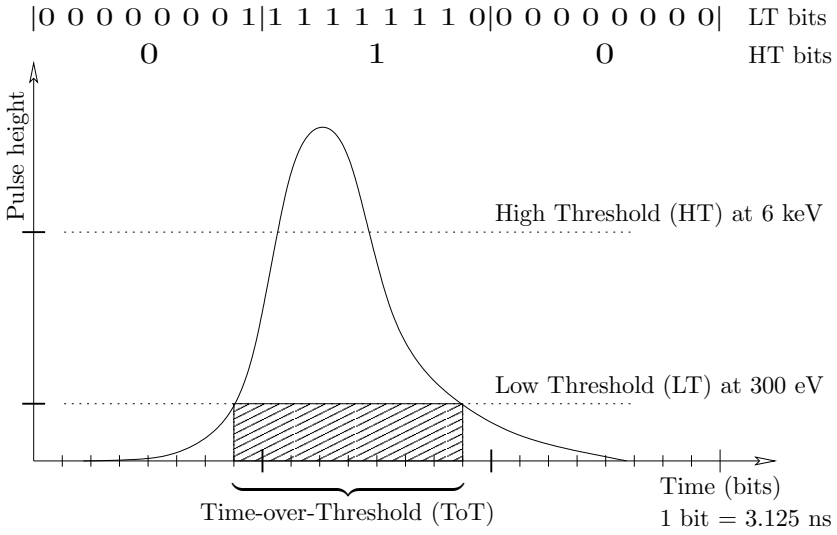

Figure 9.9: *Illustration of the scheme used for the encoding of the TRT digits. Source: [MP06].*

10 Noise in the 2004 TRT Standalone Test Beam

Even though the noise level in the TRT is known to be quite low, $\mathcal{O}(1\%)$, noise is still a relevant effect and as such needs to be both simulated along with the rest of the electronics response in the digitisation phase as well as accounted for in reconstruction algorithms. The study presented in this section provides the basis for the modelling of noise and threshold fluctuations in the digitisation code, which is the subject of section 11, by analysing data from the 2004 TRT standalone test beam [LR04].[1]

In section 10.1 it is demonstrated how genuine noise hits are reliably extracted from the test beam data and the general characteristics of those are shown in section 10.2. Finally, section 10.3, will briefly investigate the possibility of noise induced cross-talk between channels, as well as touch upon the subject of correlations between the performance of channels.[2]

10.1 2004 Test Beam Data

In July 2004, 6 functional and tested TRT barrel modules, corresponding to 2 out of 32 phi sectors, were placed in the H8 test beam in the SPS North Area at CERN Prévessin. Using various targets and magnet field selectors, beams of either pions, electrons or muons were delivered. This was the first test beam with the front-end electronics (including the ASDBLR chips) essentially of the final design, while the readout chain and data acquisition system were final prototypes, very similar to the ones that will be used in ATLAS.

In figure 10.1 the test beam setup is shown, with the beam entering from the left. The beam passed through various detectors in addition to the barrel modules. These additional detectors possessed abilities to do high quality particle identification, tracking and shower

[1] The contents of this section are essentially similar to sections 1, 2 and 3.4 of [KK06], and are the results of a study performed in collaboration with E. Klinkby. Thanks are in order to V. Tikhomirov for the effort of making the test beam results easily accessible by filtering, selecting, storing and documenting the data.

[2] Note that this and the following section will use the terms "channel" and "straw" interchangeably. In reality, each barrel straw is electrically separated at the middle and instrumented at both ends, thus corresponding to two different channels. However, only one end was illuminated and instrumented in the 2004 test beam.

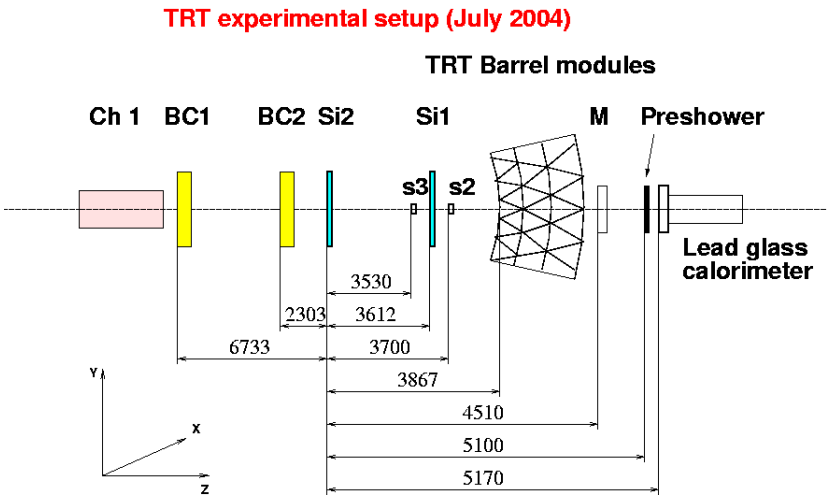

Figure 10.1: *Basic setup of the 2004 TRT standalone test beam. Indicated distances are in mm.*

Run number	Beam Energy	Electrons	Number of events
3183	80 GeV	no	90K
3240	20 GeV	yes	60K
3241	20 GeV	yes	80K

Table 10.1: *Beam types and number of events for the three runs of data taking used in this section. All of the runs were taken using low threshold values of approximately 300 eV, high threshold values around 7 keV and with the beam passing the modules roughly in the position shown by the track in figure 10.2. Note that the beams were not pure electron or pion beams in the strict sense, but contaminated to some extent – see figure 10.3.*

vetoing. Figure 10.2 shows an example of a recorded pion event. The details of the setup and its performance are described in e.g. [Tik04]. The setup thus provides events, where the identity of the passing particle is well known and with good external knowledge of the parameters of the passing track, thus making them suitable for detailed studies of detector performance.

For the present noise studies the type of beam should in principle not matter, but for cross-checking purposes three runs with different beam compositions and energies are selected as summarised in table 10.1 and figure 10.3.

10.1.1 Initial Selection of Noise Hits

The goal of the study in this section is to understand the rate and composition of the noise digits, i.e. the digits associated with the hits in figure 10.2, whose appearance are clearly not due to the passing track. Since the number of such noise digits over an entire run is rather high, quite luxurious cuts can be afforded, while keeping the statistics at a sufficient level.

The first and primary selection cut is to only consider digits originating from straws with wires located at a certain minimum distance away from the passing track. In addition to removing hits on the reconstructed track this cut must also ensure a high rejection of digits caused by secondary particles and "cross-talk" – i.e. wire currents induced through capacitive

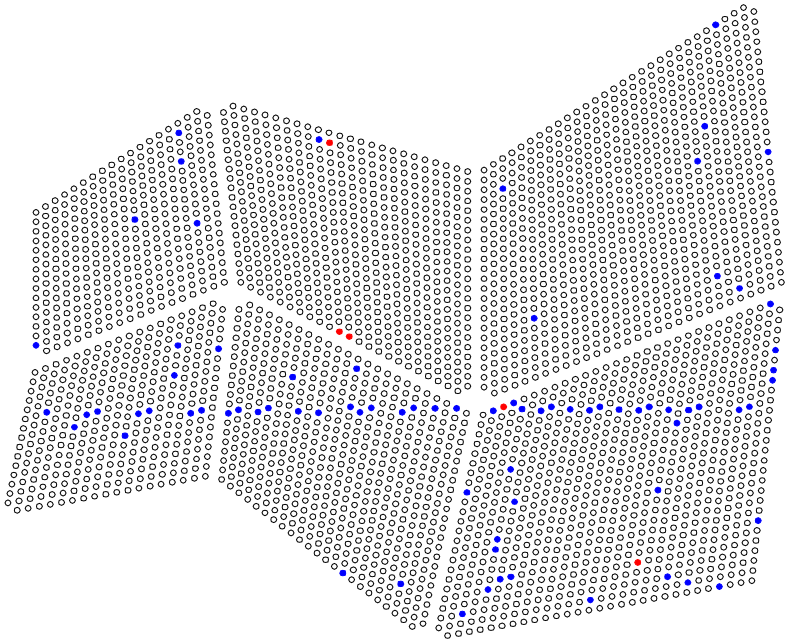

Figure 10.2: *Example pion event from run 3183. Straws with HT passed are shown in red, while those with only a passed LT are shown in blue.*

couplings to a genuine large current in a neighbouring wire. Finally, some events will have more than one beam particle, and since the typical vertical spread of the beam was around 3 cm, the cut must be much larger than this. The final value of the cut is determined by looking at the actual distribution of the distance between the wires of the digits and the reconstructed track as shown in figure 10.4. A cut of 8 cm is chosen as the tails from the track peak at zero seems to be negligible after this distance.

An additional way to ensure that the single track in the event is indeed single and well reconstructed, is to require the number of hits on the track not to deviate too much from the average number for a single track. As indicated by the greyed out regions in figure 10.5, the reconstructed track is required to have at least 33 and a maximum of 43 hits on track. In particular, the distribution in the 80 GeV pion run exhibits a small secondary peak around 48, which is probably due to multiple beam particles, and something the chosen cut avoids.

In addition, one should note that some digits have initially been disregarded by the requirement that it should be possible to identify at least one leading edge in the LT bit pattern. This simply means that bit patterns where the first LT bit is already on, e.g. `111111000000000000000000` are discarded as they do not allow one to identify the time of the leading edge for tracking, since it could have occurred before the considered 75 ns interval. The exception to the rule is where a second leading edge is present as in e.g. `111100000001111110000000`.

Finally, straws with a noise level higher than 15% (corresponding to a noise frequency of $0.15/75$ ns $= 2$ MHz) have been masked out. This is reasonable since such noisy channels will certainly be masked out in the final setup of the detector in ATLAS.

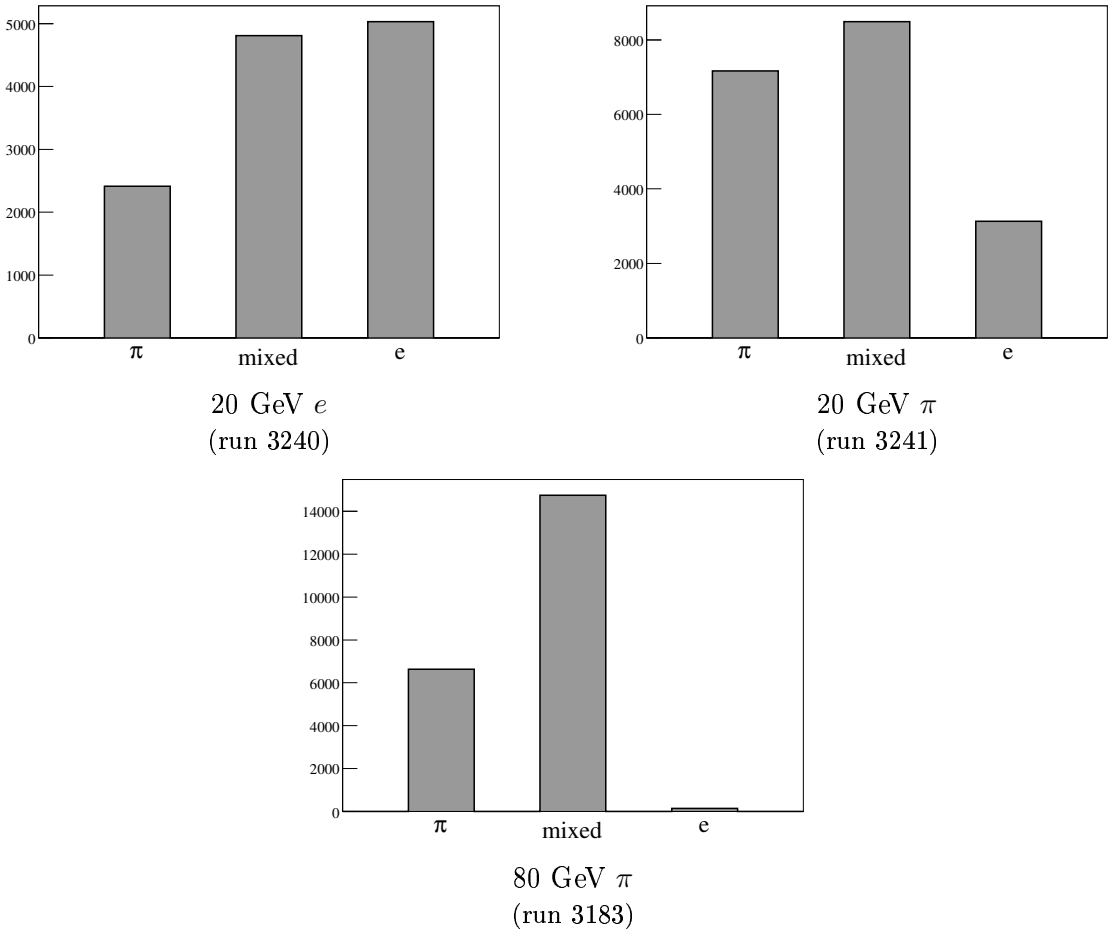

20 GeV e
(run 3240)

20 GeV π
(run 3241)

80 GeV π
(run 3183)

Figure 10.3: *Beam compositions for the three investigated runs. The electron and pion identification was performed using external detectors and the remaining impurities are thought to be at a negligible level (although this is not essential for the present study). The central column denoted* mixed *simply refers to the events that were not identified as either electrons or pions under the given cuts.*

10.1.2 Removal of Hits in Abnormal Straws

The vast majority of noise digits are thought to originate from Gaussian fluctuations in thresholds and potentials in the various channels occasionally exceeding the low threshold and producing said noise. However, looking at the data one notices a few noise digits that clearly have another origin. This is for instance digits such as,

```
0 01111111 0 11111111 0 11111111 [a]
1 00000000 1 00000000 1 00000000 [b]
0 00000001 1 00000000 0 00000000 [c]
```

Here digit [a] is from a straw where the LT is almost always exceeded, indicating a major problem in that particular straw or front-end electronics. Such a straw will probably be masked in any case in the long run. Digit [b] shows a similar structure, but here the problem is in the high threshold. Finally, digit [c] shows an otherwise innocent noise digit with a high threshold in the second 25 ns time-slice, where somewhat mysteriously the low

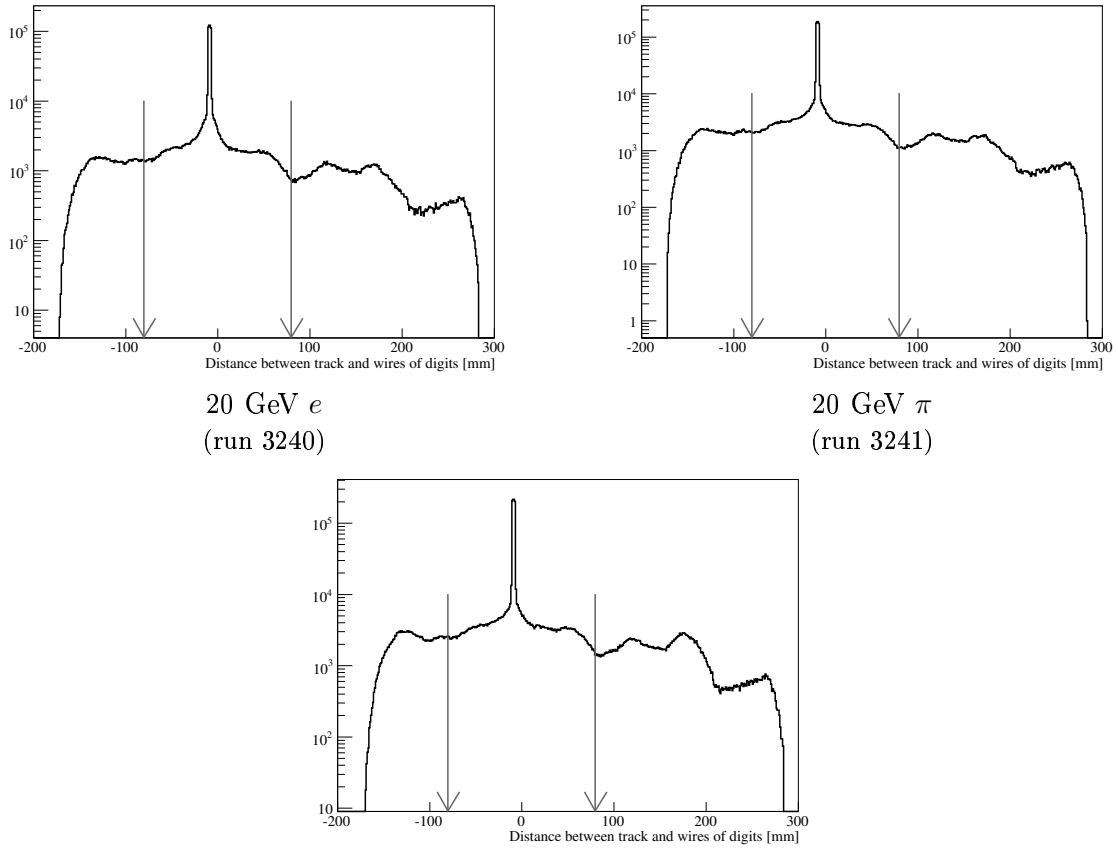

20 GeV e
(run 3240)

20 GeV π
(run 3241)

80 GeVπ
(run 3183)

Figure 10.4: *Illustration of the cut used to eliminate digits from straws too close to the track.*

threshold was never exceeded. A possible explanation is that the noise digit was in reality `0 00000011 0 00000000 0 00000000`, but that the readout driver (ROD) was somehow off by one bit – a problem known to have occurred occasionally.

While these problems are real and present in the test beam data, there is no reason to believe that they will be sufficiently similar to the ones encountered in the full ATLAS running to make a detailed study sensible. Furthermore, they represent a relatively small fraction of the total number of noise hits. For these reasons the present study will only try to understand and model the intrinsic noise of the apparatus and will not deal with these kinds of errors. In the following it will therefore be described how straws producing such "abnormal" digits can be identified and removed based on a statistical analysis of the digits they produce. To be meaningful this of course requires a minimum number of digits from each straw, and noise hits originating from straws, from which there are accumulated less than a total of 20 hits are therefore a priori ignored. This systematically throws away noise hits from the least noisy straws, but the alternative is to keep a few straws around with e.g. quite abnormal electronics. Given the low number of hits from these straws, neither choice is likely to affect the conclusions, as long as the overall noise level is not taken too literally. For simplicity the plots shown in this section will all be based on run 3240, although all runs have of course been

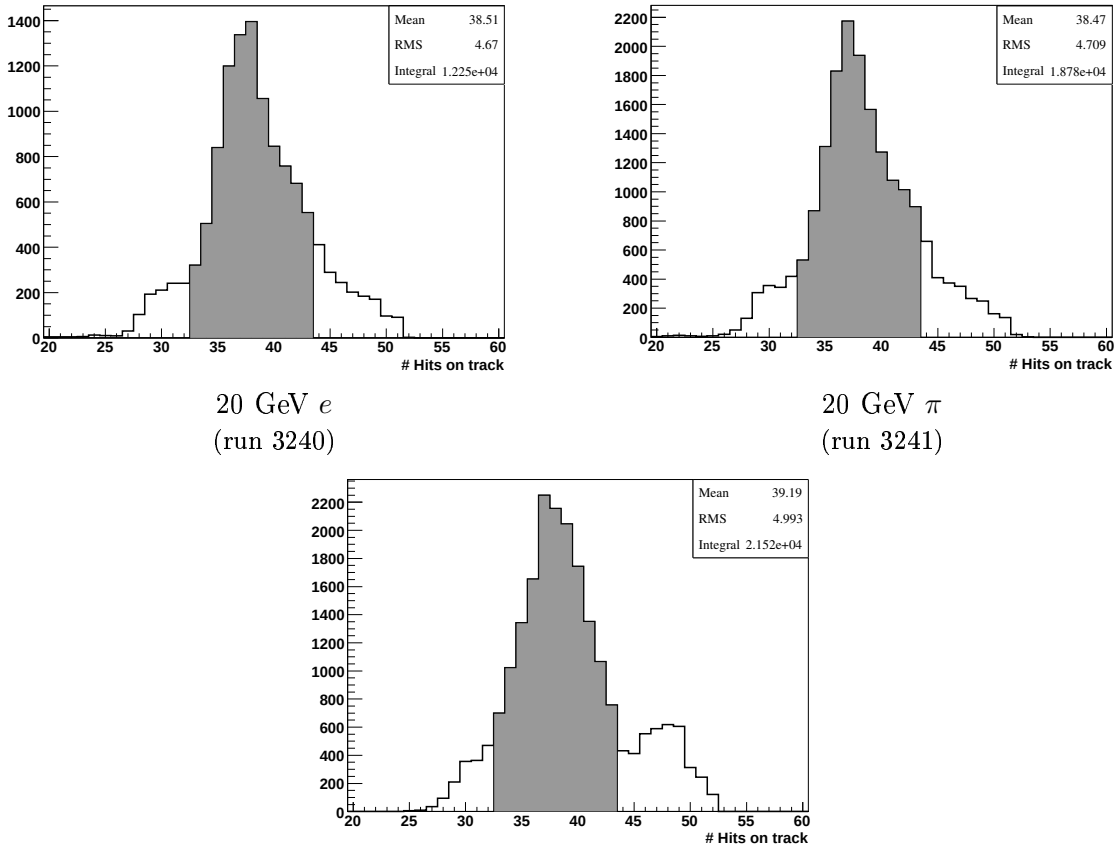

20 GeV e (run 3240)

20 GeV π (run 3241)

80 GeVπ (run 3183)

Figure 10.5: *Distribution of number of hits that are classified as being "on track". The accepted ranges are indicated in grey.*

examined.

The first quantity to consider, on a straw-to-straw basis, is the fraction of noise hits having at least one high threshold bit turned on. The high threshold is typically much higher than the low threshold (e.g. 7 keV versus 0.3 keV) and the frequency with which the low threshold is being exceeded by noise is low ($\mathcal{O}(1\%)$). This implies, given the naïve expectation of noise being generated by Gaussian fluctuations in electronics voltages and thresholds, that it is very unlikely that high thresholds will appear in noise hits, apart from those really caused by abnormal electronics and possibly stray tracks.

In figure 10.6.a the average high threshold fraction is shown for each of the remaining relevant straws. It is clear that not all straws give noise hits with negligible HT fractions, and that there is some structure caused by systematic differences between the channels (although the straw numbering is somewhat arbitrary, it is still such that straws with numbers close to each other are usually placed close to each other in the detector). In figure 10.6.b the distribution of those same fractions is shown. As expected there is a huge peak around zero with a tail towards higher values, which is likely due to secondaries (the results in the next section seems to be in agreement with such an interpretation). However, at values around 0.04

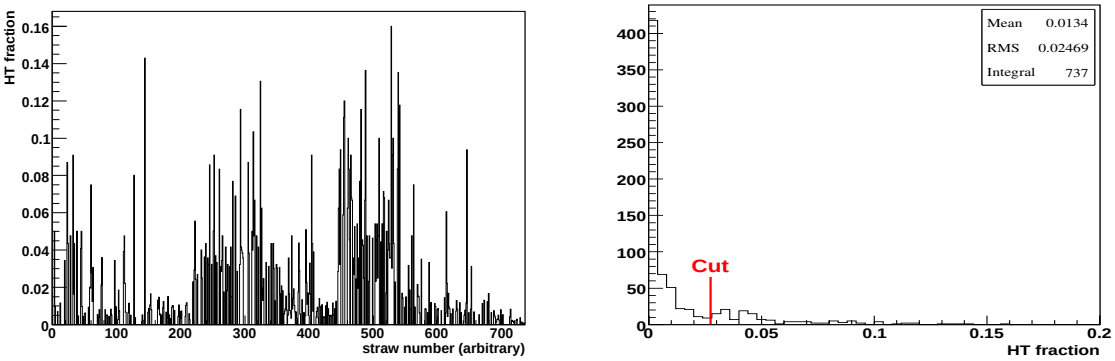

(a) (b)

Figure 10.6: *Average high threshold fractions straw by straw (a), and their distribution (b). Around 10 straws had an observed HT fraction of exactly 1.0 and are not shown in either figure.*

there is again a peak, which is hard to interpret as anything else than straws with some sort of abnormal behaviour. One might be tempted to put in an aggressive cut just around the peak at zero, but could fear that such a cut would simply remove a genuine "non-abnormal" effect. Consequently, it is chosen to place the cut just before the peak, eliminating any straw with an average HT fraction above 0.028. Not surprisingly this gets rid of almost all of straws producing digits consisting almost entirely of 1's (which really is a sign of seriously flawed electronics).

Another indication that a straw is "abnormal" is if it often produces noise digits with a relatively large number of LT bits set. In figure 10.7.a is shown the distribution of the average number of set LT bits in the selected noise digits for each straw. The grey regions show how the straws with a HT fraction higher than respectively 6% and 2.8% are distributed. It shows that there is clear correlation between straws with a relatively large number of LT bits set and straws with a large HT fraction. This supports the claim that the source of noise in these straws is somehow "abnormal", if not simply contamination due to remaining real passing particles. Based on this knowledge, channels with an average of more than 3.0 set LT bits are discarded.

Finally, the distribution of noise levels of the individual remaining straws is shown in figure 10.8. Note that the removal of straws with less than 20 hits earlier has removed entries at the lowest values.

10.2 Noise Digit Features

Using the selected sample of noise digits, with the exclusion of contributions from "abnormal" sources, the features of the noise can be extracted. Regardless of the considered feature, its distribution ought to be essentially independent of beam type and run.

The time structure of the noise can be investigated by plotting the mean occupancy of the 24 LT bits in noise digits as is done in figure 10.9. If, on average, the noise had no time structure, one would expect a flat distribution, but this is obviously not what is seen. First of all, there is a large dip in the early bins, which is easily understood as an artificial effect caused by the requirement of a leading-edge as discussed in section 10.1.1.

Apart from the dip there does seem to be some sort of plateau around 8-10% with a 8 bit

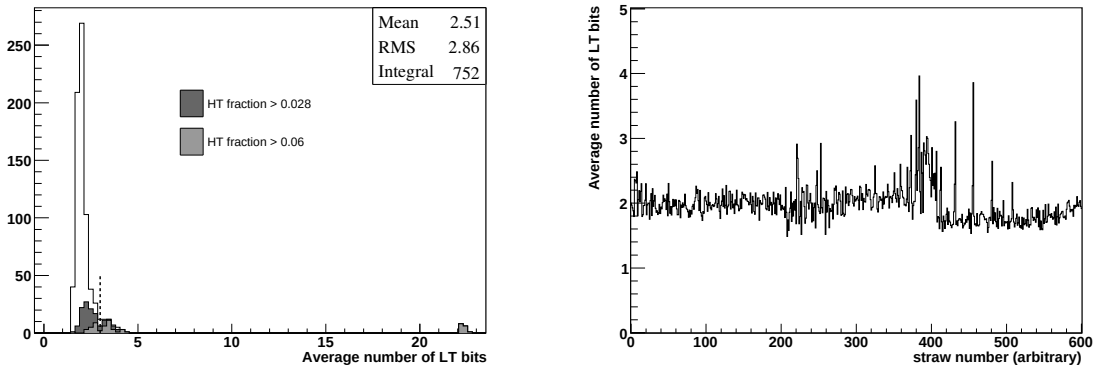

(a) (b)

Figure 10.7: *Average number of low threshold bits that are set for each straw. Note that 24 set LT bits are an impossibility due to the requirement of a leading edge as discussed in the text, but some abnormal channels with 22-23 set bits appear. In (a), the distributions of straws with a high HT content are also indicated. In (b), only the straws with HT fraction < 0.028 are shown.*

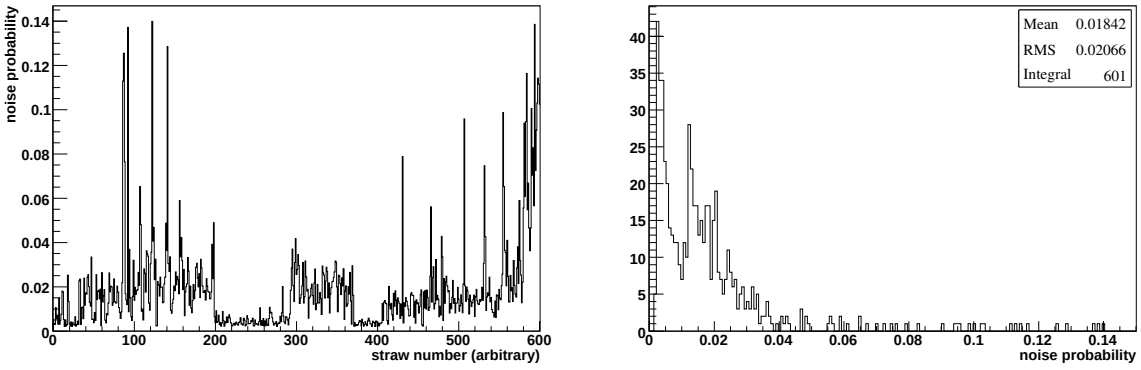

(a) (b)

Figure 10.8: *Observed noise levels versus straw number (a) and the distribution of those numbers (b).*

(25 ns) periodic structure on top of it. This is not surprising since the front-end electronics is influenced by an externally driven 25 ns clock cycle. Also shown in the figure is the contribution which is due to digits with a HT. The very low number of these means that their contribution is too small to be of significance. Secondly, they have a different and in-time shape, not too different (although a little broader) from that of beam particles. This supports the hypothesis that the noise digits with HT content are likely due to stray beam particles or secondaries.

Perhaps the most striking feature of figure 10.9 is how robust the distributions seem to be across the three runs, which underlines the consistency of the entire procedure. Furthermore, it shows the stability of the detector conditions, which means that e.g. reconstruction and digitisation can improve performance by using appropriate calibration data updated at moderate intervals.

Figure 10.10 shows another extracted quantity, namely the distribution of the number of set low threshold bits, and it is seen that, in the selected noise digits, about 99% of the entries have 4 bits or less set, and the three runs agree completely regarding the distribution in those bins.

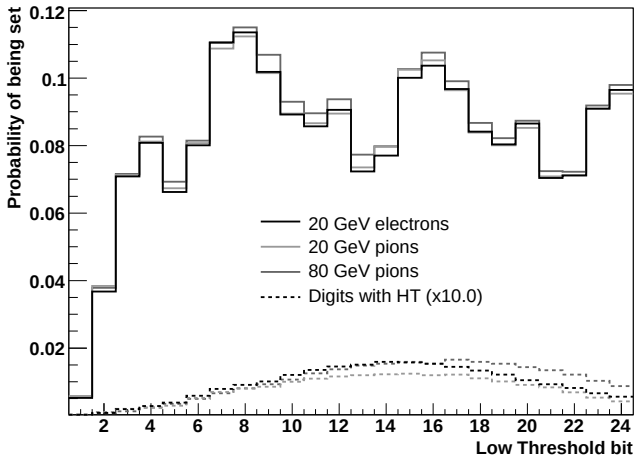

Figure 10.9: *Solid histograms show the observed mean low threshold bit occupancy in selected noise digits for the three different runs. The dashed histograms show the contribution from digits with a high threshold – scaled up by a factor of 10 for better visibility.*

Looking at the distribution for digits with a HT bit set, it again agrees with the hypothesis that there is a small and negligible contamination of stray particles and secondaries, thus explaining the present tail. If instead, for instance, the presence of set HT bits were caused by ROD off-by-one problems, one would expect the distribution of digits with HT to follow the overall distribution, apart from being shifted left by one.

The mean of the distribution corresponds to an average time above threshold around $1.96 \cdot 3.125$ ns ≈ 6 ns, which is quite low compared to the time over threshold due to pulses from traversing particles that are usually around 15-35 ns, as shown in figure 10.10.b. This is certainly a feature that could be used for reducing the impact of noise at just a small cost in tracking efficiency.

Figure 10.11.b shows the distribution of the relative fractions of number of LT bit "islands" (also known as "trains"). An island is simply a group of set LT bits surrounded by unset bits (`00000001100010000000000` thus contains two islands while `00111000000000000000000` contains only one). One notices that the relative suppression of the fraction of 2-island digits compared to 1-island digits is roughly equal to the one between 3 vs. 2 islands and 4 vs. 3 islands respectively, and that the suppression is comparable to the overall noise level of around 2%. This implies that the creation of noise fluctuations at separate times is uncorrelated to a good approximation.

10.3 Channel-to-Channel Cross-Talk and Variations

A potentially important effect is the possibility of cross-talk between straws, i.e. when a signal in one straw induces a signal in another straw through a capacitive coupling. This is potentially dangerous for track reconstruction as a signal due to a passing particle could give rise to a fake signal in a nearby straw. These issues are investigated in [TRT04]. Here it will instead be investigated whether a *noise* hit in one straw is likely to induce noise hits in nearby straws. A priori one would not expect this as the level of tracking induced cross-talk is known to be rather low and the energy deposited in a straw by a passing particle will usually be much higher than the one associated with a noise hit.

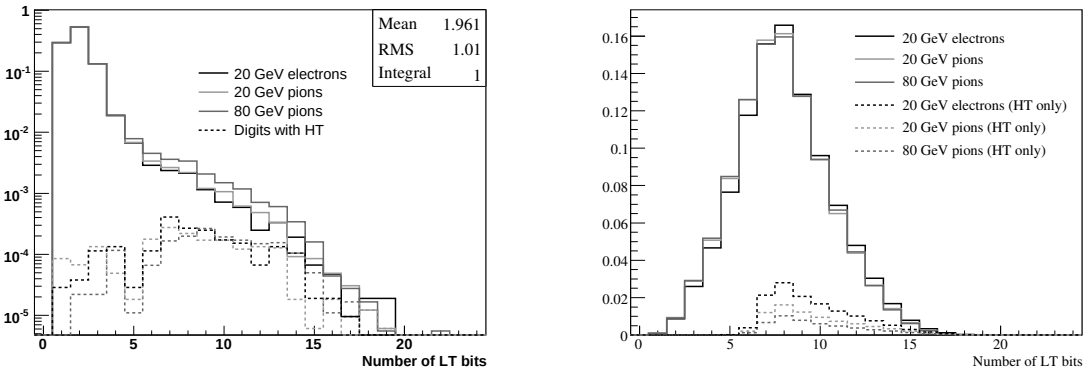

(a) (b)

Figure 10.10: *The solid histograms in (a) show the distribution of the number of low threshold bits set in the selected noise digits, with the hashed histograms showing the small contribution of digits with a HT. In (b) the same distribution is shown for on-track digits.*

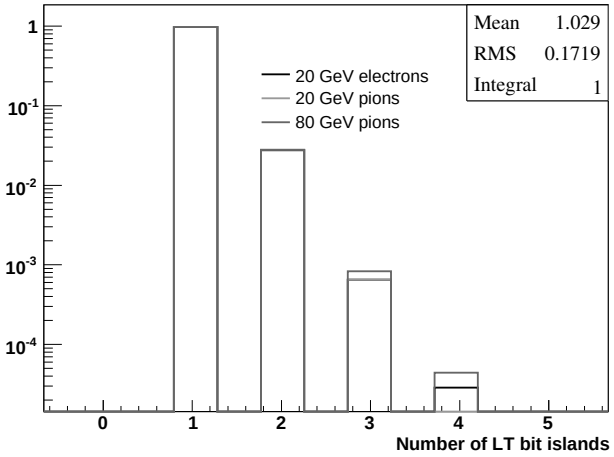

Figure 10.11: *Distributions of the number of LT bit "islands" (or "trains") in the selected noise digits.*

To estimate the magnitude of this effect from noise one can look at the number of noise hits within a given pad in a given event, of course using the fact that straws connected to a given pad are located next to each other. The distribution is shown in figure 10.12, but only for pads where all eight straws are deemed usable. Assuming no cross-talk and similar noise levels in all straws, the ratio between the content of the $'1'$ and $'0'$ bins, r_{10}, should follow from binomial distributions:

$$r_{10} = \frac{N_{\mathrm{ch}} p_{\mathrm{nl}} \left(1-p_{\mathrm{nl}}\right)^{N_{\mathrm{ch}}-1}}{\left(1-p_{\mathrm{nl}}\right)^{N_{\mathrm{ch}}}} = \frac{N_{\mathrm{ch}} p_{\mathrm{nl}}}{1-p_{\mathrm{nl}}} \quad \Rightarrow \quad p_{\mathrm{nl}} = \frac{r_{10}}{r_{10}+N_{\mathrm{ch}}} \tag{10.1}$$

Where p_{nl} is the probability for a noise hit in a given straw and $N_{\mathrm{ch}} = 8$ the number of channels per pad. Likewise the expected ratio of the $'2'$ and $'0'$ bins is found to be:

$$r_{20} = \binom{N_{\mathrm{ch}}}{2} \frac{p_{\mathrm{nl}}^2 \left(1-p_{\mathrm{nl}}\right)^{N_{\mathrm{ch}}-2}}{\left(1-p_{\mathrm{nl}}\right)^{N_{\mathrm{ch}}}} = \binom{N_{\mathrm{ch}}}{2} \frac{p_{\mathrm{nl}}^2}{\left(1-p_{\mathrm{nl}}\right)^2} \tag{10.2}$$

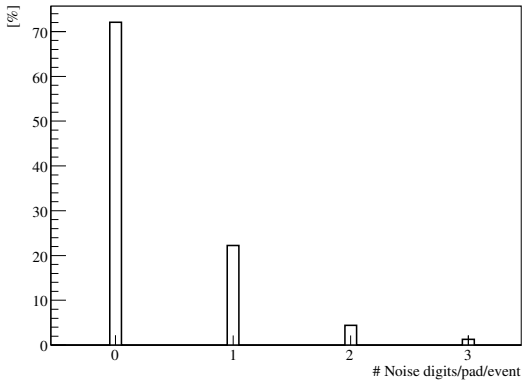

Figure 10.12: *Number of noise hits in a given pad in a given event.*

If the assumption of truly independent channel noise is valid, a direct calculation of r_{20} from the bin content in figure 10.12 should give the same result as first calculating p_{nl} via equation (10.1) and inserting the result in equation (10.2). Inserting numbers the two methods turns out to result in $r_{20}^{\mathrm{obs}} = 6.1\%$ and $r_{20}^{\mathrm{indep}} = 4.2\%$ respectively. This thus reveals a slightly enhanced tendency to get two noise hits in the same pad, which could to some extent be explained by noise induced channel-to-channel cross-talk. However, it is still true that $r_{20}^{\mathrm{obs}} \ll r_{10}^{\mathrm{obs}}$, meaning both that the effect is relatively rare and it is difficult to conclude with any certainty that it is not simply caused by e.g. a small contamination of the noise hit samples from real particle induced hits.

Another, and essential, issue is the fact that channels behave differently. This was already apparent from e.g. the distribution shown in figure 10.8.b, but in order to get a true picture of the impact of channel-to-channel variations, one also needs to investigate how the distribution of features such as noise levels are correlated between e.g. neighbouring straws (or, more realistically, between channels attached to the same front-end chip). Such studies are carried out in greater detail in [KK06], but figures 10.13 and 10.14 illustrates the importance of such variations. In figure 10.13 is shown a combined beam illumination map of the detector, using data from several runs, with the beam in different vertical positions. Malfunctioning front-end electronics will clearly impact the performance in certain parts of the detector. Figure 10.14 represents an attempt to identify dead channels, and to look for correlations with channels that are known beforehand to be unusable since they had their wire removed during pre-testing (to avoid e.g. electrical problems in those channels to affect neighbouring channels).

The larger collection of straws entirely without hits are indeed seen to correspond nicely to non-functioning ASDBLR or DTMROC chips.[1] There is a clear correlation between the channels disconnected at the test bench and the straws that hardly if ever produce noise hits. However, the main part of them *do* produce noise hits to some degree, meaning that part of the noise is produced in the front-end chip itself. This agrees with the interpretation of noise being due to capacitive coupling, with a non-zero contribution from the chips themselves, which will be discussed further in section 11.1.1.

[1] The 8 straws connected to an ASDBLR chip are typically placed as two lines of four straws separated by one line of other straws.

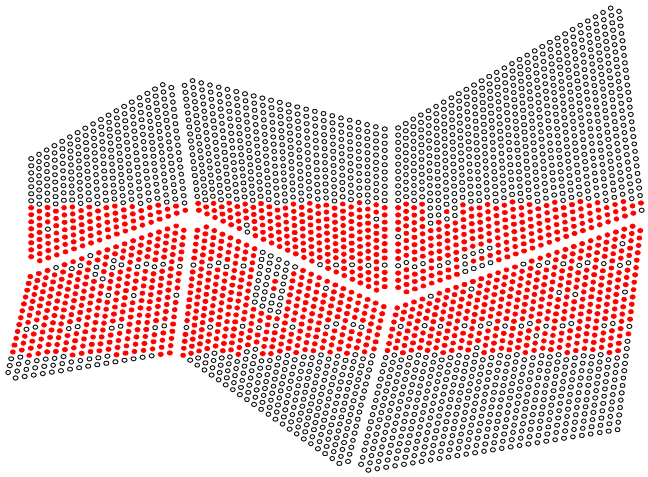

Figure 10.13: *Straws illuminated by the beam (e.g. containing on-track hits) in any of the runs considered for the channel-to-channel noise level variation studies in [KK06].*

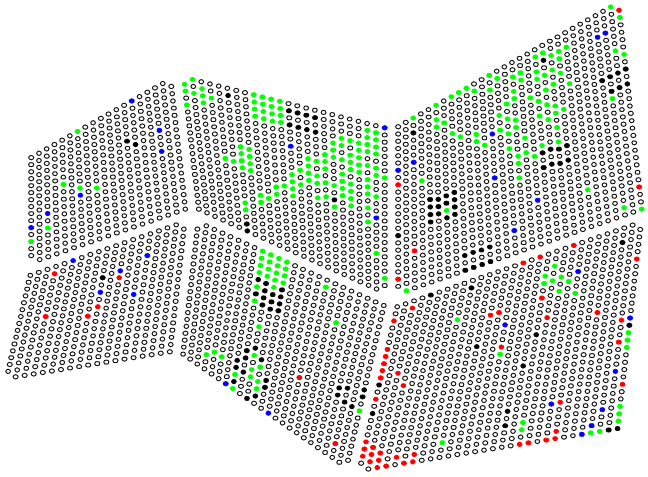

Figure 10.14: *An end view of the straws in the test beam setup. Black straws have no hits in the test beam, green straws have a noise level below* 0.01%*. Blue straws correspond to those without wires, and which had a noise level below* 0.01% *in the considered test beam runs, whereas red straws are the ones which are without wires but which nonetheless had a noise level above* 0.01%*.*

11 Modelling Noise and Channel-to-Channel Variations

Even though the noise levels and channel-to-channel variations in the TRT are known to be relatively small ($\mathcal{O}(1\%)$ and $\mathcal{O}(10\%)$ respectively), they are not entirely negligible and as such ought to be present in simulated data. Furthermore, the effects can be enhanced locally due to groupings of abnormal channels, resulting in non-negligible effects in the reconstruction.[1]

In this section, a model based on a small number of physics inspired assumptions is presented, which is able to reproduce the performance of the detector down to a channel-level granularity through variations in thresholds and background noise. In particular, it reproduces the off-beam occupancy maps after the trivial handling of channels that are not active or masked. An important point is that the noise fluctuations are assumed to be independent of whether a given straw is passed by physics particles or not. This is done both for consistency and because it is the most straight-forward way to ensure that the modelling of the noise in the hit straws is actually constrained by input from conditions data. In principle the model presented could be easily applied in any multi-channel detector with threshold discrimination.[2]

The model is presented in section 11.1. In section 11.2 it is tuned and validated against the results of the test beam study of section 10, and finally issues of implementation and performance are discussed in section 11.3.

11.1 The Model

There are several possible sources of the observed noise and non-uniformities found in the TRT, but the goal of the model presented here is not necessarily to provide a detailed understanding of these effects but rather to reproduce their combined effect in the detector. The model will thus be based on a set of simple assumptions about the workings of the detector.

As mentioned in section 10.3, one possible source of noise is the channel-to-channel cross-talk induced by very large signals from passing tracks, primarily due to the collection of TR photons associated with electrons. It has been found [TRT04] to be quite low, 3 and 0.6 extra hits for each electron and pion track, respectively, when running at a reduced low threshold of 250 eV where the effects should be enhanced. The model constructed here will thus not attempt to include such effects. It is not inconceivable, however, that one might wish to revisit this issue at a later date.

Regarding noise effects that are not track-induced, one can think of two sources: The first is so-called "white noise", thought to consist of a high number of uncorrelated and small fluctuations, which is the recipe for Gaussian behaviour. This will form the basis of the present model and is discussed in detail in section 11.1.1. Naturally one can also imagine noise caused by a low number of large fluctuations. One such source could be the "spontaneous" appearance of energetic noise clusters in the straw volumes, e.g. from radioactive decays of unstable nuclei in the active gas. Apart from the magnitude and frequency of signals due to such a source, the main feature distinguishing it from white noise is that the signals induced will always be positive, since one can not have real ionisation clusters giving negative potential fluctuations at the wire. Code for dealing with such spontaneous noise clusters has been written, but have so far not been employed since the test beam data do not seem to warrant it. It is of course

[1] The contents of this section are essentially similar to [Kit06].

[2] The author is grateful to the many persons contributing to the development of the model described in this section, through their helpful advice and interesting discussions. This includes in particular R. Van Berg, M. Dam, D. Froidevaux, P. Hansen, E. Klinkby, F. Luehring and A. Romaniouk.

possible that the environment in the full ATLAS setup will turn out to be different, so the capability has been kept.

In addition, one final related effect is being added at the same time: Allowing the actual low threshold value used for signal discrimination in a given channel to vary from event-to-event. This will be controlled by one overall parameter, r, denoting the relative level of the fluctuations. As will be discussed in section 11.2, a value of $r = 11\%$ gives the best agreement with test beam data. Specifically this means that a channel with an average low threshold of y_i, in a given event will have an effective low threshold of $y_i^{\text{eff}} = g_{i,\text{evt}} \cdot y_i$ where $g_{i,\text{evt}}$ is a random number picked independently for each channel and event from a Gaussian with unit mean and spread $\sigma = r$.

Finally, it is crucial that the model is able to reproduce channel-to-channel variations as seen in the detector and quickly adapt to drifts of the electronics parameters. This ability is discussed in section 11.1.2.

11.1.1 Basic Modelling of White Noise

In the TRT, white noise is thought to be mainly thermal in nature, originating in dissipative elements such as the amplifiers in the ASDBLR chips. These fluctuations directly results in free charges, again giving rise to potential fluctuations through capacitive coupling. This means that, on average, the amplitude of the noise in a given channel is expected to depend on the length of the attached wire:

$$\langle \mathcal{A} \rangle_l \propto c_1 l + c_2 \tag{11.1}$$

where l is the wire length, and c_1 and c_2 are constants thought to depend on the wire and chip capacitances respectively. Experience indicate that a barrel straw with $l = 70$ cm has noise of about 3000 ENC (equivalent noise charges), while a chip with no attached wire has 1200-1500 ENC ([VB05] and [B+96]). Apart from an overall constant of proportionality this implies $c_1 = (3000 - 1350)/70$ cm and $c_2 = 1350$.

In addition to the amplitude of the noise it is necessary to be able to provide simulated samples of the time-structure of white noise. The basic mechanism for doing this will be based upon a scheme due to A. Romaniouk.[1] As used here, the scheme is to generate a number of primary potential fluctuations with a magnitude given by a zero-centred Gaussian distribution of some width (amplitude). The determination of the amplitude for a given channel will be discussed in section 11.1.2. These primary fluctuations should then be distributed in time according to some chosen scheme. The next step is to perform a signal shaping on the primary fluctuations, mimicking the amplification and shaping taking place even for noise fluctuations in the ASDBLR. The shaping function used for noise is shown in figure 9.8. It is not exactly identical to the shaping used for signal, which is simply because of a slight (otherwise neglected) dependence on the magnitude of the signal that is shaped.

In figures 11.1.a-d are shown examples of the resulting noise distribution given by four different schemes: The average distance between pulses is either 1 ns or 25 ns and they are distributed with either fixed or random distances. It is clear that for average distances much less than the width of the shaping function of $\mathcal{O}(10$ ns), the fine structure is washed out by the signal shaping and there is no difference between a fixed and random distribution scheme.

[1] Implemented in FORTRAN code by F. Luehring for studies estimating the impact of noise on tracking resolution.

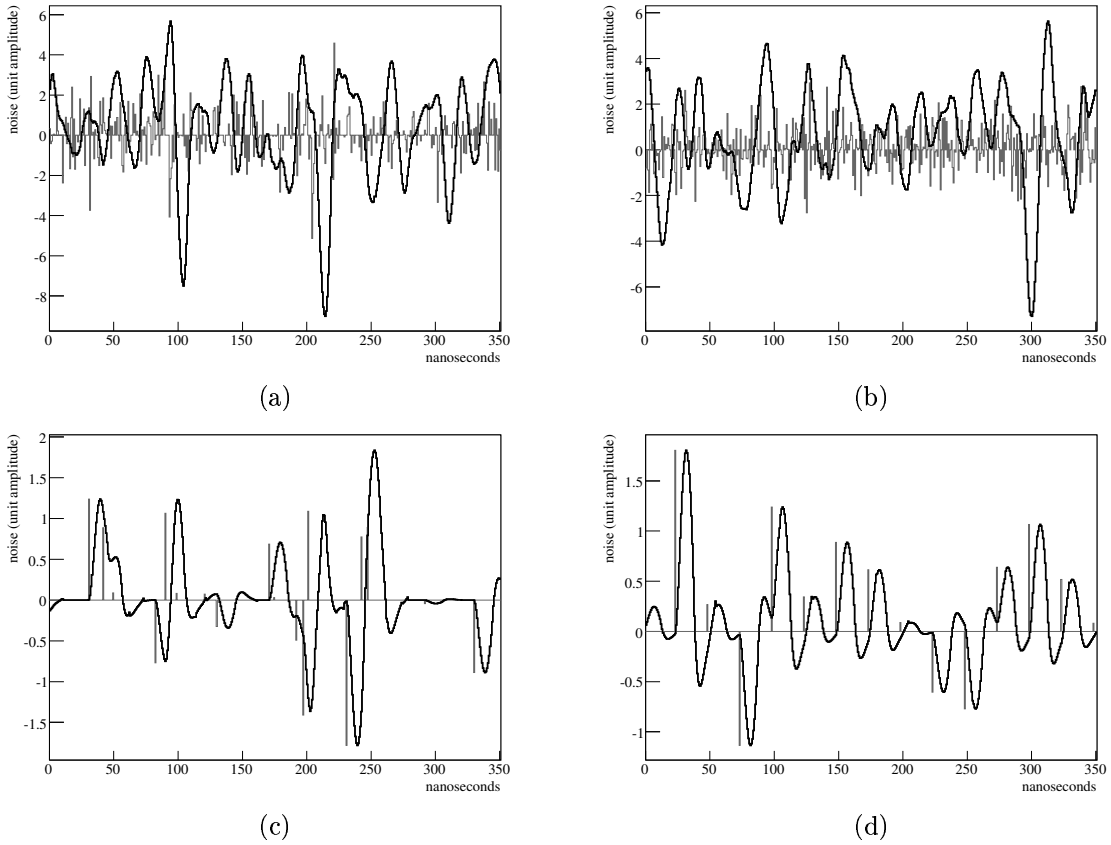

Figure 11.1: *Examples of noise signals given various schemes for distribution of unit-amplitude Gaussian noise pulses, before (grey) and after (black) shaping. One every 1 ns with random (a) and regular (b) distribution, and one every 25 ns with random (c) and regular (d) distribution.*

Furthermore, the only difference between two low values for the average distance will be an overall scale factor, which is anyway left undetermined at this point.

11.1.2 Channel-to-Channel Variations

Apart from geometrical issues of placement and length, the TRT simulation and digitisation algorithms have traditionally been treating all straws as copies of the same "ideal" straw. In the real detector, however, each straw has unique performance and settings. Some straws are outright "dead" for various reasons. Some had their wire removed due to problems found in high voltage tests and some were sealed off due to e.g. gas leaks found in gain tests. Others are effectively dead because the front-end electronics is faulty, with e.g. a blown fuse, and some are simply masked out due to intolerable behaviour like an excessive level of noise. The active straws differ from each other in several ways. Examples are: Threshold values, the magnitude of the noise fluctuations, the amplification and signal shaping in the front-end electronics, some straws might be bent, the gas densities might differ slightly, different straws are connected to different power supplies, etc. Trying to implement all of these effects in the digitisation code would be impractical and one clearly has to make some simplifying assumptions in order to keep the number of free parameters low and the requirements on

resources and conditions data reasonable. The model presented here will, in addition to a few overall free parameters, simply assume that an off-beam straw occupancy map, i.e. a map of the individual channel noise levels, is available along with a list of channels that are dead or masked. It is also important to note that the model is constructed so as to be robust enough that it can be trusted to work well over a large range of input parameters without requiring constant by-hand re-validation by experts.

As a side remark one should notice that there are straw, module and wheel misalignments which break the geometrical symmetries of the ideal setup. The simulation of those is the regime of the detector description as mentioned in section 8 and will not be discussed further here.

The goal of the model presented here is to assign each of the N active straws a unique low threshold value, LT_i, and white noise amplitude, $\mathcal{A}_i$, so as to reproduce the actual measured noise level map for a given period of running. It is clear that, with white noise as the sole source of noise, there must be a one-to-one relationship between the noise level of the i'th straw and the ratio $f_i = \mathrm{LT}_i/\mathcal{A}_i$. The discussion of how to derive this relationship is postponed until section 11.1.3. For now it will just be assumed that f_i is provided by conditions data along with the overall average value of the low threshold, $\langle \mathrm{LT} \rangle = \frac{1}{N}\sum_{i=1}^{N} \mathrm{LT}_i$, which is needed to set the scale.

This leaves $N-1$ degrees of freedom in the model which are taken to be the N values of the noise amplitudes, $\mathcal{A}_i$, apart from a common scale factor, a. With $\mathcal{A}_i = a \cdot r_i$ the distribution of r_i is thus an input to the model. As described in section 11.1.1 the average noise amplitude of a channel depends on the length of the attached wire. Consequently, the r_i's will be modelled as being distributed according to Gaussians with means, μ_i, given by formula (11.1) and a width given by $\sigma_i = w\mu_i$, where w is a free overall parameter which thus controls the relative channel-to-channel spread in noise amplitudes.

Inserting $\mathrm{LT}_i = f_i \cdot \mathcal{A}_i = f_i \cdot a \cdot r_i$ in $\langle \mathrm{LT} \rangle = \frac{1}{N}\sum_{i=1}^{N} \mathrm{LT}_i$ yields $a = \langle \mathrm{LT} \rangle / \langle f_i \cdot r_i \rangle$, and thus

$$\mathrm{LT}_i = f_i \cdot r_i \cdot \frac{\langle \mathrm{LT} \rangle}{\langle f_i \cdot r_i \rangle} \quad \text{and} \quad \mathcal{A}_i = r_i \cdot \frac{\langle \mathrm{LT} \rangle}{\langle f_i \cdot r_i \rangle} \tag{11.2}$$

Test beam experience indicates that $\langle \mathrm{LT} \rangle = \mathcal{O}\,(300\ \mathrm{eV})$ is a reasonable setting (cf. section 12), but detector condition maps containing the individual channel noise levels will obviously not be available until actual ATLAS data taking starts. Until then the simulations will of necessity have to be based on a random map of noise levels.

Based on the test beam findings in figure 10.8.b, a simple distribution of straw noise levels, x, according to $\exp[-x/0.019]$ was chosen. In particular this means that there will be no channel-to-channel correlations in noise level at neither the chip nor the module level. One could have envisioned a more ambitious distribution model including these correlations, but at the time of this writing the channel-to-chip mappings were not available at run-time in the ATHENA framework. Finally, the relative channel-to-channel spread in noise amplitudes, w is chosen to be $w = 15\%$. That this is reasonable can be seen by roughly comparing e.g. the root-mean-square of the resulting final distribution of low thresholds, shown in figure 11.2, with that expected from tests.

11.1.3 Determining Ratio between LT and Noise Amplitude from Noise Level

As described in the previous section, the model hinges on having the mapping between the noise level (i.e. off-beam occupancy) and the corresponding ratio between low threshold and

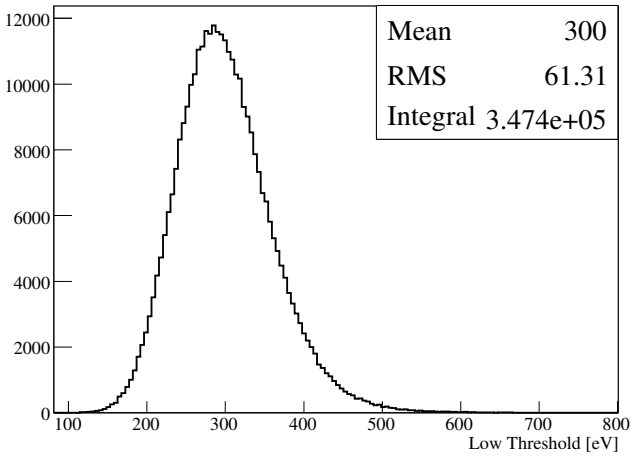

Figure 11.2: *Distribution of low thresholds given a simple distribution of straw noise levels, x, according to* $\exp[-x/0.019]$ *and a relative channel-to-channel spread in noise amplitudes of 15%.*

noise fluctuation amplitude, depicted in figure 11.3.a, readily available at initialisation time. Furthermore, it is greatly preferable to have said mapping constructed on the fly during each digitisation job.

Denoting the ratio between low threshold and noise fluctuation amplitude by f and the noise levels by x, the straight-forward procedure for finding $f(x)$ would be to select a number of values of f, and generate a large number (e.g. 10^5) of digits for each value. The fraction of those with a signal above the low threshold gives the corresponding noise level. The function $f(x)$ can subsequently be evaluated by interpolation between the obtained set of points, $\{(f, x)\}$. While this *sampling* method of determining $f(x)$ can be robust if sampling occurs at a sufficiently large number of f-values, it is also prohibitively inefficient in terms of computational time. This is because each 75 ns segment of shaped noise fluctuations is simply discriminated against a single low threshold, throwing away a lot of information in the process. It would be more efficient to reuse the *same* set of 75 ns segments of noise fluctuations for comparing against all of the f sampling values. This philosophy, taken to its limit, results in the following improved scheme, the *sorted-array* method: Generate a large number, N, of 75 ns segments of noise fluctuations with unit amplitude, and store the maximum level reached by the fluctuations within each segment into an array. Next sort the array in descending order. The i't entry in said array now contains the value of f that should be associated with the noise level $x = i/(N-1)$ (given zero-based array indices).

To test the timing and accuracy of the methods, a small standalone toy Monte Carlo study were performed with some hypothetical true $f(x)$.[1] The three lower curves in figure 11.3.b show a comparison between this "true" $f(x)$ and the same $f(x)$ reconstructed using the two aforementioned methods. The sorted-array method is seen to perform extremely well, and is orders of magnitudes faster than the perhaps more straight-forward sampling method.

Unfortunately this is not the end of the story. The procedures discussed so far in this section fail to take into account the possibility of event-to-event low threshold fluctuations in each straw. As noted in section 11.1 the low threshold in a given straw will fluctuate according

[1] The "true" $f(x)$ used as input to the toy Monte Carlo is based on a failed attempt at calculating the curves analytically, and is thought to approximate the true curve particularly well at small values of x.

