
Jet Physics at ATLAS

Daniel Robert Clements
Department of Physics and Astronomy
University of Glasgow

Thesis submitted in June 2008 for the degree of
Doctor of Philosophy
at the University of Glasgow

Abstract

The use and optimisation of integration grid techniques to generate next-to-leading order predictions of jet cross-sections, independent of parton distribution functions, was investigated. Such methods were found to provide an accurate approximation to a standard Monte-Carlo simulation (within 1%) and enable collider data to be readily included in global PDF fitting procedures. However, the benefit of including inclusive-jet cross-section data from ATLAS in global fits is only significant if the jet energy scale (JES) can be constrained to $\sim 1\%$ at high P_T .

Uncertainties in the theoretical prediction of the inclusive-jet cross-section such as PDFs and fixed-order (scale) uncertainties were studied and compared with experimental errors arising from jet energy resolution and absolute scale. These uncertainties were then considered in the context of a quark compositeness search where a sensitivity to a compositeness scale of $\Lambda \leq 10\text{TeV}$ can be achieved with 10fb^{-1} of data, if the jet energy scale can be constrained to $\sim 1\%$. An analysis using dijet angular distributions found a similar sensitivity without the dependence on the jet energy scale.

A potential method of evaluating the stability of the jet energy scale out to high P_T by ‘bootstrapping’ the calibration at low P_T by the use of multi-jet events was also investigated. This suggests that a calorimeter non-linearity can be detected for jets with $P_T > 500\text{GeV}$ at $\sim 1.5\%/500\text{GeV}$ (i.e. a 1.5% change in JES over 500GeV in P_T).

An investigation of inner-detector commissioning issues associated with the ATLAS Semiconductor Tracker (SCT), including a review of ‘noisy’ modules on the SCT Barrel (from May 2007) was carried out. In addition a tool for DCS monitoring within the online monitoring framework was developed and tested during the M5 and M6 commissioning weeks. Finally, a method of assessing the track reconstruction efficiency by track-insertion was considered for the particular case of minimum bias events.

Declaration

I have played an active role in developing the integration grids method (originally proposed in [1]) including introducing techniques for reducing computer memory consumption and optimising its implementation. This work is currently being prepared for publication in association with Tancredi Carli et al. Furthermore, I generated grids for use in PDF fitting and applied the technique to the calculation of theoretical errors on the inclusive jet cross-section as generated using NLOJET⁺⁺. This work was presented at the DIS-2006 conference and on behalf of the ATLAS collaboration (by invitation) at DIS-2007. I then used these errors to inform a quark compositeness search.

In addition I have worked in assessing performance issues of the ATLAS detector. These include a study of a ‘bootstrapping’ method of calibrating the high P_T jet energy scale and research into a means of assessing the track reconstruction efficiency of the inner detector by track insertion. This method is being developed in preparation for the first running of ATLAS. I have also contributed to the ATLAS inner detector monitoring group by writing a tool for the online monitoring of the DCS as well as acting as a monitoring expert for the online shifts in the M5 and M6 commissioning weeks.

This thesis is the result of my own work, except where explicit reference is made to the work of others, and has not been submitted for another qualification to this or any other institution.

Dan Clements

Contents

1	Introduction	1
1.1	The Standard Model Of Particle Physics	1
1.2	Physics Beyond The Standard Model	5
2	Experimental Description	8
2.1	The Large Hadron Collider - LHC	8
2.2	The ATLAS Experiment	10
2.2.1	The Inner Detector	11
2.2.2	Calorimetry	14
2.2.3	Muon Spectrometer	17
2.2.4	Trigger and DAQ	19
2.2.5	Simulation and Offline Analysis	22
3	Theory	24
3.1	Introduction to QCD	24
3.1.1	Features of QCD	24
3.1.2	Hadron-Hadron Interactions	26
3.2	QCD Calculations	28
3.2.1	The Inclusive Jet Cross-Section at Leading Order	28
3.2.2	Next-to-Leading Order Corrections	30
3.2.3	Higher Orders and Unphysical Scales	32
3.2.4	Resummation and Parton Showers	33
3.3	QCD Coupling and PDFs	34
3.3.1	Determining the Strong Coupling Constant	34
3.3.2	Determining the PDFs	35
3.4	Jet Algorithms	38
3.4.1	The role of jet algorithms	38
3.4.2	Examples of jet algorithms	39
4	Grid Methods for NLO Cross-Sections	42
4.1	Introduction to PDFs and Global Fits	42
4.2	Integration grids for QCD calculations	43
4.2.1	Using a grid to represent a PDF	43
4.2.2	Example of Grid Method	45
4.2.3	Recording cross-section weights on a grid	48
4.2.4	Grids for Hadron-Hadron interactions at NLO	49
4.2.5	Adjusting the hard scale	51

4.3	Implementing a grid design	52
4.3.1	Co-ordinate Transformations	53
4.3.2	Bin-Boundaries Optimisation	55
4.3.3	Interpolation	56
4.3.4	PDF Weighting	56
4.4	Testing Grid Implementation	62
4.4.1	Effect of Modifying Grid Size	63
4.4.2	The Effect of Interpolation Order	66
4.4.3	The Effect of PDF Weighting	67
4.4.4	The Effect of the Co-ordinate Transformation	68
4.4.5	Adjusting the Renormalisation and Factorisation Scales	69
4.5	Grids in Global PDF Fits	69
5	The Inclusive Jet Cross-Section	73
5.1	Introduction	73
5.2	Theoretical Uncertainties	75
5.2.1	Fixed-Order Uncertainties	75
5.2.2	PDF Errors	81
5.3	Jet Errors - Experimental	84
5.3.1	Issues with Jet Reconstruction	84
5.3.2	Offline Jet Reconstruction and Calibration at ATLAS	86
5.3.3	Standard Candles for Calibration	88
5.3.4	Experimental Errors on Jet Measurements	90
5.3.5	Experimental Errors on the Inclusive Jet Cross-Section	92
5.3.6	Statistical Errors	92
5.3.7	Jet Energy Resolution Errors	94
5.3.8	Jet Energy Scale Errors	97
6	Bootstrapping the Jet Energy Calibration	100
6.1	Calibration of high P_T jets	100
6.2	Jet Balancing Cuts	101
6.3	Jet Balancing using the ATLAS simulation	101
6.4	PYTHIA Generator-Level Jet Balancing	103
6.5	Determining Jet Energy Scale Uncertainty	105
7	Quark Compositeness at the LHC	109
7.1	Sensitivity to New-Physics	109
7.2	Compositeness Searches in Inclusive Jets	111
7.2.1	Analysing the Inclusive Jets Cross-Section	115
7.3	Di-Jet Angular Distributions	117
7.3.1	Theoretical Uncertainties	121
7.3.2	Experimental Uncertainties	128
7.3.3	Analysis Of Di-Jet Angular Distributions	129

8	The ATLAS Semiconductor Tracker	135
8.1	Introduction to The SCT	135
8.2	The p-n junction	135
8.3	Module Composition and Operation	137
8.4	Reading Out Modules	137
8.5	Calibration of Modules	139
8.5.1	Setting the Overall Threshold	139
8.5.2	The Trim Range Test	141
8.5.3	Three point gain test	142
8.5.4	Response Curve Test	143
8.5.5	Noise Occupancy Test	143
8.6	SCT Barrel Tests in the pit - May 2007	144
8.6.1	Aim	144
8.6.2	Sources of Data	144
8.6.3	Noise Occupancy from Monitoring, RODs and Calibration	144
8.6.4	Strip-Level Information	147
8.6.5	Known problem strips from calibration	148
8.6.6	Conditions Data	151
8.6.7	Summary and Postscript	152
8.7	Online Monitoring of the SCT and DCS	155
8.7.1	Online DCS monitoring tool	156
8.7.2	Testing in the M5 commissioning runs	157
9	Tracking Efficiency in Minimum-Bias Events	160
9.1	Criteria for Successful Track Reconstruction	162
9.2	Minimum-Bias Events	162
9.2.1	Effect Of Selection Cuts	164
9.2.2	Reconstruction Efficiency	167
10	Conclusions	171
11	Acknowledgements	189

Chapter 1

Introduction

1.1 The Standard Model Of Particle Physics

Particle physics is concerned with the fundamental constituents of matter and their interactions. At present four basic types of interaction are known: gravitational, weak, electromagnetic and strong. With the exception of gravity, the forces are described by a quantum field theory referred to as the ‘Standard Model’ (SM), which is built upon both special relativity and quantum mechanics.

The particles that constitute matter possess different ‘charges’ which determine the types of interaction they can undergo. Each interaction relates to a different charge, e.g. mass is the charge which allows objects to experience gravity. In a similar fashion ‘electric charge’, ‘weak isospin’ and ‘colour’ are the respective charges for the electromagnetic, weak and strong forces.

The different characteristics of each interaction can be demonstrated by appealing to examples where an individual dominates the others. For instance, the orbit of planets around the sun in the solar-system can be attributed solely to the gravitational force. Gravity is very feeble at a particle level but dominates at astronomical scales as the force is always attractive between bodies and increases in strength proportionally with the mass of a body involved. The other forces tend not to cause noticeable effects at such scales as

a large net ‘charge’ (e.g. electrical or colour) is not found in macroscopic objects¹.

The weak interaction is responsible for β decay of neutrons into protons and the associated radioactivity. In most situations, the weak interaction is negligible in comparison to the electromagnetic and strong forces, however, certain particle decays (those involving a flavour change) can only proceed via this channel.

The electromagnetic interaction mediates a force between electrically charged particles. The force underpins the structure of everyday macroscopic objects (such as this thesis²), as it is responsible for the formation of atoms and the interactions between them. As a result, chemistry may be considered at a fundamental level as a study of the electromagnetic interaction.

Finally, the strong force holds coloured partons together to form hadrons (e.g. protons, neutrons and pions). In addition, the force is responsible for binding protons and neutrons together to form nuclei, by a residual effect in analogy with Van der Waals forces between electrically neutral objects. The strong interaction in nuclei dominates the repulsive electromagnetic forces which would otherwise act to separate the positive charges of the protons. However, the strong force does not directly affect atomic structure as the effects of the short range interaction are contained within the nuclei and electrons do not carry colour charge.

A brief summary of the particles of the first generation of the Standard Model and their electric charge is given in Table 1.1.

Symbol	Type	Electric Charge(e)	Possible Interactions
u	Quark	$+2/3$	Strong, Weak, EM and gravity
d	Quark	$-1/3$	Strong, Weak, EM and gravity
e^-	Lepton	-1	Weak, EM and gravity
ν_e	Lepton	0	Weak and gravity

Table 1.1: First generation of particles in the Standard Model, all are fermions i.e. have $\text{spin} = 1/2$

In the Standard Model, the interactions between the spin $1/2$ fermions are mediated

¹It is possible to give a macroscopic object a small, net electrostatic charge (static electricity).

²This is true whether read in paper form or from a computer monitor.

by the exchange of integer spin particles called bosons. The exchange particles couple with the ‘charge’ on the fermion with a given coupling ‘constant’ which determines the strength of the interaction. The term constant is a misnomer as the coupling can vary as a function of the momentum transfer (or hardness) of an interaction. The exchange (or gauge) bosons and relative strengths of the couplings for the Standard Model are given in Table 1.2.

Interaction	Charge	Gauge Boson	Relative Coupling
Strong	colour	gluon	~ 1
EM	electrical charge	photon	10^{-2}
Weak	weak isospin	W^+, W^-, Z	10^{-5}

Table 1.2: Interactions of the Standard Model. The relative couplings are defined as: strong= α_s (strong coupling), EM= α (fine-structure constant) and Weak= $G_F m_p^2$ (where G_F is the four fermion coupling and m_p is the proton mass). The true strength of the interactions vary with the momentum transfer (Q^2) between participating particles. Here the low Q^2 limit is considered.

As a consequence of relativity, each particle has an anti-particle partner, which has the same mass but the opposite sign for all additive quantum numbers (e.g. electrical charge). Hence, there is an anti-up quark, anti-down quark, anti-neutrino and an anti-electron more commonly referred to as a positron. In the case of fermions the anti-particles are distinct particles, however for bosons (e.g. the photon) the particles are often their own anti-particles.

In addition to the particles listed in Table 1.1, there exists two further generations of fermions. The 2nd and 3rd generations contain particles which are similar to those of the first in that they undergo the same types of interaction and have the same electrical charges, but are at higher mass, see Table 1.3. Each lepton generation has an associated quantum number (lepton number) which is conserved in interactions whilst each quark carries a distinct ‘flavour’ which may be changed in weak decays. When all three generations are considered it brings the total number of fermions in the Standard Model to 24 and the total number of fundamental particles to 30 (including the Higgs Boson).

