The Level 1 Jet Trigger for ATLAS and a Preliminary Study of R-Parity violating Supersymmetry

Adam Connors

Thesis submitted for the degree of
Doctor of Philosophy

School of Physics and Astronomy
Faculty of Science
University of Birmingham

September 1998

Synopsis

This thesis covers work carried out by the author as part of the Level-1 Calorimeter Trigger Group for ATLAS. ATLAS (A Toroidal LHC ApparatuS) is a general purpose particle physics experiment currently being developed for use in the Large Hadron Collider (LHC) which is due to begin taking data in 2005.

Three main areas are discussed reflecting the author's main contributions.

• A range of potential jet trigger algorithms were investigated using a Monte Carlo simulation with a parameterised detector response. The performance of a number of algorithms was studied for single and multi-jet events for a range of physics processes. The baseline solution proposed in the ATLAS Technical Proposal was found to have a significantly degraded performance for multi-jet events and was discarded; two other candidate algorithms were found to provide an incorrect multiplicity count and were also discarded. The set of algorithms that gave the best performance now forms the current proposal; their overall performance was measured and the trigger efficiency for a realistic trigger rate given.

• The architecture of the Level-1 demonstrator system is described, and its performance in the September 1997 test beam discussed.

• Supersymmetry in the framework of the SUGRA model was studied in the case where explicit R_P violation causes the lightest supersymmetric particle to decay. Two decay modes

were considered: $\chi_1^0 \to eud$, and $\chi_1^0 \to \nu_e d\bar{d}$. The ability of ATLAS both to isolate an inclusive supersymmetric signal and to make precision measurements of the particle masses in this scenario was investigated using ATLFAST to model the expected performance of the ATLAS detector. It was found that ATLAS would be able to isolate an inclusive signal within the first year of low luminosity running for both scenarios, and in the case of charged lepton decays it was demonstrated that the masses of a complete decay chain could be directly reconstructed and used to constrain the parameters of the SUGRA model. In the case of neutral lepton decays precision measurements are more problematic and this is flagged as a potentially troublesome scenario for ATLAS.

The Author's Contribution

I contributed to the work of the Level-1 Calorimeter Trigger group in three main areas as described in the thesis synopsis.

The jet trigger analysis was carried out by myself under the supervision of Alan Watson. Throughout my course I have regularly presented the ongoing results of this work to the Level-1 group, and in October 1997 I presented my results to the larger trigger community at a trigger performance meeting held at CERN. The results of this analysis have been incorporated into the relevant sections of three documents: the level-1 Technical Design Report [22], the trigger menus document [25], and the trigger performance status report [29].

I have taken an active role in the level-1 demonstrator programme. During the debugging phase of new prototype modules I spent several weeks at Rutherford Appleton Laboratory assisting with the discovery and repair of hardware faults. During beam-test periods I have spent several months at CERN helping with the installation and debugging of the system. In addition I have been particularly involved in the various software packages used as part of the demonstrator programme. I was responsible for the design and development of several classes within the diagnostic software package, dealing with the modelling and interface of the cluster processing modules. During the last test-beam I was responsible for the continued development, maintenance, and debugging of the off-line analysis software, and for using this software to isolate low-level hardware faults.

During the supersymmetry analysis I regularly consulted with members of the ATLAS physics group. Whilst this analysis draws upon studies of R-parity violation in other scenarios [51], it includes several novel developments and was carried out entirely independently. Hopefully the results from this analysis will ultimately be incorporated into the ATLAS physics Technical Design Report.

Acknowledgements

- This thesis is of course dedicated to mum, David, Amy and Peter.
- Thanks are also due to my supervisor, Alan Watson, for all the hard work and effort that was needed to get me through a PhD.
- And finally to all my friends for doing their best to prevent me becoming a complete sad techie geek - with limited success. For posterity some names are recorded below (in no particular order):

Rob Whitney, Simon Carroll, Alex Robinson, Nick Drew, Jo, Brian Earl, Paul Thompson, Shaun Ashby, Michael Thompson, Paul Norman, Scott Talbot, Alun Lloyd, James Bloomer, Steve Ogelsby, Andy McGordon, Becky Dore, Mark Furness, Adrian Coles, Andy Cale, Ben Spurgen, Nigel Hart, Gail Rew, Niki Butler, Chris Burlton, Simon Lea, Phil Whitfield, James Pearson, Sophie Barber, Helen and Richard, Rachael Bailey, Gary Willington, Matthew Blackler, Iain Rhodes, Heidi, Andy Brown, Meera, Andy Cooper, Meg Riddington, Robert Last, Jessica Hunter, Dan Haagman, Caroline Pattenden, Caroline Lambert, Sophie and Chris, Chris H, Claire Wadey, Julian, Kate Dwyer, Christine Barnett, Elizabeth Kilbey, Sarah Smith, Alan Pain and Steph, Bob Khan, Duncan Percival, Steve Newman, Shaun Spain, Glynn Sweatnam, Catherine, Hannah Wolfe, and I'm sure I've forgotten some more...

I wish thee as much pleasure in the reading, as I had in the writing.

Francis Quarles.

Contents

List of Figures

List of Tables

Chapter 1

Physics Overview

1.1 Introduction

Particle physics is the study of matter at its most fundamental level; it attempts to describe the world around us in terms of the simplest set of rules possible that are compatible with observation.

One such theory that shows great promise is the Standard Model (SM) [1, 2]. The SM forms the framework for most of our current efforts in particle physics. It describes all known matter in terms of just two different types of fundamental particle (quarks and leptons), and it describes the interactions of these particles in terms of just four types of force.

This chapter briefly summarises the SM. The SM has made many significant predictions that are supported by experimental evidence, and so far there is no evidence contrary to the SM. However it is still considered an incomplete theory. The final part of this chapter describes why this is the case and looks at possible extensions to the SM. Chapter 2 presents one of the next generation of experiments (ATLAS) designed to further our understanding of the SM and physics beyond it.

1.2 The Standard Model

In describing the universe two things must be considered: the matter that makes up the world around us, and the forces through which that matter interacts. In the Standard Model both matter and force are manifest as different types of particle. Particles of matter are called fermions; they have an intrinsic spin of $\frac{1}{2}$, and obey Fermi-Dirac statistics. These fermions interact with each other by the exchange of a particle of force - called a boson; bosons have an integer intrinsic spin, and obey Bose-Einstein statistics.

There are twelve different kinds of fermion (6 quarks, and 6 leptons) and four different kinds of force - electromagnetic (EM), weak, strong, and gravitational. So far only the electromagnetic, weak, and strong forces can be successfully described within the Standard Model. Gravity is not included in the model, and can only be understood in terms of general relativity. This does not affect the predictions of the Standard Model however, since at the scale of experimental particle physics gravity can be considered negligible. Nonetheless it would be highly desirable to include gravity and theories have been formulated that attempt to do this, although none are entirely satisfactory at present.

1.2.1 The Elementary Particles and their Interactions

According to the Standard Model all matter is made up from just twelve kinds of particle. These particles are grouped, according to the types of force they couple to, into quarks and leptons. The quarks and leptons are further divided into generations of doublets as shown in table 1.1.

These particles are believed to be fundamental and hence have no structure. The properties of the first generation are replicated in higher generations but with increasing mass, except for the neutrinos which are believed to have very little or zero mass [a].

[a] Until very recently neutrinos were considered massless. However results presented at the 4th June Neutrino '98 Conference in Takayama, Japan, by the SuperKamiokande Collaboration offer the first convincing evidence that neutrinos do in fact possess a small mass [3]. The full impact of this discovery is not yet understood.

Fermion	**Generations** I	II	III	**Charge**	**Interactions**
Quarks	$\begin{pmatrix} u \\ d \end{pmatrix}$	$\begin{pmatrix} c \\ s \end{pmatrix}$	$\begin{pmatrix} t \\ b \end{pmatrix}$	$+\frac{2}{3}$ $-\frac{1}{3}$	EM, Weak, Strong, Gravity
Leptons	$\begin{pmatrix} \nu_e \\ e^- \end{pmatrix}$	$\begin{pmatrix} \nu_\mu \\ \mu^- \end{pmatrix}$	$\begin{pmatrix} \nu_\tau \\ \tau^- \end{pmatrix}$	0 -1	Weak, Gravity EM, Weak, Gravity

Table 1.1: The fundamental fermions can be grouped into quarks and leptons, which are further divided into generations of doublets. Only the first generation is stable.

The quarks combine to form all the observed hadrons [b]. A quark - anti-quark pair can combine to form a meson, or three quarks can combine to form a baryon. Most hadrons are unstable however, and can only be observed in particle physics experiments. Only baryons made up from the first generation of quarks are observed in everyday matter; for example the proton is made up from 2 up (u) type quarks and 1 down (d) type quark, and the neutron is made up from 2 down type quarks and 1 up type quark. Protons and neutrons combine to form atomic nuclei which then bind with electrons to form the atoms which make up all common matter.

The forces between the fermions are caused by the exchange of field quanta (the bosons). This can be visualised with a Feynman diagram as shown in figure 1.1a. Here two electrons scatter off each other, and the momentum exchanged between them is carried by a virtual photon. Each of the four forces is mediated by a different type of boson. The electromagnetic force as shown is mediated by a photon; the strong force is mediated by a gluon; and the weak force is mediated by either a W^+, W^-, or Z^0 boson.

[b] Hadron is a generic term for all bound quark states.

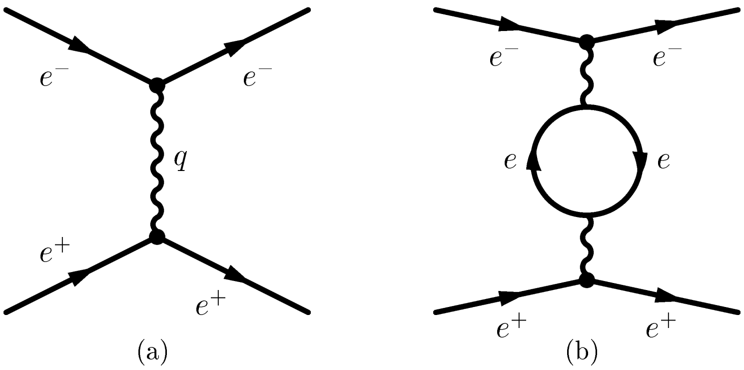

Figure 1.1: An example of e^+e^- scattering via the exchange of a photon, for a) Leading Order, and b) 2nd Order.

The influence that each force has on a particle is determined primarily by the coupling strength between that force and the particle. The electromagnetic force couples to any particle that carries electric charge (all except the neutrinos), and the weak force couples to all particles. The electromagnetic and weak force both have approximately the same coupling strength but the influence of the weak force, and hence the amplitude for weak processes, is far smaller. This is because the weak bosons have mass. In general the bosons mediating an interaction are virtual; this introduces a "propagator" term into the amplitude for a given process, which describes the probability of the boson travelling between the two vertices of an interaction. For the weak bosons this propagator is roughly inversely proportional to the boson mass and it is the presence of this factor that strongly suppresses the overall amplitude of weak processes relative to electromagnetic processes. This feature can be visualised in terms of Heisenberg's uncertainty principle; since the weak bosons produced are virtual they can exist only fleetingly in accordance with the uncertainty principle; because of their mass, their range (and hence their influence) is very small.

The gluons have a very strong coupling to the quarks, but no coupling to the leptons (hence the leptons do not feel the strong force). The gluons are massless but their range is also very short. This is caused by the "self-interaction" of the gluons and will be discussed later.

These few particles and their interactions describe all the observed complexity of our universe, but this in itself is little more than a catalogue of observations. What gives the SM its predictive powers is the theoretical framework at the root of its construction - this framework is the Gauge Field Theory.

1.2.2 Gauge Field Theories

Central to quantum physics is the relationship between conserved quantities and physical symmetries; that is, every symmetry has associated with it a conserved quantity. For example: symmetry under spatial translation leads to conservation of momentum; symmetry under translation in time leads to energy conservation; and invariance under spatial rotation leads to conservation of angular momentum [4].

Gauge field theories are an extension of this notion; the aim in constructing a gauge theory is to find a Lagrangian that describes the interactions between the relevant particles, and is invariant under an appropriate transformation. In order for this requirement to be satisfied new terms have to be introduced to the Lagrangian that can be associated with the physical gauge bosons through which the fermions interact with each other. Such a Lagrangian constitutes a quantum field theory that describes a particular kind of interaction, and can be used to make calculations about the probability of such interactions occurring.

Physical predictions are made from quantum field theories using a perturbative approach. This approach is embodied in the Feynman diagrams which represent the interactions; figure 1.1a shows the “tree level” diagram for e^+e^- scattering. An important feature of this approach is that a calculation of the amplitude for this process based on this diagram is only correct to leading order. For more accurate calculations higher-order processes have to be taken into account known as radiative corrections (eg. figure 1.1b). When calculating the contribution of such closed loop diagrams an integration is required over the full range of energy and momentum that can be carried by the virtual electron-positron pair. This integral is divergent and an infinity is encountered.

The way around this problem introduces a crucial feature into gauge field theories known as **renormalisation**. Renormalisation is a mechanism that allows us to avoid these divergent corrections by re-defining certain parameters of the Lagrangian (for example the coupling constants) in terms of their experimentally observed values. As an example of this consider the charge of the electron. In calculating the charge of the electron it is necessary to consider the fact that the electron will be continually emitting and re-absorbing photons, which in turn convert to e^+e^- pairs. This creates a virtual charge cloud around the electron as illustrated in figure 1.2. A calculation of the charge of the electron must consider all possible radiative corrections, and would become divergent. Empirically the charge of the electron is found to be finite, and so these divergent integrals must be encapsulated by substituting the measured charge of the electron into the Lagrangian, in place of the original (bare-charge) parameter and the divergent terms. The postulate is made that the bare-charge of the electron is equally divergent in such a way to cancel the divergences of the radiative corrections.

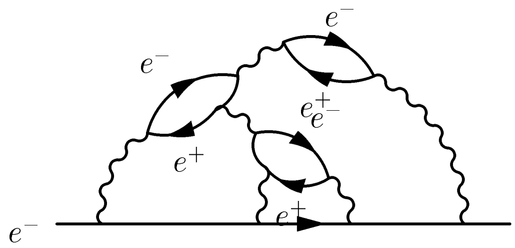

Figure 1.2: An illustration of the virtual charge cloud that surrounds any charged object.

A consequence of this notion is that the charge of the electron is never measured directly. In the charge cloud illustrated in figure 1.2 the virtual positrons will be preferentially closer to the electron than the virtual electrons due to electro-static attraction. This is termed vacuum polarisation in analogy with the electro-statics of dielectrics, and causes the bare-charge of the electron to be shielded. Experimentally the charge on the electron (or equivalently the electromagnetic coupling constant) is measured by studying the electron's interaction with another test charge. Since the test charge distance is related to the 4-momentum exchanged in this interaction (Q^2), charge shielding is seen as a "running" of the electromagnetic coupling constant as a function of Q^2.

1.2.2.1 Quantum Electro-Dynamics

The first successful gauge theory was quantum electro-dynamics (QED) [5]. This theory requires that the Lagrangian for a single free Dirac particle (i.e. a fermion) be invariant under a local phase transformation in "charge space."

Consider the Lagrangian describing a fermion of mass m propagating in free space,

$$\mathcal{L} = i\bar{\psi}\gamma^{\mu}\partial^{\mu}\psi - m\bar{\psi}\psi$$

and the transformation,

$$\psi \rightarrow \psi' = e^{iq\Lambda(x)}\psi(x)$$

Where $\Lambda(x)$ is a parameter that can take any real value and is a function of position. Transformations of this type belong to the group U(1). The fact that it is a function of x makes this a *local* gauge transformation, which is critical for generation of the electromagnetic coupling.

The first term in the Lagrangian leads to extra terms when transformed that violate gauge invariance. In order to maintain gauge invariance the derivative, ∂^{μ}, has to be redefined so that it transforms in the same way as the fields themselves, eg:

$$\mathcal{D}_{\mu}\psi \rightarrow \mathcal{D}'_{\mu}\psi' = e^{iq\Lambda(x)}\mathcal{D}_{\mu}\psi.$$

This is achieved by the introduction of a new vector field A_{μ},

$$\mathcal{D}_{\mu} \equiv \partial_{\mu} + iqA_{\mu}(x),$$

which transforms as,

$$A_\mu(x) \to A_\mu(x) - \partial_\mu \Lambda(x).$$

The new, locally gauge invariant, Lagrangian now takes the form:

$$\mathcal{L} = \mathcal{L}_{FREE} - j^\mu A_\mu$$

where the additional term, $-j^\mu A_\mu$, is familiar as the interaction between a free particle and an electromagnetic field and j^μ is the conserved electromagnetic current,

$$j^\mu = q\bar{\psi}\gamma^\mu\psi.$$

Requiring local gauge invariance has therefore forced the inclusion of a new term, which we can associate directly with the physical photon field (A_μ).

1.2.2.2 Electroweak Theory

The electromagnetic and the weak nuclear forces were originally considered to be quite separate phenomena. But in the 1960's Glashow, Weinberg, and Salam [6–8] successfully unified the weak and electromagnetic forces with a single Hamiltonian, invariant under the combined transformation $\mathrm{SU(2)_L} \times \mathrm{U(1)}$.

The transformation $\mathrm{SU(2)_L}$ is a rotation in "weak isospin space." Weak isospin is defined analogously to nucleon isospin by assigning the left-handed quarks and leptons to weak isospin doublets (table 1.2). The quark states d', s', and b' are weak isospin eigenstates for the quarks, and are in fact mixed quark mass eigenstates as described by the CKM [9] matrix. Only left-handed particles are considered to carry weak isospin since charged weak currents (i.e. $W^\pm$ interactions) don't interact with right-handed particles. The right-handed particles are written as SU(2) singlets; note that right-handed neutrinos are not found in nature.

The $\mathrm{SU(2)_L}$ transformation couples only to the left-handed particles, and is a rotation between states in a doublet; eg. an electron rotates into an electron-neutrino. Requiring local invariance to this kind of transformation requires the introduction of a triplet of scalar fields;

Left-Handed Doublets:	$\begin{pmatrix} \nu_e \\ e \end{pmatrix}$		$\begin{pmatrix} \nu_\mu \\ \mu \end{pmatrix}$		$\begin{pmatrix} \nu_\tau \\ \tau \end{pmatrix}$	$\begin{pmatrix} u \\ d' \end{pmatrix}$		$\begin{pmatrix} c \\ s' \end{pmatrix}$	$\begin{pmatrix} t \\ b' \end{pmatrix}$
Right-handed Singlets:	$(e)_R$	$(\mu)_R$	$(\tau)_R$	$(u)_R$	$(d)_R$	$(c)_R$	$(s)_R$	$(t)_R$	$(b)_R$

Table 1.2: Grouping of fundamental fermions into left-handed weak isospin doublets, and right-handed weak isospin singlets.

one negative, one neutral, and one positive. The charged components can be associated directly with the $W^{\pm}$ bosons that mediate the charged current interactions. The neutral component however cannot immediately be associated with the Z^0 (neutral current) since the Z^0 couples both to left- and right-handed particles.

It is interesting that not only is electroweak unification desirable from a conceptual viewpoint, it is necessary to fully describe the neutral current. Requiring gauge invariance under the combined transformation $\mathrm{SU(2)_L} \times \mathrm{U(1)}$ generates, in addition to the triplet of weak isospin currents, a neutral field that couples to "weak hypercharge", defined by,

$$Q = t_3 + \tfrac{Y}{2}.$$

Where Q is charge, t_3 is the 3rd component of weak isospin, and Y is the weak hypercharge. Now the physical weak neutral current and the electromagnetic field are given by linear combinations of the neutral weak isospin component and the weak hypercharge field:

$$Z_\mu = W_3 cos\theta_W - B_\mu sin\theta_W$$
$$A_\mu = W_3 sin\theta_W + B_\mu cos\theta_W.$$

Where A_μ is the electromagnetic field mediated by the photon, Z_μ is the weak neutral current mediated by the Z^0, W_3 is the neutral component of the weak isospin triplet generated by $\mathrm{SU(2)_L}$, B_μ is the weak hypercharge field generated by U(1), and θ_W is the weak mixing angle that describes how the fields mix.

The only flaw in this approach is that it predicts massless gauge bosons. Any attempt to include a mass term in the Lagrangian destroys the gauge invariance. This is at odds with experiment which has measured the masses of the $W^{\pm}$ and Z^0 bosons to be 81 GeV/c^2 and 91 GeV/c^2 respectively. Furthermore if divergent cross-sections in the production of W pairs are to be avoided, cancellations must occur between the contributing Feynman diagrams; these cancellations require that the electromagnetic and the weak and neutral current couplings must all be the same which, in order to explain the differing strengths of the EM and weak interactions, requires that the weak bosons are massive.

This problem was solved by the introduction of a new boson, the Higgs [10], which gives mass to the other particles by a process of spontaneous symmetry breaking known as the Higgs Mechanism.

1.2.2.3 Spontaneous Symmetry Breaking and the Higgs Mechanism

Four new fields are introduced to the electroweak Lagrangian, arranged as a weak isospin doublet:

$$\phi = \begin{pmatrix} \phi^+ \\ \phi^0 \end{pmatrix} \text{ where } \begin{array}{c} \phi^+ = \frac{1}{\sqrt{2}}(\phi_1 + i\phi_2) \\ \phi^0 = \frac{1}{\sqrt{2}}(\phi_3 + i\phi_4) \end{array}$$

The electroweak Lagrangian can be considered in the form $\mathcal{L} = T - V$, where T is the kinetic energy and V is the self-interaction potential and can be written as:

$$V(\phi^\dagger\phi) = -\mu^2\phi^\dagger\phi + \lambda(\phi^\dagger\phi)^2$$

Where μ^2 and λ are constants. Crucially, this potential has a minimum at $|\phi|^2 = \frac{\mu^2}{2\lambda}$, and so its vacuum expectation value is non-zero. The Lagrangian remains gauge invariant as required, but the physical ground state need not be. This arises because the Hamiltonian is symmetric in weak isospin space, but has a minimum when one component of the Higgs doublet is non-zero. To visualise this scenario more easily it can be thought of analogously to the situation in ferromagnetism. The Hamiltonian describing the spin-spin interactions in a ferromagnetic crystal is invariant under spatial rotations. Yet below the Curie temperature the spins are aligned in the direction of magnetisation. Thus, in a ferromagnet there is a preferred direction in space, although *any* direction is equally possible. An observer inside the crystal would not be aware of its rotational symmetry, although the Hamiltonian describing it is still fundamentally symmetric. This is referred to as a "hidden symmetry."

To apply perturbation theory to the electroweak Lagrangian it is necessary to expand about one of these minima. The physical vacuum is re-defined as

$$\phi_0 = \begin{pmatrix} 0 \\ \frac{\mu}{\sqrt{2\lambda}} \end{pmatrix},$$

and the Lagrangian is re-written in terms of this shifted field. It is the act of defining a physical vacuum that is referred to as "spontaneous symmetry breaking." This procedure generates new terms which correspond to the masses of the $W^{\pm}$, and Z^0, but preserves the underlying gauge symmetry responsible for generating the gauge bosons in the first place. Of the scalar fields which comprise the Higgs doublet, three disappear. One neutral component acquires a mass term itself, and it is this term that predicts the existence of a massive Higgs boson.

This Higgs boson is as yet undiscovered despite many careful searches. Theory does not predict the mass of the Higgs, although if other problems with the SM are to be avoided then it must have a mass below around 1 TeV (see section 1.3). Despite being a promising solution, the Higgs mechanism lacks direct experimental evidence, and represents a particularly critical region of the SM that requires further study.

1.2.2.4 Quantum Chromo-Dynamics

Quantum chromo-dynamics [11] is the gauge theory that describes the strong interaction of the quarks. In order to explain the observation that quarks exist only as quark anti-quark pairs in mesons, or as quark triplets in baryons, the proposed charge for the strong force is defined in analogy with colour theory. That is, there are three 'colours' to the strong charge: Red, Green, and Blue. Quarks carry 1 unit of colour charge which can be any colour; but free objects can only exist that are colour neutral. So mesons are made up from $colour - \overline{colour}$

combinations, and the baryons are colour neutral through combination of R G and B quarks. This postulate has been verified by measuring the cross-section of $e^+e^- \rightarrow q\bar{q}$ which is three times larger than would be expected if colour charge did not exist; this reflects the 3 extra degrees of freedom introduced by the colour charges.

The three colour charges of the quark form an SU(3) symmetry group, and similarly to QED and Electroweak gauge theories, gauge invariance is required under local transformations in colour. This introduces an octet of gauge bosons, named gluons, which are massless, and carry colour charge in $colour - \overline{colour'}$ combinations.

An interesting consequence of this is that the gluons are found to carry colour charge. This has a profound effect on the nature of charge shielding. Section 1.2.2 discussed the charge shielding of the electron, where a polarised cloud of virtual electrons and positrons shield the bare-charge of the electron. In the case of gluons the shielding produced is fundamentally different due to the self interaction of the gluons. It can be shown that this creates an anti-shielding where the observed influence of the gluon increases with increasing distance. This explains why quarks are never found in free states since in principle they are only "asymptotically free"; i.e. they are only free at very short distances, or equivalently very high Q^2. As a quark moves away from a colour neutral region the colour force acting on it increases, until it becomes energetically more favourable to generate new quarks which combine with the escaping quark and form a 'jet' of hadrons. The process is called fragmentation and is observed experimentally.

1.3 Beyond the Standard Model

Despite the many successes of the Standard Model, it is not considered a complete theory. There are 19 free parameters that are not predicted by the theory, and can only be measured empirically. For example the magnitudes of the fermion, and boson masses, and the size of the electroweak mixing angle, must all be put in by hand. Although the gauge part of the SM is well understood, the Higgs sector is not, and the energy range required to prove, or exclude, this theory is potentially beyond the range of current accelerators.

Furthermore, it would be desirable to more fully unify the forces into a single gauge group. The so-called unification of the electroweak force is in fact the product of two separate groups of gauge transformations. True unification of the fundamental forces would require a single, larger gauge group of which SU(2), U(1), and SU(3) would all be subgroups. In 1974 Georgi and Glashow showed that the smallest possible candidate for this hypothetical grand unified theorem (GUT) symmetry group is SU(5) [12]. Several attempts to build a theory around a group like this have been made [1] but so far none have been entirely satisfactory.

More crucially, there is a so-called "naturalness problem" which cannot be solved by the Standard Model, or by extended GUT symmetry groups. Coupling of the Higgs boson to the electroweak gauge bosons is necessary, not only to give them mass, but also to cancel divergent integrals which occur in electroweak radiative corrections to W and Z scattering. Without these cancellations the cross-sections for longitudinal W and Z scattering violates unitarity at energies above about 1 TeV.

For these cancellations to work, the mass of the Higgs itself, must be below around 1 TeV. However, radiative corrections to the Higgs mass are divergent and themselves require a cancellation to occur between the interactions of the Higgs with all other particles in the theory, and the bare Higgs mass at the Planck scale.

If the mass of the Higgs is to be of the order of the electroweak scale, then a cancellation of 1 part in 10^{16} must occur. That cancellations this finely tuned should occur between two independent effects is considered unlikely, and suggests that there might be some new principle at work.

Exactly what this new principle might be, is unclear. But these arguments all point to the same conclusion: *some kind* of new physics must be discovered at energies below $\approx$1 TeV. Understanding this new physics is the prime motivation for building the Large Hadron Collider (LHC) described in the following chapter.

There are several theories that attempt to describe physics in this region. One of the most promising is Supersymmetry which is described below.

1.3.1 Supersymmetry

Supersymmetry (SUSY) postulates the existence of superpartners for all the presently observed particles [13]. These superpartners have identical quantum numbers to their counterparts, but differ in spin by $\frac{1}{2}$. Therefore every fermion (the leptons and quarks) is associated with a boson-like superpartner (denoted sleptons, and squarks); and every gauge boson is associated with a fermion-like superpartner conventionally called a *gaugino*.

This approach is favoured for several reasons:

- It is represented by the only possible non-trivial extension to the Lorentz group which lies at the heart of quantum field theory. As such it is a very natural, and aesthetic extension to current theories.
- It provides a very elegant solution to the naturalness problem. This is because radiative corrections to the Higgs mass must now include both fermionic and bosonic corrections which have the opposite sign, and will exactly cancel.
- If supersymmetry is formulated as a local symmetry then a spin-2 graviton field must be introduced. This leads naturally to the inclusion of gravity as a quantum field theory.
- In some supersymmetric models the lightest supersymmetric partner is stable, electrically neutral, and weakly interacting. With a mass of around 100 GeV/c^2 this makes it a good candidate for cold dark matter.

In its simplest form SUSY introduces 32 new fundamental particles; this is termed the *Minimal Supersymmetric extension to the Standard Model* (MSSM) and its particles are enumerated as follows:

- 12 partners to the quarks; six of each chirality.
- 6 partners to the charged leptons; 3 of each chirality.

- 3 partners to the neutral leptons; only left-handed.
- 1 partner to the gluon.
- To give the superpartners mass 2 higgs doublets are now required resulting in a physical spectrum consisting of 3 neutral and 2 charged higgs bosons.
- 2 mixed states of the partners to the charged electroweak and Higgs bosons; collectively named charginos.
- 4 mixed states of the partners to the neutral electroweak and Higgs bosons; collectively named neutralinos.

It is assumed that supersymmetry is a broken symmetry since none of the supersymmetric particles have been observed. The exact mechanism for this symmetry breaking is unknown and several theories exist. One of the most popular and most thoroughly explored is a supergravity model (SUGRA) in which the symmetry breaking is driven by gravitation-like interactions between the MSSM particles and some higher energy superfield referred to as the 'hidden sector'.

This model, and possible ways that ATLAS could be used to investigate it, is studied in more detail in Chapter 6.

Chapter 2

The ATLAS Experiment

2.1 Introduction

In order to explore physics in the TeV energy range, CERN [(a)] in 1994 approved the Large Hadron Collider (LHC) [14,15].

The LHC is a proton-proton collider operating at a centre of mass energy of $\sqrt{s} = 14$ TeV. It will be built in the existing LEP (Large Electron Positron [16]) tunnel at CERN in Geneva, and is due to start operating in 2005. Figure 2.1 shows the layout of the LHC and its associated accelerators and detectors. The LHC consists of two storage rings which are fed by the Super Proton Synchrotron (SPS) at 0.45 TeV and then accelerate counter-rotating bunches of protons to 7 TeV in each direction. The centre of mass energy must be considerably higher than the 1 TeV region being explored since in proton collisions the quarks and gluons carry only a fraction of the proton momentum and it is these, rather than the protons themselves, which collide at high Q^2.

The LHC has a design luminosity of $10^{34} cm^{-2} s^{-1}$ but will have an initial running period (of approximately 3 years) with a lower luminosity of $10^{33} cm^{-2} s^{-1}$. In addition there is scope for an alternative mode where beams of lead ions are collided.

[(a)] European Organisation for Nuclear Research

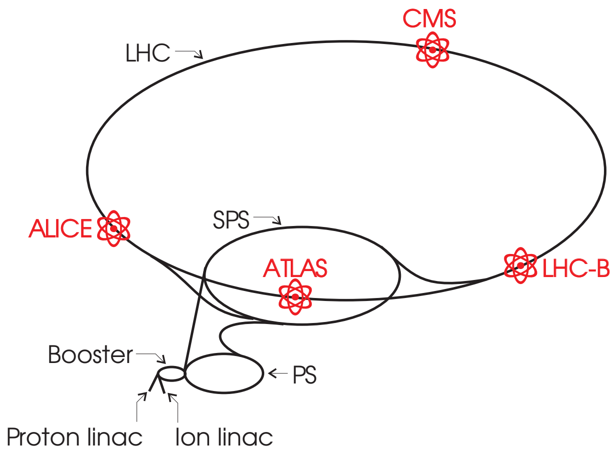

Figure 2.1: The layout of the LHC accelerator complex, and its detectors.

There are four experiments located at four interaction points around the circumference of LHC. ALICE and LHC-B are specialised experiments designed for heavy ion and B-physics studies respectively. ATLAS (A Toroidal LHC ApparatuS) and CMS (Compact Muon Solenoid) are both general purpose experiments designed to explore the full physics capabilities of the LHC.

This thesis is concerned primarily with the ATLAS experiment, and in particular with the Level-1 Calorimeter Trigger for the ATLAS experiment.

2.2 Physics Potential

As well as looking for possible extensions to the standard model, the LHC will provide many opportunities for more exact measurements of known physics. A detailed study of the physics potential of the LHC can be found for example in [17]. Listed below is a brief summary of some of the prime goals of LHC experiments:

- **Search for the SM Higgs:** The Electroweak symmetry breaking mechanism of the standard model predicts the existence of a single scalar Higgs boson. The discovery channel for this boson will depend on its mass: For $m_H < 150$ GeV/c^2 the Higgs could be observed through H $\rightarrow b\bar{b}$ or H $\rightarrow \gamma\gamma$; for higher Higgs masses H $\rightarrow ZZ$, or H $\rightarrow WW$, with either real or virtual ZZ or WW would provide very clean multi-lepton signatures.

- **B-physics:** The large number of b-quarks produced during the initial low luminosity running of the LHC will provide an excellent environment for precision B-physics measurements. Of particular interest is the search for and measurement of CP-violation, and the determination of the interior angles of the CKM unitarity triangle.

- **Top physics:** The LHC also provides an excellent opportunity for detailed top measurements: eg, more accurate measurement of top mass, and searches for top decays through supersymmetric processes.

- **Searches for supersymmetry:** The LHC will play a prime role in the discovery or exclusion of SUSY extensions to the SM.

- **Other extensions to the SM:** GUT-like extensions to the SM predict the existence of *leptoquarks* (particles carrying both lepton and baryon number). Other solutions to the SM hierarchy problem require the existence of quark compositeness, or new strong interactions.

- **Heavy ion physics:** $Pb - Pb$ collisions could lead to the creation of the predicted quark-gluon plasma (QGP). Discovery and measurement of the QGP would allow the detailed study of the strong interaction under very different conditions, further testing the validity of gauge theory.

2.3 The ATLAS Detector

Design of the ATLAS detector is driven by the desire to optimise its performance over a broad range of physics processes. This versatility provides the best chance of finding new, unexpected physics, and allows measurements to be cross-checked through a variety of channels. A detailed description of ATLAS can be found in the Technical Proposal [18], and in the relevant Technical Design Reports of its sub-detectors [19–22]. This section summarises the chief elements of the ATLAS detector.

Figure 2.2 shows a diagram of the proposed ATLAS detector. Protons will enter from the extreme left and right hand sides of the detector and collide at its very centre. As particles move out from the interaction point they will encounter a series of sub-detectors. These sub-detectors are:

- **Inner detector:** Consisting of pixel detectors, silicon strip tracking, and continuous straw tube trackers.

- **Calorimeters:** Consisting of electromagnetic and hadronic calorimetry.

- **Muon chambers.**

The inner detector is contained within a super-conducting solenoid which generates a 2 T magnetic field necessary for measurement of the charged particle momenta.

The return yoke for this solenoid is integrated into the support structure for the hadronic calorimeter. A second system of toroidal magnets lies outside the calorimeter and provides the magnetic field for the muon spectrometer. The muon spectrometer is built up from muon chambers placed throughout this magnetic volume. The overall dimensions for ATLAS are determined by the muon spectrometer: radius = 11 m, and overall length = 42 m.

The following sections describe the sub-detectors in more detail. Throughout this thesis the following convention will be used for the coordinate system: The interaction point is taken to be the origin; the z-axis is parallel to the beam direction, and x, y or the radius R and azimuthal angle ϕ are used to denote positions in the transverse plane; the polar angle may be denoted directly as θ but more usually will be expressed in terms of the Lorentz invariant pseudo-rapidity $\eta = -ln\ tan(\frac{\theta}{2})$.

2.3.1 Inner Detector

The primary requirements on the inner tracking detector are as follows:

- Tracking coverage over pseudo-rapidity range $|\eta|$ <2.5.
- Efficient momentum measurement of leptons over all this range.
- Tagging of b-jets.
- Measurement of the z-coordinate of the primary interaction vertex.
- At low luminosity running ($10^{33} cm^{-2} s^{-1}$) the inner detector is also required to reconstruct the secondary vertices from b-quark and tau decays.

To accomplish these goals with the high track densities expected at the LHC will require very high-precision measurements to be made with fine granularity detectors. Figure 2.3 shows the layout of the inner-detector at ATLAS. The inner detector combines three technologies: pixel detectors, silicon strip detectors (called the semi-conductor tracker or SCT), and continuous straw tube tracking (called the transition radiation tracker or TRT).

Semi-conductor pixel trackers offer the finest granularity and are located closest to the beam pipe. At larger radii, where lower granularity can be tolerated, silicon strip detectors are used. At the outer radius of the inner detector, tracks can be followed using continuous straw tube tracking; this detector is called the transition radiation tracker (TRT) since in addition to tracking it provides additional electron identification by measuring transition radiation photons.

This combination offers the very robust pattern recognition of a large number of continuous track measurements over a long track length, as well as a number of high-precision points. But it minimises the amount of material in front of the calorimeters that would degrade their resolution.

Mechanically the inner detector is split into a barrel region and two identical forward regions, with a transition at $|\eta| \simeq 1$. In the barrel region, layers of detectors run parallel to

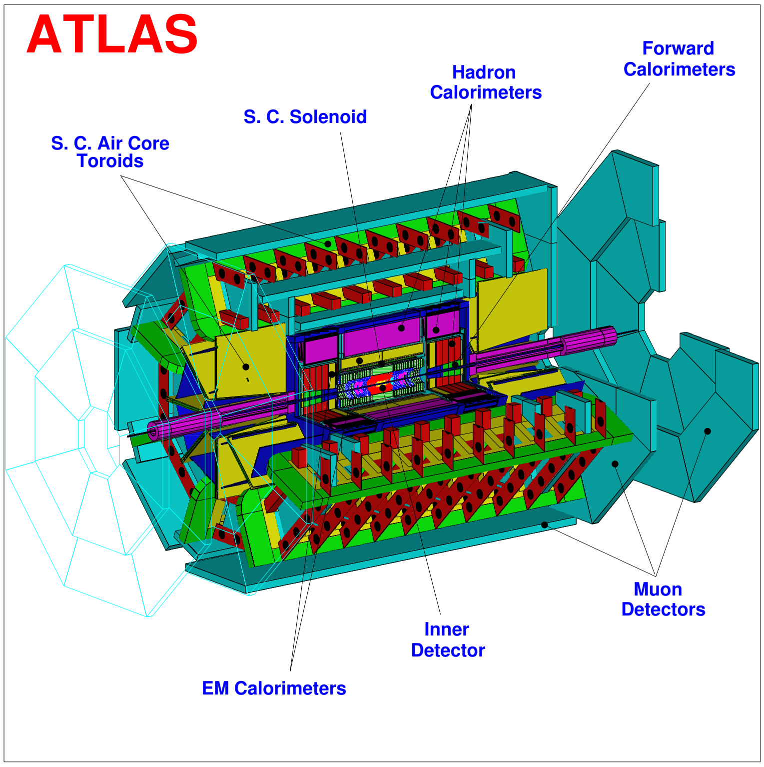

Figure 2.2: The ATLAS detector.

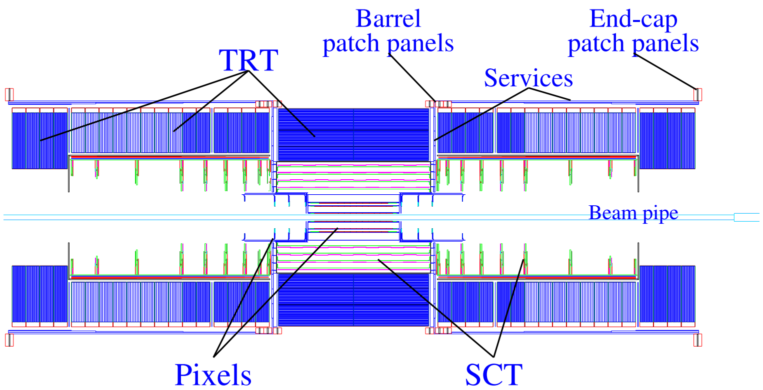

Figure 2.3: A side-on view of the ATLAS inner tracking showing the three kinds of tracking used: The innermost pixel detector, the semi-conductor tracker (SCT), and the transition radiation tracker (TRT).

the beam line; in the forward regions layers of detectors are arranged in discs perpendicular to the beam pipe.

The detectors are arranged so that each track must pass through at least three pixel layers, four strip layers, and provide around 36 TRT points. The outer radius of the tracker cavity is 115 cm, and its total length is 7 m. The barrel region extends over ±80 cm. The precision tracking elements are contained within a radius of 56 cm; in order to give uniform coverage over the full acceptance, the final TRT wheels at high z must extend inwards to a lower radius than the rest of that detector.

2.3.1.1 Pixel Detector

In order to untangle the very high track densities at the innermost radii the pixel detector is used. The pixel detector is also useful to tag short-lived particles such as b-quarks and τ-leptons. Three barrel layers, and eight discs (four in each forward region) will initially be installed. The two-dimensional segmentation of each layer means that the three dimensional coordinates of each point can be read-out directly. Resolution in the barrel is $R\phi$=12 μm by z=66 μm; resolution in the discs is $R\phi$=12 μm by R=77 μm. In total the pixel detector provides 140 million channels.

The innermost layer of the pixel detector will eventually be destroyed by radiation damage. Originally it was planned to use this layer only during initial low luminosity studies, but recent simulations have demonstrated the value of this layer for Higgs and SUSY studies and so the possibility to replace this layer at a later date has been retained.

2.3.1.2 Semi-conductor Tracker

The high cost, and large material content, of the pixel detector means it cannot be used throughout the inner detector. At larger radii lower track densities, and larger average track radii, mean that a lower granularity can be tolerated and silicon strip detectors can be used.

There are four layers of silicon microstrip detector modules in the barrel region. Each module contains two sets of 768 read-out strips each of 80 μm pitch, lying back-to-back at an angle of 40 $mrad$ relative to each other. The ϕ coordinate can be measured directly, whilst the use of "stereo angle" strips allows the z coordinate to be obtained by correlating the results between sets of strips.

The forward detectors are similar but with strips running radially to the beam line; hence the ϕ and z coordinates are measured directly and R can be reconstructed. There are nine end-cap discs on each side.

The detector generates 6.2 million read-out channels which must be stored in buffered chips directly bonded to the detector substrate while awaiting a level-1 trigger decision. In the barrel, the spatial resolution is 16 μm in $R\phi$ and 580 μm in z; whilst in the end-caps resolution is $R\phi$=16 μm by R=580 μm. Hits can be distinguished if separated by more than $\simeq$200 μm in $R\phi$.