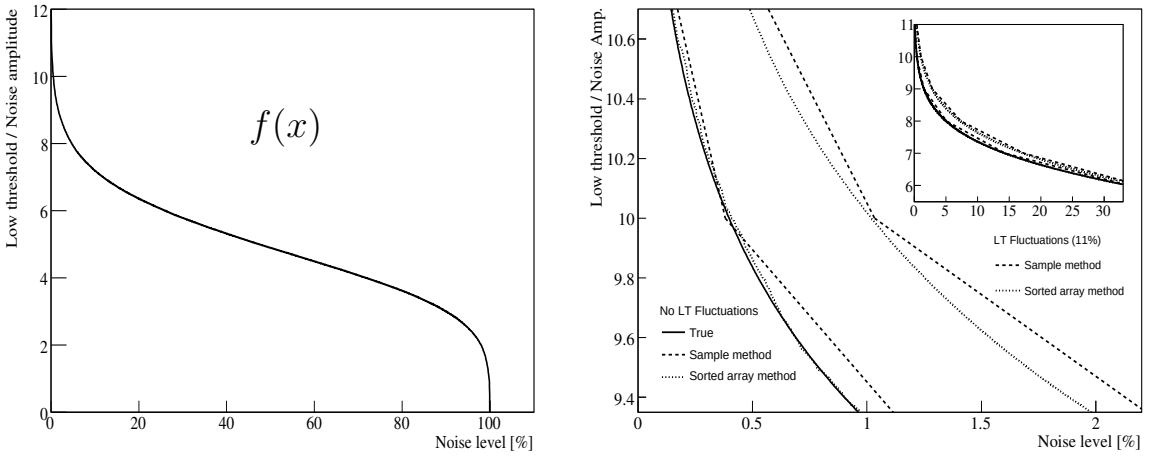

(a) (b)

Figure 11.3: *Ratio of low threshold to noise fluctuation amplitude as a function of noise level (a). In (b) is shown the results of a toy Monte Carlo study testing various schemes for reproducing said relationship, as described in the text.*

to a Gaussian from event-to-event with a relative factor of $r = 11\%$. Adding this effect to the sampling procedure is straight-forward, since the fluctuation can simply be included in the signal to threshold comparisons performed. On the other hand, obtaining the proper $f(x)$ from the sorted-array method is a little more tricky. If $f_{r=0}(x)$ is the mapping in the case of no LT fluctuations and $f_r(x)$ is the corresponding mapping one gets when the LT undergoes event-to-event fluctuations of relative size r, the following relation can be used to obtain $f_r(x)$ from $f_{r=0}(x)$:

$$f_r^{-1}(y) = \int_{-\infty}^{\infty} f_{r=0}^{-1}(y')\, G(y - y'; \sigma = ry) dy' \tag{11.3}$$

Where G is a Gaussian distribution. Thus the additional numerical work involved for the sorted array method involves inverting a look-up table twice which is negligible and a number of numerical integrations which takes a total of $\mathcal{O}(0.1\text{ s})$ after moderate tuning.

11.2 Tuning and Validation

Having thus constructed the model in the previous sections, the few free overall parameters need to be tuned to best reproduce data, in addition to reproducing the off-beam occupancy map discussed in section 11.1.2. The test beam distributions shown in this section are all the result of the study presented in section 10.

First of all one needs to decide upon a specific scheme for distribution of the pre-shaping potential fluctuations as discussed in section 11.1.1. Looking at how often each of the 24 LT bits are set in the pure noise hits found in data in figure 10.9, one sees a definite 25 ns structure on top of a flat distribution. As discussed in section 10.2, this is not surprising since the front-end electronics are driven by an external 40 MHz clock frequency. Consequently, the final composition is taken to be 1 ns and 25 ns fixed distance components, weighted by 0.4

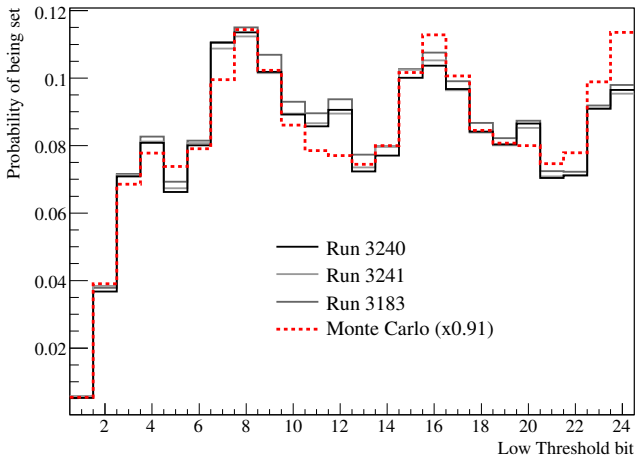

Figure 11.4: *Probability of each of the 24 low threshold bits of being set in noise digits for both data and tuned Monte Carlo simulation. To ease comparison the Monte Carlo simulation has been scaled to the data by a factor of 0.91 (a difference caused by the different means in figure 11.5). Also note that the initial dip is merely an artifact of the leading edge requirement (cf. section 10.1.1).*

and 0.6 respectively.[1] The final Monte Carlo result in figure 11.4 is quite satisfactory.

Another variable that is useful for comparison is the number of LT bits set in noise digits, i.e. the time over threshold. The distribution of this variable is shown in figure 11.5. The first thing to note is that the average time over threshold for noise appears to be around 2 bins $\times$ 3.125 ns/bin $\approx$ 6 ns, which is comparable to the width of the noise shaping function from figure 9.8 – exactly what would be expected from Gaussian fluctuations exceeding threshold. The figure was used to tune two of the existing digitisation parameters. These are the length of the enforced minimum period the threshold will register as being exceeded even when the actual period where the signal is above threshold is very short, and the additional time the discriminator will stay on after the signal drops below threshold (the settling time). The two parameters went from 6.24 ns (8 internal 0.78 ns bins) and 1.56 ns (2 internal 0.78 ns bins), respectively, to 0.78 ns for both (one internal bin). The distributions for both the original and final values are shown in the figure. Although the final distribution does differ somewhat from data, the average number of set LT bits is correct to within 10%. The difference is thought to be due to minor differences in the way discrimination is actually performed in the ASDBLR chips and the way it is implemented in code, primarily tuned for the longer track-induced signals.

Finally, it is necessary to tune the overall event-to-event level of low threshold fluctuations. It turns out that the fraction of noise digits containing two "islands" of set LT digits is a parameter which is *very* sensitive to this level, and good agreement is found by fixing $r = 11\%$ as can be seen in figure 11.6. The figure also shows a remarkable agreement in the fraction of digits having 3 islands. This can be taken as confirmation that the basic understanding and modelling of the noise is sound.

[1] Note that this gives a result whose dominating contribution comes from the 1 ns component, since more of those are generated during a given time interval. The difference in amplitude between the two components can be inferred by comparing the scales of figures 11.1.b and 11.1.d, and can be estimated to be about 2.75.

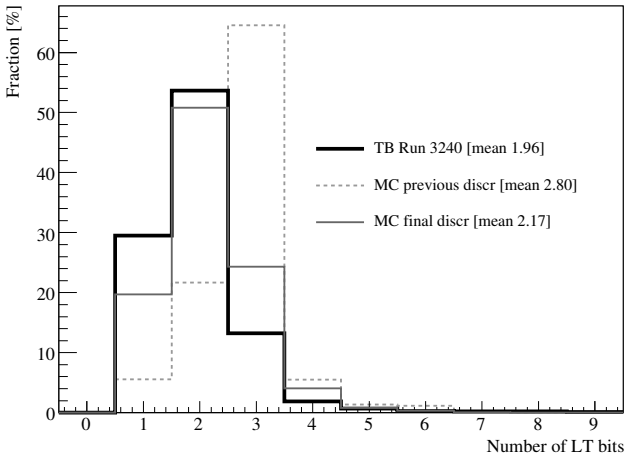

Figure 11.5: *Distribution of the number of LT bits set in noise digits for both test beam data and Monte Carlo data with two different settings of discrimination parameters.*

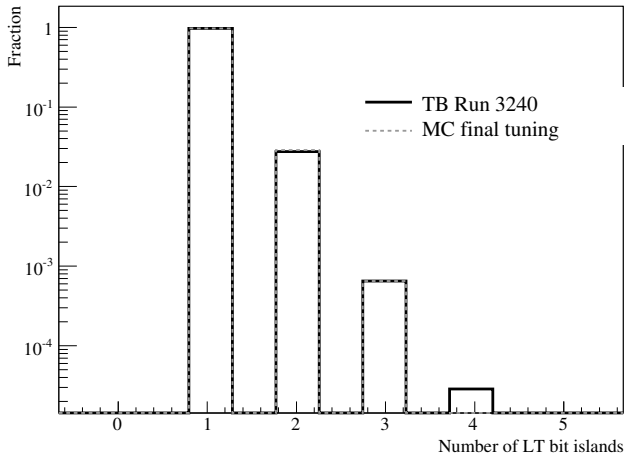

Figure 11.6: *Number of LT bit "islands" in pure noise digits for both test beam and Monte Carlo data. Islands are here intuitively defined as sets of LT bits with the LT exceeded separated by either bits where the threshold is not exceeded or the ends of the digit.*

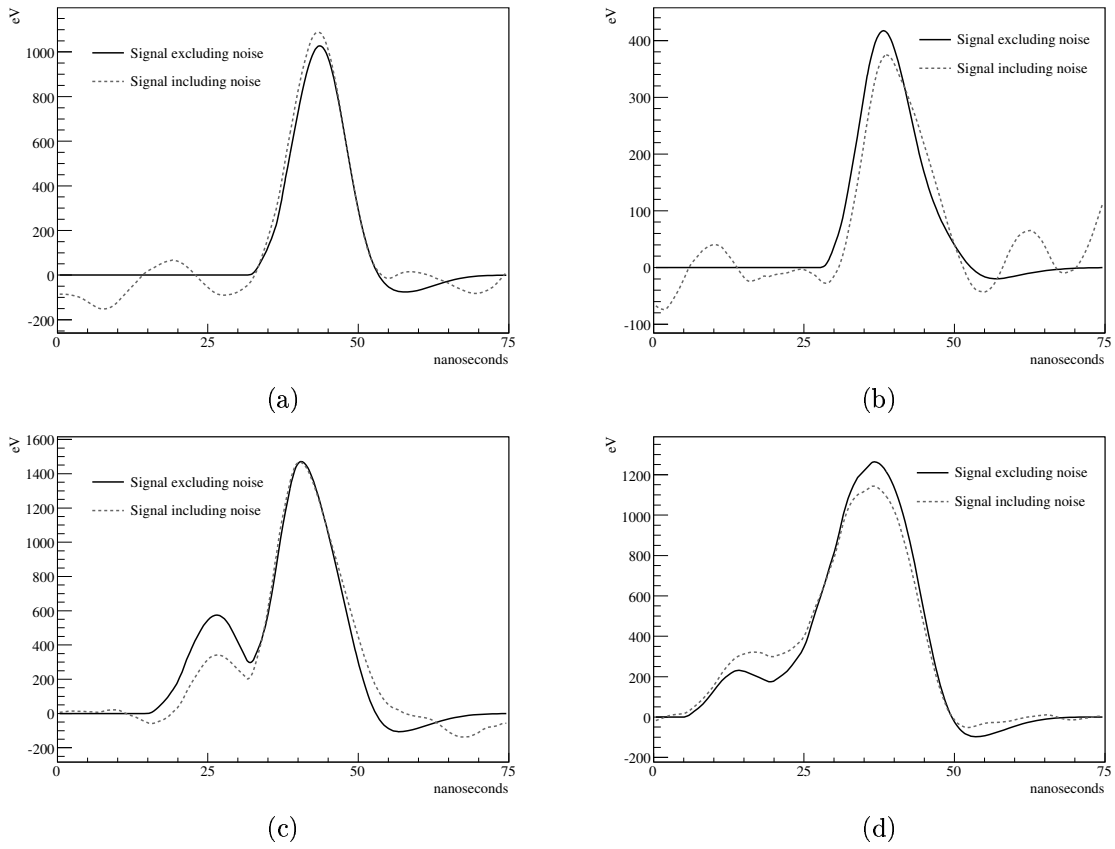

Figure 11.7: *Pre-discrimination signals due to a passing 20 GeV pion in four different straws. Solid curves show the original signals after shaping while dashed curves are after the addition of noise.*

11.2.1 Effects on Track Residuals

Apart from the addition of noise hits in unhit straws, another important effect is the influence of noise and channel-to-channel variations on the leading edge of the signal due to a genuine particle passing a straw. The effect due to noise (but not varying LT) is illustrated in figure 11.7.a-d, where the post-shaping signals found right before threshold discrimination in four straws passed by a single 20 GeV pion, are shown with and without noise fluctuations. Figures 11.7.a and 11.7.b show how the leading edge can be moved slightly forward and backwards – the latter only possible since the model allows a negative noise contribution. Figures 11.7.c and 11.7.d show a few cases where a more drastic shift could be observed, depending on the specific low threshold in the straws.

To quantify the effect on the leading edge, and thus on the tracking resolution for each straw, a simplified setup was made where high numbers of fake straw crossings were fed to the digitisation.[1] The procedure is discussed in more detail in section 12. After constructing appropriate r-t relationships, returning the most probable (as opposed to the mean) r-value

[1]Note that the version of the digitisation code used for the plots in this section did not include all currently implemented effects: Missing are for instance details such as the random spread in the drift electron arrival time and the gas gain factor. Since it is the relative impact of noise that is investigated here, this should not affect the present results. Tuning of the most current version and comparison with test beam data is the subject of section 12.

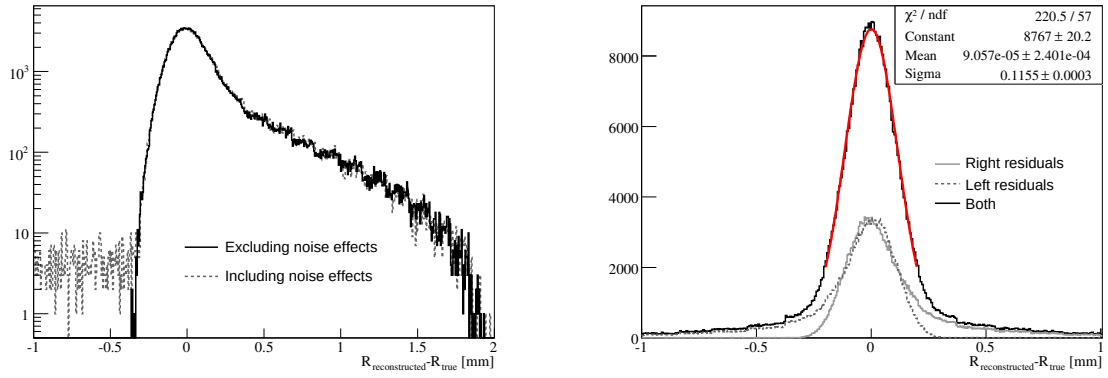

(a) (b)

Figure 11.8: *Residual distributions of reconstructed track-wire distances with and without noise effects (a). A residual appropriate for actual tracks, passing both right and left of wires, are shown in (b) along with the result of a Gaussian fit to the range* ± 200 μ*m.*

given each t-bin, the residual distribution of the track-wire distance can be found and is shown in figure 11.8.a. The only clearly visible effect of the addition of noise is the appearance of a low plateau to the left of the peak. This is easily understood to represent the cases where a large noise pulse exceeds the threshold before any signal from the track could possibly have reached the wire. Clearly this effect is sensitive to the time offset of the front-end electronics, as an offset pushing all signals to higher times will leave more room for noise fluctuations to disturb the leading edge. The version of the digitisation code used to produce these results had a low value of said offset, so the noise effect is underestimated somewhat (i.e. it was 0 ns for the plots here, rather than 15 ns used in the latest versions).

The one-sided tail in 11.8.a means that there is a bias towards overestimating the track-wire distance. In a real track fit this bias will not always pull in the same direction, but will have opposite sign depending on whether the track passed to the right or the left of the wire. To get a realistic residual distribution as experienced by a track, one should add the distribution shown in 11.8.a (corresponding to "right" passings) to itself after a sign change (corresponding to "left" passings). The resulting distribution is shown in figure 11.8.b. Due to the non-Gaussian shape of said distribution, it is ambiguous how the resolution should be defined, but here it will be taken as the width of a Gaussian fitted to the central region. Three different central intervals will be used, ± 150 μm, ± 200 μm and ± 250 μm. Figure 11.8.b includes the result of a Gaussian fit to the region given by ± 200 μm, and the results from similar fits to the other regions, and for residuals with and without effects from noise and straw-to-straw variations, are shown in table 11.1. The effect due to noise and straw-to-straw variations can thus be found by taken the quadratic difference between the two numbers, and the table shows this to be around $25 - 30$ μm. It should of course be stressed that this effect is only the *average* expected effect. Correlations in the distribution of noisy straws could lead to much higher effects for tracks in certain regions of ϕ and η space. Furthermore, the effect of pure noise hits on track pattern recognition is not considered.

Finally, there is naturally also an effect on the low threshold efficiency when the low threshold is allowed to vary from straw-to-straw as well as from event-to-event. Such efficiencies, including all implemented effects in the digitisation, are investigated in section 12.

	Gaussian width		
Fit Range	Basic model	With noise effects	Quadratic Difference
± 150 μm	108.7 ± 0.2 μm	112.6 ± 0.4 μm	29.3 ± 2.3 μm
± 200 μm	111.8 ± 0.1 μm	115.5 ± 0.3 μm	29.0 ± 1.6 μm
± 250 μm	115.0 ± 0.0 μm	117.9 ± 0.2 μm	26.0 ± 0.9 μm

Table 11.1: *Results of Gaussian fits to residuals such as the one shown in figure 11.8.b, with and without noise and straw-to-straw variation effects. Last column shows the derived quadratic difference attributed to said effects.*

11.3 Implementation

The model as described so far has been implemented in ATLAS offline software in the publicly available TRT Digitisation `C++` package, and has been used by default for productions since November 2005. For the actual details of the implementation, the reader is referred to the source code itself, while only the main concepts pertaining to the present model will be described here.

The first thing done at initialisation is to construct a list of all available channels in the detector, and associate them with individual noise levels, thresholds and noise amplitudes. The list is constructed through information made available via the GeoModel framework, and conditions data are then either looked up or simulated as mentioned in section 11.1.2, i.e. either marking a channel as dead or assigning it both a noise level and a value of the relative noise amplitude, r_i as described in section 11.1.2. The latter two are then transformed into thresholds and absolute noise amplitudes as also described in that section.

In order to avoid unreasonable execution times several optimisations take place using caching techniques. First of all, instead of generating the computationally expensive post-shaping fluctuations of the potential discussed in section 11.1.1 on demand, two very long ($\mathcal{O}(1000 \cdot 75$ ns)) time intervals of shaped fluctuations are generated - one for the 1 ns and 25 ns separated noise respectively. Whenever a noise signal is needed a random 75 ns segment is taken from each of the cached sequences, added with appropriate weights and multiplied by the actual amplitude of the straw in question. During the initialisation phase the sequences are regenerated several times in order to avoid degenerated statistics, whereas the same sequences are used repeatedly for adding noise to all straws with passing particles ("hit" straws).

Regarding noise in "unhit" straws, it is first determined which of those will actually produce a noise digit by comparing their noise level against a random number. The actual noise digit produced is not generated on the fly but is instead, for performance reasons, taken from a large ($\mathcal{O}(5000)$) pool of noise digits generated at initialisation time. This approach of course assumes that noise digits produced in straws with different values of $\mathrm{LT}_i/\mathcal{A}_i$ are "sufficiently" similar.

As an illustration of the implementation, the same event with a Higgs boson, simulated at low luminosity, is shown without and with the new noise and straw-to-straw variation effects in figures 11.9 and 11.10, respectively.

11.3.1 Performance

It is important that the addition of noise and straw-to-straw variations does not impose unreasonable demands upon especially event processing time and output file size. Table 11.2

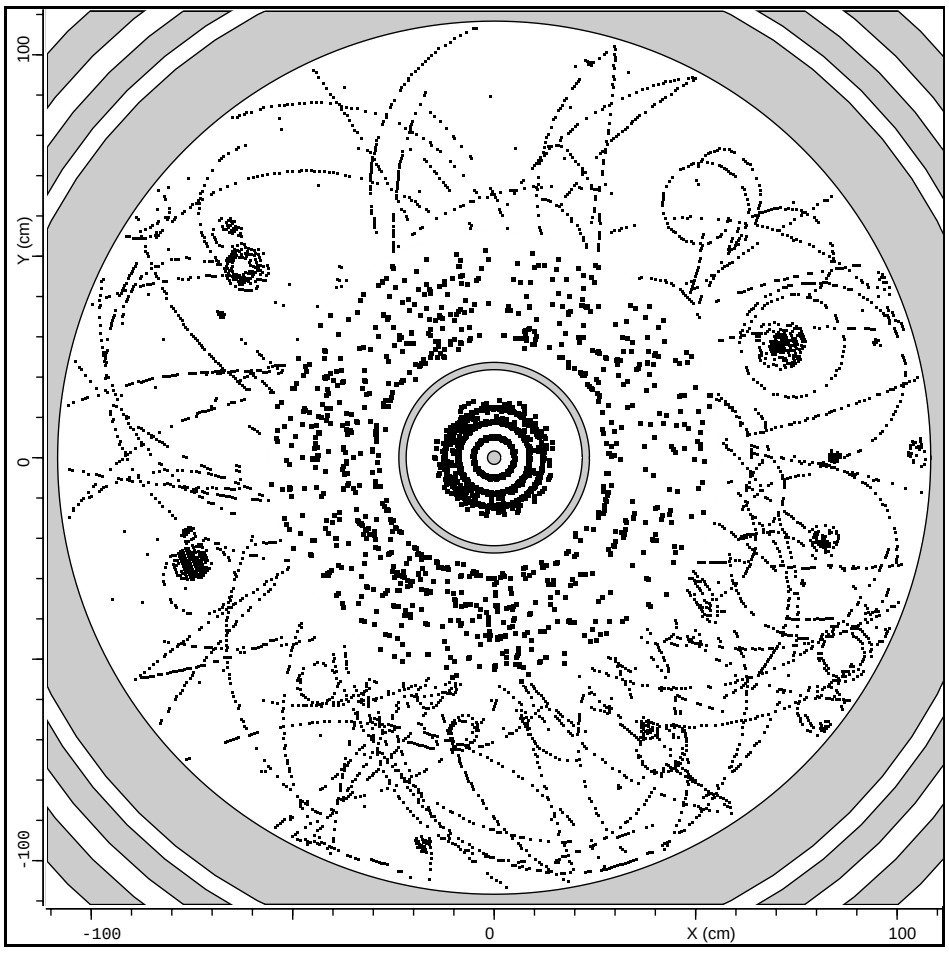

Figure 11.9: *Inner Tracker hits in simulated $H \rightarrow \gamma\gamma$ event at low luminosity with a uniform TRT detector without noise.*

shows the additional resource requirements for three different scenarios: A pathological single track event, a physics event at low luminosity (3.3 average pp interactions) and a physics event at high luminosity (25 average pp interactions). The only case where the performance is significantly compromised is the single-track case, which is easily understood since the number of pure noise hits in the TRT will, in such a scenario, be much larger than the number of track induced hits. For dedicated single track studies, in particular those involving only the inner tracker, one might consider turning off the generation of pure noise digits. As could be expected, this eliminates the performance overhead entirely as is seen in the last column. It is essential to note that table 11.2 shows the added event time for just the digitisation job. A simulation job is typically at least an order of magnitude slower for the same number of events, so the increase in the total required event processing time in large-scale productions will be negligible.

In addition to the extra resources required per event, the noise also adds about 10 s to the job initialisation time, and the job memory consumption is increased by about 10 MB (which is inevitable when assigning unique properties to $3.5 \cdot 10^5$ channels).

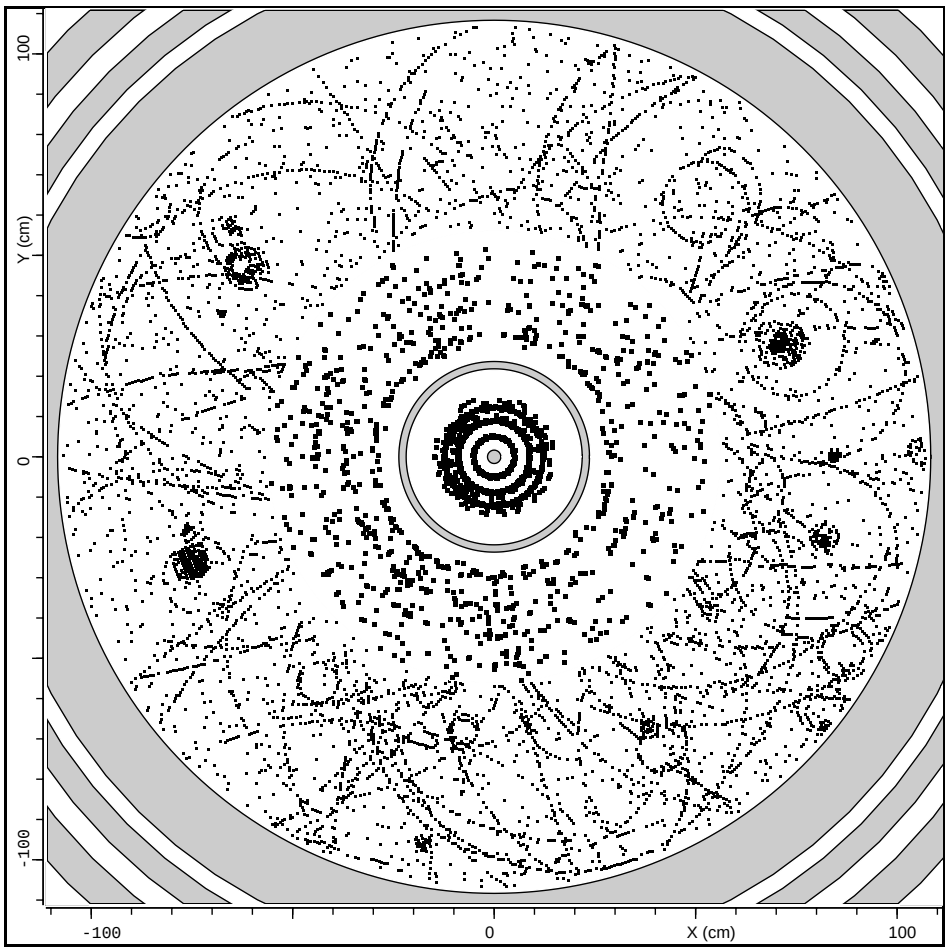

Figure 11.10: *Inner Tracker hits in simulated $H \rightarrow \gamma\gamma$ event at low luminosity with a non-uniform noisy TRT detector.*

	Physics Events		Single Tracks	
	Low lum.	High lum.	Full noise	Hit straws only
Added time/event	1 s (6%)	2 s (3.6%)	1 s (14%)	0 s (0%)
Added file size/event	73 KB (2%)	40 KB (0.6%)	89 KB (5%)	0 KB (0%)

Table 11.2: *Additional resources required for the noise and straw-to-straw variation for a digitisation job. Note that when the full chain of GEANT 4 simulation and digitisation is run the relative processing requirements goes down with an order of magnitude.*

12 Performance and $r - t$ Relationships

In order to tune and validate the performance of the TRT digitisation code, residuals of individual hits with respect to reconstructed tracks, observed in the 2004 test beam, are compared to the output of simulations in section 12.2.

First, however, it is investigated in section 12.1 how track fits should ideally use the information encoded in the output digits, in face of the large non-Gaussian effects in the leading edge measurements, mainly due to the limited number of primary ionisation clusters created by charged particles in the TRT straws.

12.1 Investigations of Optimal $r - t$ Relationships and Outlier Removal

Performance testing of the digitisation code is inherently complicated by the fact that it is executed after a GEANT4 simulation and before pattern recognition algorithms in the reconstruction code. It is thus difficult to isolate the "intrinsic" performance of the digitisation code from artifacts in these two other steps. Therefore, the results presented here will be based on a number of controlled digitisation-only simulations, as illustrated in figure 12.1: Straw crossings corresponding to a particle with a given gamma factor and with a distance at closest approach, r_{true}, distributed uniformly in the interval between ± 2 mm, are used as input to the digitisation code. The output of this procedure is a long list of digits, each associated with a value of r_{true}.

Unless otherwise noted the results are based on 2 M pions with an energy of 80 GeV, passing straws at an incidence perpendicular to the wire. The external field is taken to be $B = 0$, and the individual straws are selected at random among the (active) barrel straws present in the test beam simulation. The setup thus essentially mimics that of the test beam, although the value of the low threshold (300 eV), the beam energy, and the position along the wires where the particles pass the straw are not exactly similar. For the actual comparison with test beam data in section 12.2, the setup will of course be similar to that of the test beam in all aspects.

For tracking, the distance at closest approach is inferred from the time the LT is first exceeded. As the maximal drift time is around 40 ns, while the discriminated output is digitised in bins of 3.125 ns, around 40 ns/3.125 ns $\approx$ 13 of the 24 bits should contain the leading edge in almost all cases. This is clear from figure 12.2, which shows the distribution of r_{true} as a function of which of the 24 bits contains the leading edge. It is clear that essentially all of the entries in this setup falls into bits 7-21 (the precise subset of course depends on the overall timing offsets of the individual channels). To first approximation the majority of entries follow a pattern of a linear $r - t$ relationship, and lie approximately on a line between the two points (bit 7, 0 mm) and (bit 21, 2 mm). However, this is only approximately true, since a significant amount of entries lie below this line. These are the cases where the random distribution of primary ionisation clusters along the track happens to result in a lower than average energy deposition near point at closest approach, thus resulting in a later leading edge than the linear $r - t$ relationship would predict. Furthermore, a low number of entries lie above the central line. These represent the cases where an independent noise fluctuation exceeded the threshold before the track induced signal (this contribution was already shown in figure 11.8.a).

Figure 12.3 shows the individual distribution of r_{true} for leading edge positions between bit 7 and 22 (thus representing vertical slices of figure 12.2), along with Gaussian fits to the

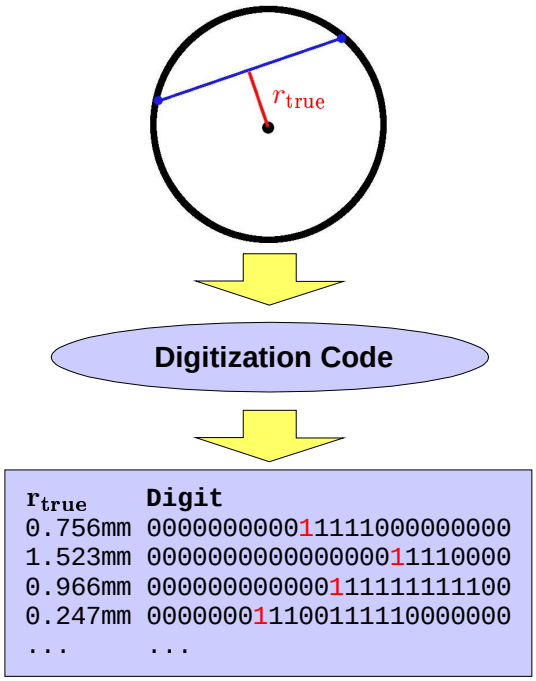

Figure 12.1: *Illustration of the procedure used to produce the data, on which investigations of the performance of the digitisation code are based: Fake straw crossings are fed to the digitisation code as discussed in the text, and the result is a long list of pairs of true distances at closest approach, r_{true}, associated with the resulting digit. For the track reconstruction it is the leading edge information, given by which of the 24 LT bits is the first to indicate an exceeded threshold (indicated in red), that is important.*

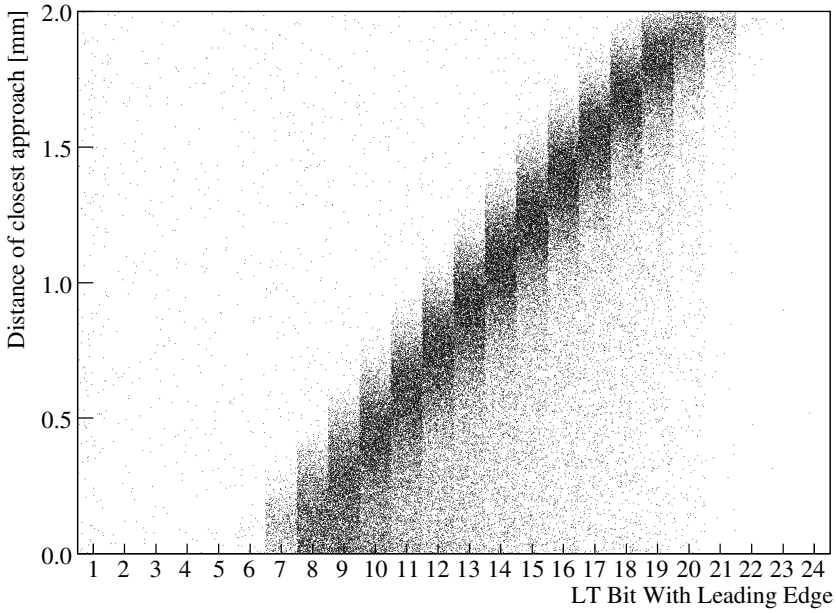

Figure 12.2: *Distribution of r_{true} as a function of which of the 24 LT bits contains the leading edge.*

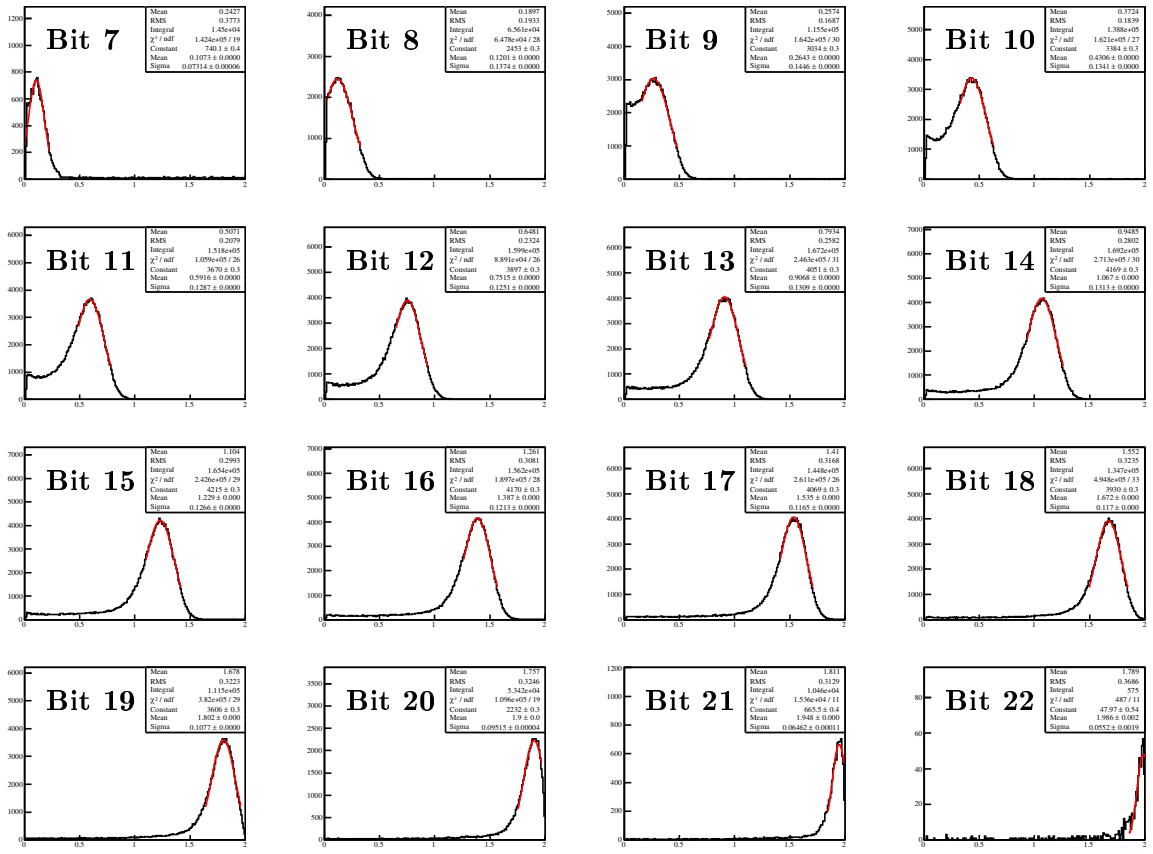

Figure 12.3: *Individual distribution of r_{true} for leading edge positions between bit 7 and 22 (thus representing vertical slices of figure 12.2), along with Gaussian fits to the peaks.*

peaks. In a perfect detector, the distributions should all be (narrow) Gaussians, since this is the assumption used in essentially all tracking algorithms, using these hits. However, the large tails towards lower values of r, makes the assumptions of Gaussian behaviour blatantly untrue. The tails are most severe at the earliest time bins (corresponding, on average, to low values of r_{true}). This is a geometrical effect: At low r the difference in arrival times for clusters spaced $\mathcal{O}(\lambda)$ apart (the mean free path length, $\mathcal{O}(165\ \mu\text{m})$), is larger by about λ/v_{drift} compared to the corresponding difference at larger r. If, for instance, two clusters are created along the track at distances d_1 and d_2 from the distance at closest approach to the wire, the difference in arrival times will be given by:

$$t_1 - t_2 \approx \frac{r_1}{v_{\text{drift}}} - \frac{r_2}{v_{\text{drift}}} = \begin{cases} \mathcal{O}\left(\frac{\lambda}{v_{\text{drift}}}\right) & r_{\text{true}} \ll \lambda \\ \mathcal{O}\left(\frac{\lambda}{v_{\text{drift}}}\frac{\lambda}{r_{\text{true}}}\right) & r_{\text{true}} \gg \lambda \end{cases} \tag{12.1}$$

Here $r_i = (r_{\text{true}}^2 + d_i^2)^{1/2}$ and the d_i's are assumed to be $\mathcal{O}(\lambda)$. At large values of r_{true}, the cascades from electrons in all of the primary ionisation clusters will thus essentially start at the same time, and the time of the leading edge will thus be highly regular (if the LT is exceeded at all) leading to lower contents in the tails in the later plots in figure 12.3.

Another feature of the distributions that is clearly visible in the distributions for leading

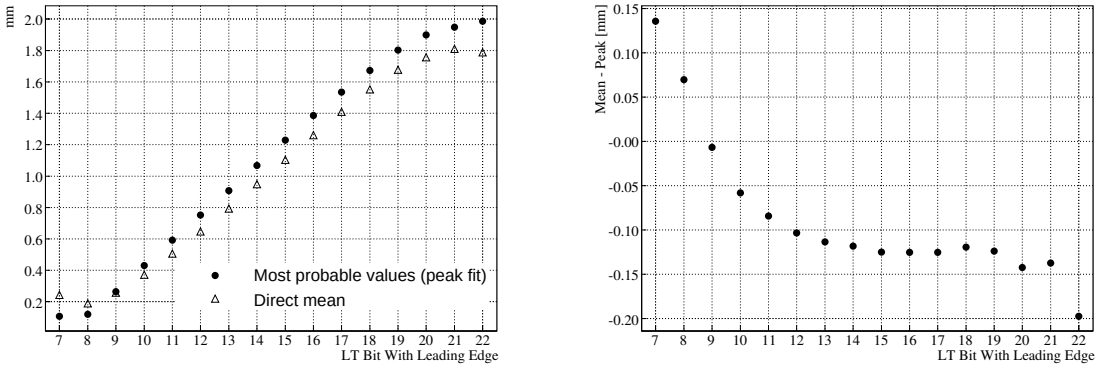

(a) (b)

Figure 12.4: *Mean and most probable (i.e. peak) values of r_{true} for leading edges in different bits (a), and the difference between the two (b).*

edges in bits 9-13, and visible to some degree for bits 14-16, is an enhancement at r_{true} less than ~ 100 μm. This means that tracks passing closer than this distance to the wire have an enhanced probability to register a leading edge "too late", larger than can be simply explained by the distribution of primary clusters. The explanation is of course that electrons starting their drift closer than 100 μm from the wire, will not develop full cascades (cf. figure 9.3).

When constructing the $r-t$ relationships used for tracking, one might naïvely use the mean value of r_{true}, for a given leading edge time bin, as the distance at closest approach estimate for hits with the leading edge in that time bin. The problem with this approach should, however, be clear from figure 12.3: The large and one-sided tails implies that the mean will be located a significant distance away from the majority of entries, contained in the peaks. Figure 12.4 illustrates the difference between the direct means and the peak locations (representing the most probable value). For most time bins the difference is between 100 μm and 150 μm, which is disastrous for the width of the residuals. If, for instance, the means rather than the peaks had been used to construct the $r-t$ relationships used to produce the plots in figure 11.8, the peak in figure 11.8.a would have become broader and have moved around 100 μm to the left. Consequently the "left" and "right" residuals in figure 11.8.b would have peaked in different locations, and the width of the resulting combined residual would become 200-300 μm (with the exact value depending on the fit range). Thus, even though $\sim 80\%$ of the entries are located in narrow peaks of width $110-130$ μm, the resulting residual would have a width much larger than this value. Using the mean as the basis for the $r-t$ relationship is clearly not appropriate, and the centre of the peaks should thus always be used instead.[1]

Instead of using the peak positions as input to track reconstruction algorithms assuming Gaussian measurements, one might wish to pursue a more ambitious strategy and use the distributions in figure 12.3 directly as probability densities for the track-to-wire distance and use likelihood based algorithms in the track reconstruction code. Although a priori "the

[1] One small improvement upon this scheme that might prove useful, however, is to use $r=0$ as the estimated radius in case of the first two or three bits. This is because the distributions in figure 12.3, shown a functions of $|r|$, are really even functions in r. For the first few bits, a Gaussian centred at $r=0$, and with a broader width than indicated by the fits in the figure, might thus better approximate the overall distribution. Whether such a scheme would really improve the tracking performance remains to be tested in detail.

right thing to do", the problem with such an approach is that, while the peak positions are relatively stable under changing detector conditions, the tails are not (which incidentally is another reason to never use the direct mean for $r-t$ relationships). Figure 12.5 shows how the distribution of r_{true} values, resulting in a leading edge at bit 12, changes when either the pion energy, the average value of the low threshold, or the level of backgrounds particle fluxes is varied. It is seen that the peak position and width are only shifted by about ± 30 μm and ± 10 μm, respectively, in the different plots. Figure 12.5.e is made assuming an average straw occupancy of 36%, meaning that about 15% of the resulting digits will contain a background track, while about 2% will contain two.[1]

Even though the width of the central peaks in figure 12.3 of $\mathcal{O}(100-150$ μm) are thus much smaller than the root-mean-squares of the same distributions, it is the latter which will impact tracking algorithms, since there is no way to know whether a given hit is a "peak hit" or a "tail hit". It is thus essential that tracking software implements rather strict outlier cuts, in order to get rid of the hits in the tails (note that this concerns tracking only, and that PID algorithms using high thresholds naturally wish to include as many hits as possible). The natural way to implement such outlier cuts is through a cut on the contribution of the hit to the chi-square of the track fit. This translates directly to a requirement on the number of standard deviations between the hit and wire. Note that this thus implies a cut-distance which depends on the leading edge bit, since the width of the Gaussian fits in figure 12.3 (which gives the estimates of hit uncertainties) are not exactly similar.[2]

A given outlier cut will impact both the root-mean-square of the accepted hits, as well as the average efficiency, defined for a given track as the number of hits on track divided by the number of straws passed by the track. For outlier cuts between 0.5 and 3.5 standard deviations, the resulting root-mean-square of the accepted hits and the efficiency are shown in figure 12.6 for three different values of "track jitter", i.e. the assumed uncertainty in the local track position used to calculate the residual on which the outlier cut decision is based. In ATLAS, track jitter of 50-100 μm is likely the most realistic, while it is slightly lower in the controlled setup of the 2004 test beam (cf. section 12.2).

The optimal value of the outlier criterion represents a trade-off between getting a low root-mean-square of the accepted hits, without impacting the efficiency too much. Figure 12.7 shows the efficiencies versus root-mean-square curves obtained by varying the outlier cut between 0.5 and 3.5 standard deviations. Such a figure, while useful for assessing the usefulness of a given outlier removal scheme, does, however, not provide information about the optimal value of the outlier cut. In order to do so, realise that the track measurements from the TRT can be approximated as being a single high quality measurement that is to be added to the tracks found in the silicon trackers. Given such an approximation, which assumes that a reasonable number of straws, N, are crossed by the straw, and neglects issues due to the spacial distribution of the straws, the resolution of this single TRT hit becomes

$$\sigma_{TRT} \approx \frac{\sigma_{\text{accepted}}}{\sqrt{\epsilon N}} \propto \frac{\sigma_{\text{accepted}}}{\sqrt{\epsilon}} \tag{12.2}$$

where σ_{accepted} is given by the root-mean-square of the accepted hits and ϵ is the efficiency. It thus makes sense to choose a value of the outlier cut which optimises (minimises) the variable

[1] For simplicity, figure 12.5.e was made by mixing an appropriate number of digits *after* discrimination. This obvious approximation should, however, not affect the present discussion.

[2] Also tested was the performance of a straight distance-to-track cut, independent of the leading edge bit. Not surprisingly such a procedure proved less optimal.

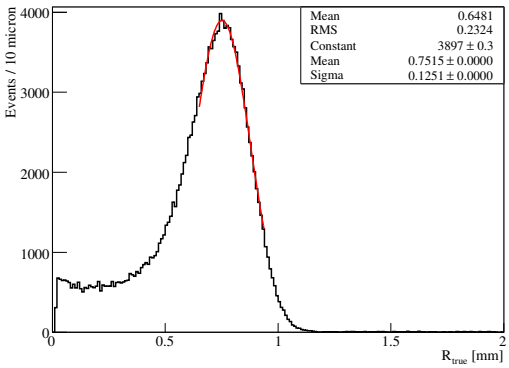

(a): 80 GeV π, LT = 300 eV

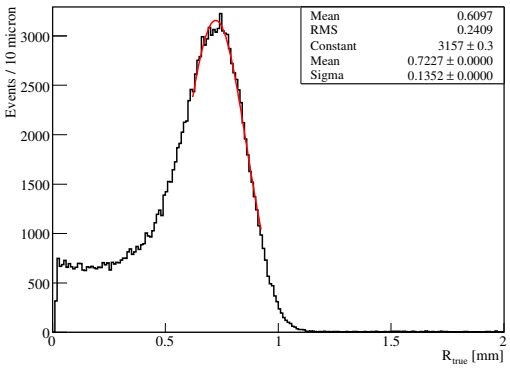

(b): 2 GeV π, LT = 300 eV

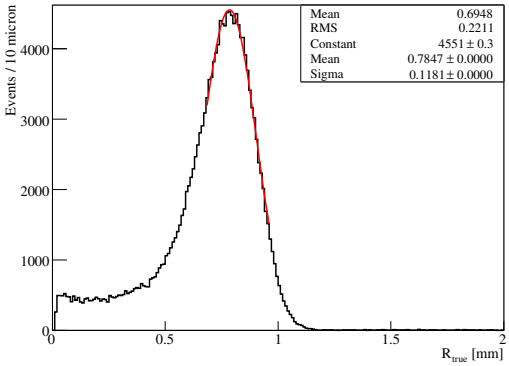

(c): 80 GeV π, LT = 200 eV

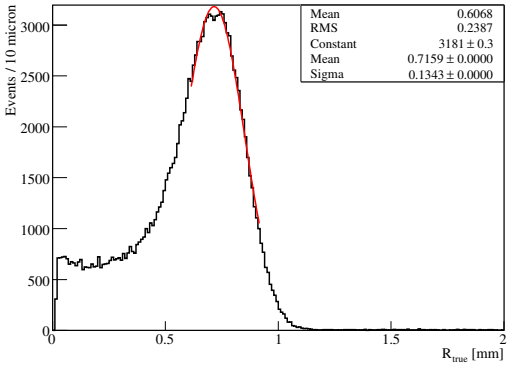

(d): 80 GeV π, LT = 450 eV

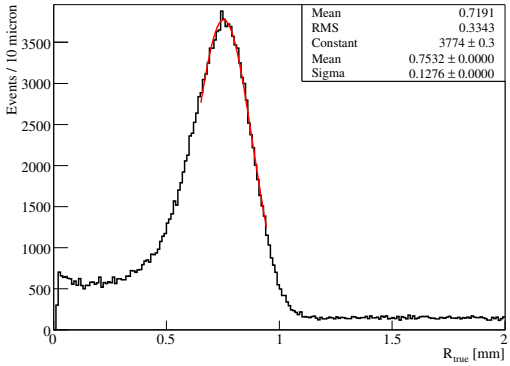

(e): 80 GeV π, LT = 300 eV, with pileup.

Figure 12.5: *Distribution of r_{true} values, resulting in a leading edge at bit 12. Figure (a) shows the distribution resulting from an 80 GeV pion crossing straws with average LT = 300 eV. In figure (b) this is instead a 2 GeV pion, and in figures (c) and (d) the average LT is instead changed to 200 eV and 450 eV respectively. Figure (e) shows the residual in face of background tracks from pileup as discussed in the text.*

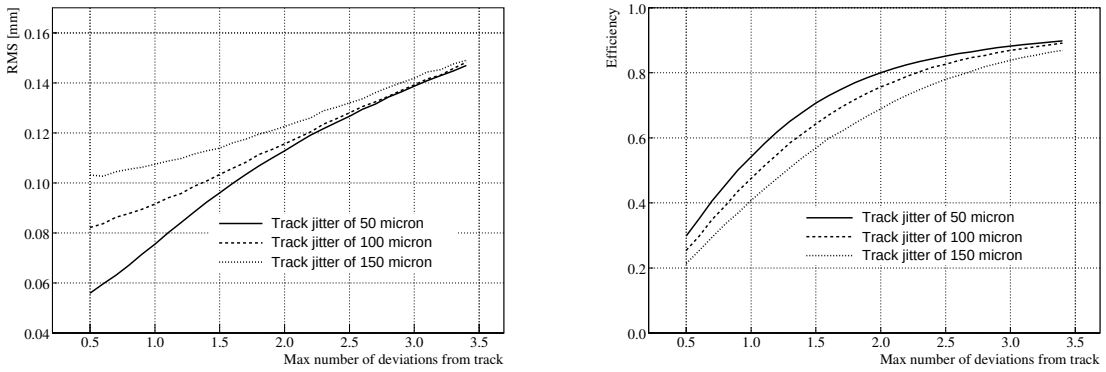

Figure 12.6: *Root-mean-square (a) and efficiency (b) for the accepted hits as a function of the value of the outlier cut, and for three different values of track jitter as discussed in the text.*

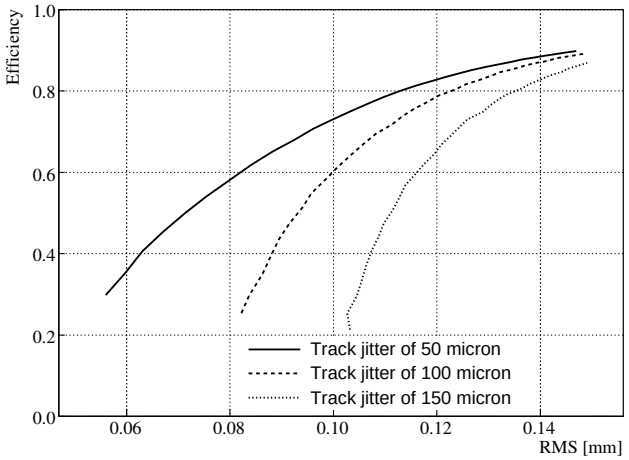

Figure 12.7: *Curves of efficiency versus root-mean-square for the accepted hits obtained by varying the outlier cut, and for three different values of track jitter as discussed in the text.*

$\sigma_{\text{accepted}}/\epsilon^{1/2}$. However, considering pattern recognition and tracks with a low number of crossed straws, a secondary concern should also be to optimise the efficiency. Of course, a tight cut also prevents pattern recognition algorithms from assigning hits to the wrong track in high occupancy environments, so a final optimisation really requires full-blown realistic event simulations.

Figure 12.8.a shows curves of the optimal variable, $\sigma_{\text{accepted}}/\epsilon^{1/2}$, versus the value of the outlier cut, while 12.8.b shows the curves of the same optimal variable versus the efficiency that result from varying the outlier cut between 0.5 and 3.5 sigma. A reasonable compromise between a low value of $\sigma_{\text{accepted}}/\epsilon^{1/2}$ and a high efficiency seems to be obtained by an outlier removal cut of between 2.0 and 2.5 standard deviations (perhaps slightly lower for a track jitter of 50 μm).

While the plots in figure 12.5 show that $r-t$ relationships based on peak positions are remarkable stable in the presence of changing detector conditions and particle fluxes, there are several other effects that can influence them significantly. Fortunately, it should be possible to

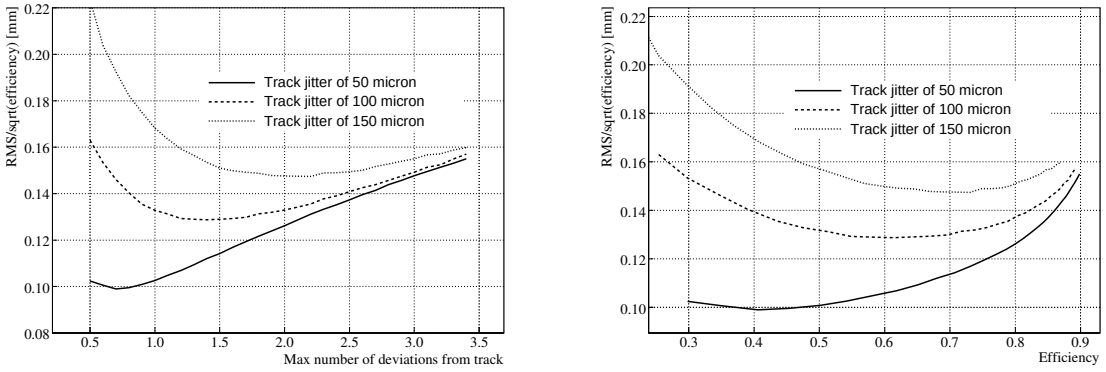

Figure 12.8: *The optimisation variable $\sigma_{accepted}/\epsilon^{1/2}$ versus the outlier cut (a) and efficiency (b) respectively, as discussed in the text.*

correct for those effects, and this is certainly necessary in order to get optimal performance of the detector. The first and possibly most significant effect is the dependency of the detector response upon the angle with which a particle crosses a track. As an example, a particle crossing a straw at $\theta = 60°$ from perpendicular incidence to the wire, will produce twice as many primary ionisation clusters for a given distance at closest approach, due to the longer path-length inside the straw. The plots in figure 12.9 show how the distribution of r_{true} values, resulting in a leading edge at bit 12, changes for different values of the crossing angle, θ.

Similarly, as mentioned in section 9.1, the $r-t$ relationship will also depend on the position along the wire where the particle crosses: In addition to effects such as wire-sagging (which are thought to be negligible, but not otherwise investigated here), the main reason for this is that cascades reaching the wire will give rise to signals travelling towards the electronics, but also to signals in the opposite direction that only reach the electronics after a reflection at the far end. Figures 12.10.a and 12.10.b shows the distribution of r_{true} values, resulting in a leading edge at bit 12, for a particle crossing a 70 cm long barrel straw at the near and far ends respectively.

The variation in peak positions due to the crossing angle and crossing position along the wire shown in figures 12.9 and 12.10 respectively, are seen to be $\mathcal{O}\,(60-100\ \mu\text{m})$. If not taken properly into account these effects will significantly impact the tracking performance of the detector (with the difference in the peak widths affecting the reliability of assumptions going into the track fits, such as chi-square estimates). Of course, unless based on tracks seeded from the silicon trackers, the z coordinate and crossing angle will not be known initially, meaning that a two stage tracking will have to be carried out, with the first step using less precise $r-t$ relationships. Another effect that is likely to be of importance is the impact of the external magnetic field on the $r-t$ relationships in the end caps (cf. figure 9.5). As the field is neither constant over the entire end cap, nor parallel to the wires, such corrections will again depend on the exact track hypothesis.

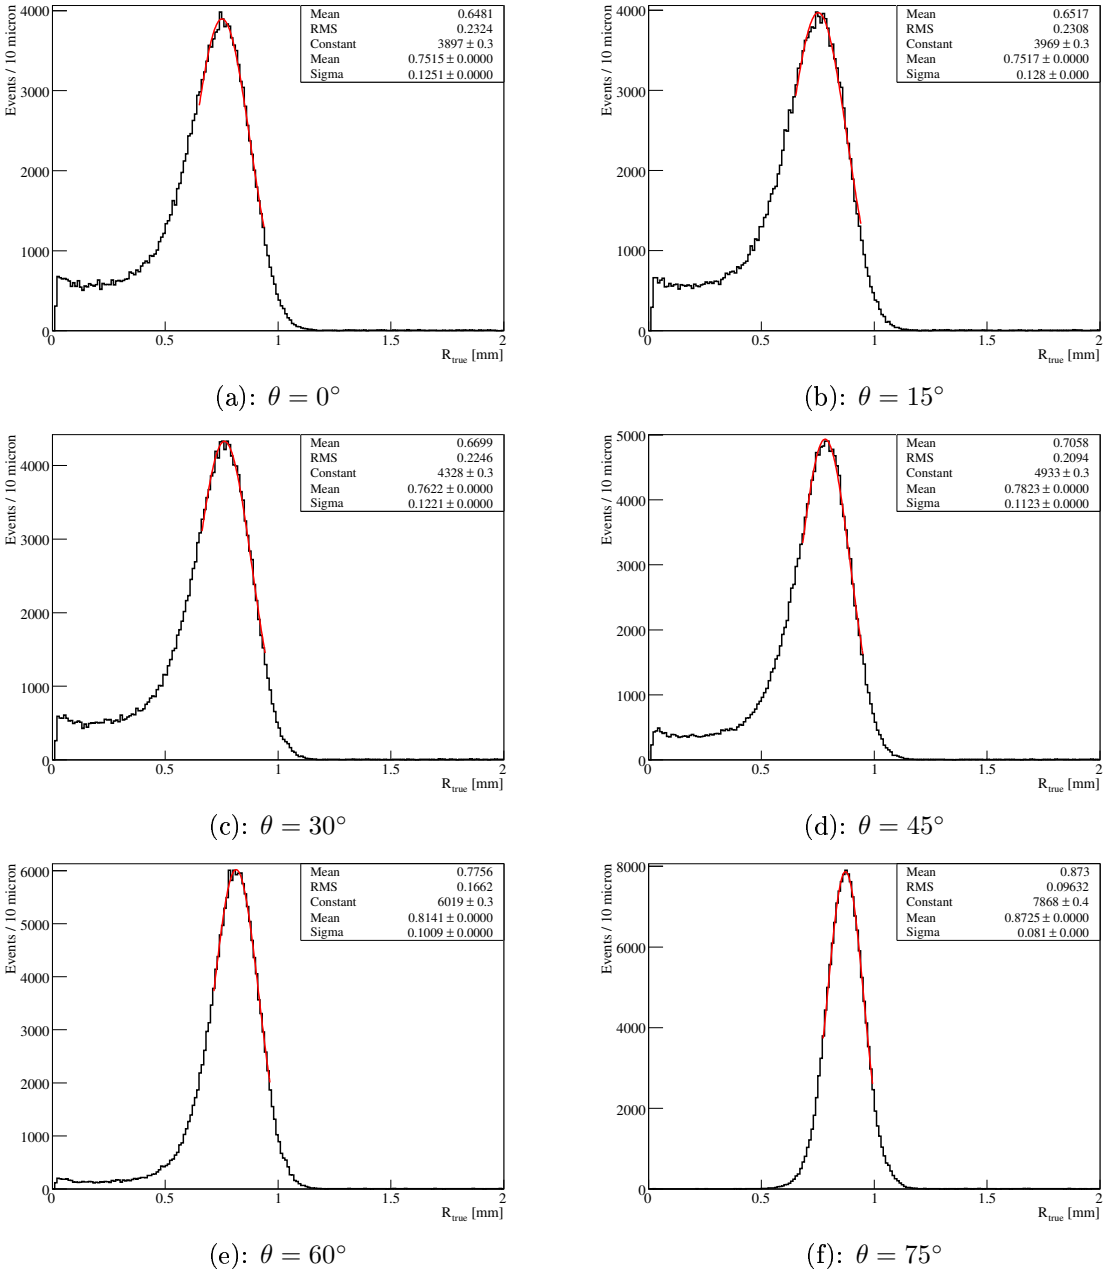

(a): $\theta = 0°$ (b): $\theta = 15°$

(c): $\theta = 30°$ (d): $\theta = 45°$

(e): $\theta = 60°$ (f): $\theta = 75°$

Figure 12.9: *Distribution of r_{true} values, resulting in a leading edge at bit 12, as a function of the angle, θ, of track incidence relative to the wire. θ is thus defined as the angle between the track and the direction perpendicular to the wire. All plots correspond to 80 GeV pions crossing straws with average threshold values of $LT = 300$ eV.*

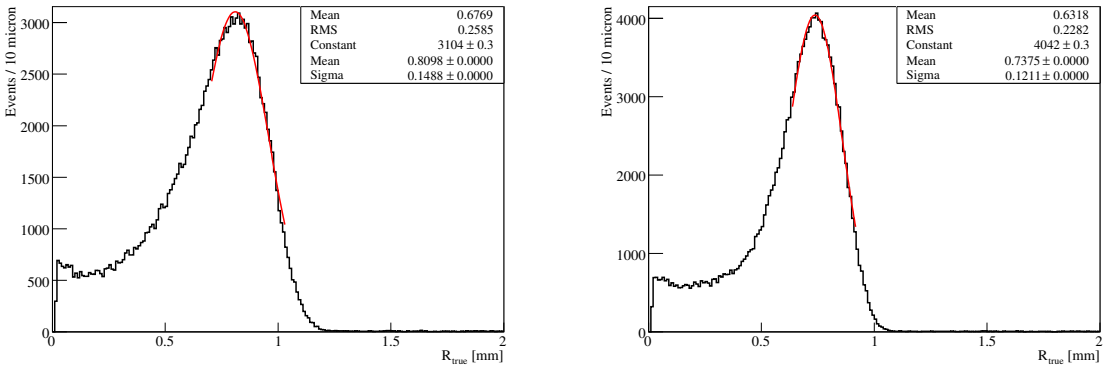

(a) (b)

Figure 12.10: *Distribution of r_{true} values, resulting in a leading edge at bit 12, for tracks crossing the 70 cm long barrel straws at perpendicular incidence, at the end of the straw close to the electronics (a) and the opposite end (b), near $z = 0$. Both plots correspond to 80 GeV pions crossing straws with an average threshold of $LT = 300$ eV.*

12.2 Comparison with Test Beam Data

Following the discussion of the optimal procedure for accessing the tracking information in the TRT digits in the previous section, this section will proceed with a comparison between the output of the digitisation algorithms and data recorded at the 2004 test beam.[1] These data represents a unique opportunity to tune and validate the digitisation code, both since they are taken in the only test beam setup with finalised barrel modules, gas composition and front-end electronics, but also because the identity and geometrical configuration of the passing physics particles are controlled to a degree which will not be possible during data taking in the full ATLAS setup. In the chosen run, 100 GeV pions pass the barrel modules around 15 cm from the front-end electronics, and at angles perpendicular to the wires. An example event from a similar run is shown in figure 10.2.