The Standard Model can be thought of as a collection of two separate theories.

	Quarks			Leptons		
Generation	Flavour	Charge(e)	Mass(GeV)	Flavour	Charge(e)	Mass(GeV)
1st	u	$+2/3$	~ 0.003	e^-	-1	0.000511
	d	$-1/3$	~ 0.006	ν_e	0	~ 0
2nd	c	$+2/3$	~ 1.3	μ^-	-1	0.106
	s	$-1/3$	~ 0.1	ν_μ	0	~ 0
3rd	t	$+2/3$	~ 175	τ	-1	1.7771
	b	$-1/3$	~ 4.3	ν_τ	0	~ 0

Table 1.3: The three generations of the Standard Model

The electroweak theory encompasses both electromagnetic³ and weak interactions, whilst quantum chromodynamics (QCD) describes the strong interaction. The theories are based on a Lorentz invariant Lagrangian which describes the material particles and the interactions between them. Symmetries in the Lagrangian play a key role in determining the nature of the interactions. In particular the application of a ‘local gauge symmetry’ allows particle interactions to appear as a natural consequence. A more detailed overview of QCD is given in chapter 3.

Currently the Standard Model is in full agreement with available experimental data accumulated over the past three decades⁴. This includes data from fixed target and colliding beam experiments in addition to more specific experiments measuring the anomalous magnetic moment of the electron⁵. The only prediction of the Standard Model for which there is currently no experimental evidence is the existence of the Higgs boson. The Higgs is required in the electroweak theory to allow the gauge bosons (W^+ , W^- and Z^0) to acquire mass and may have escaped detection thus far by being too massive to have been generated at present collider experiments. Although the Standard Model does not predict the mass of the Higgs, general arguments involving radiative corrections limit it to the TeV scale. Despite the success of the Standard Model there is, however, a great interest in possible extensions or alternatives to the theory.

³The electroweak theory contains the earlier derived quantum electrodynamics (QED) which describes electromagnetic interactions only.

⁴With the exception of measured neutrino oscillations which are not compatible with the massless neutrinos of the Standard Model. The measurement of the anomalous magnetic moment of the muon (g-2) is also in tension with the theoretical prediction ($\sim 2.6\sigma$ discrepancy).

⁵The measurement of which provides the best numerical agreement between data and theory in physics.

1.2 Physics Beyond The Standard Model

The Standard Model contains a large number of free parameters (19 excluding neutrino masses) that have to be determined by experiment. The parameters fix the rest mass of fundamental particles, coupling constants for the interactions and weak mixing angles⁶. It is often propounded on aesthetic grounds that a complete theory should not contain multiple free parameters, but this in itself does not constitute evidence of a flaw in the current formalism. There are, however, other features intrinsic to the theory itself, that suggest the presence of physics beyond the Standard Model (BSM).

The fundamental particles of the Standard Model (see Table 1.3) are arranged into three generations. The theory includes all these generations in its formalism but does not provide a motivation for this structure or explain why there are only three generations. In addition the Standard Model is unable to explain why electrical charge in the fundamental particles appears to be quantized.

The Higgs mechanism also creates problems for the Standard Model as the Higgs boson which arises due to spontaneous symmetry breaking in the electroweak theory suffers large radiative corrections to its mass. These corrections are divergent and if the Standard Model is presumed to be accurate up to the Planck scale, fine tuning is required (between the bare Higgs mass and the fluctuations) to create a Higgs with a mass of $O(100)GeV$.

One approach to addressing the fine tuning problem is to apply an additional symmetry to the Standard Model Lagrangian. In supersymmetric (SUSY) theories a symmetry transformation between a fermion and a boson is postulated. The symmetry is necessarily broken in order to reconcile the absence of boson-fermion partners occurring at the same mass (e.g for the charged leptons). By allowing the symmetry to be broken, however, the superpartners of the current particles could be higher in mass and have escaped detection at present experiments. SUSY models tame the radiative corrections of the Higgs (and other particles) as contributions from superpartners tend to cancel those of each other. For

⁶there is an additional strong CP parameter θ which is often set to zero

this cancellation to be effective the superpartners cannot be considerably more massive than the original particle and hence SUSY would suggest a rich discovery ground for new particles at the TeV scale. Stable SUSY particles are also candidates for explaining dark matter in the universe.

Another possible extension to the Standard Model is motivated from the formalism itself which relies heavily on group theory. The gauge theories of QED, the weak interaction and QCD all are based on Lie groups ($SU(3)$, $SU(2)$ and $U(1)$) which can be a subgroup of a larger group (e.g $SU(5)$). Grand unified theories (GUTS) aim to combine all the groups of the Standard Model into a single (larger) group which at a certain scale (the GUT scale) is responsible for all interactions. This is motivated by the fact that coupling constants for the interactions change (or run) with the scale of the interaction (or hardness) and would appear, if extrapolated, to approximately converge at high Q^2 . The structure of GUTS also naturally leads to an explanation for the consistency of electrical charge quantization between leptons and quarks. GUTS although theoretically appealing suffer from a lack of testable predictions and early attempts at theories (e.g. using $SU(5)$) led to a prediction of a proton lifetime which was too short when compared to experiment.

The generations of particles found in the Standard Model could be in part explained if they were considered to be composite as opposed to fundamental. The constituents of these particles ‘preons’ would be held together by very strong forces interacting with hypercolour (or metacolour) charges. The higher mass generations of the Standard Model could then be interpreted as excitations of the lowest mass state. There is, however, currently no experimental evidence for compositeness and theoretical problems are encountered in its formalism.

To make progress with the questions surrounding the Standard Model including the existence of the Higgs boson, experimental data is required. A major source of such data over the last few decades has been from colliding particle beam experiments such as ALEPH and those at the Tevatron. These experiments collide particles at high energy and allow high mass particles (such as the Higgs etc.) to be produced and their existence

inferred from their decay products. One such colliding beam experiment is ATLAS (based at the LHC) which is due to come on-line in mid-2008.

Chapter 2

Experimental Description

2.1 The Large Hadron Collider - LHC

The Large Hadron Collider (LHC) under construction at CERN is designed to accelerate and collide protons at a higher centre-of-mass energy ($\sqrt{s} = 14TeV$) and luminosity ($L = 10^{34}cm^{-2}s^{-1}$) than previous accelerators. It is hoped that this will allow the nature of electroweak symmetry breaking to be determined and an exploration of potential ‘new physics’ scenarios such as supersymmetry. The collider’s high luminosity is instrumental to the detection of rare events and for providing sufficient statistical accuracy to claim discovery of new particles.

The LHC is a synchrotron which accelerates bunches of protons from an energy of 450GeV to 7TeV by dynamically adjusting the magnetic and electric fields which keep the particles in a near circular orbit and accelerate them respectively. This adjustment of the electric and magnetic fields allows the bunches to be accelerated along the circumference of a circle at a constant radius.

Prior to being injected into the LHC, protons are accelerated from rest by a number of smaller accelerators which gradually increase their energy. These include a Linac, the Proton Synchrotron Booster (PSB - 1.4GeV), the Proton Synchrotron (PS - 26GeV) and the Super Proton Synchrotron (SPS - 450GeV). A diagram showing the accelerator complex at CERN can be found in Figure 2.1.




Figure 2.1: Accelerators at CERN

Synchrotrons are generally limited to accelerating charged, stable particles which in practice leads to a choice between electrons, protons¹ or heavy-ions. For the study of high energy interactions, colliding leptons is preferable to composite QCD objects (such as protons) as the particles are fundamental and the centre of mass energy of a particle collision is related directly to the particle's energy in the beam².

However, when a charged particle is accelerated through a magnetic field it emits radiation (synchrotron radiation) which degrades its energy. The energy loss is inversely proportional to the squared mass of the accelerating particle and is the reason why the LHC is forced to collide protons rather than electrons.

The centre-of-mass energy of a synchrotron is limited by the strength of the dipole (bending) magnets needed to keep the bunches in circular motion at a fixed radius. The large radius of the synchrotron (4.24km) reduces the curvature per unit length along the circumference allowing particles to be accelerated to higher energies for a given magnetic

¹The anti-particles: positrons and anti-protons, may also be used

² $\sqrt{s} = 2E_b$, for colliding beams whose particles have energy E_b .

field. The luminosity of the collider is a function of many parameters including how many particles are present in a bunch, the frequency of bunch crossings at an interaction point and the geometry of the collision [2].

The protons undergo collisions at four points around the circumference where the experiments ATLAS, CMS, LHCb and ALICE are situated. These experiments are designed to detect the products of the proton-proton collision in order to reconstruct the event and determine the nature of the interaction. ATLAS and CMS are general purpose detectors, whereas LHCb is a dedicated B-physics experiment and ALICE specialises in the physics of heavy-ion collisions³.

2.2 The ATLAS Experiment

The ATLAS detector is designed to be sensitive to signatures of the Higgs boson and new physics which may be produced at the LHC. The philosophy behind the detector is to identify a wide range of particle types and energies, occurring in a variety of event topologies. An ideal experiment would detect all the particles produced in an event and directly measure their charge, energy and momentum. This is unfeasible as many particles are unstable and decay before such a measurement can be made. Instead, the production of unstable particles is inferred from visible decay products⁴ which the detector is optimised to measure.

General purpose detectors are built as a series of coaxial cylinders centred on the interaction point (where the proton bunches collide) whose common axis is the beam line. Typically the innermost detector is a tracker in a solenoidal magnetic field for measuring the path and momenta of charged particles followed by electromagnetic and hadronic calorimeters, which measure the energy of e/γ and jets respectively. Calorimetry is a destructive measurement in that measured particles are stopped in the detector and their energy is dissipated (hence they are placed after the tracking). Muons, however, can

³The LHC can be used as a heavy ion Pb-Pb collider.

⁴In practice only photons, electrons, muons, pions, kaons, protons and neutrons can be directly measured at a collider-detector.

penetrate calorimeters and may be measured afterwards in a spectrometer. Neutrinos produced at the interaction point escape the detector, but may be inferred from kinematic constraints by the presence of missing transverse energy (E_T) in an event.

ATLAS (see Figure 2.2) has efficient tracking, good electromagnetic calorimetry and near hermetic jet measurement capability. These features allow the detector to perform precision measurements on leptons, jets and photons in addition to identifying events with missing E_T . A muon spectrometer is also used to measure the momentum of muons and to assist in triggering. In addition to reconstructing the event, the detector has to survive in a harsh radiation environment and operate at high speed in order to differentiate events from adjacent bunch crossings which occur every 25ns.




Figure 2.2: The ATLAS detector [3]

2.2.1 The Inner Detector

The inner detector (see Figure 2.3) is designed to accurately measure the paths and momenta of charged particles. The momentum measurement is carried out by measuring the curvature of a particle's path in a 2T solenoidal magnetic field. The solenoidal field causes

charged particles to spiral in the ϕ plane and is provided by a cylindrical superconducting magnet which surrounds the inner detector volume (length $\sim 6m$, radius $\sim 1m$). The paths of particles are determined from discrete measurements made by silicon and straw tube (gas) detectors.

High-energy charged particles have a tendency to ionise matter they traverse and detection of this forms the basis of many types of particle detector. In a semiconductor tracker, charge carriers (electrons and holes) are produced in silicon by a traversing charged particle. These charge carriers are swept (in an electric field) toward electrodes creating a measurable current. Silicon trackers can be finely segmented into strips or pixels allowing for high positional accuracy (typically μm).

The ATLAS silicon trackers consist of high resolution pixel and strip detectors. The pixel detectors are the closest to the interaction point and consist of three cylinders at radii of approximately 5cm, 9cm and 12cm [3]. Further pixel detectors are arranged on 3 disks either side of the barrel. Surrounding the pixel layers there is a semiconductor tracker (SCT) comprised of silicon strip detectors. The SCT barrel region has 4 layers of microstrip detectors supplemented by 9 end-cap disks on each side.

The SCT is surrounded by the Transition Radiation Tracker (TRT). This is comprised of a large number of small tubes filled with a mixture of Xe , CO_2 and CF_4 . They operate in a similar manner to the silicon trackers with the gas as the active medium⁵. Between the straws a radiator of polymer fibres is placed to produce transition radiation which can be used to discriminate between electrons and hadrons with low P_T . The transition radiation is produced from charged particles as they traverse a boundary between two media of different dielectric constant. The subsystems of the inner detector and their acceptance are given in Table 2.1⁶.

The discrete spatial points ('hits') measured by the inner detector subsystems are reconstructed into tracks by a tracking algorithm. A track is defined by its $1/P_T$, azimuthal

⁵Liberated electrons move toward a wire in the tube and ionised gas molecules move towards the surrounding cylinder to create a detectable current.