2.3.1.3 Transition Radiation Tracker

At the outermost radius of the inner detector continuous straw tube tracking is used. This device consists of drift tubes with a radius of 4 mm (straws). The straws run parallel to the beam line in the barrel region, and radially in the forward regions; they are arranged in such a way that any particle with $\eta < 2.5$ must pass through at least 36 straws, offering continuous tracking out to this rapidity. The ϕ coordinate is reconstructed from a drift-time measurement in each tube, giving a resolution of 170 μm.

The ability to continuously track particles with a large number of points greatly improves the pattern recognition abilities of the inner detector. In addition the straw tubes of the TRT are interleaved with radiators which cause electrons to emit transition x-rays which can be detected in order to improve electron identification; transition-radiation hits and tracking hits are distinguished by applying two independent thresholds to each channel.

2.3.2 Calorimetry

The role of the ATLAS calorimetry is to measure the energy of the photons, electrons, and jets, and allow reconstruction of the missing transverse energy (E_T^{MISS}).

The electrons and photons interact with the calorimetry electromagnetically, while the hadrons also interact via the strong force. This means that they have very different interaction lengths requiring the calorimetry to be split into electromagnetic (EM) calorimetry, and hadronic calorimetry.

Figure 2.4 shows the calorimeter set-up for ATLAS. EM and hadronic calorimetry is split into three regions: the barrel, and two end-caps. All EM calorimeters use liquid argon (LAr) technology and must be housed in cryostats. The barrel part of the hadronic calorimeter uses a less expensive iron-scintillator system and is called the "tile calorimeter." The tile calorimeter is further split mechanically into a central barrel, and two extended barrel regions. The hadronic end-cap calorimeters use LAr technology, as do the integrated forward calorimeters which provide extreme forward coverage out to $3.2 < |\eta| < 4.9$.

LAr calorimetry is contained in three cryostats: The barrel cryostat which houses just the barrel EM calorimeter and the inner solenoid; and two end-cap cryostats that house both the EM and hadronic end-cap calorimeters and the integrated forward calorimeter.

2.3.2.1 Electromagnetic Calorimeter

The barrel EM calorimeter covers rapidity $|\eta| \leq 1.4$, and the end-caps complete coverage up to $|\eta| < 3.2$ where the integrated forward calorimeter takes over.

The EM calorimeter is a sampling calorimeter. It is based on absorber plates immersed in an ionising medium, with read-out electronics located between the absorber plates. In this case the absorber plates are lead and are "accordion-shaped", and the ionising medium used is liquid argon. In the barrel the lead absorbers lie radially in ϕ; in the end-caps the absorbers lie parallel to the beam axis, and surround the beam pipe like the spokes of a wheel (see figure 2.5).

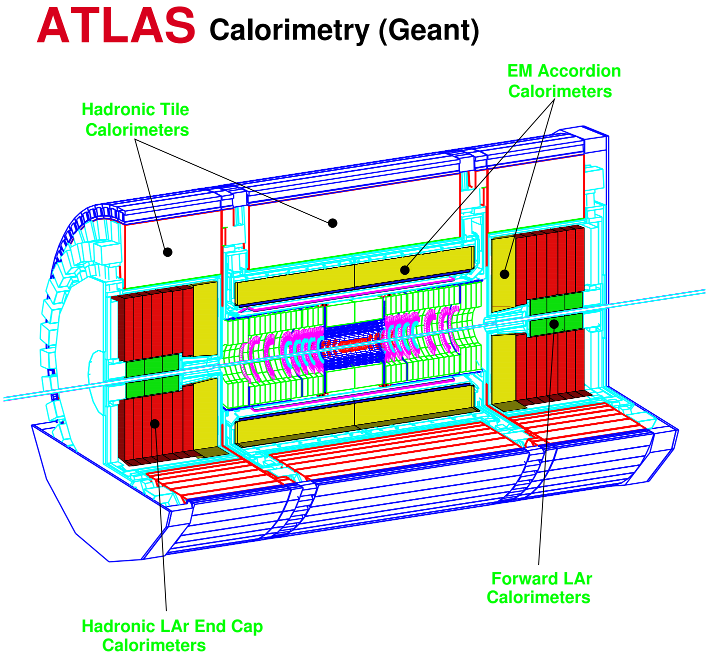

Figure 2.4: A schematic view of the ATLAS calorimetry.

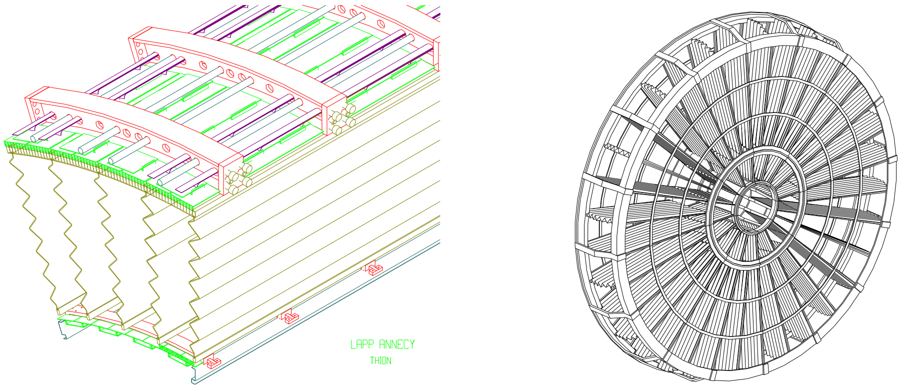

(a) (b)

Figure 2.5: A schematic view of the electromagnetic calorimeter geometry showing the "accordion" shaped absorbers in the barrel (a) and in the end-caps (b).

The dynamic range of the EM calorimeter is determined by the requirement to reconstruct both very low energy electrons ($\approx$ 2 GeV) from semi-leptonic b-decays, and also electrons with up to 5 TeV from Z' and W' decays. To meet this requirement 16-bit electronics is used for each cell. Resolution is required to be $\frac{\Delta E}{E} \leq \frac{10\%}{\sqrt{E}} \oplus 1\%$ in order to reconstruct the most challenging channels with adequate precision after a few years of running.

The segmentation (granularity) of the EM calorimeter is a balance between providing adequate precision, cost, and read-out issues. Except for the higher rapidity region of the end-cap three depth samplings are provided. Table 2.1 summarises the granularities of each of the depth samplings of the EM calorimeter. $\Delta\eta \times \Delta\phi$ is used to specify calorimeter granularity; it describes the angle subtended by a region of the calorimeter in azimuthal angle ϕ and pseudo-rapidity η. A pre-sampler is also located directly behind the cryostat wall; this has a granularity of $\Delta\eta \times \Delta\phi$=0.025×0.1, and covers up to $|\eta|$=1.8. This is used to improve the energy reconstruction by correcting for energy lost by particles showering before they reach the calorimeter.

η **range**	0 to 1.8	1.8 to 2.0	2.0 to 2.5	2.5 to 3.2
Presampler	0.025×0.1			
Sampling 1	0.003×0.1	0.004×0.1	0.006×0.1	0.1×0.1
Sampling 2	0.025×0.025	0.025×0.025	0.025×0.025	0.1×0.1
Sampling 3	0.05×0.025	0.05×0.025	0.05×0.025	

Table 2.1: Summary of granularity for EM calorimeter in $\Delta\eta \times \Delta\phi$.

2.3.2.2 Hadronic Calorimeter

The barrel hadronic calorimeter is also a sampling calorimeter but based on cheaper and less dense steel absorbers and plastic scintillator tiles. The tiles lie in planes perpendicular to the beam pipe and are read-out by optical fibres to photomultipliers from both ends. This approach improves the light yield and provides for some redundancy.

Physically the barrel hadronic calorimeter is split into a barrel, and two extended barrels in order to allow room for cabling and services. The barrel covers the region $|\eta| \leq 1.0$, and the extended barrels cover the regions $0.8 \leq |\eta| \leq 1.7$.

The main role of the hadronic calorimeter is to contribute to the reconstruction of jet energies, and provide a good E_T^{MISS} measurement. The energy resolution of these quantities is limited by intrinsic effects and not by the detector resolution, so resolution in the barrel is allowed to be considerably less than the EM calorimeter: $\frac{\Delta E}{E} \leq \frac{50\%}{\sqrt{E}} \oplus 3\%$.

In the end-caps where radiation levels are higher, liquid argon is again used as an ionising medium, but with copper absorber plates. The hadronic end-caps extend coverage up to $|\eta|=3.2$. The granularity in the range $|\eta| <2.5$ is $\Delta\eta \times \Delta\phi=0.1 \times 0.1$, and above $|\eta| =2.5$, $\Delta\eta \times \Delta\phi=0.2 \times 0.2$.

2.3.2.3 Integrated Forward Calorimeter

The integrated forward calorimeter provides both hadronic and electromagnetic calorimetry for the extreme forward regions (down to $|\eta|=4.9$). Its prime role is to provide suitable hermeticity required for a good measurement of missing E_T and total E_T.

The forward calorimeter relies on liquid argon sampling technology similar to the EM and end-cap hadronic calorimeters. However, the particle fluxes are high and a conventional LAr calorimeter would suffer from accumulating ions in the ionising region which would degrade performance. To prevent this the forward calorimeter uses a novel rod and tube structure which allows the LAr gap to be very small (250 μm compared with the more usual 1.94 mm used in the barrel calorimeter).

It is divided into three modules directly behind each other with respect to the interaction point. The first module provides EM calorimetry and is made from copper electrodes, embedded in a copper matrix which acts as an absorber. The second and third modules are the hadronic modules and use stainless steel electrodes in a tungsten-alloy absorber matrix. Tungsten is used since it has a relatively short interaction length which well contains the shower. This prevents leakage into the end-caps thus reducing pileup noise in this region.

The forward calorimeter does not have a constant granularity in eta and phi; but the segmentation in the electromagnetic layer is always equivalent to a granularity of better than $\Delta\eta \times \Delta\phi$=0.2×0.2.

2.3.3 Muon Spectrometer

In order to exploit the potentially robust signatures of high-momentum muons ATLAS will include a high-resolution spectrometer with a stand-alone trigger.

Figure 2.6 shows the muon spectrometer system. It consists of a large magnetic volume instrumented with precision muon chambers and muon trigger chambers.

The magnetic field for $|\eta|$ <1.0 is provided by a long barrel toroid. Two smaller end-cap toroids are inserted into each end of the barrel toroid to provide a magnetic field out to $|\eta|$=2.7. Each toroid is constructed of eight flat coils arranged octagonally around the beam axis. The open structure of the muon magnet system is designed to minimise the effects of multiple scattering on muon momentum measurement. Typically the magnetic field generated provides for bending powers of 3 Tm in the barrel and 6 Tm in the end-cap regions.

In the barrel, the muon chambers lie in three concentric cylindrical shells throughout the magnetic volume. In the end-caps the muon chambers are stacked vertically. Different systems of chambers are used for precision measurements and triggering, due to the differing requirements of each task. Four different technologies are used for the muon chambers:

- **Monitored Drift Tubes (MDTs):** Used by precision chambers over most of the rapidity range.

- **Cathode Strip Chambers (CSCs):** Used by the precision chambers at higher rapidities ($|\eta|$ >2) and close to the interaction point where a higher granularity is needed to cope with higher particle fluxes.

- **Resistive Plate Chambers (RPCs):** Used by the trigger chambers in the barrel

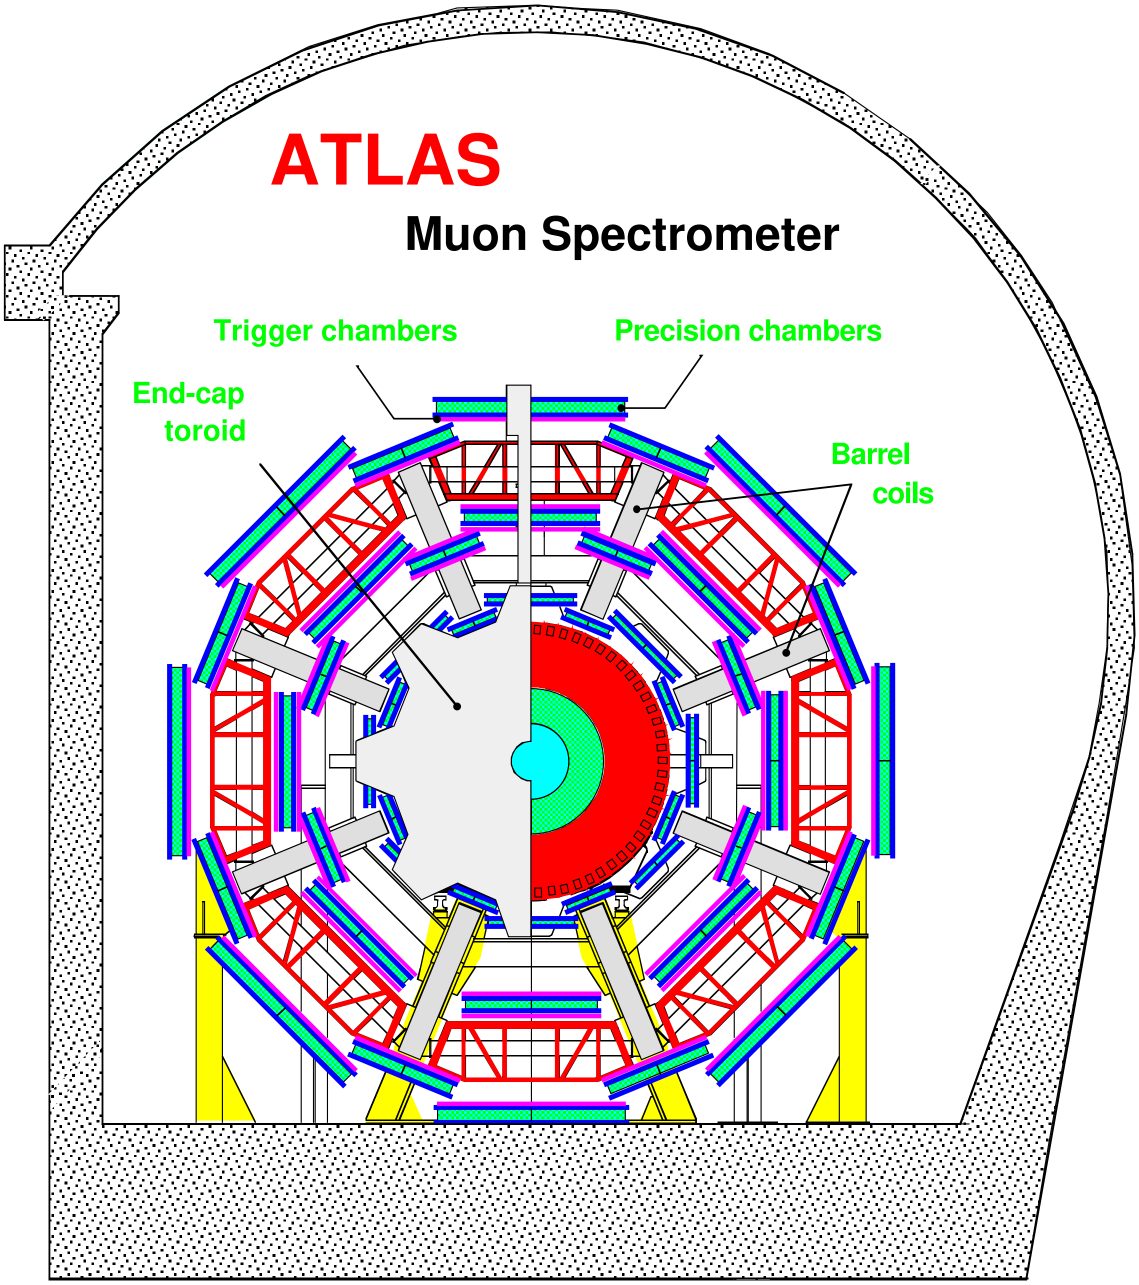

Figure 2.6: An end-on view of the ATLAS muon spectrometer system.

region.

- **Thin Gap Chambers (TGCs):** Used by the trigger chambers in the end-caps.

The precision chambers provide more accurate tracking but have read-out times much greater than the LHC bunch-crossing time. The trigger chambers provide less accurate information but are needed for triggering and to assign hits in the precision chambers to specific bunch-crossings.

The MDTs are made up from closely packed drift tubes while the CSCs consist of layers of multiwire proportional chambers; the CSCs provide faster read-out and higher granularity, but are prohibitively expensive so are only used in the high rapidity regions. The RPCs are gaseous detectors with parallel electrode plates, and the TGCs are based on multiwire chambers. For a detailed description of these technologies see [21].

2.3.4 The Trigger System

Triggering and data acquisition in ATLAS will be particularly challenging. The physics processes of most interest have very low cross-sections so the LHC must operate at very high interaction rates (around 10^9 inelastic interactions per second). Most of these events however will be soft-interactions and contain no new physics; the trigger must have sufficient rejection power to reduce this initial data rate to a level suitable for permanent storage (expected to be around 100 Hz) whilst maintaining a high efficiency for the rare new physics that is sought at ATLAS. It must do this for a detector with more than 10^7 channels.

The proton beams are divided into bunches that collide once every 25 ns (40 MHz); this is called the “bunch-crossing” rate. The trigger must accept data at the full bunch-crossing rate and form its decision with negligible dead-time. To do this the ATLAS triggering system is divided into three levels: Level-1 (LVL1), Level-2 (LVL2), and the Event Filter (EF). Figure 2.7 shows a functional view of the overall trigger system.

Each level is designed to be able to make an increasingly refined decision due to the reduced data rate from the preceding level. LVL1 must accept data at the full bunch-

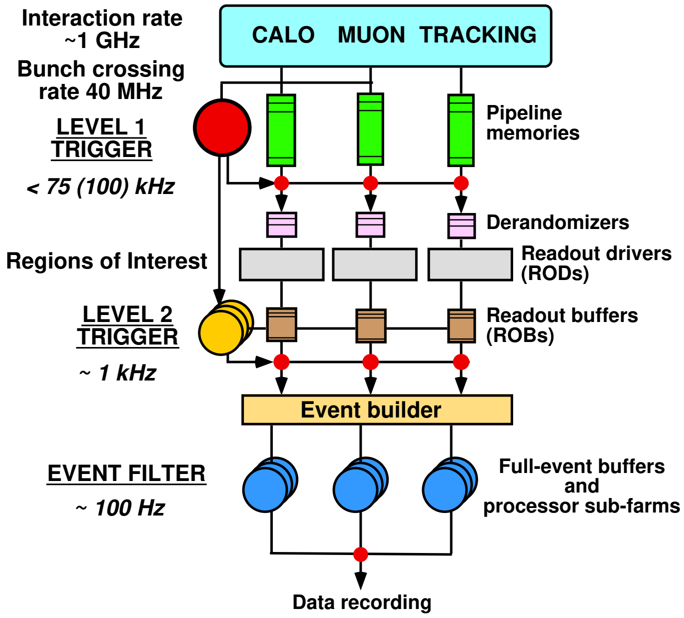

Figure 2.7: A block diagram of the ATLAS trigger and data acquisition system.

crossing rate and make an initial selection based on reduced granularity information from the calorimeters and the muon chambers. It bases its decision on fairly simple trigger objects that might be associated with potentially interesting physics processes: high P_T electrons and photons, muons, jets, isolated hadrons from τ decays, E_T^{MISS} and total E_T. A more detailed description of the LVL1 trigger is given in the next chapter.

During the time it takes LVL1 to make its decision (referred to as latency) data from all detector channels are stored in pipeline memories in the ATLAS front-end systems. These memories are usually contained in custom-built integrated circuits on or near the detector. For this reason the LVL1 latency must be kept relatively short so that these memories do not become too large; the current target is 2 μs (80 bunch-crossings), where most of this time is taken up by cable delays. On receiving a LVL1 Accept signal these data are transferred into read-out buffers (ROBs). The maximum trigger rate for LVL1 is determined by the maximum rate at which the front-end systems can accept triggers; initially this will be 75 kHz, but with possible upgrade to 100 kHz. De-randomising buffers are used to average out the high instantaneous data rate at the output of the pipeline memories to match the available input bandwidth of the read-out drivers (RODs).

Data are held in the ROBs while LVL2 makes its decision. The LVL2 latency will be variable with an average value of 10 ms. LVL2 refines the decision made by LVL1 by using full granularity data, and by using information from the inner detector. But it reduces the overall amount of data it has to process by concentrating its efforts on specific regions of the detector flagged by LVL1 as containing potentially interesting trigger objects; these regions are termed Regions of Interest (RoIs). It is envisaged that LVL1 will pass two kinds of RoI information onto LVL2: Primary RoI coordinates that are associated with the objects that actually caused the LVL1 trigger signal, and secondary RoIs associated with lower P_T objects.

The expected trigger rate from LVL2 is around 1 kHz although there is no architectural constraint on this number, and optimisation of the sharing between LVL2 and the EF is currently under study.

If LVL2 accepts an event, data from the RODs are integrated into a single event record (referred to as event building) and stored in an event buffer while the third level, the event filter (EF), makes its decision. The decision time available to the EF will be of the order of a few seconds. The EF selection is based on off-line algorithms, and reduces the rate to a level suitable for permanent storage. This is currently assumed to be $\approx$ 100 Hz for events of size $\approx$ 1 Mbyte.

Chapter 3

The Level-1 Calorimeter Trigger

3.1 Introduction

The requirements on the level-1 (LVL1) trigger can be summarised as follows:

- It must accept data from the calorimeters and muon chambers at the full bunch-crossing rate of 40 MHz.
- It must reduce this rate to a level suitable for LVL2 (around 75 kHz) whilst maintaining a high efficiency for interesting physics processes.
- It must pass Region of Interest (RoI) information on to LVL2.
- It must uniquely identify the bunch-crossing that contained the interaction of interest.

The LVL1 trigger decision can be thought of in two stages. In the first stage a set of processors work independently, and in parallel, to identify a series of basic inclusive trigger objects. These are: muons, isolated electrons and photons, isolated hadrons (mainly from tau decays), jets, missing E_T (E_T^{MISS}), and total E_T ($\sum E_T$).

Information from each sub-trigger about the energies and multiplicities of these trigger objects is then passed on to the Central Trigger Processor (CTP) which makes the overall LVL1 decision based on combinations of sub-trigger results.

Four of these sub-triggers require information from the calorimeters, so these sub-triggers are grouped together physically to form the Level-1 Calorimeter Trigger [22]. This chapter describes the architecture and the algorithms implemented in the calorimeter trigger. The remaining sub-trigger, the muon trigger, is described elsewhere [22].

3.2 Architectural Overview

Figure 3.1 shows the overall architecture of the LVL1 calorimeter trigger. Analogue sums of calorimeter signals are made on the detector to form "trigger towers" of reduced granularity, which are then transferred to the trigger system. Signals are converted to transverse energy (E_T) at the detector front-end to reduce the required dynamic range. The calorimeter trigger can be thought of as being comprised of three main components: the pre-processor, and two trigger processors. The pre-processor provides the interface to the analogue signals, and the trigger processors execute the various sub-trigger algorithms.

The design of the trigger processors was driven by two defining characteristics: the overlapping nature of the algorithms, and the requirement that the LVL1 trigger must operate synchronously at 40 MHz without introducing any significant dead-time.

All trigger algorithms divide the calorimeter into overlapping, sliding $\eta - \phi$ subspaces. These subspaces are required to overlap in order to provide uniform acceptance across the calorimeter; if they did not overlap a particle falling on the boundary between neighbouring subspaces would share its energy between processor windows and would not be recognised as a single object. This overlap requires that each trigger tower must be fanned-out to a number of processor destinations. Minimising the number of signals that need to be duplicated favours a more compact system with maximal processing density throughout.

The 25 ns interval between successive bunch-crossings is not sufficient to execute the complete trigger algorithms required to make a LVL1 decision. In order to prevent dead-time the LVL1 trigger uses pipelined processors. This means that the pre-processor digitises data at 40 MHz, and these data are clocked through the system also at 40 MHz. The trigger

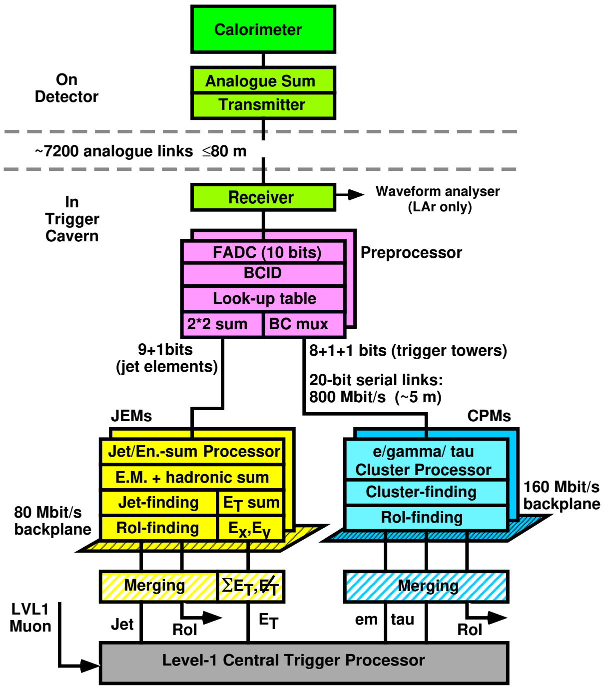

Figure 3.1: Overall architecture of the Level-1 calorimeter trigger.

algorithms are divided into a series of steps, each of 25 ns, so that data can pass through the system in a constant stream; this means that for every bunch-crossing clocked into the system, a trigger decision for some earlier bunch crossing is being clocked out. The time lag between a bunch-crossing entering the system and a decision being made is referred to as the "latency." The maximum allowable latency has been specified to be 2.5 μs; the target latency for the LVL1 trigger is 2 μs leaving a 500 ns contingency. However it should be noted that this latency is dominated by the time it takes to transfer signals approximately 60 m between the detector and the trigger hardware.

3.3 Pre-processor

The pre-processor receives one analogue signal for each electromagnetic and hadronic trigger tower from the calorimeters. In general each trigger tower covers an area of approximately

0.1×0.1 in $\Delta\eta \times \Delta\phi$ and is summed over all calorimeter depth samplings. For η >2.5 the tower size definition becomes more complicated and is dependent on the calorimeter geometry. Above η=3.2 signals are taken from the forward calorimeter for the E_T^{MISS} and $\sum E_T$ triggers. In total there are around 7200 trigger towers.

The pre-processor modules (PPMs) perform four main tasks: digitise calorimeter data; perform bunch-crossing identification (BCID); perform signal calibration, thresholding, and pre-summing; and transfer data to trigger processors.

The pre-processor modules (PPMs) digitise to 10-bit precision, over the range 0 to 256 GeV. This corresponds to a least significant bit equivalent to 0.25 GeV. A programmable delay is provided so that maximum resolution can be obtained by ensuring that each pulse is sampled with an appropriate phase. A constant off-set (pedestal) is added to the signal before digitisation to account for the fact that LAr signals have a negative lobe associated with each pulse.

The charge collection time of the calorimeters is longer than the 25 ns bunch-crossing interval of the LHC, so these signals are electronically shaped at the calorimeter front-end. Even after shaping, a LAr pulse spans around three or four 25 ns time-slices, and tile calorimeter pulses span around five or six time-slices. Figure 3.2a shows a typical pulse from the electromagnetic calorimeter taken during the 1996 test-beam period[a]. The x-axis represents time in 25 ns samples, and the y-axis represents the energy of the pulse in counts which are roughly equal to 1 GeV/count. The constant thirty count off-set is the pedestal.

If signals such as these were passed directly into the trigger processors, samples in the leading edge of the pulse could cause the trigger to fire; since the read-out electronics on the detectors cannot handle triggers in two consecutive time-slices this would result in the bunch-crossing containing the interaction of interest being missed. Bunch-crossing identification (BCID) is required to uniquely identify each calorimeter signal with the crossing that contained the particle that produced the pulse. It does this by transforming the pulse which

[a] The LVL1 test-beam programme is described in chapter 5.

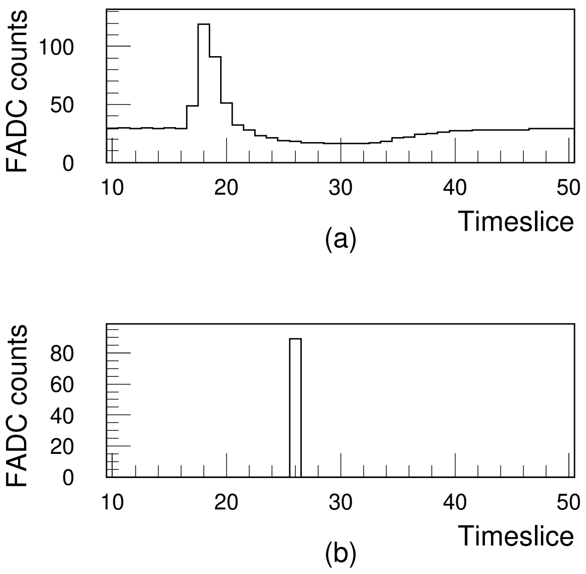

Figure 3.2: Example of an EM calorimeter pulse before (a) and after (b) bunch-crossing identification and pedestal subtraction. The x-axis represents time in 25 ns sample bins, and the y-axis represents pulse energy in ADC counts.

occupies a finite number of samples into an impulse, which occupies only one sample. Figure 3.2b illustrates the effect of BCID on the LAr pulse of 3.2a. Note that the BCID processed pulse has had the pedestal subtracted, and is delayed with respect to the actual peak of the LAr pulse; this delay represents the latency of the trigger although its exact value is a feature of the read-out method and does not reflect the true latency of the final system.

The main component of BCID is a **peak finder**: a sample (time-slice) is suppressed if it does not contain a higher energy than the samples before and after it. To improve the ability of the peak finder to resolve peaks in a noisy environment the signal is first transformed by a finite impulse response (FIR) filter; this filter takes consecutive time-slices ($y(n)$) and convolves them according to a set of coefficients (a_i):

$$y'(n) = \sum_{i=0}^{N-1} a_i y(n-i)$$

For optimum performance the coefficients should be chosen to match the shape of the incoming pulse. For a detailed description of BCID see [23].

After BCID the PPMs contain look-up tables which perform pedestal subtraction, map the data onto 8 bits (with a 1 GeV least count) for the trigger processors, and provide scope for additional calibration.

Two types of trigger processor receive data from the PPMs: the cluster processor (CP), implementing electron/photon and isolated hadron triggers; and the jet/energy-sum processor, implementing jet, E_T^{MISS}, and $\sum E_T$ triggers. The CP algorithms operate directly on trigger towers so these data can be sent directly. The jet/energy-sum algorithms are based on a step size of 0.2×0.2 in η and ϕ and so 9-bit sums of this granularity (called jet cells) are formed in the PPMs before transmission.

Due to the large number of signals it is not feasible to transfer all data bits in parallel. The baseline solution is to use Hewlett-Packard G-Link serialisers (HDMP 1012/1014) to serialise the data, although other serialiser protocols are currently under investigation. In the baseline solution each data-link operates at 800 Mbit/s allowing two 40 MHz channels of up to 10 bits each to be multiplexed onto a single link. For the jet/energy-sum system two 9-bit jet cells plus one parity bit for each are multiplexed onto a link; for the CP system two 8-bit trigger towers are sent on each link with one parity bit, and one flag bit for each.

The flag bit is used to allow a further saving in the number of channels going to the CP system by using a scheme referred to as "bunch-crossing multiplexing." The definition of the BCID algorithm means that two consecutive time-slices can never be non-zero; using this knowledge, the time-slice following a pulse in one channel can be used to transmit information from a neighbouring channel. This allows four trigger towers to be transmitted on each link, compared with two allowed with G-Link multiplexing alone. This scheme cannot be used for the jet/energy-sum processor however since BCID has already been performed independently on the four trigger towers that constitute the jet cell, and it is entirely possible in this case for consecutive samples in one jet cell to be non-zero.

3.4 Trigger Processors

The basic processing element for the trigger algorithms is a 4×4 window. For the CP system this is a 4×4 window of trigger towers, where EM towers and hadronic towers are considered separately. For the jet/energy-sum processor 4×4 processing elements are made up from jet cells (and hence have a physical size of 0.8×0.8 in η and ϕ) with the EM and hadronic contributions added. The trigger area is covered by an array of overlapping trigger windows, each neighbouring window being displaced by one trigger tower, or one jet cell, in either the η or ϕ directions.

Physically the trigger processors are made up from a number of modules, each module being concerned with a specific region (subspace) within the detector. The trigger space is partitioned according to a "$\phi-$quadrant" architecture, with each $\pi/2$ azimuthal quadrant being processed by a number of identical modules within a single electronics crate. Each module within the crate covers a small part of the quadrant in η, but spans the whole quadrant in ϕ. This is represented in figure 3.3. Hence in order to cover overlap between subspaces in the ϕ direction it is necessary to share signals between crates; whilst in order to cover overlap between subspaces in the η direction it is necessary to share signals between immediately adjacent modules within the same crate.

Sharing of signals between different crates is done by duplication of signals from the PPMs. Sharing of signals within a crate is done by a high speed, point-to-point, single ended backplane. In order to reduce the number of pins on each module the backplane operates in a semi-serial mode, multiplexing several bits from each channel onto a single link. In the cluster processor this backplane operates at 160 MHz, serialising 4 bits from each trigger tower onto each link. In the jet/energy-sum processor the lower granularity means there is a smaller volume of data to be transferred and the backplane operates at 80 MHz, serialising 2 bits from each jet tower onto each link.

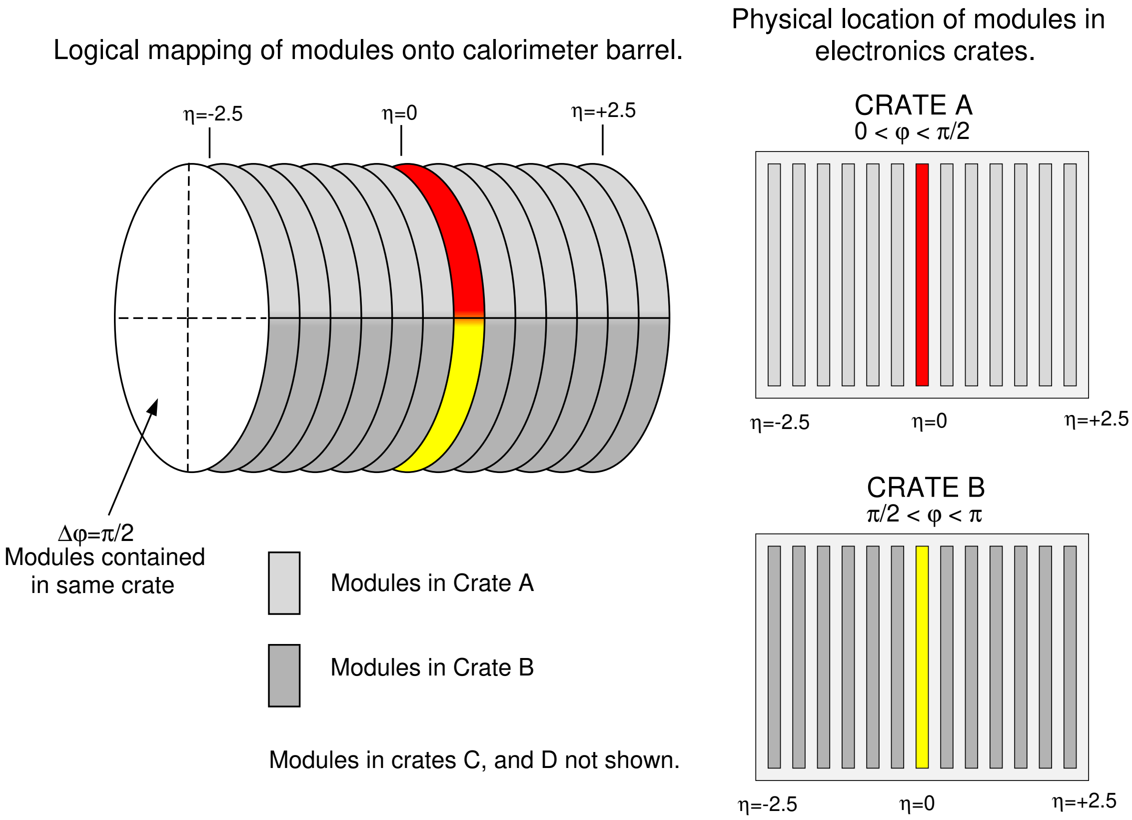

Figure 3.3: Partitioning of the calorimeter in a ϕ-quadrant architecture showing the relationship between logical subspaces processed by each module, and the physical layout of modules in a crate.

3.4.1 Cluster Processor

The cluster processor operates over the η range ± 2.5, and is implemented in 52 CP modules contained in four crates. Each crate contains an additional module which interfaces to the LHC timing and trigger control system, and two additional crates are used for merging CP data and for transferring region of interest and DAQ information.

The algorithms are implemented inside application specific integrated circuits (ASICs) to allow sufficient component density and speed. Since the algorithms for the electron/photon trigger and the isolated hadron trigger are sufficiently similar that they require identical input data, they can be implemented simultaneously within the same ASICs.

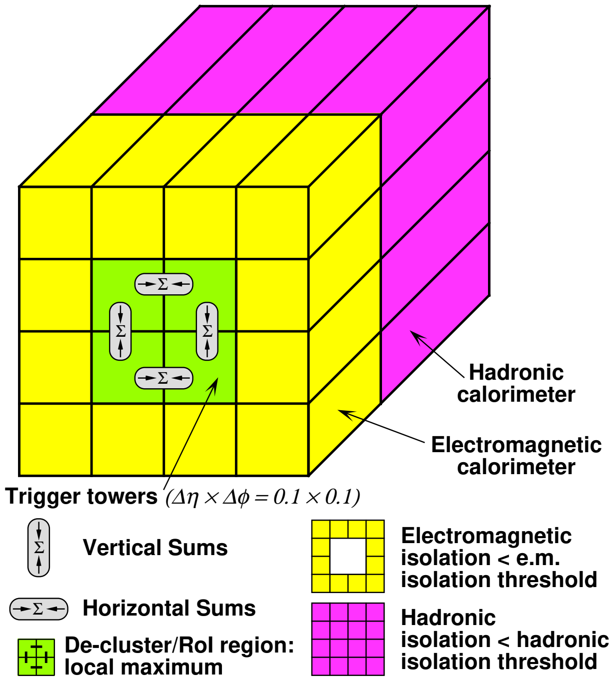

Figure 3.4: Electron/Photon trigger algorithm.

3.4.1.1 Electron/photon Trigger

The electron/photon (e/γ) trigger must recognise electron and photon candidates amongst a high jet background. Since QCD jets are the main background in the LHC the design of the e/γ trigger has been driven chiefly by the desire to provide a high rejection factor against jets, and at the same time maintain efficiency for electrons and photons. No distinction is made at LVL1 between electron and photon candidates.

Figure 3.4 shows the e/γ trigger algorithm. Its pattern recognition ability lies in distinguishing between small, isolated energy deposits in the EM calorimeter (electrons and photons), and larger objects spread across a number of trigger towers and occupying both the EM and hadronic calorimeters (jets).

The requirements for a trigger object to be found within the window are:

- The most energetic of the four horizontal (2×1) and vertical (1×2) sums (trigger clusters) must pass the EM cluster threshold.

- The total E_T in the electromagnetic isolation region must be below the EM isolation threshold.
- The total E_T in the hadronic isolation region must be below the hadronic isolation threshold.
- The de-cluster/RoI region must be a local-E_T maximum (see below).

It is necessary to consider the energy deposit in 1×2 or 2×1 trigger clusters in order to maintain efficiency for energy deposits that fall on the boundary between two trigger towers. This criterion alone would preferentially select electrons and photons over jets, but further jet rejection is achieved by vetoing trigger clusters that are *not* isolated; isolation regions are defined as a ring around the trigger clusters in the EM calorimeter, and the entire 4×4 window in the hadronic calorimeter.

These criteria are insufficient however because of the overlap between neighbouring trigger windows. A high E_T deposit in a single EM tower could for example pass cluster and isolation thresholds in four overlapping windows. This would clearly produce an incorrect multiplicity count, and make multiplicity based triggers impossible. To prevent this the additional condition is made that the central 2×2 cluster of cells in the EM calorimeter (defined as the de-cluster/RoI region) contains more energy than any of the other overlapping 2×2 clusters within the same trigger window. This stage is commonly referred to as "declustering." In reality this condition is relaxed a little to account for occasions where two neighbouring clusters have exactly the same energy: the actual local-E_T maximum condition is to require that a given RoI region has more energy than its neighbours along two adjoining edges, and at least as much energy as its neighbours along the opposite two edges.

3.4.1.2 Isolated Hadron Trigger

There are several possible roles for an isolated hadron trigger at LVL1:

- Principally it will be used to detect hadronic decays of τ leptons.
- In coincidence with a muon trigger it could be used to improve efficiency to $A \rightarrow \tau^+\tau^-$ events.
- In coincidence with a missing-E_T trigger it could be used to trigger on $W \rightarrow \tau\nu$ decays at low luminosity.
- It could be used to select high P_T hadrons to aid calibration of the hadron calorimeter.

It is not such a critical trigger as the e/γ and jet triggers, but its algorithm is sufficiently similar to the e/γ algorithm that it can be implemented in the same ASICs at very little additional cost.

An isolated hadron will leave a deposit in the calorimeter which appears very similar to an e/γ deposit, except that it will leave energy in both the electromagnetic and hadronic calorimeters. Hence, the algorithm for the isolated hadron trigger is identical to the e/γ trigger (see figure 3.4) except:

- The 2×1 and 1×2 trigger clusters in the EM calorimeter are all added to the 2×2 cluster directly behind in the hadronic calorimeter.
- The hadronic isolation region is now a 12 tower ring instead of the entire 4×4 region.
- The de-cluster/RoI region is a 2×2 sum in *both* the EM and hadronic calorimeters.

3.4.2 Jet/Energy-sum Processor

The jet/energy-sum system is contained within four crates, each containing sixteen jet/energy-sum modules (JEMs). In each crate fifteen JEMs perform the full jet algorithm and E_T summation for the range $|\eta|$ <3.2. The sixteenth JEM in each crate receives additional FCAL information and is used only to provide the means to include this information in the E_T^{MISS} and $\sum E_T$ calculations.

The reduced granularity of jet cells means that the data volume, and hence the data sharing, requirements on the jet trigger are much smaller. This means that the algorithms can be implemented inside field programmable gate arrays (FPGAs) providing a much easier and cheaper development cycle than the CP system.