In order to be able to get a more reliable comparison between track residuals in simulated and real data, the simulated data must be corrected for two known $\mathcal{O}(1\%)$ effects that enters the observed distribution for the test beam data. The first is that the individual wire positions are not known with infinite precision, but rather with a finite resolution estimated to be $\sigma_{\text{align}} \approx 27$ μm. The second effect is that the track position used to determine the residual of a specific hit in a given event is not known with infinite precision either, but is limited by the finite number of TRT hits. In the test beam setup there is no significant magnetic field in the TRT volume, and the tracks will essentially be horizontal lines apart from multiple scattering effects which should be small for 100 GeV pions. A track measurement in the test beam setup can therefore to a good first approximation be seen as a measurement of one parameter, namely the vertical coordinate, y (of course, the actual track fit *does* take multiple scattering effects into account). Assume further that the straws each provide a Gaussian measurement with $\sigma_{\text{true}} \approx 130$ μm and that all events will have a number of hits on track equal to the average of $n = 33.55$ (in the present use case it is not a problem that this is not an integer). As it is the residuals in individual straws that are the objects under investigation, the measurement of the vertical track parameter will be seen as the combination of two measurements: One from the straw under consideration, and one from the other $n-i$ straws. Figure 12.11 defines the various relevant parameters for considering the residuals for a given hit: The vertical misalignment of the straw is given by $\delta_{\text{align}} \equiv y^{\text{wire}}_{\text{align}} - y^{\text{wire}}_{\text{true}}$, and the residual of the track measurement from the other $n-1$ drift tubes, $y^{\text{meas}}_{\text{other}}$, with respect to the true vertical position of the passing particle, $y^{\text{track}}_{\text{true}}$, is denoted δ_{other}. Finally, $r(t)$ is the measured drift radius, given by the leading edge position.

Ideally, one would like to measure the true residual,

$$\Delta_{\text{true}} \equiv r(t) + y^{\text{wire}}_{\text{true}} - y^{\text{track}}_{\text{true}} \tag{12.3}$$

However, the measurements from test beam data are of the observed residual,

$$\Delta_{\text{obs}} \equiv r(t) + y^{\text{wire}}_{\text{align}} - y^{\text{track}}_{\text{recon}} \tag{12.4}$$

Let the error estimate on $r(t)$ be given by $\sigma(t) = a\sigma$ with $a = a(t)$ and $\langle a \rangle_t \approx 1$. The parameter a thus holds the t dependence on the error estimate as illustrated by the varying widths of the fitted Gaussians in figure 12.3, and σ represents the average resolution on individual

[1] The data shown in this section were recorded during the combined running of the test beam, and are taken from run number 2102355. The alignment, calibration and reconstruction of the data were carried out by P. Hansen.

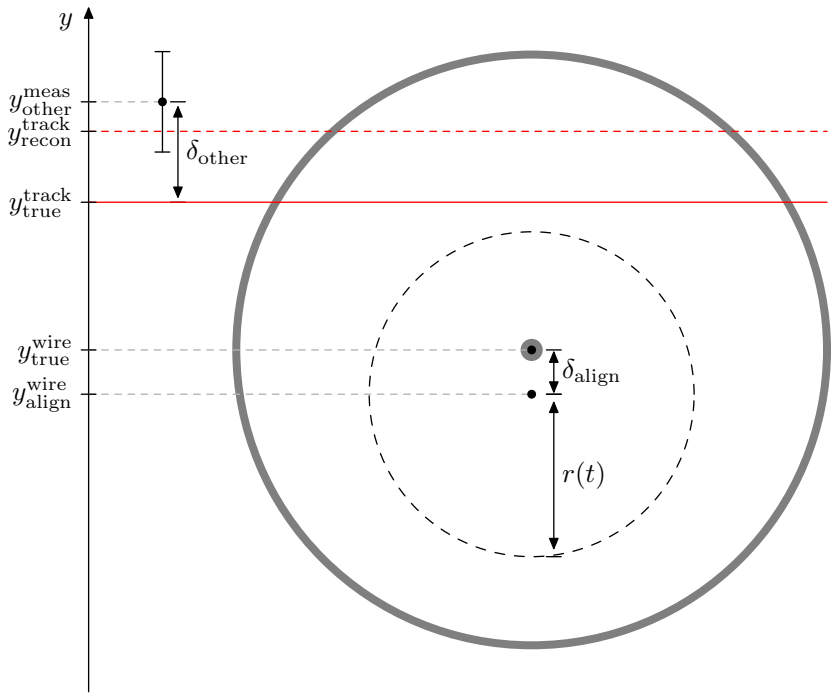

Figure 12.11: *Illustration of the variables used for the discussion of systematic effects in the relationship between the true and observed single straw track residual. Note that the figure is merely an illustration and does not represent realistic scales of the parameters.*

measurements. Assuming that $y^{\text{meas}}_{\text{other}}$ is based upon $n-1$ Gaussian measurements, its error estimate will thus be given by $\sigma_{\text{other}} = \sigma/\sqrt{n-1}$. The reconstructed track parameter thus becomes (assuming e.g. a chi-square fit for the track parameter):

$$y^{\text{track}}_{\text{recon}} = \frac{\sigma^{-2}_{\text{other}} y^{\text{meas}}_{\text{other}} + (a\sigma)^{-2}\left[r(t) + y^{\text{wire}}_{\text{align}}\right]}{\sigma^{-2}_{\text{other}} + (a\sigma)^{-2}} \tag{12.5}$$

$$= \frac{(n-1) y^{\text{meas}}_{\text{other}} + a^{-2}\left[r(t) + y^{\text{wire}}_{\text{align}}\right]}{n-1+a^{-2}} \tag{12.6}$$

And thus,

$$\Delta_{\text{obs}} = \frac{n-1}{n-1+a^{-2}}\left[r(t) + y^{\text{wire}}_{\text{align}} - y^{\text{meas}}_{\text{other}}\right] \tag{12.7}$$

$$= \frac{n-1}{n-1+a^{-2}}\left[\Delta_{\text{true}} + \delta_{\text{align}} - \delta_{\text{other}}\right] \tag{12.8}$$

Here, the factor in front of the parenthesis represents the bias from having the hit, whose residual is being calculated, contribute to the track fit. The second and third term inside the parenthesis represents the (Gaussian) fluctuations due to uncertainty in wire and track positions. A refined comparison between data and simulation can be performed by scaling the generated Δ_{true} with the bias factor, and by adding to this random values of δ_{align} and δ_{other}. The latter should obviously be generated from a zero-mean Gaussians with a width of $\sigma/\sqrt{n-1} \approx 23$ μm, while the spread of the former value is estimated to be $\sigma_{\text{align}} \approx 27$ μm

(the two effects could of course be combined by adding a single number from a single Gaussian distribution of width around 35 μm). Additionally, one might use this relationship to estimate the true resolution of the drift tubes from the width of an observed residual distribution to be (using that $a = 1$ when averaged over all t):

$$\sigma^2_{\text{obs}} = \frac{n-1}{n}\left[\sigma^2 + \left(\frac{\sigma}{\sqrt{n-1}}\right)^2 + \sigma^2_{\text{align}}\right] \tag{12.9}$$

$$\Leftrightarrow \qquad \sigma^2 = \sigma^2_{\text{obs}} - \frac{n-1}{n}\sigma^2_{\text{align}} \tag{12.10}$$

Showing a fortunate cancellation between the bias factor and the uncertainty in the track measurement from the $n-1$ other straws, but still leaving a systematic 2% effect on the width measurement due to misalignment.

In addition to these effects, the uncertainty in the determination of the timing offsets of the individual channels in the test beam are estimated to be between 0.2 ns and 0.5 ns. Each digit in the simulated data will thus include a random Gaussian distributed time-shift of 0.2 ns (0.2 ns rather than 0.5 ns was chosen in order to be conservative, since this value gives the worst agreement with data).

Figure 12.12.a shows the track residuals of the test beam data compared to those given by the digitisation code for an identical setup and an average LT of 300 eV. Both data and digitisation have removed outliers more than 5 standard deviations from the track, and the $r-t$ relationships and estimations of hit uncertainty for both data and simulation uses the procedure outlined in the previous section. The residuals from the simulation in figure 12.12.a are clearly too broad, and with an enhanced tail content compared to the data. With the distributions in figures 12.5.a-d in mind, this would seem to indicate that a threshold value of 300 eV is too high, or that the amount of ionisation is underestimated. Of course, as discussed in section 9.1, a wrong value of the average LT really indicates a wrong conversion between the offline value of the average LT, expressed as an energy deposition, and its online value, which is really given by a current in the front-end electronics. Problems in this conversion might really be related to issues with the initial ionisation, the gas gain factor, the signal shaping and amplification or the actual threshold discrimination. Fortunately, this means that it is possible to calibrate for all of these unknown effects by the simple tuning of this single parameter, the average LT. That this works well is underlined by the similarity of the distributions in figures 12.5.b (a less ionising particle than in figure 12.5.a) and 12.5.d (a higher average value of the LT than in figure 12.5.a), as well as by the similarity between the distributions in figures 12.5.c (a lower average value of the LT than in figure 12.5.a) and 12.9.d (effectively a higher number of ionisations per track length in the plane perpendicular to the wire than in figure 12.5.a).

Another indication that the average value of LT = 300 eV is too high, is given by the comparisons of the time over threshold distributions in figure 12.12.b, from which it is clear that the simulation has a tendency to provide too low time over thresholds. Finally, the LT efficiency (i.e. the fraction of straw crossings leading to digits with at least one set LT bit) is obviously also a very important parameter that the simulation should reproduce. Since some straws in the test beam were unresponsive or dead (cf. figures 10.13 and 10.14), this efficiency is difficult to assess precisely, but is estimated to be $(95 \pm 2)\%$, with the errors being systematic. The simulation with an average LT of 300 eV gives a worse LT efficiency of 87.0%, again indicating that the value of 300 eV is too high.[1]

[1] To understand why a value of 300 eV, which was the control-room value during data taking, is too high,

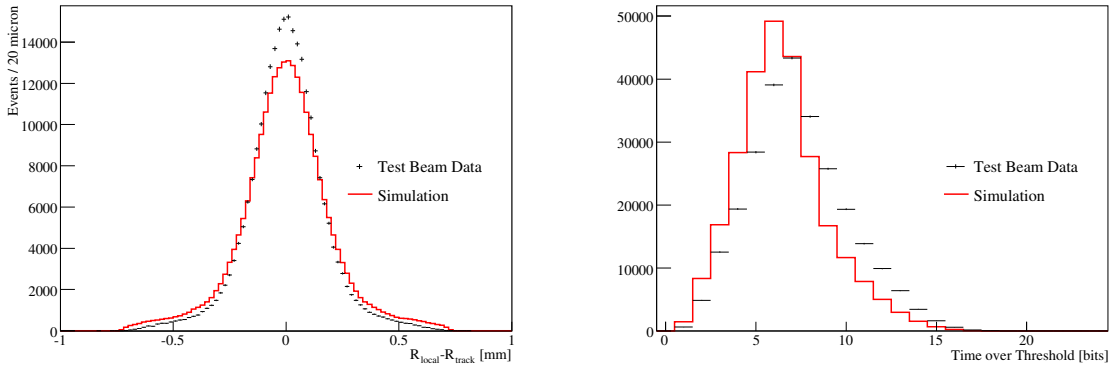

(a) (b)

Figure 12.12: *Comparisons between test beam data and the output of the digitisation for a similar setup and an average LT value of 300 eV. Figure (a) and (b) shows the distributions of track residuals and time over threshold, respectively. In both cases the high statistics simulation is scaled to match the test beam data.*

Tuning the average value of the LT to best reproduce the observed distribution of track residuals yields a value of LT = 161 eV. The resulting distribution is shown in figure 12.13.a, which is essentially indistinguishable from the observed distribution (a Kolmogorov test yields a compatibility probability of 71%). Even though the tuning of the LT guarantees that the width of the central peaks can always be made to match (for any particular definition of such, which could for instance be given by a Gaussian fit over ± 200 μm), the fact that the non-Gaussian tails are also well described certainly constitutes a non-trivial success. More importantly, however, is the fact that the corresponding LT efficiency is 94.6%, consistent with the observed value.

The time over threshold distribution following from the lowered threshold, shown in figure 12.13.b, also gives a better agreement with data, although there are still discrepancies. This is not surprising since the time over threshold will depend strongly on the signal shaping, which is known to have a scale dependency which is not simulated (cf. section 9.2). Also, drift electrons created in the straw walls by the positive ions are not simulated, even though a few will reach the wire without being reabsorbed by the gas. The presence of such electrons should have a negligible impact on the leading edge resolution, but might influence the trailing edge slightly.

In conclusion, after the necessary tuning of the average value of the LT, the digitisation code seems to be able to reproduce the details of the test beam data well. It will be interesting, albeit challenging, to perform similar detailed comparisons in more complex setups involving data from cosmic particles or real collisions at the LHC. Such comparisons, and the improvements to the code likely to follow from them, should ensure that the simulations can achieve high and increasing levels of realism. Not only will this be obviously beneficial for the development of physics analyses, but a realistic simulation is also helpful in the developments of reconstruction algorithms which strive to get the maximal performance of the detector.

one need to understand the origin of the offline to online threshold conversion factor that was used. This conversion factor was based on the result of previous versions of the digitisation code which did not include several of the smaller effects such as the spread in drift times and noise which contribute to the width of the track residual. In order to reproduce the widths of the observed track residual, the LT thus had to be put at a too high value (with a negative impact on the LT efficiency).

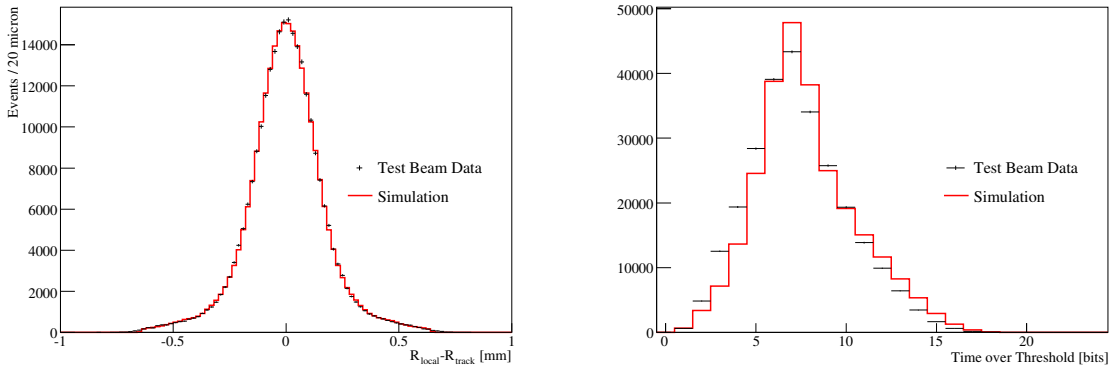

(a) (b)

Figure 12.13: *Comparisons between test beam data and the output of the digitisation for a similar setup and an average LT value of 161 eV. Figure (a) and (b) shows the distributions of track residuals and time over threshold, respectively. In both cases the high statistics simulation is scaled to match the test beam data.*

Part IV
Slepton Spin through Dilepton Signals

As described in section 6, it is likely to be straight-forward to observe physics beyond the Standard Model at the LHC at the ATLAS and CMS experiments if such phenomena indeed exists at the TeV scale. There is, however, a big step between making the first direct observations of deviations from the Standard Model and being able to identify their nature satisfactorily (in any sense of the word). Apart from spectacular phenomena such as black hole production, an identification would in general rely on precise measurements of couplings and spins of the new particle states. Due to the high backgrounds and unconstrained initial states at the LHC, such measurements could prove exceedingly difficult for all but the most prominent production and decay channels – in particular when the new physics is associated with a new global symmetry and stable invisible particles consequently are found amongst the final decay products of all new particles. It is therefore unlikely that a solid classification could be achieved before a TeV scale linear collider is operational.

This part of the thesis will present a model independent method, which can be used, at both the LHC and a future linear collider, to extract information about the spin of new particles carrying leptonic quantum numbers in models associated with new global symmetries. In such models cascade decays, where a pair of SM leptons are emitted "around" a new particle with leptonic quantum numbers, X_{lep}, as shown in figure 13.1, can occur. The emitted leptons will be denoted "near" and "far", respectively, depending on their distance from the production vertex. The observed dilepton invariant mass distribution will depend upon the spin state of X_{lep} (and if it is not a scalar its polarisation state will of course depend on the decay mothers). These decays can indeed exist in both supersymmetric models, where the corresponding particles and symmetry are the sleptons and R-parity, and universal extra dimensions (UED) where they are KK excitations of SM leptons and KK parity respectively. In the following, the method will be tested for feasibility in the scenario where the mass spectra and couplings are given by the Constrained Minimal Supersymmetric Standard Model (CMSSM) and the possibility of discriminating against models with one universal extra dimension will explicitly be addressed.

In section 13 the method will be introduced and in 14 the relevant kinematic parameters in CMSSM, along with the expected changes in light of an UED hypothesis, investigated. Section 15 will then test this method on two well known ATLAS CMSSM points at the generator level, using realistic acceptance cuts, and will discuss how to correct for efficiency effects. Based on this, section 16 will then proceed with scans of CMSSM space including high-statistics toy Monte Carlo experiments at each point, to determine how generally applicable this method is expected to be.

13 Connecting the Distributions of Dilepton Mass and Decay Angles

In this section the proposed method for extracting information about the decay angle of the new physics particle X_{lep} in figure 13.1 will be presented. Section 13.1 will set out by deriving the appropriate formulae and discuss how the invariant dilepton mass connects to the decay angle of the X_{lep} and section 13.2 will go through the details of how to extract decay angle parameters directly from an observed invariant dilepton mass distribution.

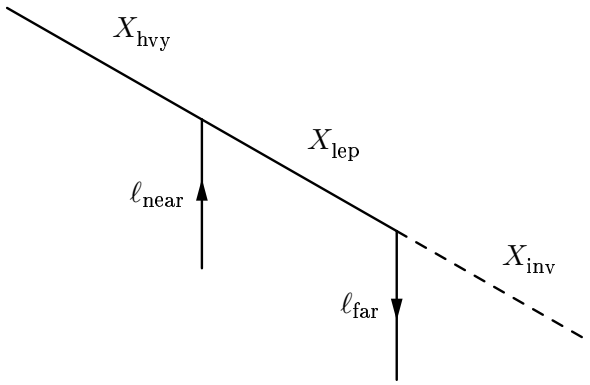

Figure 13.1: *Generic cascade decay involving new physics particles where two SM leptons are emitted along with an invisible stable particle.*

13.1 Dilepton Mass Edge Phenomenology

Figure 13.2.a and 13.2.b shows the kinematics of the $X_{\text{hvy}} \to X_{\text{lep}}\ell_{\text{near}}$, $X_{\text{lep}} \to X_{\text{inv}}\ell_{\text{far}}$ decay chain in the X_{hvy} and X_{lep} rest frames respectively. As the X_{hvy} decay is two-body, the three-momenta of ℓ_{near} and X_{lep} will be of equal magnitude and opposite direction in the X_{hvy} rest frame, and the magnitude will be determined exclusively by the X_{hvy}, X_{lep} and ℓ_{near} masses and thus be the same in all decays. This means that the boost to the X_{lep} frame will always be parallel to the ℓ_{near} direction, and of a fixed magnitude. Consequently, the magnitude of the three-momentum of ℓ_{near} in the X_{lep} frame will be the same in all decays.

As the X_{lep} decay is likewise two-body, the three-momenta of the ℓ_{far} and X_{inv} particles in the X_{lep} rest frame will similarly be equal in magnitude and opposite in direction, and with a magnitude determined by the masses of X_{lep}, X_{inv} and ℓ_{far}. This means that the magnitudes of the three-momenta of both ℓ_{near} and ℓ_{far} in the X_{lep} rest frame will be the same in all decays. The combined invariant mass of these two leptons will in a given decay thus be fully determined by the (cosine of) the angle, θ, between them in this frame.[1] Thus,

$$m_{\ell\ell} \equiv m_{\ell\ell}\left(\cos\theta, M_{\text{hvy}}, M_{\text{inv}}, M_{\text{lep}}, m_\ell\right) \approx m_{\ell\ell}\left(\cos\theta, M_{\text{hvy}}, M_{\text{inv}}, M_{\text{lep}}\right) \tag{13.1}$$

Where it has been made explicit that the masses of the observed leptons will be utterly negligible at the relevant energy scales.

To find the actual content of expression (13.1) one starts by finding the energies and momenta of X_{lep} and ℓ_{near} in the X_{hvy} rest frame, using the natural assumption that $m_\ell^2 \ll$

[1] One might wish to identify θ as being "the decay angle" of the X_{lep} particle. However, the decay angle of a particle is not actually a well defined concept in itself (even when discounting the obvious ambiguity in the selection of which daughter particle in the decay, the angle should specify the direction of): First of all, the frame of reference where the angle is given must be specified. The only obvious choice here is the rest frame of the decaying particle. However, the decaying particle does not have a well defined direction in its own rest frame. One might try to rectify this by taking as its direction, its direction in *another* frame, typically the laboratory frame or the frame of the mother of the decaying particle (X_{hvy} in this case). From a physics point of view, the latter choice is the natural one, as the wave function of the mother is likely to influence the decay of the decaying particle itself. The present definition of theta as the angle between ℓ_{near} and ℓ_{far} in the X_{lep} rest frame is consistent with the decay angle defined in this way (modulo π), so in that sense, θ can indeed be thought of as being a decay angle of the X_{lep} decay.

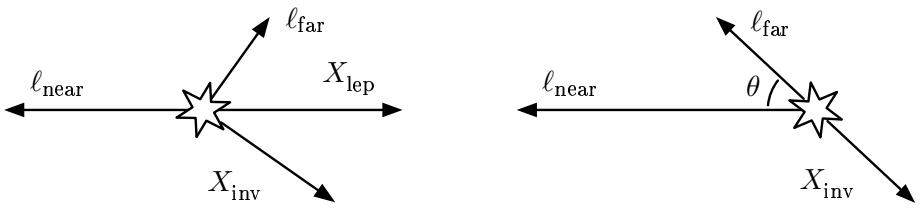

Figure 13.2: *Kinematics of the $X_{hvy} \to X_{lep}\ell_{near}$, $X_{lep} \to X_{inv}\ell_{far}$ cascade decay. Seen from the X_{hvy} (a) and X_{lep} rest frame (b) respectively.*

$M^2_{\text{hvy}}, M^2_{\text{lep}}$:[1]

$$E'_{\text{lep}} = \frac{M^2_{\text{hvy}} + M^2_{\text{lep}}}{2M_{\text{hvy}}} \tag{13.2}$$

$$E'_{\ell_{\text{near}}} \approx \left|\vec{p}'_{\ell_{\text{near}}}\right| = \left|\vec{p}'_{\text{lep}}\right| = \frac{M^2_{\text{hvy}} - M^2_{\text{lep}}}{2M_{\text{hvy}}} \tag{13.3}$$

The boost to the X_{lep} rest frame has

$$\begin{aligned} \beta \equiv \frac{\left|\vec{p}'_{\text{lep}}\right|}{E'_{\text{lep}}} = \frac{M^2_{\text{hvy}} - M^2_{\text{lep}}}{M^2_{\text{hvy}} + M^2_{\text{lep}}}, \quad \gamma \equiv \frac{E'_{\text{lep}}}{M_{\text{lep}}} = \frac{M^2_{\text{hvy}} + M^2_{\text{lep}}}{2M_{\text{hvy}}M_{\text{lep}}} \\ \Rightarrow \quad \gamma\,(1+\beta) = \frac{M_{\text{hvy}}}{M_{\text{lep}}} \end{aligned} \tag{13.4}$$

and the energy and momenta of ℓ_{near} in the X_{lep} rest frame thus becomes,

$$E_{\ell_{\text{near}}} \approx \left|\vec{p}_{\ell_{\text{near}}}\right| = \gamma\left|\vec{p}'_{\ell_{\text{near}}}\right| + \gamma\beta E'_{\ell_{\text{near}}} \approx \gamma\,(1+\beta)\left|\vec{p}'_{\ell_{\text{near}}}\right| = \frac{M^2_{\text{hvy}} - M^2_{\text{lep}}}{2M_{\text{lep}}} \tag{13.5}$$

Whereas the corresponding parameters for ℓ_{far} are given directly from the two-body X_{lep} decay:

$$E_{\ell_{\text{far}}} \approx \left|\vec{p}_{\ell_{\text{far}}}\right| \approx \frac{M^2_{\text{lep}} - M^2_{\text{inv}}}{2M_{\text{lep}}} \tag{13.6}$$

Thus,

$$\begin{aligned} m^2_{\ell\ell} &= \left(E_{\ell_{\text{near}}} + E_{\ell_{\text{far}}}\right)^2 - \left(\vec{p}_{\ell_{\text{near}}} + \vec{p}_{\ell_{\text{far}}}\right)^2 \approx 2\left|\vec{p}_{\ell_{\text{near}}}\right|\left|\vec{p}_{\ell_{\text{far}}}\right|\left(1 - \cos\theta\right) \\ &\approx \frac{\left(M^2_{\text{hvy}} - M^2_{\text{lep}}\right)\left(M^2_{\text{lep}} - M^2_{\text{inv}}\right)}{M^2_{\text{lep}}}\frac{1-\cos\theta}{2} \end{aligned} \tag{13.7}$$

[1] From the two-body decay formulae 38.15 and 38.16 in [PDG06], with $m_\ell = 0$.

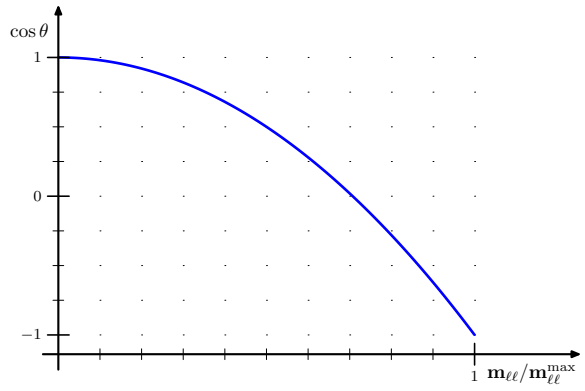

Figure 13.3: *The one-to-one relationship between $m_{\ell\ell}$ and $\cos\theta$ given by equation (13.9).*

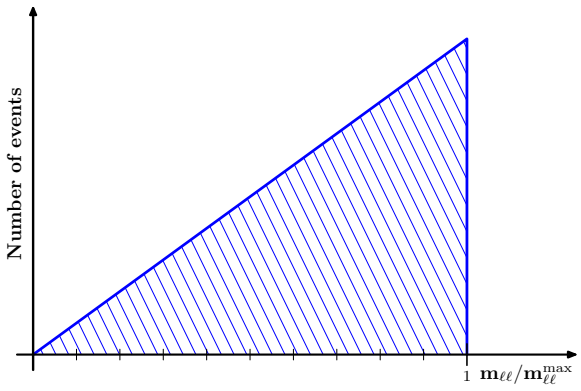

Figure 13.4: *Distribution of $m_{\ell\ell}$ in the case where X_{lep} is a scalar.*

And since $\cos\theta$ can only attain values in $[-1,1]$, the distribution of $m_{\ell\ell}$ will have an upper kinematic limit of:

$$m_{\ell\ell}^{\max} = \frac{\sqrt{M_{\text{hvy}}^2 - M_{\text{lep}}^2}\sqrt{M_{\text{lep}}^2 - M_{\text{inv}}^2}}{M_{\text{lep}}} = M_{\text{lep}}\sqrt{\frac{M_{\text{hvy}}^2}{M_{\text{lep}}^2} - 1}\sqrt{1 - \frac{M_{\text{inv}}^2}{M_{\text{lep}}^2}} \tag{13.8}$$

and its particular value in a given event will be fully determined by θ:

$$m_{\ell\ell} = m_{\ell\ell}^{\max}\sqrt{\frac{1-\cos\theta}{2}} = m_{\ell\ell}^{\max}\sin\theta/2 \tag{13.9}$$

The simplicity of this formula is rather extraordinary. It implies that the cosine of the decay angle can be measured directly from the momenta of the two hard leptons in the final state as shown in figure 13.3. This even in a complex event at a hadron collider where initial and final states have unconstrained momenta and where the X_{hvy} particle might have an arbitrary boost with respect to the laboratory frame. As long as the right events can be reliably singled out and the backgrounds kept under control, it would a priori seem like a rather reliable method of accessing the spins of the new particles in the decay chain. Of course, as $m_{\ell\ell}$ is the directly measured quantity – and as $\cos\theta$ is not defined for background events – it is the $m_{\ell\ell}$ distribution that should be measured directly.

When X_{lep} is a scalar particle (such as in supersymmetric models) ℓ_{far} will be emitted with equal probability in all directions in the X_{lep} rest frame, i.e. with $dP/d\Omega = (4\pi)^{-1}$. Ignoring ϕ which does not enter into (13.9), this means that $\cos\theta$ will be distributed flatly in $[-1,1]$,

$$\frac{dP_{\text{scalar}}}{d\cos\theta} = \frac{1}{2} \tag{13.10}$$

This means that $m_{\ell\ell}$ values can be generated as $m_{\ell\ell}^{\max}\sqrt{R}$ with R uniformly distributed in the unit interval. Familiarity with the transformation method for Monte Carlo generation immediately reveals that the $m_{\ell\ell}$ distribution has the characteristic triangular shape depicted in figure 13.4.[1] Of course, the derivation here assumes that the new particles X_{hvy}, X_{lep} and

[1] In short, the inverse of $\sqrt{y}$ is x^2 which is an anti-derivative of x. This implies that the transformation method predicts a linear distribution shape in the allowed interval.

X_{inv} all have negligible widths. This is a rather robust and model independent assumption as long as these particles do not carry colour-charge, as e.g. widths smaller than $\mathcal{O}(0.1\text{ GeV})$ correspond to lifetimes greater than 10^{-23} s – a limit far below the typical values associated with weak and electromagnetic interactions. Of course, if X_{inv} is stable it will have exactly zero width.

13.1.1 Extending to Non-Scalar X_{lep}

If X_{lep} is not a scalar, then the distribution of $\cos\theta$ will not be flat in $[-1,1]$, but will have a different shape depending upon the spins and polarisations of the particles in the decay chain. The central analysis idea presented here is based upon the recognition of the fact that, due to the forms of the lowest level spherical harmonics, the possible distributions of $\cos\theta$ will in most conceivable scenarios contain only a limited type of terms: Either constant or proportional to $\cos\theta$ or $\cos^2\theta$. Without further assumptions regarding the properties of the underlying physics scenario it should thus be possible to extract the coefficients of these non-constant terms from data. The extracted values can then subsequently be compared to predictions of the competing models. The studies in the present thesis will limit themselves to the cases of having either a $\cos\theta$ term or a $\cos^2\theta$ present, but not both at once. It is of course perfectly conceivable that one might wish to consider both terms at once in a more complicated analysis. In any case, if the data have enough sensitivity to distinguish a given value of b from zero, then a goodness-of-fit estimate should also be able to tell the two distributions $1+b\cdot\cos\theta$ and $1+b\cdot\cos^2\theta$ from each other.[1]

Properly normalised the two distribution families chosen for investigation thus looks like:

$$\frac{dP_1}{d\cos\theta} \equiv \frac{1}{2}\left[1+b\cdot\cos\theta\right], \quad b\in[-1,1] \tag{13.11}$$

$$\frac{dP_2}{d\cos\theta} \equiv \frac{3}{6+2b}\left[1+b\cdot\cos^2\theta\right], \quad b\in[-1,\infty[\tag{13.12}$$

Here b is the "decay parameter", that is to be measured experimentally and subsequently compared to the predictions from different physics models. An important point is that a measurement of $b\neq 0$ will rule out a supersymmetric X_{lep} even without an alternative hypothesis. Inverting and differentiating equation (13.9) with respect to $m_{\ell\ell}$ gives:

$$\cos\theta = 1-2\left(m_{\ell\ell}/m_{\ell\ell}^{\max}\right)^2 \implies \frac{d\cos\theta}{dm_{\ell\ell}} = -\frac{4m_{\ell\ell}}{(m_{\ell\ell}^{\max})^2} \tag{13.13}$$

and thus the general formula for the shape of the $m_{\ell\ell}$ mass edge for a given $dP/d\cos\theta$:

$$\frac{dP}{dm_{\ell\ell}} = \frac{4m_{\ell\ell}}{(m_{\ell\ell}^{\max})^2}\frac{dP}{d\cos\theta}\Theta\left[m_{\ell\ell}\right]\Theta\left[m_{\ell\ell}^{\max}-m_{\ell\ell}\right] \quad \text{with} \quad \cos\theta = 1-2\left(m_{\ell\ell}/m_{\ell\ell}^{\max}\right)^2 \tag{13.14}$$

where the fact that the distribution is vanishing outside $[0, m_{\ell\ell}^{\max}]$ has been included explicitly in step functions. Note how the factor of $m_{\ell\ell}$ always forces the distribution towards zero for small values of $m_{\ell\ell}$. This is a consequence of the horizontal asymptote of the curve in figure

[1] If both distributions fail to describe the data adequately it might of course be because $M_{\text{hvy}} < M_{\text{lep}}$, with the decay proceeding through a virtual M_{lep} (or a virtual Z boson, cf. [CST03]). In such a case the decay is really three-body and the formula derived in the previous section will not be valid. The only model-independent prediction in such scenarios is that $m_{\ell\ell} < M_{\text{hvy}} - M_{\text{inv}}$.

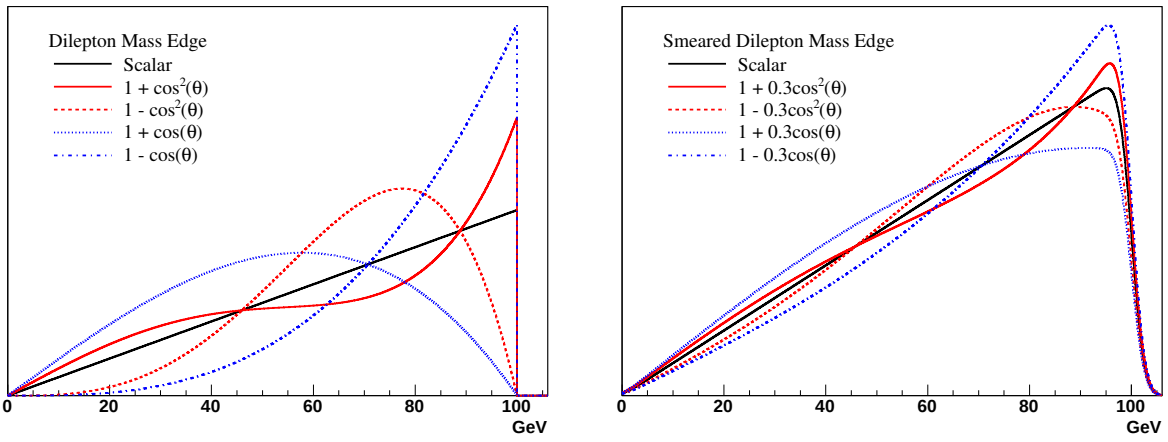

(a) (b)

Figure 13.5: *Hypothetical dilepton invariant mass distribution with* $m_{\ell\ell}^{max} = 100$ *GeV and various decay angle distributions corresponding to equations (13.15)–(13.16). Figure (a) shows raw distributions for* $b = \pm 1$ *while figure (b) shows distributions for* $b = \pm 0.3$ *after a (realistic) Gaussian smearing with* $\sigma = 2$ *GeV.*

13.3 at $m_{\ell\ell} = 0$. Defining $\hat{m}_{\ell\ell} \equiv m_{\ell\ell}/m_{\ell\ell}^{\max}$, the two distributions in equations (13.11) and (13.12) yield respectively:

$$\frac{dP_1}{d\hat{m}_{\ell\ell}} = 2\hat{m}_{\ell\ell}\left[1 + b\left(1 - 2\hat{m}_{\ell\ell}^2\right)\right], \quad 0 \le \hat{m}_{\ell\ell} \le 1 \tag{13.15}$$

$$\frac{dP_2}{d\hat{m}_{\ell\ell}} = \frac{6\hat{m}_{\ell\ell}}{3+b}\left[1 + b\left(1 - 2\hat{m}_{\ell\ell}^2\right)^2\right], \quad 0 \le \hat{m}_{\ell\ell} \le 1 \tag{13.16}$$

In figure 13.5.a is shown the shape of the mass edge for $b = \pm 1$ and $m_{\ell\ell}^{\max} = 100$ GeV for both P_1 and P_2 as well as the regular triangular shape associated with a scalar decay corresponding to $b = 0$ in both cases. When the decay angle is distributed according to $1 + b \cdot \cos^2\theta$, positive values of b results in more central decays around $\theta = \pm\pi/2$ while the opposite is true for negative values of b. As $\theta = \pm\pi/2$ is equivalent to $m_{\ell\ell} = m_{\ell\ell}^{\max}/\sqrt{2}$ this corresponds to either a depletion or an augmentation of the $m_{\ell\ell}$ distribution around 70 GeV. When the decay angle on the other hand is distributed according to $1 + b \cdot \cos\theta$, positive and negative values of b will lead to either more forward or backward decays, favouring $\cos\theta \approx -1$ and $\cos\theta \approx 1$ respectively. This leads to either an enhancement or a depletion of the peak near $m_{\ell\ell} \approx m_{\ell\ell}^{\max}$.

To give a better picture of what might be observed experimentally the edges are shown for $b = \pm 0.3$ and with a slight smearing in figure 13.5.b. Note how the position of the edge at $m_{\ell\ell}^{\max}$ is quite stable under the different configurations. This is fortunate as the correlations between measurements of $m_{\ell\ell}^{\max}$ and of the actual shape itself can thus be expected to be minimal.

13.2 Fitting Procedure

To actually extract information about the decay parameter b from a given observed distribution of $m_{\ell\ell}$ like those shown in figure 13.5.b, one will need to fit it with the theoretical distributions of (13.15) and (13.16), using methods such as unbinned log-likelihood or binned chi-square

fits. It is important that such a procedure takes experimental resolution and backgrounds into account.

This subsection will deal with issues specific to the signal distribution: The information it contains about decay parameters (section 13.2.1), how it should be modified to include effects of $m_{\ell\ell}$ resolution (section 13.2.2) and how to deal with the situation where the allowed range of a fitted decay parameter extends into unphysical regions (section 13.2.3).

13.2.1 Sensitivity Estimate in the Ideal Limit

Ultimately the resolution that can be obtained on the decay parameter, b, will thus depend upon several factors including signal purity, the uncertainty on the individual observations of $m_{\ell\ell}$, correlations between b and other parameters, the available statistics and how sensitive the distribution is to changes in b at a given value of that parameter. For comparison it is useful to use an approximation to estimate the latter sensitivity. Given a normalised probability density function, $f(x;\alpha)$, where x is the observable and α the parameter under investigation. Assume that there are no background events, perfect x measurements, vanishing correlations between α and other parameters, and that f is well approximated by its second order expansion. The resulting resolution on α, σ_α, based on N observed values of x will then be given by the following formula for sufficiently high N [Cow98, chap. 6.6]:

$$\sigma_\alpha^{-2} = N \int dx \frac{1}{f(x;\alpha)} \left(\frac{\partial f(x;\alpha)}{\partial \alpha} \right)^2 \tag{13.17}$$

Denoting by $I_1(b)$ and $I_2(b)$ the value of the integral in equation (13.17) in the case where f is given by the probability densities for the $m_{\ell\ell}$ distributions corresponding to $1 + b \cdot \cos\theta$ and $1 + b \cdot \cos^2\theta$, respectively (cf. equations (13.15) and (13.16)), they can be solved to yield:

$$I_1(b) = \frac{\operatorname{arctanh} b - b}{b^3} \tag{13.18}$$

$$I_2(b) = \frac{3 \arctan\sqrt{b}}{b^{5/2}(b+3)} - \frac{9}{b^2(b+3)^2} \tag{13.19}$$

The latter expression can be evaluated for $b < 0$ by using $\sqrt{b} = i\sqrt{-b}$ and $\arctan ix = i \operatorname{arctanh} x$. The resulting resolutions on b for $N = 10^4$ are shown by plotting $(NI_i(b))^{-1/2}$ in figure 13.6. The curves goes to zero near the boundaries of valid values of b, as expected from equation (13.17) since the probability densities acquires ranges of $m_{\ell\ell}$ values with almost vanishing density there. This should of course not be taken to mean that one could expect fantastic measurements in these cases, but rather that the assumptions leading to the use of equation (13.17) breaks down near those boundaries. How to deal with boundaries will be discussed in more detail in section 13.2.3.

13.2.2 Accounting for Resolution Effects

No matter what method is used to extract b from the data it will be necessary to evaluate the probability density to observe a given value of $m_{\ell\ell}$ in signal events, *including* effects of the detector resolution. If the difference between the true and reconstructed value of the dilepton invariant mass in a given event, $\delta m_{\ell\ell}$, is distributed according to a Gaussian with parameters

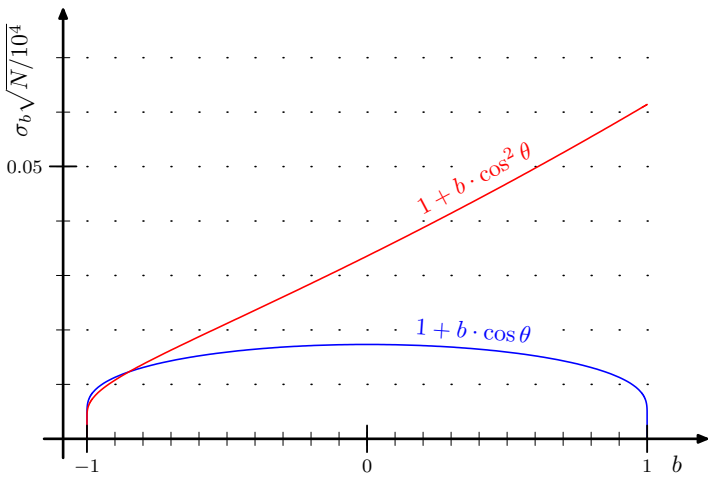

Figure 13.6: *Analytical predictions for decay parameter resolution.*

σ and μ, then the observed distribution of $m_{\ell\ell}$ will be given by:

$$f_{\text{edge}}^{\text{obs}}(m_{\ell\ell}; b, m_{\ell\ell}^{\max}, \sigma, \mu) = \frac{dP}{dm_{\ell\ell}}(b, m_{\ell\ell}, m_{\ell\ell}^{\max}) \otimes G(m_{\ell\ell}; \sigma, \mu) \tag{13.20}$$

$$= \int_0^{m_{\ell\ell}^{\max}} \frac{dP}{dm_{\ell\ell}}(b, m'_{\ell\ell}, m_{\ell\ell}^{\max}) \frac{1}{\sqrt{2\pi}\sigma} \exp\left[-\frac{(m_{\ell\ell} - m'_{\ell\ell} - \mu)^2}{2\sigma^2}\right] dm'_{\ell\ell} \tag{13.21}$$

$$= \int_0^1 \frac{dP}{d\hat{m}_{\ell\ell}}(b, \hat{m}'_{\ell\ell}) \frac{1}{\sqrt{2\pi}\hat{\sigma}} \exp\left[-\frac{(\hat{m}_{\ell\ell} - \hat{m}'_{\ell\ell} - \hat{\mu})^2}{2\hat{\sigma}^2}\right] d\hat{m}'_{\ell\ell} \tag{13.22}$$

with $\hat{\sigma} \equiv \sigma / m_{\ell\ell}^{\max}$ and $\hat{\mu} \equiv \mu / m_{\ell\ell}^{\max}$. If dealing with binned data one should furthermore integrate $f_{\text{edge}}^{\text{obs}}(m_{\ell\ell}; b, m_{\ell\ell}^{\max}, \sigma, \mu)$ over the bin range in order to minimise systematical effects. The expected number of events in a bin $[x_{\text{low}}, x_{\text{high}}[$ is thus,

$$N_{\text{edge}}^{\text{obs}}(x_{\text{low}}, x_{\text{high}}; b, m_{\ell\ell}^{\max}, \sigma, \mu) = \int_{x_{\text{low}}}^{x_{\text{high}}} f_{\text{edge}}^{\text{obs}}(m'_{\ell\ell}; b, m_{\ell\ell}^{\max}, \sigma, \mu)\, dm'_{\ell\ell} \tag{13.23}$$

The integrals in equations (13.22) and (13.23) must be evaluated either numerically or analytically. A numerical evaluation is easiest to implement but suffers from imprecision and long computation times. The latter is of course not a problem if one only has to fit the observed data a couple of times, but a full-blown analysis will invariably lead one to repeat the procedure on a very large number of simulated data sets. As an example, the results presented in section 16 used 3 CPU years of computation, but would have used several thousands of CPU years if they would have been based upon numerical evaluation of (13.22). However, for some parameter choices the analytical expressions can be such that their evaluation is numerically sensitive (such as when the expression requires cancellation between two very large terms). For this reason a hybrid approach was used for the present analyses, with the analytical expression being evaluated when possible and with a fall-back to a numerical Romberg integration (using Neville's algorithm) [P$^+$92] in those rare cases where that would give a more precise result. The validity of the approach was tested by systematically comparing the predicted densities against the results of *very* high statistics toy Monte Carlo experiments, and for broad ranges of the input parameters $(b, \hat{\sigma}, \hat{\mu})$.

The actual integrations of equations (13.22) and (13.23) are rather tedious, and only the essential details will be reviewed here. Inspection of (13.15) and (13.16) shows that (13.22) requires one to solve integrals of the form $u^k \exp[-u^2]$ for $k \geq 0$. They can be solved recursively using partial integration and the error function, $\mathrm{erf}(x) \equiv \frac{2}{\sqrt{\pi}} \int_0^x e^{-u^2} du$. It is, however, possible to find closed form expressions for the results. For odd k they are:

$$\int u^{2n+1} e^{-u^2} = -e^{-u^2} \frac{n!}{2} \sum_{i=0}^{n} \frac{u^{2i}}{i!}, \quad n \geq 0 \tag{13.24}$$

and the analogous result for even k are:

$$\int u^{2n} e^{-u^2} = \frac{F(n)}{2^{n+1}} \left[\sqrt{\pi} \mathrm{erf}(u) - e^{-u^2} \sum_{i=1}^{n} \frac{2^i u^{2i-1}}{F(i)} \right], \quad n \geq 0 \tag{13.25}$$

Where $F(j) \equiv \prod_{i=1}^{j} (2i-1) = 1 \cdot 3 \cdot 5 \cdots (2j-1)$ for $j \geq 1$ and $F(0) \equiv 1$. The validity of these results is straight-forward to check by differentiation of the right-hand sides.

It might turn out that a single Gaussian does not describe the resolution adequately, but in most cases a resolution function which is a sum of multiple Gaussians should suffice. It would be trivial to extend the formula and routines described here to account for that eventuality.

13.2.3 Extending the Validity Domains of the Decay Parameters

It is obvious that the decay angle distributions in equations (13.11)–(13.12) are not defined outside the ranges $b \in [-1, 1]$ and $b \in [-1, \infty[$ respectively, as the probability would otherwise attain negative values for some $\cos\theta$. However, when the "best value" of b is found from data by a generic heuristic search algorithm such as a straight down-hill minimisation, simulated annealing, etc., it is not rare for the algorithm to try to sample the distribution for these "forbidden" values of b. This might be solved, as is done in the program MINUIT [Jam94], by mapping the physics variable b to a dummy variable such as $b' = 2\pi^{-1} \tan(\pi b/2)$ which is unbounded while b remains in $[-1, 1]$. The problem with this method is twofold. First of all, when the fluctuations in the data really do prefer a value of b outside the well defined interval, the real variable will invariably end up pinned to the boundary, e.g. $b = -1.0000$, and a nonsensical reported error estimate. The second problem is that even when this is not the case, the mapping between b' and b is non-linear, meaning that the error estimates on b' returned by the chosen algorithm can not be related to those on b without great care.

A solution to this problem which ensures perfect linearity at sensible values of b while avoiding the problems of the naive approach is to formally extend the distributions in equations (13.11)–(13.12) outside the immediate valid domains for b. There is of course no unambiguous way to do that, so the choice here will simply be the most straight-forward: Keep the same analytic form of the distributions but force them to vanish where they would otherwise attain negative values. With appropriate normalisation factors, the resulting distributions thus

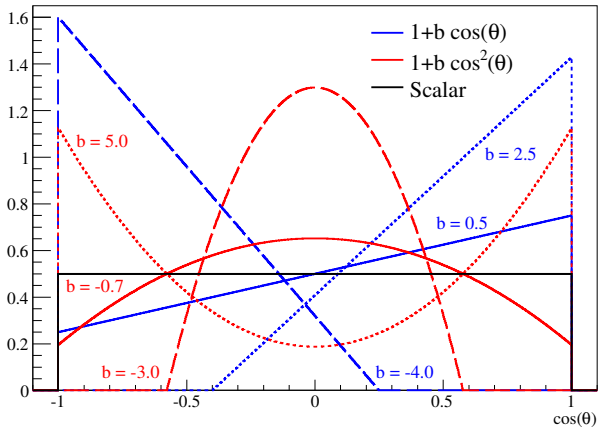

Figure 13.7: *Examples of the decay angle distributions covered by the two decay families* $1 + b \cdot \cos\theta$ *and* $1 + b \cdot \cos^2\theta$ *when the valid ranges of* b *are extended as discussed in the text.*

becomes:

$$\left.\frac{dP_1}{d\cos\theta}\right|_{b<-1} \equiv \frac{-2b}{(1-b)^2}\left[1 + b \cdot \cos\theta\right], \quad \cos\theta \in [-1, -\tfrac{1}{b}] \tag{13.26}$$

$$\left.\frac{dP_1}{d\cos\theta}\right|_{b>+1} \equiv \frac{2b}{(1+b)^2}\left[1 + b \cdot \cos\theta\right], \quad \cos\theta \in [-\tfrac{1}{b}, 1] \tag{13.27}$$

$$\left.\frac{dP_2}{d\cos\theta}\right|_{b<-1} \equiv \tfrac{3}{4}\sqrt{-b}\left[1 + b \cdot \cos^2\theta\right], \quad \cos\theta \in [-\sqrt{\tfrac{-1}{b}}, \sqrt{\tfrac{-1}{b}}] \tag{13.28}$$

Examples of the resulting distributions are shown for various values of b in figure 13.7. It is clear that while there is no plausible physics that would give a distribution proportional to e.g. $1 - 2\cos\theta$, that the general qualitative consequences of changes in values of b are preserved. For $1 + b \cdot \cos\theta$, higher and lower values of b still favours more forwards and backwards decays respectively, and for $1 + b \cdot \cos^2\theta$, lower values of b will result in more central decays.

14 SUSY and UED predictions in the case of a CMSSM-Like Spectrum

This section will introduce the relevant dilepton mass edges in the cases where either supersymmetry or universal extra dimensions represents the new physics at the TeV scale. This will be done in the context of mass spectra and branching ratios given by the CMSSM, but with investigations of the difference in decay angle distribution which would happen if the model was really UED. In this sense the study presented in this and the following sections can be said to by SUSY-centric. One could equally well have used UED to predict particle spectra and branching ratios and then have investigated the phenomenological consequences of changing certain decay angle distributions to follow SUSY predictions. The study presented here thus addresses the following question: If the new physics is really SUSY (in the CMSSM scenario), how well can one determine the decay parameter b and how likely is the method to be able to rule out UED at the LHC.

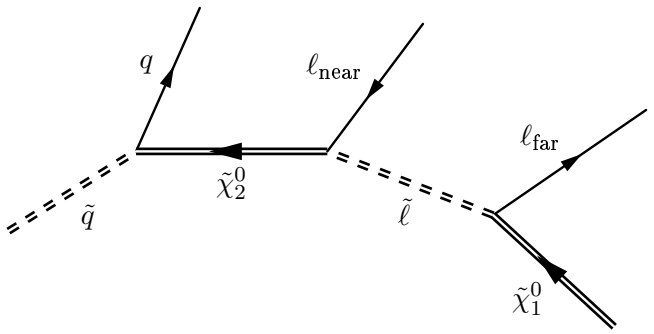

Figure 14.1: *A $\tilde{\chi}_2^0 \to \tilde{\ell}\ell_{near}$, $\tilde{\ell} \to \tilde{\chi}_1^0\ell_{far}$ cascade decay with an emitted dilepton pair surrounding a slepton.*

Section 14.1 will review the relevant decay chain in SUSY, including the relevant masses and the frequency with which it appears at the LHC. Section 14.2 will then discuss the changes to the decay angle distribution in the case of universal extra dimensions.

14.1 The Pure CMSSM Scenario

When the new physics is supersymmetry the only candidates for X_{lep} are the sleptons. When kinematically allowed the decay chain $X_{\text{hvy}} \to X_{\text{lep}}\ell_{\text{near}}$, $X_{\text{lep}} \to X_{\text{inv}}\ell_{\text{far}}$ will consequently become $\tilde{\chi}_2^0 \to \tilde{\ell}\ell_{\text{near}}$, $\tilde{\ell} \to \tilde{\chi}_1^0\ell_{\text{far}}$. Of course other possibilities exists such as substituting $\tilde{\chi}_2^0$ with $\tilde{\chi}_3^0$ or $\tilde{\chi}_4^0$ above, but being significantly heavier these should be a lot less frequent (to the point of being non-existing). The decay chain is illustrated in figure 14.1, where the dominant source of the $\tilde{\chi}_2^0$ at a hadron collider is also shown – namely the decay of a squark.

In principle there are six different slepton mass eigenstates, but in order to reconstruct $m_{\ell\ell}$ fully and to get as clean a signal as possible the stau states are not considered. Furthermore, as discussed in sections 5.4.2 and 5.4.4 it is generally expected that the superpartners of the right and left-handed muons and electrons are degenerate such that $m_{\tilde{\ell}_L} \equiv (m_{\tilde{e}_R} \approx m_{\tilde{\mu}_R})$ and $m_{\tilde{\ell}_R} \equiv (m_{\tilde{e}_L} \approx m_{\tilde{\mu}_L})$. In the CMSSM this is certainly the case. When observing a dilepton mass edge it is of course straight-forward to verify this assumption by checking whether the position of $m_{\ell\ell}^{\max}$ is the same for the $m_{e^+e^-}$ and $m_{\mu^+\mu^-}$ distributions.

The decay chain in figure 14.1 is often referred to as the "golden" SUSY decay chain. Not only does it have clear experimental signatures in the form of hard jets, two leptons and large amounts of missing p_t, but due to the associated sharp edge it offers a very precise constraint of SUSY masses through $m_{\ell\ell}^{\max}$. It has thus been the subject of several detailed studies (see for instance [ATL99b]). In figure 14.2 is shown an example of such an extracted signal. Note in particular that most of the background is due to SUSY itself, with only a rather small remaining fraction from SM processes (usually from $t\bar{t}$ events).

If both of the mass hierarchies $m_{\tilde{\chi}_2^0} > m_{\tilde{\ell}_R} > m_{\tilde{\chi}_1^0}$ and $m_{\tilde{\chi}_2^0} > m_{\tilde{\ell}_L} > m_{\tilde{\chi}_1^0}$ holds, the e^+e^-

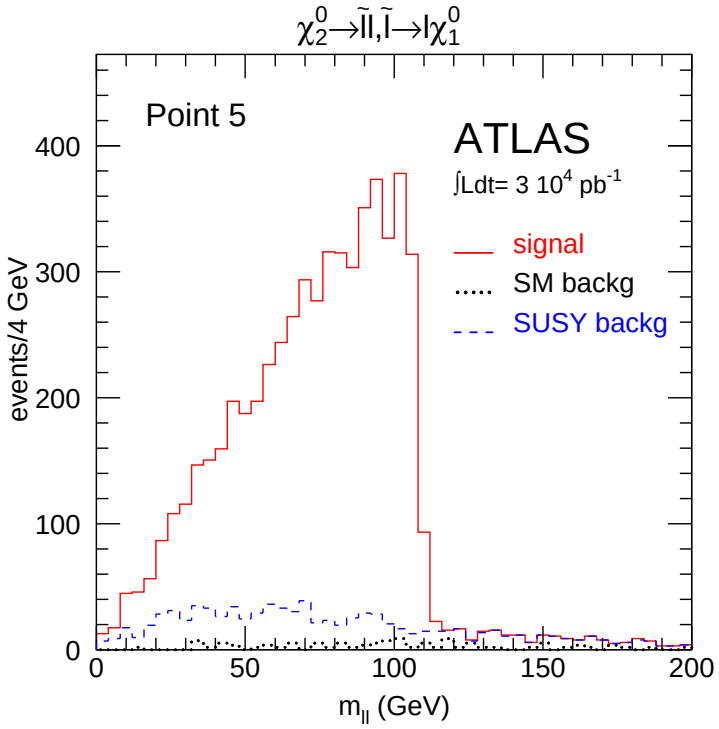

Figure 14.2: *Example of observed dilepton signal and backgrounds for* $\mathcal{L} = 30\ fb^{-1}$ *and CMSSM parameters* $(m_0, m_{1/2}) = (100\ GeV, 300\ GeV)$ *and* $(\tan\beta, A_0, \mathrm{sign}\mu) = (2.1, 300\ GeV, +1)$. *Source: [ATL99b].*

and $\mu^+\mu^-$ signal will have two distinct mass edges at two different $m_{\ell\ell}^{\max}$.

$$m_{\ell\ell}^{\max,1} \equiv \min\left(m_{\ell\ell}^{\max,\mathrm{R}},\, m_{\ell\ell}^{\max,\mathrm{L}}\right) \tag{14.1}$$

$$m_{\ell\ell}^{\max,2} \equiv \max\left(m_{\ell\ell}^{\max,\mathrm{R}},\, m_{\ell\ell}^{\max,\mathrm{L}}\right) \tag{14.2}$$

where

$$m_{\ell\ell}^{\max,\mathrm{R}} \equiv \frac{\sqrt{m_{\tilde{\chi}_2^0}^2 - m_{\tilde{\ell}_R}^2}\sqrt{m_{\tilde{\ell}_R}^2 - m_{\tilde{\chi}_1^0}^2}}{m_{\tilde{\ell}_R}} \tag{14.3}$$

$$m_{\ell\ell}^{\max,\mathrm{L}} \equiv \frac{\sqrt{m_{\tilde{\chi}_2^0}^2 - m_{\tilde{\ell}_L}^2}\sqrt{m_{\tilde{\ell}_L}^2 - m_{\tilde{\chi}_1^0}^2}}{m_{\tilde{\ell}_L}} \tag{14.4}$$

An example of the resulting $m_{\ell\ell}$ distribution in such a case with two distinct mass edges is shown in figure 14.3. When only one of the decay chains is kinematically allowed, it will be the one involving $\tilde{\ell}_R$, as the mass hierarchy $m_{\tilde{\ell}_R} < m_{\tilde{\ell}_L}$ is generally true in the CMSSM (cf. equation (5.12)).