⁶The acceptance is given in terms of pseudorapidity η . This is related to the azimuthal angle θ by $\eta = -\ln(\tan\frac{\theta}{2})$. Pseudorapidity can be considered as an approximation to rapidity in the case of a highly relativistic particle [4].




Figure 2.3: The ATLAS inner detector, incorporating the pixel detector, SCT and TRT[3]

System	Description	Resolution (μm)	coverage $ \eta $
Pixels	3 barrel layers	$R\phi = 10, z=115$	0-2.5
	3 end-cap disks (on each side)	$R\phi = 10, R=115$	
SCT	4 barrel layers	$R\phi = 17, z=580$	0-2.5
	9 end-cap wheels (on each side)	$R\phi = 17, z=580$	
TRT	axial barrel straws	130	0-2.0
	radial end-cap straws	130	

Table 2.1: Inner detector subsystems, resolution and acceptance (from [3])

(θ) and polar (ϕ) angles in addition to the transverse (d_0) and longitudinal (z_0) impact parameters. By use of a full detector simulation it is possible to estimate the resolution of these tracking parameters and this is shown in Table 2.2 [3]. The resolution is given in the form:

$$\sigma_X = \sigma_X(\infty)(1 \oplus P_X/P_T) \quad (2.1)$$

where $\sigma_X(\infty)$ is the resolution of the track parameter as $P_T \rightarrow \infty$. P_X represents a track P_T where contributions to the error on a track parameter (X) from intrinsic and multiple scattering are equal.

Track Parameter (X)	$0.25 < \eta < 0.50$		$1.50 < \eta < 1.75$	
	$\sigma_X(\infty)$	$P_X(GeV)$	$\sigma_X(\infty)$	$P_X(GeV)$
Inverse transverse momentum ($1/P_T$)	$0.34 TeV^{-1}$	44	$0.41 TeV^{-1}$	80
Azimuthal angle (ϕ)	$70 \mu rad$	39	$92 \mu rad$	49
Polar angle ($\cot\theta$)	0.7×10^{-3}	5.0	1.2×10^{-3}	10
Transverse impact parameter (d_0)	$10 \mu m$	14	$12 \mu m$	20
Longitudinal impact parameter ($z_0 \times \sin\theta$)	$91 \mu m$	2.3	$71 \mu m$	3.7

Table 2.2: Anticipated track-parameter resolutions (from [3]) for muons (momentum and angular) and pions (impact parameters). $\sigma_X(\infty)$ relates to the resolution at infinite track momentum and P_X relates to a momentum where contributions from multiple scattering are equal to the intrinsic detector resolution.

2.2.2 Calorimetry

Surrounding the inner detector (and solenoid) there is an electromagnetic calorimeter. An incoming particle interacts with material in the calorimeter leading to the production of a particle shower which degrades its energy. The shower is then detected in active regions of the calorimeter and used to reconstruct the original energy of the incident particle. An electromagnetic calorimeter is designed primarily to detect electrons and photons. The incoming particles undergo bremsstrahlung and pair-production creating a growing multi-particle state or electromagnetic (EM) shower. This shower develops until the mean particle energy drops below the pair-production threshold and the particles instead lose energy by ionising the surrounding media. Typically the ionisation is measured by scintillation counters or gas/liquid ionisation detectors.

Both the longitudinal and transverse size of an EM shower are dependent on the material in which it is created and are parameterised by the radiation length (χ_0) and Molière radius (R_M) respectively. The longitudinal size is governed by bremsstrahlung and pair production processes whilst the lateral spread of an EM shower is determined by multiple scattering of low energy shower electrons and the transverse spread of photons whose energy is below the pair-production threshold.

The ATLAS EM calorimetry is a heterogeneous detector with a lead absorber and a liquid argon (LAr) active region. The detector is split into a barrel region, end-caps and a forward detector to provide coverage up to $\eta = 5$ (See Figure 2.4). Heterogeneous detectors are comprised of a multi-layer sandwich of absorbers which create a particle

shower and active regions (samplings) which measure its longitudinal development in order to determine its energy. The active-regions themselves are read out in small regions in η , ϕ known as calorimeter cells which determine the granularity of the detector. The ATLAS EM calorimetry has three longitudinal samplings except in the crack-region between the barrel and end-cap detectors which have two. In order to prevent projective azimuthal cracks and to allow uniform ϕ coverage, the absorbers and active regions are configured in an accordion geometry. The EM calorimeter has a thickness of $\sim 25\chi_0$ (radiation lengths) over its acceptance.

The next layer of detectors beyond the EM calorimetry is the hadronic calorimetry. The basic principle of hadronic and EM calorimetry is the same, the difference being that hadronic calorimetry is designed to measure the energy of hadrons (e.g. pions) particularly in the form of jets (collimated streams of hadrons created by QCD confinement). Rather than undergoing bremsstrahlung the incoming hadrons preferably interact with the nuclei present in the absorber which undergo spallation. The spallation products then undergo further interactions with other nuclei to develop a hadronic shower until the mean energy of particles drops to a point where energy loss through ionisation of the surrounding medium dominates. As a complication, neutral pions (π^0) may be produced which decay into photons that can initiate EM showers within the hadronic one.

The longitudinal size of a hadronic shower is again dependent on the absorber material which is parameterised by the nuclear interaction length (λ). Typically nuclear interaction lengths are an order of magnitude greater than the radiation length for the same material and leads to a requirement for hadronic calorimeters to be thicker in order to contain the particle shower. The lateral size of a hadronic shower is governed by the mean transverse momentum in the hadronic interactions.

Hadronic calorimeters typically have poorer resolution than their EM counterparts due to nuclear binding effects and the presence of neutrons which are poorly detected. The calorimeter's response to the electromagnetic component of the shower is typically greater than that of the hadronic component. This is sometimes corrected for in a compensating

calorimeter by adding nuclei unstable to fission (e.g. U^{238}) under slow neutron capture into the absorber in order to increase detection of the hadronic component.

The hadronic calorimetry at ATLAS is non-compensating and is, like the EM calorimetry, based on a heterogeneous design and split into a barrel, extended barrel, end-caps and a forward detector (see Figure 2.4). In the barrel and extended barrel a tile calorimeter is used with a scintillator as the active medium, whereas the end-caps and forward calorimeter use liquid argon (like the EM calorimeter) which is radiation hard. The hadronic calorimeter has three samplings except for the end-caps which have four, and a thickness of $\sim 10\lambda$. Iron is used as an absorber in the tile calorimeter, copper in the end-caps and copper/tungsten in the forward calorimeter.




Figure 2.4: The ATLAS EM and hadronic calorimetry [3]

Any material can act as an absorber and start an EM or hadronic shower. Typically dense, high Z (atomic number) materials like lead have smaller radiation and nuclear interaction lengths and are used out of choice for absorbers in calorimeters. At ATLAS it is possible for EM and hadronic showers to begin due to interactions of particles in the inner detector volume and solenoid. In order to account for this energy loss, a presampler

(active region) is placed before the absorbers in the EM calorimetry. A summary of the calorimeter systems and their acceptance is given in Table 2.3.

In a sampling calorimeter a particle shower is measured at different points along its longitudinal development. As the development of a shower is a stochastic process a sampling error is introduced to the reconstructed energy of an object. The sampling error decreases with the number of measured particles and hence with the energy of the incident object. An outline of the expected resolutions of the ATLAS calorimetry is found in Table 2.4. In addition to energy resolution, the absolute energy scale of objects reconstructed in the calorimeter has to be calibrated. The ATLAS Technical Design Report [5] states that the jet energy scale is expected to be controlled to $\sim 1\%$ below a transverse momentum (P_T) of 1 TeV and to within $\sim 10\%$ for jets with transverse momentum around 3 TeV⁷. The EM scale is better calibrated and at ATLAS is expected to be known to within $\sim 0.3\%$ [5].

System	Description	η coverage	Samplings
Presampler	Barrel	$ \eta < 1.52$	1
	End-Cap	$1.5 < \eta < 1.8$	1
EM Calorimeter	Barrel	$ \eta < 1.35$	3
	Barrel	$1.35 < \eta < 1.475$	2
	End-Cap	$1.375 < \eta < 1.5$	2
	End-Cap	$1.5 < \eta < 2.5$	3
	End-Cap	$2.5 < \eta < 3.2$	2
Hadronic Tile	Barrel	$ \eta < 1.0$	3
	Extended Barrel	$0.8 < \eta < 1.7$	3
Hadronic LAr	End-Cap	$1.5 < \eta < 3.2$	4
Forward Calorimeter	Forward	$3.1 < \eta < 4.9$	3

Table 2.3: Calorimeter subsystems and acceptance for jets (from [3]).

2.2.3 Muon Spectrometer

Beyond the calorimetry there is a toroidal magnetic field designed to cause charged particles to spiral in the η (θ) plane. As this system is placed beyond the calorimetry it is expected that the only charged particles to traverse this system will be muons. By

⁷Although the precise method of determining the scale at high P_T is not detailed.

System	Region	Energy Resolution
EM Calorimeter	Barrel	$10\%/\sqrt{E} \oplus 0.7\%$
	End-Cap	$10\%/\sqrt{E} \oplus 0.7\%$
Hadronic Calorimeter (Jets)	Barrel	$50\%/\sqrt{E} \oplus 3\%$
	End-Cap	$50\%/\sqrt{E} \oplus 3\%$
Forward Calorimeter		$100\%/\sqrt{E} \oplus 10\%$

Table 2.4: Design specifications for resolution of calorimeter subsystems (from [6],[3])

measuring the curvature of muon paths in the magnetic field their momentum can be determined in an analogous manner to the inner detector.

The toroidal field is provided by a set of superconducting air-core toroids, eight over the barrel region and one for each end-cap. The toroidal field strength varies with pseudorapidity due to the different magnet systems and their overlap. A typical bending power between $2Tm$ and $8Tm$ is achieved. The arrangement of muon detectors (see Figure 2.5) is similar to that of the inner detector. In the barrel region ($|\eta| < 1.0$) the muons are detected in three layers of Monitored Drift Tubes (MDTs) arranged as coaxial cylinders of increasing radius ($5m$, $7.5m$, $10m$). The end-cap chambers are comprised of 4 disks of MDTs and Cathode Strip Chambers (CSC) at distances of 7, 10, 14 and $22m$ from the interaction point.

The MDTs and CSCs are designed to allow for precision measurements of muon momentum, however the muon spectrometer has another important role in the triggering of events (see Section 2.2.4). As an initial trigger decision is needed very rapidly, fast detectors are required, which are provided by a series of Resistive Plate Chambers (RPC) and Thin Gap Chambers (TGC). The RPCs are placed in three layers in the barrel region, two on the middle cylinder (either side of the MDT layer) and one on the outermost cylinder. In the barrel the TGCs are placed in three layers close to the MDT layer at $14m$ from the interaction point. The muon spectrometer is designed to allow for triggering up to a pseudorapidity of $|\eta| < 2.4$.




Figure 2.5: The ATLAS Muon-Chambers [7]

2.2.4 Trigger and DAQ

The proton-proton cross-section for high energy interactions is dominated by elastic, diffractive and low Q^2 non-diffractive inelastic scattering. The processes indicative of new physics (e.g supersymmetry) tend to be considerably less common by many orders of magnitude. At LHC energies ($\sqrt{s} = 14\text{TeV}$) the predicted total proton-proton cross-section is approximately 100mb (based on PYTHIA [8] estimates) which is comprised of elastic scattering, diffractive⁸ and non-diffractive processes. The inelastic non-single diffractive events account for $\sim 70\text{mb}$ of this cross-section which at design luminosity ($L = 10^{34}\text{cm}^{-2}\text{s}^{-1}$) leads to an interaction rate of $\sim 700\text{MHz}$ ⁹. This leads to an average of ~ 23 interactions per bunch crossing which occur every 25ns (assuming a bucket (bunch slot) occupancy of 80%). The presence of multiple non-single diffractive (minimum-bias) events in a single bunch crossing is referred to as ‘pile-up’.

Storing the detector information from a bunch crossing takes around 1MB of disk

⁸The diffractive processes can be split into single and double diffractive processes

⁹The products of an elastic collision tend to escape the ATLAS detector via the beam-pipe and hence do not contribute to pile-up.

space. This leads to a required data storage rate of $40TBs^{-1}$ if all the ATLAS data were to be kept. The quantity of data this represents cannot, therefore, be handled efficiently with current technology and instead a selection process or trigger is used to reduce the number of events stored by a factor of 10^7 . The basic premise is to discard low Q^2 minimum bias events which make up the majority of the interactions at the LHC and retain high Q^2 events which may show signs of new physics.