3.4.2.1 Missing E_T and Sum E_T Triggers

E_T^{MISS} and $\sum E_T$ calculations are made over the full FCAL eta range, $|\eta| < 4.9$. The full jet cell granularity of $\Delta\eta \times \Delta\phi = 0.2 \times 0.2$ is used for $|\eta| < 3.2$, and FCAL data are used above this. Each JEM forms E_T, E_x, and E_y sums from the data it receives from the PPMs. This information is then passed to a sum merger module (SMM) in each crate where crate-wide E_x, E_y, and $\sum E_T$ sums are made. SMMs from three of the JEM crates pass their information to the SMM in the fourth crate, which performs final merging before passing its result onto the CTP. It is envisaged that the $\sum E_T$ trigger will be used independently in order to provide a trigger that is not biased by the event topology; the E_T^{MISS} trigger is expected to be used primarily in conjunction with other triggers, for example a combined jet and missing energy trigger.

3.4.2.2 Jet Trigger

Jet production is expected to be the dominant background at the LHC. For this reason the jet trigger is essentially background free, and so its task is not to distinguish between different kinds of object, but to provide optimal E_T resolution for a minimal rate.

Figure 3.5 illustrates the components of the jet trigger. Similarly to the e/γ algorithm the jet algorithm contains two main components: a jet cluster window used to measure E_T; and an RoI region used to prevent multiple counting of energy deposits.

The requirements for a trigger object to be found within the window are:

- The RoI region must be a local-E_T maximum as defined in the e/γ section.
- The jet cluster window must pass the jet cluster threshold.

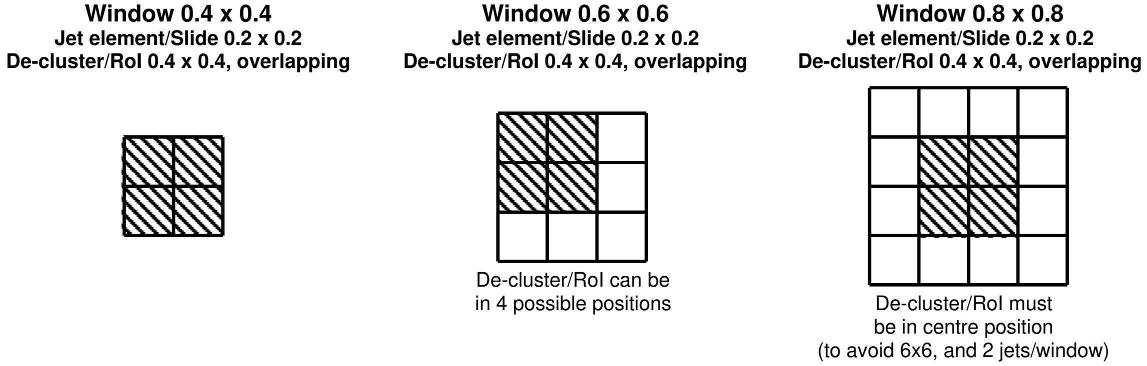

Figure 3.5: Jet trigger algorithms. The window size (clear) can be varied between 0.4×0.4, 0.6×0.6, and 0.8×0.8 in η and ϕ. The RoI region (shaded) size is 0.4×0.4; in the 0.6×0.6 algorithm it can occupy any corner of the jet cluster window.

The optimum size of jet cluster window will depend on the kind of trigger being used. For single jet triggers, at high E_T, containment of the jet E_T deposit is the main issue, and a larger window is preferred. In multi-jet triggers however, where jet E_Ts are small, the effects of noise and pileup are more significant; in this case summing over a smaller calorimeter area, and hence having a small jet cluster window, gives a better overall performance. To introduce a degree of flexibility into the system three window sizes are implemented in the same environment, and the choice of which one is used for any given trigger can be made during running.

As with the e/γ trigger, the overlapping nature of the algorithm means that an additional de-cluster phase is needed. As before, this is done by requiring that a 2×2 jet cell cluster (which in this case has a physical size of 0.4×0.4 in η and ϕ) is a local-E_T maximum. When the jet trigger is used with a jet cluster window of 0.4×0.4 there is obviously only one place the RoI region can go. For a jet cluster window of 0.6×0.6 there is some ambiguity about which window to associate with a given RoI region; this is solved by comparing the jet cluster window that has the *highest* E_T that contains the RoI region with the trigger threshold. Or to put this another way: the RoI region can occupy any corner in the jet cluster window. Finally, in the case of 0.8×0.8 windows it is also possible, in principle, to associate a number of jet cluster windows with the RoI region. However, in practice this approach is not feasible and so the only valid RoI region associated with a given jet cluster window is the central RoI region as shown in the diagram. This feature is discussed in more detail in chapter 4.

The simulation studies that led to this choice of jet trigger algorithm are described in detail in the next chapter.

Chapter 4

Simulation Studies of the Level-1 Jet Trigger for ATLAS

4.1 Introduction

The ability to trigger from high P_T jets is an essential part of the LVL1 trigger. Amongst other processes it will be necessary for many searches for supersymmetric particles, leptoquarks, and quark compositeness.

Jet production is expected to be the dominant hard process at the LHC, and so the main feature of merit for the LVL1 jet trigger is not its ability to distinguish between different types of object (as is the case in the e/γ trigger) but its ability to discriminate on the basis of E_T and multiplicity of jets. The tasks of the jet trigger can be summarised as follows:

- Identify hadronic jets in the barrel and end-cap calorimeters ($|\eta|$ <3.2).
- Classify jets according to E_T and provide an accurate multiplicity count of jets found in each E_T range.
- Provide accurate region of interest (RoI) information to the LVL2 trigger.

This chapter describes the simulation studies that led to the current proposal for the jet trigger algorithm [22], and gives estimates of its expected performance.

4.2 Algorithms

Jets are characterised by a dispersed, roughly conical, energy deposit in both the electromagnetic and hadronic calorimeters. The trigger architecture essentially forces the jet algorithm to be based on square, fixed-sized windows, sliding in $\eta - \phi$ space. The jet energy is reconstructed by summing the energy within a window, and the trigger decision made by comparing this energy with a set of programmable thresholds.

Two further choices about the nature of the jet algorithm were made prior to detailed simulation studies:

- Only algorithms based on overlapping trigger windows have been considered.
- Only algorithms using a "local-E_T maximum de-clustering" scheme have been considered.

The reasoning behind these choices is described below.

4.2.1 Non-overlapping Windows

In principle a jet trigger could be made by partitioning the calorimeter into a series of *non-overlapping* trigger windows, and comparing the energy in these windows with thresholds [24]. This kind of algorithm is considerably simpler to implement since data do not need to be fanned-out to more than one processor. However, a jet falling on the boundary between two windows will share its energy between them. This has several adverse effects:

- A jet above a given trigger threshold could be treated as two or more clusters of lower energy, none of which would pass that threshold. This reduces the P_T resolution of the trigger, lowering its overall performance (i.e. lowering its efficiency for a given trigger rate).
- A single high E_T jet could be reconstructed as two or more lower E_T clusters both above a given threshold, hence generating an incorrect multiplicity count at that threshold.

- Since a single jet can be reconstructed as two or more *adjacent* clusters additional processing would be required in order to provide a reliable RoI coordinate to LVL2.

For these reasons it was felt that algorithms based on non-overlapping windows could not meet the requirements of the LVL1 trigger, and so they are not considered any further.

4.2.2 De-clustering Algorithms

Using overlapping windows has an important effect on the trigger algorithm: a single deposit of energy can cause several overlapping windows to be above threshold. This would clearly generate an incorrect multiplicity count and greatly increase the multi-jet trigger rate. To prevent this an additional stage is required in the algorithm that resolves cases where two or more overlapping windows have detected the same energy deposit. This process is called "de-clustering"; two potential de-clustering schemes are described next but only the "local-E_T maximum de-clustering" scheme is simulated in detail.

4.2.2.1 The Bitmap Veto De-cluster

A simple form of de-clustering algorithm that was initially proposed as the baseline solution is the "bitmap veto" de-cluster: the total energy in a trigger window is compared with a set of thresholds. Bit patterns representing the number of thresholds passed by each window are then passed on to a separate de-clustering processor which vetoes a window if any of its neighbours along two edges are above the same threshold (figure 4.1).

This approach effectively removes multiple counting caused by overlapping windows and allows an RoI coordinate to be directly associated with unvetoed windows. However, it is not well suited to the high multiplicity environment at ATLAS since it contains several potential instabilities. Most crucially, there is an undesirable coupling between the choice of trigger thresholds and the ability of the algorithm to resolve nearby jets; this is because the bitmap veto bases its decision on which thresholds have been passed, rather than window energy. This feature is illustrated in figure 4.2. The left-hand box shows a set of jet window energies

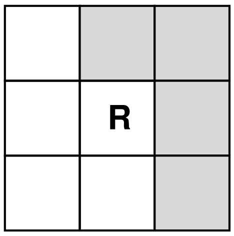

Figure 4.1: Diagram of bitmap veto style of de-clustering. A given trigger window, R, is vetoed if any of its neighbouring windows along two edges (shaded) contain a hit above the same threshold.

for an event that appears to contain two jets (i.e. two clear maxima can be observed at 100 GeV and 120 GeV). The middle box shows the results of applying a threshold of 100 GeV followed by a bitmap veto de-cluster; a "1" indicates a window that passed threshold and a shaded square indicates a window that passed the de-clustering conditions; clearly both jets are correctly de-clustered. In the far right box however the bitmap veto de-clustering algorithm is used with a 50 GeV threshold, and in this case it incorrectly identifies only the *less* energetic of the two jets.

A further factor is the ability of the trigger to provide an RoI coordinate to LVL2. There are two main problems in generating RoIs with a bitmap veto de-cluster:

- If a jet deposits its energy over a large area, the RoI associated with that jet will always drift to the top right hand corner. This instability in the location of the RoI means that LVL2 would have to process a much larger area of the calorimeter to be certain of locating the jet; hence increasing the processing load on LVL2.

- The location of the RoI coordinates associated with a given set of trigger objects will be different for each trigger threshold. For this reason additional logic would be needed to transform a series of RoI coordinates related to each threshold into a set of absolute RoI coordinates.

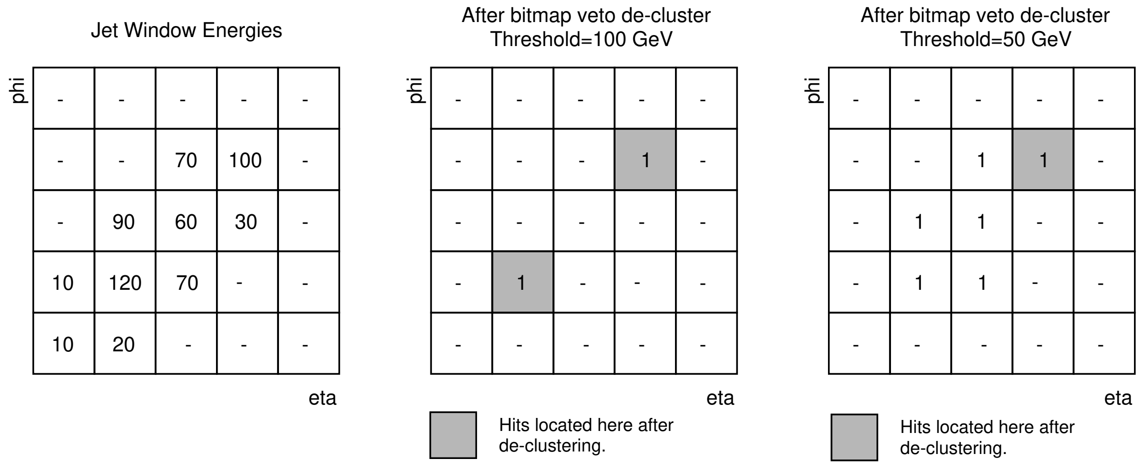

Figure 4.2: Diagram illustrating the main failings of the bitmap veto de-clustering algorithm. The left-hand box shows the energy deposits in a set of jet windows before de-clustering; the middle box shows the effect of applying a bitmap veto de-cluster based on a trigger threshold of 100 GeV, which correctly identifies both jets in the event; the far right box shows the effect of applying a bitmap veto de-cluster based on a trigger threshold of 50 GeV, which vetoes all except the *less* energetic jet in the event.

4.2.2.2 Local-E_T Maximum De-cluster

The preferred solution which constitutes the basic form for all algorithms simulated in this analysis is called the local-E_T maximum de-cluster.

The calorimeter area is covered by an array of "jet cluster windows", each displaced from each other by one "jet cell" in either the eta or phi directions. Within each jet cluster window another window is defined, referred to as the "RoI region". The method adopted is to require that any trigger window considered must also contain an RoI region that is a local-E_T maximum. Figure 4.3 illustrates this approach for the 0.8×0.8 window of the current jet trigger proposal (see chapter 3). Note that although this diagram is drawn for a specific algorithm the same basic components apply to all algorithms considered.

The conditions for a specific jet cluster window to generate a LVL1 trigger are:

- The total energy within the jet cluster window must be above a given threshold.

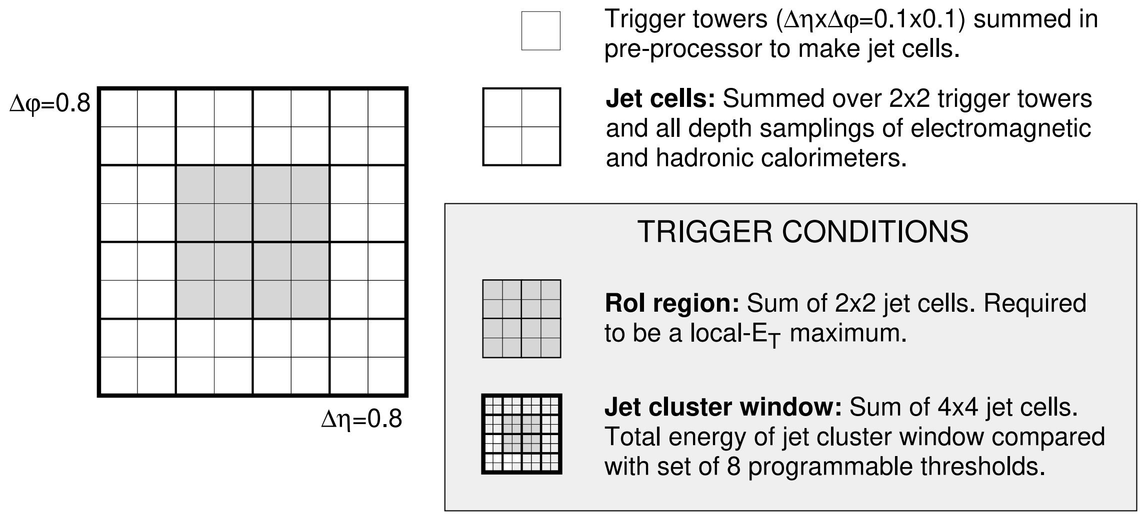

Figure 4.3: Schematic diagram representing components of a jet algorithm based on a local-E_T maximum de-cluster, using the 0.8×0.8 window of the current jet proposal as an example.

- The RoI region must be a local-E_T maximum with respect to the other possible, *overlapping*, RoI regions within the jet cluster window[(a)].

In the example given the jet cluster window is made up from 4×4 jet cells, each of which has a size of 0.2×0.2 in $\Delta\eta \times \Delta\phi$. The RoI region is made up from 2×2 jet cells. The definition of a local-E_T maximum is the same as that described in chapter 3 and is illustrated in figure 4.4.

This scheme is able to count accurately the number of jets in the event because, for a given jet energy deposit, there can only be one local maximum RoI region, and hence there can only be one trigger object. The RoI coordinate is then derived from the location of the RoI region associated with a given jet cluster window.

It should be noted here that in the example shown the RoI region only has one valid position (i.e. centred on the cluster window), so only a cluster window containing a local maximum RoI region at its centre can constitute a valid trigger. This is however a special

[(a)] Note that although RoI regions overlap in the example given, this is not always necessary and one algorithm considered later compares non-overlapping RoI regions for the local-E_T maximum condition.

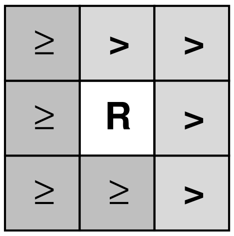

Figure 4.4: The local-E_T maximum test for an RoI candidate. The potential RoI (R) is compared with each of its overlapping neighbours; it must have greater energy than its neighbours along two edges (light grey), and at least as much energy as its neighbours along the remaining two edges (dark grey). This condition ensures that if two equal RoI regions are compared, one and only one of these will always be selected.

case and in general the RoI region is allowed to occupy any position within the jet cluster window. Naively it may appear that this means there can be more than one jet cluster window associated with a given RoI, and that the algorithm could again provide an incorrect multiplicity count. This is prevented by associating with a given RoI region only the *most energetic* jet cluster window that contains it, and using this in the trigger condition. This is discussed in more detail later on. Of course, valid algorithms can also be made where the RoI region is the same size as the jet cluster window in which case there is no ambiguity in its location.

All the jet algorithms considered here are based on this kind of algorithm, although simulation studies were needed to optimise the parameters of this algorithm in order to determine which combination of sizes for jet cluster window, step, and RoI region gave the best overall performance.

4.2.3 Performance Issues

The overall performance of the jet trigger is defined by its physics acceptance for a given rate. This is controlled by two factors: the sharpness of the efficiency threshold curve, and the rate produced for a given trigger threshold. The lower the rate produced for a given threshold, the lower the trigger threshold that can be set before the available rate budget is exceeded. In general, the sharper the efficiency curve the higher the acceptance for a given threshold. Threshold sharpness will be controlled by the degree of containment of the jet within the jet cluster window, the amount of noise and pileup summed, and the ability of the algorithm to resolve jets that lie close to each other in the calorimeter. Rate will depend chiefly on the area of calorimeter being summed, and hence on the size of the jet cluster window. There are three parameters which need to be optimised, (refer to figure 4.3 for their definitions):

Size of Jet Cluster Windows: This will determine how large an area of the calorimeter is summed for the trigger decision. Since jets are large compared to the size of trigger towers we might expect a larger window to more fully contain the jet, and have a sharper threshold behaviour. However, a larger window will also be more susceptible to "pileup" and detector noise leading to a higher rate for a given trigger threshold. At low luminosity noise caused by the physical detector components will be dominant; at high luminosity "pileup noise", caused by several interactions existing in each bunch-crossing, will become significant.

Size of Jet Cells: This will set the *step size* for the algorithm, (i.e. the distance between neighbouring windows), and determines how much overlap there is between neighbouring jet cluster windows and RoI regions. A small step size might improve the threshold sharpness of algorithms that have small jet cluster windows since it allows the window to be more precisely centred on the jet. A small step size will also improve the ability of the algorithm to resolve nearby jets.

Size of RoI region: This will directly affect the ability of the algorithm to resolve nearby jets, since the requirement that an RoI region must be a local-E_T maximum means that in cases where two jet cluster windows have neighbouring RoIs, only one window can contribute to the trigger decision.

In order to study the effect of these parameters on the performance of the jet trigger several algorithms have been simulated with a range of parameters. The intention is to study the merits of all possible algorithms without being biased by a particular architecture. However, since the ultimate aim is to build a working system there is little point simulating algorithms that could not possibly be built. The main limits on practical implementation of these algorithms are the number of input channels to the jet trigger (determined by the size of the jet cells) and the size of the processing environment required. Discarding impractical solutions, 7 algorithms have been simulated in detail. These are outlined in table 4.1.

Label	**Cell Size (Step)**	**Cluster Window**		**RoI Region Size**		**Comments**
	$\Delta\eta \times \Delta\phi$	Cells	$\Delta\eta \times \Delta\phi$	Cells	$\Delta\eta \times \Delta\phi$	
(a)	0.1×0.1	4×4	0.4×0.4	4×4	0.4×0.4	RoI region
(b)	0.2×0.2	3×3	0.6×0.6	3×3	0.6×0.6	is entire
(c)	0.4×0.4	2×2	0.8×0.8	2×2	0.8×0.8	jet cluster window.
(d)	0.2×0.2	2×2	0.4×0.4	2×2	0.4×0.4	
(e)	0.2×0.2	3×3	0.6×0.6	2×2	0.4×0.4	
(f)	0.2×0.2	4×4	0.8×0.8	2×2	0.4×0.4	Figure 4.3.
(g)	0.4×0.4	2×2	0.8×0.8	1×1	0.4×0.4	RoI region is single, non-overlapping, jet cell.

Table 4.1: Summary of jet algorithms studied.

Throughout this analysis the following shorthand will be used to describe which algorithm is being considered: Clus=A, Step=B, RoI=C will represent an algorithm with a jet cluster window size of A×A, a step size of B×B, and an RoI region size of C×C, in $\Delta\eta \times \Delta\phi$, where $\Delta\eta \times \Delta\phi$ describes the angle subtended by a region of the calorimeter in eta and phi directions.

Note that algorithm f (Clus=0.8, Step=0.2, RoI=0.4) is the algorithm drawn in figure 4.3. The original solution proposed in the ATLAS Technical Proposal [18] had the window and step sizes of algorithm c, (Clus=0.8, Step=0.4, RoI=0.8), but used the bitmap veto de-cluster.

4.3 Physics Requirements

In general the LVL1 trigger conditions are designed to meet the physics requirements with as simple and inclusive a set of trigger items as possible. Three kinds of jet trigger are currently foreseen:

- Single jet, (inclusive jet trigger).
- Multi-jet trigger (2, 3, and 4 jets).
- Jet trigger plus missing E_T.

Broadly speaking maximising physics acceptance requires the trigger to be efficient for the lowest P_T jets possible, within the limits of available rate budget. Some of the main new physics processes that will rely heavily on the jet trigger are:

- **Supersymmetry:** Single and multi-jet triggers will be used for initial inclusive searches for supersymmetry. In the case of R_P conserving SUSY [b] jets plus missing E_T can be used; in the R_P violating case single and multi-jet triggers can be used for inclusive studies, although most searches in R_P violating scenarios rely on leptons in the final state.
- **Leptoquarks:** Searches for leptoquarks decaying into jets and neutrinos, as predicted in some GUT models, will rely on a jet and missing E_T trigger.
- **Quark compositeness:** Triggering from jets with very high P_T will be required to test for quark compositeness.

[b] For a detailed description of SUSY, and R-Parity (R_P) see [41], or chapter 6.

- **Top Physics:** Although lepton triggers are currently expected to provide the best way to make measurements on the top quark, in many cases a multi-jet trigger may offer an alternative approach. This might be particularly useful in some non-Standard Model decays, for example in the case of a charged Higgs decay, $t \to bH^+$.

- **Higgs Physics:** Many searches for both Standard and non-Standard Model Higgs bosons will probably include a jet trigger; for example, the MSSM channel $hh \to b\bar{b}b\bar{b}$.

Based on physics processes such as these a rough idea of the jet P_Ts the trigger must be efficient for can be obtained. Illustrative jet P_Ts around these energies are summarised in table 4.2. Throughout the comparison studies in section 4.5 the jet trigger performance is often measured for jets of these energies in order to test the trigger performance in a region roughly representative of the P_T region the trigger will ultimately be required to operate in.

	Low Luminosity $10^{33}cm^{-2}s^{-1}$	**High Luminosity** $10^{34}cm^{-2}s^{-1}$
1 jet	100 GeV	200 GeV
3 jet	50 GeV	100 GeV
4 jet	50 GeV	100 GeV

Table 4.2: Illustrative jet P_Ts used in the initial testing and comparison of the jet trigger algorithms. In measurements of explicit jet trigger performance it was the efficiency of the trigger to these jet P_Ts that was measured.

In reality however, it is the actual trigger performance that is responsible for determining what jet P_Ts the trigger can be efficient for without exceeding the constraints of a given rate budget. A balance has to be struck between the physics requirements and the practical rate constraints of the trigger hardware.

Section 4.6 describes the calculation of the expected overall performance of the current jet trigger proposal from which the lowest jet P_Ts that could feasibly be triggered from are derived. It is these figures that have been used in the specification of the "trigger menus" given in reference [25].

4.4 Simulation Tools

To evaluate the performance of the trigger algorithms Monte Carlo simulations were made. The simulation chain contains two main components: Event generation, i.e. the simulation of physics processes at the 4-vector level; and detector response.

Physics events were generated with PYTHIA 5.7 [26] Monte Carlo for Standard Model processes and PYTHIA 6.1 [31] for Supersymmetric processes. PYTHIA generates events using a "parton shower" model; it is widely used within the ATLAS collaboration, and is considered to be the best understood and most versatile tool for describing generic hadron collisions.

Detector simulation can be performed in various levels of detail. The most detailed level is termed a "full simulation" and is carried out within the GEANT [30] package. This method uses a detailed model of the ATLAS detector [32] to track the passage of all stable particles through the sensitive regions of the detector. It accounts for detailed effects of material and detector geometry; for example, overlap regions between the various sub-detectors and the possibility of particles showering in sub-detectors or support structures within the experiment. GEANT simulations are considered the most accurate estimation of the ultimate response of ATLAS available at this time; they are however very slow, with a single di-jet event taking around 1000 seconds to simulate on an HP735 [33]. The analysis presented here contains several data sets each of around 20000 events; simulating this with GEANT would take several years, and is clearly unfeasible.

Fortunately the jet trigger simulations are not very sensitive to details of detector geometry and shower profile, and so a much simpler parameterised detector simulation can be used.

For the studies presented here a fast simulation was used that was written specifically for comparison of LVL1 calorimeter algorithms. It contains the following features:

- The calorimeter space was assumed to be a 64×64 grid of trigger cells representing trigger towers of $\Delta\eta \times \Delta\phi$=0.1×0.1 extending to rapidity $|\eta| <3.2$. The E_T contained in each cell was calculated from the 4-vector information in PYTHIA assuming a simple barrel shaped calorimeter.

- Each trigger cell represents summation of all calorimeter depth samplings, with the relative energy in electromagnetic and hadronic cells being calculated from a parameterised shower profile.

- Trigger cells were smeared by calorimeter resolution equal to $\frac{10\%}{\sqrt{E}}$ in the electromagnetic calorimeter and $\frac{50\%}{\sqrt{E}}$ in the hadronic calorimeter.

- The resolution of the trigger front-end electronics was simulated by adding 5% gain variations between channels, 400 MeV of electronic noise, and digitising to 8 bits with a calibration of 1 GeV/count.

- Electromagnetic and hadronic trigger cells were summed together, and groups of these sums used to form appropriately sized jet cells.

- At high luminosity ($10^{34}cm^{-2}s^{-1}$) each bunch-crossing will contain many interactions. Most of these interactions are very soft collisions however, and result only in a general low P_T hadronic background. This background is termed “pileup noise” and is very significant when attempting to trigger on low energy jets. It is simulated by superimposing 20 minimum bias events on top of every signal event.

- A full pulse history of up to 25 preceding bunch-crossings was recorded so that the true time evolution of the calorimeter signals could be estimated. This was necessary for correct simulation of BCID.

Since the jet trigger uses only relatively coarse calorimeter information, and sums over all calorimeter depth samplings, it is believed that this crude level of detector simulation is sufficient. To confirm that this is indeed the case, GEANT and fast simulation results are compared later in this section.

4.4.1 Data Sets

Events were simulated at both low ($10^{33}cm^{-2}s^{-1}$) and high ($10^{34}cm^{-2}s^{-1}$) luminosity.

For most studies, simple events containing only QCD interactions were used as a source of jets. To investigate the behaviour of the jet trigger with other kinds of topologies other physics process were simulated: top events, SUSY events, and $hh \rightarrow b\bar{b}b\bar{b}$ events.

Since the QCD jet cross-section falls exponentially with rising P_T, generating a suitable number of events with high P_T jets would require very large samples. To avoid this a lower limit on the 4-momentum exchange of the hard parton-parton collision can be artificially set in the PYTHIA datacards; by convention this parameter is often referred to as the *minimum hard scatter*, and labelled ${P_T}^{2-2}$ due to the way PYTHIA simulates this kind of process. This increases the number of events with high energy jets, but introduces a potential systematic error since the low P_T part of the jet spectrum is missing. To avoid this several QCD samples were generated with a range of minimum hard scatters; for any given measurement high statistics were gained by using as high a minimum hard scatter as possible, while systematics were avoided by ensuring that the minimum hard scatter was well below the lowest trigger threshold being considered.

The processes generated can be summarised as follows:

- **QCD:** 20000 events for each minimum hard scatter: 10 GeV, 15 GeV, 35 GeV, 70 GeV, and 140 GeV. QCD events generate mostly isolated jets, i.e. single jet events, or di-jet back-to-back events.

- **Top:** 20000 events top pair production, decaying freely and hence mostly hadronically. In general each event should contain six high energy, closely spaced jets.

- **SUSY:** 20000 events generated for both low and high mass SUSY. LHC points 3 and 4 were used [42], see chapter 6 for a detailed description. SUSY processes produce events containing many moderate energy jets, spread isotropically throughout the detector.

- **Higgs to b-quarks:** The channel $hh \rightarrow b\bar{b}b\bar{b}$ was generated specifically since it was of particular interest to some physics studies.

4.4.2 Jet Reconstruction

In order to determine the efficiency of the trigger algorithms some definition of which jets were actually present in the event is required. In PYTHIA, the transformation of initial coloured partons to jets of colourless hadrons is governed by an iterative fragmentation model. In this process final state particles are not associated uniquely with the initial partons and so PYTHIA's internal particle tables cannot be used in a jet definition. A jet definition could be made on the basis of the initial partons in the event, but this would not provide a realistic measure of trigger performance. Instead we are most interested in the performance of the trigger algorithms relative to the kinds of jet finding algorithms that will eventually be used in the off-line analysis. For this reason clustering algorithms are used to define which jets were contained in the event. Two kinds of jet definition have been considered in this analysis:

- **CONE:** This is the most widely used jet definition algorithm. It considers all particles in an event that are above some threshold as jet initiators. Each jet initiator is considered in E_T order and joined with all particles that fall within the same region of the detector (defined by a cone of tunable radius). If the candidate jet formed is above some minimum jet energy then it is retained, and all particles associated with it removed from further consideration. This process continues until no jet initiators remain [26].

- **$\mathbf{K_T}$CLUS**: This is a clustering algorithm similar to the type first introduced by the JADE collaboration. It uses an algorithm referred to as "binary joining" - all possible pairs of particles are considered, and combined if their relative transverse momentum is below some threshold. This process is repeated, treating newly formed clusters in the same way as particles, until no pairs of clusters remain that are below the combination threshold. [27].

Both algorithms operate directly on PYTHIA 4-vector information, before detector response and the effect of pileup has been calculated. The crucial difference between these algorithms is that K_TCLUS clusters according to the relative momentum of particles, whilst the CONE algorithm clusters according to the relative position of particles. Clearly the CONE algorithm more closely resembles the algorithms used for the jet trigger. For this reason it was feared that the CONE algorithm might bias results of a comparison study by falsely enhancing the predicted performance of algorithms with a similar window size as the CONE radius.

To avoid this problem the comparison analysis presented here used the K_TCLUS algorithm to define reference jets. K_TCLUS results have been cross-checked throughout with the CONE algorithm; in general the two definitions give very similar results, with the CONE algorithm giving a slightly more favourable prediction of performance since the trigger and CONE response are more similar. The section on overall performance uses the CONE algorithm to obtain absolute measurements of the overall rate since this is the definition most widely used in the ATLAS collaboration, and as such makes comparison of results between different institutes, and different studies, easier.

4.4.3 Analysis Method

In order to determine the performance of potential trigger algorithms the rate and efficiency must both be known as functions of trigger E_T threshold.

For each trigger algorithm, rate as a function of trigger threshold was calculated directly, independent of jet definition, from a sample of QCD events with suitably low minimum hard scatter.

Efficiency measurement depends on which kind of object we are trying to measure the efficiency for. Three different types of efficiency measurement have been used in this analysis:

- **Explicit efficiency** for the n^{th} jet at X GeV was measured by selecting a signal sample containing the n^{th} jet in an energy range X GeV $\rightarrow$ X+5 GeV, as defined by $\mathrm{K_T}$CLUS.
- **Sample efficiency** measured the overall efficiency for an entire sample of events of a particular process: eg, top or SUSY production. In this way the efficiency for a given sample as a function of trigger threshold was measured independently of the jet definition.
- **Physics efficiency** attempts to measure the efficiency for the kinds of events that will make it through the rest of the selection chain and be used in the final off-line analysis. Samples containing a particular process were generated, and then $\mathrm{K_T}$CLUS used to make very loose inclusive cuts dependent on the kind of process being considered - this is detailed later.

For a given type of measurement, efficiency as a function of trigger threshold can be measured and the overall performance of the algorithm (i.e. its efficiency for a given rate) found. In most studies the explicit efficiency was measured in order to study the jet threshold curve for jets of a specific P_T, and hence better understand the trigger algorithm performance in a range of conditions. Sample efficiency and physics efficiency measurements were useful mostly as cross-checks to ensure that the jet definition was not biasing the conclusions.

4.4.4 GEANT Cross-checks

To ensure that the fast simulation is an adequate estimation of the behaviour of the jet trigger, some comparisons were made with full GEANT simulation.

Figure 4.5 shows the E_T of the leading jet as measured by trigger algorithm (f) after detector simulation as a function of the leading jet E_T measured by $\mathrm{K_T}$CLUS based on PYTHIA 4-vectors. This is shown for both full GEANT simulation, and the parameterised "fast" simulation described earlier. They are both in good agreement, indicating that the full detector simulation does not have a large effect on the energy measured by the jet trigger algorithms. However, the jet trigger energy in the full simulation is systematically slightly lower than in the fast simulation; this is understood to be the effect of material absorbing jet energy before it reaches the calorimeters.

As a further cross-check that the fast simulation is giving reasonable results the threshold efficiency curves for three cluster window sizes were calculated for a single jet at 100 GeV. These are shown in figure 4.6 for both GEANT and fast simulation. The systematic lowering of the reconstructed energy for GEANT simulation can be seen as an off-set in the GEANT threshold curve of about 10 GeV compared with fast simulation. This effect will not affect the predicted algorithm performances however. The critical issue is the threshold sharpness for these curves. In both cases the threshold difference between 5% and 95% efficiency is around 30 GeV indicating that the fast simulation is correctly modelling the resolution of the trigger algorithms; crucially, the threshold softening exhibited by smaller jet cluster windows caused by poor containment of the jet (discussed later) is correctly modelled by the fast simulation.

The differences between GEANT and fast simulation are small, and can be neglected provided the jet P_Ts being considered are reasonably high. The only region where the fast simulation might be considered untrustworthy would be in a very low P_T study, for example when looking a RoI flagging efficiency.

These cross-checks add confidence that the fast simulation studies presented here represent a good estimate of both the relative performances of the proposed jet algorithms, and the absolute rates presented in the final section.

4.5 Comparison Studies

This section details the comparison studies that led to the current proposal for the jet trigger algorithm as described in chapter 3.

The performance of the trigger algorithms is investigated for isolated jets using a single

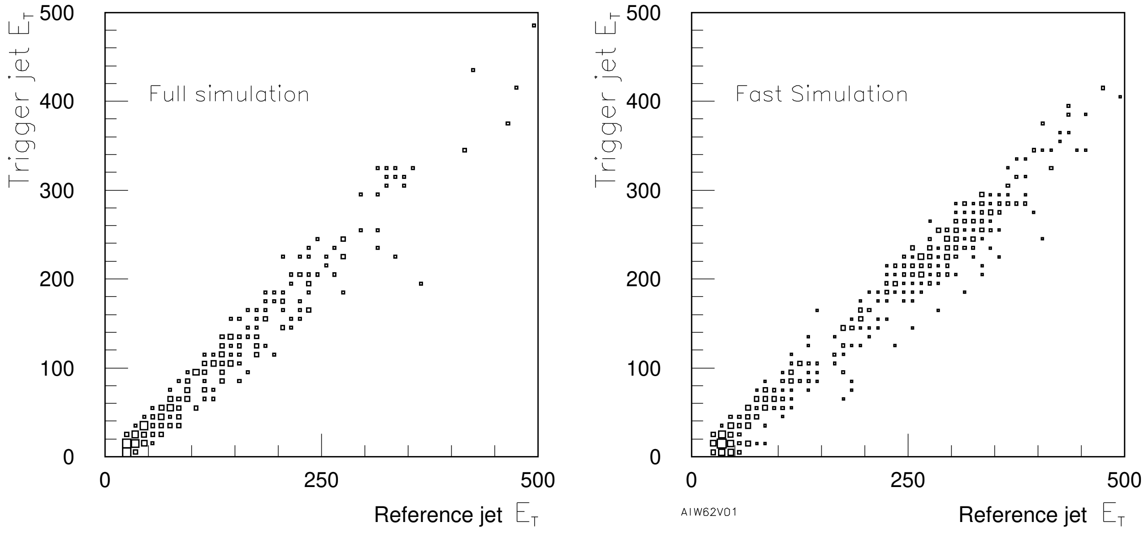

Figure 4.5: Trigger jet E_T as a function of reference jet E_T, for full GEANT simulation, and fast simulation. Note that different event samples were used for each plot.

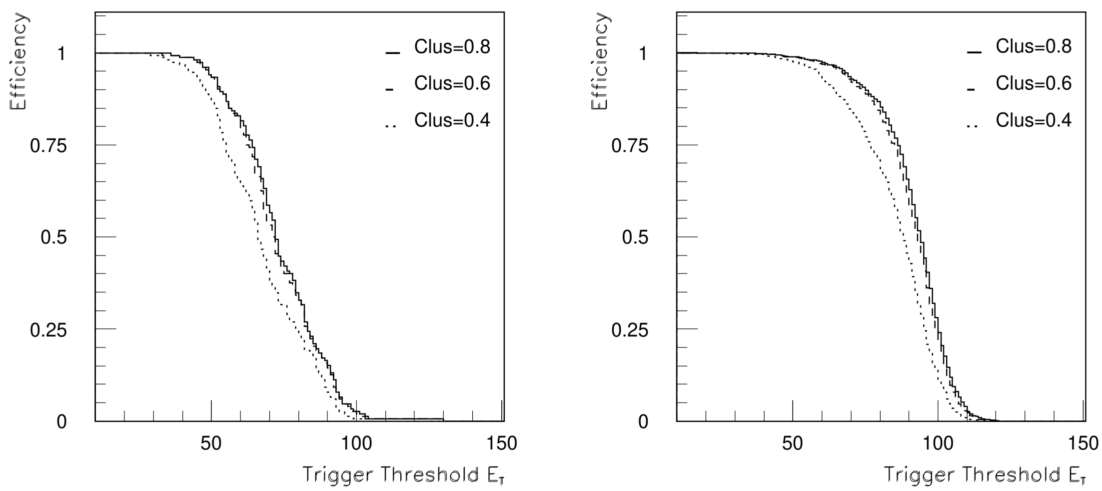

(a) Full Simulation (b) Fast Simulation

Figure 4.6: Explicit efficiency as a function of trigger threshold for a single 100 GeV jet at low luminosity for full GEANT simulation and fast simulation. A sample of QCD events were used.

jet inclusive trigger. The effect of closely spaced jets on performance is investigated in the context of a multi-jet trigger, and the low P_T region is investigated by looking at the algorithms' abilities to flag low P_T RoIs for LVL2.

4.5.1 Single jet Triggers

To investigate the performance of the algorithms for isolated jets, QCD events were simulated and the explicit efficiencies of the algorithms for a single jet at a certain energy were measured.

For isolated jets performance depends primarily on how well the jet energy is contained within the jet cluster window, and hence is controlled chiefly by the algorithm's step and jet cluster window size. Figure 4.7 shows the efficiency thresholds and rates for three different window sizes:

- $\Delta\eta \times \Delta\phi = 0.4 \times 0.4$
- $\Delta\eta \times \Delta\phi = 0.6 \times 0.6$
- $\Delta\eta \times \Delta\phi = 0.8 \times 0.8$

All three of these algorithms have a step size of $0.2{\times}0.2$, and an RoI region of $0.4{\times}0.4$. The efficiency curve shows the efficiency for a single 100 GeV jet at low luminosity. Both the $0.6{\times}0.6$ and the $0.8{\times}0.8$ sized windows appear to well contain the jet and have a similar threshold sharpness. The $0.4{\times}0.4$ windows however have a notably softened threshold curve. To reach a given efficiency (say 95%) the algorithm based on $0.4{\times}0.4$ windows would have to have a much lower threshold: 56 GeV, compared with 69 GeV for an algorithm based on $0.8{\times}0.8$ windows. For a given threshold however, the rate produced by $0.4{\times}0.4$ windows is considerably lower since it sums over a smaller calorimeter area.

Ultimately we are most interested in the trigger rate of each algorithm for a given physics acceptance. To see this the threshold and rate curves can be combined, and rate as a function of efficiency plotted (figure 4.8).

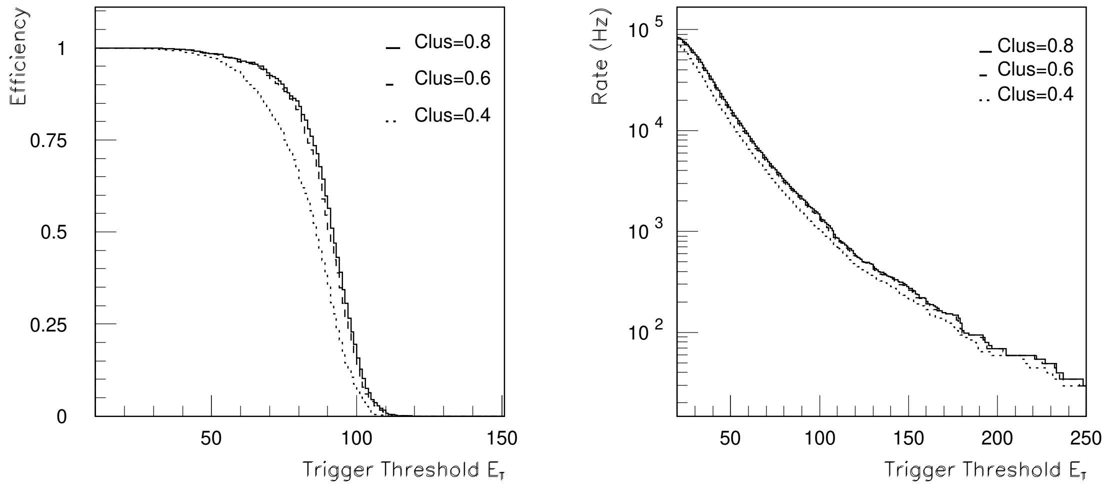

(a) Efficiency vs Trigger Threshold (b) Rate vs Trigger Threshold

Figure 4.7: Efficiency for a single isolated jet at 100 GeV, and rate vs trigger threshold for 3 different jet cluster window sizes at low luminosity.