The four relevant masses are shown in the $(m_0, m_{1/2})$–plane in figure 14.4 for reasonable values of $\tan\beta$, A_0 and $\mathrm{sign}\mu$. Also indicated in figures 14.4.c and 14.4.d are the regions where $m_{\tilde{\chi}_2^0} > m_{\tilde{\ell}_L}$ and $m_{\tilde{\chi}_2^0} > m_{\tilde{\ell}_R}$ respectively. The resulting values of the edge positions are shown

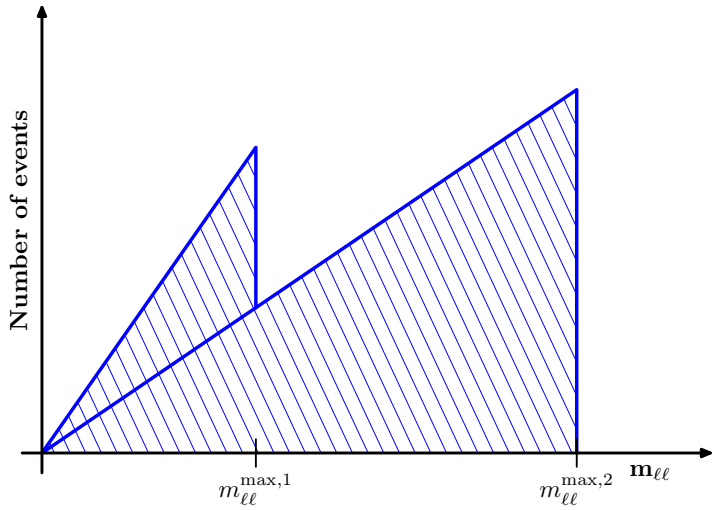

Figure 14.3: *Distribution of $m_{\ell\ell}$ in the case where two decay chains, leading to distinct dilepton mass edges, are possible. The shown distribution corresponds to the case where 15% of the decays contributes to the edge with the lower mass, and $m_{\ell\ell}^{max,1}/m_{\ell\ell}^{max,2} = 0.3$.*

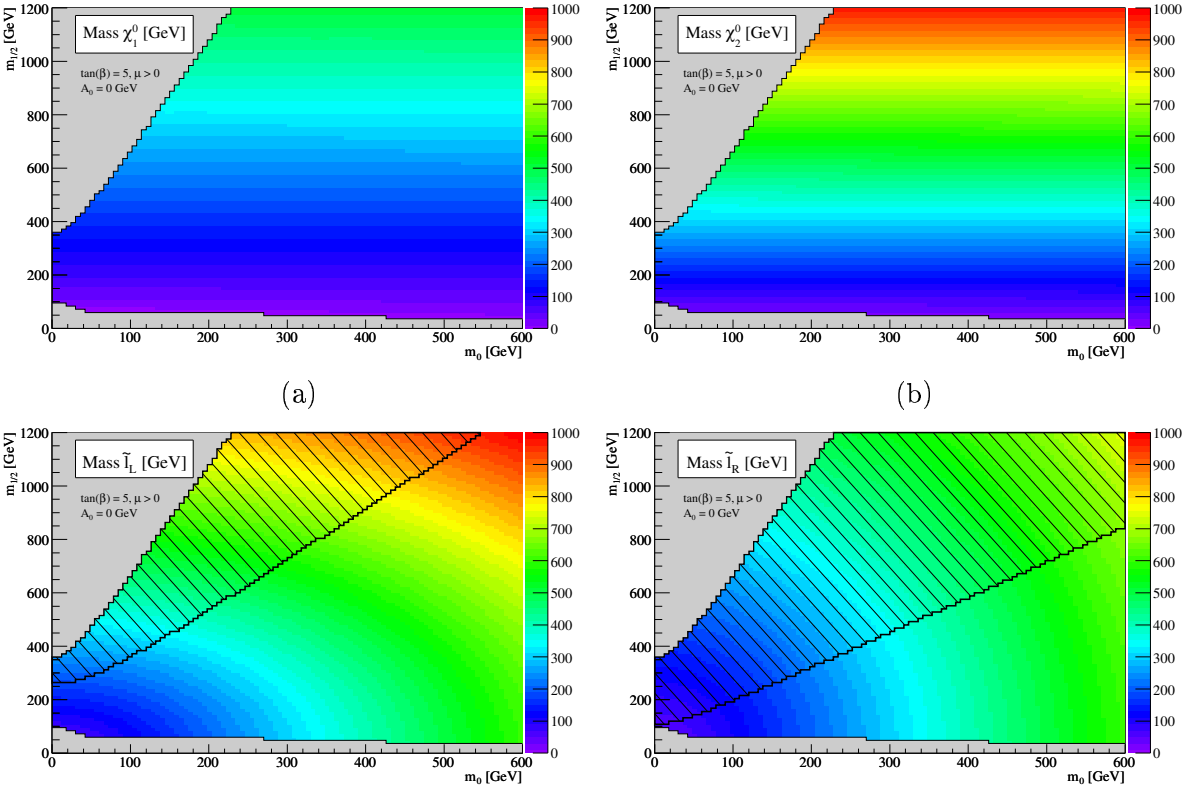

Figure 14.4: *Sparticle masses relevant for the dilepton mass edge kinematics as a function of $(m_0, m_{1/2})$, here shown for $(\tan\beta, A_0, \mathrm{sign}\mu) = (5, 0\ GeV, +1)$. Also indicated as hatched areas in the two lower plots are the regions where the decay of $\tilde{\chi}_2^0$ to the respective slepton is kinematically allowed.*

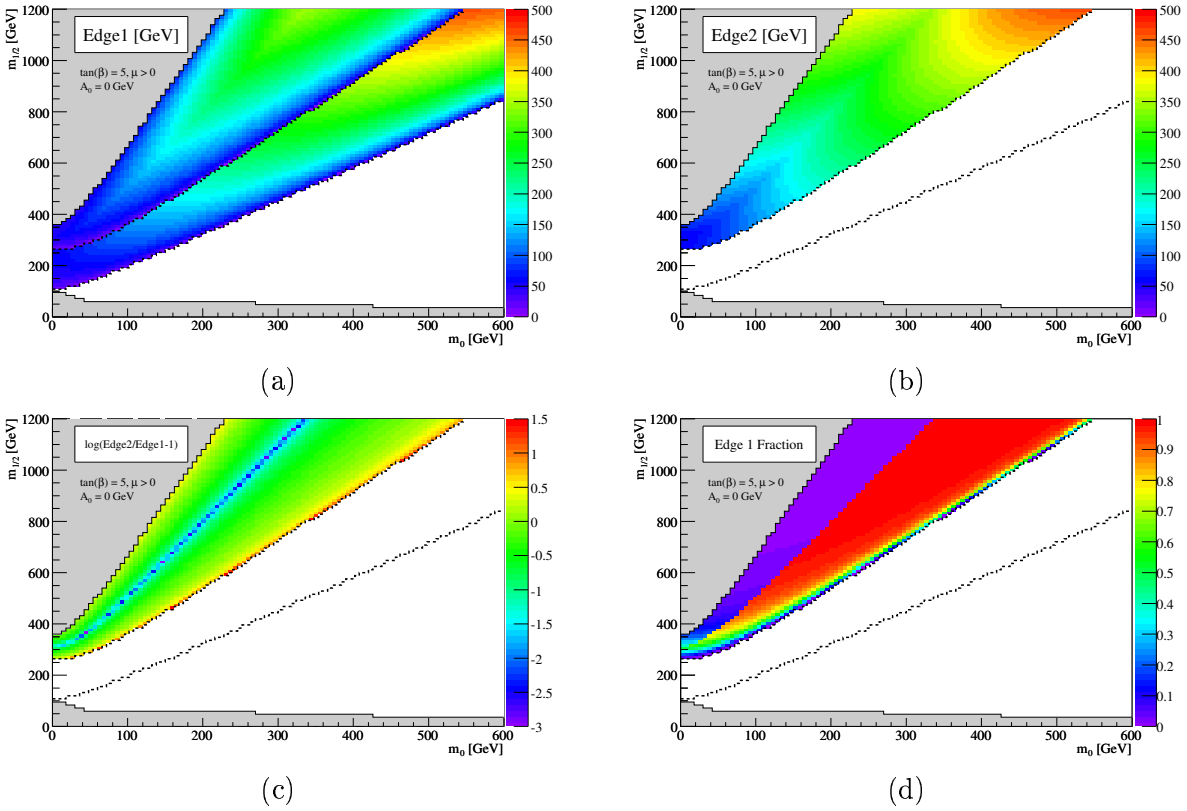

Figure 14.5: *Mass edge parameters as a function of* $(m_0, m_{1/2})$ *and for* $(\tan\beta, A_0, \text{sign}\mu) = (5, 0\ GeV, +1)$*, as discussed in the text.*

in figures 14.5.a and 14.5.b. "Edge1" and "Edge2" in the figure texts refers to $m_{\ell\ell}^{\max,1}$ and $m_{\ell\ell}^{\max,2}$, respectively, and the hatched lines separate the regions in the $(m_0, m_{1/2})$–plane where (from low to high values of $m_{1/2}$) respectively zero, one and two distinct slepton mass edges are kinematically allowed. Figure 14.5.c shows the ratio between the two edge positions in the case where two such are allowed. Notice how there is a thin region where the two edges are almost at the same position, $m_{\ell\ell}^{\max,1} \approx m_{\ell\ell}^{\max,2}$. It follows from equations (14.3)–(14.4) that this happens when $m_{\tilde{\chi}_1^0} m_{\tilde{\chi}_2^0} \approx m_{\tilde{\ell}_L} m_{\tilde{\ell}_R}$ (or $m_{\tilde{\ell}_L} \approx m_{\tilde{\ell}_R}$ obviously), and this configuration is rather unfortunate as it is impossible to fit two separate edges in such a scenario. However, it is still a prediction of SUSY models that such a "merged edge" should be entirely triangular, i.e. have $dP/dm_{\ell\ell} \propto m_{\ell\ell}$, and this can thus in principle still be used to distinguish SUSY from other models barring unfortunate cancellations of e.g. $\hat{m}_{\ell\ell}^3$ and $\hat{m}_{\ell\ell}^5$ terms between the two edges in such models. Figure 14.5.d shows, based on the $\tilde{\chi}_2^0$ branching ratios, the fraction of events going into the lower edge, f_1. For high values of m_0 more or less all of the decays of the $\tilde{\chi}_2^0$ to sleptons will go to $\tilde{\ell}_L$ (since the main component of the $\tilde{\chi}_2^0$ is the neutral wino, which does not couple to the $\tilde{\ell}_R$), resulting in either $f_1 \approx 1$ or $f_1 \approx 0$ depending on whether $m_{\ell\ell}^{\max,\mathrm{R}} = m_{\ell\ell}^{\max,1}$ or $m_{\ell\ell}^{\max,\mathrm{R}} = m_{\ell\ell}^{\max,2}$ respectively. The sharp cross-over between the two regions happens along the line where $m_{\ell\ell}^{\max,1} = m_{\ell\ell}^{\max,2}$.

Regarding the question of how numerous the relevant decay chain will be at the LHC, it will mainly be determined by the branching ratios of the $\tilde{\chi}_2^0$ and the general SUSY production cross section. This is because $\tilde{\chi}_2^0$'s will appear in at least a third of all SUSY events, as shown

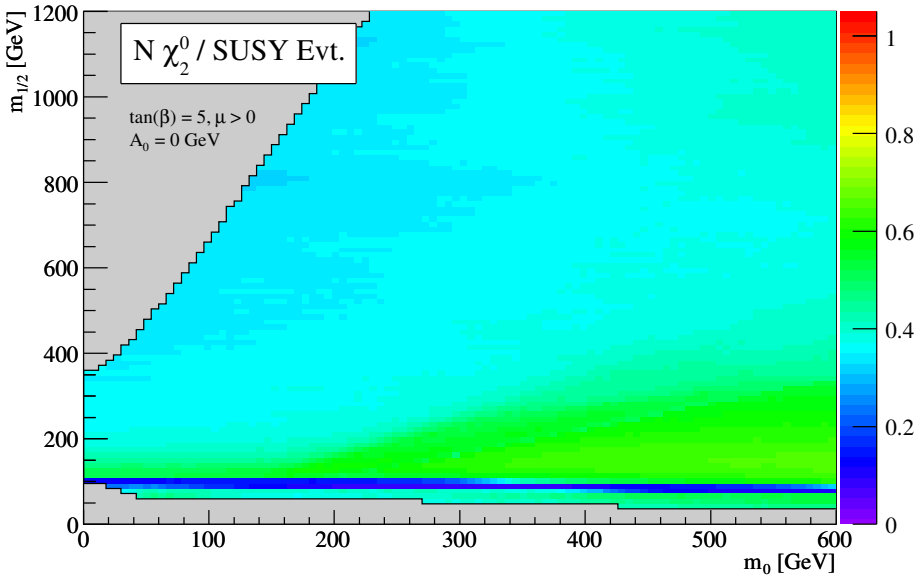

Figure 14.6: *Number of produced $\tilde{\chi}_2^0$'s per SUSY event at LHC. According to HERWIG and ISASUGRA.*

in figure 14.6. In principle the slepton branching ratio could also be important, and one might for instance wonder whether decays such as $\tilde{e}^{\pm} \to \tilde{\chi}_1^{+}\nu_e$ could play a role. However, when $m_{\tilde{\ell}} < m_{\tilde{\chi}_2^0}$ it will in the CMSSM also be true that $m_{\tilde{\ell}} < m_{\tilde{\chi}_1^{\pm}}$, thus ruling out such decays (cf. equation (5.12)). The branching ratio of $\tilde{\chi}_2^0$ to sleptons is shown in figure 14.7, and other dominant decays are indicated. Generally, the main contenders that takes branching ratio from $\tilde{\chi}_2^0 \to \tilde{e}e$ and $\tilde{\chi}_2^0 \to \tilde{\mu}\mu$ in regions where $m_{\tilde{\ell}} < m_{\tilde{\chi}_2^0}$ are $\tilde{\chi}_2^0 \to \tilde{\nu}\nu$, $\tilde{\chi}_2^0 \to \tilde{\chi}_1^0 h$ and obviously $\tilde{\chi}_2^0 \to \tilde{\tau}_1\tau$. The latter especially becomes a problem at very high values of tanβ such as tan$\beta = 30$.

Taking SUSY production cross sections and branching ratios into account, the number of decay chains $\tilde{\chi}_2^0 \to \tilde{\ell}\ell_{\text{near}}$, $\tilde{\ell} \to \tilde{\chi}_1^0\ell_{\text{far}}$ occurring at an integrated luminosity of 100 fb^{-1} (corresponding to one year of running at full LHC luminosity) are shown in figure 14.8.

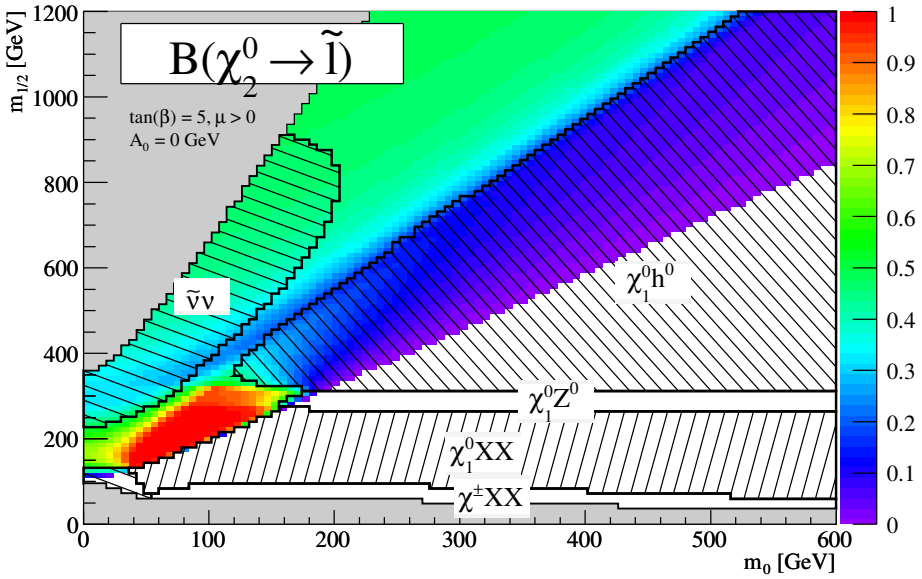

Figure 14.7: *Dominant $\tilde{\chi}_2^0$ branching ratios according to ISASUGRA.*

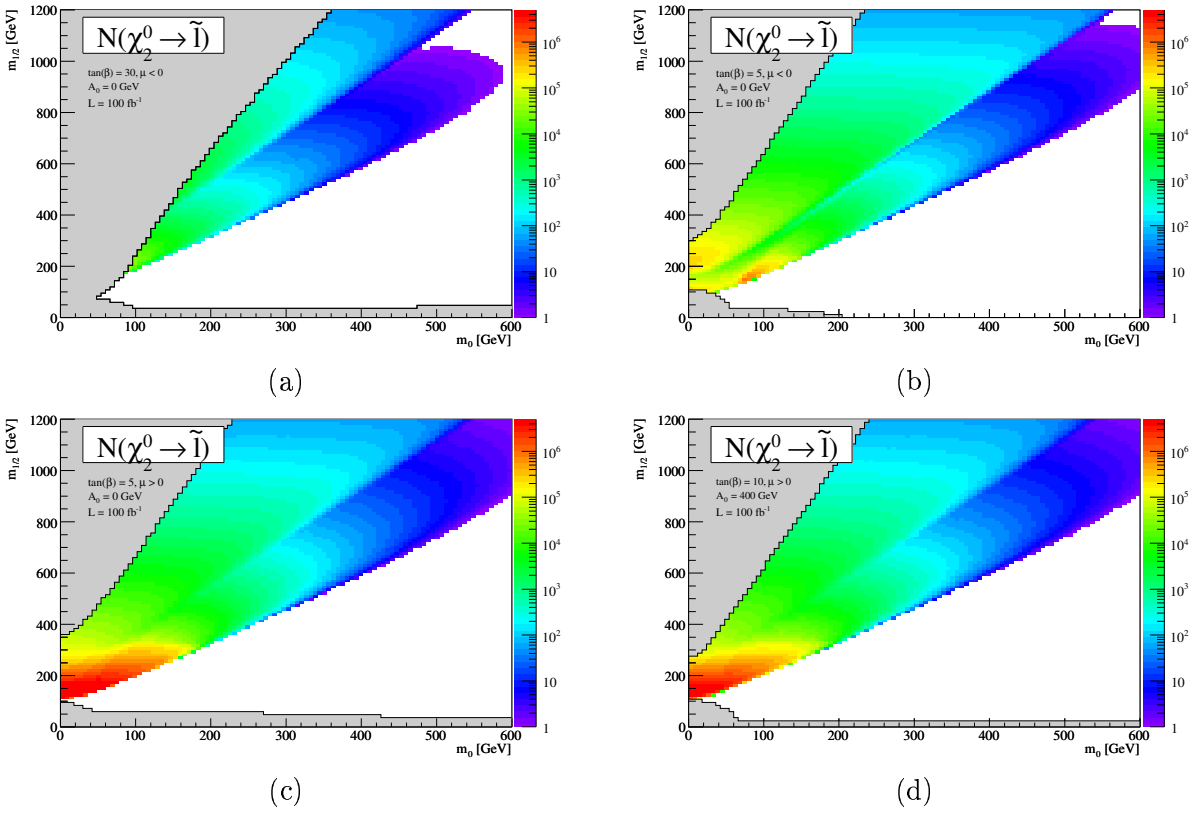

Figure 14.8: *Number of produced $\tilde{\chi}_2^0$ particles decaying to $\tilde{\ell}\ell_{near}$ followed by $\tilde{\ell} \to \tilde{\chi}_1^0 \ell_{far}$ at LHC for an integrated luminosity of 100 fb^{-1}.*

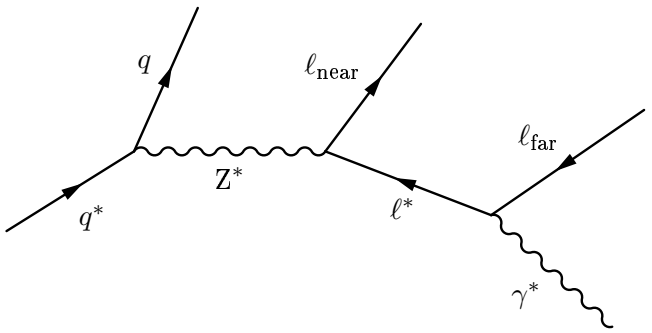

Figure 14.9: *Equivalent cascade decay in a model with universal extra dimensions. Particles marked with an asterisk are the first level Kaluza-Klein excitations.*

14.2 Expectations for Models with Universal Extra Dimensions

The analogy of the SUSY decay chain of figure 14.1 in models of universal extra dimensions is the decay chain shown in figure 14.9 where the new particles are now Kaluza-Klein excitations with the same spin as their Standard Model partners. Again the new particles with electron and muon lepton numbers, e^* and μ^*, will be almost degenerate in mass, so the dilepton mass edge phenomenology will be very similar to that of SUSY. As the ℓ^* is not a scalar the dilepton invariant mass distribution will no longer be triangular. In [SW05] the tree-level prediction for the invariant mass distribution of the two emitted leptons is found to be:

$$\frac{dP}{d\hat{m}_{\ell\ell}} = \frac{4\hat{m}_{\ell\ell}}{(2+y)(1+2z)}\left[y + 4z + (2-y)(1-2z)\hat{m}_{\ell\ell}^2\right], \quad 0 \leq \hat{m}_{\ell\ell} \leq 1 \tag{14.5}$$

Where y and z are defined as:

$$y \equiv \left(\frac{m_{\text{lep}}}{m_2}\right)^2 \tag{14.6}$$

$$z \equiv \left(\frac{m_1}{m_{\text{lep}}}\right)^2 \tag{14.7}$$

and the masses $(m_1, m_2, m_{\text{lep}})$ represents either $(m_{\tilde{\chi}_1^0}, m_{\tilde{\chi}_2^0}, m_{\tilde{\ell}})$ or $(m_{\gamma^*}, m_{Z^*}, m_{\ell^*})$ depending on whether the fields are interpreted as being SUSY or UED particles, respectively. After some manipulations equation (14.5) can be rewritten as:

$$\frac{dP}{d\hat{m}_{\ell\ell}} = 2\hat{m}_{\ell\ell}\left[1 + \frac{4z + y - 2yz - 2}{(2+y)(1+2z)}\left(1 - 2\hat{m}_{\ell\ell}^2\right)\right], \quad 0 \leq \hat{m}_{\ell\ell} \leq 1 \tag{14.8}$$

And comparison with equation (13.15) thus reveals that the corresponding decay angle distribution of ℓ^* is

$$\frac{dP}{d\cos\theta} = \frac{1}{2}\left[1 + b_{\text{UED}} \cdot \cos\theta\right] \tag{14.9}$$

with

$$b_{\text{UED}} = b_{\text{UED}}(y,z) = \frac{4z + y - 2yz - 2}{(2+y)(1+2z)} \tag{14.10}$$

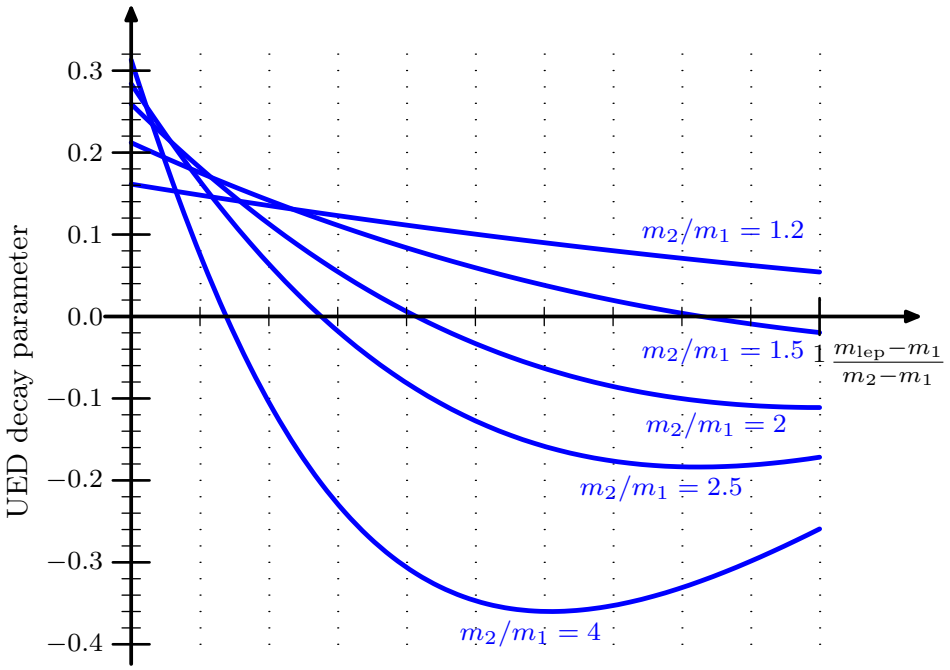

Figure 14.10: *Predicted values of b_{UED} for different values of m_2/m_1 as a function of $(m_{lep}-m_1)/(m_2-m_1)$, i.e. the relative position of m_{lep} in the interval $[m_1, m_2]$.*

The fact that (14.5) can be rewritten as (14.8) is of course not trivial, as it requires coefficients in constant terms and terms proportional to $\hat{m}_{\ell\ell}^2$ to match up. It was, however, expected since equation (14.5) indeed has its origin in particle spin effects in the decay chain.

It is clear that $b_{\rm UED}$ only vanishes when $z = 1/2$, i.e. when $m_{\rm lep} = \sqrt{2}m_1$. This is of course not possible when $m_2 < \sqrt{2}m_1$, which is a likely scenario in an UED-like mass spectrum. In a CMSSM-like mass spectrum on the other hand, $m_2 \approx 2m_1$, so here $b_{\rm UED}$ can vanish. The value of $b_{\rm UED}$ as a function of different ratios of m_2/m_1 and of different positions of $m_{\rm lep}$ in the interval $[m_1, m_2]$ is shown in figure 14.10 and the specific values associated with the two edges in a CMSSM like spectrum are shown in figures 14.11 and 14.12. Viewed in connection with figures 13.6 and 14.8 one would already at this stage expect that a measurement of decay parameters in some regions should be able to distinguish between SUSY and UED. But it is also clear that it will be impossible to do so in the regions where $b_{\rm UED} \approx 0$. Of course, in such a case a measurement might still exclude other models of new physics not considered here.

Finally, the situation where two decay chains gives "merged edges", i.e. two edges with indistinguishable $m_{\ell\ell}^{\rm max}$, can be investigated in the case where the decay angle distribution is given by equations (14.9) and (14.10). As already noted, merged edges happens when either $m_{\rm lep1} = m_{\rm lep2}$ or $m_{\rm lep1}m_{\rm lep2} = m_1 m_2$. In the first case the $m_{\ell\ell}$ distribution will of course be indistinguishable from the edge resulting from one of the decay channel (just as the $\tilde{\mu}_L$ and $\tilde{e}_L$ edges are already merged). In the second case $y_1 = z_2$ and $z_1 = y_2$. The combined observed $m_{\ell\ell}$ mass distribution will in that case thus be given by

$$\frac{dP^{\rm comb}}{d\hat{m}_{\ell\ell}} = 2\hat{m}_{\ell\ell}\left[1 + b_{\rm UED}^{\rm comb}\left(1 - 2\hat{m}_{\ell\ell}^2\right)\right], \quad 0 \leq \hat{m}_{\ell\ell} \leq 1 \tag{14.11}$$

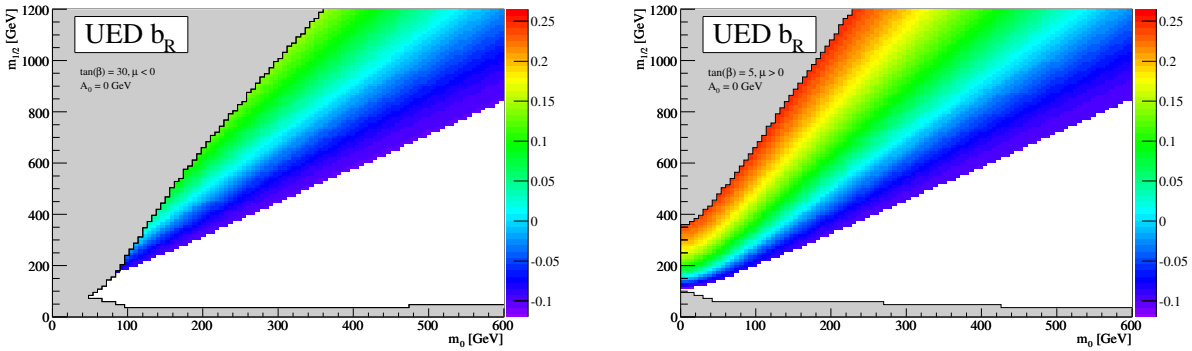

(a) (b)

Figure 14.11: *Predicted values of the decay parameter b_{UED} corresponding to decays through $\tilde{\ell}_R$ for different CMSSM-like particle spectra.*

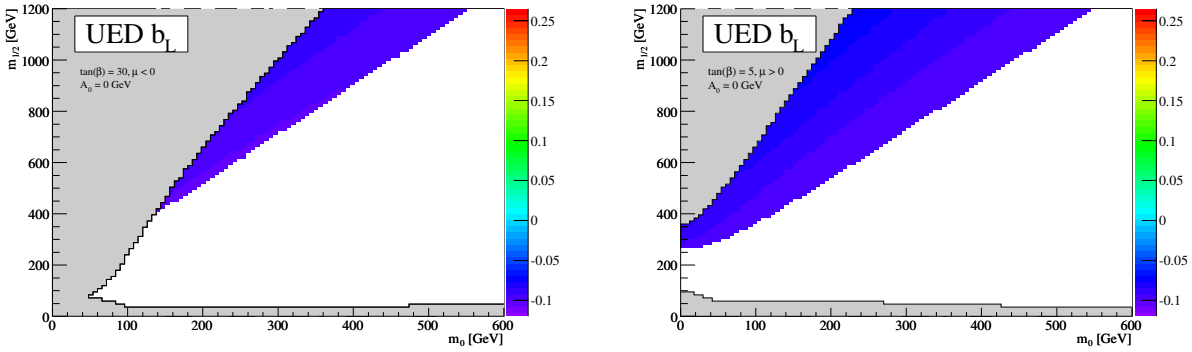

(a) (b)

Figure 14.12: *Predicted values of the decay parameter b_{UED} corresponding to decays through $\tilde{\ell}_L$ for different CMSSM-like particle spectra.*

Where

$$b_{\mathrm{UED}}^{\mathrm{comb}} \equiv f_1\, b_{\mathrm{UED}}(y_1, z_1) + (1 - f_1)\, b_{\mathrm{UED}}(y_2, z_2) \tag{14.12}$$

$$= f_1\, b_{\mathrm{UED}}(y_1, z_1) + (1 - f_1)\, b_{\mathrm{UED}}(z_1, y_1) \tag{14.13}$$

With f_1 being the relative fraction of decays happening through the first decay channel. When $f_1 \approx 0$ or $f_1 \approx 1$ the edge is really just one regular edge rather than a merged one. The expected value of $b_{\mathrm{UED}}^{\mathrm{comb}}$ for $f_1 = 1/2$ is plotted in figure 14.13, again for various values of m_2/m_1 and for different positions of m_{lep} in the interval $[m_1, m_2]$. Due to the symmetry of the problem, it does not matter whether m_{lep1} or m_{lep2} is used in the formula.

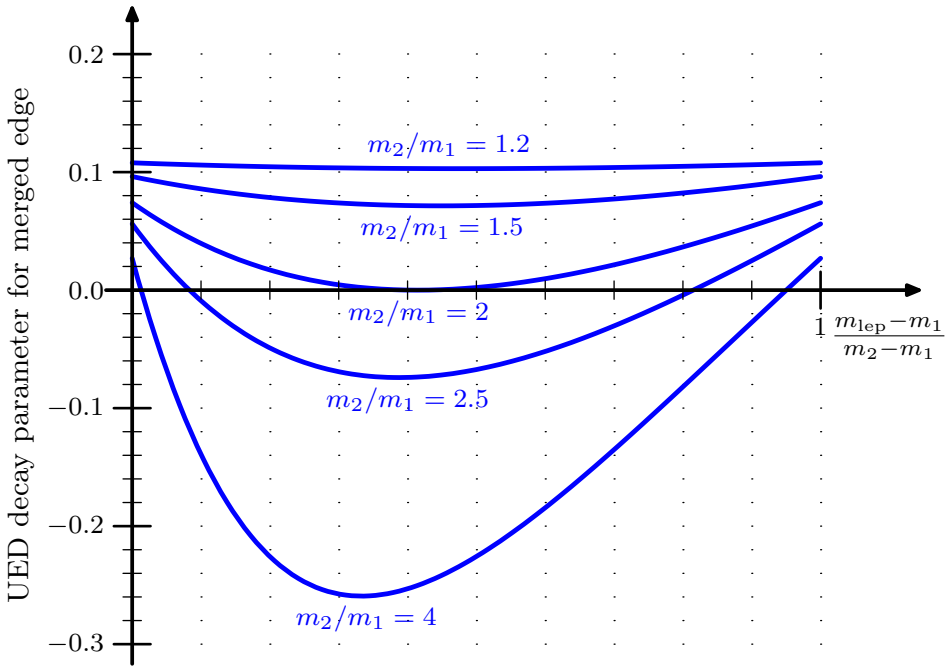

Figure 14.13: *Predicted values of the combined decay parameter, b_{UED}^{comb} of a merged edge, when both edges contribute with equal statistics (i.e. $f_1 = 1/2$. The values are shown for different ratios of m_2/m_1 as a function of $(m_{lep} - m_1)/(m_2 - m_1)$, i.e. the relative position of m_{lep} in the interval $[m_1, m_2]$ (m_{lep} can be either m_{lep1} or m_{lep2}).*

15 SUSY Generator Level Studies

Having thus discussed the phenomenology of dilepton mass edges and the expectations of SUSY versus UED, this section now aims to investigate whether or not these mass edges will be robust enough to preserve the information contained in their shape after the experimental procedure of event selection and signal treatment. Numerous other analyses such as the ones accompanying figure 14.2 and others in [ATL99b] have already concluded that the SM backgrounds can in most cases be made practically negligible by demanding two hard leptons, several hard jets and a high level of missing transverse momentum or energy. It is not the intent of the present thesis to redo those analyses.

It is also clear that the experimental resolution on the dilepton momenta, and thus on $m_{\ell\ell}$, should not completely wash out the features of a given mass edge. This is because the scale of the features of the mass edges is given by $m_{\ell\ell}^{\max}$ which is $\mathcal{O}(100\text{ GeV})$, while the resolution effects are going to be no more than $\mathcal{O}(1\text{ GeV})$.

The main remaining concerns are whether event selection criteria and presence of background will distort the shape of the mass edges to a degree that prevents the reliable extraction of information about the decay angle distribution from them. This will be investigated here by generating SUSY events with different slepton decay angle distributions, and subsequently impose reasonable event selection cuts (at the generator level) and test the feasibility of the method and the impacts of those cuts. The performed generation is based upon the SUSY generation code set up by the ATLAS Monte Carlo group. This uses the HERWIG generator in the ATLAS ATHENA framework.[1] The generated event record of events containing sleptons is then subsequently changed by a custom made algorithm which overwrites the momenta of the slepton daughters based upon an appropriately distributed value of $\cos\theta$. Such a post-generation "fix" of the event record is of course only possible since the slepton daughters are colour-neutral and thus colour-disconnected from the rest of the event.

15.1 Event Selection and Background Treatment

Signal candidate events are required to have total missing transverse momentum of at least 180 GeV (assuming all particles to be observed unless they are neutrinos, $\tilde{\chi}_1^0$'s or fall outside the detector acceptance at $\eta > 5$). They must have 4 jets at $\eta < 2.5$ – three with a transverse momentum greater than 50 GeV and one greater than 250 GeV.[2] Finally, they must have at least two leptons (muons or electrons) at $\eta < 2.5$ and transverse momenta greater than 30 GeV and 15 GeV respectively. It is important to note that these cuts were chosen simply by inspiration from existing analyses and as such no tuning for the present analyses were performed. This approach is deliberate so as not to arrive at unreasonable expectations due to cut fine-tuning. Nevertheless, strategies for optimising the cuts will be discussed along the way. An additional important point is that all of these requirements are well above the foreseen trigger thresholds, meaning that the events can be expected to be triggered with large redundancy and with a resulting high efficiency.

Two different SUSY spectra were analysed, both singled out as benchmark points for

[1] For reference the version of HERWIG was 6.510 and JIMMY [BFS96] was used for the underlying event. Parton distribution functions used the so-called "CTEQ6L" versions [L$^+$95]. Finally, the top mass was set to 173 GeV.

[2] For reference the jet algorithm was a cone algorithm with cone size of 0.4 (based entirely on code from the ATLAS generator level jet filter by I. Hinchliffe).

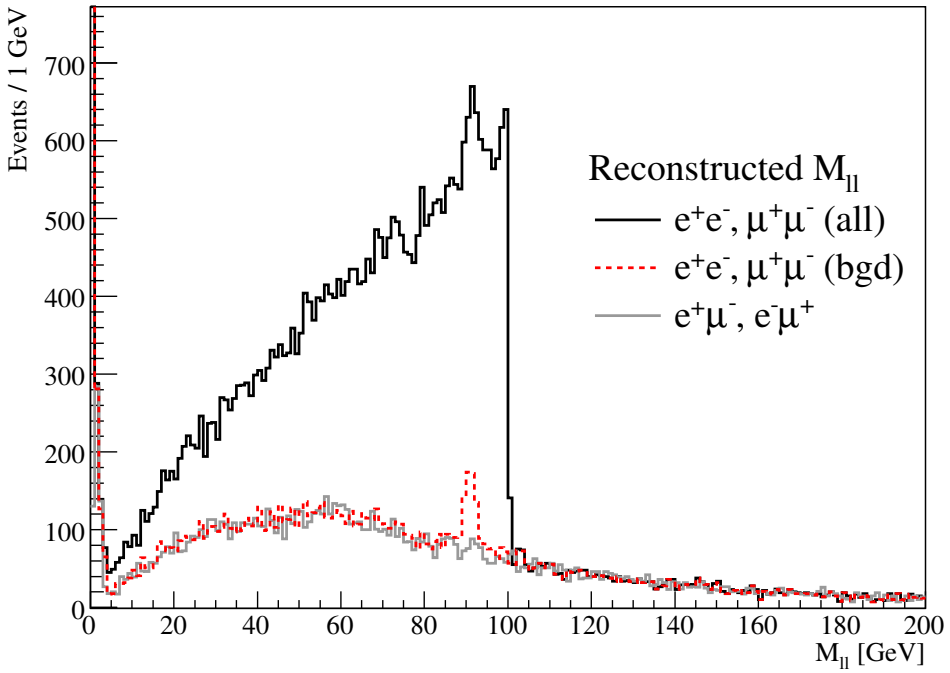

Figure 15.1: *Generator level dilepton samples after event selection cuts at the ATLAS SU3 point for 300 fb^{-1}.*

ATLAS. At the first, the "SU3" point, the relevant sparticle masses are:

$$\begin{aligned} m_{\tilde{\chi}_2^0} &= 218.6 \text{ GeV}, \quad m_{\tilde{\chi}_1^0} = 117.9 \text{ GeV} \\ m_{\tilde{\ell}_L} &= m_{\tilde{e}_L} = m_{\tilde{\mu}_L} = 230.5 \text{ GeV} \\ m_{\tilde{\ell}_R} &= m_{\tilde{e}_R} = m_{\tilde{\mu}_R} = 155.5 \text{ GeV} \end{aligned} \tag{15.1}$$

This means that only the decay $\tilde{\chi}_2^0 \to \tilde{\ell}_R^{\pm}\ell^{\mp}$ is open, leading to one dilepton mass edge with $m_{\ell\ell}^{\text{max,R}} = 100.2$ GeV. According to equation (14.10), an analogous UED mass spectrum would have a decay parameter $b_{\text{UED}} \approx 0.0415$, which is not very different from the SUSY case.

The distributions of the invariant mass of the two leptons with highest p_t in the SUSY events surviving these cuts are shown in figure 15.1 for the case of a scalar slepton. Only the samples where the two leptons have opposite charge are shown, meaning the sample of signal candidates (e^+e^- and $\mu^+\mu^-$ events) and the $e^{\pm}\mu^{\mp}$ control sample. The background level in the $e^+e^-/\mu^+\mu^-$ sample is also shown based upon Monte Carlo truth information. Several conclusions can be drawn from this plot. First of all it is seen that the $e^{\pm}\mu^{\mp}$ control sample can be used to very accurately estimate the background level in the $e^+e^-/\mu^+\mu^-$ signal sample, apart from a small contamination of events with $Z \to \ell^+\ell^-$ events. Secondly, a large background peak is present at very low masses which is mainly due to background leptons from photons, π^0's, and hadrons with c and b quarks. This peak is removed by removing all lepton pairs combining invariant masses less than 4 GeV from the events. Obviously this affects the signal events at $m_{\ell\ell} < 4$ GeV, but due to the phase space factor of $m_{\ell\ell}$ in equation (13.14) the total effect on the signal sample will be negligible, as the distribution should tend to zero at those values in any case.

It is not a coincidence that the $e^{\pm}\mu^{\mp}$ control sample so well approximates the component of the $e^+e^-/\mu^+\mu^-$ sample that is not due to dileptons being produced correlated in both charge and flavour (such as when they are produced by the signal decay chain or Z decays). This background will consist mainly of electrons and muons from sources that produce dileptons of opposite charge but uncorrelated in flavour, such as for instance the uncorrelated decays of the two taus in the stau version of the decay chain in figure 14.1.[1] Branching ratios to electrons and muons will be similar to a high degree since both masses are negligible at the relevant mass scales, and the reconstruction efficiencies should in both cases be close to 1. If the combined effect of branching ratios and reconstruction efficiencies at a given $m_{\ell\ell}$ is that on average $(1-\epsilon)$ muons meeting the selection criteria are observed for every 1 electron, then the relative number of events expected from such charge-correlated sources will satisfy:

$$\frac{N(e^+e^-)+N(\mu^+\mu^-)-N(e^{\pm}\mu^{\mp})}{N(e^+e^-)+N(\mu^+\mu^-)} = \frac{1^2+(1-\epsilon)^2-2\cdot 1\cdot(1-\epsilon)}{1^2+(1-\epsilon)^2} = \frac{\epsilon^2}{2}+\mathcal{O}\left(\epsilon^3\right) \qquad (15.2)$$

Where not only the $\mathcal{O}(1)$ terms cancelled as expected, but also those of $\mathcal{O}(\epsilon)$ disappeared which is very fortunate. Another smaller source of dilepton events surviving the lepton cuts are events where the source of the two leptons are uncorrelated in both charge and flavour. Typical sources of these leptons are decays of W bosons (themselves coming from decays of top quarks or $\tilde{\chi}^{\pm}$), decays of taus produced in $\tilde{\tau}\to\tilde{\chi}_1^0\tau$ or $\tilde{\chi}_2^0\to\tilde{\tau}\tau$, decays from B mesons, misidentified hadrons, or simply a signal event where one lepton did not satisfy the selection criteria. Again, assuming that the combined effects of production rates and reconstruction efficiencies are such that $(1-\epsilon')$ muons from such sources are observed for each electron, then the relative number of events expected from such uncorrelated (combinatorial) sources will satisfy:

$$\frac{N(e^+e^-)+N(\mu^+\mu^-)-N(e^{\pm}\mu^{\mp})}{N(e^+e^-)+N(\mu^+\mu^-)} = \frac{2\cdot 1^2+2\cdot(1-\epsilon')^2-4\cdot 1\cdot(1-\epsilon')}{2\cdot 1^2+2\cdot(1-\epsilon')^2} = \frac{(\epsilon')^2}{2}+\mathcal{O}\left((\epsilon')^3\right) \qquad (15.3)$$

Once again the very fortunate cancellation of $\mathcal{O}(\epsilon')$ occurred. The effect of this is that the background subtraction method will be quite robust even in the face of non-negligible electron-muon asymmetries. If, for instance, the rate of observed background electrons and muons is asymmetric by 5%, the relative difference between the number of background events and the number of events in the $e^{\pm}\mu^{\mp}$ control sample will be just 0.1%.

As an extra cross-check that $(\epsilon')^2$ is low, one could use the same sign samples $e^{\pm}e^{\pm}$, $\mu^{\pm}\mu^{\pm}$ and $e^{\pm}\mu^{\pm}$ to get a handle on ϵ', as e.g.:

$$\frac{N(e^{\pm}e^{\pm})-N(\mu^{\pm}\mu^{\pm})}{N(e^{\pm}e^{\pm})+N(\mu^{\pm}\mu^{\pm})} = \frac{2\cdot 1^2-2\cdot(1-\epsilon')^2}{2\cdot 1^2+2\cdot(1-\epsilon')^2} = \epsilon'+\mathcal{O}\left((\epsilon')^2\right) \qquad (15.4)$$

A similar check on ϵ can subsequently be done, but only outside the region of signal (and Z) events (i.e. in the area above ~ 100 GeV in figure 15.1).

[1] Since taus are relatively long lived it might be interesting to investigate whether the tau background could be reduced by a requirement that the backwards extrapolated tracks of the lepton candidates should both originate in the vertex of one of the jets in the event. The main obstacle with this idea is obviously that the taus are in general highly boosted, with its daughter leptons consequently pointing back at the vertices where the taus originated.

As long as the differences in production rates and reconstruction efficiencies between electrons and muons in the background samples are reasonable. the $e^{\pm}\mu^{\mp}$ control sample can thus be expected to provide a very accurate handle on the background in the signal sample up to the level of statistical fluctuations. By using this information it can thus be assured that the presence of background events should never introduce systematic effects in a decay parameter measurement, but might reduce the statistical sensitivity.

The most straight-forward way to proceed when analysing the data, and the one pursued here, is simply to subtract the binned distribution of $e^{\pm}\mu^{\mp}$ from that of $e^{\pm}e^{\mp} + \mu^{\pm}\mu^{\mp}$, and subsequently perform a chi-square fit to the resulting histogram. Apart from being straightforward this method has the advantage that it requires no prior assumptions regarding the shape of the background distributions. The downside is that the statistical fluctuations in the $e^{\pm}\mu^{\mp}$ sample are imposed on the signal sample. If N_S and N_B, respectively, denotes the average expected number of signal and background events in a given bin, then the resulting error on the contents of the corresponding bin in the background subtracted histogram will be:

$$\sigma = \sqrt{\sqrt{N_S + N_B}^2 + \sqrt{N_B}^2} = \sqrt{N_S}\sqrt{1 + 2\frac{N_B}{N_S}} \tag{15.5}$$

This means that observations based on N_S signal events in the presence of N_B background events will be statistically comparable to those performed without backgrounds on $N_{\text{eff}} = N_S/(1 + 2N_B/N_S)$ events. The statistics power of the sample will thus be highest if the event selection criteria are tuned so that

$$N_{\text{eff}}^{1/2} = \frac{N_S}{\sqrt{N_S + 2N_B}} \tag{15.6}$$

is optimised rather than simply N_S.

Instead of simply using the control sample to directly subtract the background bin for bin, the background shape might instead be parametrised by an (empirical) function based upon a few free parameters. These parameters could be determined with a fit to the $e^{\pm}\mu^{\mp}$ sample, and the background level in the signal sample could thus be taken into account with negligible impact due to fluctuations in this control sample. In such a case the resulting error on the estimated number of signal events in a given bin would reduce to:

$$\sigma \approx \sqrt{\sqrt{N_S + N_B}^2 + \sqrt{0}^2} = \sqrt{N_S}\sqrt{1 + \frac{N_B}{N_S}} \tag{15.7}$$

Meaning that the sensitivity of measured edge parameters will increase by about a factor of $\sqrt{(1 + 2\frac{N_B}{N_S})/(1 + \frac{N_B}{N_S})}$ when the background is fitted instead of merely subtracted (and that the event selection criteria should be tuned in order to optimise $N_S/\sqrt{N_S + N_B}$). Furthermore, such a treatment of backgrounds would allow the physics parameters to be extracted through a simultaneous unbinned log-likelihood fit to both samples, thus preserving the full information about the individual $m_{\ell\ell}$ measurements. It would also mean that the fit could allow each individual $m_{\ell\ell}$ value to enter with its own error estimate (based on e.g. the track fit parameters) instead of simply using a global σ in the resolution function. Thus this unbinned method is clearly the preferred way to go in a full-fletched analysis, but as already stated, the results presented here will be based upon the simpler binned chi-square fit to a background-subtracted distribution.

SM backgrounds are not included in figure 15.1, but they should be low due to the selection cuts, and even if not entirely negligible they will automatically be accounted for by the $e^{\pm}\mu^{\mp}$ control sample just as backgrounds from SUSY events are. SM background must originate from a source which generates hard leptons of opposite sign, jets and missing p_t. This essentially only leaves $t\bar{t}$ events, where the two W's from the $t \to b\mathrm{W}$ decays each decay into leptons, i.e. $t\bar{t} \to b\bar{b}\nu\bar{\nu}\ell^+\ell^-$.

The background due to $\mathrm{Z} \to \ell^+\ell^-$ decays is not approximated by the $e^{\pm}\mu^{\mp}$ sample, and it is necessary to include a Z-peak directly in the fit, thus introducing one additional free parameter, f_{Z}, for the fraction of events in this peak. As the Z bosons are produced on mass shell, from the decays of neutralinos, charginos and stops, the actual shape of this peak will be given by [PDG06]:

$$\frac{dP}{dm_{\ell\ell}} = N_{\mathrm{Z}} \frac{m_{\ell\ell}^2 \Gamma_{\mathrm{Z}}}{\left(m_{\ell\ell}^2 - m_{\mathrm{Z}}^2\right)^2 + \left(\Gamma_{\mathrm{Z}}^2/m_{\mathrm{Z}}^2\right) m_{\ell\ell}^4} \tag{15.8}$$

Where N_{Z} is a normalisation constant. This distribution must of course also be convoluted with the resolution function when fitting to data, which is done numerically. Finally it is important to note that the presence of a background Z-peak in the SUSY sample does not imply a significant SM background due to Z bosons. The Z bosons in SUSY events are special in that they appear in events where the event selection criteria of several hard jets and a high level of missing transverse momentum are naturally satisfied.

15.2 Uniformity of Event Selection Efficiency

A priori one would not expect a large dependency of the event selection efficiency on $m_{\ell\ell}$. The value of $m_{\ell\ell}$ is after all simply given by the slepton decay angle, meaning that its value should only be correlated with the energies and momenta of ℓ_{far} and the $\tilde{\chi}_1^0$. And as the slepton direction and energy itself depends upon the $\tilde{\chi}_2^0$ decay as well as the momenta of the $\tilde{\chi}_2^0$, this correlation will tend to be washed out. Figure 15.2 shows the distributions of the transverse momenta of the four hardest jets as well as the distribution of missing transverse momentum for various ranges of $m_{\ell\ell}$. As expected the jet distributions are seen to not depend upon the value of $m_{\ell\ell}$ at all, and the same seems to be true for the missing transverse momentum.

Similarly, figure 15.3 shows the distributions of the relevant kinematic variables for the two leptons (the definition of "lepton 1" and "lepton 2" is such that "lepton 1" has the higher transverse momentum). Figures 15.3.a and 15.3.b shows the η distribution of the two lepton momenta. It is clear that the requirement of $\eta < 2.5$ (defined by the extent of the inner ATLAS tracking system) only excludes very few events, and that this cut does not introduce significant correlations between selection efficiency and $m_{\ell\ell}$. Finally figures 15.3.c and 15.3.d shows the transverse momentum distributions of the two leptons. Unfortunately the requirement of $p_1^t > 30$ GeV *does* seem to introduce some degree of correlation between selection efficiency and $m_{\ell\ell}$. If, by raising other requirements, the requirements upon the transverse momenta of the two leptons could be allowed to be reduced to e.g. 10 GeV in both cases, then this correlation could obviously be reduced greatly. This should be kept in mind when, in a given analysis scenario, a more detailed cut optimisation investigation is carried out. Of course, for very low p_t requirements, trigger efficiencies might become important, although the redundancy due to jet and missing energy triggers will alleviate this issue somewhat.

In order to study the resulting combined effects of the event selection criteria upon the shape of the dilepton mass-edge, the distribution of $m_{\ell\ell}$ values of the selected signal events

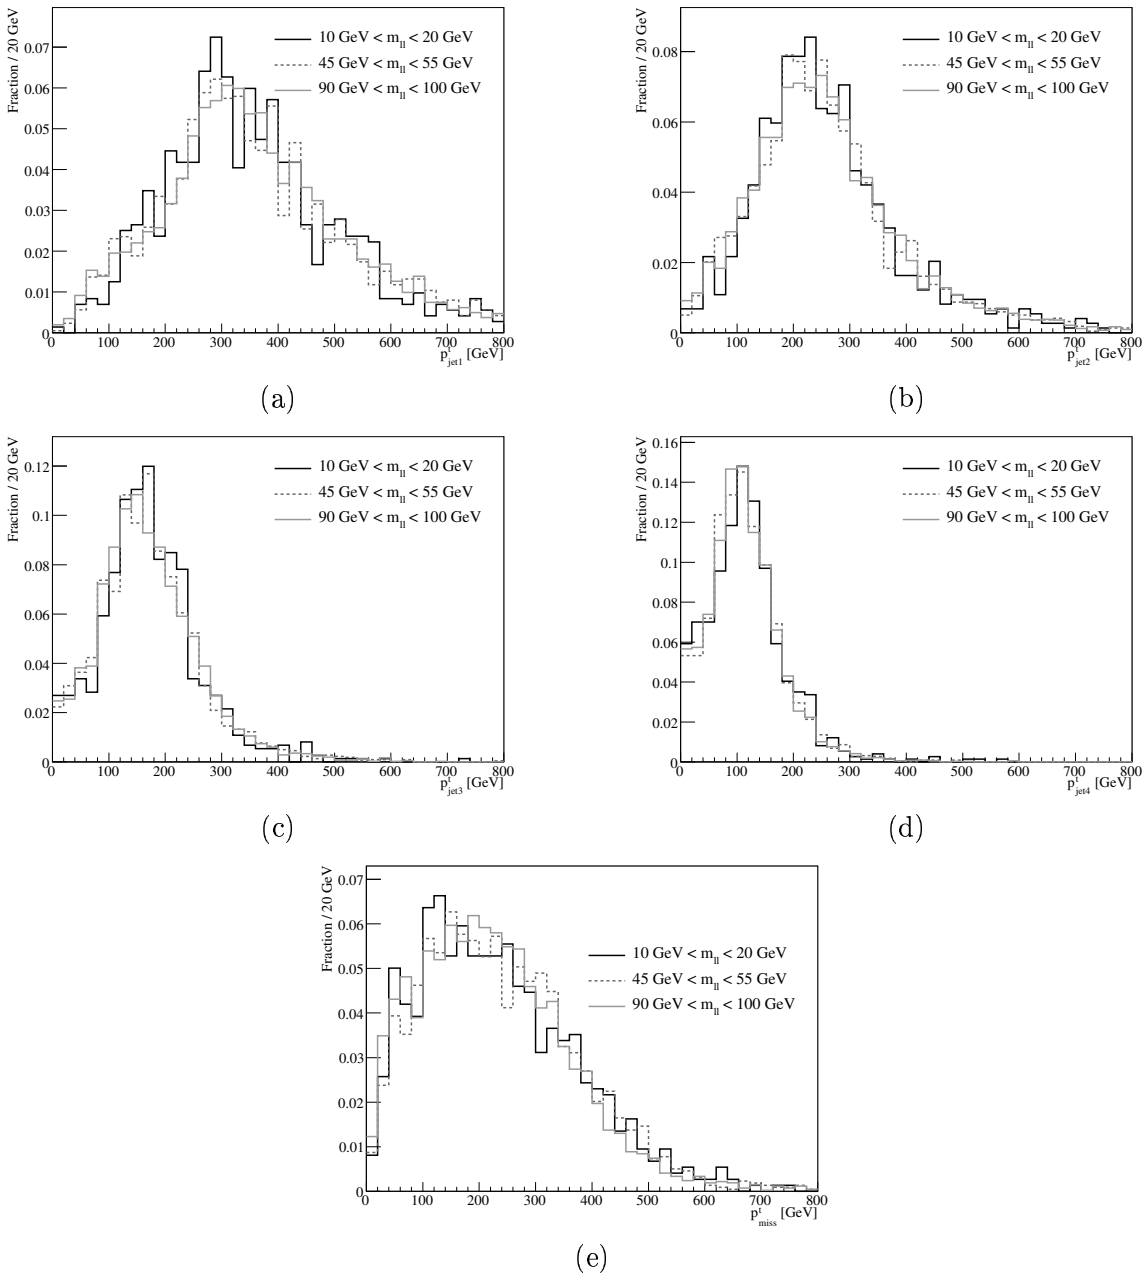

Figure 15.2: *Distribution of transverse momenta of the four hardest jets (a-d) and of missing transverse momentum (e) in signal events at the ATLAS SU3 point for different ranges of* $m_{\ell\ell}$.