A trigger needs to be able to make a decision whether to keep an event quickly and in a consistent manner. This is achieved at ATLAS by use of a multi-level trigger where event data is held in pipelines whilst a decision to discard or pass the event up a hierarchy of levels is made. The combined set of three triggers: Level 1, Level 2 and the event filter (LVL1, LVL2, EF) reduce the $40MHz$ bunch crossing to $\sim 100Hz$ for storage (See Figure 2.6).




Figure 2.6: The multi-level trigger [7]

The first trigger, LVL1, is primarily a set of hardware subsystem thresholds and considers events from every bunch crossing in the detector. The time taken to reach a decision takes $\sim 2.5\mu s$ during which the events are held in pipeline memories (usually integrated

circuits situated close to the detector). The LVL1 latency is kept short in order to reduce the quantity of pipeline memories (and hence cost) and means the decision can only be made on a subset of the full event information. LVL1 has access to data from the trigger muon chambers (RPC, TGC) and the entire calorimetry (although with reduced granularity). Within these detector subsystems a search for high P_T objects (e.g. leptons or jets) and missing transverse momentum is carried out. A decision is then made based on a particular object or objects passing a set of basic requirements. LVL1 is expected to pass events at a rate of $\sim 75kHz$, which are then stored in a further buffer before undergoing the LVL2 trigger.

LVL2 is afforded a greater time ($1 - 10ms$) in which to make its decision and is hence implemented as a set of software triggers. Unlike LVL1, LVL2 has access to the entirety of the event data at full precision, however, only data which is necessary to make the LVL2 decision is accessed. This is achieved by passing information on ‘regions of interest’ from the LVL1 to the LVL2 trigger which includes the location of candidate objects in $\eta - \phi$ space. LVL2 is expected to pass events at a rate of $\sim 1kHz$ which are in turn directed to the event filter where the last stage of online selection is made.

The event filter (EF) makes use of the complete event data and more sophisticated algorithms (including vertex and track fitting) in order to confirm the LVL2 decision and to make more subtle cuts. In addition, the event filter has access to information on the detector calibration which can be used owing to reduced time constraints in comparison with LVL2. The EF writes events to mass storage for offline analysis at a rate of $\sim 100Hz$, corresponding to a data rate of $\sim 100MBs^{-1}$.

Despite the strong rejection of events by the trigger it is anticipated that $\sim 1PByte$ of raw event data will be stored by ATLAS each year. In order to make this data more amenable to analysis, a series of processing steps is carried out to reduce its volume. The digitised record of event data is reconstructed to form event summary data (ESD) which is later summarised to analysis object data (AOD) for analysis [9].

2.2.5 Simulation and Offline Analysis

In order to understand the reconstructed events that will be created by ATLAS it is necessary to have a full understanding of the detector's efficiency and any artefacts it may introduce to the data. In order to do this a sophisticated simulation of the detector including its material, geometry and subsystems is developed and used in conjunction with an event generator (such as PYTHIA [8]) to model event data. The simulated output can then be passed through the same reconstruction algorithms as in offline analysis to produce Monte Carlo 'detector' event data which can be compared to experimental results.

The simulation of the detector is a vital part of the experimental concept as it allows the imperfect response of the detector to be taken into account and for a comparison of experimental data with theoretical predictions. At ATLAS the process of simulating events and detector response is contained within the ATHENA framework [10] and is carried out as follows:

Generation An event generator such as PYTHIA is used to simulate a proton-proton collision. The resultant particle's flavour and kinematics are output as a record of 4-vectors.

Simulation The trajectory of each generated particle and the energy deposited in each of the detector subsystems is simulated using GEANT [11].

Digitization The detector subsystem's response to the deposited energy is simulated to give voltages and times as would be measured in a 'real' event. This is in effect the raw event data and is stored in 'DIGI' files.

Reconstruction Algorithms are applied to the raw data in order to form 'physics' objects such as tracks, jets and associated energies. This step is carried out both on simulated and experimental events. The output is typically to ESD or AOD files.

The output of reconstruction does not provide an unambiguous definition of each 'physics' object and as an example an electron candidate may also be a photon candidate etc.

The process of simulating and reconstructing events is CPU intensive. In addition the data volume of reconstructed events alone is expected to be too large for individual scientists to have local access to the entire data set. In order to circumvent these difficulties, the ATLAS collaboration has moved toward a distributed computing framework known colloquially as the ‘grid’ [12]. The philosophy is to provide a framework by which computing facilities (including processors and storage) in remote locations can communicate and share tasks.

The grid is structured as a hierarchy of tiers, where tier 0 is CERN, tier 1 deals primarily with reconstruction of events and tier 2 with analysis. Tiers 1 and 2 are not in a single geographical location and are instead comprised of a number of large computer clusters (or sites) across the world. The aim is to allow a user to define an analysis job which is submitted to the grid. The grid machinery is then expected to take care of allocating which sites perform the job and the process of locating the necessary data which the job requires. After the job is completed the results alone are returned to the user. The advantage with this system is that a more efficient use of computing resources can be made (less dead-time) and a user is generally not constrained by local facilities.

Chapter 3

Theory

3.1 Introduction to QCD

Quantum ChromoDynamics (QCD) is one of three fundamental interactions described in the Standard Model. Also known as the strong force, QCD describes the interaction between particles carrying a ‘colour charge’ called partons, which comprise composite objects known as hadrons (examples of which include the proton, neutron and pion). In what follows a basic introduction to some features of QCD is given in an attempt to encompass aspects relating to jet production. More detailed accounts may be found elsewhere [13, 4, 14].

3.1.1 Features of QCD

The QCD charge ‘colour’ occurs in three distinct types: red, green and blue (r, g, b) and is believed to be a conserved quantity (like electric charge). Colour conservation can be considered as the product of a global symmetry involving the colour degree of freedom within the Lagrangian which describes the dynamics of partons. When this symmetry is promoted to the level of a local symmetry (i.e. dependent upon space-time) it is found that to maintain invariance a gauge field has to be introduced.

The gauge field couples to coloured partons and mediates interactions between them. In QCD this field is identified with massless spin 1 particles known as gluons, which are

exchanged between massive spin 1/2 partons called quarks. This can be seen in analogy to QED where the gauge field consists of neutral photons exchanged between electrically charged electrons. QCD is distinct from QED in that gluons (unlike photons) carry the charge of the interaction (i.e. colour). This is a result of the non-abelian nature of the gauge theory and causes the gluons to undergo self-interactions (i.e. gluons may couple to each other).

Gluon self-interactions fundamentally affect the nature of the strong force and the behaviour of its coupling constant which sets the strength of the interaction. In QCD the coupling strength changes depending on the distance over which the interaction takes place. This can be viewed in the conjugate description of momentum transfer, where a high momentum transfer between particles (Q^2) corresponds to an interaction between the particles at a short distance¹

The running or evolution of the coupling constant occurs in other renormalised gauge theories (such as QED), however in QCD it transpires that as the ‘hardness’ (Q^2) of an interaction increases (i.e. a closer interaction) the coupling constant decreases. In the limit of $Q^2 \rightarrow \infty$ the coupling constant α_s tends to zero, a feature described as asymptotic freedom. Conversely when the interaction occurs with small Q^2 (i.e. a distant interaction) the coupling constant becomes very large due to gluon self interactions. As the coupling constant grows with distance, particles with net colour are effectively trapped by the strong force, a phenomenon referred to as colour-confinement.

Confinement can be understood at a qualitative level by considering strong interactions between a quark and anti-quark carrying a colour and anti-colour respectively. If the quark and anti-quark were gradually separated the strength of the strong interaction between them would increase owing to the running of the coupling constant. The attractive force between the partons increases with distance (unlike the other fundamental forces) preventing either parton from emerging as a free particle. The individual coloured partons are hence confined in colour neutral hadrons. If sufficient energy is introduced to the quark-antiquark system, it can be transferred into particle creation along the axis of

¹This can be considered as a consequence of Heisenberg’s uncertainty principle $\Delta x \Delta P \geq \frac{\hbar}{2}$.

interaction between the partons (see Figure 3.1). The separating quark and anti-quark are observed experimentally as a collimated beam of many colourless particles (predominantly mesons) referred to as a jet.




Figure 3.1: Confinement and particle production due to QCD. The curly lines represent gluons.

3.1.2 Hadron-Hadron Interactions

Hadron-hadron colliders (such as the Tevatron and LHC) are frequently used to explore high energy physics. The colliding hadrons (usually protons) are not fundamental particles but are instead composite objects made of partons interacting via QCD. As may be anticipated from the feature of confinement, the hadron itself carries no net-colour charge² and hence does not undergo direct QCD interactions with other hadrons. However, in high energy collisions between composite objects it is the constituents that undergo interactions rather than object as a whole (e.g. Rutherford scattering). As a result in a hadron-hadron collision, the coloured partonic constituents can undergo strong interactions despite the hadrons being colour neutral.

Colliders create hard (high Q^2) parton-parton interactions where QCD is perturbatively calculable owing to the small coupling constant. This allows elastic scattering cross-sections for parton-parton interactions to be calculated for a given initial state of particles.

²The colour-singlet nature of hadrons can also be explained in terms of the QCD gauge theory [13].

Following a hard parton-parton interaction, a parton is elastically scattered from a hadron. The parton and the system it has been scattered from both have a colour charge and hence non-perturbative confinement effects take place. If the energy of the scattering is sufficient, confinement leads to particle creation between the scattered parton and the system it was scattered from. As mentioned earlier this ‘hadronization’ is seen experimentally, as a collimated jet of particles emerging in the direction of the initial scattered parton.

Hadronization is by its nature a low Q^2 effect and cannot be described by perturbative QCD. Instead sophisticated phenomenological models such as the ‘Lund string model’ are used in Monte Carlo simulation programs (e.g PYTHIA [8]) to take such effects into account when creating theoretical predictions.

Although the elastic scattering of the partons may be calculated perturbatively, this demands a knowledge of the initial state (i.e. the incoming partons and their kinematics). This is not known *a priori* at a hadron collider as hadrons themselves are QCD objects consisting of multiple partons in a bound state under the strong interaction. The bound state again involves low Q^2 interactions between partons which cannot be described perturbatively. Hence the partonic composition of the hadron must be parameterised.

The structure of a hadron is non-trivial, for example in a proton there are three valence quarks (u,u,d) that supply the overall quantum numbers of the hadron which can be considered to reside in a ‘sea’ of quarks, antiquarks and gluons all of which are undergoing low Q^2 non-perturbative QCD interactions. The other category of hadrons, the mesons, are in effect identical except the valence partons are a quark and an antiquark. As a result of the interactions the partons are in a constant state of change as regards the fraction they carry of the energy or momentum of the overall hadron.

The hadron structure is parameterised by introducing parton distribution functions or PDFs. The PDFs describe the composition of a proton in terms of the probability of finding a parton of a given flavour carrying a certain fraction of the hadron’s momentum. In a similar manner to the strong coupling constant the PDFs also change (or evolve)

depending on the hardness (Q^2) of an interaction. This effect known as ‘scaling violation’ can be described using perturbative techniques by the Altarelli-Parisi or DGLAP equations [13]. By combining the PDFs with the elastic scattering parton cross-sections and a hadronization model the hadron-hadron jet cross-section may be calculated³.

PDFs cannot be calculated from first principles using QCD and instead have to be parameterised using experimental data. As PDFs only concern proton structure it is sufficient to determine the PDFs in a set of experiments independent to the one in which they are applied to make a theoretical prediction.

3.2 QCD Calculations

3.2.1 The Inclusive Jet Cross-Section at Leading Order

To demonstrate some of the features of QCD previously described we shall consider the specific process of inclusive jet production at hadron colliders and outline how a cross-section may be calculated. The inclusive jet cross-section is the probability of finding a jet in an event at a given transverse momentum (P_T) and is often defined within a given range in pseudorapidity (e.g $0 < |\eta| < 1$).

As described earlier, hadrons are composite objects however a high Q^2 interaction of protons can be approximated as an interaction between two partons (one from each incoming hadron). The selection of the flavour and momentum of the interacting partons may be made by appealing to the PDFs which parameterise the structure of each hadron. This leads to the possibility of many different initial states each of which must be individually taken into account when considering the total jet cross-section.

Assuming that the selection of incoming partons and their kinematics have been made the next step is to calculate the cross-section for their hard scattering. This is traditionally carried out by performing a fixed order calculation using Feynman diagrams. The technique involves drawing all possible Feynman diagrams that connect the initial to the

³It is possible to neglect a parton shower and hadronization model to arrive at a jet cross section at the ‘parton level’.

final state. For an inclusive jets cross-section calculation, this involves all final states with coloured partons (that could form jets) and the range of possible initial states as defined by the PDFs. The formalism relates the diagrams to an expression for a matrix element of the overall process. The dominant contributions to jet production at leading order (LO) are shown in Figure 3.2.