Initially, for a low efficiency, the 0.4×0.4 windows have a lower rate. But the efficiency does not rise as quickly, so for higher efficiencies which we are most interested in there is an advantage to using larger jet cluster windows. Table 4.3 gives the trigger rates that would be produced in order to be 95% efficient to a single isolated jet at 100 GeV (200 GeV) at low (high) luminosity.

At high luminosity pileup noise is much higher, and the disadvantages of summing over a larger calorimeter area become more significant. This effect is shown in figure 4.9; here the requirement is made that the leading jet energy according to the K_TCLUS algorithm is less than 20 GeV in order to select events where the effects of pileup are significant. The difference in jet cluster window energy between 0.8×0.8 and 0.4×0.4 windows is shown for events at both low and high luminosity. At high luminosity the presence of pileup noise makes the energy difference somewhat higher, and the additional rate produced by 0.8×0.8 windows compared with 0.4×0.4 windows for a given threshold will be larger. This difference

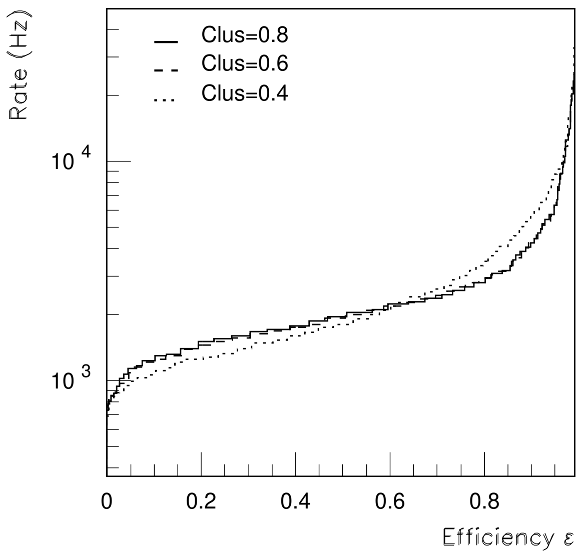

Figure 4.8: Rate vs efficiency for a single isolated jet at 100 GeV for 3 different window sizes at low luminosity.

	0.4 x 0.4	0.6 x 0.6	0.8 x 0.8
Low Luminosity	8.8 kHz	6.3 kHz	6.3 kHz
High Luminosity	3.8 kHz	3.0 kHz	3.0 kHz

Table 4.3: Trigger rates as a function of window size, for 95% efficiency for a single jet at 100 GeV for low luminosity, and 200 GeV for high luminosity.

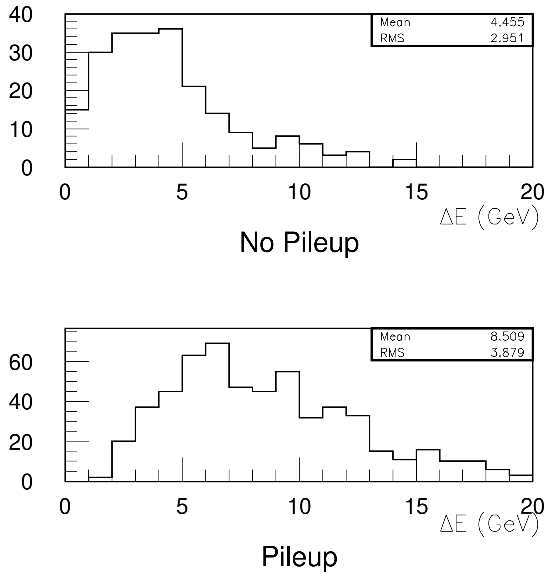

Figure 4.9: Energy difference ${E_T}^{Clus=0.8} - {E_T}^{Clus=0.4}$ distributions for high and low luminosity events containing mostly background noise: ${P_T}^{\mathrm{K_TCLUS}}$ <20 GeV.

is small compared with typical trigger jet energies and so the sharper threshold of the 0.8×0.8 windows means that there is still an overall performance advantage in using larger jet cluster windows (figure 4.10).

Note that there is apparently no performance advantage between Clus=0.6 and Clus=0.8 algorithms. It should be noted that this is not entirely an intrinsic effect related to the cluster size, but is influenced by the fixed RoI region of algorithm (f); this is discussed later.

The algorithm based on jet cluster windows of 0.4×0.4 has a softer efficiency threshold because jet energy is leaking out of the window. In principle this can be somewhat compensated for by allowing it to slide by a finer step; this allows the jet cluster window to be positioned more finely in the calorimeter, and so there is more chance that a window will exist that contains a sufficient amount of the jet energy to pass threshold. Table 4.4 shows the effect of varying the step size on algorithms based on both 0.8×0.8 and 0.4×0.4 jet cluster windows.

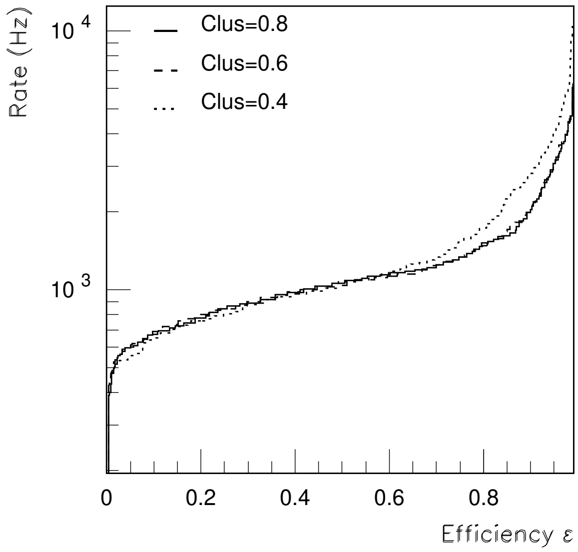

Figure 4.10: Rate vs efficiency for a single isolated jet at 200 GeV for 3 different window sizes at high luminosity.

Cluster Window	Step Size	Low luminosity Rate	High luminosity Rate
0.8 × 0.8	0.4 × 0.4	5.98 kHz	2.73 kHz
	0.2 × 0.2	6.33 kHz	2.99 kHz
0.4 × 0.4	0.2 × 0.2	8.76 kHz	3.84 kHz
	0.1 × 0.1	8.43 kHz	3.64 kHz

Table 4.4: Trigger rates as a function of step size, for 95% efficiency for a single jet at 100 GeV for low luminosity, and 200 GeV for high luminosity. All algorithms use an RoI region size of 0.4×0.4.

The intrinsic effect, related to varying the step size, is that a finer step size does slightly improve performance, but not enough to significantly compensate for a smaller window size. This is apparent when comparing the results from the Clus=0.4 algorithms.

For the Clus=0.8 algorithms this effect is still present, but it is masked by the effect of the fixed RoI region of algorithm (f).

4.5.1.1 Effect of a Fixed RoI Region in Algorithm (f)

Algorithm f (Clus=0.8, Step=0.2, RoI=0.4) differs from the rest of the algorithms considered in that its RoI region is fixed within the jet cluster window. Figure 4.11 illustrates this difference; for a given RoI candidate, algorithm g (Clus=0.8, Step=0.4, RoI=0.4) generates four possible jet windows, and takes the highest of these to be the jet energy. In (f) only the jet window centred at the same coordinate as the RoI region is generated and this is taken to be the jet energy. The remaining eight possible locations for the RoI region within the window are not considered.

This is purely an implementation issue. Allowing the RoI region to occupy any position in (f) is undesirable for two reasons: firstly, it would require a basic processing window of 6×6 jet cells, compared with 4×4 which is considered to be the maximum feasible environment before demands on signal fan-out become prohibitively complicated. Secondly, it would be possible for two local maximum RoI regions to exist within a single jet cluster window, causing the algorithm to give an incorrect multiplicity count. This constraint lowers the efficiency of (f) since it essentially requires the jet to have a specific shape. This effect is small however, and the large jet cluster window size and fine step size mean that (f) is still expected to have a better performance than most other algorithms.

This effect is mentioned here in order to flag (f) as a special case, so that the effect of a fixed RoI region is not confused with the intrinsic behaviour of the jet algorithms related to window, step, and RoI size.

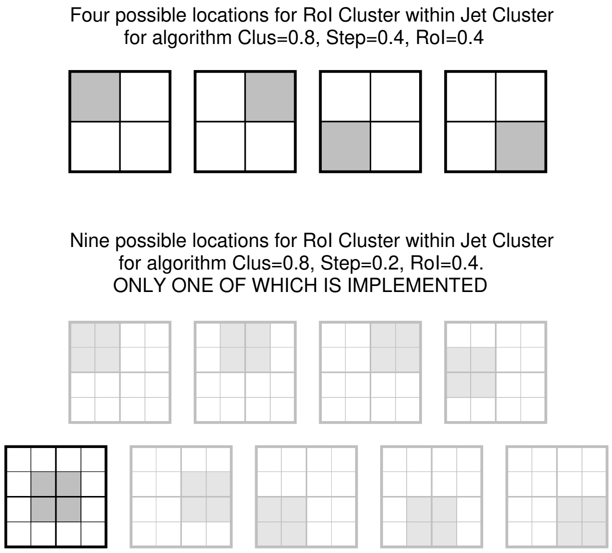

Figure 4.11: Diagram comparing algorithm g (Clus=0.8, Step=0.4, RoI=0.4) which allows the RoI region to occupy any cell in the jet cluster window, with algorithm f (Clus=0.8, Step=0.2, RoI=0.4) which has a fixed RoI region. The fixed RoI region will introduce an addition small inefficiency compared with an identical algorithm with a moving RoI region.

4.5.2 Double Counting of Jets

Before discussing the performance of the jet algorithms for multi-jet triggers, it should be noted here that simulation studies uncovered a flaw in a particular class of jet algorithm.

In cases where the RoI regions overlap by more than one step in each direction there is the possibility that a single narrow energy deposit can be counted more than once. This is because the de-clustering condition has a “greater-than-or-equal-to” condition on two sides, and only compares nearest neighbours. Figure 4.12 illustrates this problem: the left-hand box shows the energies in a set of jet cells of size $\Delta\eta \times \Delta\phi$=0.1×0.1, and the right-hand box shows the energies that would be measured by a set of overlapping 4×4 windows equivalent to the algorithm Clus=0.4, Step=0.1, RoI=0.4. Clearly this is the energy deposit of a single jet with an E_T≈50 GeV, but two separate hits have been recorded by the algorithm.

The significance of this effect can be seen by looking at the difference in the ϕ coordinate between the first and second jets in an event. Figure 4.13 shows this for two algorithms: (c)

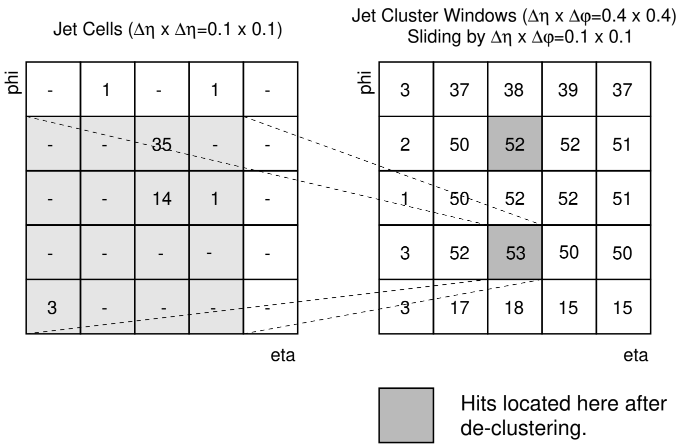

Figure 4.12: Diagram illustrating the flaw in algorithm a: Clus=0.4, Step=0.1, RoI=0.4, which causes it to double count some jets. The left-hand box shows the E_Ts in a set of jet cells of size $\Delta\eta \times \Delta\phi$=0.1×0.1, and the right-hand box shows the E_Ts that would be recorded by the jet cluster windows of algorithm (a). Because the RoI regions overlap by four steps in each direction the local-E_T maximum de-cluster is no longer sufficient to prevent the algorithm falsely generating two hits.

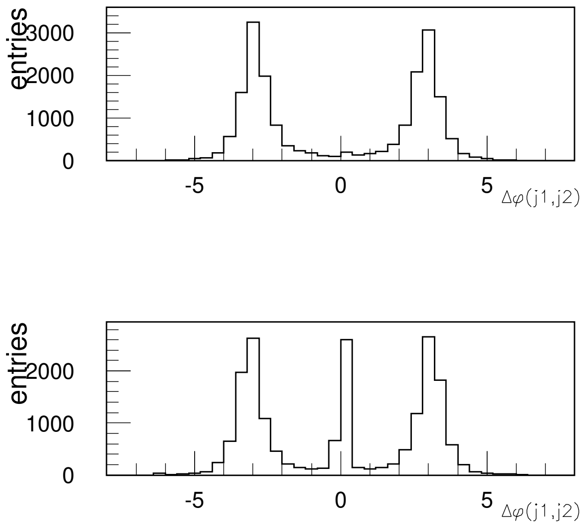

Figure 4.13: Difference in phi coordinates of first two jets found by algorithm c (top) and algorithm a (bottom) showing effect of double counting. Results shown for QCD events.

Clus=0.8, Step=0.4, RoI=0.4, and (a) Clus=0.4, Step=0.1, RoI=0.4.

In the first plot most jet pairs are back-to-back as would be expected for QCD events; for the second plot the large central peak indicates that the first jet is often counted twice. This double counting would be very dangerous, and would significantly increase the rate of multiple jet triggers. This problem affects two of the algorithms considered, and these have been discounted from further analysis:

- **a:** Clus=0.4, Step=0.1, RoI=0.4
- **b:** Clus=0.6, Step=0.2, RoI=0.6

4.5.3 Multi-jet Triggers

In multi-jet events it is possible that two jets will lie very close to each other in the detector. In this case the ability of an algorithm to resolve nearby jets gets folded into the threshold sharpness for that algorithm. The size of the RoI region and the step size are expected to control most critically an algorithm's resolving power. Figure 4.14 shows the efficiency of three algorithms based on a jet cluster window of 0.8×0.8 to identify a reference jet as a function of its separation ΔR from a more energetic neighbour, where $\Delta\mathrm{R}=\sqrt{\Delta\eta^2+\Delta\phi^2}$.

As expected, the larger step and RoI region sizes have more difficulty resolving closely spaced jets. Qualitatively, it appears that the size of the RoI region has the most significant effect in this case.

Figure 4.15 shows the effect of the step size and RoI region size on the threshold sharpness for 1 jet events (a), 3 jet events (b), and 4 jet events (c) at low luminosity. For 1 jet events a QCD sample was used. For 3 and 4 jet events top samples were generated in order to improve statistics.

To trigger on a single isolated jet (figure 4.15a) the RoI region size and the step size make little difference. Two of the threshold curves in this plot are identical; the apparent difference between Step=0.2, and Step=0.4 is caused by the fixed RoI region in algorithm f.

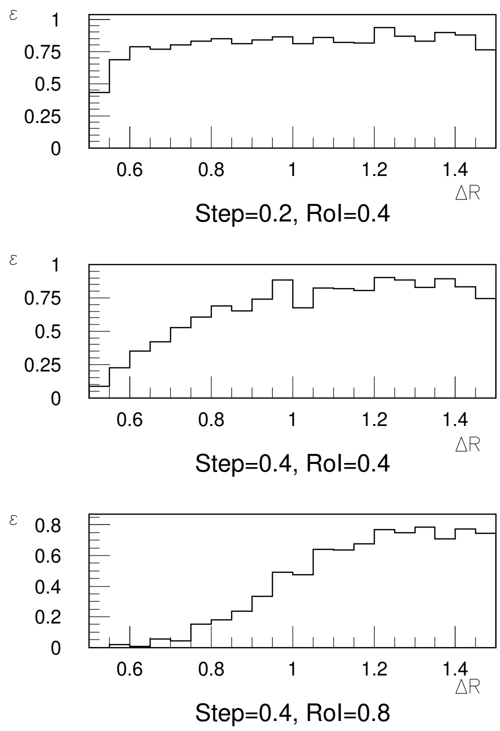

Figure 4.14: Efficiency to recognise two separate jets as a function of their separation, for three RoI/Step sizes. Shown for algorithm f (top), algorithm g (middle), and algorithm c (bottom). All algorithms have Clus=0.8.

For multi-jet events the algorithm based on RoI=0.8 becomes increasingly degraded with respect to the RoI=0.4 algorithms (figures 4.15b and 4.15c). The algorithm based on Step=0.4 shows a small loss of threshold sharpness relative to Step=0.2 for 4 jet events, but this is far less significant.

The overall effect on performance is shown in table 4.5. The mean statistical error is quoted to give an estimate of the significance of the relative rates. It should be noted that since the rates are calculated for each algorithm from the same events the efficiencies are correlated which makes small differences between algorithms more significant than the error would suggest.

In two cases the RoI=0.8 algorithm shows a clear performance penalty: 3 jet QCD events,

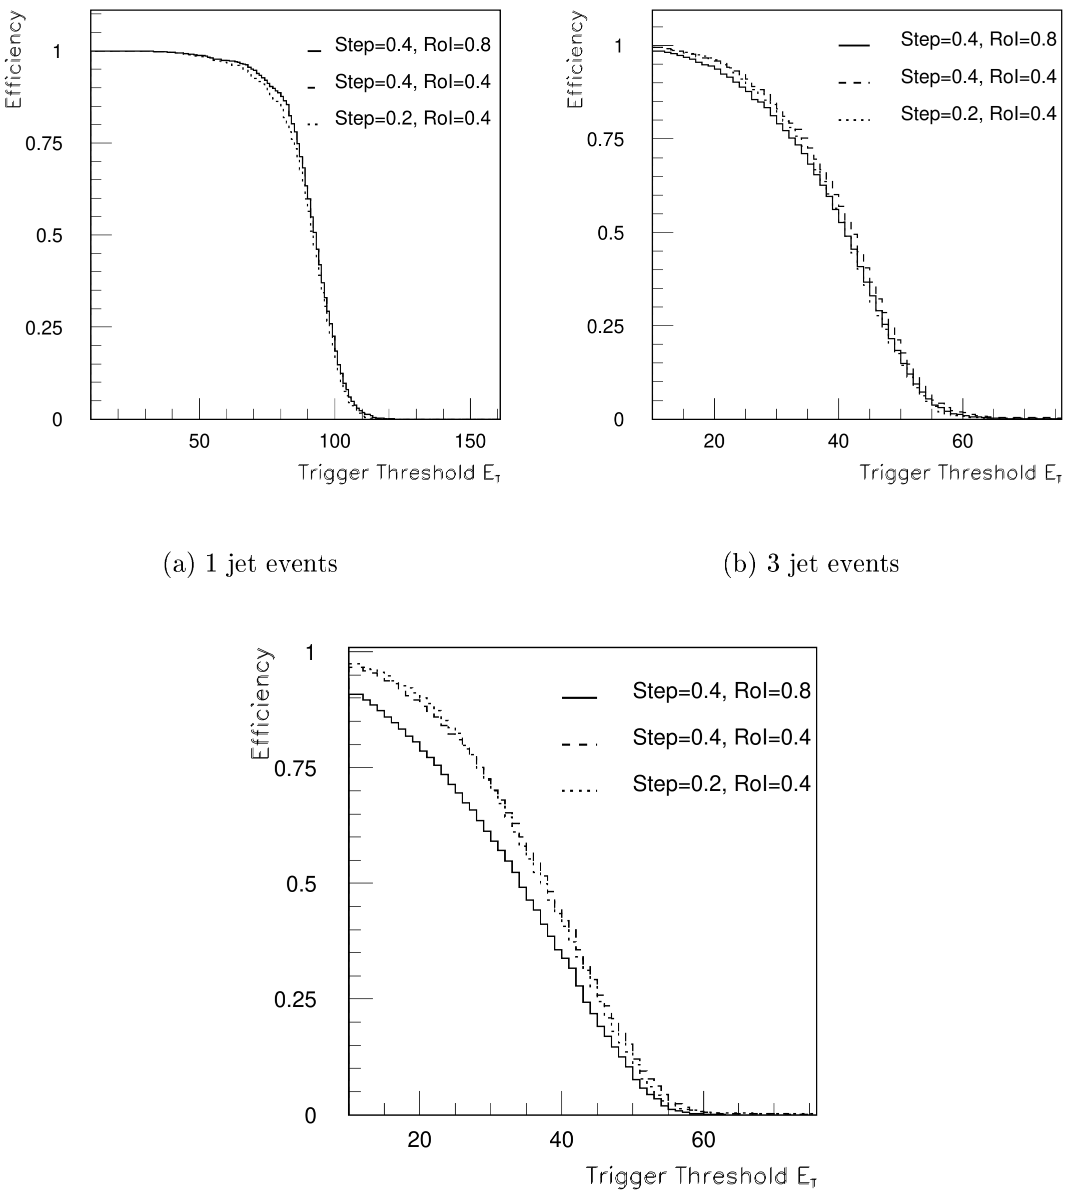

(a) 1 jet events

(b) 3 jet events

(c) 4 jet events

Figure 4.15: Efficiency vs trigger threshold for 1 jet at 100 GeV (a), 3 jets at 50 GeV (b), and 4 jets at 50 GeV (c) at low luminosity. Showing the threshold softening caused by poor resolving power algorithms based on larger RoI regions.

Step Size	**RoI Region Size**	**QCD**		**Top**	
$\Delta\eta \times \Delta\phi$	$\Delta\eta \times \Delta\phi$	1 jet	3 jet	3 jet	4 jet
0.4 × 0.4	0.8 × 0.8	68 GeV	20 GeV	23 GeV	11 GeV
		6.0 kHz	11.2 kHz	7.3 kHz	22.0 kHz
0.4 × 0.4	0.4 × 0.4	68 GeV	28 GeV	25 GeV	18 GeV
		6.0 kHz	5.1 kHz	7.3 kHz	4.7 kHz
0.2 × 0.2	0.4 × 0.4	66 GeV	26 GeV	24 GeV	20 GeV
		6.3 kHz	5.8 kHz	7.1 kHz	3.7 kHz
Statistical Error		±0.2 kHz	±0.7 kHz	±0.8 kHz	±0.6 kHz

Table 4.5: Thresholds and rates for 90% efficiency for 1, 3, and 4 jet events at low luminosity at 100 GeV, 50 GeV, and 50 GeV respectively, showing the effect of resolving power on performance for algorithms based on 0.8×0.8 jet clusters.

and 4 jet top events. The fact that this algorithm appears to perform adequately for 3 jet top events but not for 3 jet QCD events is a feature of the event topology. In top pair production, most events contain at least four jets (i.e. at least one W decays hadronically); in this case it is quite possible for the trigger to find three jets that are not necessarily close to one another. Note that no test is made in this analysis to ensure that the trigger algorithms are pointing at the *same* jets in the event. In QCD processes three jet events are often produced when one out-going quark radiates a gluon; in this case an efficient 3 jet trigger relies on resolving all three jets in the event, two of which are likely to be closely spaced.

In requiring four jets in top events the resolving power of the algorithm is more likely to have a significant effect; consequently the RoI=0.8 algorithm exhibits a dramatic loss in performance (a rate increase of around 17 kHz compared with RoI=0.4 algorithms). The overall performance effect in this case is illustrated in figure 4.16.

Since multi-jet triggers are expected to be a very significant part of the trigger requirements this is unacceptable and algorithm c (Clus=0.8, Step=0.4, RoI=0.8) is rejected.

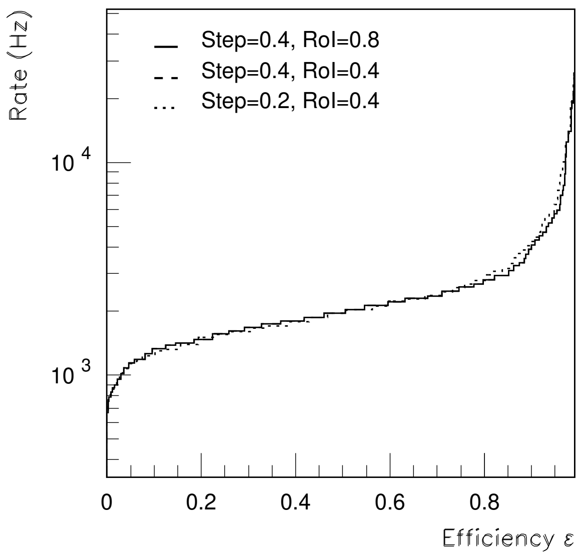

(a) 1 jet events

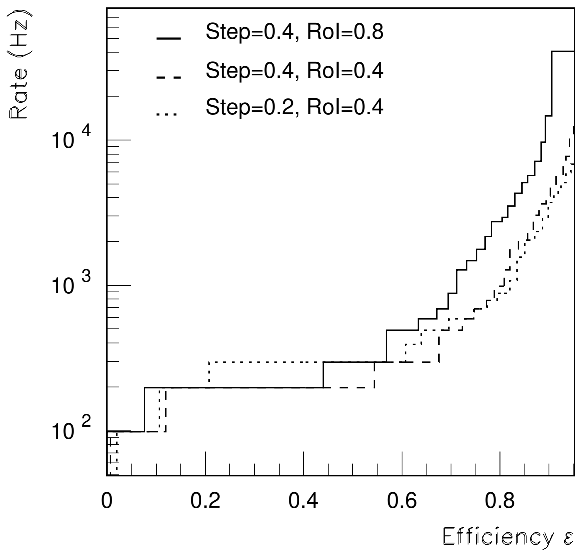

(b) 4 jet events

Figure 4.16: Rate vs efficiency for events containing 1 jet at 100 GeV (a), and 4 jets at 50 GeV (b) at low luminosity. Showing the effect of poor resolving power on algorithms based on larger RoI regions. Clus=0.8 algorithms compared in all cases.

A large step size also degrades performance in multi-jet events, but far less dramatically. In many cases algorithm Step=0.4, RoI=0.4 has a slightly better performance than Step=0.2, RoI=0.4 due to the fixed RoI region in Step=0.2, RoI=0.4. For 4 jet top events however this effect is reversed as the effects of closely spaced jets begin to become more significant. This is looked at in detail next.

4.5.4 The Programmable Jet Algorithm

The remaining algorithms that have not yet been ruled out are summarised in table 4.6. Two important notes should be made at this point:

- All remaining solutions are based on RoI regions of 0.4×0.4 . Larger RoI regions have been ruled out due to performance considerations. Smaller RoI regions are impractical due to implementation issues, and the requirement to contain a significant fraction of the jet E_T.

- Three of the remaining solutions are based on a step size of 0.2×0.2 and require the same ASIC environment for processing. This similarity means that all three of these algorithms can be implemented simultaneously with very little overhead in a *programmable algorithm*. See figure 3.5 in chapter 3.

Label	**Cluster Window** $\Delta\eta \times \Delta\phi$	**Step Size** $\Delta\eta \times \Delta\phi$	**RoI Region Size** $\Delta\eta \times \Delta\phi$
(d)	0.4 × 0.4	0.2 × 0.2	0.4 × 0.4
(e)	0.6 × 0.6	0.2 × 0.2	0.4 × 0.4
(f)	0.8 × 0.8	0.2 × 0.2	0.4 × 0.4
(g)	0.8 × 0.8	0.4 × 0.4	0.4 × 0.4

Table 4.6: Summary of jet algorithms still under consideration after initial simulation studies.

Since all three of the algorithms based on a step size of 0.2×0.2 can be implemented in the same hardware, the performance differences between them are less important. More critical is the performance cost of using an algorithm based on 0.4×0.4 steps; as can be seen in table 4.5 this solution is already starting to show a small loss of performance for 4 jet top events. However, the cost and complexity of this solution is far smaller, and if the performance loss were to be small then the cheaper option might be preferable. Due to the importance of these issues in making a decision about the jet trigger algorithm, and the marginal performance differences involved, this problem is looked at in detail in the next section.

4.5.5 Effect of a 0.4 × 0.4 Step Size on Multi-jet Performance

In principle algorithm g (Step=0.4) is expected to be less able to resolve closely spaced jets than the Step=0.2 algorithms. This is because by definition the closest two RoIs can be in (g) is ΔR >0.8, compared with ΔR >0.4 in the programmable algorithm.

Table 4.7 summarises how often this happens for a range of processes. This gives the percentage of events where the algorithm Clus=0.6, Step=0.2 successfully identifies a pair of jets that Clus=0.8, Step=0.4 counts as a single jet. This was calculated by comparing the reconstructed energy of the second most energetic jet for each algorithm; for a given object the Clus=0.8 algorithm will usually reconstruct a higher energy than Clus=0.6, hence if the reconstructed energy of the 2nd jet in Clus=0.8 is less than the reconstructed energy of the 2nd jet in Clus=0.6 this can be taken as an indication that Clus=0.8 has failed to resolve the two most energetic jets and so is in fact pointing at a third jet in the event.

Clearly for 4 jet events there is more opportunity for two jets to be merged. Some dependence on energy can also be seen particularly for top and SUSY production. This is because in order to produce a fourth jet at 100 GeV from the decay of a top or a SUSY particle the initial parent particle must be strongly Lorentz boosted causing the resulting jets to be closely separated.

Event Sample	QCD	Top	SUSY	Higgs $\rightarrow$ bbbb
3 jets @ 50 GeV	5%	3%	-	3%
3 jets @ 100 GeV	5%	7%	8%	6%
4 jets @ 50 GeV	9%	9%	4%	9%
4 jets @ 100 GeV	-	20%	13%	9%

Table 4.7: Percentage of events where Clus=0.8, Step=0.4 failed to resolve jets that Clus=0.6, Step=0.2 counted separately for low luminosity. A low mass SUSY sample (Pt 3) was used.

Figures 4.17 and 4.18 show the effect this has on threshold sharpness and overall performance for top events at low and high luminosity respectively.

The first thing to notice is that the performance advantage originally observed with inclusive triggers for Clus=0.8, Step=0.4 over Clus=0.8, Step=0.2 due to the fixed RoI region in the latter has been lost now that resolving power has become the dominant factor.

For high luminosity the performance of the Step=0.4 algorithm is an order of magnitude worse than the Step=0.2 solutions for 90% efficiency. This is caused by two independent effects:

- Attempting to trigger from higher energy jets, 100 GeV compared with 50 GeV for low luminosity, means that for the process considered the jets will be much more closely spaced, making the effect of a larger step more significant (figure 4.18a). This threshold softening is apparent even at low luminosity when attempting to trigger from very high P_T jets.
- Pileup noise is manifest in the form of a number of low energy jets in every event. As the trigger threshold is forced lower in order to maintain efficiency not only do the fluctuations in these pileup jets become more significant, but the pileup jets themselves begin to pass the trigger conditions.

On the basis of these results the Step=0.4 algorithm appears to be strongly disfavoured,

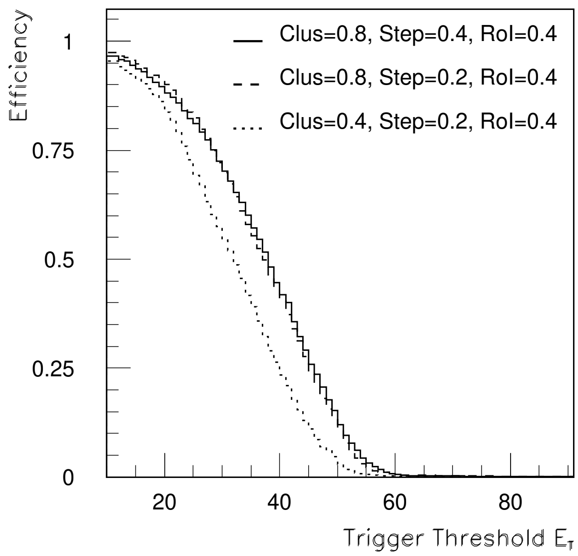

(a) Efficiency vs Threshold

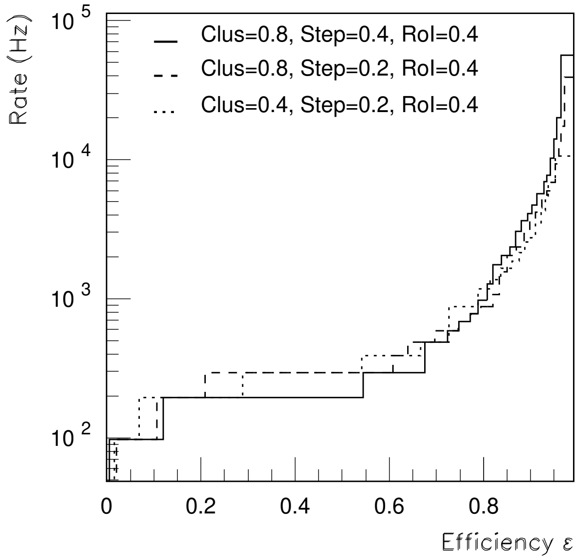

(b) Rate vs efficiency

Figure 4.17: Comparison of trigger performance between algorithms based on a 0.2×0.2 step and a 0.4×0.4 step. Results shown for 4 jet top events at low luminosity where the fourth jet is at 50 GeV.

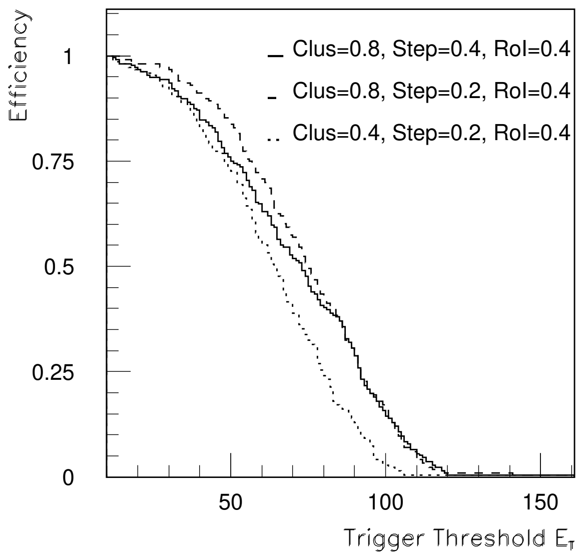

(a) Efficiency vs Threshold

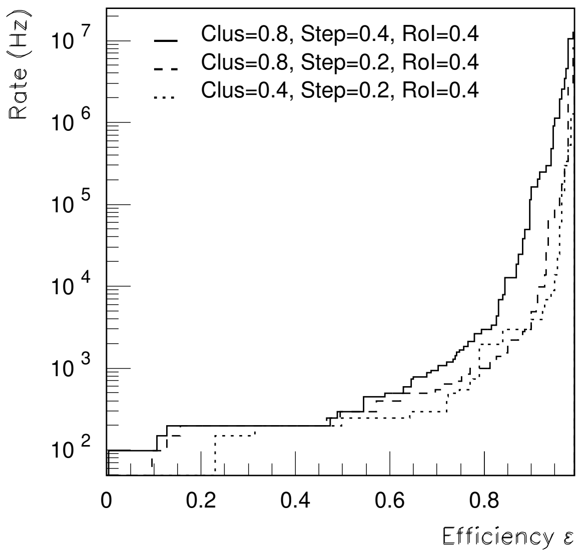

(b) Rate vs efficiency

Figure 4.18: Comparison of trigger performance between algorithms based on a 0.2×0.2 step and a 0.4×0.4 step. Results shown for 4 jet top events at high luminosity where the fourth jet is at 100 GeV.

but this is not a fair test. This example is useful to illustrate that a coarser step size can have a significant effect on the performance of the jet trigger, but trigger conditions are particularly extreme and none of the proposed algorithms are useful under these conditions. A fairer test is to look at the acceptance of the algorithms for a given rate. These results are summarised in tables 4.8, 4.9, and 4.10.

Here the efficiency of each potential trigger algorithm is given for a specific rate (1 kHz, 3 kHz, and 5 kHz). These rates are chosen to be around the maximum trigger rate reserved for jet triggers. Results are shown for a range of processes chosen to test multi-jet performance for a variety of event topologies.

Under these conditions algorithm g (Step=0.4) consistently shows a slightly poorer performance than algorithms based on Step=0.2. As has been mentioned this effect is most significant in high luminosity top events (table 4.8); here the combination of very closely spaced jets and low thresholds causes an efficiency loss of around 10% for all rates considered. Independent of topological factors, the effect due mostly to pileup noise can be seen for Higgs events in table 4.9. At high luminosity there is a slight additional loss of performance in the Step=0.4 algorithm of the order of a few percent compared with Step=0.2, but the most significant factor is the effect of the jet cluster window size. Under these conditions the containment advantage afforded by larger jet cluster windows is less significant, and the lower noise benefits of smaller windows become apparent; this can be seen as a 10% performance difference between Clus=0.4 and Clus=0.8 algorithms when requiring a rate of 5 kHz for the Higgs sample. These trends are mirrored in SUSY events (table 4.10) but less significantly, and they more difficult to compare since different SUSY samples are used for high and low luminosity.

All these effects are small and appear to be of the same magnitude as statistical errors. Since the same events were used for each algorithm the results are correlated which makes the differences more significant than the errors would naively suggest. Nevertheless the performance differences observed here are still not large from a physical point of view, but they do represent the symptoms of an algorithm being affected by its ability to resolve closely

Top sample, 4th jet @ 50 GeV, low luminosity				
	Cl=0.4, St=0.2	Cl=0.6, St=0.2	Cl=0.8, St=0.2	Cl=0.8, St=0.4
1 kHz	82%	81%	80%	79%
3 kHz	93%	91%	91%	90%
5 kHz	94%	94%	94%	93%

Top sample, 4th jet @ 100 GeV, high luminosity				
	Cl=0.4, St=0.2	Cl=0.6, St=0.2	Cl=0.8, St=0.2	Cl=0.8, St=0.4
1 kHz	81%	79%	80%	69%
3 kHz	90%	89%	88%	80%
5 kHz	93%	92%	90%	83%

Table 4.8: Explicit efficiency for a given rate using a four jet trigger on top events at high and low luminosity. Cl ≡ Cluster window Size, St ≡ Step size, RoI=0.4. Average statistical error on the efficiency is 3% (9%) at low (high) luminosity.

hh → bbbb, physics efficiency, low luminosity				
	Cl=0.4, St=0.2	Cl=0.6, St=0.2	Cl=0.8, St=0.2	Cl=0.8, St=0.4
1 kHz	75%	76%	77%	73%
3 kHz	86%	86%	87%	83%
5 kHz	90%	90%	91%	88%

hh → bbbb, physics efficiency, high luminosity				
	Cl=0.4, St=0.2	Cl=0.6, St=0.2	Cl=0.8, St=0.2	Cl=0.8, St=0.4
1 kHz	25%	24%	24%	18%
3 kHz	46%	39%	42%	35%
5 kHz	53%	48%	43%	39%

Table 4.9: Physics efficiency for a given rate using a four jet trigger for Higgs production at low and high luminosity. To measure physics efficiency an initial cut on the signal sample is made requiring $P_T(j_1)$ >100 GeV, and $P_T(j_4)$ >15 GeV. Cl ≡ Cluster window Size, St ≡ Step size, RoI ≡0.4. Average statistical error on the efficiency is 2% (1%) at low (high) luminosity.

Low mass (Pt3) SUSY, sample efficiency, low luminosity				
	Cl=0.4, St=0.2	Cl=0.6, St=0.2	Cl=0.8, St=0.2	Cl=0.8, St=0.4
1 kHz	65%	70%	71%	67%
3 kHz	81%	78%	79%	76%
5 kHz	85%	86%	85%	82%

High mass (Pt4) SUSY, sample efficiency, high luminosity				
	Cl=0.4, St=0.2	Cl=0.6, St=0.2	Cl=0.8, St=0.2	Cl=0.8, St=0.4
1 kHz	91%	91%	91%	87%
3 kHz	96%	95%	95%	92%
5 kHz	97%	96%	96%	93%

Table 4.10: Sample efficiency for a given rate using a four jet trigger for two SUSY points at high and low luminosity. Cl ≡ Cluster window Size, St ≡ Step size, RoI ≡ 0.4. Average statistical error on the efficiency is 1% for both high and low luminosity.

spaced jets.

It should be noted that these kind of events are particularly challenging for the jet trigger, and it is not clear at this stage whether the performance shortfall observed in the Step=0.4 algorithm will have any real impact on the ultimate physics acceptance of ATLAS. However, given the limits of this analysis and the exploratory nature of the LHC, it would be foolhardy to conclude that a step of 0.4×0.4 can perform as well as a step of 0.2×0.2 for all kinds of triggers that will ultimately be required in the final design.

4.5.6 RoI Flagging Performance

In addition to providing an RoI associated with each trigger object (primary RoIs) it is envisaged that the LVL1 trigger will also provide a set of secondary RoIs related to much lower E_T trigger objects. These will be used by LVL2 to refine the trigger decision.

For some processes it has been suggested that jets around 15 GeV (40 GeV) should be

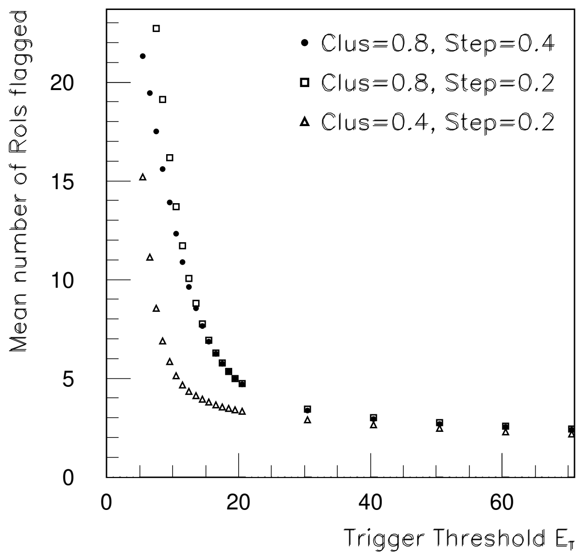

Figure 4.19: Mean number of RoIs flagged per event as a function of jet cluster window threshold at high luminosity.

flagged as secondary RoIs for low (high) luminosity [28]. This is a particularly challenging task for the jet trigger since thresholds this low will be strongly affected by electronic and pileup noise.

In order to study this a sample of QCD events was used. Events were selected that contained at least one trigger jet with E_T >100 GeV. Figure 4.19 shows the mean number of RoIs flagged as a function of secondary RoI threshold.

This diagram illustrates the main challenge; for much of the trigger threshold range the mean number of RoIs per event rises slowly with falling E_T. At some critical point however, around 20 GeV, noise/pileup fluctuations become comparable with the trigger threshold and the mean number of RoIs rises dramatically. Since the demands on LVL2 are strongly tied to the mean number of RoIs generated by LVL1, and a mean greater than about five is strongly undesirable, this region must be avoided.