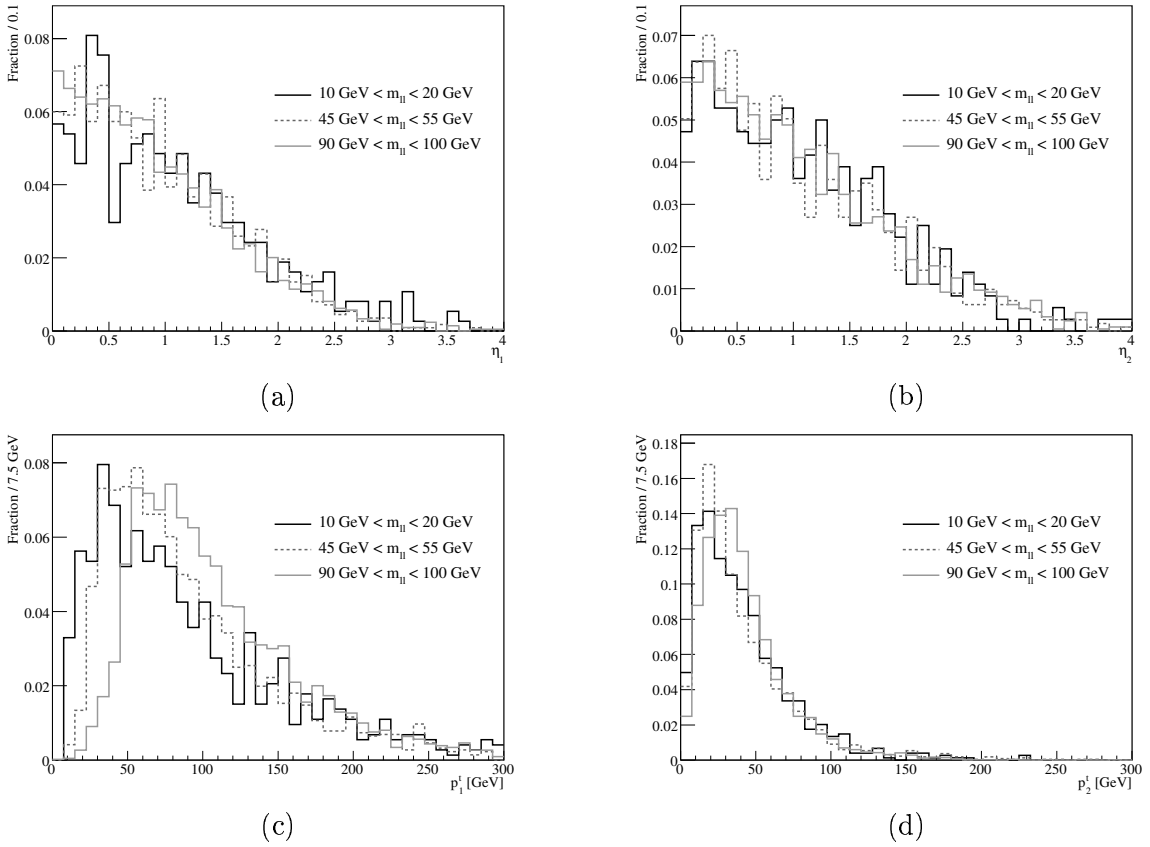

Figure 15.3: *Distribution of pseudo-rapidity (a-b) and transverse momentum (c-d) of the two hardest leptons in signal events at the ATLAS SU3 point for different ranges of* $m_{\ell\ell}$.

are shown in figure 15.4.a along with the values for all of the generated signal events. These can be divided to find the dependency of efficiencies on $m_{\ell\ell}$, and this is shown in figure 15.4.b along with the result of a first-order polynomial fit. The efficiency is seen to rise from $\sim 63\%$ at $m_{\ell\ell} \approx 0$ to $\sim 75\%$ at $m_{\ell\ell} \approx m_{\ell\ell}^{\max}$, and while the first order polynomial fit is not perfect, it does capture this general trend.

The result of this fit must be used to correct the observed data before the decay parameter is extracted. One way to achieve this is by multiplying the contents and associated error of each bin of the observed background-subtracted histogram by a correction factor, $C_{\text{eff}}(m_{\ell\ell})$. To avoid changing the appearance of the distributions more than necessary, this factor is normalised to an average value of 1. Also, as the fit range in figure 15.4.b does not cover all $m_{\ell\ell} > 0$, the interval $[10\text{ GeV}, 100\text{ GeV}]$ will be used for this normalisation, and the values of $C_{\text{eff}}(m_{\ell\ell})$ outside this interval will simply be fixed to the values at the edge. Based upon the fit in figure 15.4.b, $C_{\text{eff}}(m_{\ell\ell})$ will thus be proportional to $(0.6269 + 0.00124 m_{\ell\ell}/\text{ GeV})^{-1}$ inside the central interval, and the resulting curve is shown in figure 15.5.a. Figure 15.5.b shows the background subtracted distribution of $m_{\ell\ell}$ values before and after this correction factor is applied (including an ad hoc smearing of $\sigma = 2$ GeV of the $m_{\ell\ell}$ values).

The main issue facing this procedure of dealing with correlations between event selection efficiency and $m_{\ell\ell}$ is obviously that one needs to use Monte Carlo truth information to obtain a plot like the one in figure 15.4.b. This does of course not need to be a show stopper, given

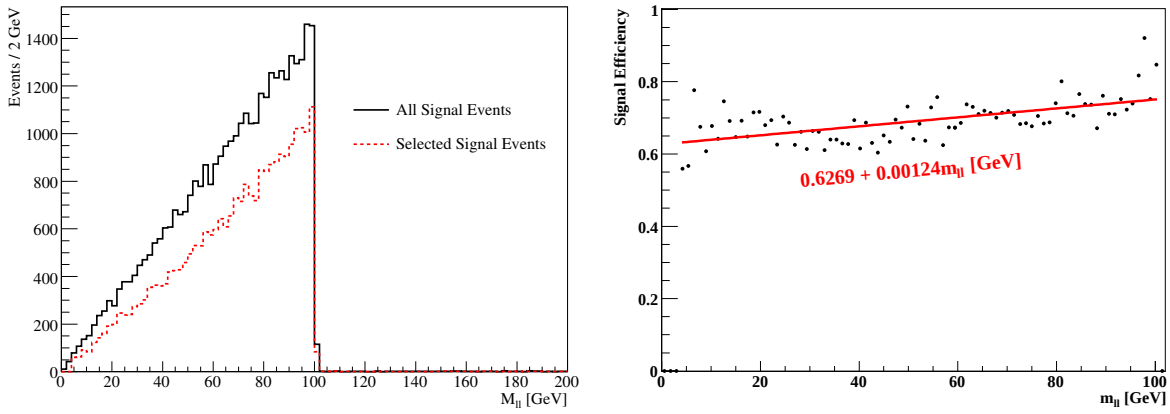

(a) (b)

Figure 15.4: *Distribution of $m_{\ell\ell}$ for all and selected signal events at the ATLAS SU3 point (a), and the resulting event selection efficiency as a function of $m_{\ell\ell}$ (b).*

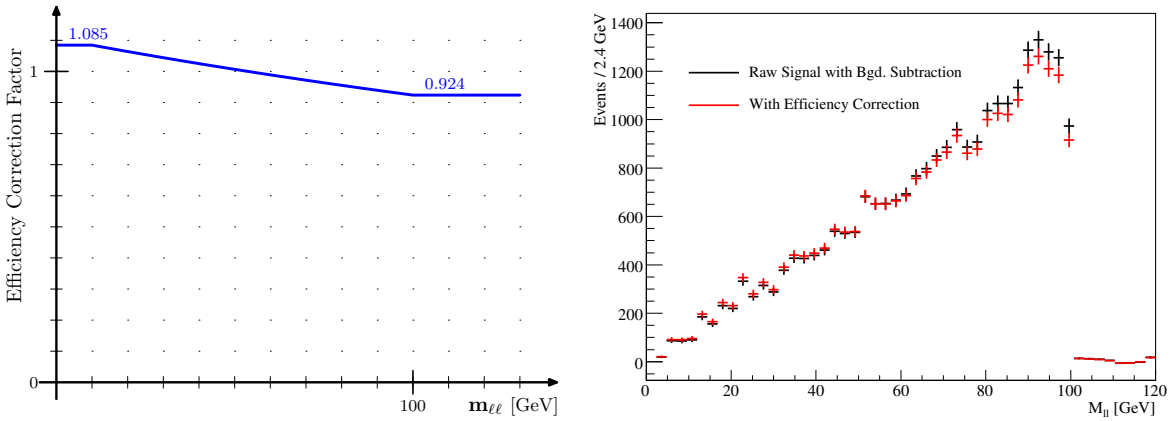

(a) (b)

Figure 15.5: *Efficiency correction factor at the ATLAS SU3 point (a) and the effect of applying it to the background subtracted data sample (b).*

that the data samples required for a useful decay parameter extraction are large enough that it is likely that most aspects of the Monte Carlo simulations will have been carefully tuned by the time they become available. And similarly a lot of information about the spectra of new particles should likewise be at hand. In any case the final decay parameter measurement will certainly contain a systematic error due to these limitations.

It is, however, interesting to note that, given enough statistics, it is in principle possible to extract the efficiency factor directly from data without resorting to input from Monte Carlo simulations. This follows directly from equations (13.9) and (13.14), which taken together means that the expression for $dP/dm_{\ell\ell}$ will only contain odd powers of $m_{\ell\ell}$ as long as the expression for $dP/d\cos\theta$ only contains a finite number of terms, all proportional to non-negative powers of $\cos\theta$ (such as 1, $\cos\theta$ or $\cos^2\theta$). This means that any even component of the observed $m_{\ell\ell}$ distribution must be due to efficiency effects. As an example, if the event selection efficiency is proportional to $1+\epsilon_1\hat{m}_{\ell\ell}$, then the observed distribution will be related

to the true distribution by:

$$\frac{dP^{\text{obs}}}{d\hat{m}_{\ell\ell}} \propto (1+\epsilon_1\hat{m}_{\ell\ell})\frac{dP}{d\hat{m}_{\ell\ell}} \tag{15.9}$$

Assuming for clarity that b is close to the SUSY case so that $b, \epsilon_1 \ll 1$, then the two decay angle distributions given by equations (13.15) and (13.16) will give rise to observable distributions:

$$\frac{dP_1^{\text{obs}}}{d\hat{m}_{\ell\ell}} \propto \hat{m}_{\ell\ell} + \epsilon_1\hat{m}_{\ell\ell}^2 - 2b\,\hat{m}_{\ell\ell}^3 \tag{15.10}$$

$$\frac{dP_2^{\text{obs}}}{d\hat{m}_{\ell\ell}} \propto \hat{m}_{\ell\ell} + \epsilon_1\hat{m}_{\ell\ell}^2 - 4b\,\hat{m}_{\ell\ell}^3 + 4b\,\hat{m}_{\ell\ell}^5 \tag{15.11}$$

And it should thus be possible to fit both ϵ_1 and b directly in the same fit. It is straightforward to use equation (13.17) to obtain an estimate for the expected resolution on the ϵ_1 parameter in the ideal limit, which is

$$\sigma_{\epsilon_1} \approx 0.04\sqrt{\frac{10^4}{N}} \tag{15.12}$$

Note that the numbers in figure 15.4 correspond to $\epsilon_1 \approx 0.2$. Since the ideal σ_{ϵ_1} is thus not much larger than the ideal σ_b, it would not be unreasonable to expect that it should be possible to fit both b and ϵ_1 simultaneously without ruining the b measurement due to correlations with ϵ_1. Of course, it is not a requirement of this method that the efficiency is modelled as $1+\epsilon_1\hat{m}_{\ell\ell}$. Any expression could be used as long as it (1) does not introduce more than one or two free parameters and (2) does not introduce terms that are allowed to vary independently in a way which mimics the variation due to b. As an example, the latter requirement means that a form such as $1+\epsilon_1\hat{m}_{\ell\ell}+\epsilon_2\hat{m}_{\ell\ell}^2$ would not work due to correlations between b and ϵ_2, but forms such as $1+\epsilon_{12}(\hat{m}_{\ell\ell}+\hat{m}_{\ell\ell}^2)$, $1+\epsilon_1\hat{m}_{\ell\ell}+\epsilon_3\hat{m}_{\ell\ell}^3$ or even $\exp[\epsilon_e\hat{m}_{\ell\ell}]$ would. The ideal parametrisation for this purpose should be found from Monte Carlo studies.

In a full-fletched analysis one should obviously perform the decay parameter extraction twice, once where the data are corrected for efficiency based on Monte Carlo simulations, and once where efficiency effects are included directly in the fit through a free parameter such as ϵ_1. The two results can then be compared and it can (hopefully) be verified that they give compatible results. The final results can then be based on a combination of these (or simply be taken to the more precise of them).

15.3 Fit Results

Results of fits to background subtracted and efficiency corrected dilepton edges at the ATLAS SU3 point corresponding to data samples of 300 fb^{-1} are shown in figure 15.6 for different input values of the decay parameter and for different fitted distributions. All $m_{\ell\ell}$ values have been smeared with $\sigma = 2$ GeV. This σ is not taken as a free parameter in the fits since, as already mentioned, it is expected to be estimated on an event by event basis in a full-fletched analysis. It has been checked, however, that letting σ be a free parameter in the fit does not change the results for the decay parameter significantly, which is to be expected since $\sigma \ll m_{\ell\ell}^{\text{max}}$.

Figures 15.6.a and 15.6.b show fits to $dP/d\cos\theta \propto 1 + b\cdot\cos\theta$ and $dP/d\cos\theta \propto 1 + b\cdot\cos^2\theta$, respectively, in the case where the slepton decay is taken to be pure scalar in nature (the data

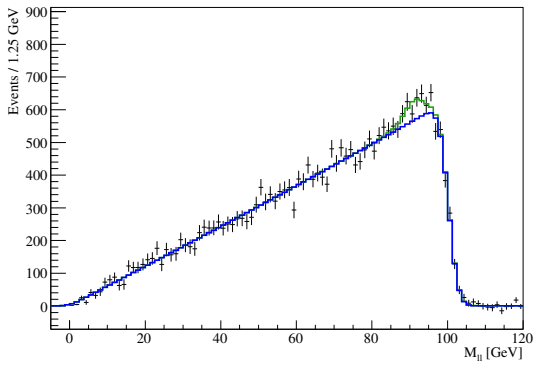

b	=	-0.019 ± 0.02
M	=	100.22 ± 0.11 GeV
f_Z	=	(1.5 ± 0.5) %
N	=	$(2.51 \pm 0.02) \cdot 10^4$
Fit reduced $\chi^2 = 1.37$		

(a): $dP/d\cos\theta \propto$ const (scalar). Fitted as $dP/d\cos\theta \propto 1 + b \cdot \cos\theta$.

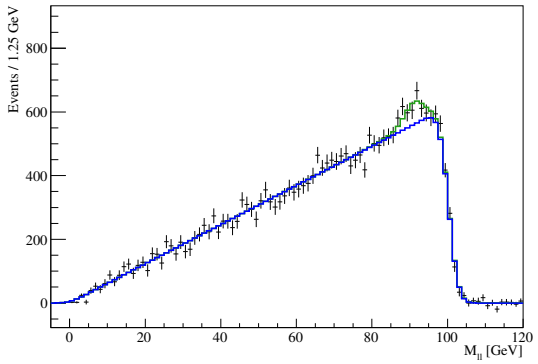

b	=	0.002 ± 0.033
M	=	100.27 ± 0.11 GeV
f_Z	=	(1.7 ± 0.41) %
N	=	$(2.52 \pm 0.02) \cdot 10^4$
Fit reduced $\chi^2 = 1.38$		

(b): $dP/d\cos\theta \propto$ const (scalar). Fitted as $dP/d\cos\theta \propto 1 + b \cdot \cos^2\theta$.

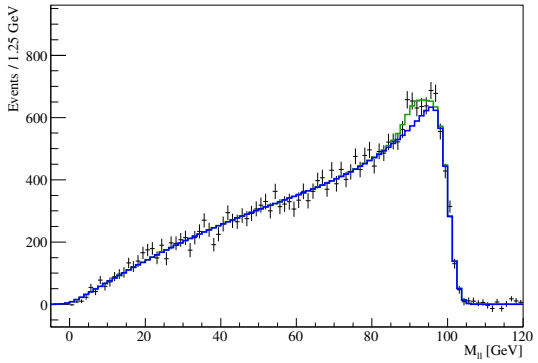

b	=	0.259 ± 0.039
M	=	100.14 ± 0.10 GeV
f_Z	=	(1.6 ± 0.4) %
N	=	$(2.53 \pm 0.02) \cdot 10^4$
Fit reduced $\chi^2 = 1.38$		

(c): $dP/d\cos\theta \propto 1 + 0.3\cos^2\theta$

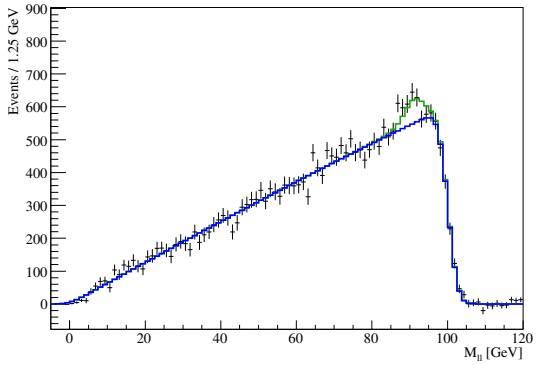

b	=	0.028 ± 0.02
M	=	100.08 ± 0.11 GeV
f_Z	=	(1.8 ± 0.5) %
N	=	$(2.51 \pm 0.02) \cdot 10^4$
Fit reduced $\chi^2 = 1.53$		

(d): $dP/d\cos\theta \propto 1 + 0.0415\cos\theta$ (UED like)

Figure 15.6: *Results of fits at the ATLAS SU3 point for 300 fb^{-1} as discussed in the text. The blue curves shows the extracted dilepton mass edges, while contributions from $Z \to \ell^+\ell^-$ are shown in green.*

are not exactly similar in the two histograms since the ad hoc $m_{\ell\ell}$ smearing is not the same in the two cases). The first important thing to note is that the extracted values of the decay parameters are compatible with the input parameters. The reported sensitivity on the decay parameter is 0.02 and 0.033, respectively, in the two cases. This can be compared to the predictions for the ideal cases shown in figure 13.6, which are 0.0173 and 0.0335, respectively, for $N = 10^4$. The fits report $N \approx 2.5 \cdot 10^4$, but due to the statistical fluctuations in the background events this number should effectively be reduced by a factor of $1 + 2N_B/N_S$ according to equation (15.5). Estimating $N_B/N_S \approx 0.4$ from figure 15.1, and scaling the errors by $\sqrt{10^4 \cdot (1 + 2N_B/N_S)/N}$ thus predicts sensitivities of 0.015 and 0.028 respectively. This means that the observed sensitivities are higher than the ideal limit by 36% and 16% respectively. It seems reasonable that this minor difference can be attributed to fit correlations and the imposed finite $m_{\ell\ell}$ resolution.

Figure 15.6.c shows a fit to data where the slepton decay angle follows the distribution $dP/d\cos\theta \propto 1 + 0.3\cos^2\theta$. Again the extracted value is compatible with the true value, and the sensitivity is now 0.039 – in line with what would be expected from a scaling of the sensitivity of 0.033 of figure 15.6.b by the ratio of expected sensitivities at the two values of b in the ideal case as given by figure 13.6.

Finally, figure 15.6.d shows a fit to the distribution for a decay parameter given by the UED prediction. Apart from a shift in the central value of b, the results are similar to those of figure 15.6.a, which is not surprising as the expected sensitivity on b is essentially identical for $b = 0$ and $b = 0.0415$. The average discriminating power between SUSY and UED for this scenario thus appears to be around 2σ. That this is not higher is obviously due to the smallness of b_{UED}, since the sensitivity to b is actually rather good.

If one looks at all four sub-figures in figure 15.6, there is a general trend for the fit to provide slightly low values of b (and this trend is consistent with other fits not shown here). This is not inherently a problem, however, as it is not unusual for fits to provide a slight bias in the returned value since $N < \infty$. One might of course argue that N of $\mathcal{O}\left(10^4\right)$ is already a large number, but such an argument is not necessarily valid as long as the shift is comparable to the statistical sensitivity. The presented procedure for extraction of decay parameters should thus be augmented by a final correction shift based upon a high number of (toy) Monte Carlo experiments. This issue will be discussed further in section 16.1.

To investigate a more troublesome scenario the ATLAS "SU1" point is also investigated. At this point the relevant sparticle masses are:

$$\begin{aligned} m_{\tilde{\chi}_2^0} &= 263.6 \text{ GeV}, \quad m_{\tilde{\chi}_1^0} = 137.0 \text{ GeV} \\ m_{\tilde{\ell}_L} &= m_{\tilde{e}_L} = m_{\tilde{\mu}_L} = 255.1 \text{ GeV} \\ m_{\tilde{\ell}_R} &= m_{\tilde{e}_R} = m_{\tilde{\mu}_R} = 154.1 \text{ GeV} \end{aligned} \tag{15.13}$$

Which means that both of the decays $\tilde{\chi}_2^0 \to \tilde{\ell}_L^{\pm}\ell^{\mp}$ and $\tilde{\chi}_2^0 \to \tilde{\ell}_R^{\pm}\ell^{\mp}$ are open, leading to dilepton mass edges with $m_{\ell\ell}^{\text{max},1} = m_{\ell\ell}^{\text{max,L}} = 56.0$ GeV and $m_{\ell\ell}^{\text{max},2} = m_{\ell\ell}^{\text{max,R}} = 97.9$ GeV respectively.

Once again, equation (14.10) can be used to find the decay parameters for these two mass edges that an analogous UED mass spectrum would have. They are $b^1_{\text{UED}} = b^L_{\text{UED}} = -0.097$ and $b^2_{\text{UED}} = b^R_{\text{UED}} = 0.159$ respectively. These values are a lot further from zero than they are at the SU3 point, but unfortunately the decay parameter measurements are also harder.

First of all, the overall SUSY production cross section at the SU1 point is about a third of the one at the SU3 point, while the event selection efficiency for events going through $\tilde{\ell}_L$

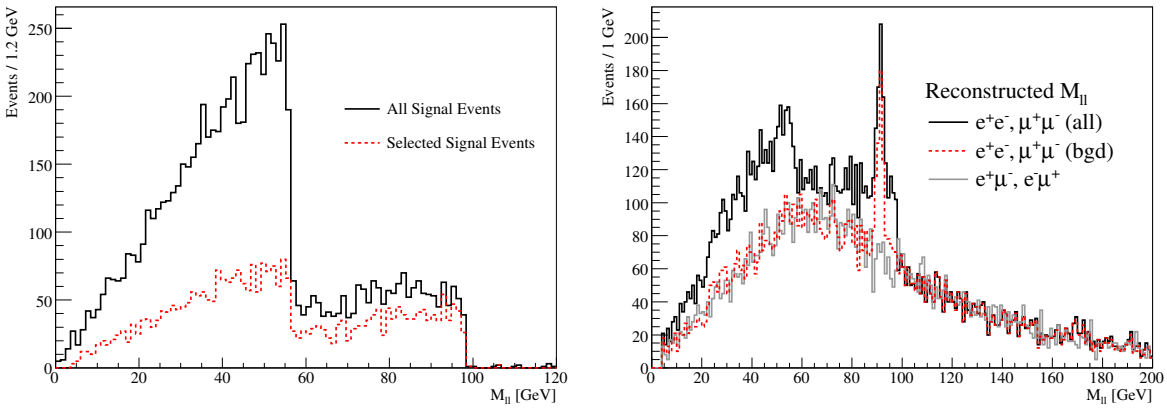

(a) (b)

Figure 15.7: *Distribution of $m_{\ell\ell}$ for all and selected signal events at the ATLAS SU1 point for 300 fb^{-1} (a) and dilepton samples after event selection cuts (b).*

is also lower ($\sim 30\%$) as can be seen in figure 15.7.a. This is because of the small value of $m_{\tilde{\chi}_2^0} - m_{\tilde{\ell}_L} = 8.5$ GeV which makes it hard for ℓ_{far} to obtain a momentum high enough to meet the necessary criteria. Furthermore, the relative amount of background, $N_B/N_S \approx 2$, is rather high as shown in figure 15.7.b, and also f_Z is a lot higher than at the SU3 point.

Results of fits to background subtracted dilepton edges corresponding to data samples of 300 fb^{-1} are shown in figure 15.8. This time, however, no efficiency correction has been applied. As expected the results are somewhat worse, and it is really only the decay parameter of the first edge that can be extracted to some degree. Again the results can be compared to those expected for the ideal case, given the effective number of events in each edge resulting from f_Z, f_1 and N_B/N_S, by scaling them by $\sqrt{10^4[1+2N_B/N_S]/[(1-f_Z)f_i N]}$. The expected decay parameter sensitivities, $(\sigma_{b1}, \sigma_{b2})$ for the four sub-figures in figure 15.8 are $(0.09, 0.10)$ $(0.18, 0.20)$, $(0.28, 0.15)$ and $(0.09, 0.11)$ respectively. The reported errors in the fits in figure 15.8 on b_1 are thus generally 22-33% larger than in the ideal case which is very reasonable. The errors on b_2 on the other hand are roughly four times larger than the ideal case which is mainly attributed to correlations between the two edges and with the Z peak.

The results can of course be expected to improve somewhat if the event selection cuts are optimised for the point and if the background is not subtracted but fitted as discussed in the previous section. If these improvements are not enough it might, however, still be possible to obtain a reasonable discrimination between SUSY and UED at the point due to the slepton decay angle distribution by sacrificing the model independence of the method. The general approach used here assumes that two edges are fitted for four physics parameters: $m_{\ell\ell}^{\max,1}$, $m_{\ell\ell}^{\max,2}$, b_1 and b_2. But by using equation (14.10) explicitly, thus assuming an UED prediction for the decay parameters, the four physics masses $m_{\tilde{\chi}_2^0}$, $m_{\tilde{\ell}_L}$, $m_{\tilde{\ell}_R}$ and $m_{\tilde{\chi}_1^0}$ can be fitted directly. If one or more of these are measured (simultaneously in a global fit or otherwise) in other analyses, the number of free parameters can be brought down.

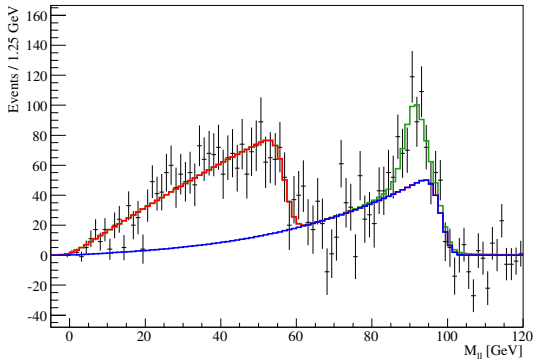

b_1	=	0.11 ± 0.12
b_2	=	-0.62 ± 0.48
M_1	=	57.0 ± 1.1 GeV
M_2	=	98.2 ± 0.8 GeV
f_1	=	(55.3 ± 7.2) %
f_Z	=	(9.6 ± 2.8) %
N	=	3428 ± 126
Fit reduced $\chi^2 = 1.05$		

(b): $dP/d\cos\theta \propto$ const (scalar). Fitted as $dP/d\cos\theta \propto 1 + b \cdot \cos\theta$.

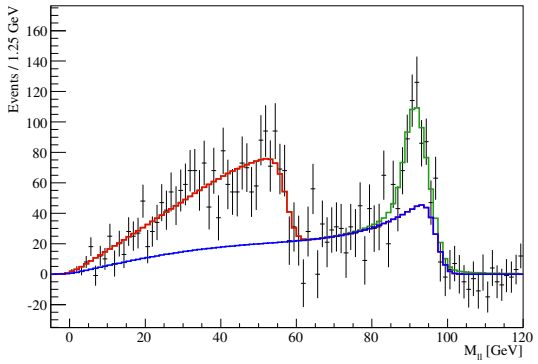

b_1	=	-0.15 ± 0.24
b_2	=	$0.61 + 0.75 - 0.61$
M_1	=	57.3 ± 0.9 GeV
M_2	=	96.9 ± 0.9 GeV
f_1	=	(47.4 ± 6.1) %
f_Z	=	(12.0 ± 3.0) %
N	=	3433 ± 125
Fit reduced $\chi^2 = 0.84$		

(a): $dP/d\cos\theta \propto$ const (scalar). Fitted as $dP/d\cos\theta \propto 1 + b \cdot \cos^2\theta$.

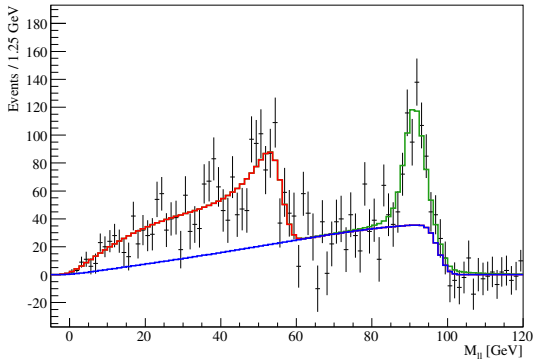

b_1	=	$0.72 + 0.40 - 0.33$
b_2	=	$-0.11 + 0.53 - 0.46$
M_1	=	56.0 ± 0.8 GeV
M_2	=	97.6 ± 1.3 GeV
f_1	=	(48.8 ± 6.0) %
f_Z	=	(14.8 ± 2.7) %
N	=	3503 ± 126
Fit reduced $\chi^2 = 1.05$		

(c): $dP/d\cos\theta \propto 1 + 0.5\cos^2\theta$ (edge 1), $dP/d\cos\theta \propto 1 - 0.3\cos^2\theta$ (edge 2)

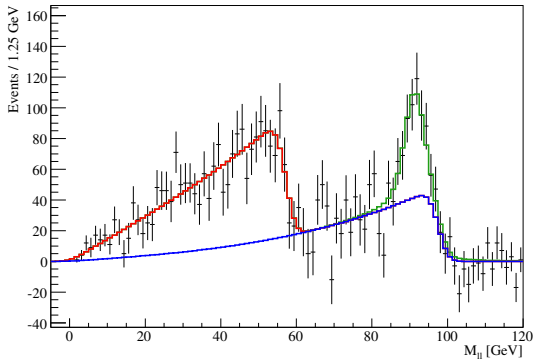

b_1	=	-0.04 ± 0.11
b_2	=	$-0.49 + 0.53 - 0.63$
M_1	=	57.2 ± 0.8 GeV
M_2	=	97.5 ± 1.0 GeV
f_1	=	(57.5 ± 7.5) %
f_Z	=	(12.4 ± 2.9) %
N	=	3416 ± 125
Fit reduced $\chi^2 = 0.90$		

(d): $dP/d\cos\theta \propto 1 - 0.097\cos\theta$ (edge 1), $dP/d\cos\theta \propto 1 + 0.159\cos\theta$ (edge 2) (UED like)

Figure 15.8: *Results of fits at the ATLAS SU1 point for 300 fb^{-1} as discussed in the text. The blue and red curves shows the two extracted dilepton mass edges, while contributions from $Z \to \ell^+\ell^-$ are shown in green*

16 CMSSM Scans and High Statistics Toy Monte Carlo Investigations

In order to investigate how likely the method presented in the previous sections is to be applicable within the context of the CMSSM and how powerful it is likely to be when it is, this section will present the results of toy Monte Carlo studies performed on a large number of SUSY spectra with each spectrum corresponding to one point in CMSSM space.

As explained in section 5.4.4 and 6.2, the sparticle spectrum and branching ratios at each CMSSM point is found by the ISASUGRA program and the direct SUSY production cross sections via a modified version of HERWIG. With this information the number of $\tilde{\chi}_2^0 \to \tilde{\ell}_L$ and $\tilde{\chi}_2^0 \to \tilde{\ell}_R$ decays (if any) that is expected in the detector for a given integrated luminosity are estimated by starting from the directly produced sparticles and traversing all possible decay channels while appropriately multiplying branching ratios until the desired decay chains are found. This information is then used to generate signal-only toy Monte Carlo samples of appropriate statistics (assuming scalar sleptons in all cases). It was shown in the last section that the presence of background can effectively be incorporated simply by lowering the available number of events. Therefore a flat and relatively low event selection efficiency of 10% is assumed. This also implies that at spectra where two edges are kinematically accessible the relative fraction of events in each edge will be given purely by the relative branching ratios of $\tilde{\chi}_2^0$ into $\tilde{\ell}_R$ and $\tilde{\ell}_L$ respectively.

In figure 16.1 is shown examples of generated toy Monte Carlo samples along with the results of their fitted shapes. Figures 16.1.a and 16.1.b shows the result of fits to one and two edges respectively. Once again, the extracted decay parameter sensitivities can be compared to the ideal expectations of figure 13.6. The sensitivity in figure 16.1.a is only 10% worse than the ideal case while the sensitivities for b_1 and b_2 in figure 16.1.b are 22% and 72% worse respectively. The results are thus compatible with the generator level results apart from the sensitivity on the second edge being somewhat better than when compared to the ATLAS SU1 studies in the last section. This difference can be attributed to the presence of the Z peak just at the second edge of the ATLAS SU1 point.

Figure 16.1.c illustrates the situation where two edges happen to have $m_{\ell\ell}^{\max,1} \approx m_{\ell\ell}^{\max,2}$. As expected, the correlation causes the errors on the two decay parameters to blow up. As mentioned in section 14 it would still be possible to gain information from such a merged edge by fitting it as a single edge instead. Such a procedure will not be attempted here, however.

Note that no background Z peak is included in these toy Monte Carlo studies, which should only affect the results significantly when one of the edges in a two-edge fit happens to have $m_{\ell\ell}^{\max}$ around $80 - 100$ GeV like at the ATLAS SU1 point. That the fit to a given edge is only affected when it has $m_{\ell\ell}^{\max}$ around $80 - 100$ GeV should be clear – if $m_{\ell\ell}^{\max}$ is lower than ~ 80 GeV, then the Z peak is positioned at $m_{\ell\ell}$ values where $dP/dm_{\ell\ell}$ vanishes for all sensible values of the fitted $m_{\ell\ell}^{\max}$ and decay parameter. In such a scenario a Z peak would even be beneficial in that it could be used to cross-check the overall energy scale and the estimates of the resolution of the individual $m_{\ell\ell}$ measurements. When an edge has $m_{\ell\ell}^{\max}$ higher than ~ 100 GeV, its shape will be visible at both sides of the Z peak, meaning that the only effect of the Z peak will be that the signal events directly below it will contribute less strongly to overall mass edge fit. Finally, the claim that the presence of a Z peak should not significantly affect the results for data samples where only one edge is present is verified by noting that the results at the ATLAS SU3 point, a one-edge point with a Z peak close to the edge, and

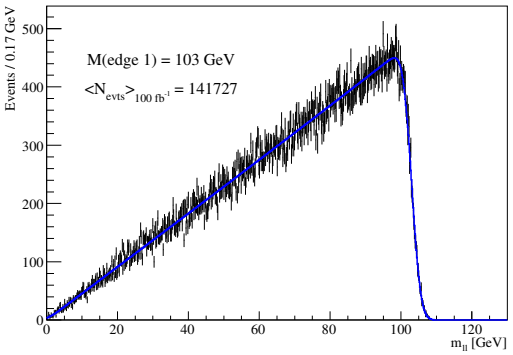

b	=	-0.0068 ± 0.0049
M_1	=	103.01 ± 0.03 GeV
N	=	$1.41 \cdot 10^5 \pm 376$
Fit reduced $\chi^2 = 0.891$		

(a): $(m_0, m_{1/2}) = (100\text{ GeV}, 300\text{ GeV})$

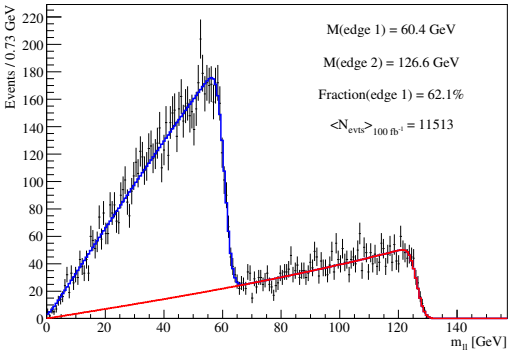

b_1	=	0.031 ± 0.025
b_2	=	-0.098 ± 0.046
M_1	=	60.68 ± 0.15 GeV
M_2	=	126.25 ± 0.18 GeV
f_1	=	63.3 ± 0.7 %
N	=	$1.136 \cdot 10^4 \pm 107$
Fit reduced $\chi^2 = 0.893$		

(b): $(m_0, m_{1/2}) = (100\text{ GeV}, 400\text{ GeV})$

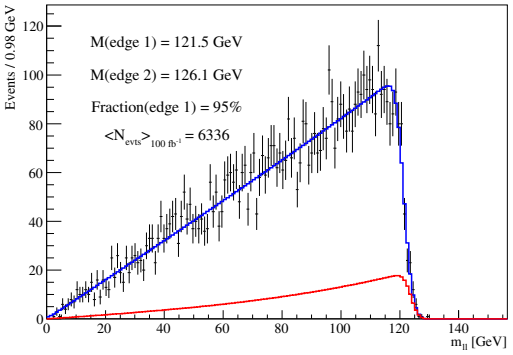

b_1	=	$0.03 + 0.44 - 2.3$
b_2	=	$-0.3 + 3.2 - 1.5$
M_1	=	$121.0 + 0.86 - 9.4$ GeV
M_2	=	$123.7 + 2.8 - 2.2$ GeV
f_1	=	84.7 ± 36.5 %
N	=	6119 ± 78
Fit reduced $\chi^2 = 0.854$		

(c): $(m_0, m_{1/2}) = (100\text{ GeV}, 500\text{ GeV})$

Figure 16.1: *Examples of fits to Toy Monte Carlo experiments at three qualitatively different SUSY points. All three experiments share* $(\tan\beta, A_0, \text{sign}\mu) = (5, 0\text{ GeV}, +1)$ *and represent experiments at 500 fb*$^{-1}$ *with event selection efficiency at 10% and an* $m_{\ell\ell}$ *resolution of 2 GeV.*

that of figure 16.1.a both agrees with the ideal expectations of figure 13.6 (and thus with each other) at the 10% level.

A total of 1000 experiments were simulated in this way at each CMSSM point at the Steno cluster at the Danish Center for Scientific Computing. For each simulated experiment the binning and histogram ranges used in the fit was adjusted according to the available statistics and when two edges were kinematically accessible, the fit was for both edges unless the content of one of them was negligible. The definition of when the content of an edge is negligible was taken to be whether or not it contains less than 1% of the events, although this is obviously a somewhat arbitrary criteria.

16.1 Fit Systematics

One important point regarding the parameter extraction is that the inherent limitations of likelihood fits and, by extension, chi-square fits means that the extracted values of parameters and associated errors will only be entirely trustworthy in the limit of infinite statistics. This is because they are based upon a set of assumptions regarding the probability density function in question and its derivatives, even when, as here, the improved MINOS method in MINUIT is used for the fit (for more information regarding likelihood based estimators see [Cow98], [EDJ+71] and [Kit03, chap. 7.1]). In order to reliably correct for this flaw one should examine in a systematic way how the extracted parameter value can be expected to relate to the true input value and what information is contained in the extracted associated error. The only method which can reliably be used as the basis for such an examination is that of studying the results of the fit on a high statistics ensemble of (toy) Monte Carlo experiments.

Considering a directly extracted parameter, α, one would obviously desire that its value averaged over the ensemble should equal the true input value, α_{true}, i.e. that it would be an unbiased estimator for α_{true}. If this turns out not to be the case it is straight-forward to define an improved unbiased estimator,

$$\alpha' \equiv \alpha + C_\alpha(\alpha_{\text{true}}, N_{\text{evt}}), \qquad C_\alpha(\alpha_{\text{true}}, N_{\text{evt}}) \equiv \alpha_{\text{true}} - \langle\alpha\rangle \tag{16.1}$$

Where $\langle\alpha\rangle$ represents the extracted parameter value averaged over the simulated ensemble. Although not shown explicitly in the definition it is of course true that C_α in principle can depend on all the parameters of the problem whether free or fixed, resolution parameters or physics parameters, etc.

Regarding the directly extracted error on α, σ_α, the situation is not always as straight-forward. It would be preferable if it could be interpreted as a Gaussian error, i.e. such that the distribution of α' for those experiments in the ensemble associated with a given reported value of σ_α would be a Gaussian distribution centred on α_{true} and with a spread of σ_α. In the ideal limit that this is indeed the case, the "pull" parameter $(\alpha' - \alpha_{\text{true}})/\sigma_\alpha$ should be distributed according to a unit Gaussian of zero mean. Unfortunately this ideal situation is not always realised even approximately: The resulting distribution might be asymmetric, have enhanced tails, etc. In many cases, however, the resulting distribution can indeed be approximately described as a Gaussian distribution, albeit with non-unit mean. In these cases an improved

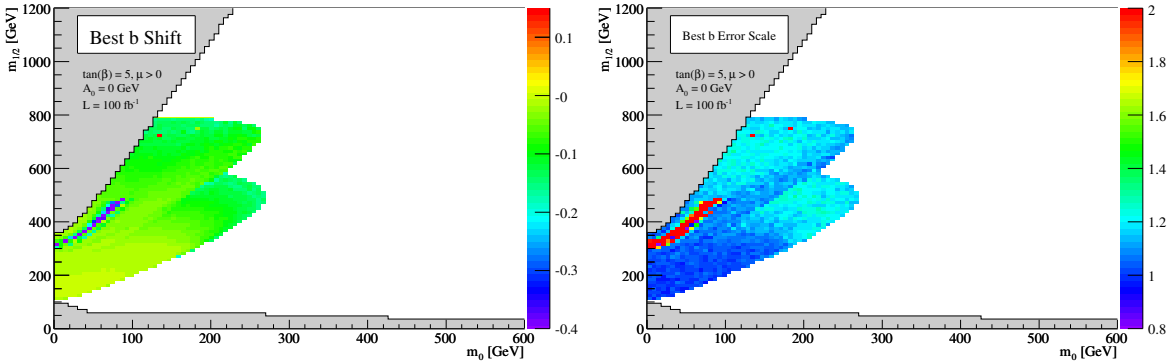

(a) (b)

Figure 16.2: *Systematic shift in fitted b values (a) and necessary scaling of reported fit errors on b due to fit systematics (b). Results are based upon 1000 simulated experiments at each point. At points where two edges are present, it is the values associated with the best determined of the two decay parameters that are plotted.*

estimator of the Gaussian error can be defined as:[1]

$$\sigma'_\alpha \equiv S_\alpha(\alpha_{\text{true}}, N_{\text{evt}})\,\sigma_\alpha \tag{16.2}$$

Where S_α is defined appropriately based on e.g. the root-mean-square of, or a Gaussian fit to, the pull distribution of α' and σ_α. If $S_\alpha > 1$ ($S_\alpha < 1$) it means that the original extracted error was under (over) estimated. If this procedure works the distribution of the pull of the final refined estimators, $(\alpha' - \alpha_{\text{true}})/\sigma'_\alpha$ should thus be a proper unit Gaussian. Of course, in many cases the result of this entire procedure might simply be that $C_\alpha \approx 0$, $S_\alpha \approx 1$ and that the pull distribution is indeed Gaussian. In such cases the extracted parameter estimates and associated errors can be used directly. In fact, in the limit of infinite statistics this should always become the case – which can be used for a cross-check of the implemented fitting procedure.

Applying this procedure to the mass edge fits at the CMSSM points results in correction terms C_b and factors S_b as shown in figures 16.2.a and 16.2.b, respectively, (actually, figure 16.2.a shows $-C_b$ rather than C_b). Note that when two edges are fitted the results for the most accurately measured decay parameter are reported and that the values of S_b are chosen such that the root-mean-square of the pull distribution is fixed at unity. The first thing to point out in these figures is the thin region approximately running between $(m_0, m_{1/2}) = (0\text{ GeV}, 320\text{ GeV})$ and $(m_0, m_{1/2}) = (90\text{ GeV}, 490\text{ GeV})$ where the corrections become very large. This is not surprising since this is the region where a two edge fit to a merged edge such as the one in figure 16.1.c is attempted. As has already been discussed the two edge fit strategy is really inadequate at this point and a one edge fit to the merged edge should instead be attempted. The reason the region stops around $(m_0, m_{1/2}) = (90\text{ GeV}, 490\text{ GeV})$ instead of continuing as one would expect from figure 14.5.c is that the contribution from one of the two edges becomes negligible at higher values of m_0, with the precise endpoint determined by the (somewhat arbitrary) requirement of at least 1% of the events.

[1] When interesting physics parameters in the fit are correlated to each other it is of course important to note that the final result should contain more than a list of results for the individual parameters as discussed here, but must also contain information about parameter correlations.

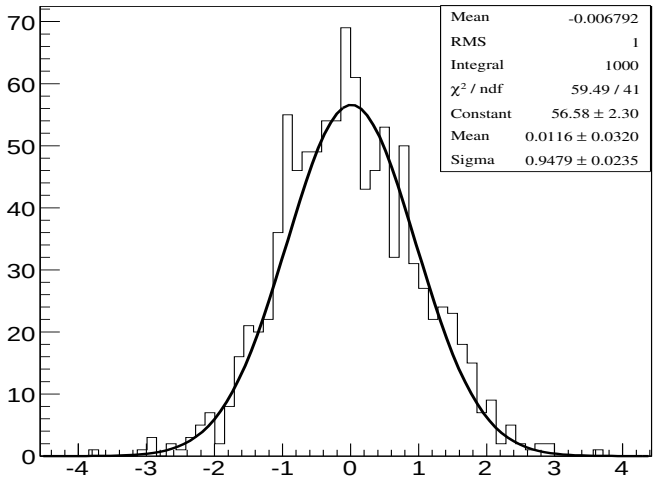

Figure 16.3: *Pull distribution of the corrected decay parameter measurement at the CMSSM point* $(\tan\beta, A_0, \mathrm{sign}\mu) = (5, 0\ GeV, +1)$ *and* $(m_0, m_{1/2}) = (100\ GeV, 300\ GeV)$ *for experiments at 500* fb^{-1} *with 10% event selection efficiency. The fitted* $dP/dm_{\ell\ell}$ *corresponds to* $dP/d\cos\theta \propto 1 + b \cdot \cos^2\theta$.

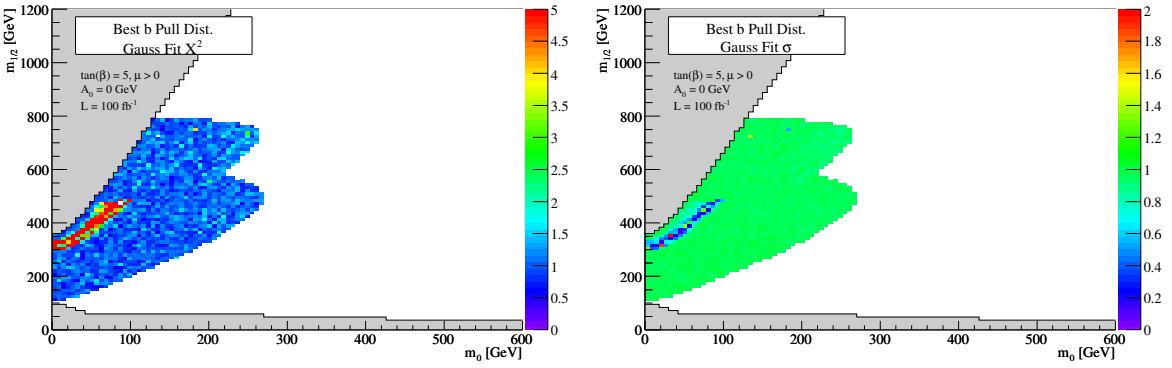

(a) (b)

Figure 16.4: *Distribution of reduced* χ^2 *(a) and* σ *(b) of the Gaussian fits to the pull distribution of the decay parameters based on 1000 simulated experiments. At points where two edges are present, it is the values associated with the best determined of the two decay parameter that are plotted.*

Apart from this special case the conclusion of figure 16.2 is otherwise clear: At low values of m_0 where, according to figure 14.8, the statistics is plentiful, $C_b \approx 0$ and $S_b \approx 1$. At larger m_0 the reported value of the decay parameter on the other hand receives a systematic shift towards lower values up to $C_b \approx 0.1$ and the reported associated error becomes slightly underestimated up to $S_b \approx 1.2$. In a sense these two corrections are not large, but they become very important given the small values of the statistical error and small difference between the UED and SUSY predictions.

To confirm that the resulting corrected decay parameter estimates are sensible, it is checked whether the distribution of their pull is compatible with being a unit Gaussian of non-zero mean. An example of such a distribution and the result of a Gaussian fit to it is shown in figure 16.3. Although not a perfect unit Gaussian the result is close enough for all practical purposes. The fact that the reduced chi-square is slightly high and the fitted Gaussian width slightly low reveals that the tails are slightly enhanced compared to a Gaussian distribution, which implies that fixing the root-mean-square rather than the width of the Gaussian fit to unity constitutes a slightly conservative choice if the errors later are interpreted as Gaussian. Figures 16.4.a and 16.4.b show the values of the reduced chi-square and widths of fitted Gaussian that are obtained from similar histograms at all of the CMSSM points. They show that, apart from the special region of merged edges, the results are everywhere pull distributions similar to the one in figure 16.3, confirming that it is sensible to interpret errors in terms of Gaussian probabilities.

16.2 Results

Figure 16.5 shows the decay parameter sensitivity in $dP/d\cos\theta \propto 1 + b \cdot \cos\theta$ fits, defined as the root-mean-square of the decay parameter values extracted in the toy Monte Carlo experiments, as a function of CMSSM point and integrated luminosity. Where two edges are fitted, the best of the two decay parameter sensitivities is shown. The areas where the second decay parameter can also be extracted with a reasonable sensitivity, defined as $\sigma_b < 0.2$, are additionally indicated with hatching. Results are only shown here for two different points in $(\tan\beta, A_0, \mathrm{sign}\mu)$ space, but, for reference, results for a total of 16 different points in $(\tan\beta, A_0, \mathrm{sign}\mu)$ space are shown in appendix A (this will similarly be true for the remaining figures in this section). As expected, the results in figure 16.5 are mainly determined by the available statistics as shown in figure 14.8, again apart from the special region where two merged edges are fitted.

Similar results for $dP/d\cos\theta \propto 1 + b \cdot \cos^2\theta$ fits are shown in figure 16.6, which are worse by about a factor of two. This is explained by the ratio of the two curves at $b = 0$ in figure 13.6.

Figure 16.7 shows the (best) extracted decay parameter in $dP/d\cos\theta \propto 1 + b \cdot \cos\theta$ fits in light of the direct, indirect and cosmological limits on the CMSSM parameter space discussed in section 5.4.4. They are certainly rather fortunate for the present analysis which in the case of the dark matter bound is not entirely coincidental since the allowed regions are those where sleptons are relatively light.

16.2.1 SUSY Versus UED Discrimination

Based upon the expected decay parameter sensitivities in figure 16.5 and the decay parameter values expected for a model of universal extra dimensions with a mass spectrum similar to the

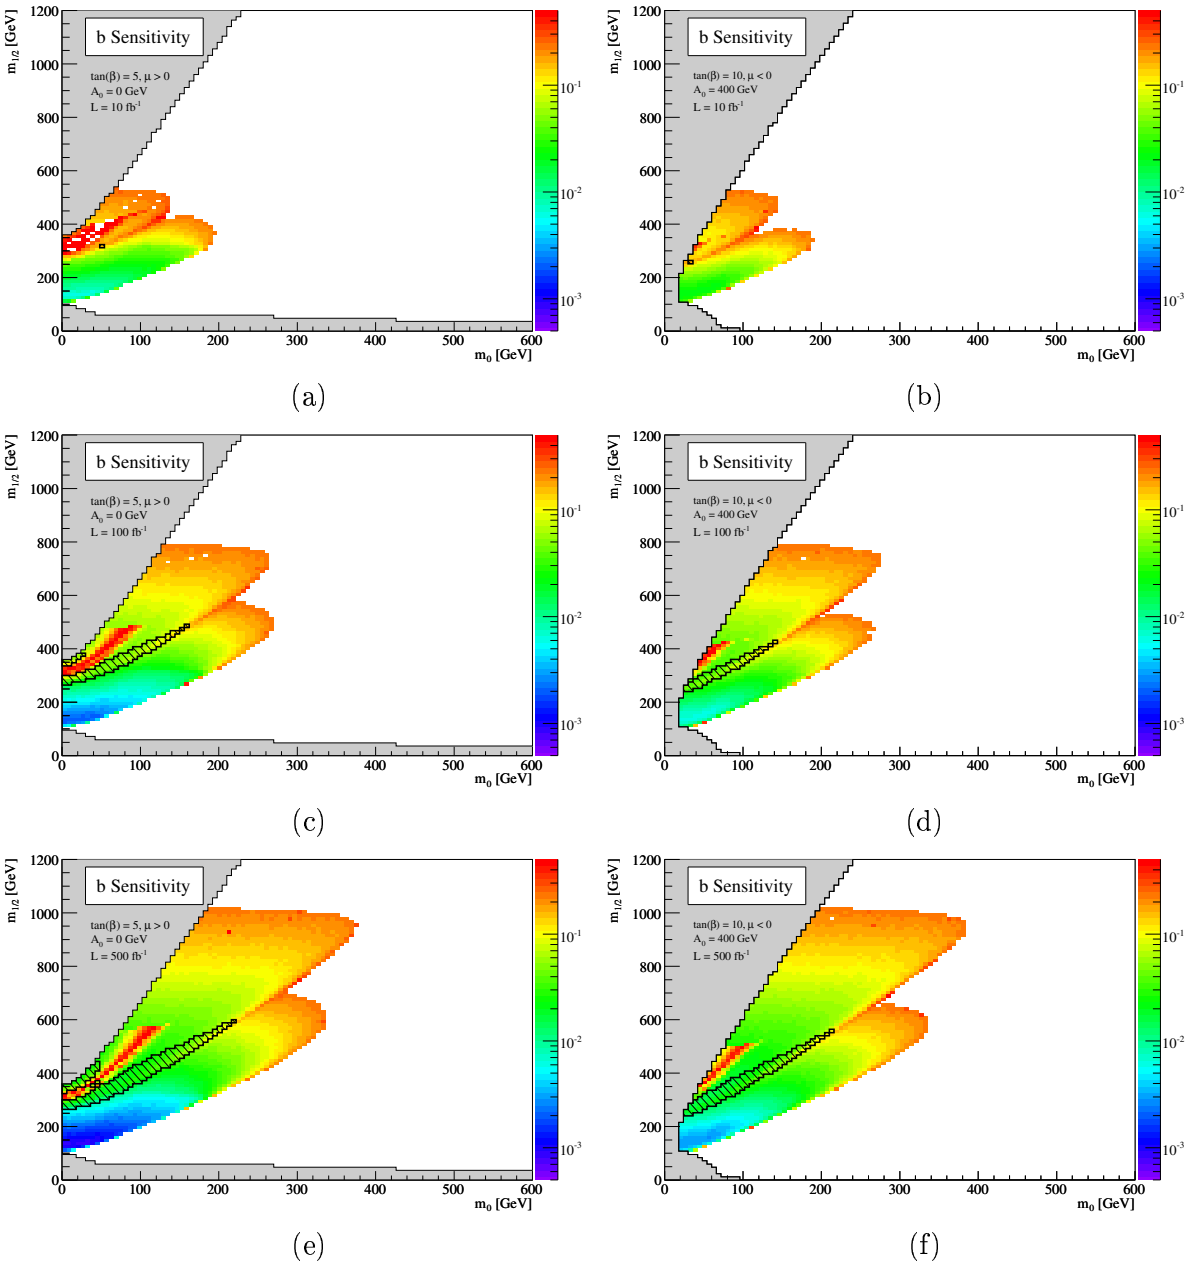

Figure 16.5: *Decay parameter sensitivity when the fitted $dP/dm_{\ell\ell}$ corresponds to $dP/d\cos\theta \propto 1 + b \cdot \cos\theta$ based upon 1000 simulated experiments at each point. At points where two edges are present, it is the values associated with the best determined of the two decay parameters that are plotted. The hatching indicates the regions where the sensitivity to the second decay parameter is better than 0.2.*

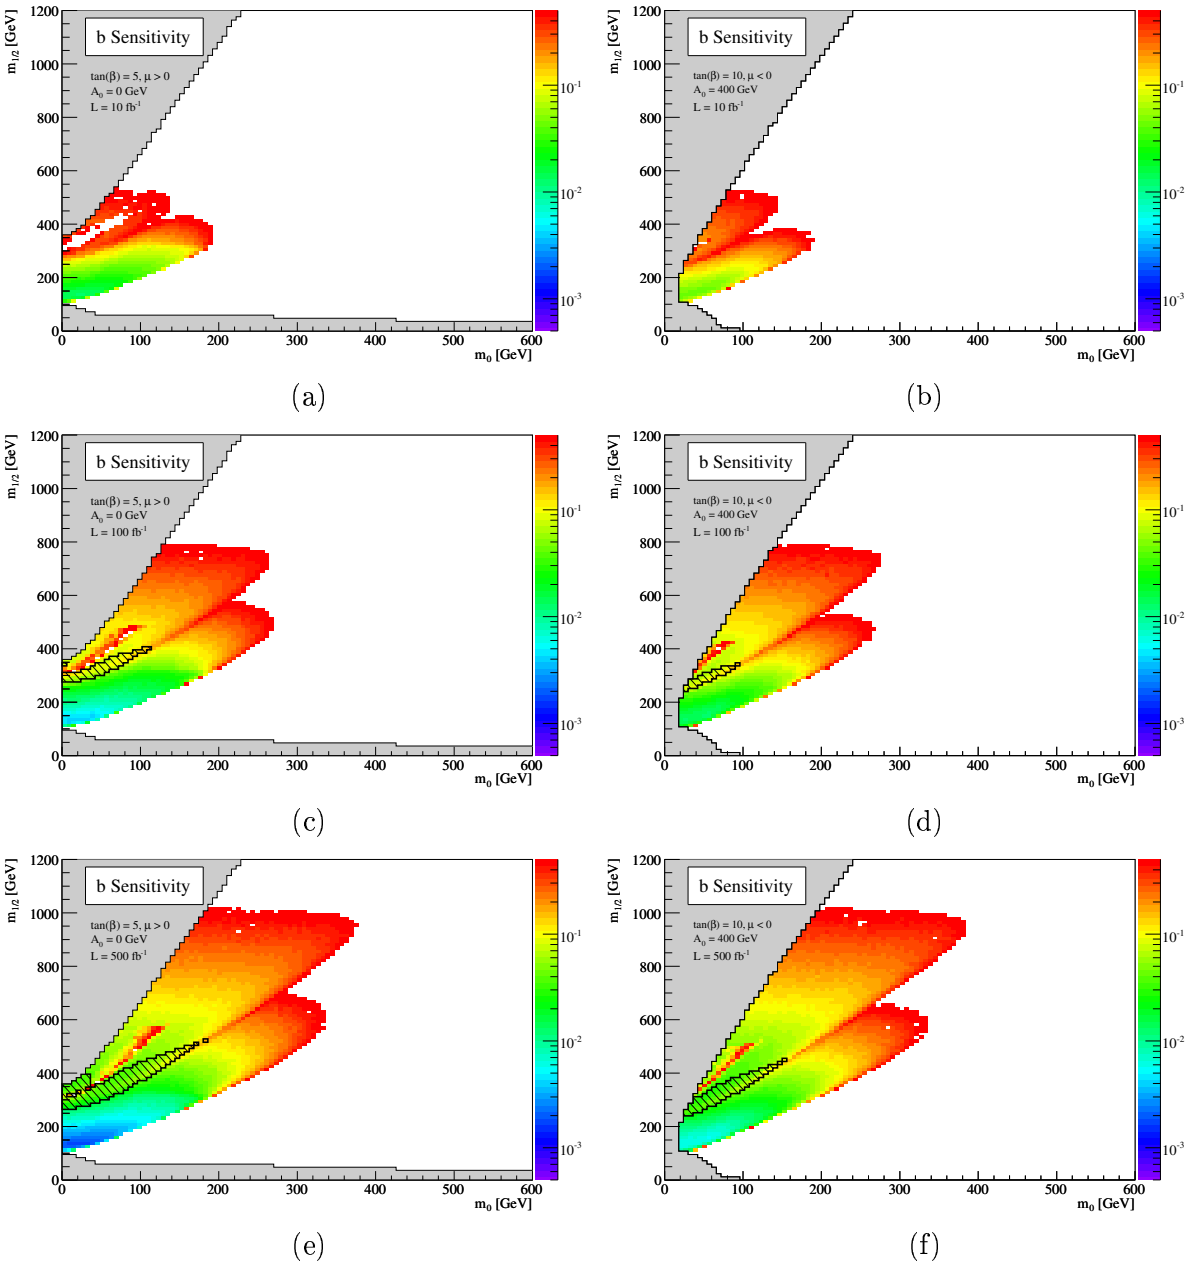

Figure 16.6: *Decay parameter sensitivity when the fitted $dP/dm_{\ell\ell}$ corresponds to $dP/d\cos\theta \propto 1 + b \cdot \cos^2\theta$ based upon 1000 simulated experiments at each point. At points where two edges are present, it is the values associated with the best determined of the two decay parameters that are plotted. The hatching indicates the regions where the sensitivity to the second decay parameter is better than 0.2.*

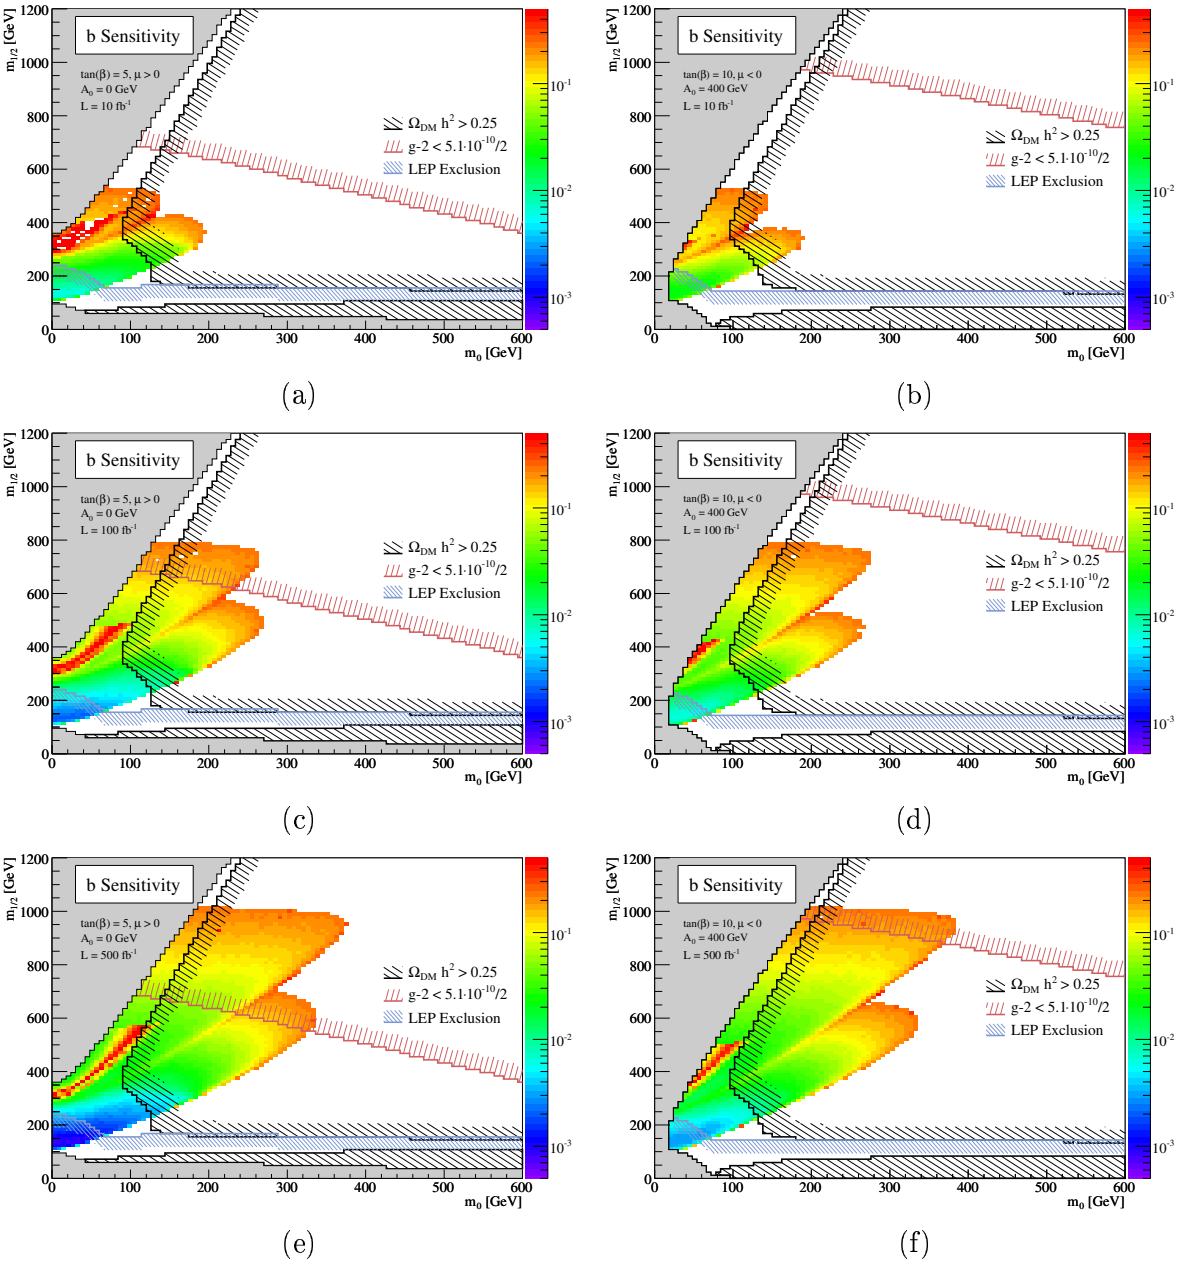

Figure 16.7: *Decay parameter sensitivity when the fitted $dP/dm_{\ell\ell}$ corresponds to $dP/d\cos\theta \propto 1 + b \cdot \cos\theta$ based upon 1000 simulated experiments at each point. At points where two edges are present, it is the values associated with the best determined of the two decay parameters that are plotted. Direct and indirect limits on the allowed regions of CMSSM space are indicated.*

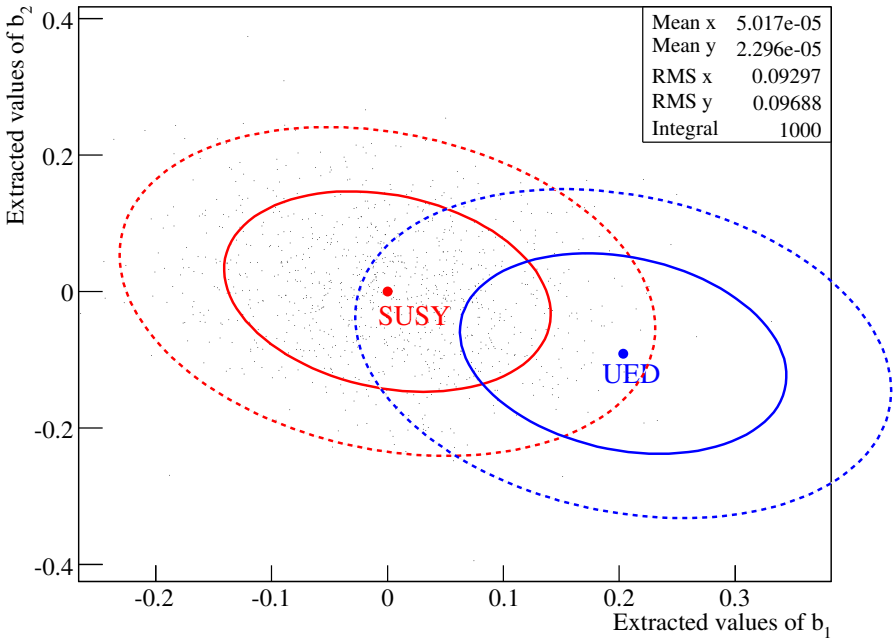

Figure 16.8: *When information about two decay parameters can be simultaneously extracted, the level of discrimination between SUSY and UED will depend on their correlation and must be inferred from error ellipses as illustrated. The shown ellipses corresponds to confidence levels of 68.3% and 95.4%, respectively. In this example there is 15.4% chance to observe a discrimination at the 3σ level.*

supersymmetric one at a given CMSSM point, it should thus be possible to get an indication of how likely the present method is to provide a given level of discrimination between these two models of new physics. In order to be able to directly compare levels of discrimination irrespectively of whether one or two edges are fitted, the results will be given as the probability for an analysis of a given experiment at the given CMSSM point and integrated luminosity to observe a separation between SUSY and UED at a given significance level or better. Significance levels will be expressed in terms of equivalent standard deviations ($2\sigma \sim 4.55\%$, $3\sigma \sim 0.27\%$, etc.).