Figure 3.2: Dominant QCD amplitudes for jet production at LO (t-channel), curly lines represent gluons, plain lines are quarks or antiquarks.

Figure 3.2 is by no means an exhaustive list of contributions (even at LO). Additional QCD contributions from s-channel processes occur, but these are typically smaller at hadron colliders than their t-channel equivalents from kinematic considerations. The QCD coupling between partons is typically several orders of magnitude greater than the weak or electromagnetic interactions allowing the latter to be neglected as a first approximation⁴.

The t-channel contributions to the squared matrix element given in Figure 3.2 are all of a similar form (Eqn 3.1):

$$|M_{i,j}|^2 \propto \frac{\hat{s}^2 + \hat{u}^2}{\hat{t}^2} \quad (3.1)$$

where $\hat{s}$, $\hat{t}$ and $\hat{u}$ are the Mandelstam variables for the partons undergoing the interaction. The form is in fact similar to Rutherford scattering in that $\hat{t}^{-2} \sim \sin^{-4}(\theta/2)$, where θ is the angle to the beam axis that the outgoing parton form in their centre of

⁴As a result, diagrams of photon exchange between quarks, for example are not considered.

mass frame⁵.

The schematic LO calculation can then be expressed as follows, first the partonic cross-sections are calculated by multiplying the squared matrix elements with the phase space for the final state.

$$\hat{\sigma}_{i,j} \propto |M_{i,j}|^2 \times \text{phasespace} \quad (3.2)$$

Where i and j denote the flavour of final and initial state partons. The partonic cross-sections are then convoluted with the PDFs of the incoming hadrons to take into account the spectrum of possible initial states.

$$\sigma(P_1, P_2) = \sum_{i,j} \int \int dx_1, dx_2 f_i(x_1, \mu_f^2) f_j(x_2, \mu_f^2) \hat{\sigma}_{i,j}(x_1 P_1, x_2 P_2, \alpha_s(\mu_r^2), Q^2/\mu_f^2) \quad (3.3)$$

Where $P_{1,2}$ are the momenta of the incoming hadrons and $x_{1,2}$ are the fraction of this momentum carried by the partons undergoing the interaction (μ_r and μ_f are defined in section 3.2.3). It may be noted from Eqn 3.1, that the contribution from the processes in Figure 3.2 have a pole as $t \rightarrow 0$ which corresponds to the outgoing particles being collinear with the incoming ones. The related s-channel processes (not shown) have a similar pole as $s \rightarrow 0$ which corresponds to a soft-collision between the incoming particles. These divergences are referred to as infra-red (IR) singularities and are the product of soft processes. The problem of IR divergences can be avoided by a phase-space cut on the kinematics e.g. applying a minimum cut on the P_T of outgoing particles. Such cuts can be justified on the grounds that experiments are generally insensitive to soft and collinear scatterings.

3.2.2 Next-to-Leading Order Corrections

There are more QCD contributions to be considered than are shown in Figure 3.2 besides the neglected s-channel processes. Some examples are shown in Figure 3.3. The contri-

⁵In a colliding beam experiment the parton-parton centre of mass frame may be boosted with respect to the lab-frame along the beam axis.

butions in this figure contain an additional gluon line, either as an additional outgoing particle (real emission), or as an internal line creating a loop (virtual correction). These next-to-leading order (NLO) diagrams contain at least one additional QCD vertex and hence are expected to be suppressed by a factor close to the coupling constant α_s in comparison with the LO terms in Figure 3.2.




Figure 3.3: Some QCD amplitudes for Jet production at NLO, curly lines represent gluons, plain lines are quarks or antiquarks.

The loops in the first two diagrams in Figure 3.3 have an undefined momentum (k) associated with them, which needs to be integrated. The integral is however logarithmically divergent in k and there is no natural way to avoid the $k \rightarrow \infty$ limit. These divergences which occur at high k are referred to as ultra-violet (UV) divergences. Some of the virtual divergences cancel with contributions from real emission (e.g. the last diagram in Figure 3.3) but to remove the remaining UV singularities a process of renormalisation has to be used.

The process of renormalisation relies on the fact that observables such as the charge and mass of a particle undergo corrections due to diagrams like those found in Figure 3.3. As a result the ‘bare’ parameters of mass and charge are not measurable properties and can hence be divergent. Renormalisation attempts to recast the problem in terms of measurable quantities (i.e. measured charge and mass) and to push the divergences onto the unmeasurable bare parameters in a consistent manner.

As a general rule a cross-section calculation is more accurate if a greater number of contributions (or orders) are considered. However, as higher order corrections are considered, the number of Feynman diagrams increase in number and complexity. The calculational difficulties mean that for a large number of processes only LO and NLO predictions are available.

3.2.3 Higher Orders and Unphysical Scales

The fact that only a few orders of corrections can be calculated in QCD is a problem as the interaction is strong and higher-order corrections may not be negligible. In fact the corrections are infinite and result in UV (ultra-violet) divergences which have to be treated by a process of renormalisation. This process leads to the introduction of a renormalisation scale μ_r which then appears in logarithmic terms relating to the correction of the QCD coupling due to loop corrections. These corrections are of the form:

$$\propto \ln(Q^2/\mu_r^2) \quad (3.4)$$

When predicting a physical quantity using the fixed-order theory a choice of the renormalisation scale has to be made⁶. The form of equation 3.4 suggests that a choice of $\mu_r^2 \sim Q^2$ reduces the strength of the loop corrections and hence will improve the accuracy of the fixed-order calculation by reducing the effect of neglected terms.

The formalism of perturbation theory requires that the interaction being described is of a weak nature in order that higher-order corrections can be neglected. As a consequence, accurate predictions may be made in QCD when the coupling constant is small (i.e. high Q^2) but not when it is large (i.e. low Q^2). It is not possible to circumvent this difficulty by arbitrarily choosing the renormalisation scale μ_r (i.e. to be much larger than Q^2) as this leads to large logarithmic corrections to predictions from neglected terms.

In a typical interaction between hadrons in a collider experiment both high and low Q^2

⁶An all orders calculation would not be dependent on μ_r as all quantum corrections would be taken into account.

processes are present. In order to benefit from the predictive power of perturbative QCD, the non-perturbative low Q^2 components are parameterised away in phenomenological models (such as PDFs). This approach relies on the factorisation theorem which allows the hard and soft components of an interaction to be separated by introducing an unphysical scale, μ_f (the factorisation scale) that defines the boundary between the two.

The factorisation scale again results from the need to regularise divergences from higher-order corrections. In this case the higher-order terms lead to ‘scaling violations’ which result in a change in the momentum spectrum of partons that make up hadrons (parton evolution). These corrections take a similar form to that of equation 3.4 with their impact again being reduced by choosing a factorisation scale $\mu_f^2 \sim Q^2$ as before.

3.2.4 Resummation and Parton Showers

In certain circumstances it is possible to model multiple soft emissions by the use of resummation techniques such as the leading-log approximation (LLA). This is often implemented in Monte Carlo simulations by use of a parton shower which models the effects of additional soft (gluon) emissions⁷. Soft radiation emitted in a direction close to the beam line (in a collider) is referred to as ‘Initial State Radiation’ (ISR) with its counterpart ‘Final State Radiation’ (FSR) being identified as soft radiation around the axis of the outgoing hard-scattered particles. However, it should be noted there is no formal distinction between the two, as the process is intrinsically a quantum mechanical one.

A parton shower allows a parton at a high energy scale to be evolved down to a scale close to where QCD becomes non-perturbative by emitting further partons. At this point a process referred to as fragmentation or hadronization (involving low Q^2 QCD interactions) takes over to arrange the partons into colourless hadrons. The fragmentation process generally involves the production of extra quarks and anti-quarks as seen before in Figure 3.1. As fragmentation necessarily involves non-perturbative physics, it is implemented in Monte Carlo simulations by the use of phenomenological models such as colour strings or

⁷Additional hard-scatterings are not modelled in a ‘Leading-Log Approximation’ (LLA).

clusters.

Due to the complexity of the integrals involved, numerical techniques are often required to obtain numerical results from cross-section calculations especially if cuts for experimental acceptance are considered. The most efficient numerical method for integrals of high dimension is the Monte Carlo technique which relies on random sampling of an integrand to approximate its value. There are a number of available computer packages which combine a variety of different physical processes at LO with a parton shower (e.g. PYTHIA [8] and HERWIG [15]). NLO processes are harder to generalise for numerical calculations and often a bespoke program is required for a given process (e.g. inclusive jet production). This has led to the development of a large number of process-specific NLO cross-section generators (e.g. NLOJET⁺⁺ [16], JETRAD [17]). Applying a parton shower approach to NLO calculations can prove problematic as there is a risk of ‘double counting’, i.e. that a NLO semi-hard gluon emission is already accounted for in the parton shower. Attempts to resolve the double counting issue have been made, notably MC@NLO [18] which combines NLO calculations with a parton shower and hadronization for a few processes.

3.3 QCD Coupling and PDFs

The theoretical framework as outlined above has two elements that require experimental measurements before a quantitative prediction can be made. These are the strong coupling constant α_s and the non-perturbative PDFs.

3.3.1 Determining the Strong Coupling Constant

Local gauge theories such as QCD describe an interaction up to a particular coupling constant, which determines the intrinsic strength of the interaction. QCD predicts how this constant evolves in Q^2 but an experimental measurement has to be made at a given Q^2 before this is possible.

There are many possible ways of determining α_s from experiment. A measurement of α_s can be made at e^+e^- colliders from rates of jet production, τ decay, and event shapes. In addition at lepton-hadron colliders (like HERA) α_s may be determined from scaling violations and momentum sum rules. A summary of α_s measurements and predictions is given in Figure 3.4 and a full description of these methods is available from the Particle Data Group [19].




Figure 3.4: Summary of $\alpha_s(M_Z)$ from Particle Data Group (2006) [19]. The errors shown include theoretical uncertainties and an average gives $\alpha_s(M_Z) = 0.1176 \pm 0.0009$.

3.3.2 Determining the PDFs

Unlike α_s , the parton distribution functions may not be measured directly, instead a parameterised set of equations describing the probabilities of finding different flavours of parton in a hadron (generally a proton) is postulated and the physical consequences of this choice calculated in terms of cross-sections. The basic form of the PDF parameterisation for quarks and gluons common to most analyses is given below:

$$F(x, Q_0) = A_0 x^{A_1} (1 - x)^{A_2} \cdot P(x; A_3 \dots) \quad (3.5)$$

Where x refers to the momentum fraction of the proton carried by the parton and A_0 ,

A_1 etc are free parameters which are adjusted in a global fit to provide the best agreement with experimental data. For valence quarks and gluons A_0 , A_1 and A_2 tend to be positive for $Q \sim 1\text{GeV}$. The poles at $x = 0$ and $x = 1$ relate to Regge behaviour and quark counting rules respectively. The $P(x; A_{3..})$ is a smooth function designed to introduce flexibility in the PDF fitting process and tends to be unique to a particular set of PDFs, fit by a given group (e.g. CTEQ or MRST). The parameterised PDFs can be evolved via the Altarelli-Parisi equations to different values of Q^2 as required to create a prediction of experimental data.

An example set of PDFs for the proton's up, down, anti-up and gluon partons is given in Figure 3.5. As can be seen the constituents of the proton generally carry a small fraction of its momentum, and the PDFs all tend rapidly to zero as $(x \rightarrow 1)$. The gluon distribution has a tendency to dominate the PDFs except at high x .




Figure 3.5: PDFs for the up, down, anti-up quarks and the gluon from CTEQ6.1. $Q_0 = 10\text{GeV}$ and the gluon distribution has been scaled by 0.1 for ease of comparison [20]

PDFs show that the centre of mass energy of two colliding hadrons (e.g. $\sqrt{s} = 14\text{TeV}$ at the LHC) is not the same as that of the two partons undergoing the interaction, which, with reference to Figure 3.5, will typically be much lower. To a first approximation the shape of the inclusive jet cross-section is determined from basic considerations of scattering (e.g. Rutherford) and PDFs. The production of high P_T jets is suppressed

as they require a large-angle parton scattering (high Q^2) where the interacting partons carry a large fraction of the hadron's momentum. This explains the exponential fall in the inclusive jet cross-section as P_T increases.




Figure 3.6: Inclusive jet cross-section for a proton-proton collider with $\sqrt{s} = 14\text{TeV}$ (LHC), CTEQ6.1 PDFs for three regions in pseudorapidity, generated by NLOJET⁺⁺.