As can be seen this critical point is not strongly related to the step size, but it is

Cluster Window	Step Size	Low luminosity P_T>15 GeV	High Luminosity P_T>40 GeV
0.8×0.8	0.4	4.8	3.8
0.8×0.8	0.2	5.5	4.2
0.6×0.6	0.2	4.8	3.9
0.4×0.4	0.2	3.7	3.7

Table 4.11: Mean number of secondary RoIs produced in order to be 95% efficient for a jet of 15 GeV (40 GeV) at low (high) luminosity.

related to the jet cluster window size; smaller windows sum less area of the detector, giving a lower average noise, and hence generating a lower average RoI multiplicity. Of course, this is balanced by the lower efficiency of smaller windows. The overall effect this has on performance is shown in table 4.11. This table shows the mean number of RoIs that would be generated per event for a secondary RoI threshold chosen to be 95% efficient for a jet at 15 GeV (40 GeV) at low (high) luminosity as defined by the $\mathrm{K_T}$CLUS algorithm.

In general Step=0.2 algorithms give a slightly higher mean RoI multiplicity since they are able to resolve RoIs that are closer together. However there is also an advantage in having smaller clusters, which compensates for the finer step size.

These results are another argument in favour of the programmable algorithm. At very low P_T the threshold sharpness is dominated by noise fluctuations and so a smaller window size is favoured; this is opposite to the situation at high energies with isolated jets where cluster containment has the largest effect on trigger resolution. The ability to adjust the jet cluster window size according to the triggering conditions provides a very desirable additional level of flexibility.

A word of warning should be noted here. Whilst these results are useful to investigate the relative performances of the trigger algorithms they should not be taken too literally. At these energies the simulation is expected to be sensitive to detector effects and a full

GEANT simulation would be needed to determine the absolute ability of the jet trigger to flag low energy RoIs. Some preliminary studies of this kind have been made elsewhere [29] which suggest that the task of flagging RoIs is somewhat more problematic than indicated here. This is an on-going study.

4.5.7 Conclusions

Of the original seven algorithms proposed as possible jet trigger solutions, three have been unambiguously excluded for the following reasons:

- **a: Clus=0.4, Step=0.1, RoI=0.4** Has RoI regions that overlap by several steps in each direction. This makes it possible for a single jet to be counted several times, causing the algorithm to provide an inaccurate multiplicity count.
- **b: Clus=0.6, Step=0.2, RoI=0.6** Excluded for the same reason as **a**.
- **c: Clus=0.8, Step=0.4, RoI=0.8** It was found that the large RoI regions of c gave it a poor resolving power for closely spaced jets. This had a dramatic impact on overall performance.

Of the remaining algorithms three (d, e, and f) are all based on a step of 0.2×0.2 and can be implemented simultaneously in a programmable algorithm. This solution however is considerably more complicated and expensive to implement than algorithm g (Clus=0.8, Step=0.4, RoI=0.4) and so simulation studies were used to determine the level of performance loss incurred in using the simpler Step=0.4 solution.

In multi-jet events a slight preference was seen towards both a finer step size and possibly a smaller window size, although this was only for events with particularly closely spaced jets that were not considered to be realistic triggering situations. When considering the ability of the jet trigger to flag low energy jets for secondary RoIs a slight preference was observed towards a smaller jet cluster window.

Although both these effects are small it was felt that the programmable algorithm was able to incorporate both of them, and so this has been chosen as the preferred solution. Despite simulation studies indicating that the performance loss of a Step=0.4 is small, it is a persistent loss of performance apparent in several scenarios. Given the exploratory nature of the LHC it has been a general philosophy of ATLAS to lean on the safe side, and select algorithms that will ultimately provide the best possible performance, and greatest flexibility in the final solution.

4.6 Overall Performance

This section looks at the expected overall performance for the programmable algorithm, using the CONE algorithm as the jet definition in order to provide a consistent definition between LVL1 and LVL2. Whilst the $\mathrm{K_T}$CLUS algorithm was considered necessary to avoid biasing results in a comparison study, the CONE algorithm is more widely used and so is a more convenient choice when defining the absolute performance of the jet trigger.

Figure 4.20 shows the trigger rate as a function of trigger threshold for single and four jet triggers at low luminosity. This measurement is independent of jet definition.

The rate budget for the LVL1 jet trigger is not yet clearly defined. Originally it was proposed to be of the order of a few kHz which would then be refined by LVL2. But recent studies in LVL2 suggest that LVL2 may be able to do little to reduce jet rate, and so the budget at LVL1 would have to be correspondingly lowered; for the purposes of the LVL1 trigger TDR [22] and the trigger performance document [29], a budget of 200 Hz was allowed for 1 jet, 3 jet, and 4 jet triggers separately.

Table 4.12 shows the trigger thresholds required to produce a suitable rate (200 Hz, or 2 kHz). For a given trigger threshold the corresponding jet P_T where the trigger would be 90% efficient is given. This is calculated in the same way as in the comparison studies except for two points:

- The CONE algorithm is used as the jet definition. The cone used has a radius of

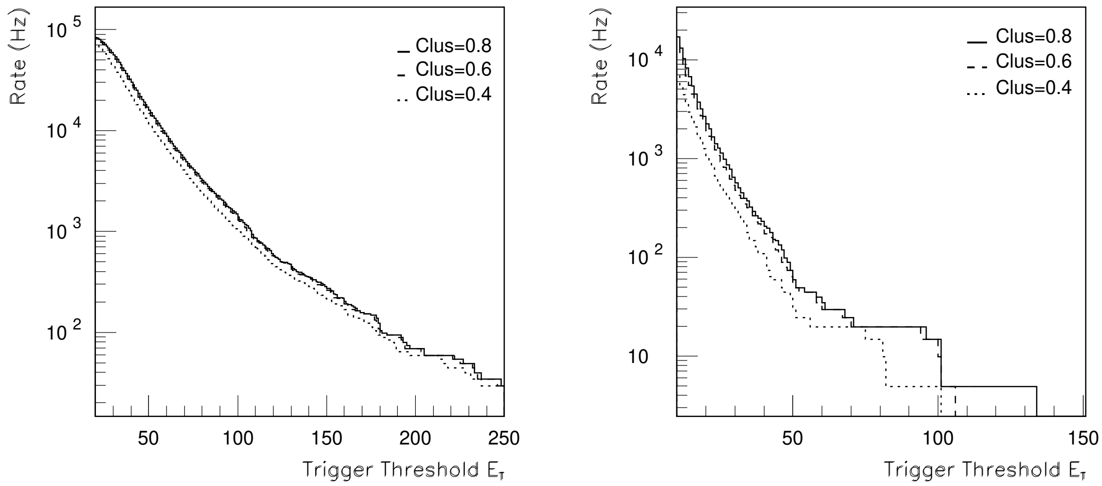

(a) 1 jet trigger. (b) 4 jet trigger.

Figure 4.20: Trigger rate vs threshold for programmable jet algorithm at low luminosity for 1 and 4 jet triggers.

ΔR=0.4.

- Although the trigger E_T is quoted for a 1, 3 or 4 jet trigger, the corresponding P_T given is the P_T for which a single jet trigger at the given threshold would be 90% efficient. In this way the performance of the jet algorithm is assessed only in terms of its threshold sharpness, without folding in the additional factor of resolving power.

This is the approach used in most other studies, and so it is used here to allow easy comparison. In general using K$_\mathrm{T}$CLUS as a jet definition and measuring the efficiency for the third and fourth jets would predict a slightly worse performance for all algorithms; mostly driven by the fact that the K$_\mathrm{T}$CLUS definition is less similar to the trigger algorithms than the CONE algorithm is. The required threshold is given for the jet cluster window size which gave the best performance, which for all energies considered was the 0.8×0.8 window.

Clearly if a 2 kHz trigger rate is allowed then LVL1 will be able to efficiently trigger

Luminosity and Rate	**1 Jet**		**3 Jet**		**4 Jet**	
	Trigger E_T	Jet P_T	Trigger E_T	Jet P_T	Trigger E_T	Jet P_T
Low, 200 Hz	161 GeV	175 GeV	62 GeV	74 GeV	41 GeV	51 GeV
Low, 2 kHz	92 GeV	106 GeV	35 GeV	43 GeV	22 GeV	28 GeV
High, 200 Hz	280 GeV	293 GeV	114 GeV	124 GeV	78 GeV	86 GeV
High, 2 kHz	166 GeV	175 GeV	71 GeV	80 GeV	50 GeV	58 GeV

Table 4.12: Trigger thresholds and equivalent jet P_Ts for which the trigger would be 90% efficient, for a given rate. Results given for Clus=0.8, Step=0.2, RoI=0.2.

from much lower P_T jets (about 50% lower) than if a 200 Hz rate is allowed. But even if it turns out that only 200 Hz is available for the LVL1 jet trigger rate, the resulting trigger menus [25] indicate that the trigger will still be able to comfortably meet the necessary physics requirements. However it should be noted that these are preliminary trigger menus to help in determining the trigger specifications at the time of the LVL1 TDR [22]. As such they are expected to evolve considerably before the final system is built.

Chapter 5

Review of Beam-Tests for the Level-1 Calorimeter Trigger

5.1 Introduction

In order to meet the demanding requirements of triggering at ATLAS, the level-1 (LVL1) calorimeter trigger relies on several critical state-of-the-art technologies. To demonstrate the feasibility of these technologies, a demonstrator system has been developed and tested using real data from prototype ATLAS calorimeters in the CERN test-beam.

The philosophy of the demonstrator programme [34] was to produce a small-scale version of the trigger system which contained all the most important elements of the final design. The LVL1 calorimeter trigger group has performed beam-tests every year since 1992. Beam-tests prior to 1995 were termed "phase-1", and concentrated on demonstrating:

- Bunch-crossing identification (BCID).
- Cluster processing algorithms, and their implementation in application specific integrated circuits (ASICs).

These tests showed that digital BCID was feasible, and demonstrated that it could work efficiently down to very low energies. They also demonstrated that current ASIC technologies were sufficient to implement trigger algorithms of the required complexity and at the necessary speeds. The most important conclusions from phase-1 beam-tests are summarised in [35].

Test-beams after 1995 used the "phase-2" demonstrator system which was designed to simulate the various data transmission and connectivity issues associated with the final LVL1 system as it was then conceived[(a)]. The problems of data transmission and connectivity are critical at LVL1 due to the large number of high speed signals that need to be processed, and the fact that they need to be duplicated in many parts of the system in order to support the overlapping nature of the trigger algorithms.

This chapter describes in detail the October 1997 test-beam activities which concentrated on achieving the necessary level of data integrity in a set-up that replicated the main technologies that will be needed to build the final system. For a broader overview of the test-beam activities see reference [22]. The required data integrity is not unambiguously defined but for the purposes of this study the bit error rate (BER) is required to be sufficiently low that no more than 1% of triggers in the final system would be attributed to a bit error; very roughly this corresponds to around 1 error in 10^9 bits.

5.2 Critical Technologies

The demonstrator programme has been driven by the desire to prove the feasibility of several critical technologies necessary for the implementation of the final system, summarised below.

[(a)] It should be noted that since the test-beam, the architecture of the LVL1 system has become considerably simpler, and as a consequence the backplane fan-out requirements have become less demanding.

5.2.1 Calorimeter Data Transmission

The data from the calorimeters must be transferred about 60 m from the detector to the LVL1 trigger hardware located in the electronics barrack. In order to address issues of distortion and noise, 1996 beam-tests included extensive comparison of electrical, optical, digital, and analogue solutions. The preferred solution is now for simple analogue electrical links.

5.2.2 Bunch-crossing Identification

Three phases of BCID modules have been used in various test-beams to evaluate the performance of BCID algorithms with real signals in a realistic noise environment. In 1993 a single channel of BCID was implemented using discrete hard-wired logic. In 1995 36 channels of BCID processing were implemented in field programmable gate arrays (FPGAs)[b]. In both years the focus was on the behaviour of the BCID algorithm itself; once this was understood the same FPGAs were moved to transmitter modules (described later) and used in their proper context as part of the phase-2 programme.

5.2.3 Inter-Crate Data Transmission

The algorithmic requirement to process overlapping windows means that each trigger tower must be fanned-out to several processor locations.

Minimising the volume of data duplication favours a compact system with maximal processing density. The proposed solution implements the cluster processors and jet/energy-sum processors in four crates each. This however generates another problem; with around 6400 trigger towers, each CP crate has to receive >1600 channels at 8 bits per channel (plus a parity bit). Transferring these data in parallel would require around 14400 connections on the crate front panel which is clearly unfeasible. Therefore a high degree of data serialisation

[b] XiLinx XC4008 devices were used.

and multiplexing is required.

The current solution is to use HP G-Link serialisers (HDMP-1014D) operating at 800 Mbit/s; this corresponds to two 40 MHz channels at 10 bits per channel. Serial links operating at this speed were well tested as part of the phase-2 programme by multiplexing pairs of 8-bit channels onto each link. In the final design additional bits will be transmitted as part of the jet cell sums, parity bits, and bunch-crossing multiplexing, which will make up the remainder of the 10-bit G-Link frame; these features were not tested during the test-beam.

5.2.4 Intra-Crate Data Transmission

Within crates data must be shared between neighbouring modules. It is proposed that this will be done with a high speed, point-to-point, single ended backplane in each crate. In order to keep the pin-count down the cluster processing backplanes are expected to operate at 160 MHz, while the jet/energy-sum backplanes will operate at 80 MHz. Demonstrating that this system can provide high data integrity at these speeds with a densely packed backplane was one of the prime goals of the phase-2 demonstrator.

5.2.5 Miniaturisation

The high processing density of each module means that a high degree of miniaturisation will be necessary for the implementation of the final system. Principally this involves two technologies:

- Application Specific Integrated Circuits (ASICs) will be required to provide the necessary component density for the pre-processor and the cluster processing trigger. Also, to keep the pin-count low on the cluster processing ASICs, i/o is required to be in a semi-serial mode at 160 MHz; this is the limit set by the capabilities of the preferred ASIC technology.

- Multi-chip modules (MCMs) allow the bare dies for many ICs to be bonded directly onto either a laminate or ceramic substrate. This greatly reduces the footprint required for implementing many channels of processing and confines high speed signals to a small area, thus reducing transmission line requirements on the PCBs and providing better data integrity.

Several ASICs were built and tested in the phase-2 demonstrator. An MCM was built but due to production problems was not functioning in time for the test-beam period. It is currently still under test.

5.2.6 Integration

For the first time in October 1997 it was possible to test the feasibility of the system as a functioning trigger. It was important to prove that the phase-2 demonstrator system could be integrated with the central trigger processor, and this system could be used to provide a reliable and correct trigger with constant latency.

This is not so much a proof of any one technology, but a demonstration of the feasibility of the overall trigger design. This test is termed the "complete slice test" and is described in section 5.8.

5.3 October 1997 Test-beam Set-up

Figure 5.1 shows the phase-2 demonstrator system built to test these technologies. It is designed to perform a simplified e/γ algorithm on 36 channels of data in a 6×6 array. Only one depth sampling is considered, and only the part of the e/γ algorithm associated with the EM calorimeter is implemented.

In 1997, data were taken from the prototype tile calorimeter module [36]. Referring back to figure 3.1 in chapter 3 the modules of the demonstrator can be related to the modules in the final system:

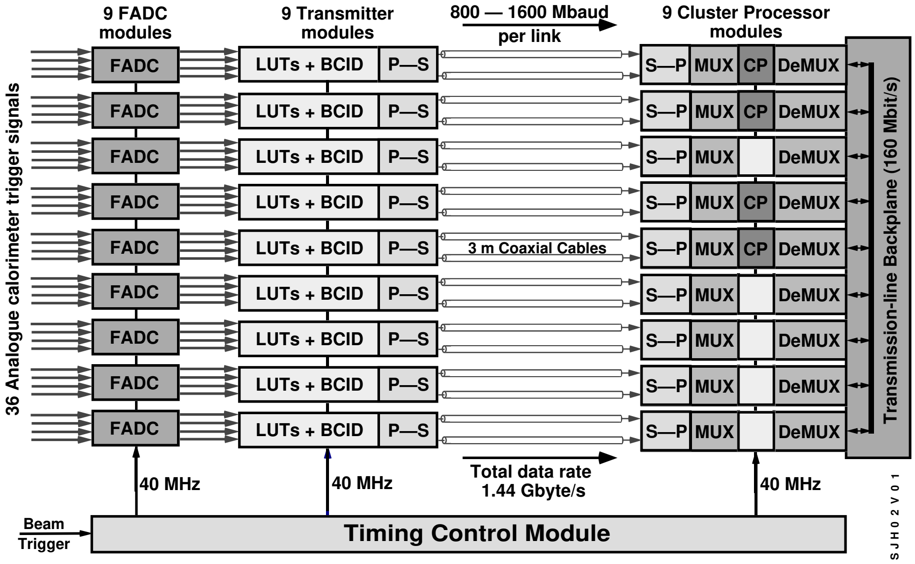

Figure 5.1: Schematic diagram of the 36-channel phase-2 demonstrator.

- **FADCs:** Flash analogue to digital converters digitise analogue signals at 40 MHz to 8-bit precision, with a least count of 1 GeV. In the final system these will be contained in the pre-processors and will work at a higher precision.

- **Transmitter modules:** (TXMs). Implementing the remaining components contained in the pre-processors of the final design: look-up tables (LUTs), BCID, and parallel to serial conversion.

- **Cluster processor modules:** (CPMs). Implementing the e/γ algorithms of the CPMs of the final system. Testing the receiving end of high speed inter-crate transfer, the data integrity of a high speed semi-serial backplane, and the reliability of custom designed ASICs operating with i/o speeds of 160 MHz.

- **Timing control module:** Provides necessary clock and control signals to all modules. Ultimately this module's function will be provided by the overall LHC timing and control system.

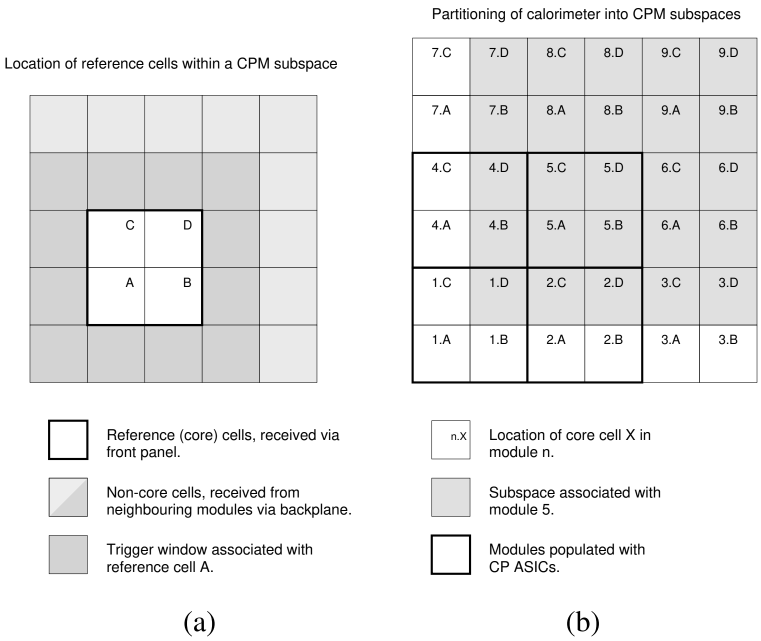

Figure 5.2: Diagram showing (a) the location of reference cells and processing windows within a single CPM subspace, and (b) eta-phi partitioning of the calorimeter into CPM subspaces corresponding to the nine transmission chains.

Although not shown, a very preliminary prototype of the jet processing module was made by the Stockholm group and is discussed later.

Each CPM considers a calorimeter subspace of 5×5 cells. Since the e/γ algorithm is based on a 4×4 sliding window (see chapter 3) this provides sufficient information to process a 2×2 array of overlapping trigger windows. Each trigger window is associated with a specific "reference cell" as illustrated in figure 5.2a.

The system consists of nine parallel data chains. Each data chain carries four trigger towers of data to the CPMs representing the reference cells of the four e/γ windows processed by each CPM. These channels enter the CPMs via the front panel and are referred to as "core cells." The remaining 21 cells needed by each CPM are received via neighbouring CPMs across the backplane. Figure 5.2b illustrates the partitioning of the 36 calorimeter cells into CPM subspaces. Note that only four CPMs have data available from enough channels to fully process any of the e/γ algorithm windows, so only these CPMs are populated with CP ASICs; only CPM 5 has sufficient information to fully process all its algorithm windows so only this board is fully populated by ASICs. The remaining boards do not have CP ASICs but are needed to receive data from the front panel and make it available to the other boards via the backplane.

A crucial difference between this and the final design should be noted here: The calorimeter is partitioned in an "eta-phi" layout (figure 5.2b), compared with the ϕ-quadrant layout of the final design (see chapter 3). This is because an eta-phi layout was the preferred solution at the time the phase-2 demonstrator was designed. Most crucially this affects the complexity of the backplane itself, since in an eta-phi layout a given board can share signals with upto eight neighbouring boards (e.g. module 5); in a crate all eight boards cannot physically be placed directly adjacent to the central board and so signals must be shared between modules that are several slots apart. This results in a backplane containing differing signal lengths and crossing signal lines, with obvious consequences for timing and cross-talk. In a ϕ-quadrant layout signals need only be shared between neighbouring modules in a crate. For this reason any conclusions based on the test-beam backplane represent a particularly pessimistic appraisal of the performance of the backplane proposed in the final system.

Clearly this represents only a tiny fraction of the processing required for the final system. There are only 36 channels instead of $\approx$6400, and each CPM processes only four trigger windows compared with 64 in the final design. However the demonstrator is designed so that in critical areas (such as the CP backplane) the signal densities that will be needed to implement the final system are accurately reproduced.

In order to test the performance of this hardware, circular spy-memories were placed in the data stream before and after each significant processing stage. When an interesting event was detected by a scintillator beam trigger the whole system could be stopped and the data in the spy-memories read-out. In this way the progress of a particular pulse could be tracked along the processing chain, and data integrity studied. This is clearly not the way read-out will be performed in the final system, but it gave a simple way of testing the integrity of data without developing buffered read-out.

5.3.1 Software

Analysis of the system performance relies on three main programs written specifically by members of the LVL1 calorimeter group in order to meet the requirements of the demonstrator programme.

5.3.1.1 Data Acquisition Software

The data acquisition (DAQ) software has evolved gradually from the CERN Spider system [38]. It handles control of all modules, automatically setting them into appropriate modes for running and read-out. In addition it provides simple histogramming and limited on-line analysis. On-line analysis takes the form of comparison histograms, where it is possible to ensure that a pulse has propagated through the transmission chain without distortion. Comparison histograms can be watched in real time so that gross errors in the system set-up can be eliminated quickly and easily, or they can be set to accumulate over a number of events so that lower level errors can be detected.

5.3.1.2 Diagnostic Software

The diagnostics package [39] was designed to provide a user friendly way for physicists and engineers to manipulate and test individual modules, and ultimately the entire system. It is characterised by two main features:

- It can give a graphical representation of the entire system, and the registers within each module, thus allowing direct manipulation of modules in an intuitive way.
- It contains a detailed model of the expected behaviour of each module so that any errors in the behaviour of the system can be exposed automatically.

Figure 5.3 shows a screen shot of the register window for the transmitter module. The status of the various control registers is represented graphically, allowing intuitive understanding of the TXM state and manipulation of its registers.

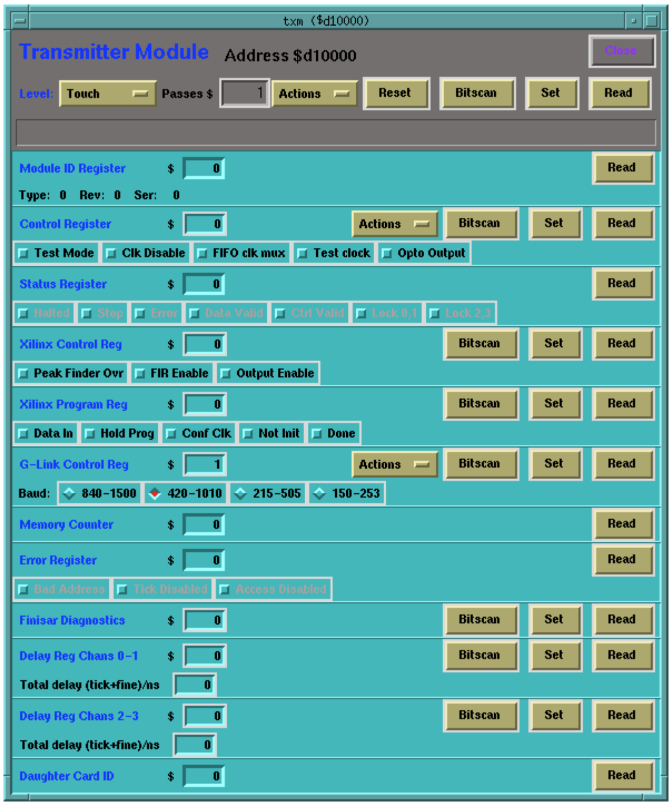

Figure 5.3: Diagnostic software representation of the state of a single transmitter module.

The internal state of each module is simulated at the level of module *entities*. So for example, the state of the TXM is simulated by defining abstract entities representing: memories, BCID processors, G-Links, registers, etc. These entities are connected together by signals that represent the functional connections between processing elements within the module. In this way the overall behaviour can be simulated at a very detailed level with very efficient code use by tracking the propagation of signals within a module, and the state of each of the module entities.

By maintaining a model of the expected state of each module, and updating the state of the model whenever the real module is accessed, the state of the model and the real module should remain in step. Any difference between the model and the real module indicates an error, and the diagnostic software can alert the user. This allows scope for very aggressive testing of the modules in all modes of operation.

The diagnostic software proved to be most useful for initial testing and setting up of the modules. It was invaluable for detecting gross errors in the overall functionality of modules, but was not very useful for lower level errors since it focuses on the instantaneous state of each module, and does not allow accumulation of data over many events.

5.3.1.3 Off-line Analysis Software

Many thousands of events were taken during the test-beam period, and stored on DAT tapes. In order to analyse these data an off-line analysis program was used.

This program also contains a model of the expected state of each of the module memories, but at a much less detailed level than the diagnostic software. It calculates the expected data transformation of each module, and uses this to detect automatically errors in the system.

In general there were several kinds of low-level error happening in a few percent of events. By correlating which memories in each module saw discrepancies with the model, the software was able to classify each event in terms of the kinds of error that caused the discrepancy. In this way very low-level bugs (of the order of 1% or less) could be identified, and this information could be used to adjust the parameters of the system.

Many errors that were detected in initial runs were the result of problems with the set-up, and could be cleared up by appropriate adjustment of the experimental system. This done, the remaining task was to measure the intrinsic bit error rates in the system caused by irreducible cross-talk and noise. Using longer sets of data the off-line software could also be used to obtain a rough measurement of this once the system was working in an optimum way; this method is however unavoidably statistics limited.

5.4 Transmitter Module Performance

Prior to the October 1997 test-beam period the transmitter modules (TXMs) had been thoroughly debugged. Internally there were no problems clocking data into the TXMs or with the BCID performance. Statistics limited tests during the test-beam calculated bit error rates (BER) $< 10^{-8}$ within the TXMs, and laboratory tests in 1998 reduced this limit to $BER < 10^{-13}$. Figure 5.4 shows a tile calorimeter pulse in two adjacent channels being processed by two separate TXM modules. Each bin in the x-axis represents a single 25 ns sample in time, and the y-axis shows the number of FADC counts measured. The BCID processing converts the extended pulse into a single impulse occupying only one time-slice. Critically this can be done reliably across many separate modules with a constant latency.

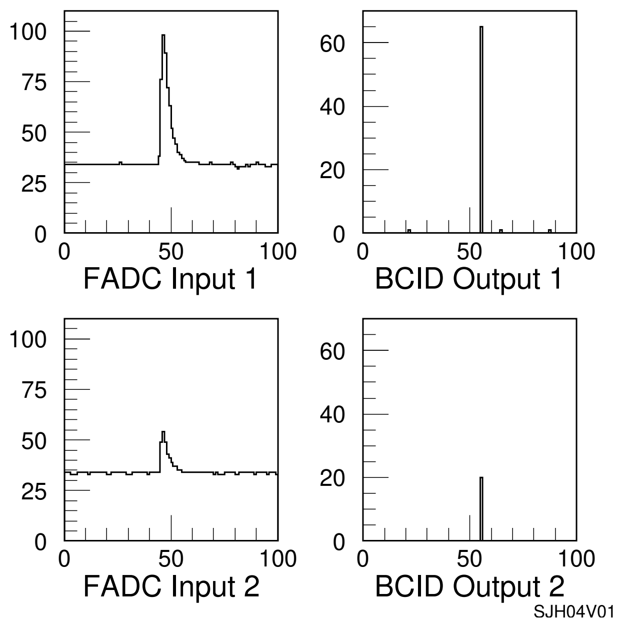

Figure 5.4: Two TXMs performing BCID on the same particle, seen in two adjacent channels of the tile calorimeter. Number of FADC counts is represented on the y-axis, against time on the x-axis.

In general G-Link operation also fulfilled required data integrity criteria but with some reservations. Two modes of operation were provided: 800 Mbit/s multiplexing two channels onto each link; and 1600 Mbit/s multiplexing four channels onto each link. The 800 Mbit/s mode worked reliably with $BER < 10^{-15}$. The 1600 Mbit/s mode however proved to be unstable, and the links failed whenever data were read-out from the TXM modules via the DAQ system. In principle this is not a critical problem for two reasons:

- The 1600 Mbit/s mode was being studied as a contingency plan to reduce the number of G-Links required if the overall power requirements became too large. With the advent of bunch-crossing multiplexing it is unlikely that this mode of operation will ever need to be used in the final system.

- Errors only occurred while the TXMs' internal memories were being read-out by the DAQ system. As has been mentioned before, this method is only relevant for the demonstrator, and will never be used in the final system.

However, these problems indicate that operation of the G-Links is not fully understood. The fact that potential problems have surfaced in the demonstrator system implies that further understanding is required to ensure that problems don't arise in the much more complicated final system. In particular, the problems with the 1600 Mbit/s mode were apparent in the 800 Mbit/s mode during the 1996 run; these problems were eliminated by improving the G-Link power supply isolation, indicating that better understanding of the G-Link power supplies is still needed. Further studies are currently underway in an attempt to improve the robustness of the G-Links.

5.5 Front-End Module Performance

Although not shown in the diagram, a front-end module (FEM) [40] was designed by the Heidelberg group that includes more of the functionality that will ultimately be required in the final pre-processor design. Two FEMs were built and used in place of TXMs. They have the same functionality as the TXM and were found to function with equal reliability.

In addition they included a preliminary implementation of "buffered read-out" that will be required for the final design. In the final system it is envisaged that diagnostic information about the content of the trigger systems will still be required. However, it will not be possible to periodically stop the system as in the demonstrator, and so an internal buffering system will be needed so that a few time-slices either side of a trigger event can be stored and read-out periodically without stopping the trigger pipeline.

Evaluation of this mode was difficult since not many modules had this capability, so cross-checks were limited. However, preliminary results indicate that buffered read-out can work reliably well above the expected LVL1 trigger rate. An example of the operation of buffered read-out mode is given as part of the discussion on the complete slice tests in section 5.8.

5.6 CPM Performance

Classification of the errors appearing in the CPMs is a more complicated affair due to the large number of operations performed in each CPM. Within a single board the following tasks are performed:

- Data are de-multiplexed from 800 Mbits/s serial to 40 MHz parallel with HP G-Link receiver chips.

- These data are multiplexed back up to 160 MHz suitable for transmission on the high speed backplane and input to the CP ASICs. This is done on custom designed multiplexing/de-multiplexing ASICs referred to as RAL163 dual purpose ASICs [37]. These ASICs can operate in either multiplexing (40 MHz $\rightarrow$ 160 MHz) mode, or de-multiplexing (160 MHz $\rightarrow$ 40 MHz) mode.

- As a purely practical measure, the ASICs containing the e/γ algorithm were originally designed to accept data at 40 MHz. For this reason an extra layer of de-multiplexing is required to convert 160 MHz data from the backplane to 40 MHz suitable for the

ASICs. This still fulfils the requirement to test ASIC i/o at 160 MHz however since the de-multiplexing ASICs use the same technology as the CP ASICs; in the final system this function will be integrated within the CP ASICs.

Figure 5.5 shows the data flow through a single CPM board. Each board receives data from the front panel (core cells) and from neighbouring modules via the backplane (non-core cells) as indicated in figure 5.2. To investigate the data integrity of each stage in the CPM processing, data are captured in spy-memories at three stages: after G-Link de-multiplexing but before 160 MHz multiplexing (FIFOs); after 160 MHz multiplexing (input memories); and after e/γ processing (output memories).

The off-line analysis software was used extensively to deduce which part of the system was causing errors by correlating which memories showed discrepancies with the expected data. Expected data were calculated for each memory by copying the real data from a previous memory in the data stream, and applying an appropriate transformation. The locations of source memories used to calculate the expected data at each stage in the CPM are listed as follows:

- **Expected FIFO data** are calculated from the appropriate TXM output memory.
- **Expected input memory** core cells are similarly calculated from the same TXM output memories.
- **Expected input memory** non-core cells are calculated from the input memory core cells of a neighbouring CPM. These are the channels that are transmitted across the backplane in the real system.
- **Expected output memory** data are calculated from the appropriate input memory locations having applied a simulation of the e/γ finding algorithm.

Table 5.1 summarises the relationship between which memories disagree with the expected data and the probable location of the cause of the error. The input memories are

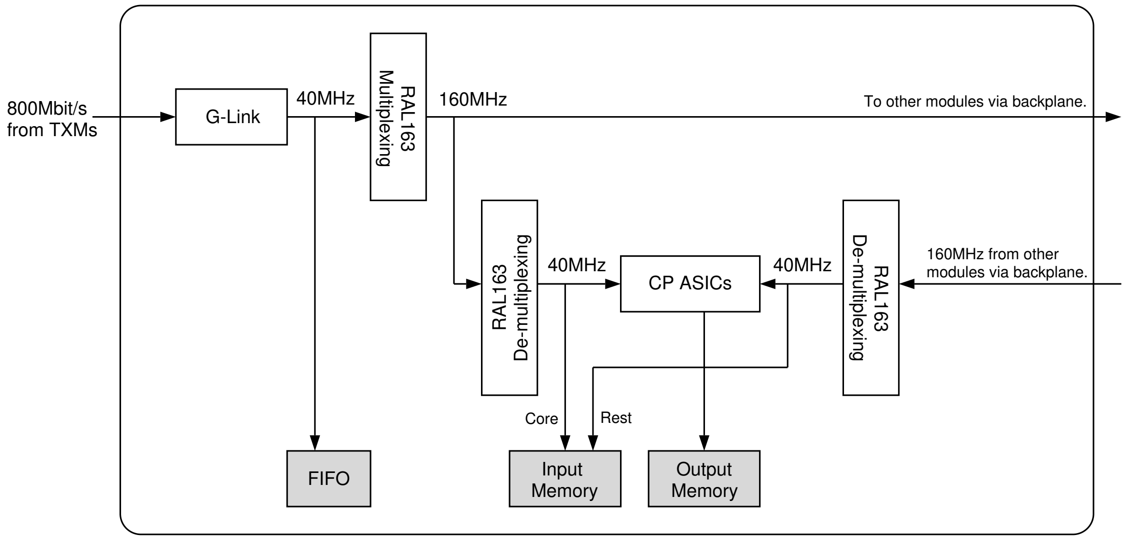

Figure 5.5: Data flow and location of spy-memories in the demonstrator cluster processing modules.

Error Location	**Mismatch Location**			
	FIFO	Input Memory (Core)	Input Memory (Rest)	Output Memory
G-Link Receiver	X	X	0	0
FIFO Capture	X	0	0	0
160MHz MUX/De-MUX	0	X	0	0
Backplane	0	0	X	0
Input Memory	0	X	X	X
Ouput Memory	0	0	0	X

Table 5.1: Summary of relationships between location of off-line analysis memory mismatches, and suspected error source. "X" indicates a mismatch between the recorded data and the expected data calculated by the off-line software. "0" indicates no mismatch.

divided into locations associated with core cells, having arrived directly via the front panel, and the rest which have arrived from a neighbouring CPM via the backplane. An “X” indicates that a mismatch between real and expected data has occurred. The correlation that can be seen is driven by the way in which the software works. Since the expected data are calculated from the previous memory in the transmission chain, errors that affect the capture of data, but not the transmission chain itself, will appear as mismatches that propagate throughout the remainder of the system; for example, if the input memories themselves are failing to latch data correctly. Conversely, errors that affect the transmission chain but not the capture of data appear as mismatches which are *not* propagated throughout the rest of the system; for example, a G-Link receiver error. The exception to this is the FIFOs which are not used to calculate the expected data for any other memories. Of course some events can contain several such errors simultaneously, in which case the classification of errors becomes ambiguous; this is not a significant problem however since in general errors only occurred in a few percent of events.

There had not been time to thoroughly debug the CPMs before the test-beam period, and so in general the intrinsic errors related to noise and cross-talk were swamped by other problems. Most of the test-beam was spent understanding these errors, so that they could be resolved in time for laboratory tests later. Known problems during the test-beam were:

- A problem with FIFO control caused these data to be mis-timed with respect to the rest of the data, causing periodic latency variations of several time-slices.

- Periodic loss of synchronisation in the RAL163 multiplexing/de-multiplexing ASICs caused by bad timing of the RAL163 clock strobes. This could be fixed by resetting the RAL163 ASICs and so was not considered an intrinsic error.

- Errors in the input memories themselves caused by poor termination and control problems.

- Fan-out of signals onto the backplane was slower than designed due to unexpected capacitances. This resulted in signals arriving over the backplane being mis-timed, and requiring a temporary timing fix that was not entirely robust.

Under these conditions the average error measured over all CPM modules for data in the core cells was $BER \approx 3 \times 10^{-5}$. In the non-core cells (i.e. data having come over the backplane) the BER was approximately 5×10^{-5}.

This is clearly not satisfactory, but it was felt that this was caused by resolvable errors, and was not related to any particular technology choice. Studies in the previous year had measured the BER on the multiplexing/de-multiplexing ASICs to be better than 10^{-8} (statistics limited), and detailed observation of the signals on the backplane gave no reason to suspect error free data could not be transmitted in this way.

Figure 5.6 shows an example of a signal pulse being successfully captured in the core cells of a CPM, and simultaneously fanned-out to a neighbouring module across the backplane. The delay in the signal position between the FADC and the CPM represents the latency of the system; however the apparent latency visible here is controlled by read-out issues and does not represent the real latency of the final system. By definition no latency can exist between CPM core cells and backplane data since these data must be operated on simultaneously in order to process the trigger algorithms. Note that in this case the BCID processing in the TXM was turned off.

5.6.1 Subsequent Laboratory Tests

To provide confidence that the system was feasible further tests were still required. These were carried out back at the Rutherford laboratory (RAL) during the beginning of 1998. The errors observed in the test-beam were resolved and an alternative way to measure errors was used. The problem is that off-line analysis of bit error rates is limited by the length of time it takes to read and analyse events. A bit error rate tester module was built that operates in real time by comparing digitised pseudo-random data that have been passed directly to

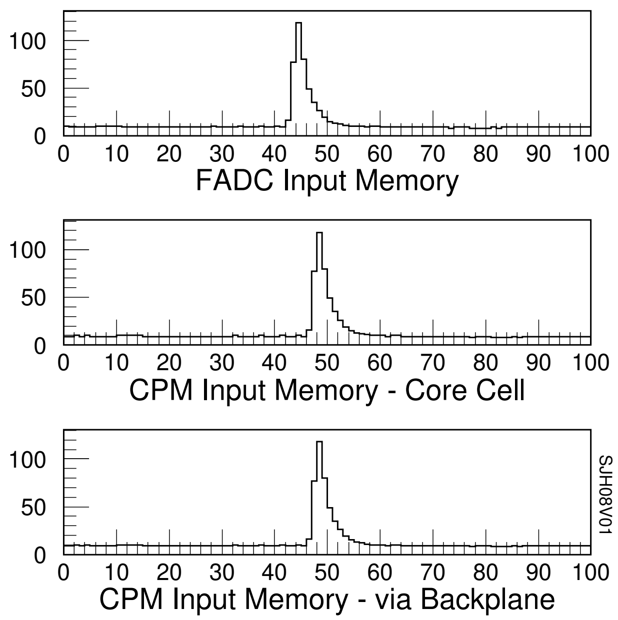

Figure 5.6: Example of tile calorimeter data from one FADC flowing perfectly onto two CPM boards, one directly and one via the backplane.

the tester module with the same data once they have travelled along the usual test-beam transmission chain. This relies on the fact that each CPM makes its input data available on its front panel as 40 MHz parallel ECL signals.

This system operated in two modes: firstly data directly from the FADCs were compared with data that had been transmitted via a TXM into a CPM; this tests the TXM module, the G-Link transmitter system, and 160 MHz multiplexing and de-multiplexing. In the second mode the core cell data of one CPM were compared before and after transmission to a neighbouring CPM across the backplane; hence this tests backplane transmission and 160 MHz multiplexing and de-multiplexing.

Since this system can measure bit error rates in real time without stopping the system it can accumulate statistics in a fraction of the time of software based analyses. It should be noted however that this method is not useful for pin-pointing the position of errors, and can only be performed once the diagnostic software and the off-line software have been used to thoroughly debug the modules.

Using this technique bit error rates have now been measured to be better than 10^{-13} across the whole system. This is well within the required limits of data integrity. Studies are still continuing at RAL to better understand the functioning of the G-Links, and to test the backplane under a range of noise conditions and data patterns.

5.7 Jet Processing Module Performance

The Stockholm group produced a jet processor demonstration module (JPMD). The module was not intended to implement a realistic jet algorithm, but to gain experience using field programmable gate arrays (FPGAs) which will be used in the final system. In line with an earlier trigger design each CPM makes available four 8-bit sums on its front panel representing the total energy in each 4×4 trigger window considered by that CPM. These data, which are fully parallel signals at 40 MHz, were received by the JPMD and a XiLinx XC4000EX FPGA used to perform simple summing and threshold comparisons. These data were then made available to the CTP.

In the test-beam the JPMD showed an unacceptable level of bit errors ($\approx 10^{-6}$). This was caused primarily by the choice of XiLinx which was not able to run reliably at 40 MHz, and by insufficient timing control on the board. With these design issues in mind, a new jet processor module is currently being built called the *jet/energy-sum 'minus-2' demonstrator*. As well as using a more robust mechanism for distribution of on-board clocks, and FPGAs which should run reliably up to 100 MHz, it will also implement a more complex algorithm more closely resembling the final jet/energy-sum proposal. This module is due to be completed before the end of 1998.