Given that values of b_{UED} are not far from zero, figure 13.6 indicates that it is reasonable to treat the sensitivity observed in figure 16.5 as being constant with respect to the true value of b. In the case where one edge is fitted, the probability to observe an $n\sigma$ separation between SUSY ($b = 0$) and UED ($b_{\mathrm{UED}} \neq 0$) is thus given by

$$\int_{n\sigma_b - |\Delta b|}^{\infty} \frac{1}{\sqrt{2\pi}\sigma_b} e^{-\frac{t^2}{2\sigma_b^2}} dt = \frac{1}{2} + \frac{1}{2}\,\mathrm{erf}\left[\frac{1}{\sqrt{2}}\left(\frac{|\Delta b|}{\sigma_b} - n\right)\right] \tag{16.3}$$

With $\Delta b \equiv b_{\mathrm{UED}} - b_{\mathrm{SUSY}} = b_{\mathrm{UED}}$. When $n = |\Delta b|/\sigma_b$ there is thus a 50% chance for a measurement of b to provide a discrimination of n standard deviations.

When information about two decay parameters, b_1 and b_2, are extracted simultaneously the level of discrimination must be found by considering confidence regions in the (b_1, b_2) plane. In the approximation of Gaussian errors, which was found in the previous section to be rather good, the confidence regions become confidence ellipses. In addition to the extracted

single parameter error estimators, σ_{b1} and σ_{b2}, one must also extract the covariance factor (or, equivalently, the correlation) of these two parameters from the ensemble of Monte Carlo experiments, which is straight-forward. An example of such an ensemble and the extracted error ellipses is shown in figure 16.8. This time the discrimination level must be found by integrating the probability density centred at e.g. the SUSY point over the error ellipse around the UED point corresponding to a given significance level. This is ultimately done numerically.

Figure 16.9 shows the probability to obtain a 3σ separation at different CMSSM points (for reference appendix A contains plots for more $(\tan\beta, A_0, \mathrm{sign}\mu)$ points and significance levels). As expected not much is likely to be doable with 10 fb^{-1}, but with 100 fb^{-1} or 500 fb^{-1} the prospects are rather good in the regions preferred by the indirect limits. There are four different regions where it is not, and the method fails in each of them for different reasons: First, the region at high values of m_0 is obviously plagued by statistics, while no $\tilde{\chi}_2^0 \to \tilde{\ell}$ decay is kinematically accessible in the small region near $(m_0, m_{1/2}) = (100\ \mathrm{GeV}, 150\ \mathrm{GeV})$. The thin region that almost slices the one-edge region in two is plagued by having $b_{\mathrm{UED}} = 0$, and it is as such impossible for the present method to provide any discriminatory power whatsoever at such points. Finally the, by now well known, region where two edges are merged but still attempted fitted as two distinct edges is also clear. As already mentioned previously, however, it is quite possible that some discrimination could still be possible in this region.

These result ignore the fact that there might be systematics errors (most importantly from the event selection efficiency issues discussed in section 15.2). To get an idea of how important these will be, three different level of systematic uncertainties, σ_b^{sys}, on the decay parameter measurements are tested: 0.005, 0.01 and 0.03. The "final" sensitivity is found by adding σ_b^{sys} in quadrature to the statistical sensitivity already found from the Monte Carlo ensembles,

$$\sigma_b^{\mathrm{tot}} = \sqrt{\left(\sigma_b^{\mathrm{sys}}\right)^2 + \left(\sigma_b^{\mathrm{stat}}\right)^2} \tag{16.4}$$

Using these values for the decay parameter sensitivities, the probability of 3σ discrimination at 500 fb^{-1} of integrated luminosities is shown again in figure 16.10. It is clear that the results are essentially unchanged when $\sigma_b^{\mathrm{sys}} = 0.005$ and only slightly different when $\sigma_b^{\mathrm{sys}} = 0.01$. When $\sigma_b^{\mathrm{sys}} = 0.03$, however, the performance is significantly degraded. This is of course not surprising since b_{UED} in many cases is not much larger than this value. However, a value of $\sigma_b^{\mathrm{sys}} = 0.03$ is, in the opinion of the author, rather conservative due to the various handles on event selection efficiency discussed in section 15.2. Especially given that $\sigma_b^{\mathrm{stat}} \ll 0.03$ at many points, meaning that it is possible to sacrifice a lot of statistical power in order to gain a smaller systematic error if necessary (which would be automatic if efficiency parameters are included in the fit). At the ATLAS SU3 point for instance, the event selection efficiency would become essentially constant if the minimal required transverse momentum of the two leptons were lowered to 10 GeV– but doing so would likely increase the relative background level, with corresponding degradation of the statistical power of the sample.

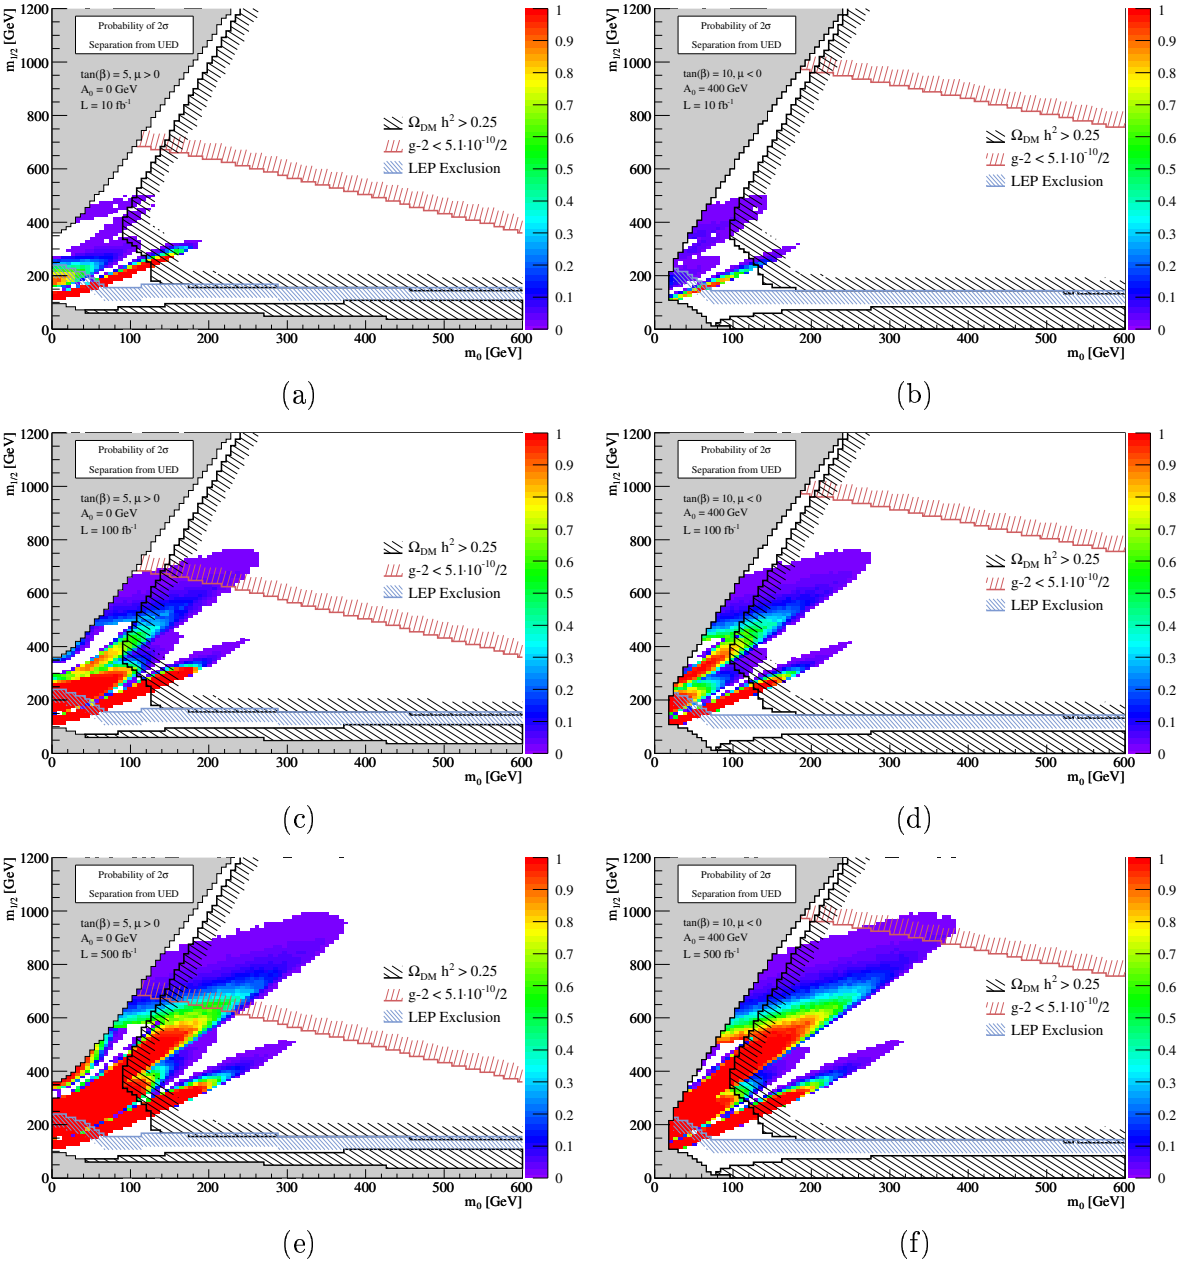

Figure 16.9: *Chance to observe a 3σ separation between SUSY and UED as discussed in the text.*

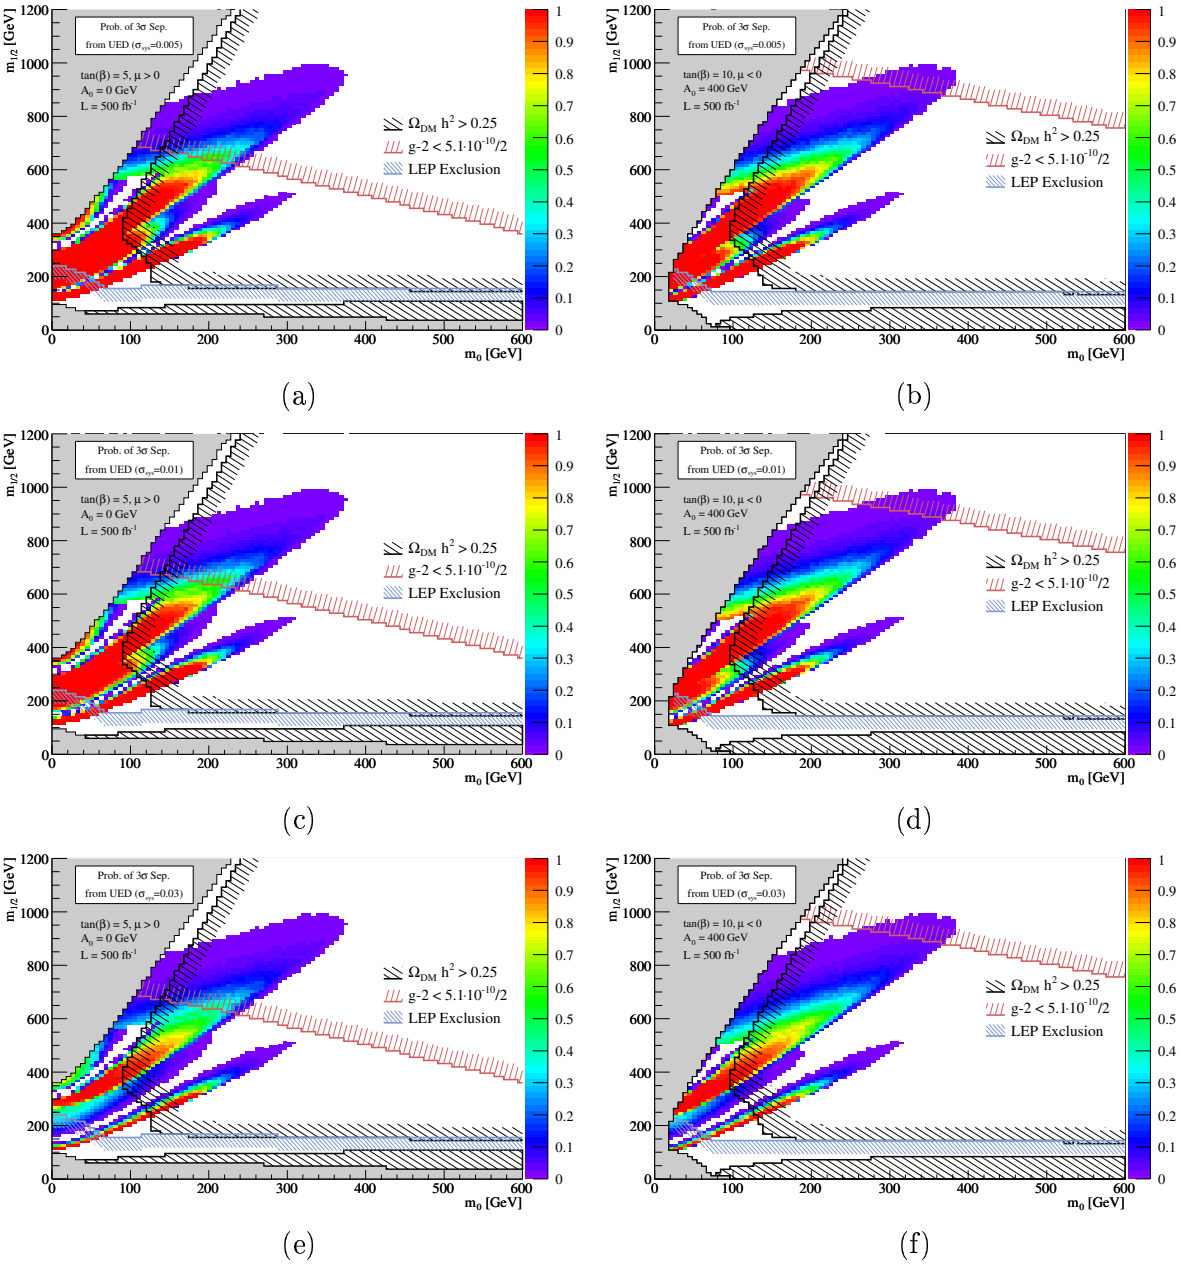

Figure 16.10: *Influence of systematic error on probability of a SUSY versus UED discrimination at the 3σ level as discussed in the text.*

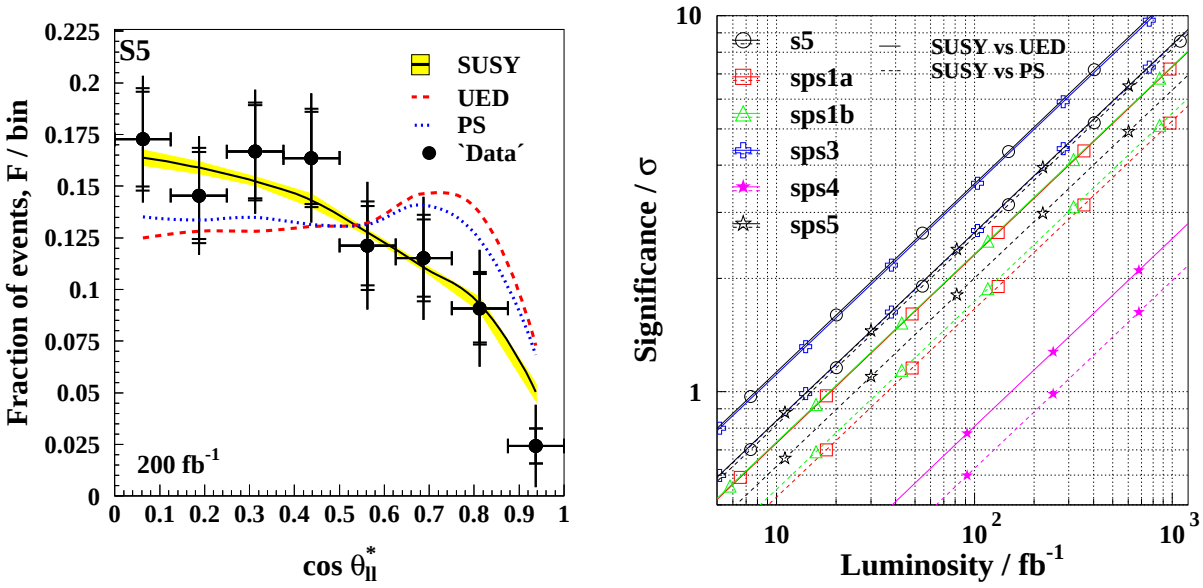

Figure 17.1: *Observation of the final-state leptons from the decays of Drell-Yan produced sleptons contains information about the slepton spin as discussed in the text. Source: [Bar06].*

17 Concluding Remarks

A method which can be used to get a handle on the spin of new physics particles with leptonic quantum numbers such as sleptons or KK excitations of SM leptons has been introduced, and the relevant formulae presented, and specifically the potential for using it to discriminate between SUSY and UED has been assessed.

The feasibility of the method will depend on the exact nature and mass spectra of new physics, but many of the challenges that can be envisioned to be important have been analysed and various methods for overcoming them have been discussed.

Other methods for accessing spins of SUSY particles have of course been suggested and studied [Bar04][DKM05][Bar06][SW05]. One might for instance use the same decay chain studied here, (cf. figure 14.1), and study the $q\ell_{\text{near}}$ invariant mass which should contain information about the spin of $\tilde{\chi}_2^0$. A challenge in such an analysis is of course to select ℓ_{near} among the two observed leptons, as well as the right jet. A different and complementary method for slepton spin determination uses Drell-Yan produced slepton-pairs (as in figure 6.3.i). Here, the decay angle in the $Z/\gamma \to \tilde{\ell}^+\tilde{\ell}^-$ decay depends on the slepton spin, and this dependency is inherited by the leptons originating from the two slepton decays. In figure 17.1.a is shown an example of how the distribution of a variable related to the rapidities of the two final state leptons, $\cos\theta^*_{\ell\ell} \equiv \tanh\left(\eta_{\ell^+} - \eta_{\ell^-}\right)/2$, is expected to differ for SUSY, UED and a pure phase space decay for reference. Figure 17.1.b shows the expected average level of discrimination between SUSY and UED (solid lines) and between SUSY and phase space (dashed lines) at six different SUSY points. The advantage of such an analysis is that it does not require the slepton to be lighter than the $\tilde{\chi}_2^0$.

Fortunately these methods are all complementary and might possibly by applied in parallel in order to solidify the understanding of the nature of new physics. The next step will hopefully be the observation of an excess of $e^+e^-/\mu^+\mu^-$ compared to $e^\pm\mu^\mp$ in events associated with jets and missing momentum at the LHC during the next couple of years.

Acknowledgements

Since the first day of my Ph.D. studies, I have enjoyed the help, wisdom and good company of a large number of individuals, to whom I am grateful. It is impossible to mention everyone here who deserves it. For this, I apologise in advance.

First of all, I would like to thank my supervisor John Renner Hansen, for offering guidance throughout these studies, and for his unwavering support whenever any obstacle has presented itself along the way.

Additionally, several other people have taken a special interest in guiding me through my studies. Thanks to Mogens Dam for helping me find a suitable project in the world of TRT offline software, and to Jakob Langgaard Nielsen for showing me the ropes when first I encountered the, at times unwieldy, ATLAS software framework. Likewise, Peter Hansen has been a constant source of constructive criticism and encouragement along the way.

In fact, I am especially indebted to Mogens and Peter for the many interesting discussions about the TRT in particular, and particle physics in general, that we have had over the last three years. Thanks for allowing my constant interruptions, whenever I felt discussing some, more or less crazy, idea at the nearest black board, or just needed an explanation of some "simple" subject.

Of course, several other people were also available for many helpful discussions. Thanks to Troels Petersen, Jørn Dines Hansen, Stefania Xella, Chafik Driouichi, and in particular Jørgen Beck Hansen whose highly engaged and enthusiastic comments I always enjoyed.

Special thanks goes to fellow Ph.D. students Esben Bryndt Klinkby, Katrine Facius and Rasmus Mackeprang, who shared my office in Copenhagen in addition to many memorable moments during the various excursions to conferences, summer schools and bars around the world, where we went together. Whether undertaking the tasks of teaching, preparing outreach talks, debugging code, solving complex formulae, or going to festivals, swimming in Esben's lake, showing disdain for Darwin by jumping out of air planes, road-tripping, preparing Christmas lunch for the group, or just working on improving our knowledge of the finer beers of the world, it was a great pleasure to share the experiences with you. I am looking forward to continuing such serious undertakings in the future.

Of course, many of the social events included participation from the rest of the group, and often with good representation from "chambre de bébé", consisting of, among others, Sascha, Martin, Henrik, Katrine, Claire, Mathieu, Zofia, and last, but not least, Jakob "Champ" Groth-Jensen.

It was a pleasure to experience the comfortable and informal atmosphere surrounding the High Energy Physics group at the Niels Bohr Institute, which is rather remarkable. Everyone, from first year students to professors, are treated with respect, and the insecurities of new members of the group are quickly vanquished by social events such as the ritual and punctual serving of beers and wine every single Friday. Such events are remarkably effective when it comes to getting everyone, whether physicists, engineers, secretaries or members of the GRID/Computing group, to work together, and are additionally highly enjoyable.

Thanks goes to Björn Nilsson and Anders Wäänanen for keeping the local computer systems in top shape, and to Daniel Kalici for help in getting my rather complex SUSY jobs to run optimally at the Steno Cluster, and thanks are also in order to Poul Henrik Damgaard, for helping an experimental physicist straightening out his understanding of the delicate points of Quantum Field Theory.

During my work on TRT software I had the pleasure of interacting with a number of

pleasant people at CERN, within and outside the TRT community. In particular, I would like to mention Fred Luehring who made me feel welcome in the big collaboration that is ATLAS, and whose kind and helpful encouragements and comments have meant a lot for my work in the TRT software community. Thanks are also in order the many helpful people I have interacted with during my work: Manuel Gallas, Joe Boudreau, Vakho Tsulaia, Andrea Dell'Acqua, Grant Gorfine and Davide Costanzo to mention but a few. Special thanks goes to Daniel Froidevaux for good advice and for the several, sometimes heated, but always interesting and constructive, discussions regarding the TRT. Special thanks also goes to Vladimir Tikhomirov, for providing the test beam data used for the noise studies in an easily accessible format, and to Peter Cwetanski, for running the special Garfield simulations used as input to the development of digitisation algorithms, and for kindly answering my various questions about the complex physics taking place in the active gas of the TRT straws.

Another special thanks goes to the people who undertook the considerable task of proof-reading parts of the present thesis, some even during their Christmas vacation: John, Esben, Jørgen, Jakob L., Troels, Peter and Fred. Without these efforts, the quality of the present thesis would certainly have suffered greatly.

Last, but far from least, I am grateful for the support that both friends and family have shown during my studies. In particular, I would like to thank my parents, Jens and Vibe, for their ever present understanding and constant support.

Aha, so it's about Big Bang is it? But that's easy! Why didn't you just say so right away instead of going on about those particles all the time?

My dad

Appendix

A Additional CMSSM Scan Results

In the pages of this appendix results similar to those presented in section 16.2 are shown, but for a total of 16 different values of the CMSSM parameters $\tan\beta$, A_0 and $\mathrm{sign}\mu$. Results are only shown for an integrated luminosity of 500 fb^{-1}.

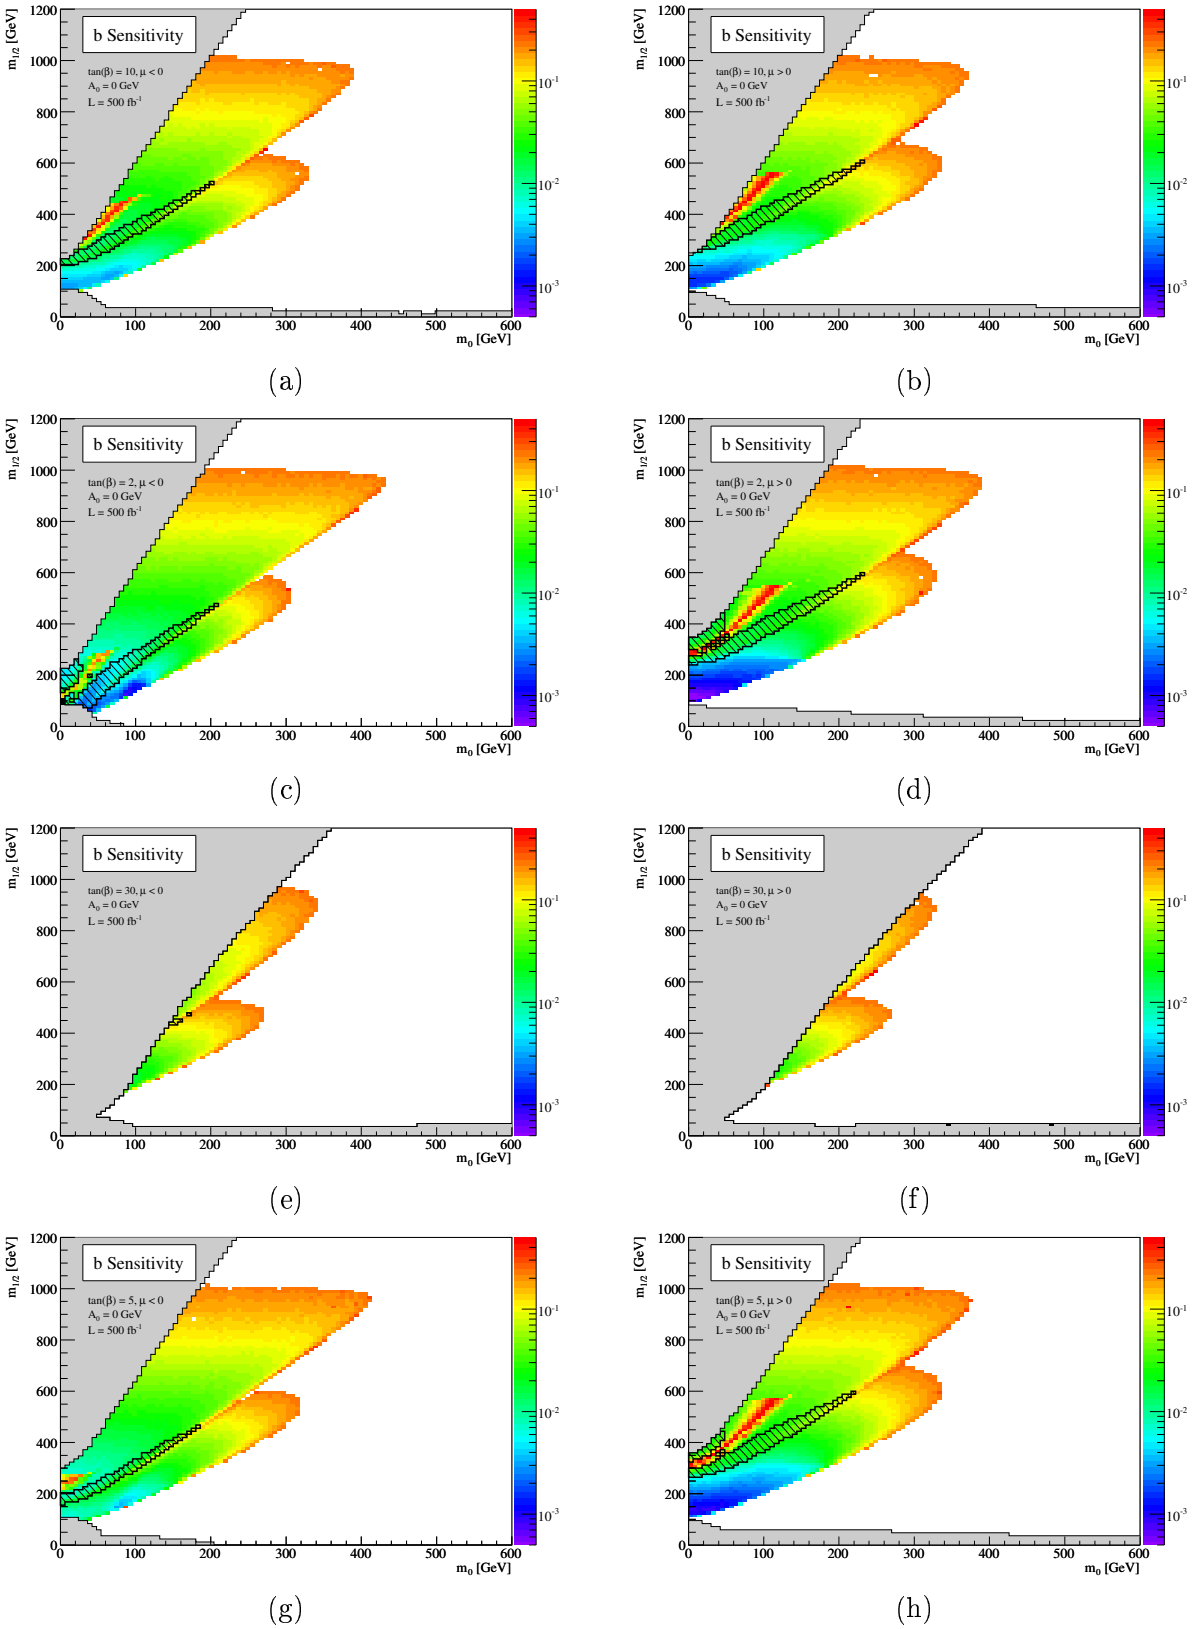

Figure A.1 [1/2]: *Decay parameter sensitivity for when the fitted $dP/dm_{\ell\ell}$ corresponds to $dP/d\cos\theta \propto 1 + b \cdot \cos\theta$ based upon 1000 simulated experiments at each point. At points where two edges are present, it is the values associated with the best determined of the two decay parameters that are plotted. The hatching indicates the regions where the sensitivity to the second decay parameter is better than 0.2.*

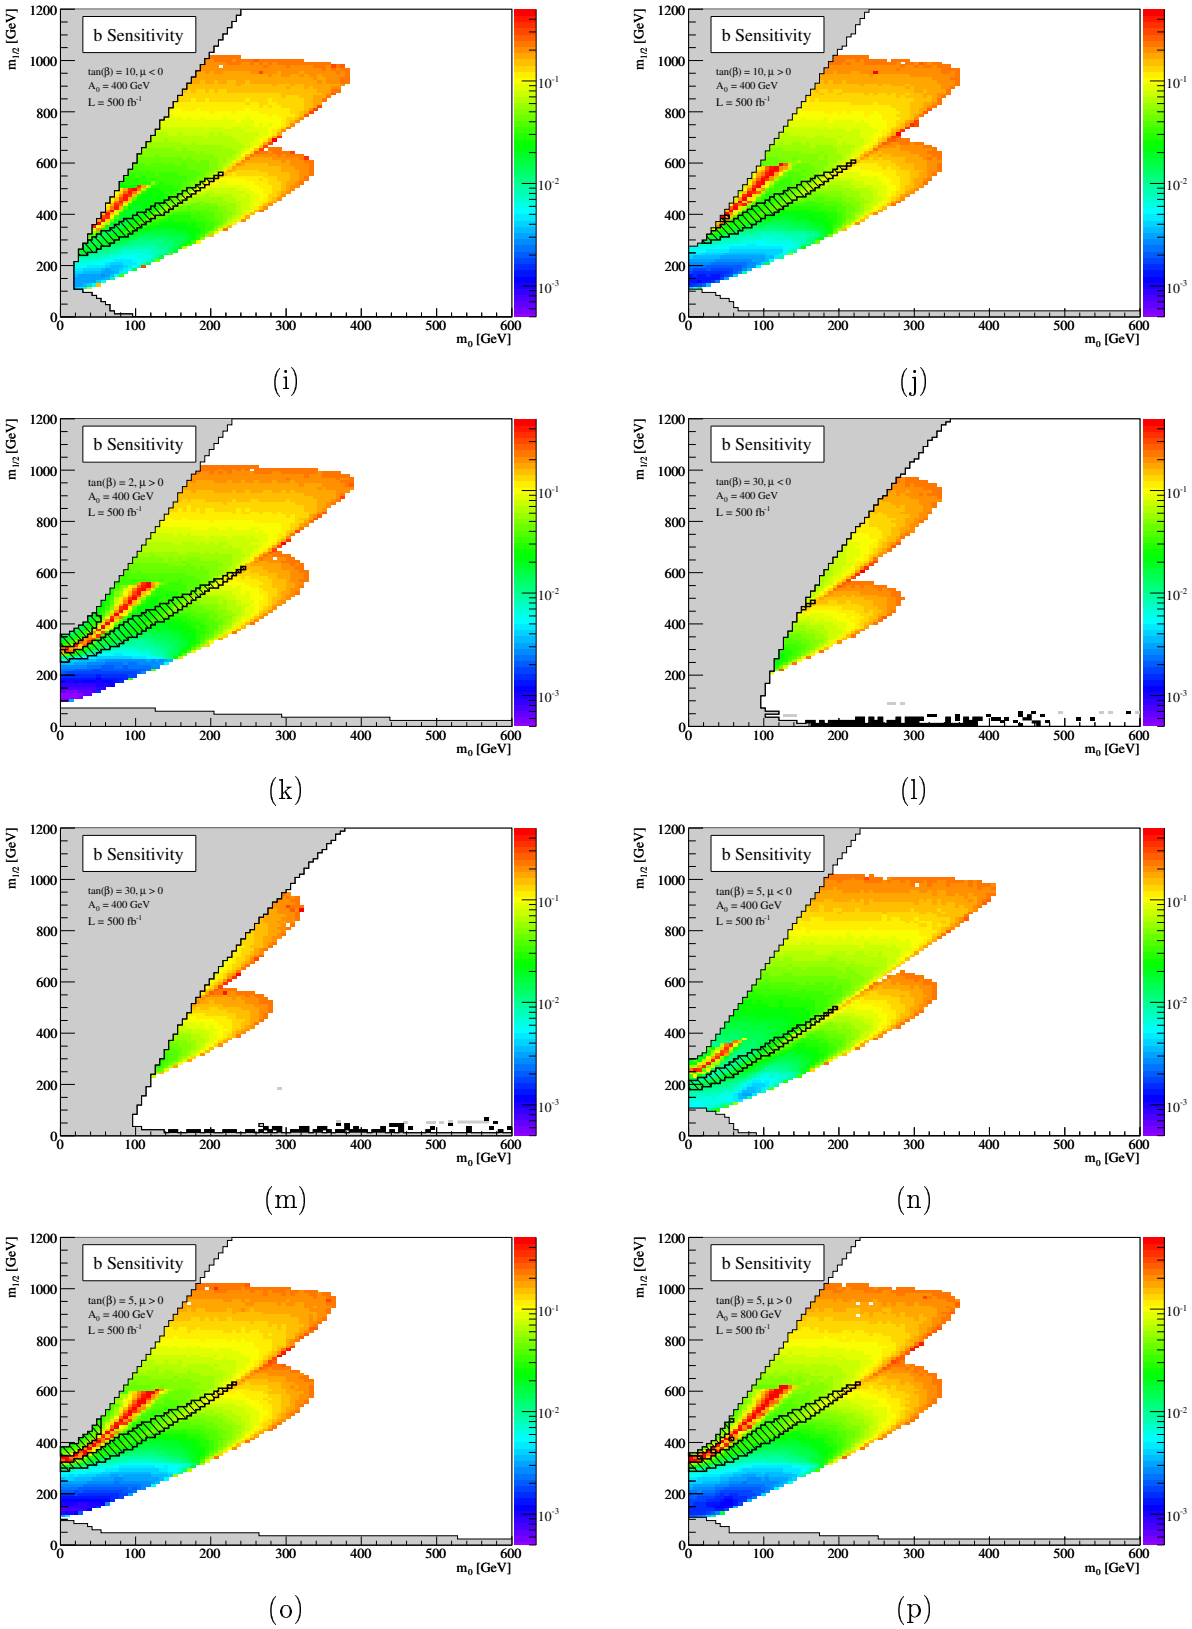

Figure A.1 [2/2]: *(continued)*

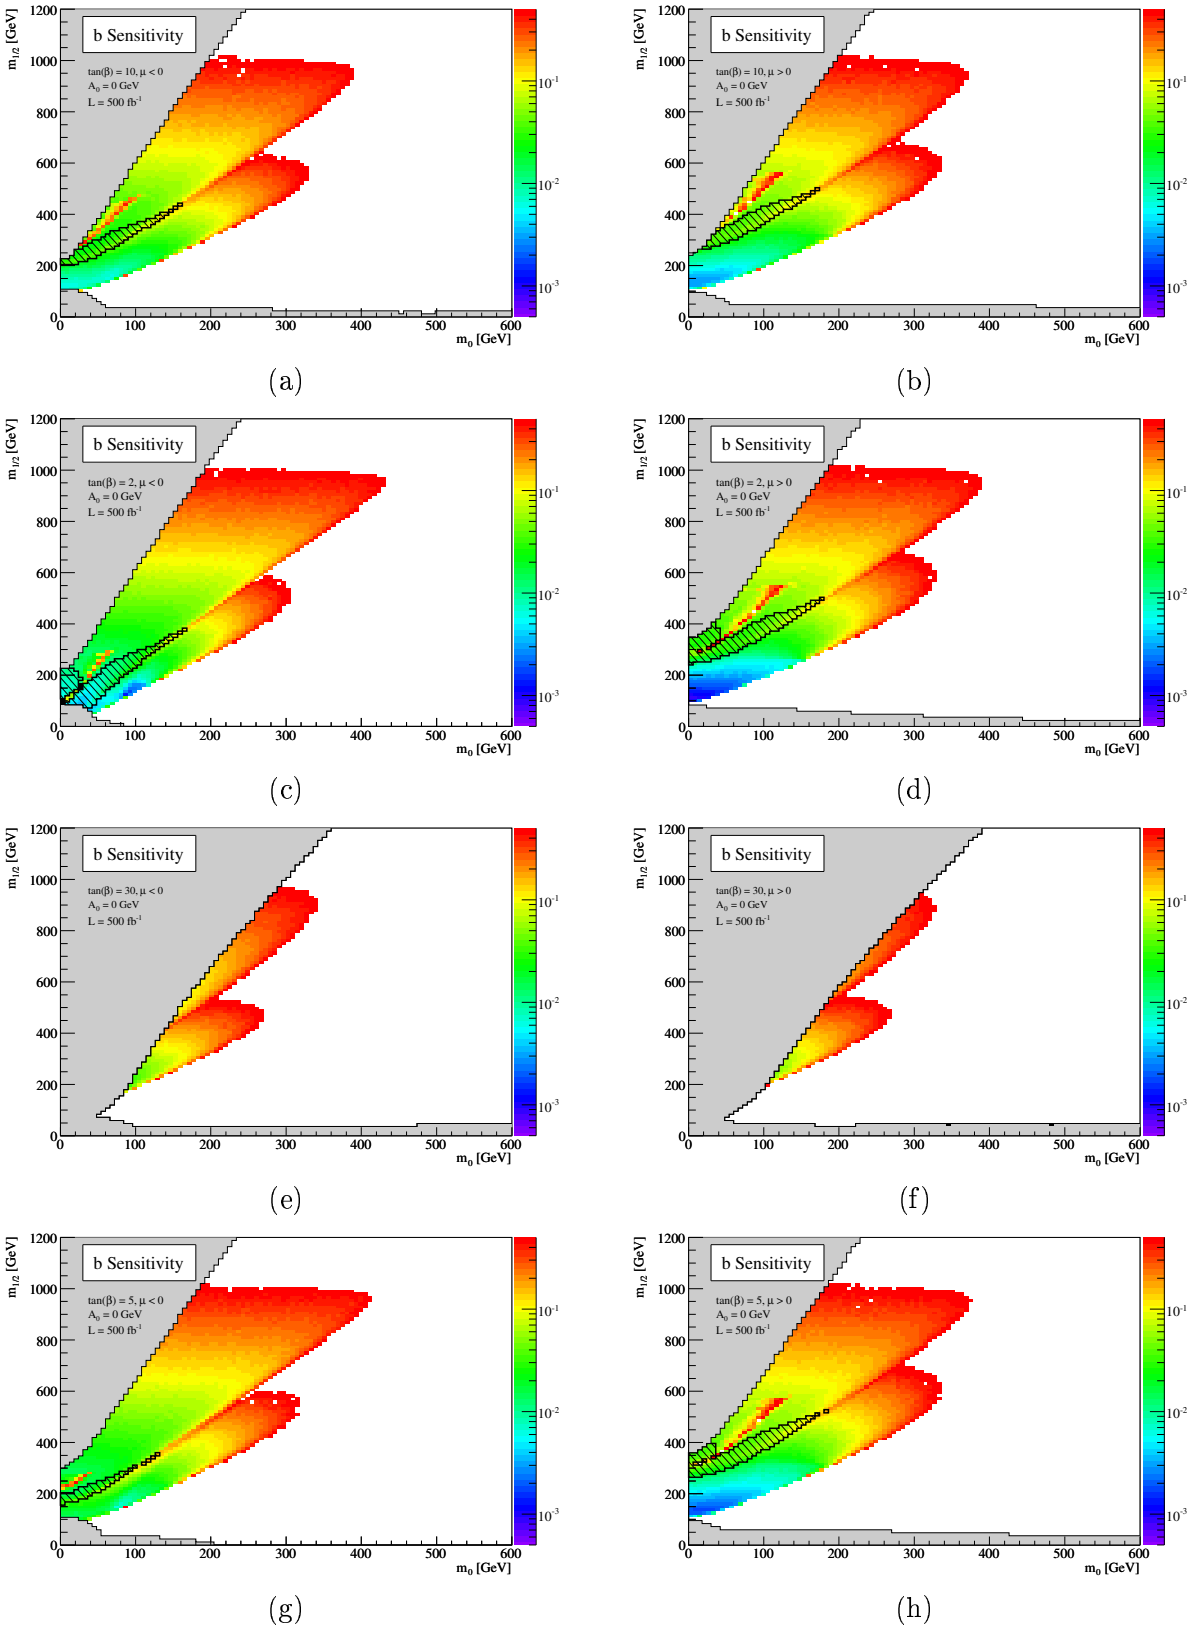

Figure A.2 [1/2]: *Decay parameter sensitivity for when the fitted $dP/dm_{\ell\ell}$ corresponds to $dP/d\cos\theta \propto 1 + b \cdot \cos^2\theta$ based upon 1000 simulated experiments at each point. At points where two edges are present, it is the values associated with the best determined of the two decay parameters that are plotted. The hatching indicates the regions where the sensitivity to the second decay parameter is better than 0.2.*

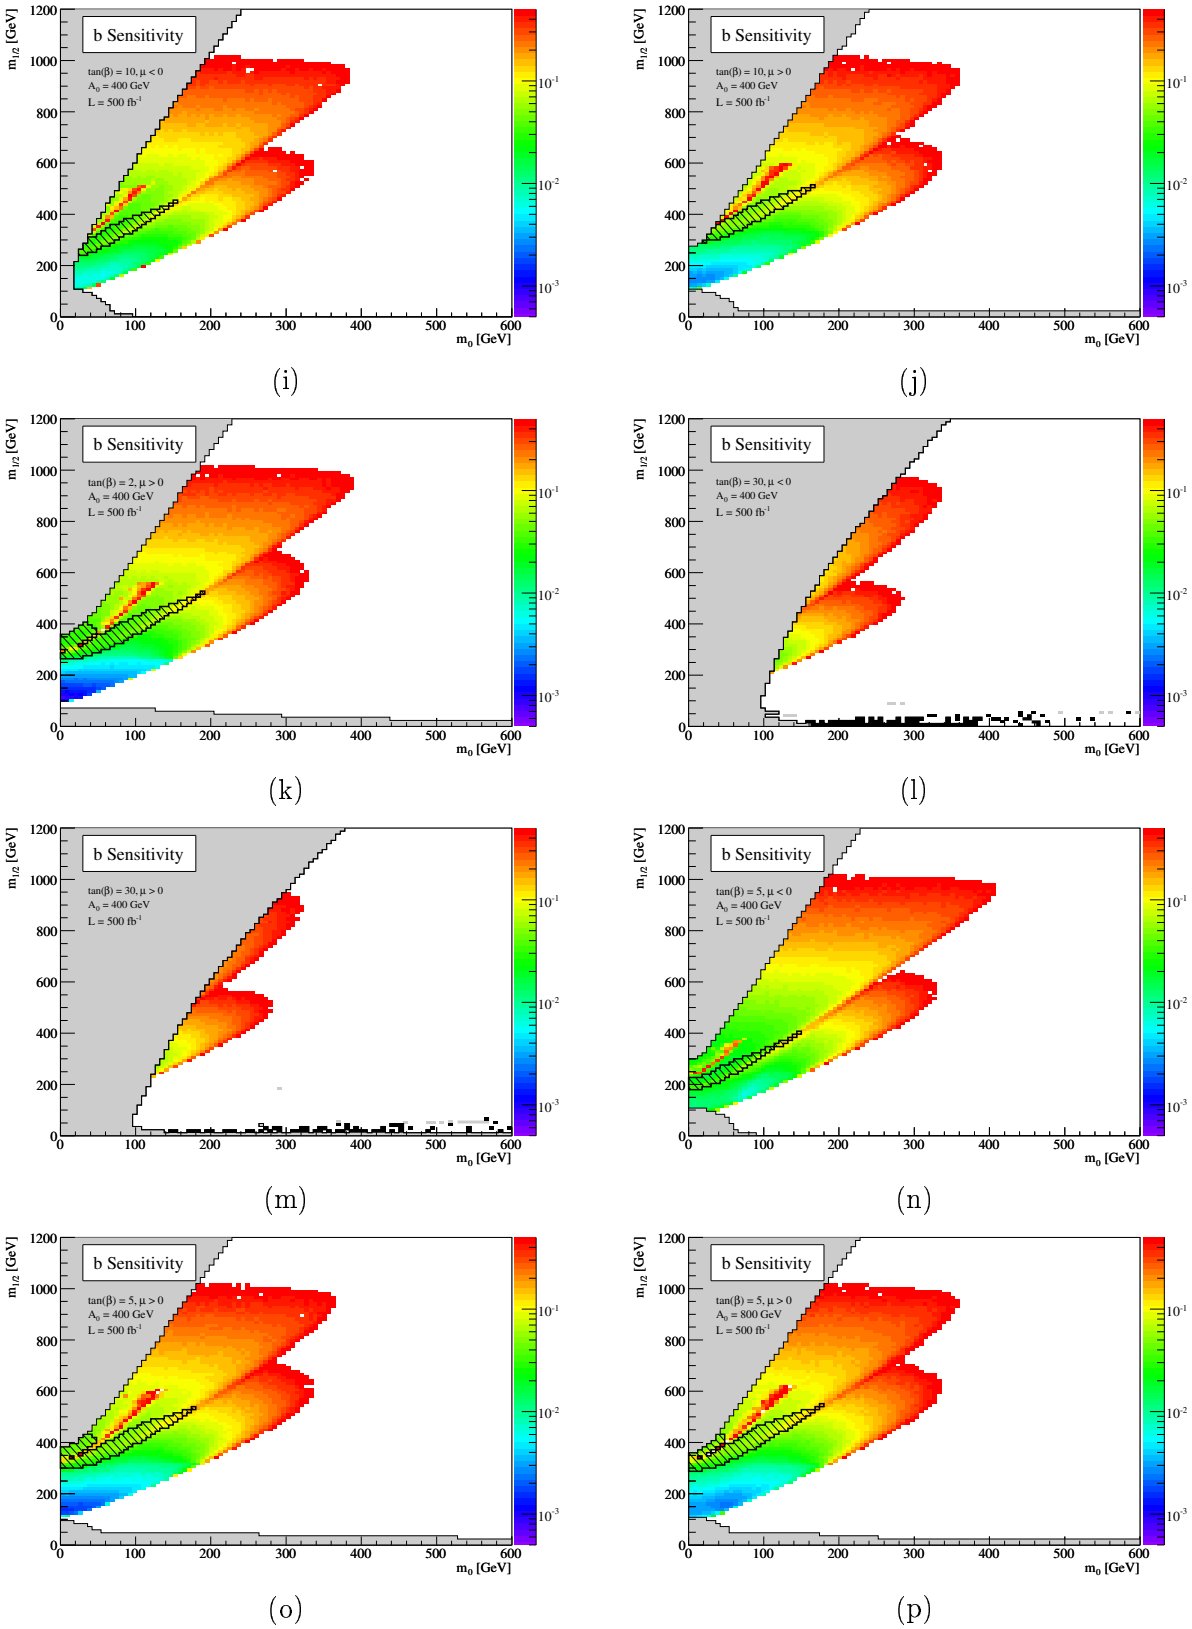

Figure A.2 [2/2]: *(continued)*

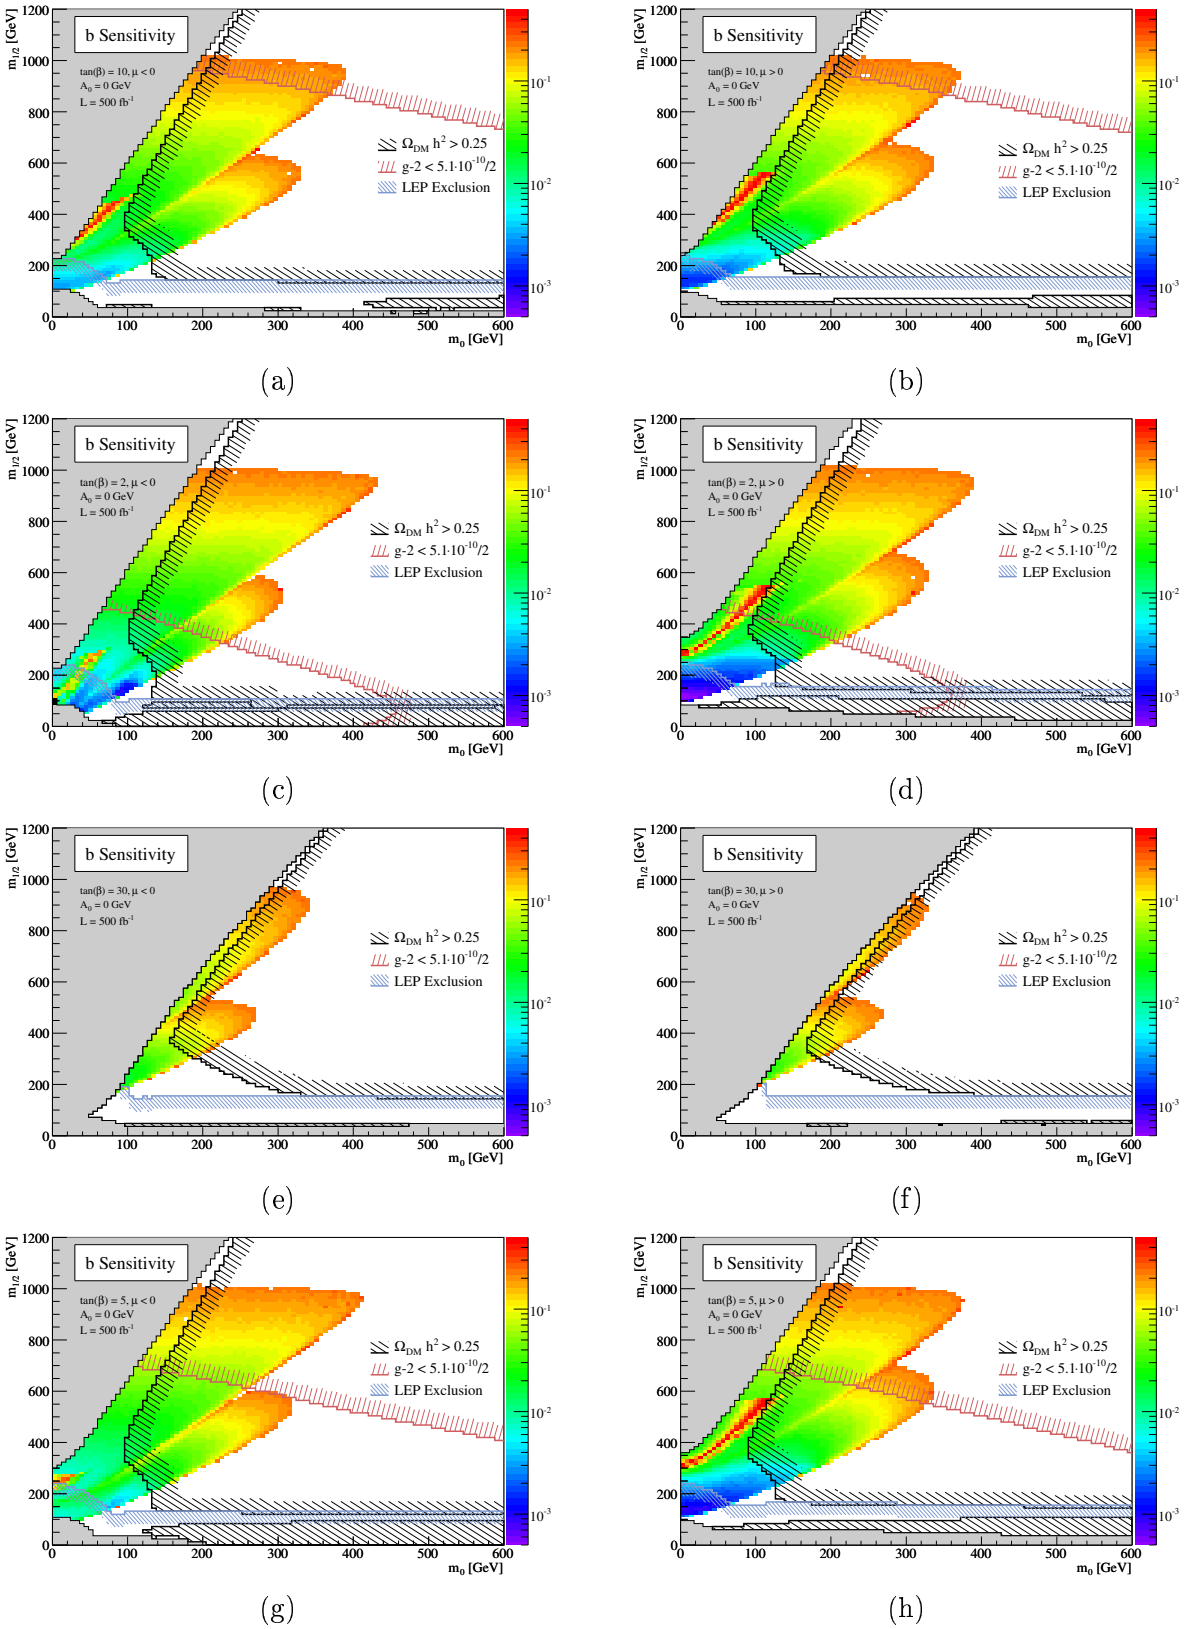

Figure A.3 [1/2]: *Decay parameter sensitivity when the fitted $dP/dm_{\ell\ell}$ corresponds to $dP/d\cos\theta \propto 1 + b \cdot \cos\theta$ based upon 1000 simulated experiments at each point. At points where two edges are present, it is the values associated with the best determined of the two decay parameters that are plotted. Direct and indirect limits on the allowed regions of CMSSM space are indicated.*