The data most commonly used to constrain PDFs is gathered from deep-inelastic scattering (DIS) experiments, which involve scattering leptons from the hadron in question (proton). The interaction is electromagnetic in nature i.e. the electric charge of the lepton interacts with that of the quarks within the hadron. Measurement of the hadron's electromagnetic structure functions is useful as they can be parameterised in terms of its PDFs. Similar experiments may be carried out with neutrino scattering from hadrons, however such measurements are considerably more difficult owing to the necessity of detecting weakly interacting neutrinos.

DIS experiments carried out at fixed target and colliding-beam experiments have provided accurate data on hadron structure over a range of Q^2 values. Hadron collider data can extend the range to higher Q^2 , however historically it has been considered too inaccurate to provide constraints on PDFs in global fits. The data does provide useful information on the gluon distribution which is poorly constrained in DIS experiments as the gluon is not electrically charged. Indeed, it was found that in order to explain jet cross-

section results at the Tevatron collider an adjustment to the gluon's high- x distribution was required [21].

The determination of PDFs is important for the LHC as they enter into the theoretical prediction of all cross-sections. As a result, uncertainty in the parton content of the proton leads directly to an uncertainty in the precision of theoretical predictions and limits the conclusions that can be drawn with comparisons to data. The subject of PDF uncertainties and constraining the gluon PDF with collider data will be discussed in later chapters.

3.4 Jet Algorithms

3.4.1 The role of jet algorithms

As described earlier a parton scattered with sufficient energy from its parent hadron will lead to the production of a multi-particle final state (due to QCD confinement), referred to as a jet. As complete information on a jet is both difficult to obtain experimentally and predict theoretically, a jet algorithm is used to summarise the broad features of this complex state into a few key parameters. The role of a jet algorithm is to identify which particles (either reconstructed experimentally or from Monte Carlo simulation) belong to jets and to provide a formalism for adding their energies and momenta to arrive at an overall jet energy and direction.

The parton model predicts jets as the longitudinal energy (as measured along the scattered parton's trajectory) is free to increase with the centre-of-mass energy of a collision. However, the transverse momenta of fragmentation products (with respect to the hard-scattered parton) arise from soft QCD processes and tend to remain small. This results in a multiple particle state that becomes increasingly collimated in $\eta - \phi$ as the energy of the interaction increases (a jet). The overall jet energy and trajectory of a jet reflect that of the parton which originally emerged from the hard scattering as the fragmentation products effectively share its energy.

At colliders, the most common events are dijet-like events where two back-to-back jets (in ϕ) are produced and reflect the paths of the two hard-scattered partons (in a $2 \rightarrow 2$ scattering). Due to higher order contributions (e.g. NLO and above) there are also events with a greater number of jets in the final state. The role of the jet algorithm in both the experimental and theoretical context is hence to provide a consistent method of reconstructing jets regardless of the event topology.

The jet algorithm has to identify which particles belong to a given jet and afterwards apply a method to reconstruct the kinematics of the jet given its constituents (recombination). An ideal jet algorithm should be well understood theoretically and insensitive to the effects of soft radiation (infra-red safety). In experiments, only a subset of all final state particles are identified and reconstructed. As a result, jet measurements focus on calorimetry in order to measure the sum of energy emitted in a given region (a segment) in $\eta - \phi$ space. An ideal jet algorithm must be able to function when these calorimeter segments are used as an input instead of reconstructed particles.

3.4.2 Examples of jet algorithms

Two basic types of jet algorithm have emerged, the cone and K_T . The cone algorithm relies on the clustering of jet particles within a narrow region in $\eta - \phi$ space. A circle⁸ of fixed radius ' R ' is first defined in $\eta - \phi$ space (typically $R \sim 0.4 - 0.7$). Particles (or calorimeter segments) within the circle are considered to be members of the jet, whereas those that remain outside are not, but may be members of other jets. The direction of a jet is found by starting with a trial circle (in $\eta - \phi$) and calculating an energy weighted centroid by considering all the particles contained within. This centroid is then used as the focus of a new circle and the process is iterated until a stable solution is found. Once a stable jet is found the kinematics of the constituents are combined (recombination) to obtain global values for the jet.

The K_T algorithm works by combining particles of similar momentum. It begins with

⁸This circle forms the base of a cone with the interaction point as the apex.

a list of energy clusters (or preclusters) which are assigned a transverse momentum (P_T), pseudorapidity (η) and azimuthal angle (ϕ). For each cluster a value d_i is defined as:

$$d_i = P_{T_i}^2 \quad (3.6)$$

For each pair of preclusters (i,j) a value $d_{i,j}$ is calculated.

$$d_{i,j} = \min(P_{T_i}^2, P_{T_j}^2) \frac{\Delta R_{i,j}}{D^2} \quad (3.7)$$

Where $R_{i,j}$ is the separation of the clusters in $\eta - \phi$ space and D^2 is a parameter of the K_T algorithm akin to the cone size of the Cone algorithm. The minimum d_{min} of all d_i and $d_{i,j}$ is calculated. If $d_{min} \in d_{i,j}$ then the clusters i and j are merged by adding their 4-momenta. Otherwise, $d_{min} \in d_i$ and the cluster is removed from the list of clusters and defined as a jet. The process is iterated until no clusters remain.

Practical jet algorithms suffer from a range of complications. In cone algorithms jet cones may overlap and a procedure has to be defined for splitting/merging the two objects. Furthermore, if the cone size in the algorithm is too small then jet energy can be lost (out-of-cone corrections). On the other hand if the cone is too large, then the jet algorithm can cluster energy which belongs to the underlying event (separate from the hard scattering) or merge distinct jets. Underlying event problems are also present in the case of K_T algorithms and can be harder to correct for.

In certain circumstances cone algorithms can fail to cluster energy deposits that occur nearby into a more energetic deposit. In the absence of the neighbouring energy cluster however, the ‘dark-tower’ would have been reconstructed as a jet. These ‘dark-towers’ are the result of the strong attraction of cone centres to the larger energy deposit and the spatial resolution of a jet’s energy due to fragmentation/hadronization. Basic cone algorithms can also suffer from being infra-red unsafe, as the addition of soft radiation between cones can lead to their merger. These problems can, however, be partially circumvented by using midpoints between jets as seeds for cones (e.g. the Midpoint algorithm).

Despite its drawbacks, the cone algorithm has traditionally been preferred in hadron

collider experiments as it was perceived that its jets are easier to calibrate and that the algorithm converged more quickly than its K_T counterpart. A standard implementation of the K_T algorithm has a convergence time which scales as N^3 where N is the number of particles or preclusters considered. However, recent advances [22] have reduced the convergence time to scale as $N\ln(N)$ which is typically faster than the cone algorithm.

By combining the elements of a perturbative QCD parton scattering with PDFs, a hadronization model and a jet algorithm to simplify the final state; a full and quantitative prediction of a jet cross-section can be made at a hadron collider. All of these ingredients have associated errors which need to be understood when comparing theory to experiment and some of these will be investigated in later chapters. In addition to uncertainties on predictions, there are also difficulties associated with the length of CPU time that a fixed order (NLO) cross-section (at parton level) entails. One such problem, arising from the need to quickly recalculate cross-sections in global PDF fits, is discussed in the following chapter.

Chapter 4

Grid Methods for NLO

Cross-Sections

4.1 Introduction to PDFs and Global Fits

Jet cross-sections at LHC energies play an important role in tests of QCD as well as searches for new physics such as compositeness. Experimental results are compared with Standard Model predictions and discrepancies can be indicative of new physics, provided that both experimental and theoretical errors are sufficiently understood. The dominant source of theoretical error (for high Q^2 events) arises from the use of PDFs needed to describe the parton content of the colliding protons.

Predictions of jet cross-sections at hadron colliders typically require NLO calculations in order to accurately recreate experimental results. The calculations are implemented as Monte Carlo programs in order to take account of experimental cuts and can take days (or even weeks) of CPU time. If the PDF is changed, the entire cross-section must be recalculated, which has historically limited the use of jet cross-section data in global PDF fitting schemes.

One method of avoiding this difficulty is to use leading order (LO) calculations with ‘k’ factors ($\sigma(NLO)/\sigma(LO)$) to approximate the NLO result. As LO calculations are simpler than NLO, this can reduce the CPU time in a global PDF fit to an acceptable

level. However, ‘k’ factors are dependent on PDFs and quickly produced LO cross-sections remain susceptible to statistical errors. An alternative method of providing a true NLO prediction for an arbitrary PDF is to employ an integration grid.

4.2 Integration grids for QCD calculations

An integration grid attempts to separate the PDF from the calculation of the jet cross-section. The method below follows that given in [1].

A typical cross-section for a $2 \rightarrow 2$ hadron scattering was given in Eqn 3.3 and is repeated below:

$$\sigma(P_1, P_2) = \sum_{i,j} \int dx_1 dx_2 f_i(x_1, \mu_f^2) f_j(x_2, \mu_f^2) \hat{\sigma}_{ij}(x_1 P_1, x_2 P_2, \alpha_s(\mu_r^2), Q^2/\mu_r^2) \quad (4.1)$$

Where f_i is the PDF for parton flavour i and $\hat{\sigma}_{i,j}(\dots)$ is the parton level cross-section. The $x_{1/2}$ relate to the momentum fraction of the colliding protons carried by the two partons undergoing the scattering, Q^2 is the momentum transfer between the partons (hardness) and $\mu_{r/f}$ is the renormalisation or factorisation scale¹. As this is a QCD cross-section, α_s is the strong coupling constant and its dependence on μ_r is a result of the running of this constant.

As can be seen the PDFs are inside the integral over the (Bjorken) x values of the partons undergoing the scattering. The first step to removing the PDFs from this integral is to represent the PDF in terms of a discrete set of points on a grid.

4.2.1 Using a grid to represent a PDF

The PDFs depend on the flavour of the parton in question, its x value and the factorisation scale which is chosen to be around the Q^2 of the parton interaction². To describe a

¹The choice of scale is always taken to be the same (i.e. $\mu_r = \mu_f$)

²An explanation of this is given in section 3.2.3.

particular parton PDF, a two dimensional array is created in terms of transformed x and Q^2 co-ordinates. The grid (over a given range of x and Q^2) is subdivided into bins that are equidistant in terms of the new transformed variables. The aim of this transformation is to allow the finite grid to better describe the kinematic regions likely to be encountered in a hadron-hadron collision.

The co-ordinate transformations used (from [1]) are:

$$y(x) = \ln \frac{1}{x}, \quad (4.2)$$

$$\tau(Q^2) = \ln \ln \frac{Q^2}{\Lambda^2}. \quad (4.3)$$

Where Λ is of order Λ_{QCD} which represents the Q^2 scale where QCD interactions become non-perturbative. These transformations result in an increased representation of the PDF at low x and low Q^2 which is characteristic of the majority of jet production at the LHC. The PDF is now defined (from [1]) as³:

$$f(y, \tau) = \sum_{i=0}^n \sum_{j=0}^{n'} f(y_{k+i}, \tau_{k'+j}) I_i^n \left(\frac{y - y_{min}}{\delta y} - k \right) I_j^{n'} \left(\frac{\tau - \tau_{min}}{\delta \tau} - k' \right) \quad (4.4)$$

Where $I_i^n(u) = 1$ for $u = i$ and otherwise given by:

$$I_i^n(u) = \frac{(-1)^{n-i}}{i!(n-i)!} \frac{u(u-1)\dots(u-n)}{u-i} \quad (4.5)$$

and,

$$k(y) = \text{int} \left(\frac{y - y_{min}}{\delta y} - \frac{n-1}{2} \right), k'(\tau) = \text{int} \left(\frac{\tau - \tau_{min}}{\delta \tau} - \frac{n'-1}{2} \right) \quad (4.6)$$

$\text{int}(m)$ is defined as the largest integer smaller than m , i.e. $\text{int}(m) \leq m$

The sum in equation 4.4 is over a region of discrete grid points around the desired y and τ . The number of grid points considered depends on the order of interpolation

³The interpolation is based on Lagrange polynomials.

used (in each dimension) n and n' . The grid's granularity is defined by δy and $\delta\tau$ which denote grid spacings. The lower edge of the finite grid is given by y_{min} and τ_{min} in the variables y and τ respectively. The $I_{i/j}^n$ are interpolation functions whose role is to weight the contribution of a given discrete grid point according to how close it is to the desired y and τ that the PDF is being evaluated at.

As the interpolation should ideally use only grid points that are defined the following transformations are applied:

$$k \rightarrow \max(0, \min(N_y - 1 - n, k)), k' \rightarrow \max(0, \min(N_\tau - 1 - n', k')) \quad (4.7)$$

Where N_y and N_τ are the number of y and τ bins in the grid respectively, which are indexed from 0 to $N_{y/\tau} - 1$. The basic procedure is shown graphically in Figure 4.1.