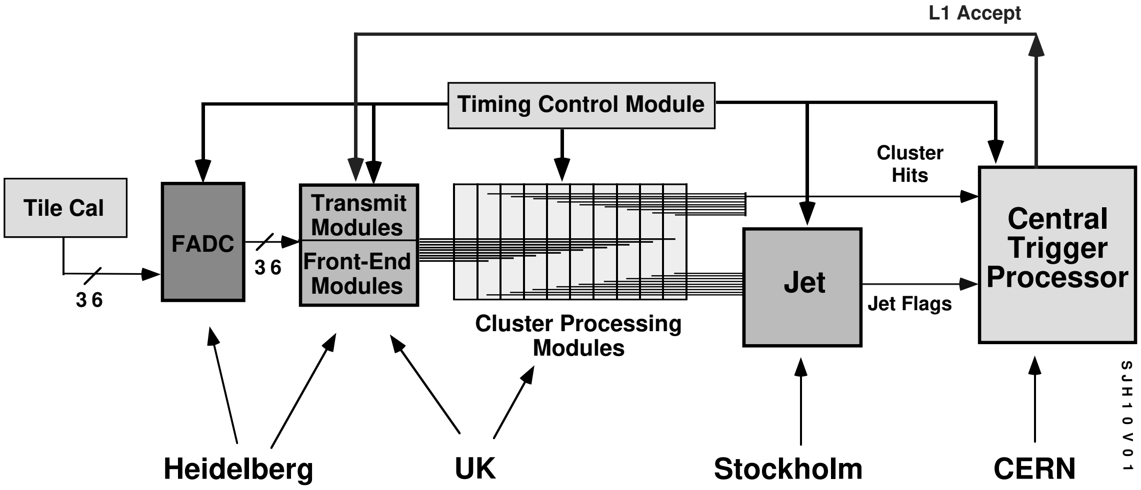

Figure 5.7: Schematic diagram of the complete slice set-up, designed to test the feasibility of the overall trigger design.

5.8 Complete Slice Test Performance

The complete slice test was designed to test the functionality of the trigger system as a whole. Figure 5.7 shows the system used. The transmission chains used throughout the test-beam period remained unchanged, but hit data from the CPMs and the JPMD were transmitted to the central trigger processor demonstrator (CTPD). Crucially, the CTPD used this information to generate an overall LVL1 accept (L1A) signal, which could be used as a genuine beam trigger.

This system was used in two modes:

- Instead of using the scintillator beam trigger to stop the demonstrator trigger pipelines, the L1A signal was used. In this way the system operated as a true trigger for the first time. Only when pulses above a certain threshold were encountered were data read-out. Lowering the input analogue signal caused the system to stop producing triggers.

- This system was also used to investigate the buffered read-out mode on the FEMs.

Instead of using the L1A signal to stop the trigger pipelines, it was used to initiate read-out of the FEM buffered memories. In this way eight time-slices surrounding each L1A were captured in the FEM buffers without stopping the system. Read-out of the FEM buffers by the DAQ system (still without stopping the trigger pipeline) was initiated by an FEM flag indicating that its buffers were full. Of all the test-beam set-ups investigated so far this system most closely resembles the final operation of the LVL1 trigger.

Figure 5.8 shows the output from one channel running in buffered pipeline mode. Each spike visible is equal to the peak sample of the calorimeter pulse that was responsible for the L1A signal. Accompanying each spike four time-slices either side were also read-out, but these appear as zero since they are suppressed by BCID processing. Note that the spacing between spikes is equal to the number of slices read-out by the FEM; since read-out is only triggered by a L1A the time between successive triggers is not represented by this plot.

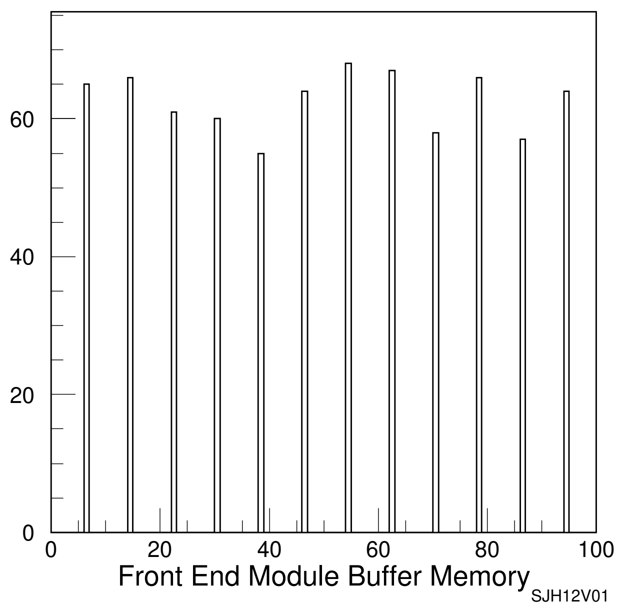

Figure 5.8: BCID processed data captured in the FEM buffer memory using L1A signal to initiate buffering.

5.9 Conclusions

Schedules for getting modules into the test-beam have always been very tight, and this has often led to modules being used in the test-beam before they have been adequately tested and debugged. This means that few of the rigorous bit error rate measurements that the test-beam programme was intended for have actually been carried out during the test-beam period. Nonetheless, the test-beam programme has proved an invaluable part of the LVL1 calorimeter trigger development cycle. Taken with the laboratory tests that invariably follow, the experience gained during the development of the demonstrator has led to several fundamental changes in the design of the trigger:

- In the 1996 test-beam the backplane used BTL (backplane transmission logic) technology. The failure of this system led directly to the decision to use ECL technology instead.
- Understanding of the backplane timing issues led to the change from eta-phi partitioning of the calorimeter to ϕ-quadrant partitioning.
- The continuing requirement to understand further the operation of the G-Links has arisen from persistent low-level errors during test-beams.
- Overall experience of getting an increasingly large and complex system to work has generated recognition for the challenging timing issues that will be encountered in the design of the final system.

The latest set of laboratory tests following the 1997 test-beam have shown that the whole system can function reliably with an overall bit error rate better than 10^{-13}, giving confidence that the proposed technologies will be able to meet the requirements of the LVL1 trigger.

With the publication of the LVL1 trigger Technical Design Report [22] no more major changes in the trigger architecture are anticipated. The main task now for the group is to continue work pinning down final specifications of the system, and to begin to tackle the challenging issues of timing in the final system. One aspect where the demonstrator has so far not accurately simulated the final system has been the timing of signals. The demonstrator contained a large number of programmable delays to ensure that data could be latched reliably. In most cases during the test-beam these delays were adjusted by hand in order to optimise the system. This was a tricky, and long-winded task in the demonstrator, and will be completely unfeasible for the final system, even in the light of the simplified backplane layout. A great deal of work is required in the following years to understand the requirements on automated calibration and timing systems.

Chapter 6

Consequences of Lepton R-parity Violation in Supersymmetry

6.1 Introduction

Chapter 1 introduced supersymmetry (SUSY) as a very appealing model that elegantly solves the problems of electroweak symmetry breaking and might describe physics in the TeV energy range. Supersymmetry proposes a new symmetry between the fundamental particles of matter (fermions) and the fundamental particles of force (bosons). The simplest practical SUSY model is called the *minimal supersymmetric extension to the standard model* (MSSM). The particle spectrum for the MSSM is summarised in table 6.1. The number of particles is approximately doubled; every fermion is associated with a boson-like partner, and every boson is associated with a fermion-like partner[a]. In addition, the Higgs sector must now contain at least two doublets [41]. After electroweak symmetry breaking there are five physical Higgs boson states: $h^0, H^0, A^0, H^\pm$. The supersymmetric partners of the

[a] The naming convention for SUSY particles is this: the partners to the fermions are named by prefixing the SM particle name with an 's', eg. the partner to the electron is the *s*electron. The partners to the bosons are named by suffixing the SM particle name with 'ino', eg. the partner to the photon is the phot*ino*. All SUSY particles are labelled with a *tilde*, eg: $\tilde{e}, \tilde{\gamma}$.

	SM: Spin-$\frac{1}{2}$ SUSY: Spin-0			SM: Spin-0,1 SUSY: Spin-$\frac{1}{2}$		
SM	$[u,d,c,s,t,b]_{L,R}$	$[e,\mu,\tau]_{L,R}$	$[\nu_e,\nu_\mu,\nu_\tau]_L$	g	$\underbrace{W^\pm\ H^\pm}$	$\underbrace{\gamma\ ZH_1^0H_2^0H_3^0}$
SUSY	$[\tilde{u},\tilde{d},\tilde{c},\tilde{s},\tilde{t},\tilde{b}]_{L,R}$	$[\tilde{e},\tilde{\mu},\tilde{\tau}]_{L,R}$	$[\tilde{\nu_e},\tilde{\nu_\mu},\tilde{\nu_\tau}]_L$	$\tilde{g}$	$\tilde{\chi}^\pm_{1,2}$	$\tilde{\chi}^0_{1,2,3,4}$

Table 6.1: Supersymmetric counterparts to the SM particles including the extended higgs sector of MSSM.

Higgs and the partners of the electroweak bosons mix to give two charged (chargino) and four neutral (neutralino) mass eigenstates.

Hence supersymmetry introduces a large spectrum of currently unobserved particles which, if supersymmetry is to be relevant to the electroweak scale, will have masses that will be observable at the LHC. Understanding this region is a critical goal of the LHC. When the LHC comes on-line investigation will proceed initially with an inclusive analysis designed to establish the existence of supersymmetric processes. If this is successful then attempts will be made to understand the nature of supersymmetry by making precision measurements of the super-particle (sparticle) masses.

The MSSM is very poorly constrained, having many free parameters, and within this model only very limited studies of SUSY at the LHC can be made. Initial studies of the reach of the LHC were made in the ATLAS Technical Proposal [18], but were unable to say anything about the potential of ATLAS to measure sparticle properties. To overcome this problem the ATLAS collaboration chose to study a relatively well constrained model called "supergravity."

Supergravity (SUGRA) refers to the mechanism responsible the breaking of supersymmetry that gives the sparticles different masses to their SM counterparts. The advantage of SUGRA inspired models is that the number of free parameters in the model is reduced to just five. The ATLAS SUSY working group picked 5 'LHC-Points' in this parameter space designed to produce a wide range of different phenomenologies. These 5 LHC-Points were studied in detail in a series of papers [42–46] that investigate the potential of ATLAS both to discover SUSY if nature has chosen such a model, and to make precision measurements capable of constraining the parameters of a SUGRA model.

In all these studies however it was assumed that a discrete symmetry exists called R-parity (R_P). R_P is defined as $R_P = (-1)^{(3B+L+2S)}$ where L is the lepton number, B is the baryon number, and S is the spin. Therefore all SM states have $R_P = +1$, and all SUSY states have R_P = -1. R_P is usually assumed to be conserved since this prevents the existence of SUSY particles affecting the phenomenology of SM processes. Most notably, without R_P conservation the proton could decay with a very short lifetime by the emission of a virtual squark, which is clearly at odds with experimental observations.

There are two main consequences of R_P conservation in MSSM models (R_P-MSSM). Firstly, SUSY particles can only be produced in pairs. At the LHC the dominant production will be pairs of gluinos and squarks, which will then cascade decay into the lighter SUSY particles. Secondly the lightest supersymmetric particle (LSP) will be stable. The missing energy carried away by two escaping LSPs per event is one of the most distinctive SUSY signatures.

However, R_P conservation is not required for internal consistency of the model, and SM observations can be preserved equally well by requiring the conservation of some other quantum number (for example lepton, or baryon number). A scenario such as this could have very different phenomenology to the R_P-MSSM case. Therefore it is necessary to investigate the potential of ATLAS to measure SUSY parameters in scenarios where R_P is violated. This chapter looks at the case where R_P is violated via explicit symmetry breaking terms causing the LSP to decay into two jets and a lepton; this scenario will be referred to as $\not{R}_P-MSSM$.

The following sections discuss the SUGRA framework used for these studies and describe the phenomenology of this model in an R_P conserving scenario. The consequences of R_P violation are described, and the effect that this has on ATLAS' ability to make precision measurements of SUSY particles is investigated for a representative LHC-Point (point 5).

6.2 Supergravity Phenomenology

Since supersymmetric particles have not yet been observed it is assumed that supersymmetry is a broken symmetry. Within the framework of the MSSM there are several proposed mechanisms for this symmetry breaking. To date one of the most promising models that is being studied extensively within the ATLAS collaboration is supergravity [47].

The minimal supergravity (SUGRA) model has the same particle content as the MSSM. It postulates the existence of extra superfields (called the 'hidden sector') which couple to the MSSM sparticles only through gravitation-like interactions. It is this coupling that causes the symmetry breaking responsible for giving the SUSY particles different masses to their SM counterparts. The advantage of this scheme is that all sparticles couple universally to the hidden sector, reducing the number of SUSY breaking terms from well over one hundred, to just five. The sparticle spectrum, and hence the model phenomenology, is strongly dependent on the values of these parameters, which are:

- $\mathbf{m_0}$: Universal scalar mass.
- $\mathbf{m_{\frac{1}{2}}}$: Universal gaugino mass.
- $\mathbf{A_0}$: Trilinear breaking term.
- $\mathbf{tan}\beta$: Ratio of vacuum expectation values of the two Higgs fields at the weak scale.
- $\mathbf{sign}\mu$: Sign associated with this ratio.

The parameter $m_{\frac{1}{2}}$ is almost completely responsible for determining the gluino mass. The squark masses are determined primarily by the larger of m_0 and $m_{\frac{1}{2}}$. Unless m_0 is very small the slepton masses are most sensitive to this. The parameters $tan\beta$ and $sign\mu$ control the electroweak mixing, and as such affect the masses and couplings of the higgs bosons. Low energy processes are insensitive to A_0 and many choices of A_0 will result in the same sparticle spectrum.

The 5 LHC-Points studied are summarised in table 6.2.

LHC-Point	m_0 (GeV)	$m_{\frac{1}{2}}$ (GeV)	A_0 (GeV)	$tan\beta$	$sign\mu$	**Total Cross-section (pb)**
1	400	400	0	2.0	+	3.4
2	400	400	0	10.0	+	3.3
3	200	100	0	2.0	-	1365.0
4	800	200	0	10.0	+	27.1
5	100	300	300	2.1	+	19.2

Table 6.2: Parameters describing the five LHC-Points in SUGRA parameter space chosen for precision study in ATLAS. Cross-sections quoted were calculated using ISAJET.

At the LHC gluinos and squarks will be most copiously produced. These particles then decay into other SUSY particles, which in turn decay to lighter SUSY particles, causing a cascade decay chain ending in the production of two χ_1^0s. Since these are the lightest supersymmetric particles in R_P-MSSM they are unable to decay and leave the detector.

Table 6.3 shows some representative sparticle masses for LHC-Point 5, and table 6.4 shows some of the dominant branching ratios. The cascade decays of these SUSY particles, and the escape of two χ_1^0s per event lead to the characteristic signatures of SUSY: events with multiple jets, leptons, and missing E_T.

In order to study the consequences of R_P violation point 5 is used throughout for comparison since it offers a wide variety of signals. Where not stated all studies were performed using point 5; however it is important to determine whether or not these techniques remain applicable across a range of SUGRA parameters so in specific sections examples of these analyses with other LHC-Points can be found.

Sparticle	Mass (GeV/c^2)	Sparticle	Mass (GeV/c^2)
u_L, c_L	643	u_R, c_R	618
d_L, s_L	646	d_R, s_R	616
t_1	490	t_2	609
b_1	579	b_2	616
e_L, μ_L	241	e_R, μ_R	157
ν_{e_L}, ν_{μ_L}	232		
τ_1	157	τ_2	241
ν_{τ_1}	232		
χ_4^0	570	$\chi_2^\pm$	567
χ_3^0	550	$\chi_1^\pm$	228
χ_2^0	229	$\tilde{g}$	729
χ_1^0	119		

Table 6.3: Sparticle masses for point 5 calculated using PYTHIA.

Decay	Branching Ratio	Decay	Branching Ratio
$\tilde{g} \rightarrow \tilde{q}_L q$	21%	$\chi_4^0 \rightarrow \chi_1^\pm W^\pm$	18%
$\tilde{g} \rightarrow \tilde{q}_R q$	16%	$\chi_4^0 \rightarrow \chi_2^\pm \tau \nu_\tau$	18%
$\tilde{g} \rightarrow \tilde{b}_L \bar{b}$	15%	$\chi_4^0 \rightarrow \chi_2^0 h^0$	14%
$\tilde{g} \rightarrow \tilde{b}_R \bar{b}$	14%		
$\tilde{g} \rightarrow \tilde{t}_1 \bar{t}$	14%		
$\tilde{q}_L \rightarrow \chi_1^- q$	65%	$\chi_3^0 \rightarrow \chi_1^0 Z^0$	10%
$\tilde{q}_L \rightarrow \chi_2^0 q$	32%	$\chi_3^0 \rightarrow \chi_1^\pm W^\pm$	31%
$\tilde{q}_R \rightarrow \chi_1^0\ d$	99%	$\chi_3^0 \rightarrow \chi_1^\pm W^\pm$	31%
		$\chi_3^0 \rightarrow \chi_2^0 Z^0$	27%
		$\chi_2^0 \rightarrow \chi_1^0 h^0$	76%
		$\chi_2^0 \rightarrow \tilde{l} l$	12%
$\tilde{l}_L \rightarrow \chi_1^- \nu_e$	13%	$\chi_1^+ \rightarrow \chi_1^0\ W^+$	99%
$\tilde{l}_L \rightarrow \chi_1^0\ e$	81%	$\chi_2^+ \rightarrow \chi_2^0 W^+$	26%
$\tilde{l}_R \rightarrow \chi_1^0\ e$	100%	$\chi_2^+ \rightarrow \chi_1^+ Z^0$	23%
		$\chi_2^+ \rightarrow \chi_1^+ h^0$	20%
		$\chi_2^+ \rightarrow \tilde{t}_1 \bar{b}$	18%

Table 6.4: Dominant decay modes for LHC-Point 5 calculated using PYTHIA.

6.3 R_P Violating Interactions

Conservation of R_P is elegant since it removes any problems of rapid proton decay, and causes the LSP to be stable and hence a good candidate for cold dark matter. However, it is an *ad hoc* symmetry that has been imposed by hand to preserve the observations of the standard model. Since there is no gauge symmetry that requires R_P conservation it is sensible to consider other models that are compatible with experiment.

In models with a minimal field content, the R_P violating part of the superpotential can be written as [48]:

$$W_{\not R_P} = \sum_{ijk} \overbrace{\lambda_{ijk}\hat{L}_i\hat{L}_j\hat{E}_k}^{\chi_1^0 \to lll} + \overbrace{\lambda'_{ijk}\hat{L}_i\hat{Q}_j\hat{D}_k}^{\chi_1^0 \to qql} + \overbrace{\lambda''_{ijk}\hat{U}_i\hat{D}_j\hat{D}_k}^{\chi_1^0 \to qqq}$$

where: $\hat{L}, \hat{Q}, \hat{E}, \hat{U}, \hat{D}$ are the superfields representing lepton and quark doublets, charged lepton, up and down quark singlets, respectively. The indices i, j, and k represent the generations. The λs represent the Yukawa couplings responsible for R_P violation. A non-zero λ causes lepton number violating couplings that allow the LSP to decay into three leptons. A non-zero λ' also causes lepton number violation, but the LSP decays into two jets and a lepton. Finally, non-zero values for λ'' allow baryon number violation and the LSP decays into three quarks.

To avoid couplings that allow the proton to decay faster than experimental bounds it is usually assumed that either the L-violating (λ, λ') terms or the B-violating (λ'') term exist, but not both at once. A range of SM processes have been used to place upper limits on the values of these couplings; a summary of present experimental limits for the R_P violating couplings can be found in [49].

Other studies have considered the cases where λ [50], and λ'' [51] are non-zero. This chapter focuses on the remaining case where λ' is non-zero.

R_P violation can profoundly change the phenomenology of the SUGRA model; if these couplings are large enough they could entirely change the sparticle spectrum. For the purposes of this study, several assumptions have been made about the nature of R_P violation;

these assumptions are arbitrary and it is easy to construct other phenomenologies that might require quite different search strategies. However, in order to avoid testing endless permutations of phenomenologies this scenario is chosen to explore ATLAS' capabilities of measuring this kind of R_P violation in a representatively challenging case. These assumptions are:

- It is assumed that the Yukawa couplings are small and that the SUSY gauge couplings dominate. Thus, the *only* effect on the SUGRA phenomenology is to allow the LSP to decay into 2 jets and a lepton via the diagram shown in figure 6.1.

- The lifetime of the LSP is assumed to be short enough *not* to produce a displaced vertex in the detector. Although it is feasible that a displaced vertex would be measurable this would produce a very distinct signature making detection and measurement even easier.

- It is assumed that the LSP is still the lightest neutralino (χ_1^0) as is the case in SUGRA. This is necessary only from cosmological considerations however and is therefore not strictly required in an R_P violating model.

- There is no model constraint on the couplings of the LSP to either charged or neutral leptons, so for example the decays $\chi_1^0 \to eud$[b], and $\chi_1^0 \to \nu_e d\bar{d}$ are both possible. The relative coupling depends on the exact composition of the χ_1^0. If it is photino-like then it will decay predominantly into charged leptons. If it is mostly zino-like then it will decay preferentially to neutral leptons [52]. Since any assumption of this branching ratio would be arbitrary both scenarios are considered separately.

- Only LSP decays to the 1st generation are considered. That is, $\chi_1^0 \to eud$ and $\chi_1^0 \to \nu_e d\bar{d}$. χ_1^0 decays to muons would have very similar search strategies and so have not been considered. χ_1^0 decaying to taus would have a final state resembling either the all hadronic decay of the χ_1^0 via λ'' or the leptonic decays discussed in this chapter. In any case,

[b] Note that in the case of charged lepton decay $\chi_1^0 \to eud$ is used to signify both possible charge states: $\chi_1^0 \to e^- u\bar{d}$ and $\chi_1^0 \to e^+ \bar{u} d$. These were simulated with a 50% branching ratio between them.

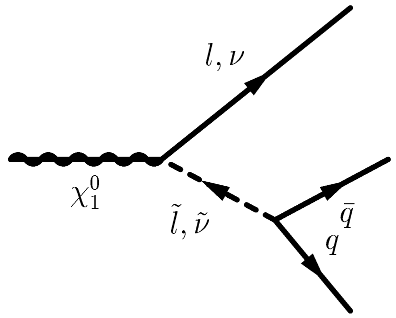

Figure 6.1: Decay of χ_1^0 due to the R-parity violating interaction from the term $\hat{L}\hat{Q}\hat{D}$. The LSP decays to two jets and a lepton via the emission of a slepton.

the neutral lepton decay of the χ_1^0 represents a more challenging case and has been selected for detailed analysis.

6.3.1 Phenomenology of R_P Violating SUGRA

Figure 6.2 shows the P_T distributions of the 9 leading jets for LHC-Point 5 in the case where the LSP decays to eud. The most obvious change to the standard SUSY signature if the LSP decays is the absence of missing E_T. Figure 6.3 compares the missing E_T spectrum for $R_P - MSSM$ and $\not{R}_P - MSSM$ for both charged and neutral LSP decays. The E_T^{MISS} spectrum is clearly much softer and cannot be used to suppress SM backgrounds. In the case where $\chi_1^0 \rightarrow \nu_e d\bar{d}$ the E_T^{MISS} spectrum is slightly harder than the $\chi_1^0 \rightarrow eud$ scenario but is still not likely to be a useful approach in itself for suppressing backgrounds.

The other changes are shown in figures 6.4 and 6.5 which compare the jet and lepton multiplicity for $R_P - MSSM$ and $\not{R}_P - MSSM$ with the LSP decaying to electrons. Clearly there is an increase in both jet and lepton multiplicity in this scenario. The increased multiplicity of hard leptons coming from the two χ_1^0 decays makes for a clear approach for suppressing SM backgrounds.

In the case where the χ_1^0 decays to a neutrino this signature is degraded and this approach will have a much lower acceptance. The lepton multiplicity in this case is shown later in figure

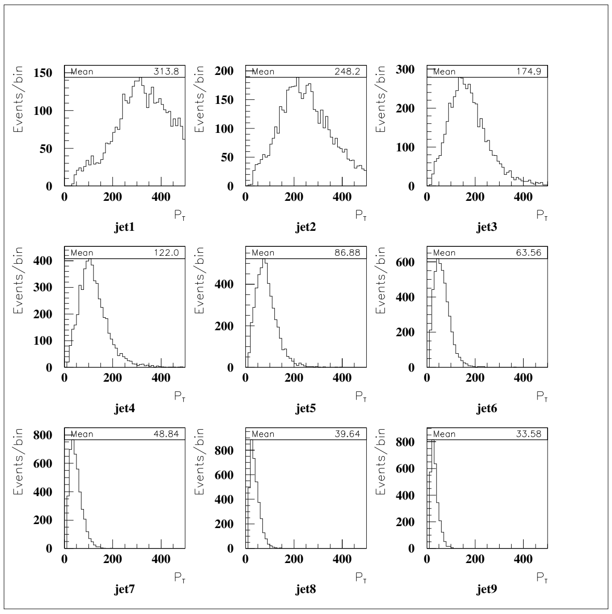

Figure 6.2: P_T spectrum of 9 leading jets where $\chi_1^0 \rightarrow eud$.

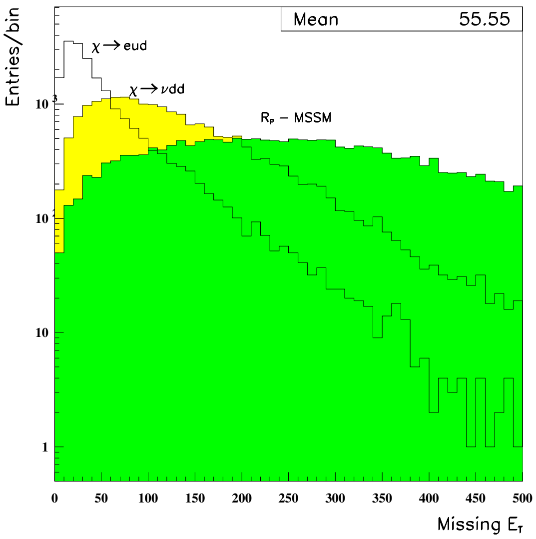

Figure 6.3: Effect of R_P violation on E_T^{MISS} spectrum.

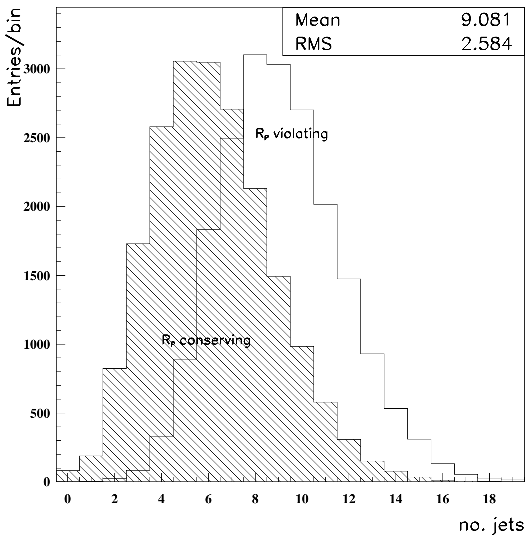

Figure 6.4: Effect of R_P violation on jet multiplicity where $\chi_1^0 \to eud$. Jet multiplicity measured as number of jets above 15 GeV.

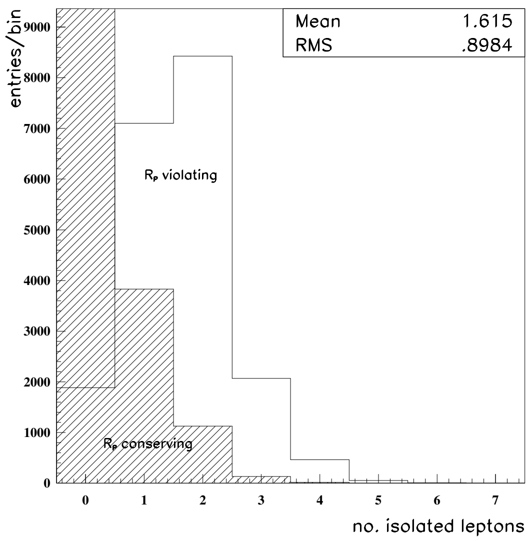

Figure 6.5: Effect of R_P violation on lepton multiplicity where $\chi_1^0 \to eud$. Lepton multiplicity measured as number of leptons (e, μ) above 5 GeV.

6.10. This fact combined with the lower E_T^{MISS} spectrum makes this scenario particularly challenging for ATLAS.

6.4 Sample Generation and Detector Simulation

Event samples were generated using PYTHIA 5.7 [26] for SM processes, and PYTHIA 6.1 [31] for SUSY processes. 4-vectors generated by PYTHIA were then run through ATLFAST 2.0 (described later) to simulate the performance of the ATLAS detector. The generated ntuple was then processed with ATLFAST-B which includes parameterisations of b-tagging efficiency, jet recalibration etc.

All results are quoted for $10^4 pb^{-1}$ which is the expected integrated luminosity after the first year of LHC running at low luminosity. It is impractical to generate the huge numbers of events expected after this time so smaller samples were generated and scale factors applied to obtain the correct signal-background ratio.

Table 6.5 summarises the samples generated. The cross-section for SUSY production at point 5 is relatively small (24 pb calculated by PYTHIA), and so it is not sufficient to generate a minimum bias sample and use this as background. Samples of events were generated that were thought likely to be significant backgrounds. In the case of $t\bar{t}$, $W + jet$, and $Z + jet$ production, gaining sufficient statistics proved difficult. To overcome this the decay of the relevant particle was altered to decay with a 100% branching ratio into leptons. In the cases where a lepton cut was used this channel was the dominant source of background, and so the precision of results that include a lepton cut could be improved by several orders of magnitude using this method.

In order to obtain sufficient QCD statistics a lower limit was placed on the 4-momentum exchange of the initial parton-parton collision; this can be done in PYTHIA by setting the *minimum hard scatter* parameter, ${P_T}^{2-2}$, (see chapter 4 for a discussion of this). High statistics were only accumulated for QCD events with ${P_T}^{2-2} > 300$ GeV. To ensure that this would not bias results by underestimating the size of the QCD background smaller

Process	**Generated Events**	**Scale Factor**
SUSY $\chi_1^0 \rightarrow eud$	200 000	1.185
SUSY $\chi_1^0 \rightarrow \nu_e d\bar{d}$	200 000	1.185
QCD ${P_T}^{2-2}$ >300 GeV	500 000	230
W+jet (leptonic decays)	500 000	158
Z+jet (leptonic decays)	100 000	143
$t\bar{t}$ (leptonic decays)	20 000	3.5
W+jet (all decays)	300 000	2444
Z+jet (all decays)	200 000	1408
$t\bar{t}$ (all decays)	40 000	147
WZ pairs (all decays)	20 000	45

Table 6.5: Event samples generated. Scale factor indicates the necessary scaling required for the sample to represent $10^4 pb^{-1}$ of data.

QCD samples were generated in a range of ${P_T}^{2-2}$ bins: 200 GeV $\rightarrow$ 250 GeV, 250 GeV $\rightarrow$ 300 GeV, 300 GeV $\rightarrow$ 350 GeV, 350 GeV $\rightarrow$ 400 GeV, and 400 GeV $\rightarrow$ 450 GeV.

A minimal set of cuts were applied to these samples: $P_T(j_1)$>300 GeV, $P_T(j_7)$>20 GeV, and m_{Tcent}>1000 GeV. This is a considerably looser event selection than anything used in subsequent analyses. The m_{Tcent} parameter represents the total central transverse energy in an event, and is defined in section 6.5.

Figure 6.6 shows the number of events expected to pass these cuts after $10^4 pb^{-1}$ of data in each ${P_T}^{2-2}$ bin. Only QCD events that have a high minimum hard scatter produce the multiple hard jets required to pass event selection. For minimum hard scatters below 300 GeV the number of events passing the cuts drops sharply, indicating that these events would not significantly contribute to the predicted background rate. Since the following inclusive analysis relies on tighter cuts the influence of these events will be even less, and so it was concluded that accumulating high statistics only for events with a minimum hard

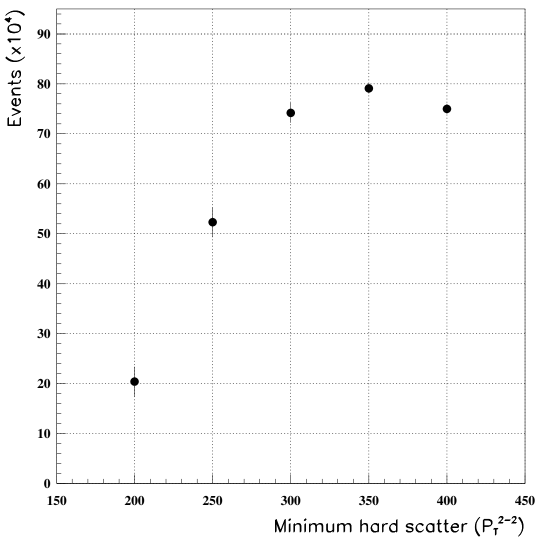

Figure 6.6: Number of events expected after $10^4 pb^{-1}$ as a function of minimum hard scatter that pass cuts: $P_T(j_1)$>300 GeV, $P_T(j_7)$>20 GeV, and m$_{Tcent}$>1000 GeV.

scatter above 300 GeV was sufficient to measure the QCD backgrounds in an unbiased way.

Note that in the charged lepton analysis the m$_{Tcent}$ cut is not explicitly used. However, as will be seen later, the SUSY signal is always centred at higher m$_{Tcent}$ and so the addition of this cut would not significantly affect the expected signal-background calculation.

6.4.1 Detector Simulation

Detector simulation is performed in two packages: ATLFAST 2.0, and ATLFAST-B. These packages describe themselves as a particle level simulation; this being an intermediate step between a very simple parton level or 4-vector analysis, and a complicated and CPU intensive full simulation such as GEANT.

ATLFAST uses a set of parameterisations to reproduce the mass resolutions, lepton and photon isolation efficiencies, jet reconstruction and E_T^{MISS} resolution expected from full simulation. Results from ATLFAST simulations show good agreement with full GEANT simulation [54]. For a detailed description of the parameterisations used by ATLFAST see reference [53]. ATLFAST was used with its default parameterisations; summarised below are some of the basic definitions used in the following analysis.

6.4.1.1 Calorimetric Clusters

Calorimetric clusters are defined for use in subsequent definitions. The energies of all undecayed particles (except neutrinos, muons, and the SUSY LSP), are first projected into calorimeter cells of granularity 0.1×0.1 for $|\eta| < 3$, and 0.2×0.2 for $|\eta| > 3$. All cells with $E_T > 1.5$ GeV are considered as potential initiators and considered in order of E_T. For each initiator considered all cells within a cone of radius=0.4 are summed. If this sum is above the minimum threshold (10GeV) the cluster is accepted and all cells associated with it removed from further consideration.

6.4.1.2 Isolated Leptons

Lepton candidates are searched for in the standard PYTHIA 4-vector table. Each lepton candidate is smeared according to a Gaussian energy resolution parameter. Lepton candidates must pass selection criteria of $P_T > 5$ GeV (6 GeV), and $|\eta| < 2.5$ for electrons (muons). For electrons the calorimetric cluster associated with each candidate is identified. The candidate is retained if it passes two isolation conditions: distance to the nearest neighbouring calorimetric cluster, $\Delta R > 0.4$, and $E_T < 10$ GeV in a cone of radius=0.2 around the lepton, where $\Delta R = \sqrt{(\eta_1 - \eta_2)^2 + (\phi_1 - \phi_2)^2}$. Note that the current version of ATLFAST does not attempt to parameterise lepton mis-identification.

6.4.1.3 Jet Reconstruction

Calorimetric clusters are used for jet reconstruction. All calorimetric clusters that are not already associated with isolated electrons are smeared according to a Gaussian energy resolution parameter. If the resulting cluster E_T is greater than 15 GeV it is labelled as a jet and removed from further consideration.

6.4.1.4 Flavour Tagging

In ATLFAST jets are labelled as b-jets if a b-quark of $P_T > 5$ GeV is found in a cone of radius = 0.2 around the reconstructed jet axis for jets with $|\eta| <2.5$. This initial definition is refined in ATLFAST-B, where the probability of mis-identifying a b-jet is parameterised in terms of the jet P_T.

6.5 Inclusive Search for Charged Lepton Decays

This section concentrates on the simpler case where the LSP decays to charged leptons: $\chi_1^0 \to eud$. The presence of at least 2 high E_T isolated leptons should make detection very easy. The following section looks at the more problematic case where $\chi_1^0 \to \nu_e d\bar{d}$.

The main role of an inclusive analysis is to demonstrate the presence of new physics, not described in the standard model, and to estimate the mass scale of this new physics.

In order to estimate the mass scale, a global variable is usually defined. In [51] this is defined as the sum of the visible transverse energy within a cut on rapidity of $|\eta| < 2$:

$$\mathrm{m}_{Tcent} = \sum_{|\eta|<2} {P_T}^{jet} + \sum_{|\eta|<2} {P_T}^{lep}$$

The value of m_{Tcent} at which SUSY processes become dominant can be used to estimate the mass scale of supersymmetry [42]. The distribution of m_{Tcent} is also centred at higher energies for SUSY processes than it is for SM processes and cuts on m_{Tcent} can be used to improve signal purity.

Inclusive analysis now proceeds in an attempt to find a set of cuts that suppress backgrounds sufficiently compared with the MSSM signal. The aim of this analysis is simply to show that there is an excess of events that exhibit the characteristic signatures of SUSY production.

6.5.1 Initial Jet Selection

An initial selection was made using the following cuts:

- at least 1 jet with $P_T > 300$ GeV
- at least 8 jets with $P_T > 20$ GeV
- circularity > 0.2, thrust < 0.9

These are very inclusive cuts designed to see if SM backgrounds can be suppressed solely with jet cuts. The cuts on circularity and thrust are very loose cuts on the overall event shape, designed to preferentially select events that are less collimated.

For events passing these cuts, m_{Tcent} was calculated and is shown in figure 6.7. These very simple cuts eliminate most sources of background, but unsurprisingly QCD is by far the dominant background, and is many times larger than the signal. Table 6.6 summarises the contributions from various sources of background for a range of jet cuts. Where no SM events passed the cuts, the signal-background is shown to 95% confidence limits.

Requiring very high jet multiplicities is a good way of eliminating most of the dominant backgrounds; for example requiring 7 jets $>$ 20 GeV instead of 6 jets $>$ 20 GeV reduces the signal by $\approx$8%, but cuts the background due to top production by 26%. Requiring that the 7th jet has a $P_T > 40$ GeV further reduces the top contribution by 50% and the QCD contribution by 73%, while reducing the signal by only 27%.

Even with the tightest multi-jet cut, leaving only 14% of the original SUSY signal, the QCD background is still $\approx$8 times larger than the signal. The additional requirement of at least 1 lepton is therefore needed to improve the signal-background ratio.

6.5.2 Lepton Multiplicity

Events with both leptons and multiple jets are a typical and robust SUSY signature. Previous studies have shown that multiple lepton signals have the best discovery reach in SUGRA models [55].

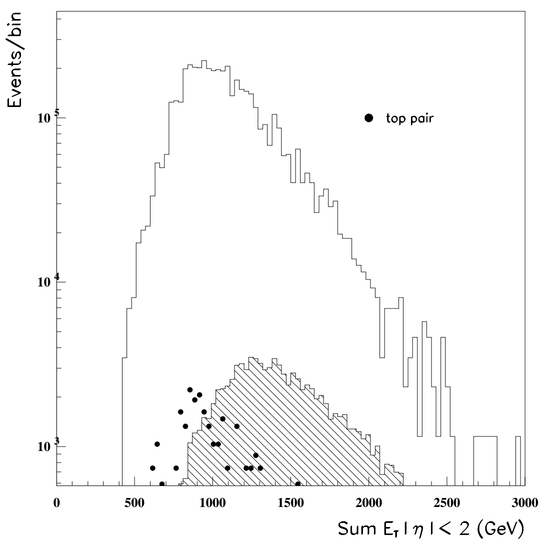

Figure 6.7: Distribution of m_{Tcent} expected for $10^4 pb^{-1}$ for SUSY (hatched), QCD (clear) and top (dots) events passing initial jet cuts.

Cut	SUSY	QCD	W+jet	Z+jet	$t\bar{t}$	WZpair	S/B
$j_7 > 20$	120000	8680000	34200	19700	45500	364	0.014
$j_7 > 40$	86200	2290000	7330	4230	20400	136	0.037
$j_8 > 20$	101000	4720000	19600	8450	30000	227	0.021
$j_8 > 40$	57900	842000	2440	<4220	7820	<135	0.068
$j_9 > 20$	79900	2350000	12200	4230	16400	136	0.034
$j_9 > 40$	34000	277000	2440	<4220	2660	<135	0.12

Table 6.6: Number of events expected for $10^4 pb^{-1}$ for SUSY and backgrounds passing a range of jet cuts. Additional cuts used were: $P_T(j_1) > 300$ GeV, circularity > 0.2, and thrust < 0.9.

Figure 6.8 shows the lepton multiplicity distributions for events that pass the same initial jet cuts listed in the previous section.

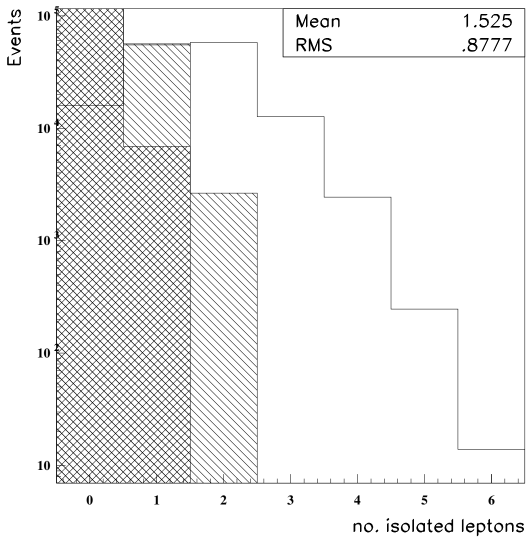

Figure 6.8: Lepton multiplicity for $10^4 pb^{-1}$ for events passing initial jet selection for SUSY (clear), $t\bar{t}$ (hatched), and QCD (cross hatched). Note that the QCD sample has been clipped to emphasise the key features of this plot.