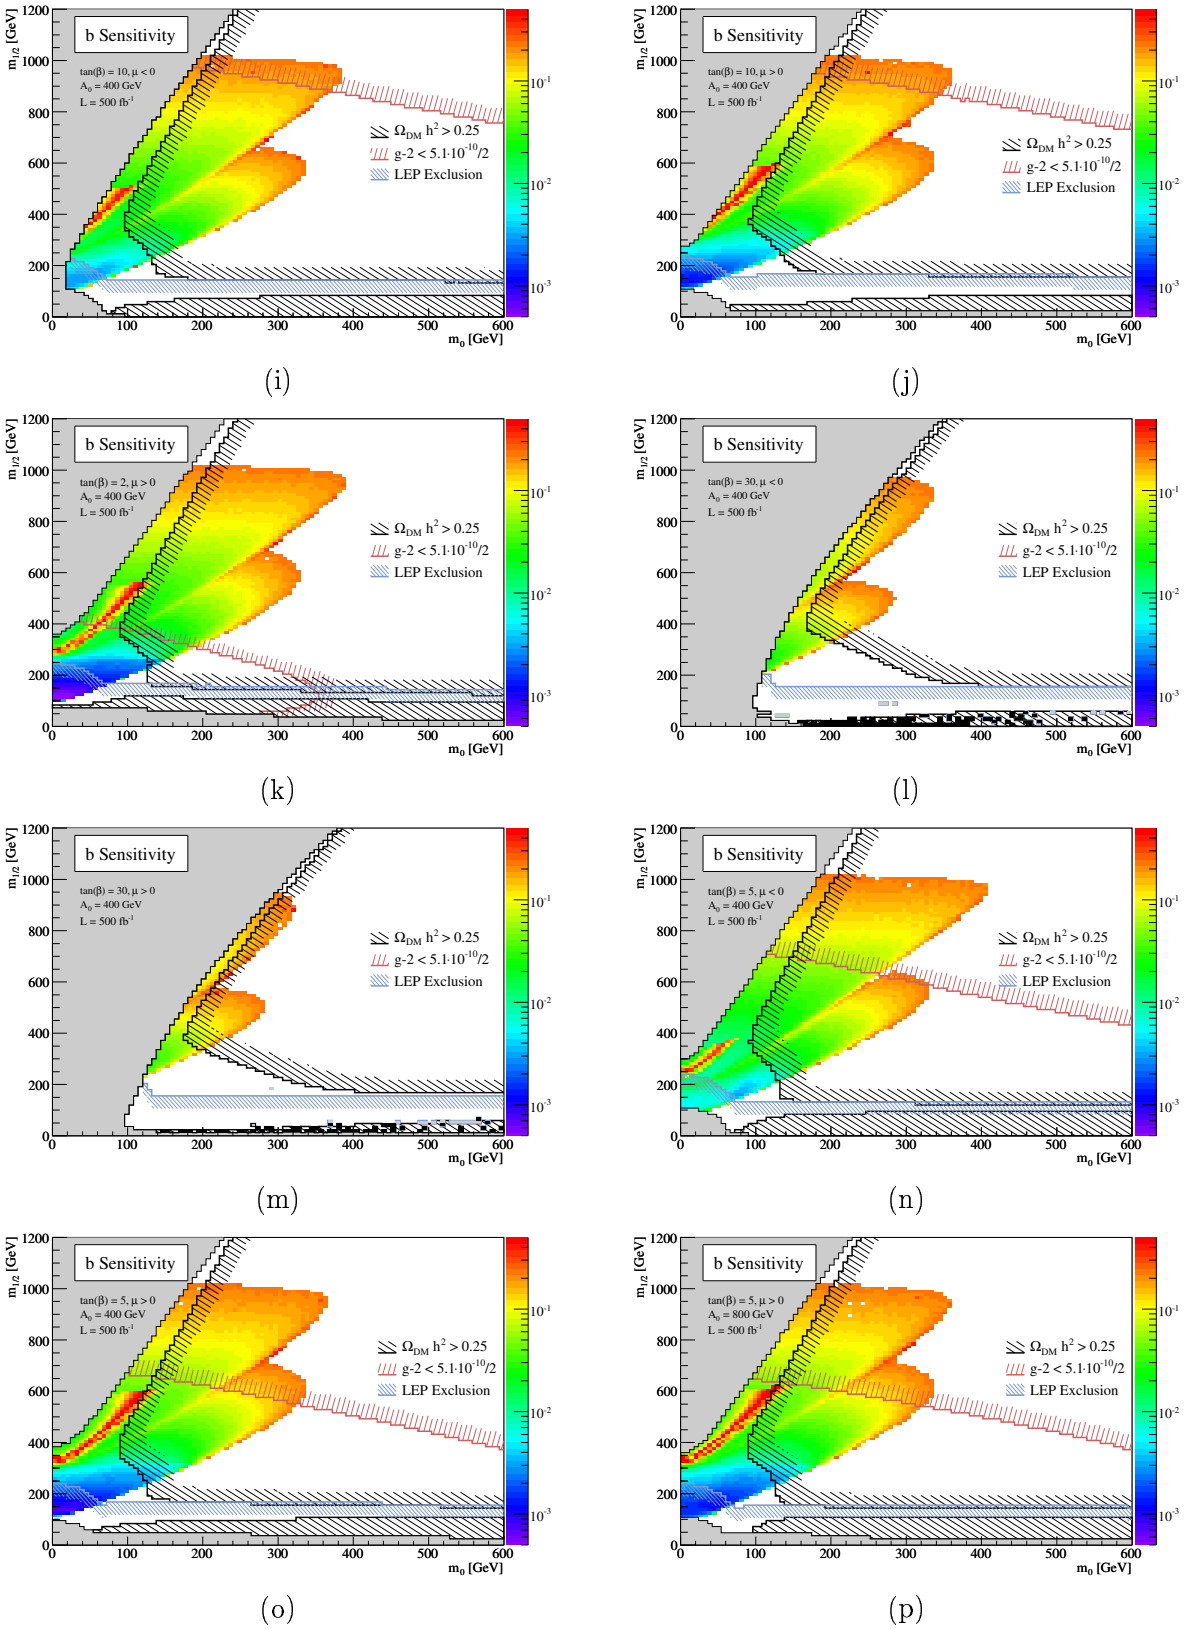

Figure A.3 [2/2]: *(continued)*

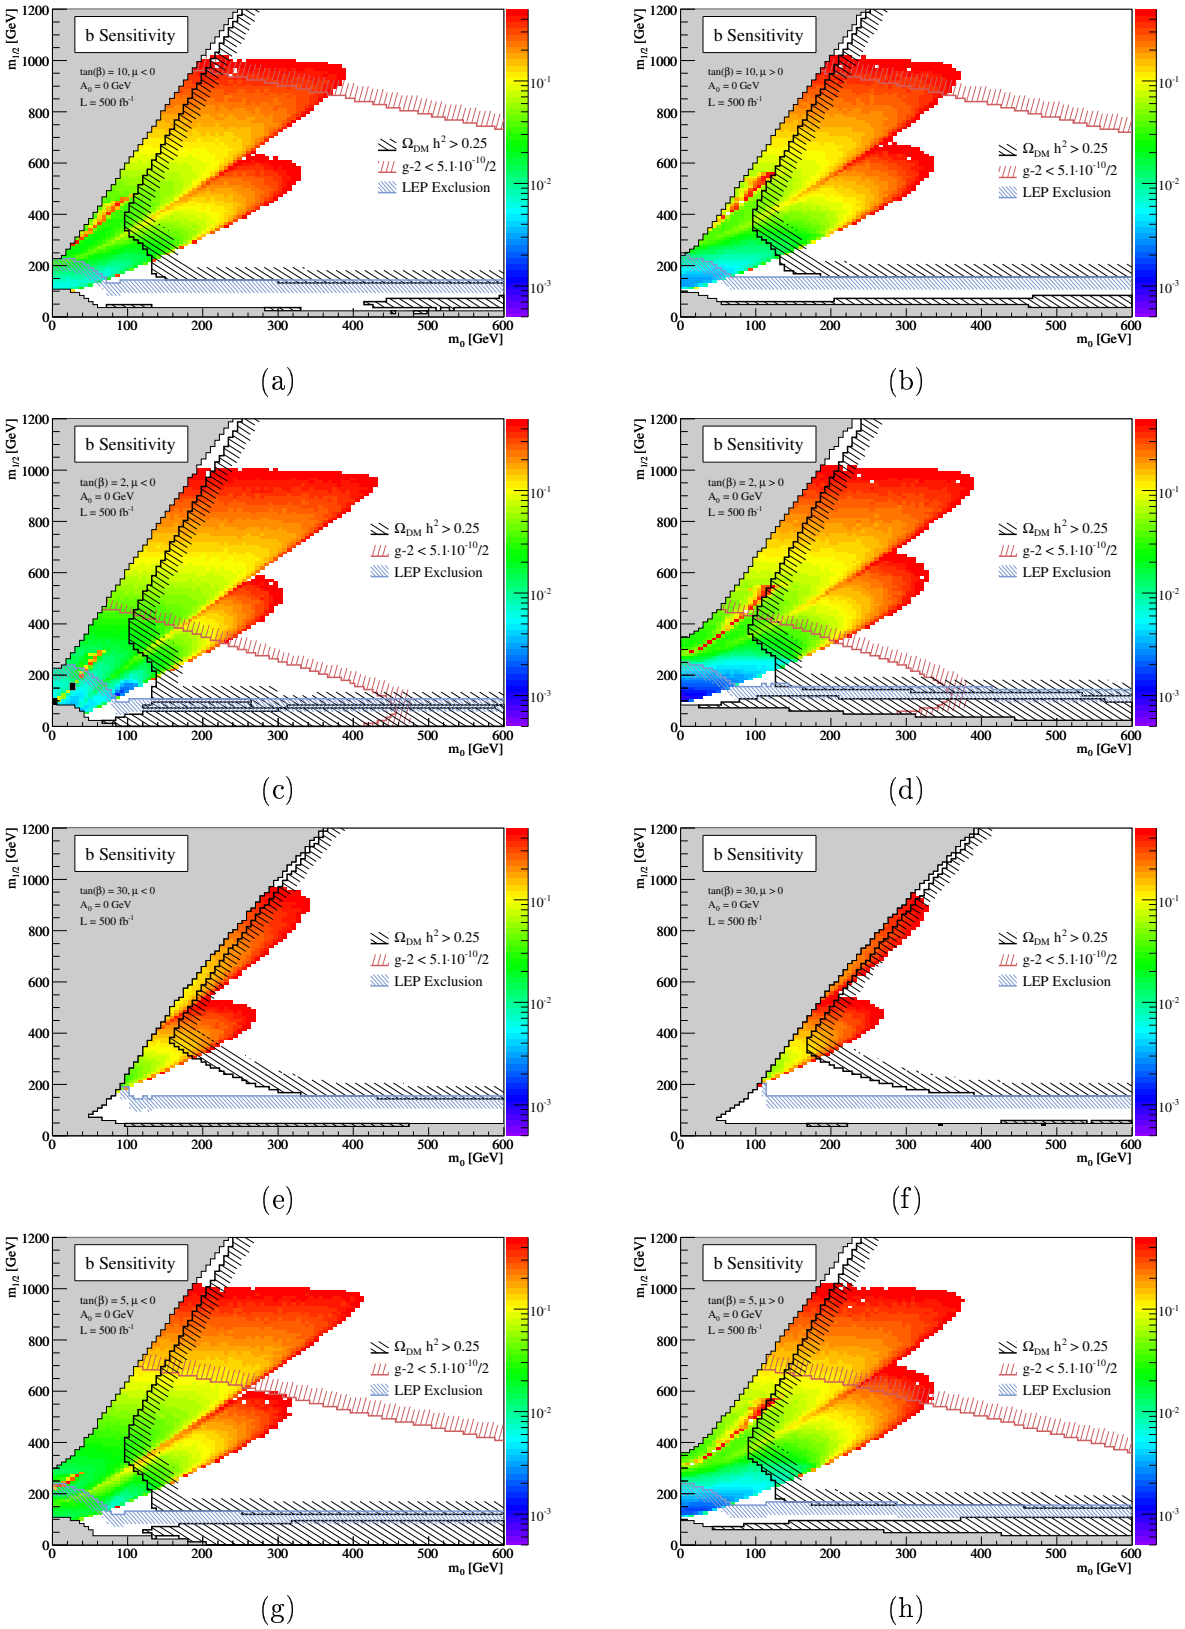

Figure A.4 [1/2]: *Decay parameter sensitivity when the fitted $dP/dm_{\ell\ell}$ corresponds to $dP/d\cos\theta \propto 1 + b \cdot \cos^2\theta$ based upon 1000 simulated experiments at each point. At points where two edges are present, it is the values associated with the best determined of the two decay parameters that are plotted. Direct and indirect limits on the allowed regions of CMSSM space are indicated.*

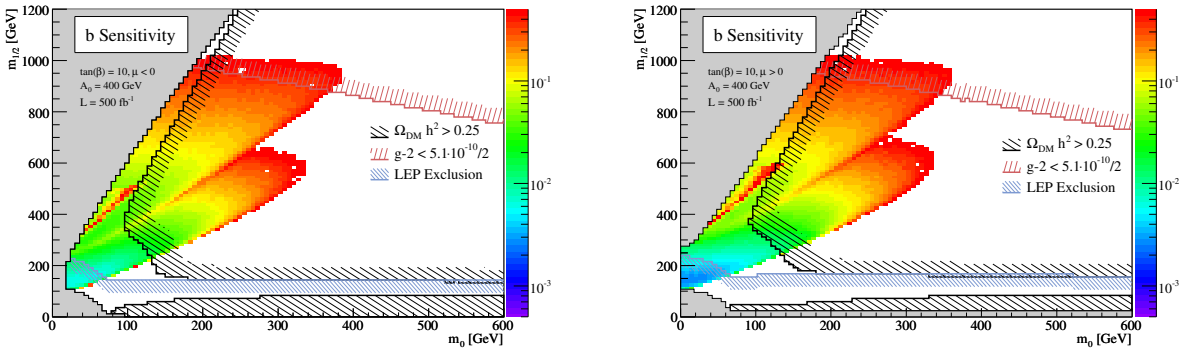

(i)

(j)

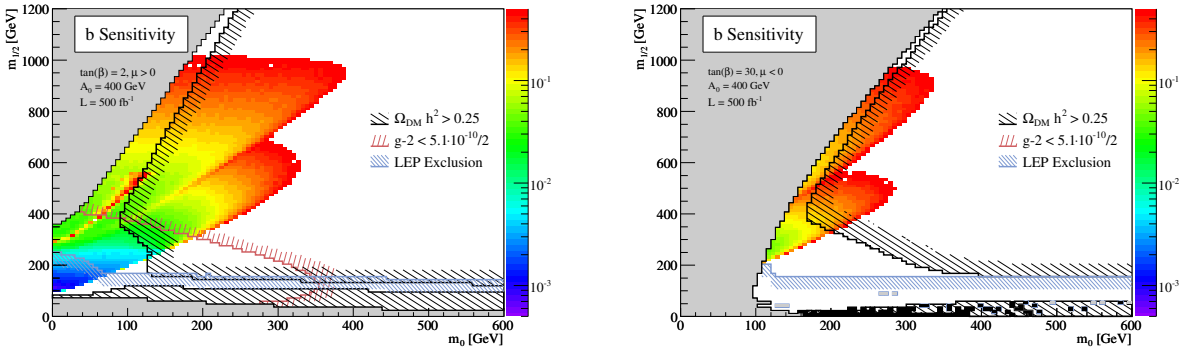

(k)

(l)

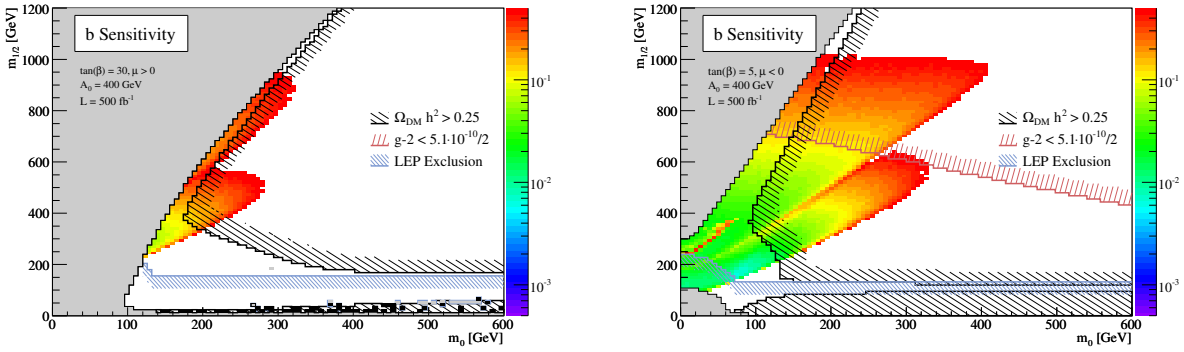

(m)

(n)

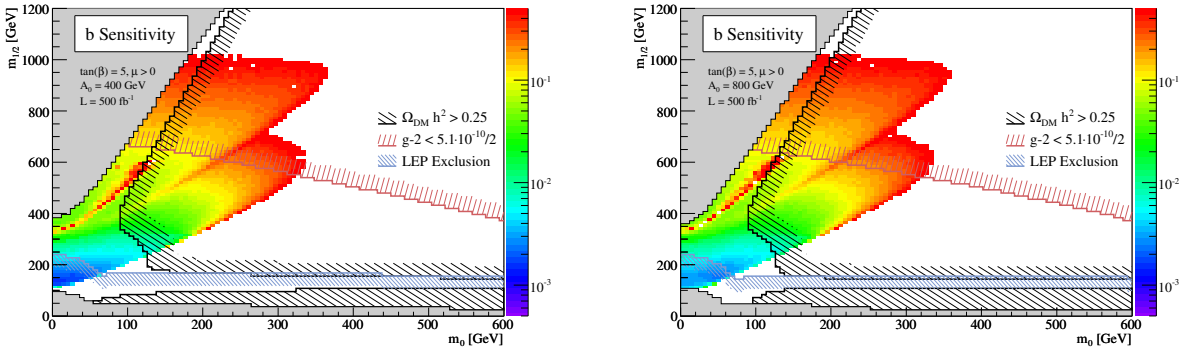

(o)

(p)

Figure A.4 [2/2]: *(continued)*

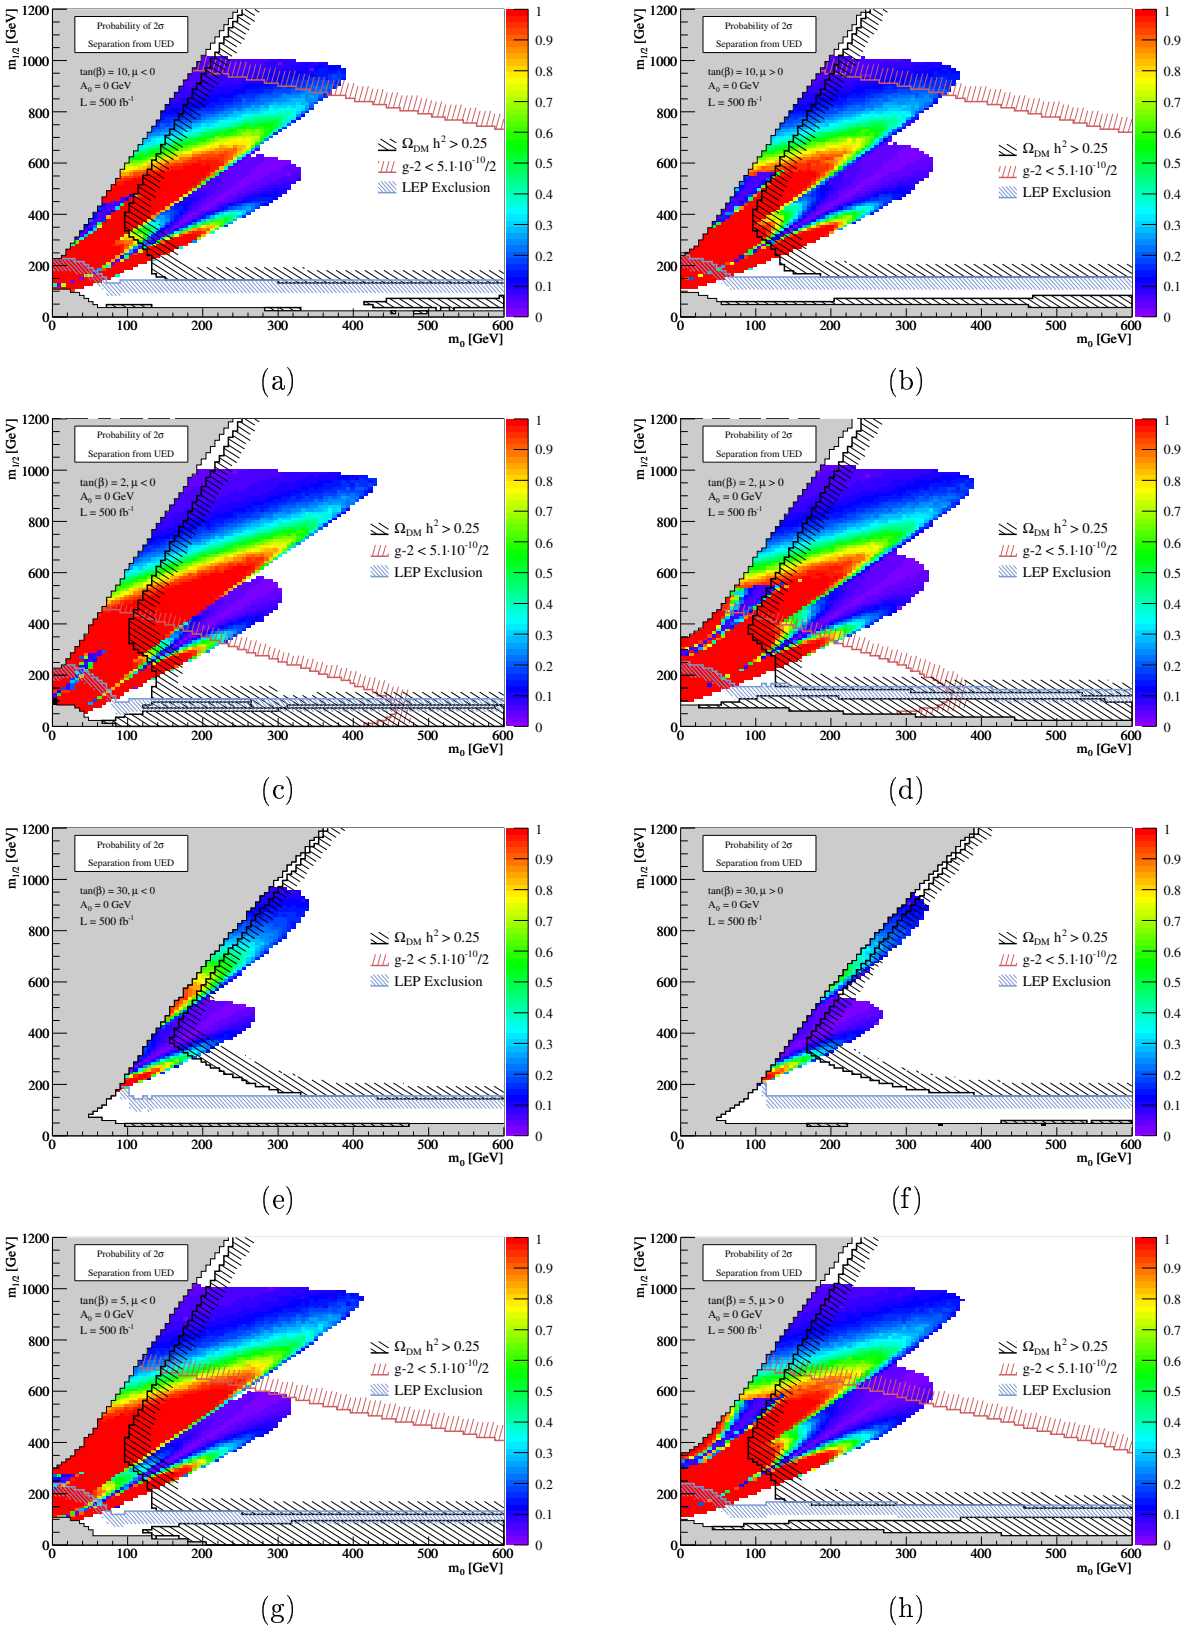

Figure A.5 [1/2]: *Chance to observe a 2σ separation between SUSY and UED.*

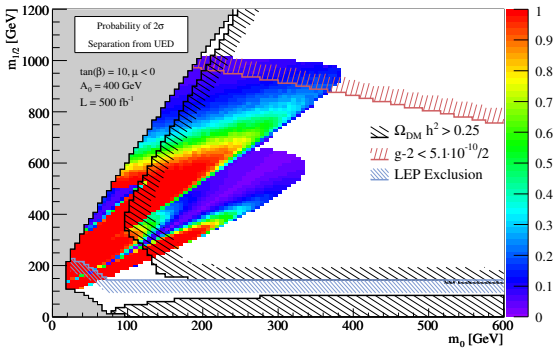

(i)

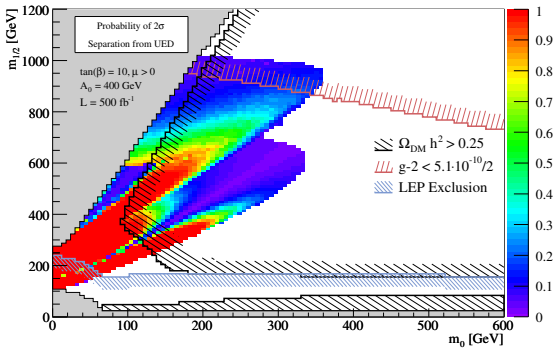

(j)

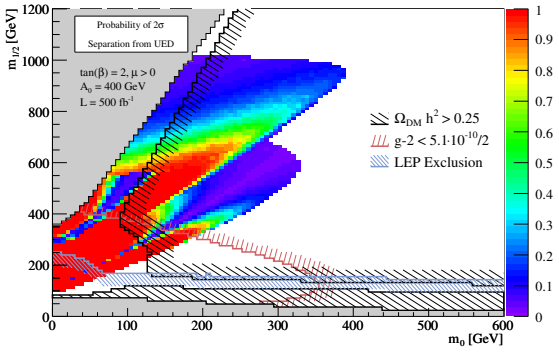

(k)

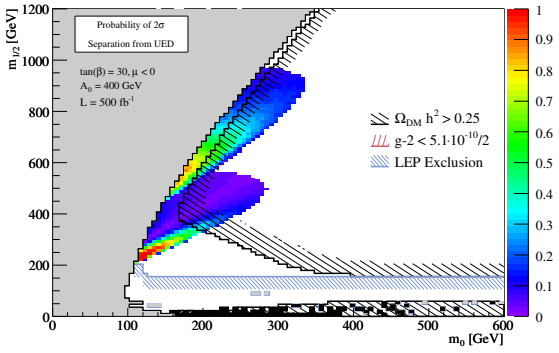

(l)

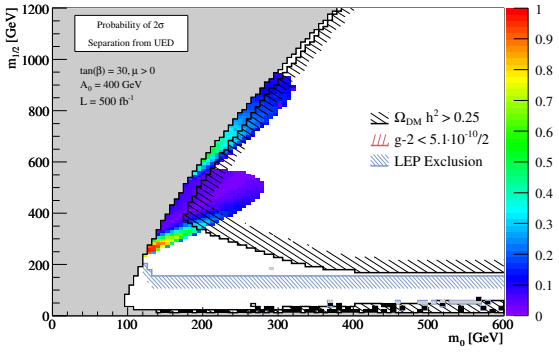

(m)

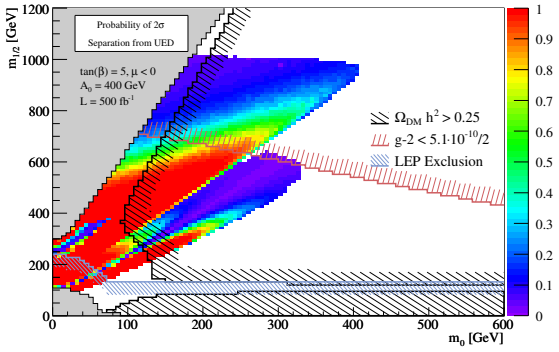

(n)

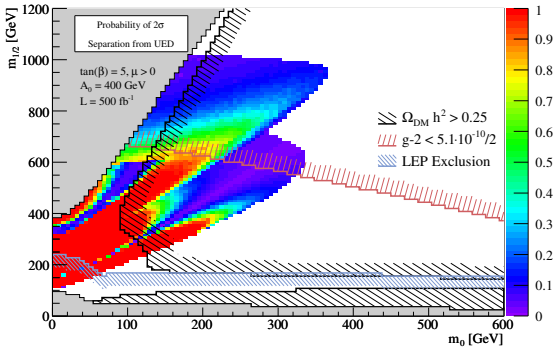

(o)

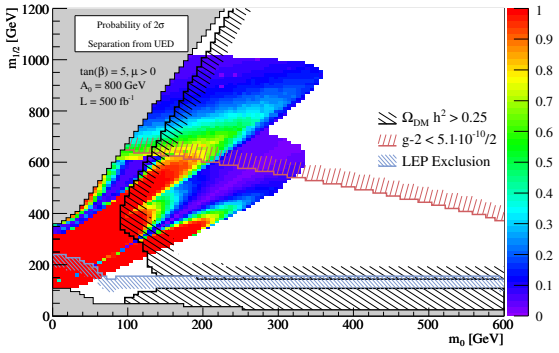

(p)

Figure A.5 [2/2]: *(continued)*

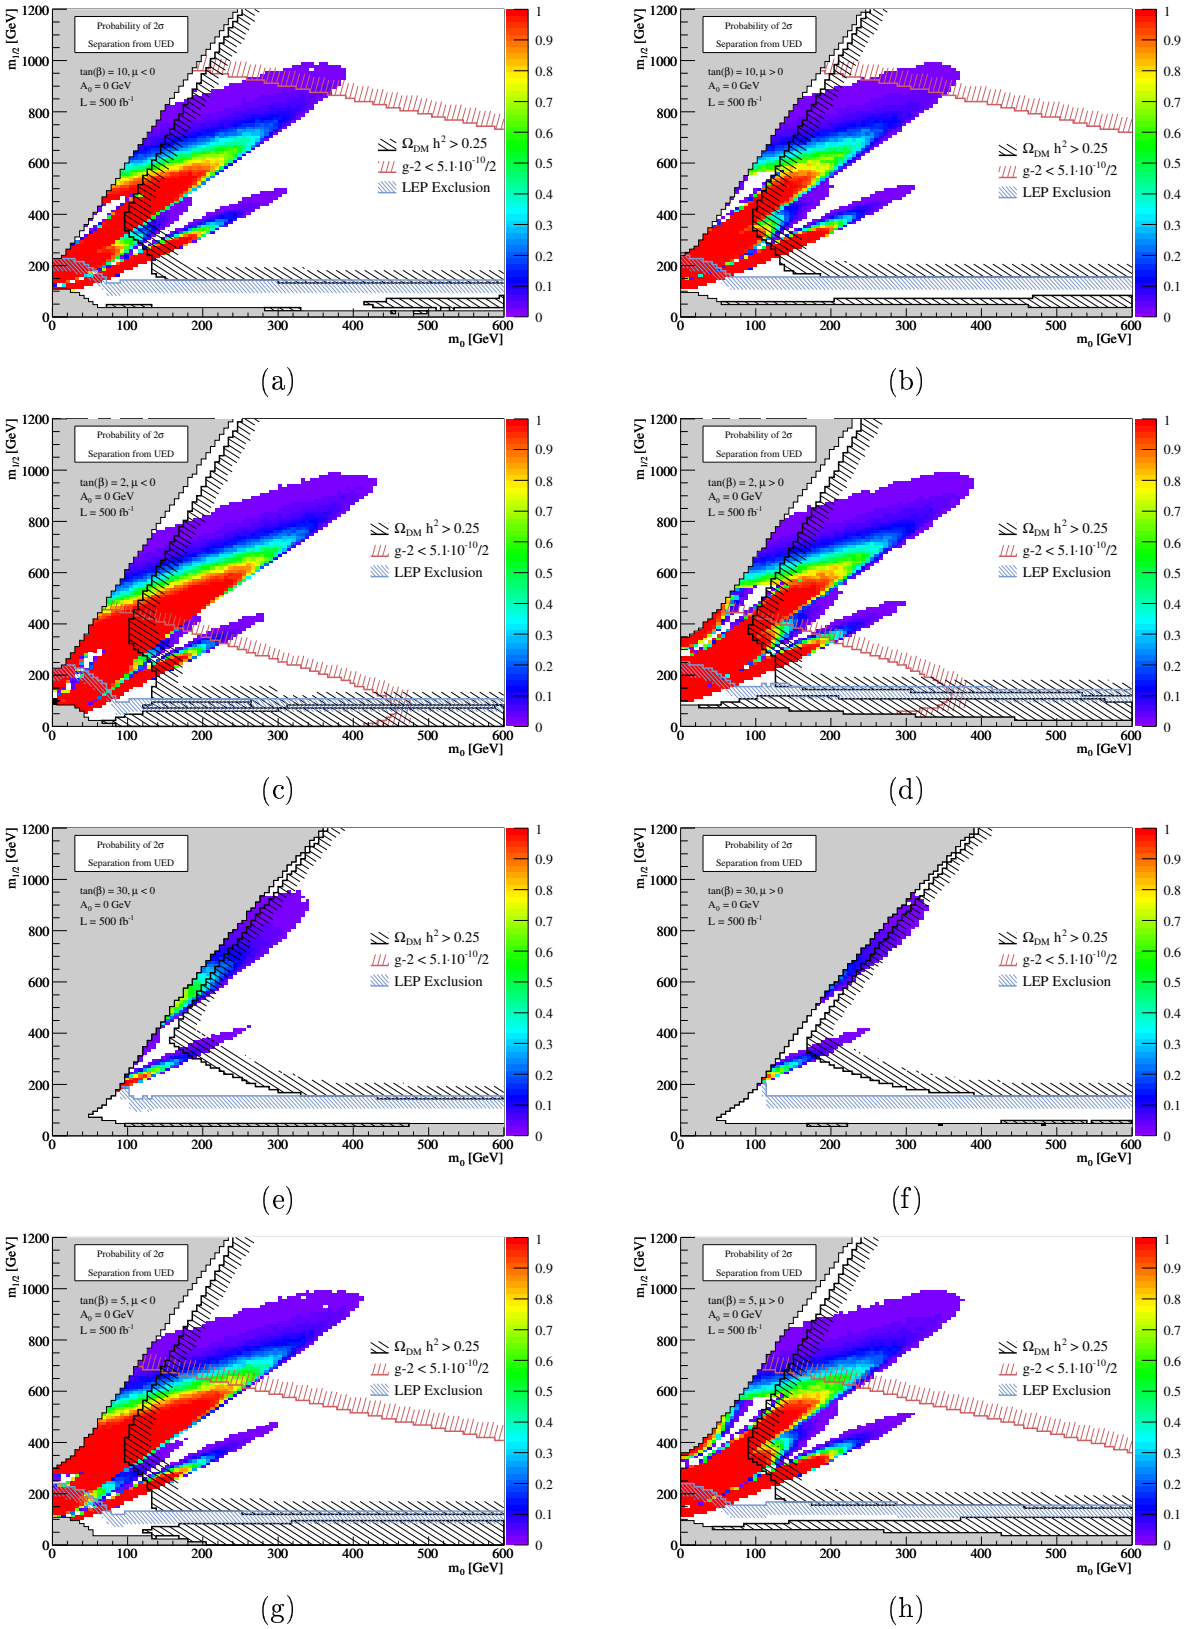

Figure A.6 [1/2]: *Chance to observe a 3σ separation between SUSY and UED.*

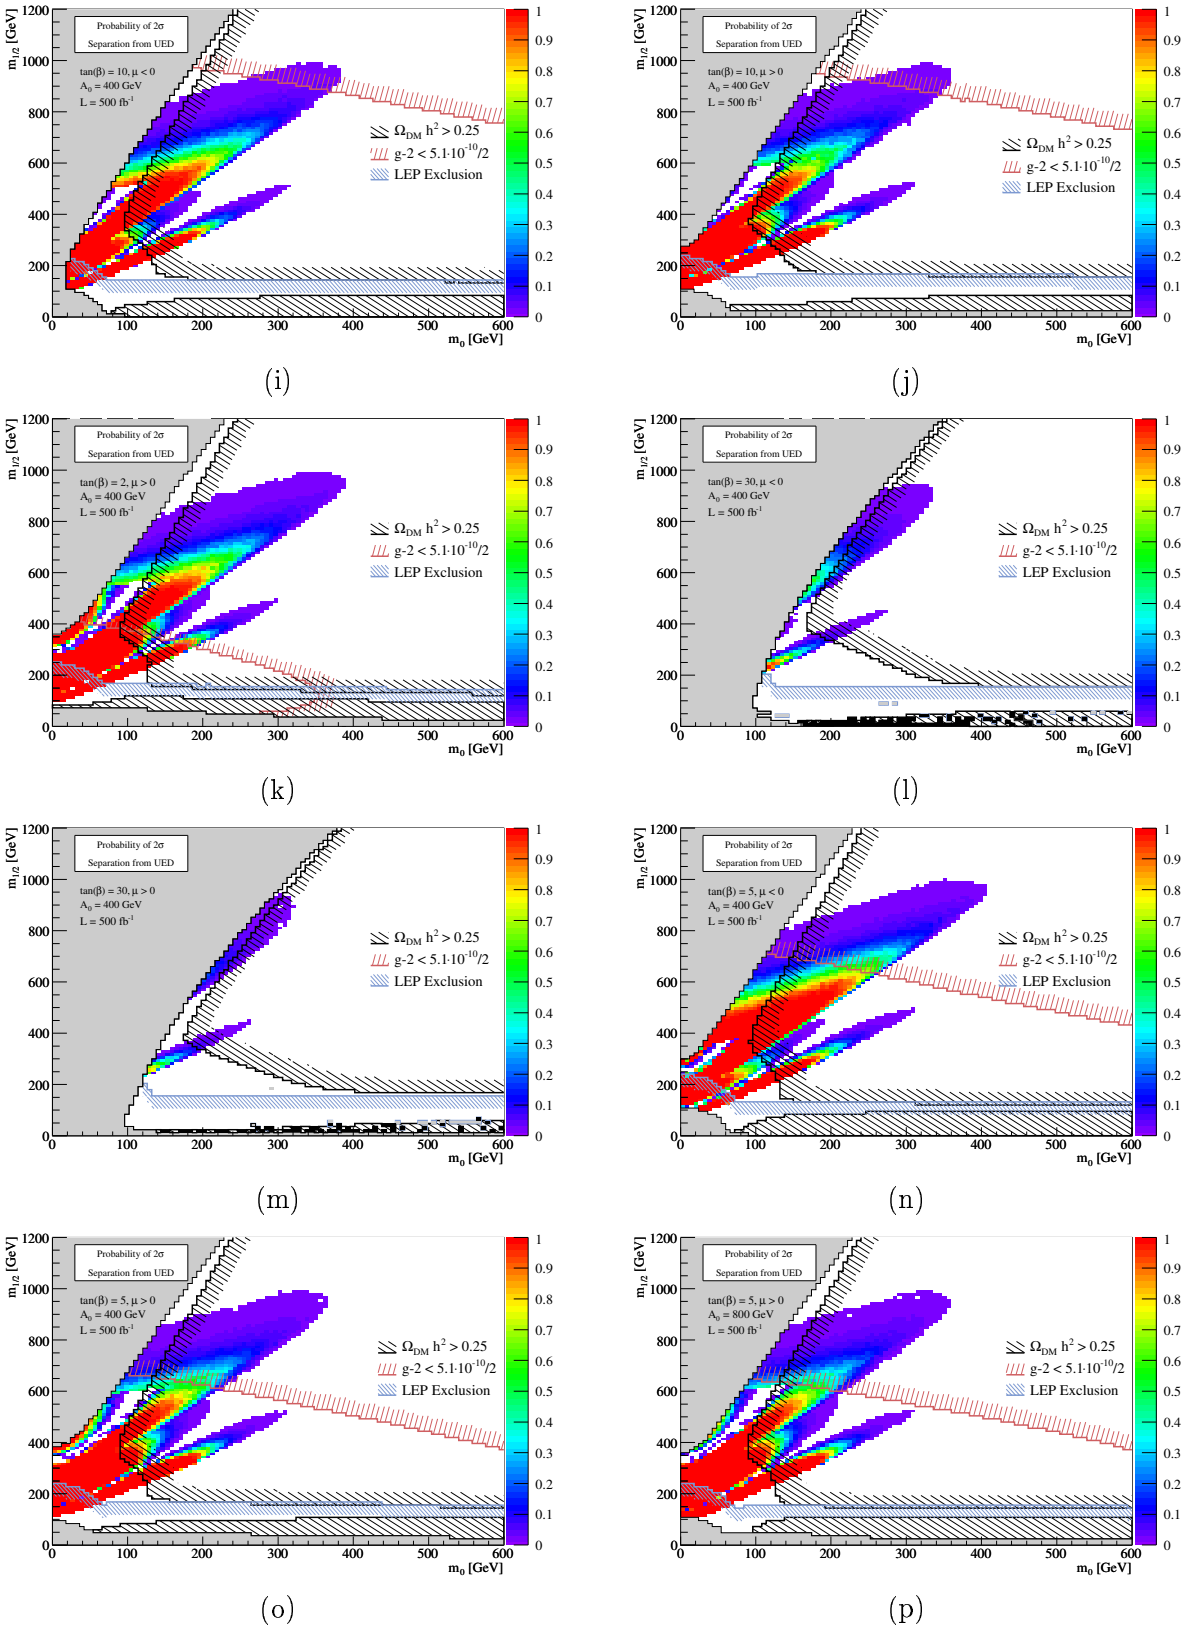

Figure A.6 [2/2]: *(continued)*

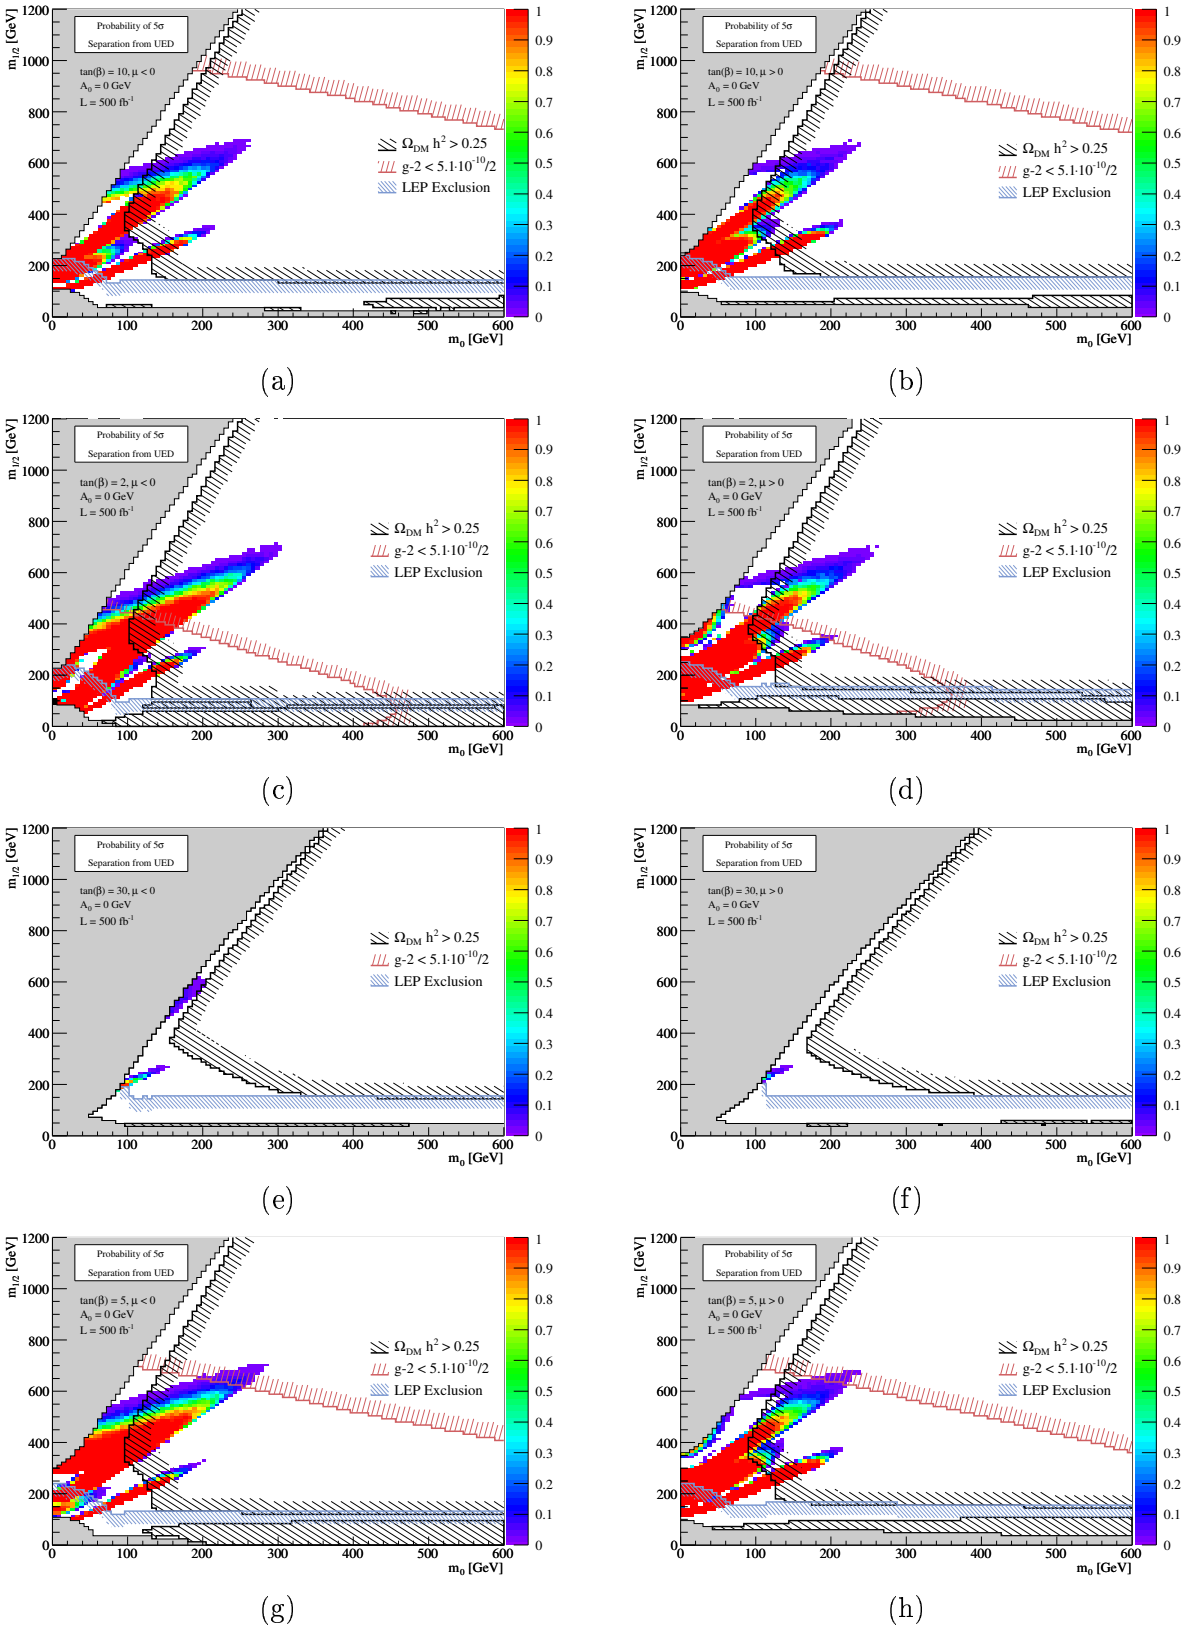

Figure A.7 [1/2]: *Chance to observe a 5σ separation between SUSY and UED.*

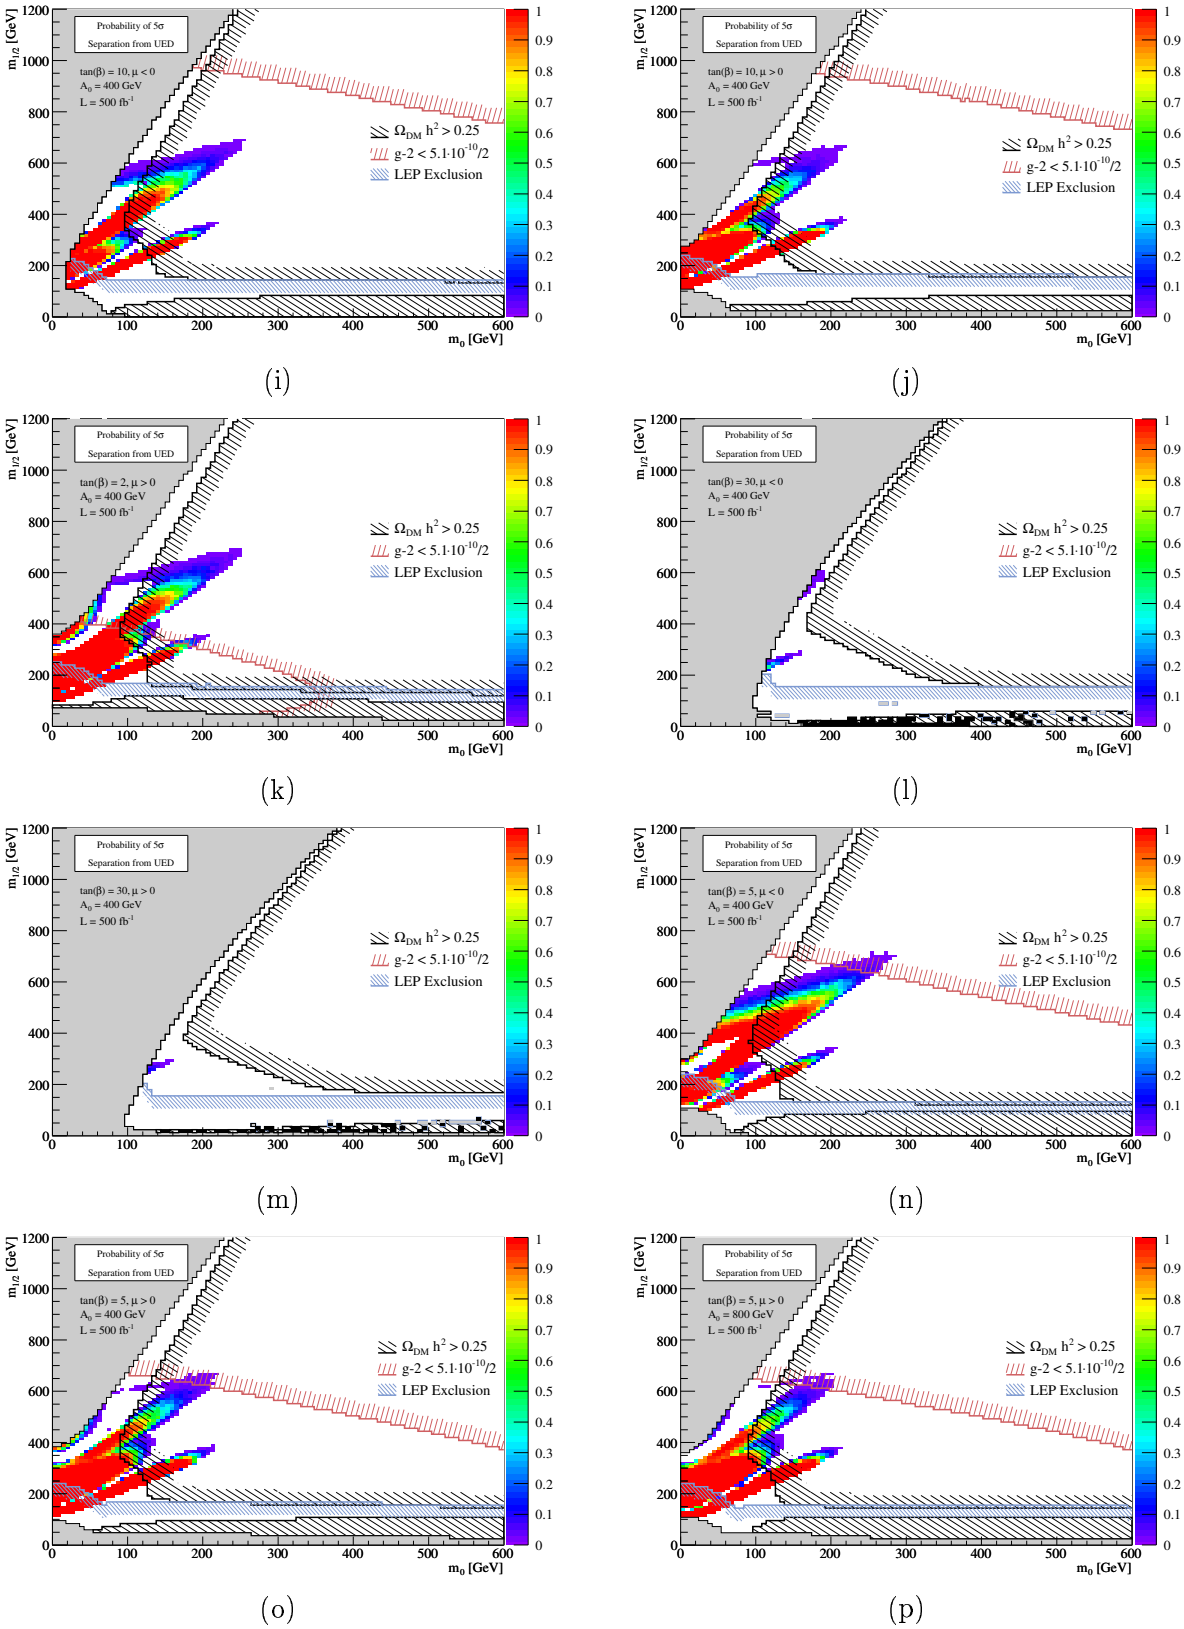

Figure A.7 [2/2]: *(continued)*

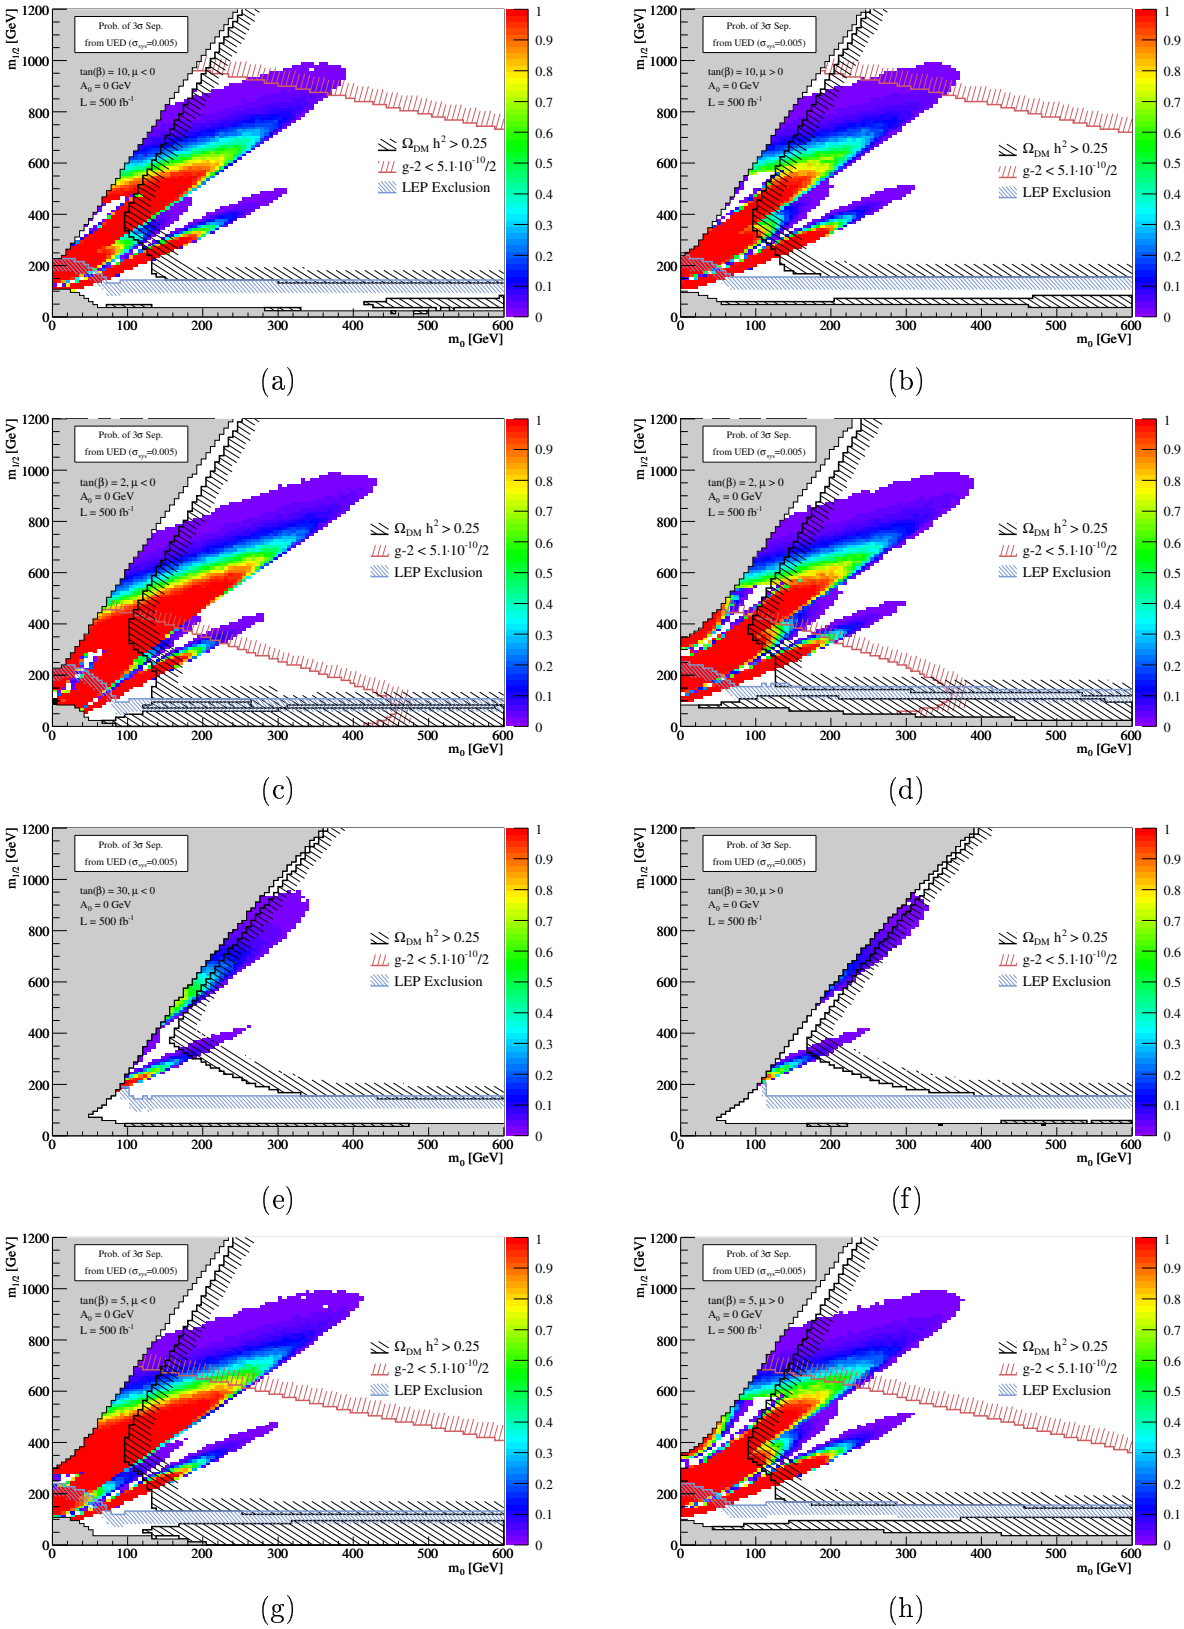

Figure A.8 [1/2]: *Chance to observe a 3σ separation between SUSY and UED assuming a systematic error of 0.005.*