Figure 4.1: A PDF described by interpolation on a grid.

4.2.2 Example of Grid Method

In order to demonstrate the grid procedure, it is instructive to look at a worked example. For simplicity, a one dimensional example using only the y variable is considered. As a result all terms relating to τ and its index j can be ignored in eqn 4.4.

Consider a smooth function $f(y) = 2y$ to be described by a finite grid with $y_{min} = 3$,

$N_y = 4$ with a grid spacing $\delta y = 1$. This is shown in Figure 4.2.




Figure 4.2: A one-dimensional grid with $N_y = 4$, $\delta y = 1$ and $y_{min} = 3$ to describe the function $f(y) = 2y$.

Suppose that we want to determine the value of $f(y)$ at $y = 4.5$, a point between the grid spacings. It is clear in this example that the value is simply $2y = 9$, but to demonstrate the grid machinery we will try to estimate a value using the interpolation method with $n = 1$ (2nd-order).

The calculation we wish to perform is:

$$f(4.5) = \sum_{i=0}^1 f(y_{k+i}), I_i^1 \left(\frac{4.5 - y_{min}}{\delta y} - k \right) \quad (4.8)$$

From equation 4.6 we can calculate $k(y)$ as:

$$\begin{aligned} k(y) &= \text{int} \left(\frac{y - y_{min}}{\delta y} - \frac{n-1}{2} \right) \\ &= \text{int} \left(\frac{4.5 - 3}{1} - \frac{1-1}{2} \right) \\ &= \text{int}(1.5) \\ &= 1 \end{aligned}$$

Thus I_0^1 can be found using $k(y)$ and equation 4.5:

$$\begin{aligned}
I_0^1(u) &= I_0^1 \left(\frac{4.5 - y_{min}}{\delta y} - k \right) \\
&= I_0^1 \left(\frac{4.5 - 3}{1} - 1 \right) \\
&= I_0^1(0.5) \\
&= \frac{(-1)^1}{0!(1-0)!} \frac{0.5(0.5-1)}{0.5} \\
&= 0.5
\end{aligned}$$

Similarly for I_1^1 :

$$\begin{aligned}
I_1^1(u) &= I_1^1(0.5) \\
&= \frac{(-1)^0}{1!(1-1)!} \frac{0.5(0.5-1)}{-0.5} \\
&= 0.5
\end{aligned}$$

The final ingredients to Eqn 4.8 are the grid values at the discrete points $f(y_{k+i})$. As $k = 1$ the grid indices $k + i$ in the sum are 1 and 2. Referring to Figure 4.2 and recalling that the grid-indices are numbered from 0 to $N - 1$ we can see that the grid points being considered are those at $y = 4$ and $y = 5$. This is to be expected as these are the nearest discrete grid points to $f(4.5)$. From the figure it is clear that $f(y_{k+i}) = 8$ for $i = 0$ and 10 for $i = 1$.

Combining all the elements into equation 4.8 we obtain:

$$\begin{aligned}
f(4.5) &= \sum_{i=0}^1 f(y_{k+i}), I_i^1 \left(\frac{4.5 - y_{min}}{\delta y} - k \right) \\
&= (8 \times 0.5) + (10 \times 0.5) \\
&= 9
\end{aligned}$$

So in this case the interpolation procedure accurately determines the function $f(y)$ at the point $y = 4.5$. This is perhaps unsurprising as in this example $f(y)$ is a simple linear function and the interpolation was carried out at 2nd order.

4.2.3 Recording cross-section weights on a grid

Consider a Monte Carlo calculation of a deep-inelastic scattering (DIS) process at leading order. This simplifies the problem in that only one PDF is required (to describe the target hadron) and NLO corrections can be neglected. The Monte Carlo generates N events, with weight $w = w_m$, at a given $x = x_m$ and $Q^2 = Q_m^2$. The (x, Q^2) can be transformed into (y_m, τ_m) using Eqns 4.2 and 4.3; the cross-section is then given by:

$$\sigma = \sum_{m=1}^N \left(\frac{\alpha_s(\tau_m)}{2\pi} \right)^2 w_m f(y_m, \tau_m) \quad (4.9)$$

This equation continues to face the original problem of the cross-section formula (Equation 4.1), in that the PDFs must be calculated on a per-event basis. However, the PDFs in Eqn 4.9 can be replaced by the version based in terms of the integrated grid points (Equation 4.4), giving:

$$\sigma = \sum_{m=1}^N \left(\frac{\alpha_s(\tau_m)}{2\pi} \right)^2 w_m \left(\sum_{i=0}^n \sum_{j=0}^{n'} f(y_{k+i}, \tau_{k'+j}) I_i^n \left(\frac{y_m}{\delta y} - k \right) I_j^{n'} \left(\frac{\tau_m}{\delta \tau} - k' \right) \right) \quad (4.10)$$

At this stage it is possible to define a two dimensional weight-grid W_{i_y, i_τ} in analogy to the discrete grid used for the PDF earlier (where the dimensions relate to y and τ). where $i_y = k + i$ and $i_\tau = k' + j$ which is independent of the PDFs:

$$W_{i_y, i_\tau} = W_{k+i, k'+j} = \left(\sum_{m=1}^N \sum_{i=0}^n \sum_{j=0}^{n'} w_m I_i^n \left(\frac{y_m}{\delta y} - k \right) I_j^{n'} \left(\frac{\tau_m}{\delta \tau} - k' \right) \right)_{k+i, k'+j} \quad (4.11)$$

The indices of the weight grid are defined as $i_y = k + i$ and $i_\tau = k' + j$. Crucially this weight-grid is independent of the PDF. A corresponding two-dimensional PDF grid can be defined as before in terms of the general indices

$$f_{i_y, i_\tau} = f(y_{i_y}, \tau_{i_\tau}) \quad (4.12)$$

Thus for a given co-ordinate (i_y, i_τ) the corresponding y and τ values are defined to be identical to those in the weight-grid.

This allows Eqn 4.10 to be recast as:

$$\sigma = \sum_{i_y} \sum_{i_\tau} W_{i_y, i_\tau} \left(\frac{\alpha_s(\tau_{(i_\tau)})}{2\pi} \right)^2 f_{i_y, i_\tau} \quad (4.13)$$

Where the sum is now over all grid points. Equation 4.13 forms the basis of the integration grid method. The parton cross-section information is contained within the weight-grids W_{i_y, i_τ} , which can now be calculated independently of the PDFs whose information is stored within the f_{i_y, i_τ} . The process is demonstrated in Figure 4.3. In effect the Monte Carlo weight is spread over a number of discrete grid-points according to an interpolation function and later combined with an arbitrary PDF defined on an identical grid.




Figure 4.3: Combining a weight-grid with a PDF in a fast step.

4.2.4 Grids for Hadron-Hadron interactions at NLO

The situation becomes more complicated for hadron-hadron collisions as two PDFs have to be considered, leading to 7 possible generalised initial state configurations (e.g.

qg = quark-gluon, $q\bar{q}$ = quark-antiquark etc.):

$$\begin{aligned}
gg : F^{(0)}(y_1, y_2, \tau) &= G_1(y_1)G_2(y_2) \\
qg : F^{(1)}(y_1, y_2, \tau) &= (Q_1(y_1) + \bar{Q}_1(y_2))G_2(y_2) \\
gq : F^{(2)}(y_1, y_2, \tau) &= G_1(y_1)(Q_2(y_2) + \bar{Q}_2(y_2)) \\
qr : F^{(3)}(y_1, y_2, \tau) &= Q_1(y_1)Q_2(y_2) + \bar{Q}_1(y_1)\bar{Q}_2(y_2) - D(y_1, y_2) \\
qq : F^{(4)}(y_1, y_2, \tau) &= D(y_1, y_2) \\
q\bar{q} : F^{(5)}(y_1, y_2, \tau) &= \bar{D}(y_1, y_2) \\
q\bar{r} : F^{(6)}(y_1, y_2, \tau) &= Q_1(y_1)\bar{Q}_2(y_2) + \bar{Q}_1(y_1)Q_2(y_2) - \bar{D}(y_1, y_2)
\end{aligned}$$

In the above, g =gluon, q =quark and r is a quark of a different flavour⁴ i.e. $q \neq r$.

The G , D and Q are generalised PDFs defined below:

$$\begin{aligned}
G_H(y) &= f_{0/H}(y, \tau) \\
Q_H(y) &= \sum_{i=1}^6 f_{i/H}(y, \tau) \\
\bar{Q}_H(y) &= \sum_{i=-6}^{-1} f_{i/H}(y, \tau) \\
D(y_1, y_2) &= \sum_{i=-6, i \neq 0}^6 f_{i/H_1}(y_1, \tau) f_{i/H_2}(y_2, \tau) \\
\bar{D}(y_1, y_2) &= \sum_{i=-6, i \neq 0}^6 f_{i/H_1}(y_1, \tau) f_{-i/H_2}(y_2, \tau)
\end{aligned}$$

where $f_{i/H}$ is the PDF of flavour i ($i = -6 \dots 6$ following the PDG MC numbering scheme [19]) for hadron H (H_1 denotes the first hadron, H_2 the second).

⁴There is no combination ‘rq’ as this is implicit in the definition of ‘qr’. Similarly $\bar{q}q$ and $\bar{r}q$ are included in the definitions of $q\bar{q}$ and $q\bar{r}$ respectively.

Three dimensional grids are required as the $y(x)$ values in each hadron and the momentum transfer $\tau(Q^2)$ must be considered. A further set of grids is required to describe the NLO contributions as these require a different coupling constant (i.e. α_s^3). The grid version of the hadron-hadron cross-section becomes:

$$\sigma = \sum_p \sum_{l=0}^6 \sum_{i_{y_1}}^{N_{y_1}} \sum_{i_{y_2}}^{N_{y_2}} \sum_{i_\tau}^{N_{i_\tau}} W_{i_{y_1}, i_{y_2}, i_\tau}^{(p)(l)} \left(\frac{\alpha_s(\tau^{(i_\tau)})}{2\pi} \right)^p F_{i_{y_1}, i_{y_2}, i_\tau}^{(l)} \quad (4.14)$$

Where:

$$F_{i_{y_1}, i_{y_2}, i_\tau}^{(l)} = F^{(l)}(y_{i_{y_1}}, y_{i_{y_2}}, \tau_{i_\tau}) \quad (4.15)$$

The sum over p takes into account LO and NLO contributions and the sum over l accounts for the 7 combinations of initial state partons. The remaining sums are over the three dimensional grids (x_1, x_2, Q^2) of weights (W) and the generalised PDFs (F).

In analogy with the DIS case, the computationally intensive Monte Carlo weight information is contained within the weight-grids (W), separate from the PDF information stored within the generalised PDFs (F). As a result the Monte Carlo weights can be reused for any PDF. Given a previously generated weight-grid the summation as represented by Eqn 4.14 may be carried out rapidly (i.e in a fraction of a second) opening the possibility of using this technique in global PDF fits.

4.2.5 Adjusting the hard scale

A bonus of separating the Monte Carlo weights and PDFs on to y, τ grids for each order (e.g. LO and NLO) is that the dependence of the overall cross-section on the choice of hard-scale (μ_r, μ_f) can be easily investigated. In the calculation of the Monte Carlo weights a choice of scale has to be made based on the kinematics of the event. Typically a choice is made around the presumed Q^2 of the event (see Section 3.2.3) e.g. by taking the P_T of the highest P_T jet in the event. This choice is implicit in the weight-grid and defines the τ co-ordinate of the grids over which the MC weight is spread.

The choice of the hard-scale affects both the strength of the strong coupling and the

PDF (through scaling violations). In the LO case, simply adjusting the value of τ in Eqns 4.14 and 4.15 which relate to α_S and the PDFs is sufficient. This simply means altering the definition of τ in Eqn 4.3:

$$\tau(Q^2) = \ln \ln \frac{\epsilon^2 Q^2}{\Lambda^2}. \quad (4.16)$$

Where ϵ relates to a multiplicative factor by which the hard scale has been varied. The weight-grid $W_{i_{y1}, i_{y2}, i_\tau}$ can be left unchanged as at LO it has no direct dependence on the scale⁵.