Several combinations of cuts using jets and leptons were used. Figure 6.9 shows the m_{Tcent} distribution for events passing the cuts:

- at least 1 jet with $P_T > 300$ GeV
- at least 7 jets with $P_T > 20$ GeV
- at least 1 lepton with $P_T > 15$ GeV
- circularity > 0.2, thrust < 0.9

The number of events expected for a variety of different cut combinations is shown in table 6.7. QCD is still by far the dominant background, but has been greatly reduced.

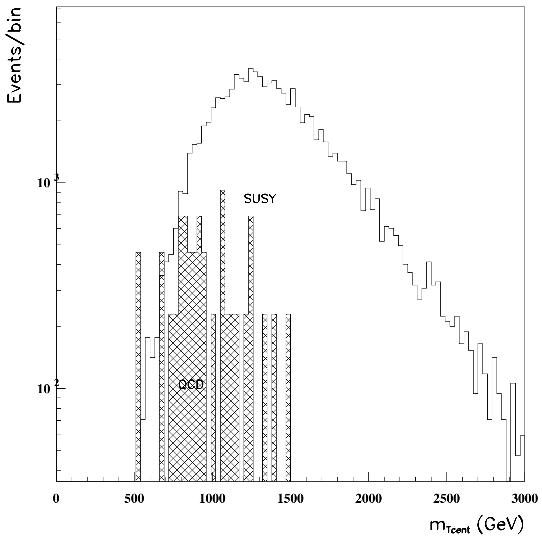

Figure 6.9: Distribution of m$_{Tcent}$ expected for $10^4 pb^{-1}$ for SUSY (clear) and QCD (cross-hatched) events passing cuts: $P_T(j_1) > 300$ GeV, $P_T(j_7) > 20$ GeV, $P_T(l_1) > 15$ GeV, circularity > 0.2, and thrust < 0.9

Cut	SUSY	QCD	W+jet	Z+jet	$t\bar{t}$	WZpair	S/B
j_7 >20 GeV	99900	8280	951	<429	94	46	>10.2
j_7 >40 GeV	70100	2990	317	<429	24	46	>18.4
j_8 >20 GeV	83900	4600	951	<429	62	46	>13.8
j_8 >40 GeV	46100	690	158	<429	10	<135	>32.4

Table 6.7: Number of events expected for $10^4 pb^{-1}$ for SUSY and backgrounds passing a range of jet cuts. Additional cuts used were: $P_T(j_1) > 300$ GeV, $P_T(l_1) > 15$ GeV, circularity > 0.2, and thrust < 0.9.

It is easy to see that S/B levels better than 10 (at 95% c.l.) are easily attainable. Furthermore tightening the lepton cut further reduces the background and provides protection from lepton mis-identification. Table 6.8 summarises the effect of tightening the lepton cuts in the case where the 7th jet is required to have > 40 GeV.

Cut	SUSY	QCD	W+jet	Z+jet	$t\bar{t}$	WZpair	S/B
$l_1 > 15$ GeV	70100	2990	317	<429	24	45	>18.4
$l_1 > 50$ GeV	48600	1380	317	-	17	45	
$l_2 > 15$ GeV	39000	460	-	-	-	-	
$l_2 > 50$ GeV	30000	230	-	-	-	-	
$l_3 > 15$ GeV	6480	-	-	-	-	-	
$l_3 > 50$ GeV	5470	-	-	-	-	-	

Table 6.8: Number of events expected for $10^4 pb^{-1}$ for SUSY and backgrounds passing a range of lepton cuts. Additional cuts used were: $P_T(j_1) > 300$ GeV, $P_T(j_7) > 40$ GeV, circularity > 0.2, and thrust < 0.9.

Requiring 2 leptons selects only around 10 - 15% of the original sample, but by lowering the cut on the 1st jet to 250 GeV this can be increased to around 15 - 20%, further increasing the significance of the signal, with very little effect on the backgrounds. Requiring that the leading jet has a $P_T > 250$ GeV improves the significance of the signal when the other cuts are quite tight (say, 8 or more jets, and 2 leptons). Where the other cuts are looser, allowing the leading jet to have a lower energy lowers the signal-background ratio since the number of background events rises more rapidly.

These results show that if this scenario exists it should be an easy matter to isolate a signal of this kind even using very simple cuts based on jets and leptons.

6.6 Neutral Lepton Decays

A more challenging channel is the case where the χ_1^0 decays via $\chi_1^0 \rightarrow \nu_e d\bar{d}$. In this case the hard lepton cuts used in the previous section will have a much lower acceptance and the signal-background ratio will be degraded. In the case of $R_P - MSSM$ E_T^{MISS} is used to suppress QCD background, but the E_T^{MISS} spectrum of the escaping ν will be much softer than from an escaping χ_1^0. However, as with charged lepton decays $\approx$4 more jets per event are expected compared with $R_P - MSSM$ and this can be used in combination with a range of other cuts to improve the signal-background ratio.

As before an initial selection is made on the basis of jets:

- at least 1 jet with $P_T > 300$ GeV
- at least 8 jets with $P_T > 20$ GeV
- circularity > 0.2, and thrust < 0.9

After these cuts the number of signal events passing is comparable with the case when the χ_1^0 decays to an electron. 41% of the signal passes, compared with only 4% of the QCD background. However, the cross-section for QCD is so large that the signal is still swamped, and the S/B is never better than 0.1 for all jet cuts tried.

Figures 6.10 and 6.11 show the lepton multiplicities and missing E_T spectra for events that pass these cuts.

Leptons are still present from squark decays into charginos (which subsequently emit a W) and squark decays into neutralinos (with subsequent decay into $\tilde{l}\, l$). This can still be used to suppress backgrounds, but the efficiency for the signal will be quite small. The E_T^{MISS} spectrum is not as well separated as the $R_P - MSSM$ scenario but will be investigated as a potential cut that could improve signal-background.

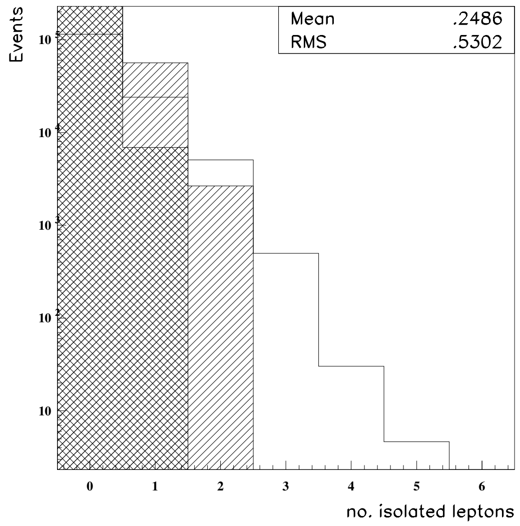

Figure 6.10: Lepton multiplicity for $10^4 pb^{-1}$ for events passing initial jet cuts, where $\chi_1^0 \to \nu_e d\bar{d}$. Key: SUSY (clear), top pair (hatched), and QCD (cross-hatched). Note that the QCD sample has been clipped to emphasise the key features of this plot.

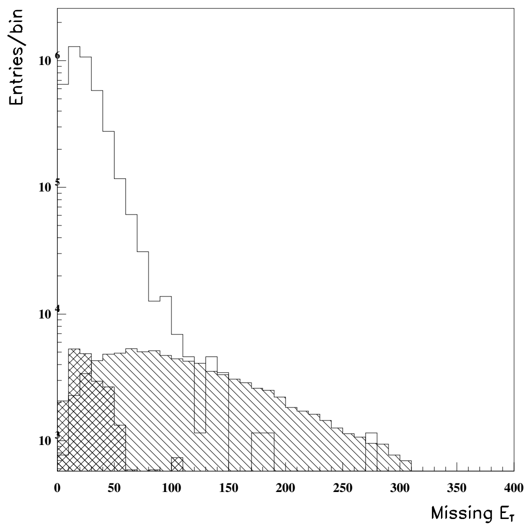

Figure 6.11: E_T^{MISS} distribution for $10^4 pb^{-1}$ for events passing initial jet cuts, where $\chi_1^0 \to \nu_e d\bar{d}$. Key: SUSY (hatched), top pair (cross-hatched), and QCD (clear).

6.6.1 Lepton Multiplicity

Many combinations of jet and lepton cuts were tried, and the results are summarised in table 6.9. For all combinations the same cuts on thrust and sphericity were used, and the leading jet was required to have a P_T above 300 GeV.

Since these cuts are the same as the charged lepton case, the backgrounds are identical. However, only around 5% of the signal is being selected so the signal-background ratio is greatly reduced.

However results still indicate that a significant signal could be extracted relatively easily within one year of low luminosity running. Furthermore the shape of the m_{Tcent} distribution itself can be used to enhance further the signal. Figure 6.12 shows the m_{Tcent} distribution for events passing cuts:

- at least 1 jet P_T>300 GeV
- at least 7 jets P_T>20 GeV
- at least 1 lepton P_T>15 GeV
- circularity > 0.2, and thrust < 0.9

All SM backgrounds are concentrated with lower m_{Tcent} and the SUSY signal is clearly dominant at high m_{Tcent}. An m_{Tcent} distribution like this would easily determine the existence of new physics. Placing an m_{Tcent} cut at 1200 GeV gives a signal-background ratio of ≈ 5.2 to 95% confidence limits.

6.6.2 Missing E_T Cuts

Although a lepton cut offers the most effective way of eliminating the dominant QCD background its poor efficiency means it is not ideal for an inclusive study. An alternative would be to use missing E_T (see figure 6.11).

Cut	SUSY	QCD	W+jet	Z+jet	$t\bar{t}$	WZpair	S/B
j_7 >20 GeV	18200	8280	951	<429	94	46	>1.9
j_7 >40 GeV	12300	2990	317	<429	28	46	>3.2
j_8 >20 GeV	15500	4600	951	<429	62	46	>2.6
j_8 >40 GeV	7850	690	158	<429	10	<135	>5.5

Table 6.9: Number of events expected for $10^4 pb^{-1}$ for SUSY and backgrounds passing a range of jet cuts. Additional cuts used were: $P_T(j_1) > 300$ GeV, $P_T(l_1) > 15$ GeV, circularity > 0.2, and thrust < 0.9.

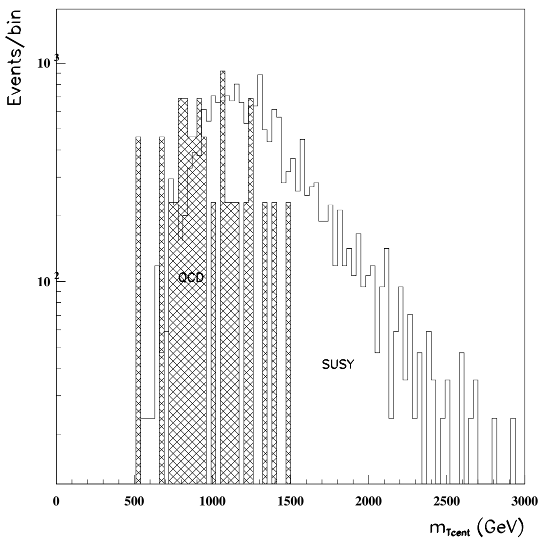

Figure 6.12: Distribution of m$_{Tcent}$ expected for $10^4 pb^{-1}$ for SUSY (clear) and QCD (cross-hatched) events passing cuts: $P_T(j_1) > 300$ GeV, $P_T(j_7) > 20$ GeV, $P_T(l_1) > 15$ GeV, circularity > 0.2, and thrust < 0.9.

A selection of jet and missing E_T cuts were tried. Figure 6.13 shows the m_{Tcent} distribution for the cuts:

- at least 1 jet $P_T > 300$ GeV
- at least 8 jets $P_T > 40$ GeV
- $E_T^{MISS} > 80$ GeV
- circularity > 0.2, and thrust < 0.9

Table 6.10 summarises background contributions for a range of missing E_T cuts. Clearly QCD is still the dominant background and is several times larger than the signal for all except the very hardest E_T^{MISS} cuts.

It appears that a very hard E_T^{MISS} cut could be used to extract a signal although its significance is not large. The m_{Tcent} distribution in figure 6.13 shows no separation between SUSY and background events and so would not help to improve the signal-background ratio, or confirm the presence of SUSY processes.

Several other combinations of cuts were tried in an attempt to improve the signal-background ratio. The following combinations are looked at in the next sections:

- jet + E_T^{MISS} + b-jets
- jet + E_T^{MISS} + 1 lepton

6.6.2.1 B-jet cut

Another feature of SUSY events is an abundance of heavy flavour quarks. These are created in gluino and squark decays. Figure 6.14 shows the number of b-jets tagged in events that passed the cuts:

- at least 1 jet > 300 GeV
- at least 8 jets > 40 GeV

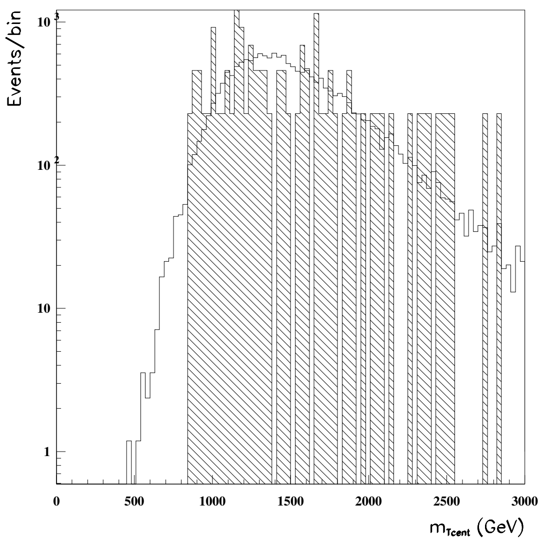

Figure 6.13: Distribution of m_{Tcent} expected for $10^4 pb^{-1}$ for SUSY (clear), and QCD (hatched) events passing cuts: $P_T(j_1) > 300$ GeV, $P_T(j_8) > 40$ GeV, E_T^{MISS}>80 GeV, circularity > 0.2, and thrust < 0.9.

Cut	SUSY	QCD	W+jet	Z+jet	$t\bar{t}$	WZpair	S/B
E_T^{MISS}>0 GeV	53400	823000	2440	-	7970	-	>0.06
E_T^{MISS}>30 GeV	49500	267000	-	-	3840	-	>0.18
E_T^{MISS}>50 GeV	44100	70800	-	-	1030	-	>0.61
E_T^{MISS}>80 GeV	35000	18000	-	-	590	-	>1.89

Table 6.10: Number of events expected for $10^4 pb^{-1}$ for SUSY and backgrounds passing a range of E_T^{MISS} cuts. Additional cuts used were: $P_T(j_1) > 300$ GeV, $P_T(j_8) > 40$ GeV, circularity > 0.2, and thrust < 0.9.

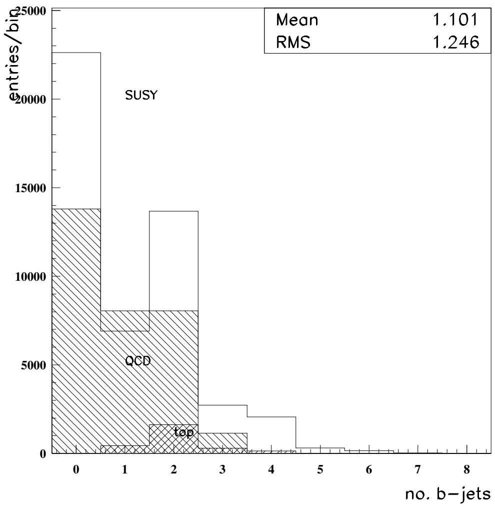

Figure 6.14: B-jet multiplicity for $10^4 pb^{-1}$ events passing cuts: $P_T(j_1)$>300 GeV, $P_T(j_8)$>40 GeV, $E_T^{MISS} > 80$ GeV, circularity > 0.2, and thrust < 0.9. Key: SUSY (clear), QCD (hatched), and $t\bar{t}$ (cross-hatched).

- $E_T^{MISS} > 80$ GeV
- circularity > 0.2, and thrust < 0.9

Like the lepton cut this will be quite inefficient, but could help to suppress QCD backgrounds. Figure 6.15 shows the m$_{Tcent}$ distribution when requiring a single b-jet with P_T>50 GeV in addition to the above cuts. Table 6.11 summarises results for a range of b-jet cuts.

As with the lepton cut the efficiency of these cuts is fairly low with between 8 and 1% of the signal passing, but the dominant QCD background is well suppressed. The poor separation between the m$_{Tcent}$ distributions indicates that an m$_{Tcent}$ cut would do little to improve the significance of the signal.

Although the signal significance is slightly less compared with the lepton cut, this is with b-tagging efficiency folded in. In the case of the lepton cut, no lepton mis-identification is included potentially giving a slightly enhanced signal-background estimate.

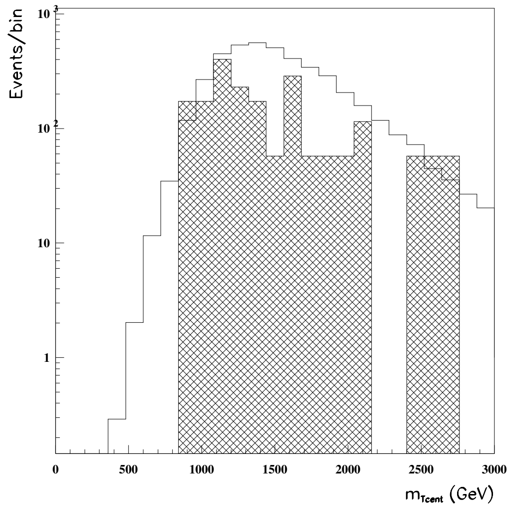

Figure 6.15: Distribution of m$_{Tcent}$ expected for $10^4 pb^{-1}$ for SUSY (clear), and QCD (hatched) events passing cuts: $P_T(j_1)$ > 300 GeV, $P_T(j_8)$ > 40 GeV, E_T^{MISS}>80 GeV, $P_T(j_1^{b-jet})$ > 50 GeV, circularity > 0.2, and thrust < 0.9.

Cut	SUSY	QCD	W+jet	Z+jet	$t\bar{t}$	WZpair	S/B
j_1^{b-jet} >15 GeV	19130	9430	-	-	590	-	>1.9
j_1^{b-jet} >50 GeV	17100	8510	-	-	590	-	>1.9
j_2^{b-jet} >15 GeV	8540	4830	-	-	148	-	>1.7
j_2^{b-jet} >50 GeV	6450	4370	-	-	148	-	>1.4
j_3^{b-jet} >15 GeV	2730	1150	-	-	148	-	>2.1
j_3^{b-jet} >50 GeV	1750	690	-	-	148	-	>2.5

Table 6.11: Number of events expected for $10^4 pb^{-1}$ for SUSY and backgrounds passing a range of b-jet cuts. Additional cuts used were: $P_T(j_1)$ > 300 GeV, $P_T(j_8)$ > 40 GeV, E_T^{MISS}>80 GeV, circularity > 0.2, and thrust < 0.9.

6.6.2.2 Lepton Cut

These results suggest that a lepton cut is still the most effective way of extracting a signal, even in the case when the χ_1^0 decays to a neutrino. The significance of the signal can be improved however by combining a lepton cut with a missing E_T cut. A range of jet and E_T^{MISS} cuts have been investigated, and combined with the requirement of a least 1 lepton with $P_T > 15$ GeV. For example consider the set of cuts:

- at least 1 jet with $P_T > 300$ GeV
- at least 7 jets with $P_T > 20$ GeV
- at least 1 lepton with $P_T > 15$ GeV
- circularity > 0.2, and thrust < 0.9

With no E_T^{MISS} cut 7% (18200 events) of the signal remains, and there are around 8000 QCD events. This corresponds to a S/B ≈ 1.9 to 95% confidence limits. Requiring $E_T^{MISS} >$ 80 GeV reduces the signal to around 5% (13200 events), but reduces the QCD background to just 3000 events; corresponding to a S/B ≈ 3.7 to 95% confidence limits.

Figure 6.16 shows the m_{Tcent} distribution produced when requiring $E_T^{MISS} > 80$ GeV in addition to the above cuts. Table 6.12 summarises the results for a range of jet and E_T^{MISS} cuts.

Although statistics are limited there is some indication that the inclusion of a lepton cut has led to an increased separation between the SUSY and QCD m_{Tcent} distributions, and hence that a cut on m_{Tcent} could be used to further increase the significance of the signal.

In conclusion, although the significance of the signal is obviously smaller it should still be possible to establish the presence of SUSY processes within the first year of LHC running, even if the LSP decays predominantly to neutral leptons.

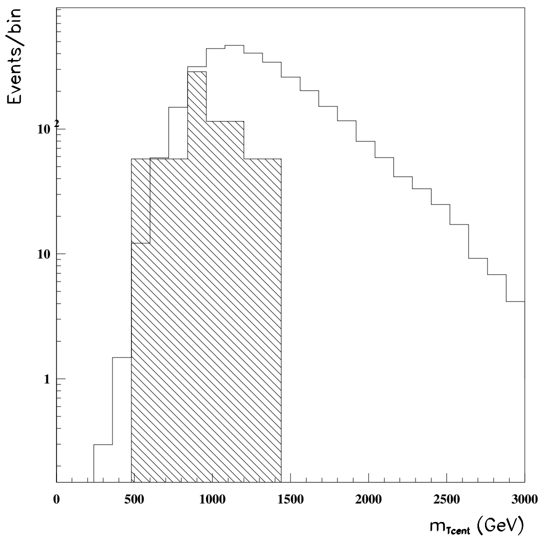

Figure 6.16: Distribution of m$_{Tcent}$ expected for $10^4 pb^{-1}$ for SUSY (clear), and QCD (hatched) events passing cuts: $P_T(j_1) > 300$ GeV, $P_T(j_7) > 20$ GeV, $P_T(l_1) > 15$ GeV, $E_T^{MISS} > 80$ GeV, circularity > 0.2, and thrust < 0.9.

Cut	SUSY	QCD	W+jet	$t\bar{t}$	S/B
$j_7 > 20$ GeV E_T^{MISS}>0 GeV	18200	8280	951	93	>1.9
$j_7 > 20$ GeV E_T^{MISS}>80 GeV	13200	2990	475	66	>3.7
$j_8 > 20$ GeV E_T^{MISS}>0 GeV	15500	4600	951	62	>2.7
$j_8 > 20$ GeV E_T^{MISS}>50 GeV	13300	2300	634	45	>4.4

Table 6.12: Number of events expected for $10^4 pb^{-1}$ for SUSY and backgrounds passing a range of jet and E_T^{MISS} cuts. Additional cuts used were: $P_T(j_1) > 300$ GeV, $P_T(l_1) > 15$ GeV, circularity > 0.2, and thrust < 0.9.

6.7 Other LHC-Points

Even if nature has chosen SUGRA as a model, there is a range of possible phenomenologies within this model that are dependent on the SUGRA parameters. As mentioned earlier, five so-called 'LHC-Points' in the SUGRA parameter space have been chosen for detailed study. These points were chosen to give a range of phenomenologies, and it is important that a search strategy can be found that will be able to extract a signal across all LHC-Points.

This section looks at inclusive searches for SUSY in the same L-violating $\not{R}_P - MSSM$ scenario studied previously, but for two other LHC points: Points 3, and 4. Points 1, 2, and 5, all have heavy sparticles, and so a search optimised for point 5 can be expected to work well for points 1, and 2. For this reason points 1, and 2 have not been investigated. Points 3, and 4 generate relatively light sparticles and it is important to determine whether a signal will still be visible in the case of L-violating $\not{R}_P - MSSM$.

It is worth noting here that although points 3 and 4 have lighter particles which will make the SUSY signal less significant, a new channel is opened up due to the different kinematics. At point 5, the channel $\chi_2^0 \rightarrow h^0\chi_1^0$ is open and the χ_2^0 decays predominantly via this channel. At points 3 and 4 this channel is kinematically forbidden, and the decay: $\chi_2^0 \rightarrow \tilde{l}\ l$ is dominant. Therefore one might expect a larger number of events to contain isolated leptons and the efficiency of lepton cuts to be improved.

6.7.1 Results

The following cuts were applied:

- at least 1 jet with $P_T > 300$ GeV
- at least 7 jets with $P_T > 40$ GeV
- at least 1 lepton with $P_T > 15$ GeV
- circularity > 0.2, and thrust < 0.9

Table 6.13 summarises the number of events passing these cuts for each of the LHC-Points considered. For these cuts 8280 QCD events are expected to pass.

LHC-Point	**SUSY**		**S/B**	
	$\chi_1^0 \to eud$	$\chi_1^0 \to \nu_e d\bar{d}$	$\chi_1^0 \to eud$	$\chi_1^0 \to \nu_e d\bar{d}$
Point 5	70100	12300	>18	>3.2
Point 4	72900	14500	>19	>3.9
Point 3	249000	90200	>65	>24

Table 6.13: Number of events expected for $10^4 pb^{-1}$ passing cuts for LHC points 3, 4, and 5. Cuts used were: $P_T(j_1) > 300$ GeV, $P_T(j_7) > 40$ GeV, $P_T(l_1) > 15$ GeV, circularity > 0.2, and thrust < 0.9.

There would be no difficulty extracting a signal using these cuts. The signal efficiency is slightly reduced at point 4 (by about 5% for charged leptons, and 2% for neutral leptons) but a higher cross-section means the overall signal-background ratio is actually slightly better. The same is also true at point 3; these cuts have very poor efficiency, selecting only 4% of charged lepton decays, and 1% of neutral lepton decays, but the cross-section is large enough that the overall signal-background is very high. It should be noted that the SUGRA parameters are not necessarily constrained to lie on any particular LHC-Point, so the true signal efficiency is likely to lie somewhere in between the points studied.

The low signal efficiency (particularly for point 3) suggests that it might be favourable to begin an initial search strategy with slightly looser cuts; but in any case it is clear that if this scenario exists a clear inclusive signal could be extracted across a wide range of parameter space. This is true even in the most pessimistic scenario where the χ_1^0 decays solely to neutrinos and quarks.

6.8 Precision Measurements

If a deviation from the standard model is observed, the next step will be to constrain the parameters of the SUSY model. This proceeds by extracting exclusive channels that can be separated from both SM background and other SUSY processes, then forming kinematical distributions sensitive to the masses of the sparticles at the bottom of the decay chain. The information about sparticles at the bottom of the decay chain can be used to reconstruct mass combinations of the particles further up the decay, and in this way the relative masses of all the sparticles can be determined.

In R_P-MSSM studies [42–46] some or all of the following combinations are usually reconstructed.

- Search for $h^0 \to b\bar{b}$ decay.
- Search for $\tilde{q}_L$ in $\tilde{q}_L \to qh\chi_1^0$.
- Search for $\tilde{q}_R$ in $\tilde{q}_L\tilde{q}_R$ production.
- Search for sleptons from $\chi_2^0 \to \tilde{l}_R l$.

If the mass of the LSP can be directly reconstructed, then this mass can be combined with particles from decays further up the chain to reconstruct the masses of the squarks and sleptons. In this case this information could be used to directly constrain the parameters of the model.

If the LSP mass cannot be directly reconstructed (for example in R_P-MSSM or if the LSP decays via invisible modes) then there is not enough kinematic information to directly reconstruct any sparticle masses, and mass combinations must be measured. This is the strategy usually used in R_P-MSSM and in most cases is sufficient to constrain the parameters of the model.

This section aims to determine whether the precision measurement methods developed for R_P-MSSM can be adapted and used in the case where the LSP decays into two jets and a lepton. In the case where the LSP decays to charged leptons the reconstruction method is most radically different. This is because the presence of extra hard leptons severely disrupt previously robust signatures, but also because the direct reconstruction of the χ_1^0 allows the masses of the other particles to be directly measured. The following sections show that by starting with the reconstructed χ_1^0 and working up, the masses of the sparticles in an entire decay chain can be reconstructed.

In the case where the LSP decays to neutral leptons, the only final state difference from R_P-MSSM will be the presence of extra jets, and a softened missing E_T spectrum. In this case analysis would proceed in much the same way as the R_P-MSSM scenario and so is not studied in detail here. However, later sections look at the effect of additional jet combinatorics on two of the most significant exclusive signatures.

6.9 χ_1^0 Mass Reconstruction

The most exciting possibility that $\not{R}_P-MSSM$ allows is direct reconstruction of the LSP. In the scenario where the LSP decays to two jets and a charged lepton reconstruction is relatively simple.

χ_1^0s will be present at the bottom of the decay cascade in all events, so cuts similar to the inclusive analysis above can be used. The dominant background will be combinatorial background from the SUSY events themselves so steps must be taken to reduce this in order to extract a clear signal.

6.9.1 Parton Level Information

Since the χ_1^0 is relatively heavy at point 5 (119 GeV/c^2), and the mass difference between the χ_1^0 and any of its parents is large, its decay products will be strongly Lorentz boosted and angular separation cuts are an obvious approach. In addition kinematic cuts on the P_T of the jets will further reduce combinatorial backgrounds.

These features have been checked at the generator level. Figure 6.17 shows the P_T distributions of the χ_1^0s generated. The harder of the two χ_1^0s has a mean P_T of 318 GeV, and the softer has a mean P_T of 168 GeV. This is reflected both in the P_T distributions of the χ_1^0 decay products and in their relative angular separation (figure 6.18), where angular separation $\Delta R = \sqrt{(\eta_1 - \eta_2)^2 + (\phi_1 - \phi_2)^2}$, provides a measure of the separation between different χ_1^0 decay products.

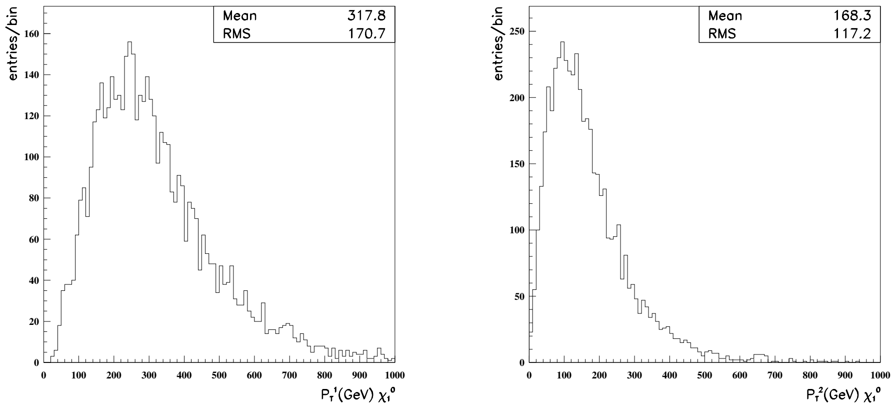

Figure 6.17: Distribution of P_Ts for χ_1^0s at generation level. Left hand plot shows harder of the two χ_1^0s, right hand plot shows the softer χ_1^0.

6.9.2 Reconstruction of χ_1^0 Decays

Reconstruction of the LSP mass proceeds as follows:

- Events are selected that pass a set of very loose inclusive cuts designed to improve the significance of the SUSY signal over SM backgrounds.
- All combinations of 2 jets and 1 lepton are considered, and their invariant mass calculated if they pass a set of tight kinematic and angular cuts designed to preferentially reconstruct the leading χ_1^0.

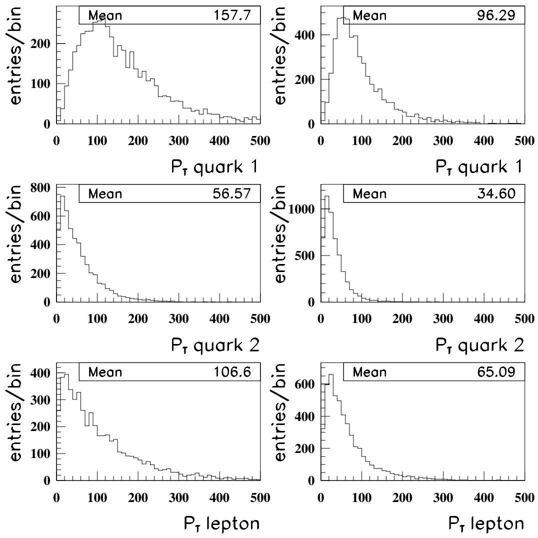

(a)

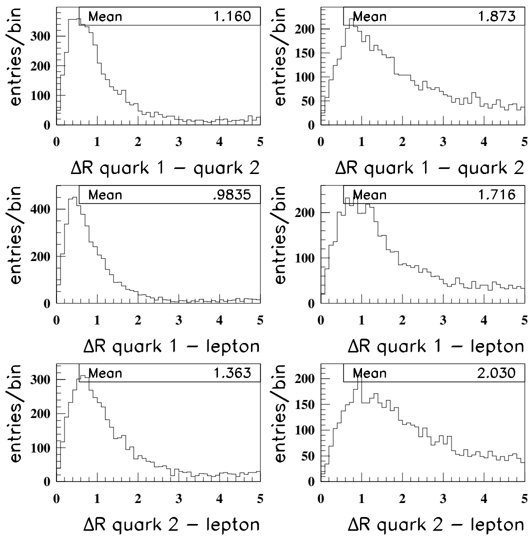

(b)

Figure 6.18: Distribution of P_T(a) and angular separation (b), for χ_1^0 decay products at generation level. The left hand columns show distributions for particles from leading χ_1^0 and the right hand column shows distributions for particles from the softer χ_1^0.

- If a set of 2 jets and 1 lepton pass the initial cuts, all *other* 2 jet + 1 lepton combinations in the event are considered to see it they pass looser cuts designed to reconstruct the decay of the softer χ_1^0.

This method minimises combinatorial background, and improves the reconstruction efficiency for the softer χ_1^0. The cuts used for reconstruction are listed in table 6.14. The inclusive cuts used for initial event selection were:

- $P_T(j_1) > 250$ GeV
- $P_T(j_6) > 20$ GeV
- $P_T(l_1) > 20$ GeV
- circularity <0.9, thrust >0.2

	Leading χ_1^0 **cuts**	**Second** χ_1^{0}**' cuts**
Angular cuts	$\Delta R(j_1 j_2) < 1$	$\Delta R(j_1 j_2) < 2$
	$\Delta R(j_1 l_1) < 1$	$\Delta R(j_1 l_1) < 2$
	$\Delta R(j_2 l_1) < 1.5$	$\Delta R(j_2 l_1) < 2$
Kinematic cuts	$P_T(j_1)$>50	$P_T(j_1)$<100
	$P_T(j_2)$<100	$P_T(j_2)$<100

Table 6.14: Angular and kinematic cuts used to extract χ_1^0 mass peak. Left hand column refers to leading χ_1^0 and right hand column refers to the softer χ_1^0.

Figure 6.19 shows the jjl invariant mass spectrum for events that pass all these selections (hatched), and for events that pass only the inclusive and angular cuts (clear). On average there are 1.4 entries per event using both cuts, which means that the second χ_1^0 is rarely reconstructed. Nevertheless, a clear peak can be seen at the mass of the χ_1^0. The kinematic

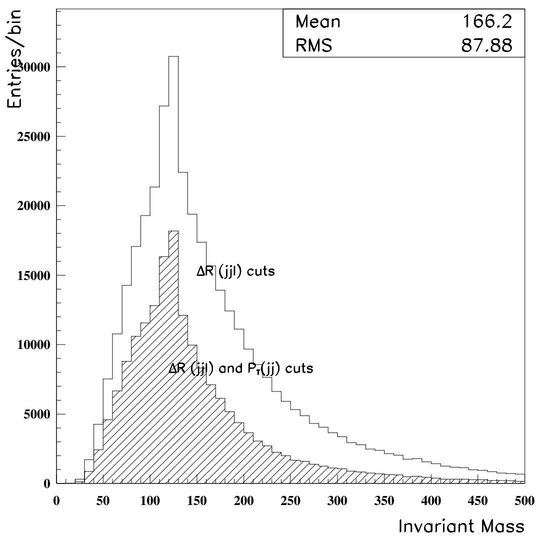

Figure 6.19: Invariant mass spectrum of jjl combinations passing initial inclusive, angular and kinematic cuts (hatched), and jjl combinations passing initial inclusive and angular cuts (clear). $10^4 pb^{-1}$ worth of events were simulated.

cuts on the reconstructed jjl combinations were designed to increase the discrimination between leading and second χ_1^0s, and to reduce the combinatoric contribution from jets produced directly in $\tilde{q}$ and $\tilde{g}$ decays. Although not clearly visible to the eye, the significance of the peak is slightly enhanced by requiring the kinematic cuts, although this is not a strong effect.

In order to quote an absolute accuracy for this measurement a detailed understanding of the backgrounds involved would usually be required; this would be premature in this analysis since the dominant backgrounds cannot be accurately predicted until real data is available. A rough estimate of the statistical accuracy of the measurement can be obtained though by varying the mass of the χ_1^0 and measuring the statistical sensitivity of the invariant mass spectrum to this movement. This method is used extensively in R_P-MSSM studies, and is detailed in the next section. Using this method with just 5000 signal events (equivalent to less than two weeks running at low luminosity) revealed that the mass of the χ_1^0 could be measured to be 119 ± 3 GeV. Since the systematic error is expected to be around 1 GeV [20] it is safe to conclude that even very early in LHC running the statistical error on this measurement would not be dominant.

Figure 6.20 shows the invariant mass spectrum for SUSY events without the initial event selection or kinematic cuts, and the same spectrum for QCD events.

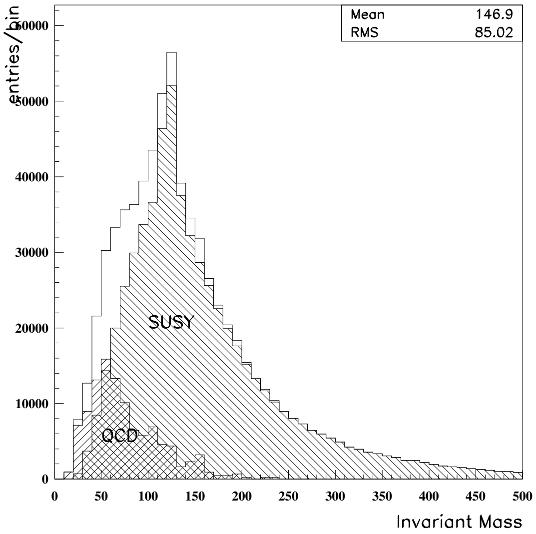

Figure 6.20: Invariant mass spectrum of jjl combinations passing angular cuts for SUSY (hatched), QCD (cross-hatched), and both (clear). Events have been normalised to $10^4 pb^{-1}$.

Even without an initial selection there are few QCD events that pass the reconstruction criteria. This does not however account for the part of the spectrum removed by the ${P_T}^{2-2}$ cut on the generated QCD events. However, from the shapes of the peaks it is still clear that the low P_T part of the QCD spectrum would not represent a significant background. Including an inclusive event selection reduced the generated QCD background to negligible levels, and would well suppress the ${P_T}^{2-2} < 300$ GeV part of the background.

Two cross-checks have been used to ensure that the observed peak is indeed due to reconstructed χ_1^0s. Firstly the mass of the χ_1^0 has been varied, both by hand and by looking at other LHC-Points below, and the invariant mass peak is seen to be sensitive to the χ_1^0 mass as expected. Secondly by using the same analysis on QCD events the absence of a reconstructed peak around the χ_1^0 mass indicates that the kinematic constraints are not artificially generating a peak.

6.9.3 Measurement of Experimental Sensitivity to χ_1^0 Signal

An estimate of the statistical uncertainty in the χ_1^0 mass measurement is obtained by varying the mass of the χ_1^0 while keeping all other masses frozen. Six samples were generated with the χ_1^0 mass varying from 110 GeV to 140 GeV. Figure 6.21 shows the invariant mass spectra for a range of χ_1^0 masses. The measurement uncertainty was calculated by testing the statistical sensitivity of the distribution to the mass of the χ_1^0 using a Kolmogorov test [56]. This technique has been used extensively in R_P-MSSM studies to obtain a fairly accurate estimate in the absence of detailed fitting procedures.

The Kolmogorov test returns a parameter (z-value) related to the probability that a given sample came from the same population as a control sample. Each mass point was compared with a reference sample at the nominal mass of the χ_1^0 of 119 GeV. Table 6.15 summarises the z-values obtained.

The z-value is related to the confidence limit α by the equation,

$$\alpha = 2\sum_{r=1}^{r=\infty}(-1)^{r-1}e^{2r^2z^2}.$$

Where α represents the probability that a sample taken from a given population will be fully contained within the confidence belt. Hence in order to quote 1σ errors assuming a Gaussian distribution the confidence belt corresponding to α=31.73% should be given. From the above equation this is equivalent to a rise in z-value of z=0.958.

Therefore the quoted error (confidence belt) is found by extrapolating from the z-values summarised above to find the confidence belt corresponding to a rise in z-value of 0.958. This yields errors at the 1σ level of ±3 GeV.

This method is also used later to estimate the accuracy of the slepton mass measurement.

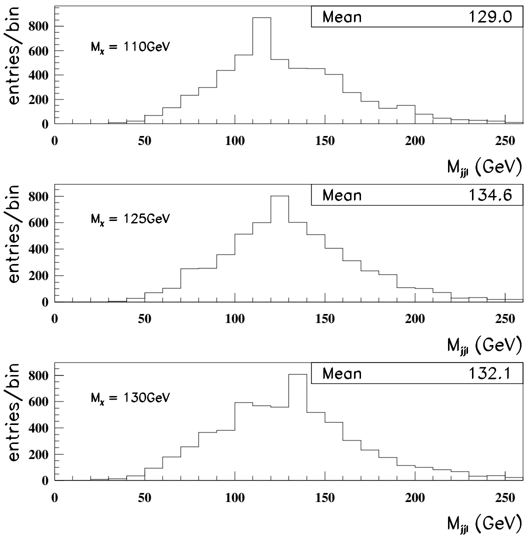

Figure 6.21: Invariant mass spectrum for jjl combinations reconstructed for 3 values of χ_1^0 mass: $M_{\chi_1^0} = 110$ GeV (top), $M_{\chi_1^0} = 125$ GeV (middle), and $M_{\chi_1^0} = 130$ GeV (bottom).

χ_1^0 mass	100 GeV	110 GeV	115 GeV	120 GeV	125 GeV	130 GeV	140 GeV
z-value	7.76	4.68	1.93	1.33	2.73	3.65	6.02

Table 6.15: Summary of χ_1^0 mass points simulated and their corresponding z-value when compared with a reference sample with the χ_1^0 at its nominal value of 119 GeV. The lower the z-value the higher the probability that the given sample came from the same population as the reference sample.

6.10 Other LHC-Points

As with the inclusive analysis the cuts used in the χ_1^0 reconstruction rely to some extent on the kinematics of the events. As such it is prudent to study the effect of some of the other LHC-Points. As before, the same analysis was run for LHC-Points 3 and 4. The lighter particles mean that the χ_1^0s will be softer, and their decay products less well collimated. As a result the efficiency will be lower and the reconstructed peak degraded.