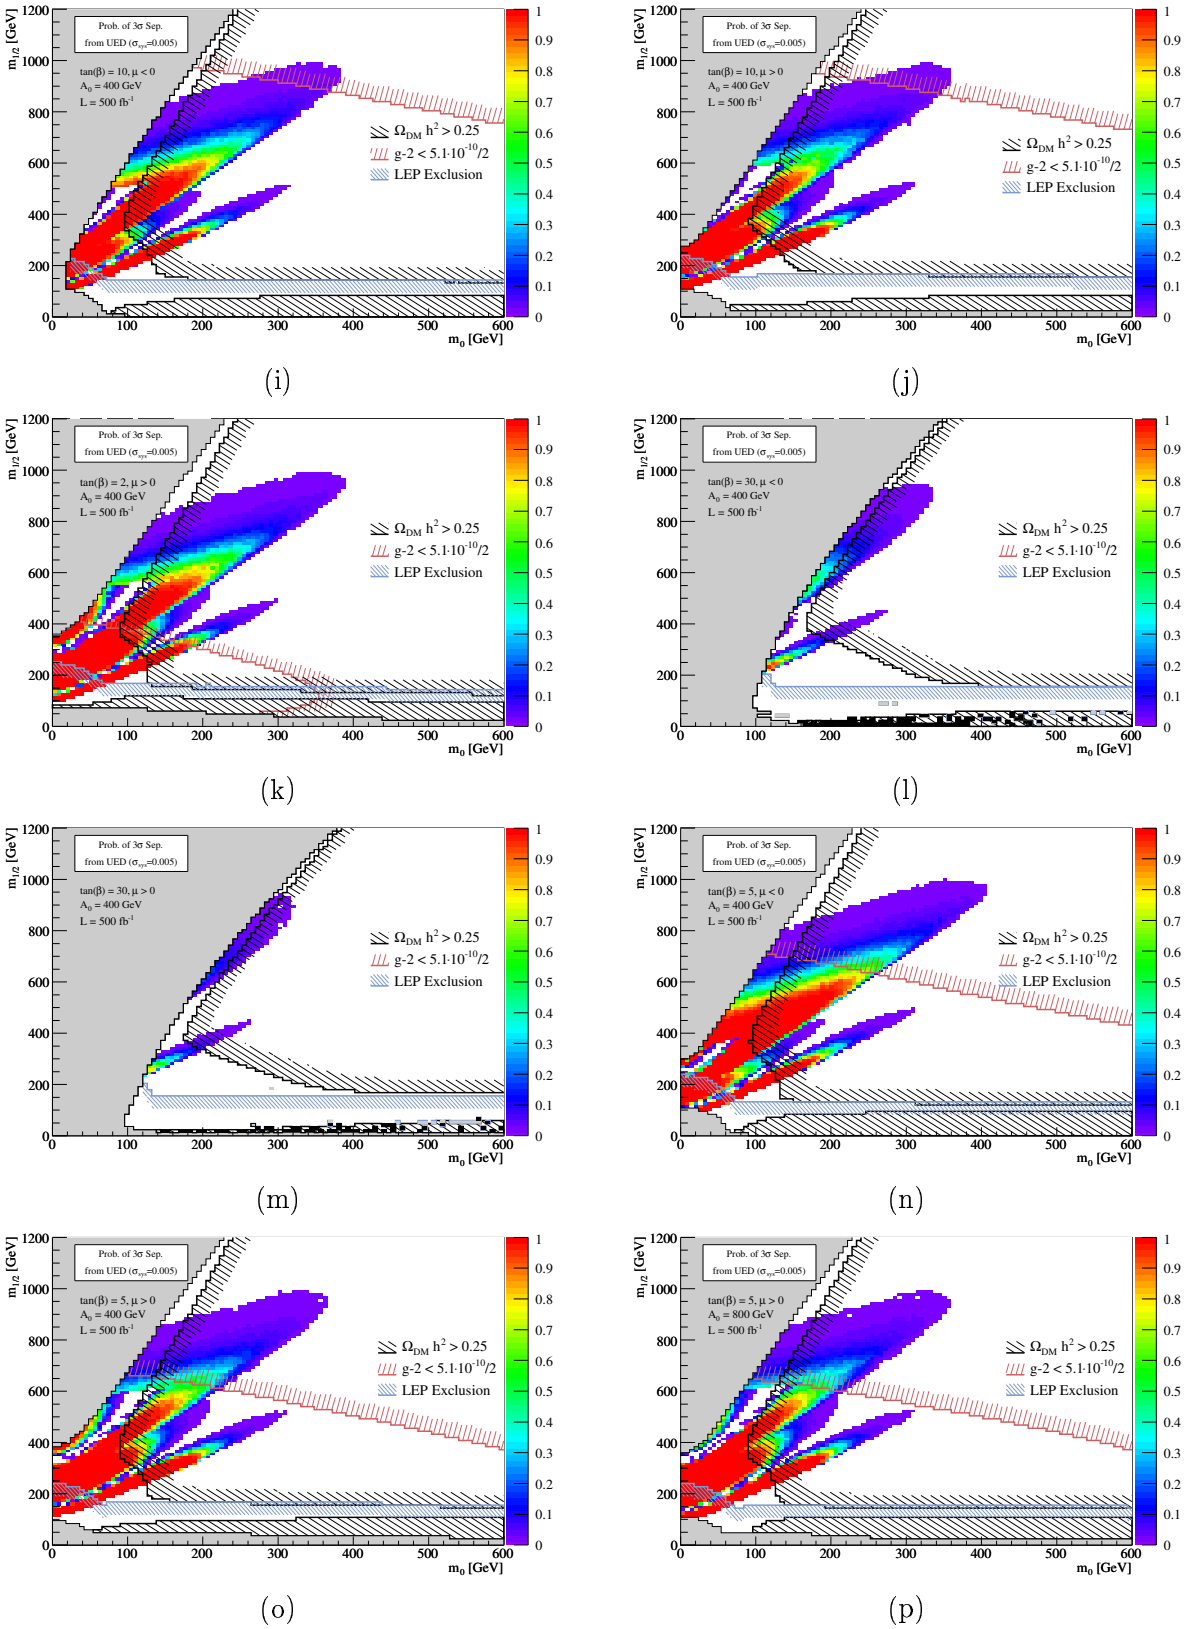

Figure A.8 [2/2]: *(continued)*

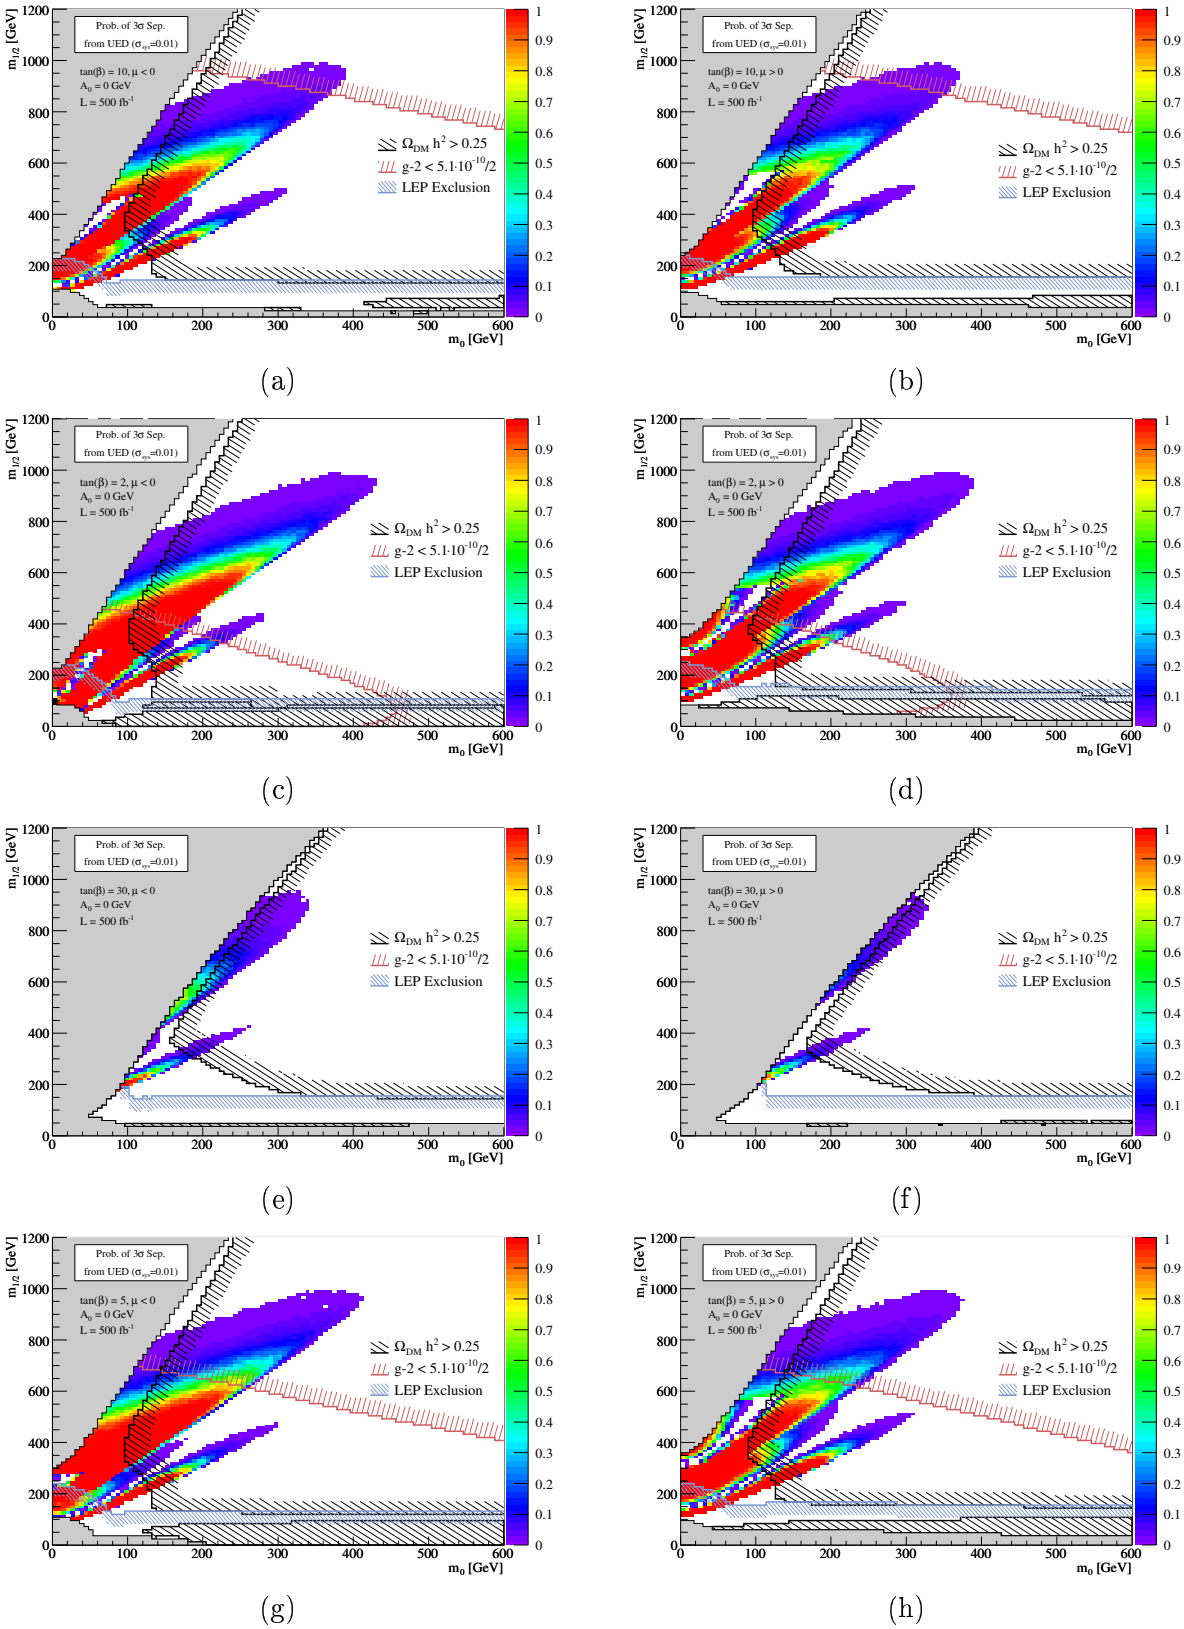

Figure A.9 [1/2]: *Chance to observe a 3σ separation between SUSY and UED assuming a systematic error of 0.01.*

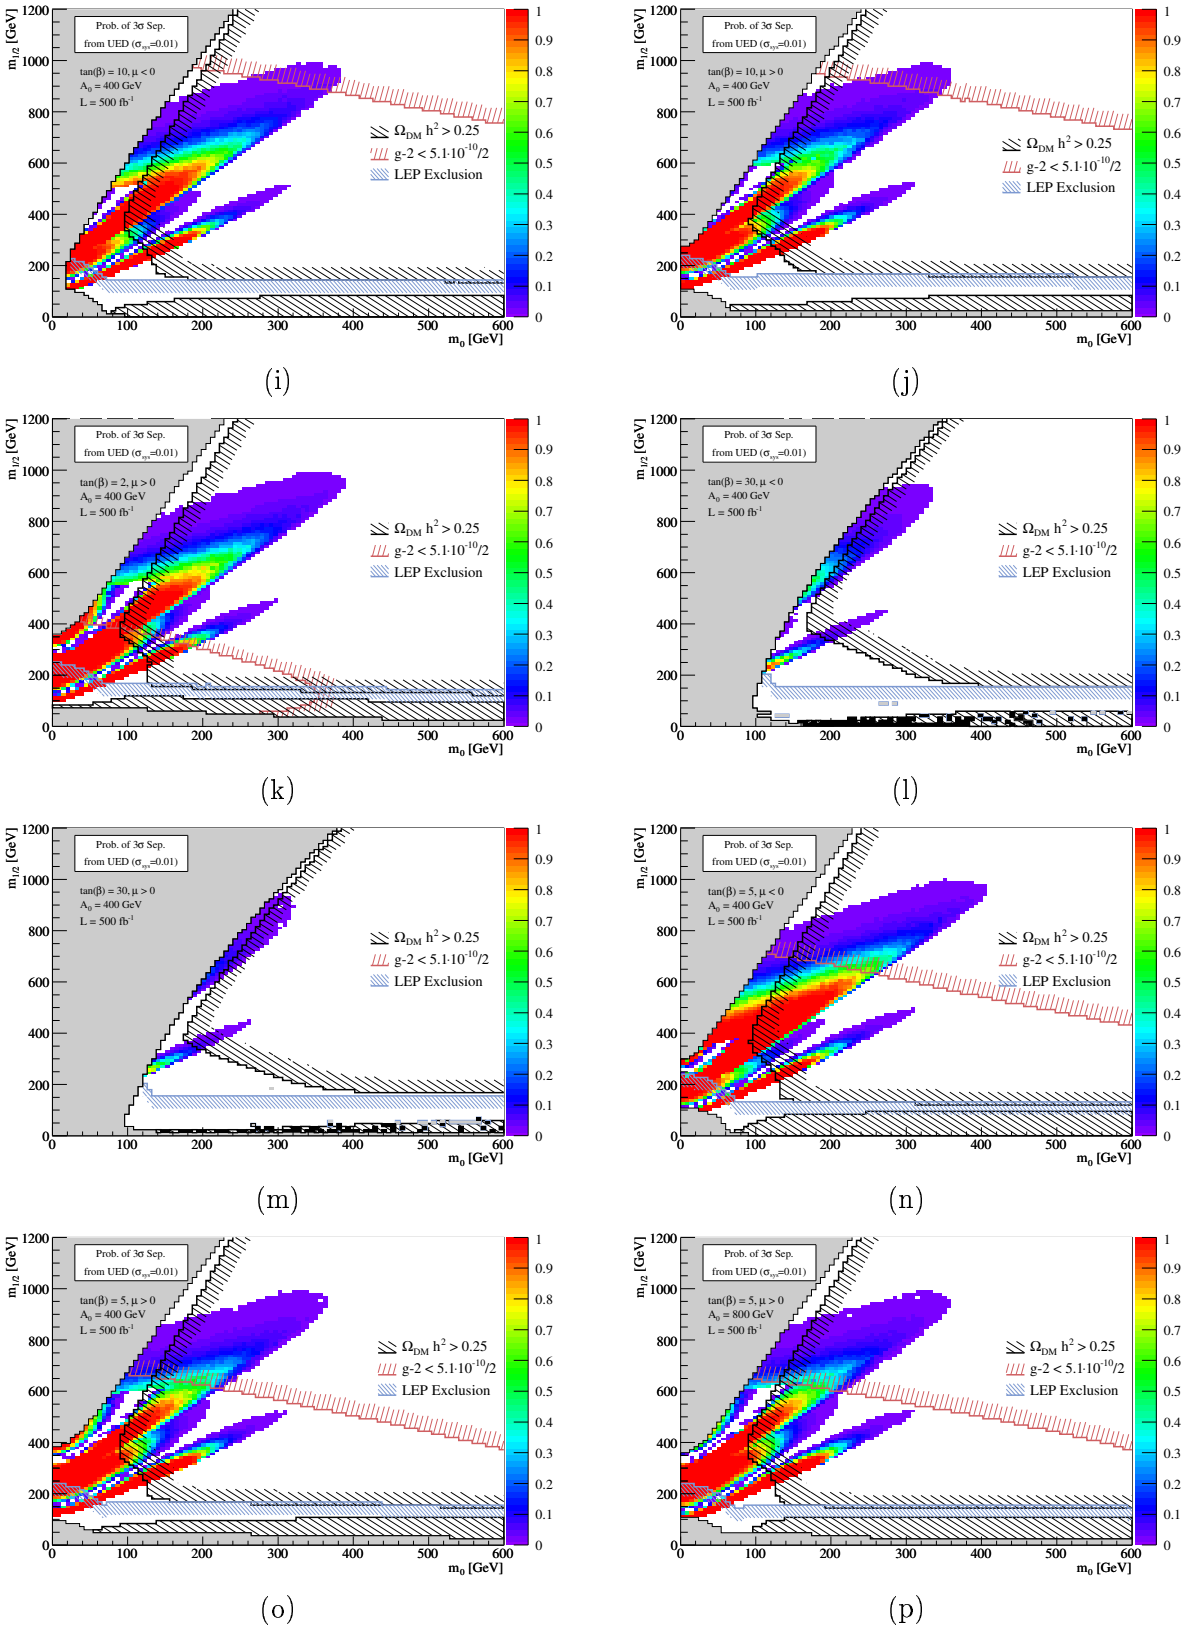

Figure A.9 [2/2]: *(continued)*

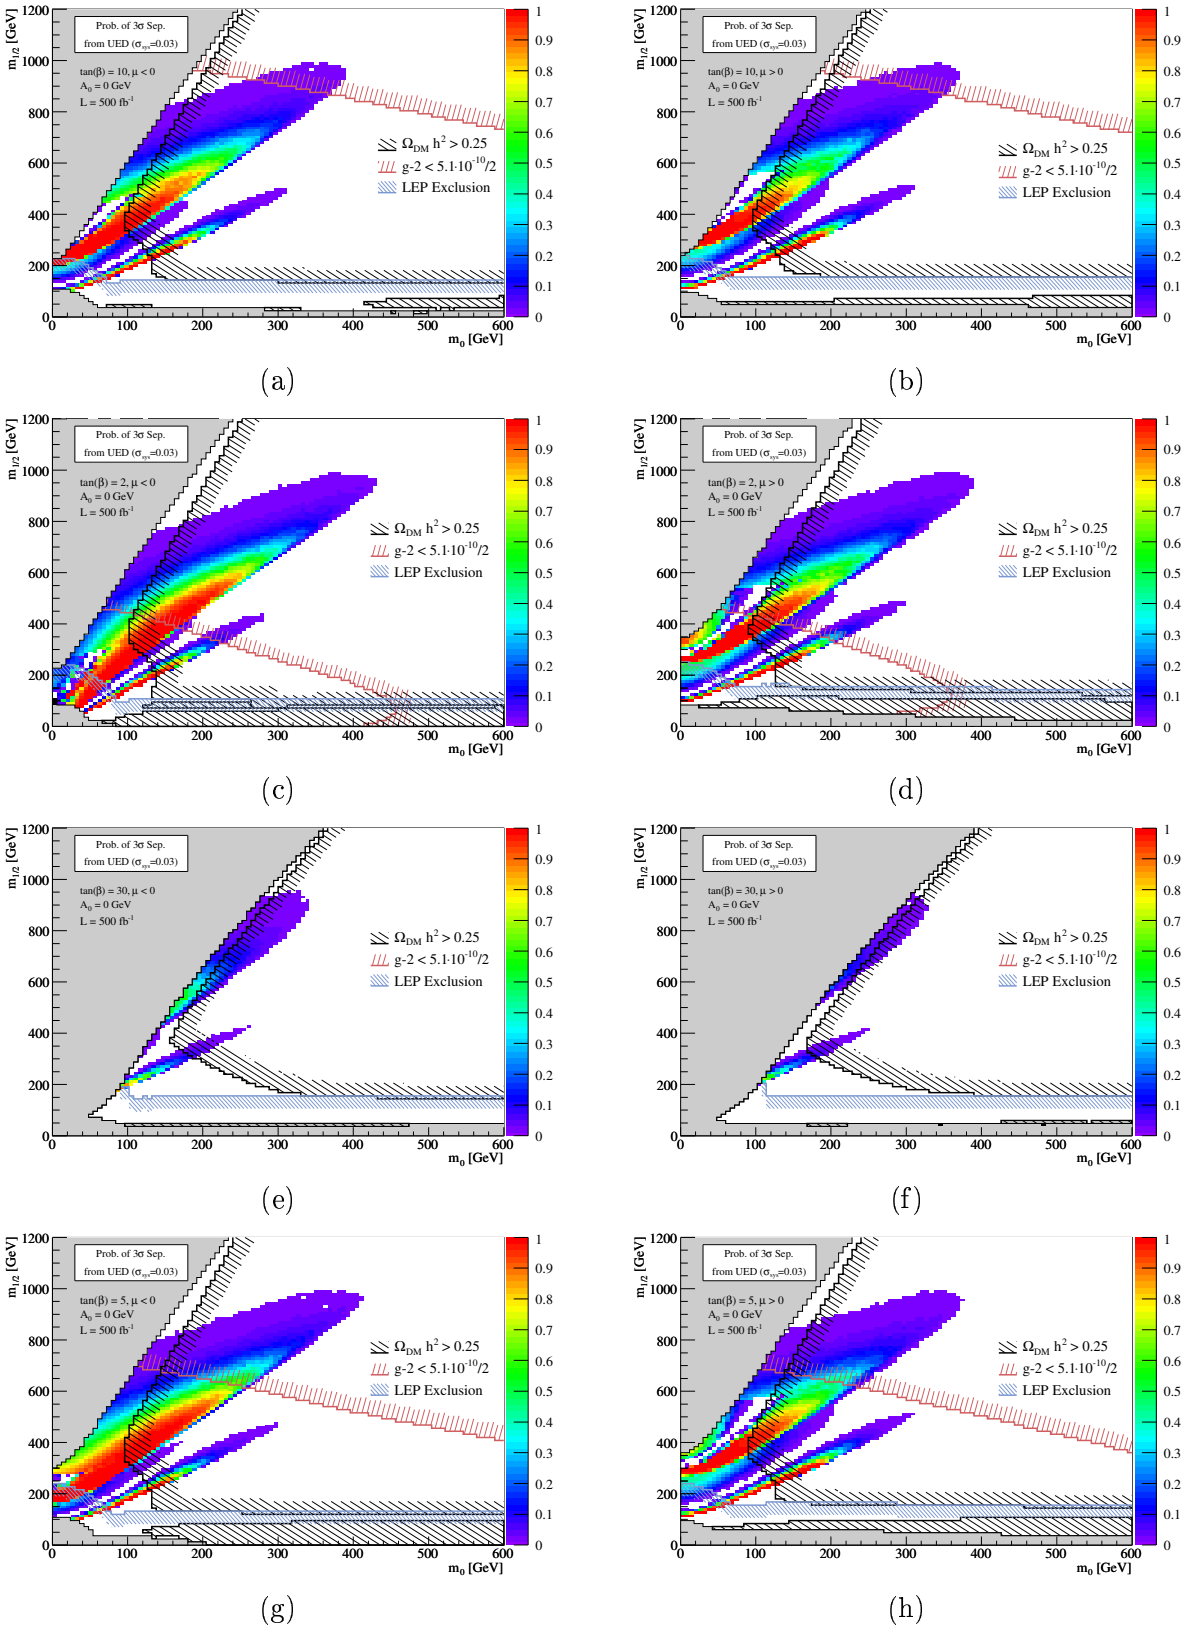

Figure A.10 [1/2]: *Chance to observe a 3σ separation between SUSY and UED assuming a systematic error of 0.03.*

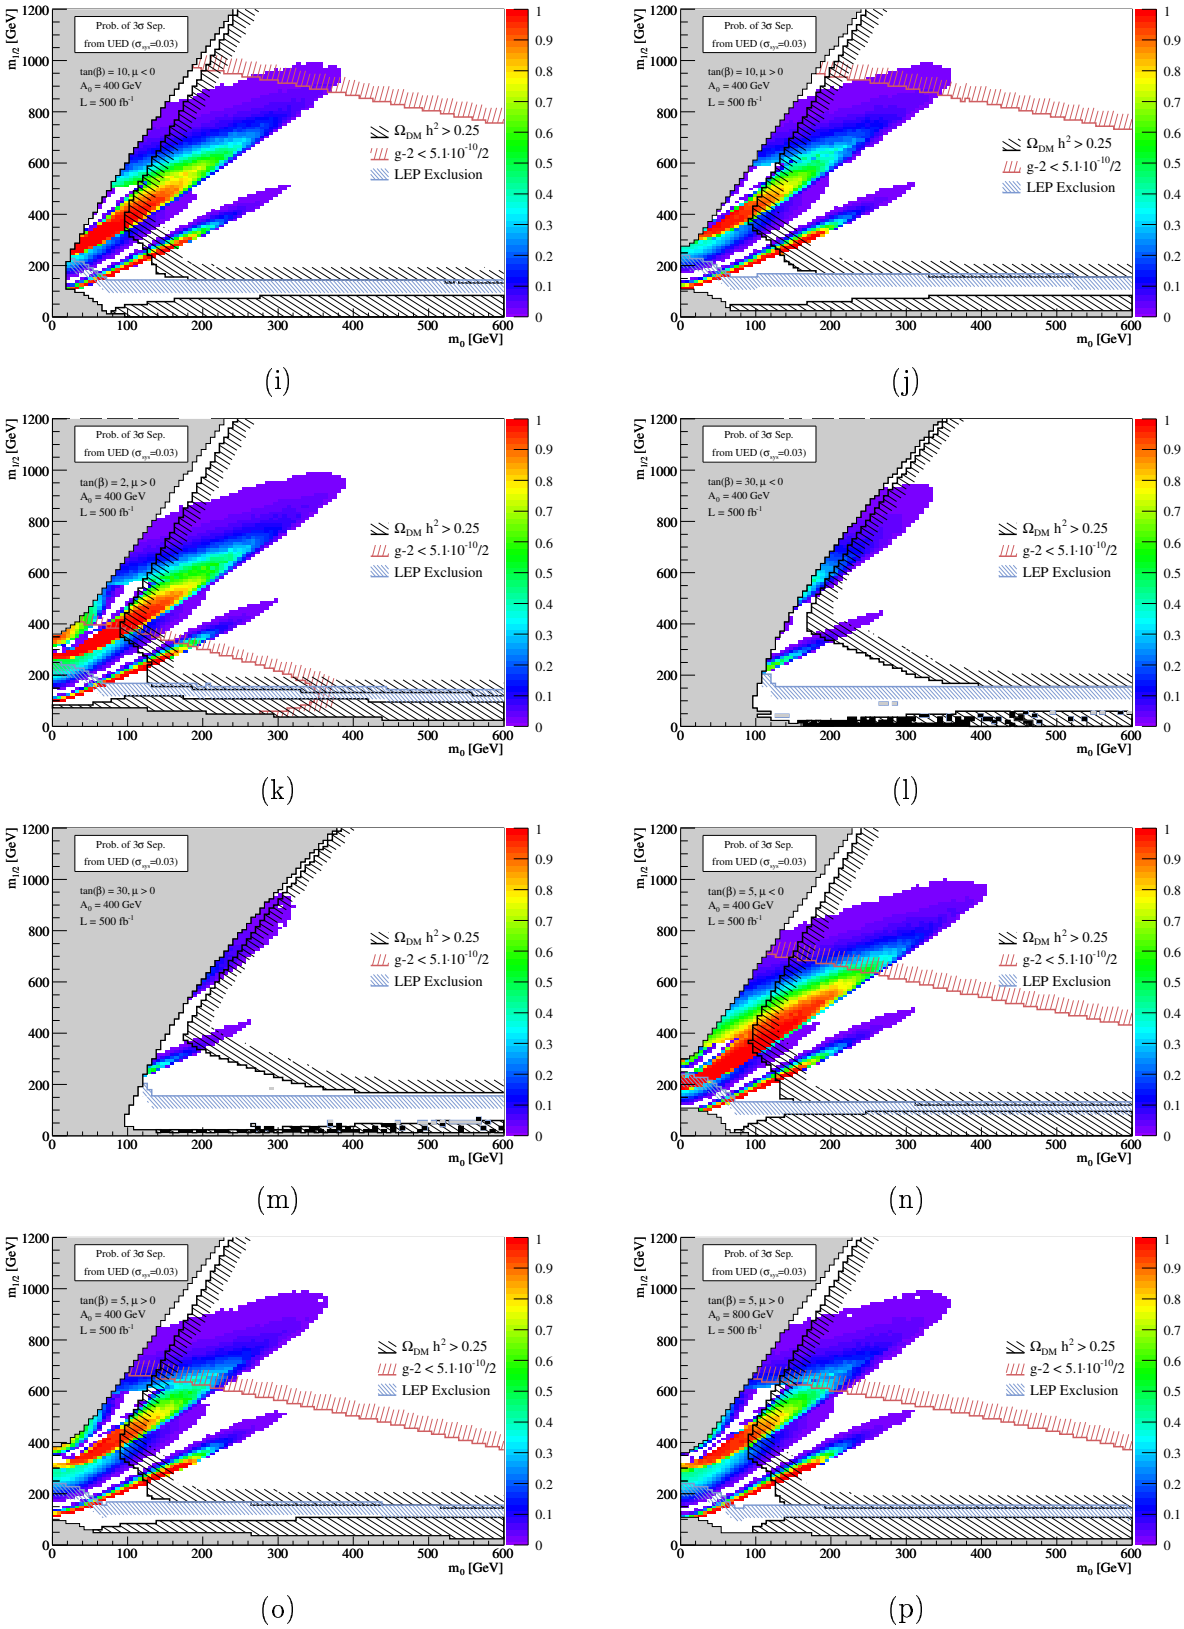

Figure A.10 [2/2]: *(continued)*

References

[A+00] N. Arkani-Hamed, H.-C. Cheng, B. A. Dobrescu and L. J. Hall, *Self-breaking of the standard model gauge symmetry*, Phys. Rev. **D62**, 096006 (2000), hep-ph/0006238.

[A+01] C. Alexander et al., *Progress in the development of the DTMROC time measurement chip for the ATLAS transition Radiation Tracker (TRT)*, IEEE Trans. Nucl. Sci. **48**, 514–519 (2001).

[AC80] W. W. M. Allison and J. H. Cobb, *RELATIVISTIC CHARGED PARTICLE IDENTIFICATION BY ENERGY LOSS*, Ann. Rev. Nucl. Part. Sci. **30**, 253–298 (1980).

[ACD01] T. Appelquist, H.-C. Cheng and B. A. Dobrescu, *Bounds on universal extra dimensions*, Phys. Rev. **D64**, 035002 (2001), hep-ph/0012100.

[ADD98] N. Arkani-Hamed, S. Dimopoulos and G. R. Dvali, *The hierarchy problem and new dimensions at a millimeter*, Phys. Lett. **B429**, 263–272 (1998), hep-ph/9803315.

[AEBC06] M. Aharrouche et al. (ATLAS Electromagnetic Barrel Calorimeter Collaboration), *Energy linearity and resolution of the ATLAS electromagnetic barrel calorimeter in an electron test- beam*, (2006), physics/0608012.

[AL00] R. Barate et al. (ALEPH Collaboration), *Observation of an excess in the search for the standard model Higgs boson at ALEPH*, Phys. Lett. **B495**, 1–17 (2000), hep-ex/0011045.

[ATL94] W. W. Armstrong et al. (ATLAS Collaboration), *ATLAS: Technical proposal for a general-purpose p p experiment at the Large Hadron Collider at CERN*, (Dec 1994), CERN-LHCC-94-43.

[ATL96a] ATLAS calorimeter performance, Dec 1996, CERN-LHCC-96-40.

[ATL96b] ATLAS liquid argon calorimeter: Technical design report, Dec 1996, CERN-LHCC-96-41.

[ATL97a] ATLAS Inner Detector: Technical Design Report. Vol. 1, Apr 1997, CERN-LHCC-97-16.

[ATL97b] ATLAS Inner Detector: Technical Design Report. Vol. 2, Apr 1997, CERN-LHCC-97-17.

[ATL99a] ATLAS Detector and Physics Performance: Technical Design Report. Vol. 1, May 1999, CERN-LHCC-99-14.

[ATL99b] ATLAS Detector and Physics Performance: Technical Design Report. Vol. 2, May 1999, CERN-LHCC-99-15.

[ATL03] ATLAS High-Level Trigger, Data Acquisition and Controls: Technical Design Report, Jul 2003, CERN-LHCC-2003-022.

[ATL05] A. F. Saavedra (ATLAS Pixel Collaboration), *Status of the ATLAS pixel detector*, Nucl. Instrum. Meth. **A541**, 130–136 (2005).

[AYM75] X. Artru, G. B. Yodh and G. Mennessier, *PRACTICAL THEORY OF THE MULTILAYERED TRANSITION RADIATION DETECTOR*, Phys. Rev. **D12**, 1289 (1975).

[B+96] B. Bevensee, F. M. Newcomer, R. Van Berg and H. H. Williams, *An amplifier shaper discriminator with baseline restoration for the ATLAS transition radiation tracker*, IEEE Trans. Nucl. Sci. **43**, 1725–1731 (1996).

[Bal79] V. Baluni, *CP VIOLATING EFFECTS IN QCD*, Phys. Rev. **D19**, 2227–2230 (1979).

[Bar04] A. J. Barr, *Using lepton charge asymmetry to investigate the spin of supersymmetric particles at the LHC*, Phys. Lett. **B596**, 205–212 (2004), hep-ph/0405052.

[Bar06] A. J. Barr, *Measuring slepton spin at the LHC*, JHEP **02**, 042 (2006), hep-ph/0511115.

[BBPS02] 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), hep-ph/0112278.

[BBPS06] G. Belanger, F. Boudjema, A. Pukhov and A. Semenov, *MicrOMEGAs: Version 1.3*, Comput. Phys. Commun. **174**, 577–604 (2006), hep-ph/0405253.

[BBR02] B. Aubert et al. (BABAR Collaboration), *The BaBar detector*, Nucl. Instrum. Meth. **A479**, 1–116 (2002), hep-ex/0105044.

[BEL98] A. Bondar (BELLE Collaboration), *The BELLE detector*, Nucl. Instrum. Meth. **A408**, 64–76 (1998).

[BFKP06] H. Baer, J. Ferrandis, S. Kraml and W. Porod, *On the treatment of threshold effects in SUSY spectrum computations*, Phys. Rev. **D73**, 015010 (2006), hep-ph/0511123.

[BFS96] J. M. Butterworth, J. R. Forshaw and M. H. Seymour, *Multiparton interactions in photoproduction at HERA*, Z. Phys. **C72**, 637–646 (1996), hep-ph/9601371.

[Bia99] S. F. Biagi, *Monte Carlo simulation of electron drift and diffusion in counting gases under the influence of electric and magnetic fields*, Nucl. Instrum. Meth. **A421**, 234–240 (1999).

[BM06] M. R. Buckley and H. Murayama, *How can we test seesaw experimentally?*, (2006), hep-ph/0606088.

[Bou03] J. Boudreau, Geometry Kernel Classes, `http://atlas.web.cern.ch/Atlas/GROUPS/DATABASE/detector_description/`, 2003.

[BT04] J. Boudreau and V. Tsulaia, THE GEOMODEL TOOLKIT FOR DETECTOR DESCRIPTION, in *Proceedings of the Conference on Computing in High Energy and Nuclear Physics (CHEP 2004), Interlaken, Switzerland*, Sep 2004.

[C+79] R. J. Crewther, P. Di Vecchia, G. Veneziano and E. Witten, *CHIRAL ESTIMATE OF THE ELECTRIC DIPOLE MOMENT OF THE NEUTRON IN QUANTUM CHROMODYNAMICS*, Phys. Lett. **B88**, 123 (1979).

[C+01] G. Corcella et al., *HERWIG 6: An event generator for hadron emission reactions with interfering gluons (including supersymmetric processes)*, JHEP **01**, 010 (2001), hep-ph/0011363.

[C+02] G. Corcella et al., *HERWIG 6.5 release note*, (2002), hep-ph/0210213.

[Cab63] N. Cabibbo, *UNITARY SYMMETRY AND LEPTONIC DECAYS*, Phys. Rev. Lett. **10**, 531–532 (1963).

[CAST05] K. Zioutas et al. (CAST Collaboration), *First results from the CERN axion solar telescope (CAST)*, Phys. Rev. Lett. **94**, 121301 (2005), hep-ex/0411033.

[CCFT64] J. H. Christenson, J. W. Cronin, V. L. Fitch and R. Turlay, *EVIDENCE FOR THE 2 PI DECAY OF THE K(2)0 MESON*, Phys. Rev. Lett. **13**, 138–140 (1964).

[CKM05] J. Charles et al. (CKMfitter Group Collaboration), *CP violation and the CKM matrix: Assessing the impact of the asymmetric B factories*, Eur. Phys. J. **C41**, 1–131 (2005), hep-ph/0406184, Figures used in the present thesis are updated versions from summer 2006 from `http://ckmfitter.in2p3.fr/`.

[CM67] S. R. Coleman and J. Mandula, *ALL POSSIBLE SYMMETRIES OF THE S MATRIX*, Phys. Rev. **159**, 1251–1256 (1967).

[Cow98] G. Cowan, *Statistical Data Analysis*, Oxford University Press, 1998, 197p.

[CST03] S. Chouridou, R. Ströhmer and T. Trefzger, Study of the three body matrix element of the neutralino decay $\tilde{\chi}_2^0 \to \tilde{\chi}_1^0 \ell^+ \ell^-$, ATLAS-PHYS-2003-034, Nov 2003.

[Cwe06] P. Cwetanski, *Straw performance studies and quality assurance for the ATLAS transition radiation tracker*, (2006), CERN-THESIS-2006-025.

[D+74] S. Deser, H.-S. Tsao and P. van Nieuwenhuizen, *NONRENORMALIZABILITY OF EINSTEIN YANG-MILLS INTERACTIONS AT THE ONE LOOP LEVEL*, Phys. Lett. **B50**, 491 (1974).

[Dam04] M. Dam, Simulation of ATLAS Transition Radiation, Talk given at TRT Collaboration Meeting, CERN, February 16-23, 2004.

[Daw06] S. Dawson, Electroweak Theory, Lecture series given at the 1st Hadron Collider Physics Summer School, Fermilab, August 9-18, 2006.

[Djo05] A. Djouadi, *The anatomy of electro-weak symmetry breaking. II: The Higgs bosons in the minimal supersymmetric model*, (2005), hep-ph/0503173.

[DKM05] A. Datta, K. Kong and K. T. Matchev, *Discrimination of supersymmetry and universal extra dimensions at hadron colliders*, Phys. Rev. **D72**, 096006 (2005), hep-ph/0509246.

[DKS90] M. Dine, A. Kagan and S. Samuel, *NATURALNESS IN SUPERSYMMETRY, OR RAISING THE SUPERSYMMETRY BREAKING SCALE*, Phys. Lett. **B243**, 250–256 (1990).

[Dri04] C. Driouichi, *B-physics potential of the ATLAS experiment and performance of the ATLAS Transition Radiation Tracker*, PhD thesis, Lund University, 2004.

[EDJ+71] W. T. Eadie, D. Dryard, F. E. James, M. Roos and B. Sadoulet, *Statistical Methods in Experimental Physics*, North-Holland, 1971.

[Ege98] U. Egede, *The search for a Standard Model Higgs at the LHC and electron identification using transition radiation in the ATLAS tracker*, (1998), CERN-THESIS-98-001.

[EN18] M. A. T. Emmy Noether, *Invariant Variation Problems*, (1918), physics/0503066, M. A. Tavel's English translation of Noether's Theorems (1918), reproduced by Frank Y. Wang (2005).

[ESW96] R. K. Ellis, W. J. Stirling and B. R. Webber, *QCD and collider physics*, Camb. Monogr. Part. Phys. Nucl. Phys. Cosmol. **8**, 1–435 (1996).

[Fri91] J. I. Friedman, *Deep inelastic scattering: Comparisons with the quark model*, Rev. Mod. Phys. **63**, 615–629 (1991).

[FS81] E. Farhi and L. Susskind, *TECHNICOLOR*, Phys. Rept. **74**, 277 (1981).

[G403] S. Agostinelli et al. (GEANT4 Collaboration), *GEANT4: A simulation toolkit*, Nucl. Instrum. Meth. **A506**, 250–303 (2003).

[GAR73] F. J. Hasert et al. (Gargamelle Neutrino Collaboration), *Observation of neutrino-like interactions without muon or electron in the Gargamelle neutrino experiment*, Phys. Lett. **B46**, 138–140 (1973).

[GG74] H. Georgi and S. L. Glashow, *UNITY OF ALL ELEMENTARY PARTICLE FORCES*, Phys. Rev. Lett. **32**, 438–441 (1974).

[GKL91] C. Giunti, C. W. Kim and U. W. Lee, *Running coupling constants and grand unification models*, Mod. Phys. Lett. **A6**, 1745–1755 (1991).

[GM204] G. W. Bennett et al. (Muon g-2 Collaboration), *Measurement of the negative muon anomalous magnetic moment to 0.7-ppm*, Phys. Rev. Lett. **92**, 161802 (2004), hep-ex/0401008.

[Gol80] H. Goldstein, *Classical Mechanics*, Addison-Wesley, 2nd edition, 1980.

[GSW62] J. Goldstone, A. Salam and S. Weinberg, *BROKEN SYMMETRIES*, Phys. Rev. **127**, 965–970 (1962).

[Gun05] J. H. Gundlach, *Laboratory tests of gravity*, New J. Phys. **7**, 205 (2005).

[H+04] C. D. Hoyle et al., *Sub-millimeter tests of the gravitational inverse-square law*, Phys. Rev. **D70**, 042004 (2004), hep-ph/0405262.

[HI93] D. J. Bird et al. (HIRES Collaboration), *Evidence for correlated changes in the spectrum and composition of cosmic rays at extremely high-energies*, Phys. Rev. Lett. **71**, 3401–3404 (1993).

[HLS75] R. Haag, J. T. Lopuszanski and M. Sohnius, *ALL POSSIBLE GENERATORS OF SUPERSYMMETRIES OF THE S MATRIX*, Nucl. Phys. **B88**, 257 (1975).

[HM84] F. Halzen and A. D. Martin, *QUARKS AND LEPTONS: An Introductory Course in Modern Particle Physics*, Wiley, 1984.

[HQ98] P. F. Harrison and H. R. Quinn (BABAR Collaboration), editors, *The BaBar Physics Book: Physics at an Asymmetric B Factory*, SLAC Report 504, Oct 1998.

[Jam94] F. James, *MINUIT: Function Minimization and Error Analysis, Reference Manual, Version 94.1*, 1994, CERN Program Library Long Writeup D506.

[Kan04] P. Kanti, *Black holes in theories with large extra dimensions: A review*, Int. J. Mod. Phys. **A19**, 4899–4951 (2004), hep-ph/0402168.

[KAT04] J. Angrik et al. (KATRIN Collaboration), *KATRIN design report 2004*, FZKA-7090.

[Kit03] T. H. Kittelmann, Putting a Limit on the Width Difference of the Neutral B-Mesons with the BaBar Experiment, Master's thesis, University of Copenhagen, Aug 2003, Available on `http://www.nbi.dk/~kittel/deltagamma/` and `http://hep.nbi.dk/`.

[Kit06] T. H. Kittelmann, Modelling of Noise and Straw to Straw Variations in the ATLAS Transition Radiation Tracker, ATL-COM-INDET-2006-016, October 2006.

[Kje04] M. Kjerulff, Discovery Potential of Large Extra Dimensions in the Di-photon Channel with the ATLAS Experiment, Master's thesis, University of Copenhagen, Apr 2004, Available on `http://hep.nbi.dk/`.

[KK06] T. H. Kittelmann and E. B. Klinkby, A Study of TRT Noise in 2004 Test Beam Data, ATL-COM-INDET-2006-015, October 2006.

[KKRW94] G. L. Kane, C. F. Kolda, L. Roszkowski and J. D. Wells, *Study of constrained minimal supersymmetry*, Phys. Rev. **D49**, 6173–6210 (1994), hep-ph/9312272.

[KM73] M. Kobayashi and T. Maskawa, *CP violation in the renormalizable theory of weak interaction*, Prog. Theor. Phys. **49**, 652–657 (1973).

[KMSS06] M. Kakizaki, S. Matsumoto, Y. Sato and M. Senami, *Relic abundance of LKP dark matter in UED model including effects of second KK resonances*, Nucl. Phys. **B735**, 84–95 (2006), hep-ph/0508283.

[KN06] T. Kinoshita and M. Nio, *Improved alpha**4 term of the electron anomalous magnetic moment*, Phys. Rev. **D73**, 013003 (2006), hep-ph/0507249.

[Kra04] A. C. Kraan, *Interactions of heavy stable hadronizing particles*, Eur. Phys. J. **C37**, 91–104 (2004), hep-ex/0404001.

[L+95] H. L. Lai et al., *Global QCD analysis and the CTEQ parton distributions*, Phys. Rev. **D51**, 4763–4782 (1995), hep-ph/9410404.

[Leo93] W. R. Leo, *TECHNIQUES FOR NUCLEAR AND PARTICLE PHYSICS EXPERIMENTS: A HOW TO APPROACH*, Springer, 2nd edition, 1993, 378p.

[LEP01] L. Pape et al., Combined LEP Chargino Results up to 208 GeV for large m0, LEPSUSYWG, ALEPH, DELPHI, L3 and OPAL experiments. Note LEPSUSYWG/01-03.1, Mar 2001, `http://lepsusy.web.cern.ch/lepsusy/`.

[LEP03] R. Barate et al. (LEP Working Group for Higgs boson searches Collaboration), *Search for the standard model Higgs boson at LEP*, Phys. Lett. **B565**, 61–75 (2003), hep-ex/0306033.

[LEP04] L. Pape et al., Combined LEP Selectron/Smuon/Stau Results 183-208 GeV, LEPSUSYWG, ALEPH, DELPHI, L3 and OPAL experiments. Note LEPSUSYWG/04-01.1, Jan 2004, `http://lepsusy.web.cern.ch/lepsusy/`.

[LEP05] The ALEPH, DELPHI, L3, OPAL, SLD Collaborations, the LEP Electroweak Working Group, the SLD Electroweak and Heavy Flavour Groups, *Precision Electroweak Measurements on the Z Resonance*, page 302 (2005), hep-ex/0509008.

[Lob03] V. M. Lobashev, *The search for the neutrino mass by direct method in the tritium beta-decay and perspectives of study it in the project KATRIN*, Nucl. Phys. **A719**, 153–160 (2003).

[LP99] K. Lassila-Perini, *Discovery potential of the standard model Higgs in CMS at the LHC*, (1999), CERN-THESIS-99-003.

[LR04] F. Luehring and A. Romaniouk, ATLAS TRT Barrel in Test Beam, Atlas eNews, Sep 2004, `http://aenews.cern.ch/aenews.php?issueno=200409`.

[Lüd57] G. Lüders, *Proof of the TCP theorem*, Annals Phys. **2**, 1–15 (1957).

[LY56] T. D. Lee and C.-N. Yang, *QUESTION OF PARITY CONSERVATION IN WEAK INTERACTIONS*, Phys. Rev. **104**, 254–258 (1956).

[M+92] G. Marchesini et al., *HERWIG: A Monte Carlo event generator for simulating hadron emission reactions with interfering gluons. Version 5.1 - April 1991*, Comput. Phys. Commun. **67**, 465–508 (1992).

[Mar97] S. P. Martin, *A supersymmetry primer*, (1997), hep-ph/9709356.

[MNS62] Z. Maki, M. Nakagawa and S. Sakata, *Remarks on the unified model of elementary particles*, Prog. Theor. Phys. **28**, 870 (1962).

[MOR+02] S. Moretti, K. Odagiri, P. Richardson, M. H. Seymour and B. R. Webber, *Implementation of supersymmetric processes in the HERWIG event generator*, JHEP **04**, 028 (2002), hep-ph/0204123.

[MP06] S. Mehlhase and T. C. Petersen, A probability based approach to PID in the TRT detector of ATLAS, ATL-COM-INDET-2006-017, October 2006.

[MSTV04] M. Maltoni, T. Schwetz, M. A. Tortola and J. W. F. Valle, *Status of global fits to neutrino oscillations*, New J. Phys. **6**, 122 (2004), hep-ph/0405172.

[MT05] P. J. Mohr and B. N. Taylor, *CODATA recommended values of the fundamental physical constants: 2002*, Rev. Mod. Phys. **77**, 1–107 (2005).

[NEMO05] R. Arnold et al. (NEMO Collaboration), *First results of the search of neutrinoless double beta decay with the NEMO 3 detector*, Phys. Rev. Lett. **95**, 182302 (2005), hep-ex/0507083.

[P+92] W. H. Press et al., *Numerical Recipes in C*, Cambridge University Press, 2nd edition, 1992, 994p.

[P+02] J. Pumplin et al., *New generation of parton distributions with uncertainties from global QCD analysis*, JHEP **07**, 012 (2002), hep-ph/0201195.

[P+06] S. Pirro et al., *Further developments in the CUORICINO experiment*, Nucl. Instrum. Meth. **A559**, 352–354 (2006).

[Pad03] T. Padmanabhan, *Cosmological constant: The weight of the vacuum*, Phys. Rept. **380**, 235–320 (2003), hep-th/0212290.

[PDG06] W.-M. Yao et al. (Particle Data Group Collaboration), *Review of Particle Physics*, Journal of Physics G **33**, 1+ (2006).

[Pec98] R. D. Peccei, *Reflections on the strong CP problem*, (1998), hep-ph/9807514.

[Pet94] J. L. Petersen, Elementarpartikelfysik (3rd revised edition), Lecture Notes, Spring 1994.

[Pia06] M. Piai, *Walking in the third millennium*, (2006), hep-ph/0609104.

[PPBT03] F. E. Paige, S. D. Protopopescu, H. Baer and X. Tata, *ISAJET 7.69: A Monte Carlo event generator for p p, anti-p p, and e+ e- reactions*, (2003), hep-ph/0312045.

[PQ77] R. D. Peccei and H. R. Quinn, *CP CONSERVATION IN THE PRESENCE OF PSEUDOPARTICLES*, Phys. Rev. Lett. **38**, 1440–1443 (1977).

[PS73] J. C. Pati and A. Salam, *Unified lepton - hadron symmetry and a gauge theory of the basic interactions*, Phys. Rev. **D8**, 1240–1251 (1973).

[PS95] M. E. Peskin and D. V. Schroeder, *An Introduction to Quantum Field Theory*, Perseus Books, 1995.

[RS99a] L. Randall and R. Sundrum, *An alternative to compactification*, Phys. Rev. Lett. **83**, 4690–4693 (1999), hep-th/9906064.

[RS99b] L. Randall and R. Sundrum, *A large mass hierarchy from a small extra dimension*, Phys. Rev. Lett. **83**, 3370–3373 (1999), hep-ph/9905221.

[RS99c] L. Randall and R. Sundrum, *Out of this world supersymmetry breaking*, Nucl. Phys. **B557**, 79–118 (1999), hep-th/9810155.

[SAB+04] S. Spagnolo, K. Assamagan, J. Boudreau, S. Baranov, V. Tsulaia, A. Nairz, I. Trigger, C. Bourdarios, G. Unal, M. Lelchuk, B. Seligman, J. Tóth, P. Strízenec, D. Costanzo, G. Gorfine, T. Kittelmann, Y. Hasegawa, D. Pomarède, A. Vaniachine and A. Zalite, The Description of the Atlas Detector, in *Proceedings of the Conference on Computing in High Energy and Nuclear Physics (CHEP 2004), Interlaken, Switzerland*, Sep 2004.

[Sak67] A. D. Sakharov, *VIOLATION OF CP INVARIANCE, C ASYMMETRY, AND BARYON ASYMMETRY OF THE UNIVERSE*, Pisma Zh. Eksp. Teor. Fiz. **5**, 32–35 (1967).

[Sal62] A. Salam, *Renormalizability of gauge theories*, Phys. Rev. **127**, 331–334 (1962).

[SDSS04] M. Tegmark et al. (SDSS Collaboration), *Cosmological parameters from SDSS and WMAP*, Phys. Rev. **D69**, 103501 (2004), astro-ph/0310723.

[Sjö06] T. Sjöstrand, Theory of Hadronic Collisions, Lecture series given at the 1st Hadron Collider Physics Summer School, Fermilab, August 9-18, 2006.

[SK98] Y. Fukuda et al. (Super-Kamiokande Collaboration), *Evidence for oscillation of atmospheric neutrinos*, Phys. Rev. Lett. **81**, 1562–1567 (1998), hep-ex/9807003.

[SMS05] A. Slosar, A. Melchiorri and J. Silk, *Did Boomerang hit MOND?*, Phys. Rev. **D72**, 101301 (2005), astro-ph/0508048.

[SMS06] T. Sjöstrand, S. Mrenna and P. Skands, *PYTHIA 6.4 physics and manual*, JHEP **05**, 026 (2006), hep-ph/0603175.

[Str06] M. Strassler, Beyond the Standard Model, Lecture series given at the 1st Hadron Collider Physics Summer School, Fermilab, August 9-18, 2006.

[STS05] M. Schmaltz and D. Tucker-Smith, *Little Higgs review*, Ann. Rev. Nucl. Part. Sci. **55**, 229–270 (2005), hep-ph/0502182.

[SW05] J. M. Smillie and B. R. Webber, *Distinguishing spins in supersymmetric and universal extra dimension models at the Large Hadron Collider*, JHEP **10**, 069 (2005), hep-ph/0507170.

[tH71a] G. 't Hooft, *RENORMALIZABLE LAGRANGIANS FOR MASSIVE YANG-MILLS FIELDS*, Nucl. Phys. **B35**, 167–188 (1971).

[tH71b] G. 't Hooft, *RENORMALIZATION OF MASSLESS YANG-MILLS FIELDS*, Nucl. Phys. **B33**, 173–199 (1971).

[Tik04] V. Tikhomirov, Barrel TRT performance and electron/pion rejection, Talk given at ATLAS TRT Meeting, CERN, August 30, 2004.

[Tok90] W. Toki, *Review of straw chambers*, (1990), Invited paper given at 5th Int. Conf. on Instrumentation for Colliding Beam Physics, Novosibirsk, USSR, Mar 15-21, 1990.

[TRT04] T. Akesson et al. (ATLAS TRT Collaboration), *ATLAS Transition Radiation Tracker test-beam results*, Nucl. Instrum. Meth. **A522**, 50–55 (2004).

[UA183a] G. Arnison et al. (UA1 Collaboration), *Experimental observation of isolated large transverse energy electrons with associated missing energy at s**(1/2) = 540-GeV*, Phys. Lett. **B122**, 103–116 (1983).

[UA183b] G. Arnison et al. (UA1 Collaboration), *Experimental observation of lepton pairs of invariant mass around 95-GeV/c**2 at the CERN SPS collider*, Phys. Lett. **B126**, 398–410 (1983).

[UA283a] P. Bagnaia et al. (UA2 Collaboration), *Evidence for Z0 -> e+ e- at the CERN anti-p p collider*, Phys. Lett. **B129**, 130–140 (1983).

[UA283b] M. Banner et al. (UA2 Collaboration), *Observation of single isolated electrons of high transverse momentum in events with missing transverse energy at the CERN anti-p p collider*, Phys. Lett. **B122**, 476–485 (1983).

[VB05] R. Van Berg, Private communications., nov 2005.

[Vee84] R. Veenhof, *Garfield - simulation of gaseous detectors*, 1984, CERN Program Library W5050.

[WAH+57] C. S. Wu, E. Ambler, R. W. Hayward, D. D. Hoppes and R. P. Hudson, *EXPERIMENTAL TEST OF PARITY CONSERVATION IN BETA DECAY*, Phys. Rev. **105**, 1413–1414 (1957).

[Wei67] S. Weinberg, *A MODEL OF LEPTONS*, Phys. Rev. Lett. **19**, 1264–1266 (1967).

[Wol83] L. Wolfenstein, *PARAMETRIZATION OF THE KOBAYASHI-MASKAWA MATRIX*, Phys. Rev. Lett. **51**, 1945 (1983).

[Yan80] T. Yanagida, *Horizontal gauge symmetry and masses of neutrinos*, Prog. Theor. Phys. **64**, 1103 (1980).