Figure 6.22 shows the jjl invariant mass spectra for point 3 (a), and point 4 (b) using the same angular and kinematics cuts described in table 6.14 . A peak, albeit badly degraded, is still visible for point 4; but for point 3 the peak is completely lost.

At point 3 the mass of the LSP is 44 GeV and so its decay leads to very soft jets and leptons. Even with modified cuts it is unlikely that the mass of the LSP could be directly reconstructed from combinations of such soft jets. From the inclusive analysis previously it is clear that SUSY processes at point 3 can be discovered. However, precision measurements rely either on certain characteristic signatures (eg. l^+l^-, $h \to bb$), or on direct reconstruction of the χ_1^0 mass. As can be seen in the following sections reconstruction of both l^+l^- and $h \to bb$ is difficult due to the additional jet and lepton combinatorics from the χ_1^0 decay. This fact potentially makes point 3 with the χ_1^0 decaying to two jets and a lepton a particularly unfortunate, and challenging scenario for ATLAS.

For point 4 squarks are heavier than gluinos so gluino production is the dominant process. This produces much more complicated event topologies, with more jets in each event further increasing combinatorial backgrounds and degrading the peak. However, the neutralinos and the charginos are light compared with the squarks and gluinos so there is a significant contribution from electroweak sparticle production. The production process $pp \to \tilde{\chi}_i^{\pm}\tilde{\chi}_i^0$ has a cross-section of 3.52 pb and can lead to,

$$\begin{array}{ll} \chi_2^0 \to & \chi_1^0 \; l^+l^- \\ & \hookrightarrow jjl, \end{array}$$

$$\begin{array}{ll} \chi_i^{\pm} \to & \chi_1^0 l\nu \\ & \hookrightarrow jjl \end{array}$$

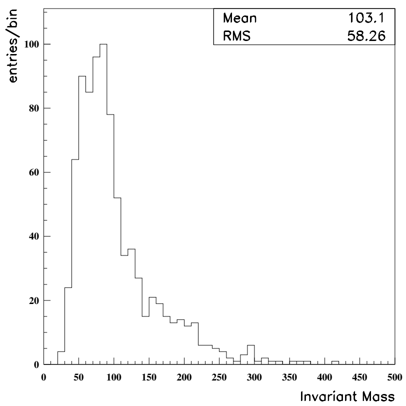

(a) LHC-Point 3

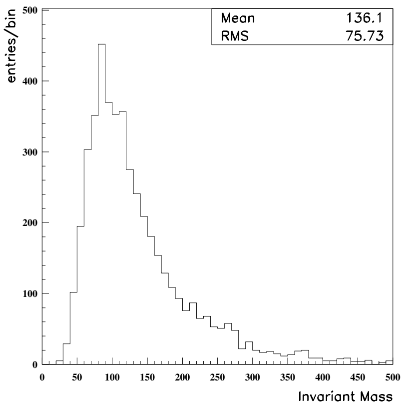

(b) LHC-Point 4

Figure 6.22: Invariant mass distributions with inclusive, angular, and kinematic cuts. For LHC-Point 3 (a), and LHC-Point 4 (b).

which has no combinatorial backgrounds from gluino or squark decay. The invariant mass spectrum using the initial selection,

- $P_T(j_5) > 15$ GeV
- $P_T(l_4) > 10$ GeV

is shown in figure 6.23. A peak can be observed at the nominal χ_1^0 mass for point 4 of 80 GeV, but its significance is not notably more than figure 6.22b.

6.11 Other Sparticle Masses

Once the mass of the χ_1^0 has been measured it is possible to combine this mass with other particles in an attempt to reconstruct sparticle masses from higher in the decay chain. As an example, consider the channels:

$$\begin{aligned} \tilde{q}_R \rightarrow \quad & \chi_1^0\ q \\ & \hookrightarrow jjl \end{aligned}$$

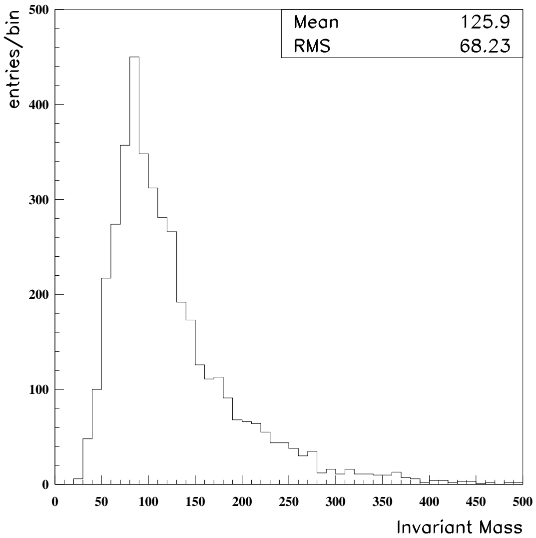

Figure 6.23: Invariant mass spectrum of jjl combinations passing angular and kinematic cuts, with an initial cut of: $P_T(j_5)$>20 GeV, and $P_T(l_4)$>15 GeV.

and

$$\begin{array}{llll}\tilde{q}_L \rightarrow & \chi_2^0\, q & & \\ & \hookrightarrow & \tilde{l}l & \\ & & \hookrightarrow & \chi_1^0\, l \\ & & & \hookrightarrow jjl\end{array}$$

Squarks are produced either directly, with a share of the total cross-section of around 35%, or via gluino decay. The $\tilde{q}_R$ decays with 99% branching fraction directly to a χ_1^0 and a quark. So combining the reconstructed χ_1^0 with an appropriate quark should reconstruct the squark mass. The $\tilde{q}_L$ has a branching fraction of 32% to $\chi_2^0\, q$ and the χ_2^0 decays to $\tilde{l}l$ with a probability of 4% for each flavour. So forming combinations of the χ_1^0 with various leptons should allow reconstruction of both the slepton mass and the χ_2^0 mass.

In addition the invariant mass spectrum of the two leptons produced in χ_2^0 and $\tilde{l}$ decays exhibits a characteristic edge related to the masses of the χ_1^0, χ_2^0 and $\tilde{l}$. This edge is investigated in the next section. Following sections investigate the potential of directly reconstructing the sparticle masses for this channel.

Background contributions come from both SM processes and from combinatorial background caused by the large number of final state objects. Throughout this investigation both sources are suppressed by assuming that a suitable measurement of the χ_1^0 mass has been possible, and using this information to require that events contain two fully reconstructed χ_1^0s. This requirement alone is sufficient to remove background contributions from SM processes, and combinatorial backgrounds can be suppressed by excluding from further reconstruction jets and leptons that are associated with one of the reconstructed χ_1^0s.

6.11.1 Search for $\tilde{l}$ in l^+l^- Mass Edge

A characteristic signature of the presence of SUSY particles is an edge in the invariant mass spectrum of pairs of opposite sign (OS), same flavour (SF) leptons. This is a robust and distinctive signature and as such is commonly used as part of an initial search strategy for SUSY processes. To investigate this, all signal events were used, and all combinations of OS SF dilepton pairs formed. Figure 6.24a shows this edge for R_P conserving SUSY where the LSP is stable. When the LSP decays to charged leptons however the presence of extra leptons severely degrades this signature (6.24b).

This situation can be improved somewhat by attempting to veto leptons coming from the LSP decay. Events were selected in the following way:

- Potential χ_1^0 combinations were formed as detailed in the previous section.
- Only events containing two reconstructed χ_1^0s with $|M_{jjl} - M_{\chi_1^0}| < 10$ GeV were considered.
- The OS SF dilepton invariant mass was calculated for all combinations of leptons *not* used in the reconstructed χ_1^0s.

The invariant mass spectrum produced is shown in figure 6.25. The efficiency for reconstructing the softer χ_1^0 is low and high statistics are needed. An edge is now just visible, but

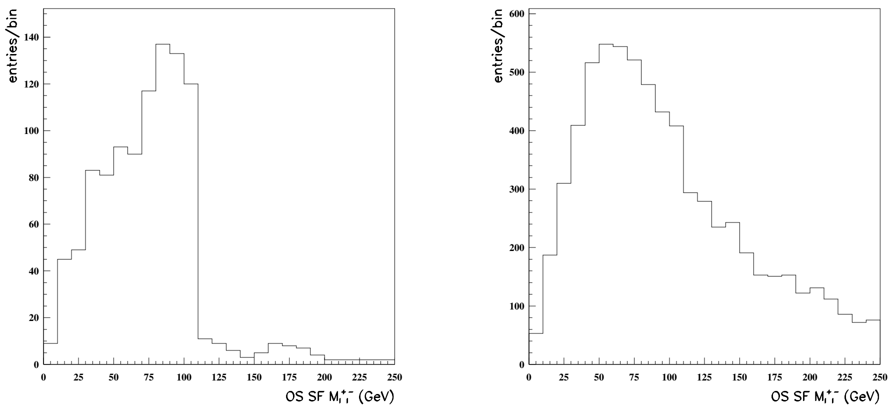

(a) Stable χ_1^0 (b) $\chi_1^0 \to eud$

Figure 6.24: OS, SF dilepton invariant mass spectrum for: (a) Stable LSP, and (b) LSP decaying to charged leptons.

it is still degraded compared with the R_P conserving case. On its own this signature is not likely to be considered very convincing.

6.11.2 Direct Reconstruction of $\tilde{l}$ and χ_2^0 Masses

The edge in the OS SF dilepton spectrum is not critical however since the $\tilde{l}$ and χ_2^0 masses can be directly measured by combining the remaining leptons in the event with the reconstructed χ_1^0.

A common source of sleptons is via the χ_2^0 decay:

$$\begin{array}{lll} \chi_2^0 \to & \tilde{l}_R l & \\ & \hookrightarrow & \chi_1^0 l \\ & & \hookrightarrow jjl \end{array}$$

In this case, the slepton mass can be reconstructed by combining the jjl combination from the χ_1^0 with one of the leptons. Events were selected with two reconstructed χ_1^0s with invariant mass $|M_{jjl} - M_{\chi_1^0}| < 10$ GeV, and at least 2 OS SF leptons in addition to the

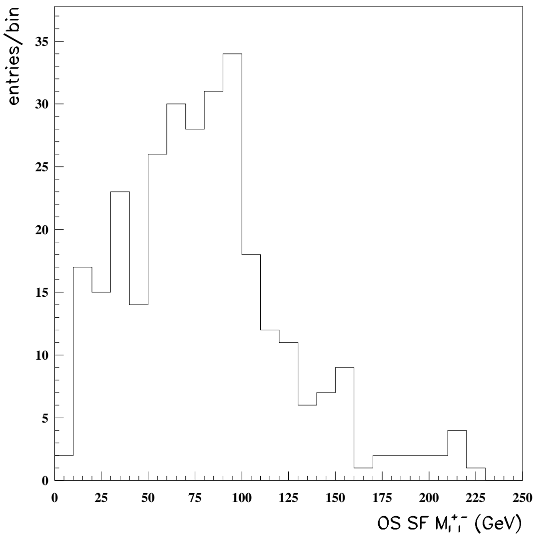

Figure 6.25: The OS SF dilepton invariant mass spectrum for $10^4 pb^{-1}$ when the LSP decays to charged leptons, using a veto on χ_1^0 leptons.

leptons associated with the χ_1^0 decays. Requiring χ_1^0 reconstruction is necessary so that χ_1^0 jets and leptons can be tagged as such and not contribute to combinatorial background; it is also an effective cut against standard model processes with none of the generated SM events passing this cut. The additional dilepton is required to minimise combinatorial backgrounds.

It is not known which χ_1^0 comes from the slepton decay and which comes directly from $q_R \to \chi_1^0 q$, or some other channel. It is assumed that the χ_1^0 coming directly from $\tilde{q}_R$ decay will have the greatest P_T, and so the soft χ_1^0 is associated with the slepton. This jjl combination is recalibrated using the constraint that the invariant mass should be equal to the nominal mass of the χ_1^0, and combined in turn with each of the remaining leptons. Due to the mass differences of the particles involved it is reasonable to expect that the lepton with the lowest P_T is associated with the χ_1^0. Efficiency is improved however if all combinations are formed and the one with the lowest invariant mass retained.

The distribution obtained for $10^4 pb^{-1}$ is shown in figure 6.26. A clear peak is observable around the nominal mass of the slepton of 157 GeV. As an additional check that this peak is indeed caused by the slepton decay the slepton mass was varied and all other mass kept the same. Figure 6.27 shows the effect of setting the slepton mass at 140 GeV and 170 GeV. As can be seen, the peak position is sensitive to the slepton mass; a Kolmogorov test on these results revealed that the mass could be measured to an accuracy of ± 15 GeV within the first year of running.

The mass of the χ_2^0 is reconstructed by combining the $jjl + l$ combination with the remaining lepton. An additional cut is made around the slepton peak of $|M_{jjll} - M_{\tilde{l}}| < 15$ GeV. In order to improve statistics the OS SF dilepton cut is no longer used. The invariant mass spectrum obtained is shown in figure 6.28. The distribution is shown for $10^4 pb^{-1}$ and statistics are clearly becoming limited. However, a clear peak is still observable around the nominal χ_2^0 mass of 228 GeV.

6.11.3 Reconstruction of $\tilde{q}_R$ Mass

To further constrain the parameters of the SUGRA model the mass of $\tilde{q}_R$ can be reconstructed by combining the harder of the reconstructed χ_1^0s with a hard jet. This attempts to reconstruct the decay $\tilde{q}_R \to \chi_1^0 q$.

Events were selected with the following cuts:

- 2 reconstructed χ_1^0s with $|M_{jjl} - M_{\chi_1^0}| < 10$ GeV
- 2 hard central jets with $|\eta| < 2$, $P_T(j_1) > 300$ GeV and $P_T(j_2) > 100$ GeV.

Two hard central jets are required in this analysis in an attempt to select preferentially events containing a $\tilde{q}_R$ and a $\tilde{q}_L$. The inclusion of this cut slightly increases the significance of the $\tilde{q}_R$ peak, but not dramatically.

In principle it is not known which jet and which χ_1^0 can be associated with the squark decay; however naively combining the leading jet with the hardest χ_1^0 provides a clear signal which can be seen in figure 6.29. This signal is clearly well separated from the purely

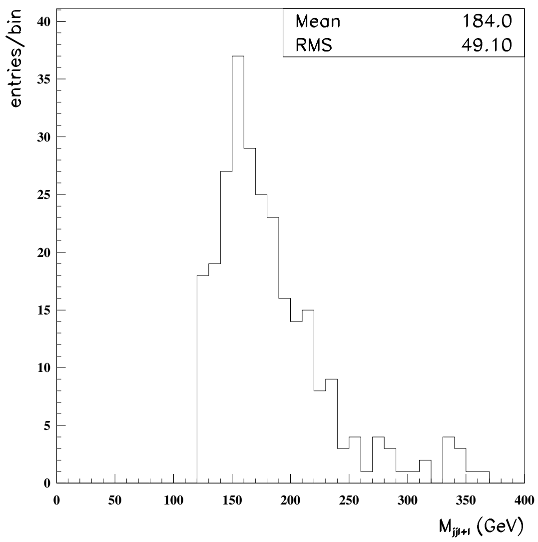

Figure 6.26: Invariant mass distribution for $10^4 pb^{-1}$ of the soft χ_1^0 candidate combined with a lepton. A peak is visible around the nominal $\tilde{l}$ mass of 157 GeV.

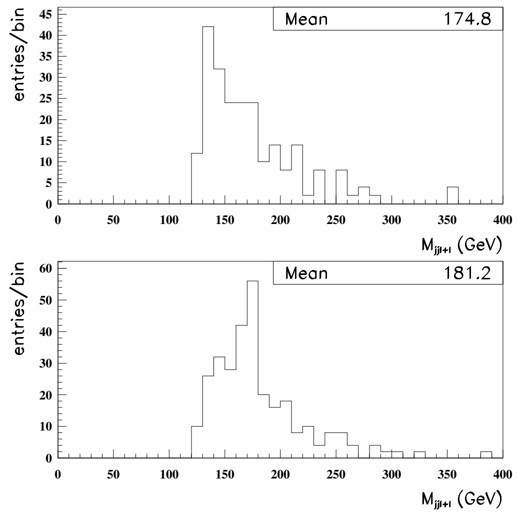

Figure 6.27: Invariant mass distribution for $10^4 pb^{-1}$ of the soft χ_1^0 candidate combined with a lepton, for slepton masses of 140 GeV (top) and 170 GeV (bottom). A Kolmogorov test indicates that slepton mass can be measured to an accuracy of ± 15 GeV.

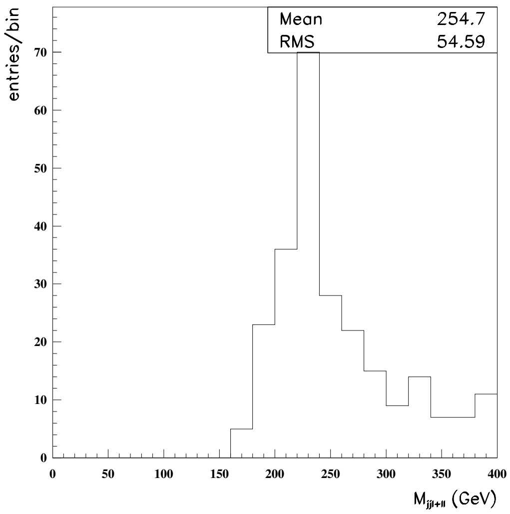

Figure 6.28: Invariant mass distribution for $10^4 pb^{-1}$ of the softer χ_1^0 candidate combined with both remaining leptons to reconstruct the χ_2^0 mass of 228 GeV.

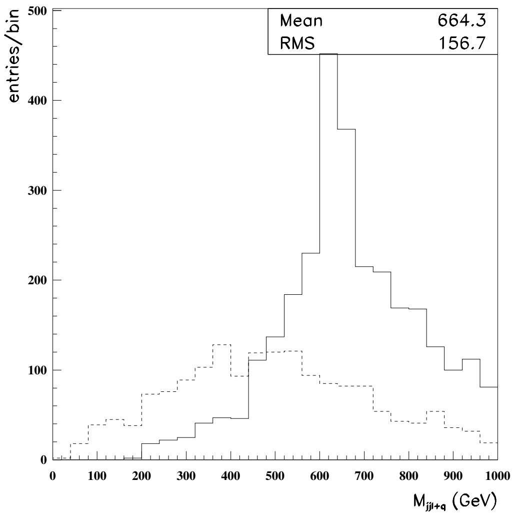

Figure 6.29: Invariant mass distribution for $10^4 pb^{-1}$ of the harder χ_1^0 candidate combined with the leading jet to reconstruct the $\tilde{q}$ mass of 618 GeV. The dotted line shows the invariant mass spectrum of the soft jet combined with the soft χ_1^0.

combinatoric contribution obtained by combining the soft χ_1^0 with the second jet (also figure 6.29).

A broad, but discernible peak can be seen from which the mass of the $\tilde{q}$ could be measured.

6.11.4 Reconstruction of $h \to b\bar{b}$

If the masses of the gauginos, the sleptons, and the squarks are known then, together with additional rate and branching ratio information, the SUGRA parameters m_0 and $m_{\frac{1}{2}}$ can be extracted.

In order to determine $tan\beta$ the mass of the light higgs needs to be measured. The light higgs' (h) are produced with large branching fractions in chargino and neutralino decays. In $R_P - MSSM$ the dominant decay of $h \to b\bar{b}$ gives a clear peak. Reconstructing the invariant mass spectrum of bb pairs when the LSP decays however reveals that the peak is badly degraded by additional jet combinatorial backgrounds (figure 6.30). Events were selected that contained two reconstructed χ_1^0s $|M_{jjl} - M_{\chi_1^0}| < 10$ GeV, and bb pairs considered if $\Delta R_{bb} < 2$.

This situation can be improved by vetoing jets that come from a decaying χ_1^0. By reconstructing the mass of bb pairs with $\Delta R_{bb} < 2$ that are *not* associated with the χ_1^0 jets the spectrum of figure 6.31 can be obtained.

With the addition of this measurement it can be concluded that all the critical parameters of the SUGRA model could be determined in this scenario.

6.11.5 Constraints on SUGRA Parameters

Once several sparticle masses are known the next step would be to perform a global fit to these masses in order to constrain the parameters of the SUGRA model.

Such a fit is not given here since it would be of limited use at this speculative stage in the analysis and it would be a duplication of work done in the case where the LSP decayed

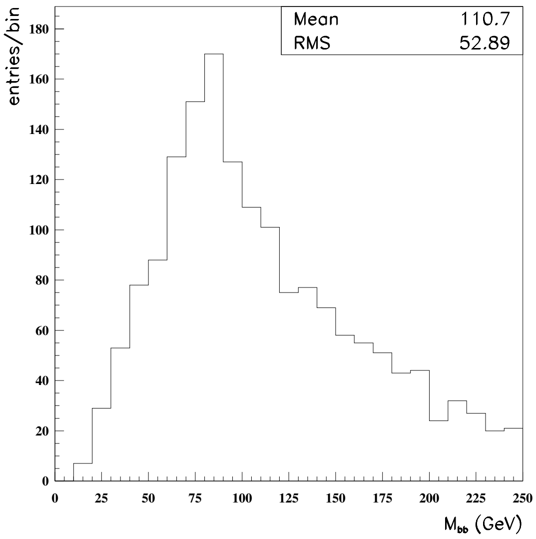

Figure 6.30: $b\bar{b}$ invariant mass spectrum reconstructing the peak from the $h \to b\bar{b}$ decay.

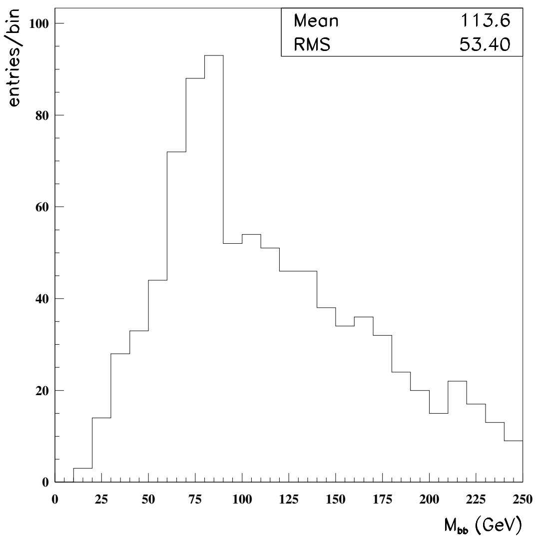

Figure 6.31: $b\bar{b}$ invariant mass spectrum reconstructing the peak from the $h \to b\bar{b}$ decay, vetoing jets that can be associated with a decaying χ_1^0.

to 3 jets [51].

However, comparing these sparticle mass resolutions in this analysis with those in [51] some of the critical conclusions from a global fit can be inferred. Kolmogorov tests have measured the accuracy of the χ_1^0 mass to be better than $\pm 1\%$ (dominated by systematic errors), and the accuracy of the slepton mass to be around ± 15 GeV. By eye it can be seen that all other measurements are accurate to between 10 - 20 GeV. These resolutions are comparable to or better than the situation quoted in [51] and so it is reasonable to conclude that the SUGRA parameters could be constrained to a similar accuracy. However some critical features of such a fit should be noted:

- m_0 and $m_{\frac{1}{2}}$ are measured to accuracies of around 10% and 2% respectively in [51]. Determination of these parameters is tied to the χ_1^0, $\tilde{l}$, and $\tilde{q}$ masses. Of these the $\tilde{l}$ mass has a slightly degraded resolution in this analysis compared with [51] due to a limited efficiency for vetoing leptons associated with χ_1^0 decay. Since the ultimate parameter resolution depends on a combination of these masses, a slightly lower resolution in one of them is not expected to significantly affect this measurement.

- A_0 is not sensitive to any of the masses measured here and cannot be constrained in either scenario. The prime effect of A_0 is to drive the mass splitting of the third generation sparticles; measurements of these mass differences (eg. $\tilde{\tau}_1 - \tilde{\tau}_2$, $\tilde{t}_1 - \tilde{t}_2$, $\tilde{b}_1 - \tilde{b}_2$) would be required to constrain this parameter.

- In [51] the parameter $tan\beta$ is constrained to the range 1.6< $tan\beta$ <4.0 in the absence of a direct measurement of the higgs mass. Further constraints on this parameter in the scenario where the χ_1^0 decays to 3 jets required measurement of the channel $h \rightarrow \gamma\gamma$, and a factor of 10 higher integrated luminosity. In the scenario considered here reduced jet background and the ability to improve higgs mass resolution by vetoing jets associated with the χ_1^0s means that a somewhat better constraint on $tan\beta$ could be made even at low luminosity.

6.12 Mass Reconstruction for Neutral Lepton Decays

Precision measurements for the case of neutral lepton decays will be much more problematic. The escaping neutrino from each χ_1^0 means that there is not enough kinematic information to reconstruct the χ_1^0 mass directly. Therefore measurement of mass combinations should proceed in the same way as for $R_P - MSSM$. However there are in principle four extra jets per event, so jet combinatorics will significantly degrade the accuracy of any measurements made, and the softer E_T^{MISS} spectrum will make Standard Model backgrounds more significant.

It would be premature at this stage in the development of ATLAS, and beyond scope of this analysis, to perform a detailed study of the potential accuracy of measurements in this scenario. For a detailed description of how analysis would proceed see reference [46].

However to investigate the severity of increased jet combinatorics, the impact on two of the most crucial SUSY signatures is investigated.

6.12.1 l^+l^- Mass Edge

The edge in the opposite sign (OS), same flavour (SF) dilepton invariant mass is related to the masses of the χ_2^0, χ_1^0 and $\tilde{l}$ and is an important first stage in a precision analysis. Events were selected that passed the cuts:

- $E_T^{MISS} > 50$ GeV
- Leading jet, $P_T(j_1) > 300$ GeV
- at least 8 jets, $P_T(j_8) > 40$ GeV

Then the invariant mass spectrum of all OS SF dilepton pairs was formed. Figure 6.32 shows the expected dilepton edge after $10^4 pb^{-1}$. Around 52000 SUSY events pass these cuts, compared with 199000 QCD events. However, none of the QCD events passing the initial cuts have any OS SF dilepton pairs and so do not present a background threat. A clear edge can be seen from which the relative masses of χ_2^0, χ_1^0, and $\tilde{l}$ could be accurately measured.

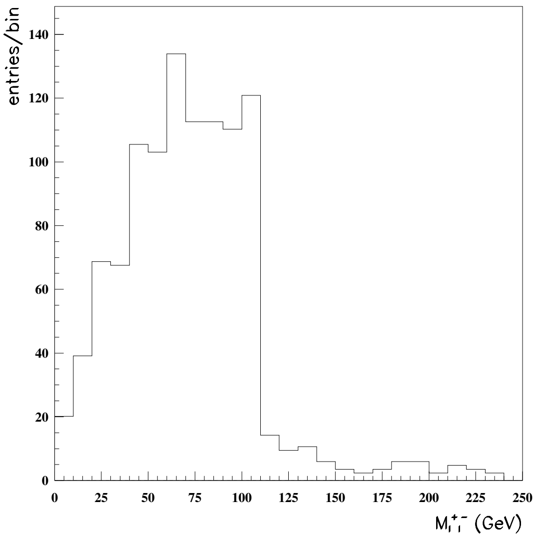

Figure 6.32: OS SF dilepton invariant mass for $10^4 pb^{-1}$.

6.12.2 $h \to b\bar{b}$ Reconstruction

The dominant decay of the light higgs, $h \to b\bar{b}$, is another important signature which gives a clear signal for R_P-MSSM. Two problems make this a particularly difficult signature in this scenario however: firstly, decays of the LSPs cause additional combinatoric backgrounds which degrade the peak. Secondly, SM backgrounds are usually suppressed by requiring a large missing energy. In $\not{R}_P-MSSM$ a lepton has previously been required to remove SM backgrounds, but the additional requirement of a lepton on top of the requirement of a $b\bar{b}$ pair results in very low statistics. Instead events were selected that passed the cuts:

- $P_T(j_1) > 300$ GeV
- $P_T(j_8) > 40$ GeV
- $P_T(j_2^{b-jet}) > 15$ GeV
- $E_T^{MISS} > 80$ GeV

- circularity > 0.2, and thrust < 0.9

After these cuts around 9000 SUSY events are expected, compared with 4000 QCD events (after $10^4 pb^{-1}$). Figure 6.33 shows the expected $b\bar{b}$ invariant mass spectrum for $10^4 pb^{-1}$. The peak is visible above combinatoric backgrounds, but since χ_1^0 jets can't be tagged its significance is badly degraded compared with the charged lepton case. Figure 6.34 shows the same plot but with broader bins (25 GeV bins compared with 10 GeV bins). This plot also shows the significance of the QCD background. Although statistics are limited it is clear that they pose a significant threat.

Statistics are not solid enough at this stage to make any firm conclusions about this channel. However it is clear that it would be a most challenging measurement, and would probably require more than $10^4 pb^{-1}$ obtained during the first year of ATLAS' low luminosity run. One note should be made at this point however; this analysis relies on light Higgs production via χ_2^0 decay. An equally valid way of producing Higgs and therefore measuring SUGRA parameters is via other processes, for instance direct $pp \to h + W$ or $pp \to h + Z$ production.

6.13 Conclusions

A SUSY model using the SUGRA framework has been investigated in the scenario where R_P violation causes the lightest supersymmetric partner to decay into two jets and a lepton. Two scenarios were considered: Charged lepton decay where $\chi_1^0 \to eud$; and neutral lepton decay where $\chi_1^0 \to \nu_e d\bar{d}$.

In the case where the LSP decays to charged leptons, it was shown that a relatively simple set of cuts based on multiple jets and leptons could extract a very clear signal within the first year of low luminosity LHC running. Furthermore it was shown that by reconstructing the invariant mass spectrum of jjl combinations and combining these with other combinations of jets and leptons, the masses of the χ_1^0, slepton, χ_2^0, squark, and lightest higgs boson could be reconstructed. The precision of these measurements is comparable with those made in

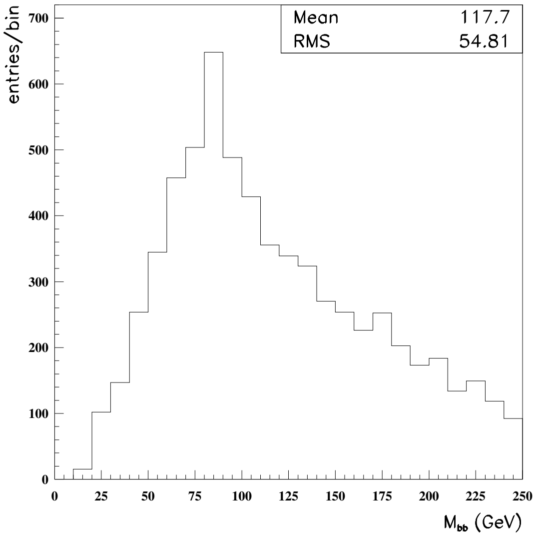

Figure 6.33: $b\bar{b}$ invariant mass spectrum for $10^4 pb^{-1}$.

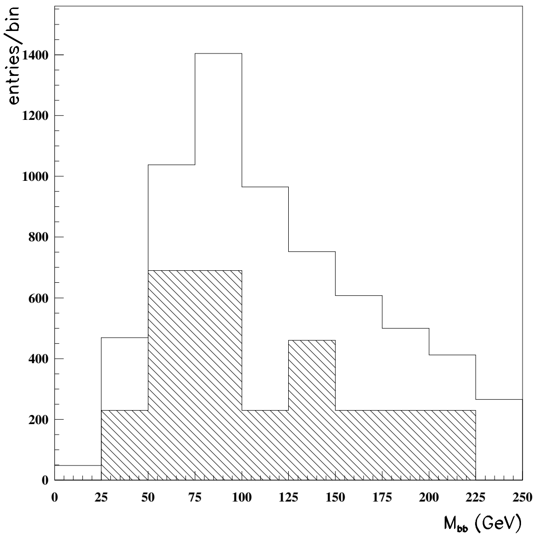

Figure 6.34: $b\bar{b}$ invariant mass spectrum for $10^4 pb^{-1}$ for SUSY (clear) and QCD (hatched).

the case where the LSP decays to three jets, and so the constraints that can be placed on the SUGRA parameters should be similar. Based on a global fit to SUGRA parameters performed in [51] it can be concluded that the values of the SUGRA parameters could be calculated with the following accuracies if the LSP decays to two jets and a charged lepton:

- $\mathbf{m_0}$: Accurate to $< 10\%$.
- $\mathbf{m_{\frac{1}{2}}}$: Accurate to $< 2\%$.
- $\mathbf{A_0}$: Not sensitive to measurements made in this analysis.
- $\mathbf{tan}\beta$: Without measurement of the h^0 mass this parameter can be constrained to 1.6 $< tan\beta <$4.0. Since in the scenario studied a measurement of h^0 is possible, albeit with quite poor resolution, the constraints on $tan\beta$ should be slightly improved.

The case where the LSP decays to neutral leptons is less favourable. An inclusive analysis showed that the acceptance of a lepton and jet based cut is greatly reduced. However, there are still several semi-inclusive channels that can be used to extract a significant signal within the first year of LHC running.

If a signal of this type were extracted then it would not be possible to directly reconstruct the mass of the χ_1^0, and precision measurements would proceed as in the R_P-MSSM case.

In this scenario the lower significance of signal over background and the increased jet combinatorics from the decaying χ_1^0s severely degrade jet based signatures; for example, the decay $h \to b\bar{b}$ is no longer easily reconstructed. For this reason this scenario would be a particularly challenging one for ATLAS, and alternative search strategies should be found.

Furthermore, the most natural scenario when neutral lepton decay could be the only available channel is if a higher generation of Yukawa coupling dominates. For example, if λ'_{333} were dominant then $\chi_1^0 \to \nu_\tau b\bar{b}$ would be allowed but the decay $\chi_1^0 \to \tau t\bar{b}$ would be kinematically forbidden. In this case the $h \to b\bar{b}$ peak would be even more severely degraded due to the high multiplicity of b jets. This could pose a serious problem for precision measurements at ATLAS, and possibly merits further analysis.

Bibliography

[1] For a general introduction see: W. E. Burcham and M. Jobes, *Nuclear and Particle Physics*, Longman Scientific and Technical, (1995).

[2] For a general introduction see: F. Halzen and A. D. Martin, *Quarks and Leptons: An Introductory Course in Modern Particle Physics,* John Wiley & Sons, (1984).

[3] `www.phys.hawaii.edu:80/~jgl/nuosc_story.html`

and references therein.

[4] P. Ramond, *Field Theory - A Modern Primer*, Addison-Wesley Pub. Co., (1989).

[5] R. P. Feynman, *Quantum Electrodynamics*, Benjamin, (1961).

[6] S. L. Glashow, *Nucl. Phys.* **22**, 579 (1961).

[7] S. Weinberg, *Phys. Rev. Lett.* **19**, 1264 (1967).

[8] A. Salam, *Elementary Particle Theory*, Almqvist and Wikell, Stokholm, (1968).

[9] M. Kobayashi and K. Maskawa, *Prog. Theor. Phys.* **49**, 282 (1972).

[10] P. W. Higgs, *Phys. Rev. Lett.* **13**, 508 (1964), and *Phys. Rev.* **145**, 1156 (1964).

[11] W. Marciano and H. Pagels. *Phys. Reports* **36C**, 137 (1978).

[12] H. Georgi and S. L. Glashow, *Phys. Rev. Lett.* **32**, 487 (1974).

[13] For a general introduction see to supersymmetry: H. E. Haber and G. L. Kane, *Phys. Rept.* **117:75**, (1985), and *Scientific American*, Vol. 254, Pg. 52, (1986).

[14] C. W. Fabjan, *LHC: Physics, Machine, Experiments*, CERN Report **CERN-PPE/95-25** (1995).

[15] *Design Study of the Large Hadron Collider (LHC)*, CERN Report **CERN/91-03** (1991).

[16] S. Myers and E. Picasso, *The Design, Construction and Commissioning of the CERN Large Electron-Positron Collider*, Contempory Physics **31**, 6 (1990).

[17] *Proceedings of the Larger Hadron Collider Workshop* (Aachen, Germany, 1990), ed. G. Jarlskog and D. Rein. **CERN 90-10/ECFA 90-133** (1990) Vol. I,II,III.

[18] W. W. Armstrong et al., *ATLAS: Technical Proposal for a General-Purpose pp Experiment at the Large Hadron Collider at CERN*, CERN Report **CERN/LHCC/94-43** (1994).

[19] ATLAS Inner Detector Community, *ATLAS Inner Detector Technical Design Report*, CERN Report **CERN/LHCC/97-16** (1997).

[20] ATLAS Collaboration, *ATLAS Calorimeter Performance Technical Design Report*, CERN Report **CERN/LHCC 96-40** (1997).

[21] ATLAS Muon Collaboration, *ATLAS Muon Spectrometer Technical Design Report*, CERN Report **CERN/LHCC/97-22** (1997).

[22] ATLAS Level-1 Trigger Group, *Level-1 Trigger Technical Design Report*, CERN Report **CERN/LHCC/98-14** (1998).

[23] J. Garvey and D. Rees, *Bunch Crossing Identification for the ATLAS Level-1 Calorimeter Trigger*, ATLAS Internal Note, **ATL-DAQ-96-051** (1996).

[24] S. Dasu and W.H. Smith, *CMS Level-1 Jet Trigger Study*, CMS note **TN/96-066**, (1996).

[25] S. George et al., *ATLAS Trigger Menus*, ATLAS Internal Note **PHYS-NO-124**, (1998).

[26] T. Sjöstrand, Computer Physics Commun., **82**, (1994), 74.

[27] S. Catani et al., *Longitudially-invariant k_T-clustering algorithms for hadron-hadron collisions*, Nuclear Physics **B406** (1993) 187.

[28] J. Bystricky et al., *ATLAS Trigger Menus at Luminosity* $10^{33}\ cm^{-2}s^{-1}$, ATLAS Internal Note **ATL-DAQ-96-054**, (1996).

[29] ATLAS Trigger Performance Group, *ATLAS Trigger Performance Status Report*, CERN Report, **LHCC 98-15**, (1998).

[30] *GEANT, Detector Description and Simulation Tool*, Application Software Group, Computing and Networks Division, CERN Program Library Long Writeup **W5013**.

[31] S. Mrenna, *SPYTHIA, A Supersymmetric Extension of PYTHIA 5.7*, **ANL-HEP-PR-96-63**, (1996).

[32] A. Artamonov et al., *DICE-95*, Atlas Internal Note **ATLAS-SOFT/95-14c**, (1995).

[33] A. Dell'Acqua et al., *1997 ATLAS Jet Production*, ATLAS Internal Note **PHYS-NO-102**, (1997).

[34] *A Demonstrator Programme for the ATLAS Level-1 Calorimeter Trigger*, RD27 note 35, ATLAS note **DAQ-NO-031**, (1995).

[35] I. P. Brawn et al., *Beam Tests of a Prototype Level-1 Calorimeter Trigger Processor for LHC Experiments*, Nucl. Instrum. Methods **A349**, (1994) 356-366.

[36] P. Grenier and L. Says, *Analysis of the 1997 Test Beam Laser Data of the Tilecal extended barrel modules 0*, ATLAS Internal Note **ATL-TILECAL-98-159**, (1998).

[37] J. Edwards and V. Perera, *User Guide for the First Level Calorimeter Trigger Processor Demonstrator Dual Function ASIC RAL163*, Hardware write-up, (1996).

[38] Level 1 Calorimeter Group, *Software Requirements for the Level 1 Calorimeter Trigger*, **LVL1-Calo-Rev-10**, (1998).

[39] J. Leake et al., *Diagnostic Software for the ATLAS Level-1 Calorimeter Trigger*, ATLAS Internal Note **DAQ-NO-72**, (1997).

[40] U. Pfeiffer et al., *Performance of the Front-end Demonstrator System of the ATLAS Level-1 Calorimeter Trigger*, **CERN/LHCC/97-60**, Third Workshop on Electronics for LHC Experiments, London, UK, September, (1997).

[41] J. F. Gunion, *A Simplified Summary of Supersymmetry*, Proceedings of the ITP Symposium, U.C. Santa Barbara, **UCD-97-9**, (1996).

[42] I. Hinchliffe et al., *Precision SUSY measurements with ATLAS: Introduction and Inclusive Measurements*, ATLAS Internal Note, **PHYS-No-107**, (1997).

[43] E. Richter-Was et al., *Precision SUSY Measurements with ATLAS points 1 and 2*, ATLAS Internal Note, **PHYS-No-108**, (1997).

[44] I. Hinchliffe et al., *Precision SUSY Measurements at LHC: Point 3*, ATLAS Internal Note, **PHYS-No-109**, (1997).

[45] F. Gianotti et al., *Precision SUSY Measurements with ATLAS: SUGRA "Point 4"*, ATLAS Internal Note, **PHYS-No-110**, (1997).

[46] G. Polesello et al., *Precision SUSY Measurements with ATLAS for SUGRA Point 5*, ATLAS Internal Note, **PHYS-No-111**, (1997).

[47] H. Baer et al., *Proceedings of the Workshop on Physics at Current Accelerators and Supercolliders*, eds. J. Hewett, A. White, D. Zeppenfield (Argonne National Laboratory, 1993).

[48] G. Bhattacharyya, *R-Parity Violating Supersymmetric Yukawa Couplings: A Mini-review*, **hep-ph/9608415**, (1996).

[49] H. Dreiner, *An Introduction to Explicit R-parity Violation*, **hep-ph/9707435**, (1997).

[50] The decay $\chi_1^0 \rightarrow ll\nu$ is studied by E. Nagy and A. Mirea, CPPM, Marseille, see ATLAS SUSY working group meeting 4th March 1998.

[51] J. Söderqvist, *Consequences of Baryonic R-Parity Violation for Measurements of SUSY Particles using the ATLAS Detector*, ATLAS Internal Note, **PHYS-No-122**, (1998).

[52] H1 Collaboration, *A Search for Squarks of R_P-violating SUSY at HERA*, **hep/ex/9604006**, (1996).

[53] E. Richter-Was, *ATLFAST 2.0 a Package for Particle-level Analysis*, ATLAS Internal Note, **PHYS-No-1xx**, (1997).

[54] Scanned copies of the SUSY 1996 workshop can be found in: `http://www.cern.ch/Committees/LHCC/SUSY96.html`

[55] Baer et al., *Phys. Rev.* **D53**, (1996) 6241.

[56] W. T. Eadie et al., *Statistical Methods in Experimental Physics*, North-Holland Pub., (1971).