At NLO the situation becomes more complicated as the parton-level cross-section represented by the weight-grid becomes dependent on the choice of scale. The relevant NLO corrections to the cross-section are sketched in [23] and [24]:

$$\sigma = [\alpha_S^2(\mu_R)\sigma_0 + \alpha_S^3(\mu_R)(\sigma_1 + b_0 \ln(\mu_R^2/Q^2)\sigma_0 - 2P_{qq} \ln(\mu_f^2/Q^2)\sigma_0)] \otimes f_q(\mu_F) \otimes f_q(\mu_F) \quad (4.17)$$

Where the P_{qq} is a quark splitting function, σ_0 relates to the LO parton-level cross-section and σ_1 is a collection of NLO corrections which do not depend on the choice of hard scale. The first term is the LO contribution and it can be seen that simply modifying the coupling α_S and the PDFs is sufficient in this case. By carefully applying the logarithmic corrections to the weight-grids it is also possible to arbitrarily vary the scale at NLO after their production, see [1] for details.

4.3 Implementing a grid design

Use of integration-grid techniques to calculate a differential cross-section (e.g. inclusive jets) requires multiple grids containing a large number of data entries to be held in a computer's memory. For each observable (e.g. histogram bin in P_T for inclusive jets), a set of seven grids is required to record the Monte Carlo weights for each order (e.g.

⁵The weight-grid is analogous to the parton-level cross-section and the renormalisation and factorisation scales only enter at NLO.

LO, NLO)⁶. As the grids are three-dimensional for hadron-hadron scattering it is clear that memory issues can quickly become problematic. An example of a weight-grid for an observable bin is given in Figure 4.4.




Figure 4.4: Example weight-grid (NLOJET⁺⁺, inclusive jets), as a function of y_1 and y_2 projected over $2.90 < \tau < 2.96$.

The most basic form of optimisation is to vary the number of y and τ bins in a given grid. In any grid system the grid-spacing ($\delta y, \delta \tau$) is important as it defines the resolution of the grid. Typically the smaller the grid-spacing the more accurate the grid can be. However a smaller grid-spacing, requires a larger number of bins to cover the same parameter space in y, τ and hence consumes more computer memory. Developing a practical implementation of an integration grid requires a compromise between obtaining sufficient resolution in x and Q^2 within the constraints of finite computer memory [25].

4.3.1 Co-ordinate Transformations

There are a few techniques that can be applied to reduce the amount of CPU memory an accurate grid uses. One of these methods is the (previously encountered) co-ordinate transformation applied to x and Q^2 (Eqns 4.2, 4.3). The effect of the transformations are given in Figures 4.5 and 4.6. The integration grid-spacings are uniform in the transformed

⁶Seven grids are necessary to take into account the seven possible initial state parton configurations.

variables (y, τ) , therefore the transformations lead to increased grid-resolution at low x and Q^2 , which reflect the kinematics of most hadron-hadron collisions at the LHC.




Figure 4.5: Co-ordinate transformation of the x -variable (Eqn 4.2).




Figure 4.6: Co-ordinate transformation of the Q^2 -variable, $\Lambda = 0.25$ (Eqn 4.3).

The co-ordinate transform can be modified by introducing an additional parameter ‘ a ’. This seeks to soften the reduction of bins at high- x in the PDF and allow a better description of high P_T forward events. This simply involves modifying Eqn 4.2 to become:

$$y(x) = \ln \frac{1}{x} + a(1 - x) \quad (4.18)$$

The only difficulty associated with using equation 4.18 is that the expression cannot be inverted analytically and numerical methods have to be applied instead. The effect of varying ‘ a ’ on the x co-ordinate transformation is shown in Figure 4.7.




Figure 4.7: Variable co-ordinate transformation of the x -variable Eqn 4.18.

4.3.2 Bin-Boundaries Optimisation

As mentioned earlier a set of grids is required for each observable (e.g. P_T bin in a differential jet cross-section). For a given observable it is likely that only a subset of the parameter space in y , τ is kinematically possible. For example, to create a central jet with large P_T requires both a high x (low y) and high Q^2 (high τ) interaction. This can be used to improve the efficiency of grids by cutting away areas of parameter space in a grid which will not be used. As there is no *a priori* method of predicting the allowed kinematic region for a general process, the simplest approach is to perform a short trial run where a grid with wide (inclusive) boundaries is filled. This grid is then examined and areas of unused parameter space are cut, to give improved grid boundaries for subsequent production. A simple implementation of this is to determine minimum and maximum values of y_1 , y_2 and τ for a particular observable bin in the trial run. The minimum and maximum values can then be used as the new upper and lower grid boundaries in a subsequent optimisation run.

4.3.3 Interpolation

The formula used to interpolate the grid was given in Eqn 4.5. The interpolation calculates the Lagrange interpolation polynomial, which is defined as the polynomial of degree n which passes through $n + 1$ points. The use of high order interpolation techniques allow precision to be maintained in a grid whilst reducing the number of discrete grid points which have to be kept in memory and recorded.

4.3.4 PDF Weighting

Another method to improve the efficiency of grids is to apply a PDF weighting. Figures 4.8 and 4.9 show the sum of the up,down and gluon PDFs and their gradients for a range of Q values. As can be seen from these plots, the PDF's gradient is steepest at low- x .




Figure 4.8: Sum of the up, down and gluon PDFs at $Q=5\text{GeV}$, 50GeV , 500GeV and 5000GeV for CTEQ6.1 best fit PDFs.

PDF weighting, first carried out by the fastNLO project [26], involves applying an x -dependent weight in order to flatten out the PDFs and hence improve the accuracy of the interpolation. The x -dependent weight is applied to a Monte Carlo weight before its introduction to the grid and then divided from the discrete PDF grid points in the final step. The form of the weighting function chosen is shown below (and is the fastNLO




Figure 4.9: Gradient of the sum of up, down and gluon PDFs at $Q=5\text{GeV}$, 50GeV , 500GeV and 5000GeV for CTEQ6.1 best fit PDFs.

choice [26]):

$$w = x^{a_1} (1 - 0.99x)^{a_2} \quad (4.19)$$

This function is an attempt to describe the basic form of a PDF and is similar to many basic PDF parameterisations, for example equation 3.5. The factor ‘0.99’ prevents a collapse in the weight as x tends to 1. The parameters a_1 and a_2 of Eqn 4.19 were found by fitting the equation to the sum of the up, down and gluon PDFs using MINUIT [27]. This can be shown for the 40 PDFs of the CTEQ6.1 sets [23] in Figures 4.10 and 4.11. Equivalent plots for best-fit PDFs (supported in LHAPDF [28]) from a range of groups CTEQ, MRST and H1 (including deprecated sets) are shown in Figures 4.12 and 4.13

The figures show that the optimum choice of parameters does not vary much with the choice of PDF within the CTEQ6.1 error sets, however, somewhat more variation is observed when considering other PDF parameterisations in particular for a_2 . There is also a dependency on Q or momentum transfer in the event, with a_1 and a_2 both increasing in magnitude with Q (and hence also Q^2). A pragmatic approach can be taken to the best choice of parameters in the case of an integration grid. Figures 4.8 and 4.9 show that the curvature of the PDFs is larger at low- x and hence this is the region where PDF-




Figure 4.10: Parameter a_1 found by fitting Eqn 4.19 to the sum of the up, down and gluon PDFs for the 40 error PDFs of CTEQ6.1 at $Q=5\text{GeV}$, 50GeV , 500GeV and 5000GeV .




Figure 4.11: Parameter a_2 found by fitting Eqn 4.19 to the sum of the up, down and gluon PDFs for the 40 error PDFs of CTEQ6.1 at $Q=5\text{GeV}$, 50GeV , 500GeV and 5000GeV .

weighting is expected to be most effective. As x relates to the momentum fraction of a proton carried by an interacting parton, a low- x implies a low Q^2 event.

This favours the use of weighting parameters that correspond to very low Q^2 . However in experiments, low Q^2 interactions often escape detection or fail trigger thresholds and can be neglected⁷. In the simple case of a 2-body scattering (in the relativistic limit), the Q^2 of the event varies from P_T^2 to $2P_T^2$ between small and large angle scatterings

⁷Low Q^2 events have outgoing particles with low transverse momentum (P_T).




Figure 4.12: Parameter a_1 found by fitting Eqn 4.19 to the sum of the up, down and gluon PDFs for 31 best fit PDFs (See Table 4.1) at $Q=5\text{GeV}$, 50GeV , 500GeV and 5000GeV .




Figure 4.13: Parameter a_2 found by fitting Eqn 4.19 to the sum of the up, down and gluon PDFs for 31 best fit PDFs (See Table 4.1) at $Q=5\text{GeV}$, 50GeV , 500GeV and 5000GeV .

respectively. The P_T in this case is the transverse momentum of one of the outgoing particles.

At hadron colliders, jets can be measured with energy down to around $10 - 30 \text{ GeV}$ and hence a Q value of approximately 50GeV is an appropriate choice. With reference to Figs 4.10 and 4.11 this leads to a choice of $a_1 \sim -1.5$ and $a_2 \sim 3$. These values are also reasonable when considering Figures 4.12 and 4.13 despite the wide variety of PDFs considered.

PDF Number	PDF Name	PDF Number	PDF Name
1	cteq6l	17	H12000lo2E
2	cteq6mE	18	MRST2004nlo
3	cteq5m	19	MRST2004nnlo
4	cteq5m1	20	MRST2003cnlo
5	cteq5d	21	MRST2003cnnlo
6	cteq5l	22	MRST2002nlo
7	cteq4m	23	MRST2002nnlo
8	cteq4d	24	MRST2001E
9	cteq4l	25	MRST2001lo
10	H12000ms	26	MRST2001nlo
11	H12000msE	27	MRST2001nnlo
12	H12000dis	28	MRST98lo
13	H12000disE	29	MRST98nlo
14	H12000lo	30	MRST98dis
15	H12000loE	31	MRST98ht
16	H12000lo2		

Table 4.1: Key for Figures 4.12 and 4.13, all PDFs from the LHAPDF 5.0 interface.

The effect of applying the weighting function with $a_1 = -1.5$ and $a_2 = 3$ to the CTEQ6.1 PDF is shown in Figure 4.14. As can be seen the curvature of the weighted PDFs is considerably reduced in comparison with the unweighted PDFs .




Figure 4.14: Weighted PDF(CTEQ6.1) (black) for $a_1=-1.5$, $a_2=3$ and unweighted PDFs (red) at $Q=5\text{GeV}$, 50GeV , 500GeV and 5000GeV . Function $g(x) = x$ in the ‘unweighted’ case and $g(x) = 1/\text{weight}'$ in the ‘weighted’ case, where *weight* is defined by equation 4.19.

In reality it is the generalised PDF grids (e.g F^n) that are combined with the weight

grids to form the cross-section. An example of such a grid is given in Figures 4.15 and 4.16. The boundaries of the grid were determined for one observable bin (in P_T) for the inclusive jet cross-section as a result of running a bin-boundaries optimisation step (see section 4.3.2). The figures show that for a given grid the variation in y is much stronger than that in τ even after weighting.




Figure 4.15: Weighted PDF grid (CTEQ6.1), $a_1=-1.5$ and $a_2=3$ for $F^{(0)}$ as a function of y_1, y_2 , projected over $2.90 < \tau < 2.96$




Figure 4.16: Weighted PDF grid (CTEQ6.1), $a_1=-1.5$ and $a_2=3$ for $F^{(0)}$ as a function of y_1 and τ , projected over $0 < y_2 < 6$

4.4 Testing Grid Implementation

The use of integration grids can, in theory, be applied to any Monte Carlo program that concerns hadron collisions. However, in order to test the implementation, a decision was made to use NLOJET⁺⁺ [16] which offers NLO jet cross-sections (at parton level) for hadron colliders in a C^{++} implementation⁸. The inclusive jet cross-section at ATLAS (proton-proton, $\sqrt{s} = 14\text{TeV}$) separated into three regions of pseudorapidity ($0 < |\eta| < 1$, $1 < |\eta| < 2$, $2 < |\eta| < 3$) was chosen as a reference process.

The integration grids code was interfaced to NLOJET⁺⁺ and used to generate weight-grids independent of a PDF. After production, the grids could be combined with an arbitrary PDF (through an LHAPDF interface) to produce a cross-section in a fast step. The same Monte Carlo weights from NLOJET⁺⁺ were also used directly with a small set of PDFs: CTEQ6.1, CTEQ4, H12000ms, MRST2004nlo and ZEUS2005 (as in the case of normal Monte Carlo cross-section generation), to provide a set of reference cross-sections for comparison with the grid generated version. The grids chosen in this case were $30 \times 30 \times 10$ (y_1 , y_2 , τ), had a co-ordinate transform parameter of $a = 5$ and were weighted. The ratio of the grid to the standard reference cross-sections for the three pseudorapidity regions are shown in Figures 4.17, 4.18 and 4.19.

As can be seen the grids are generally able to describe the cross-section to an accuracy of 0.1%. There are regions at moderate to high P_T where the accuracy decreases (for certain PDFs) but these are contained to within $\sim 1\%$ in all cases. A clear observation from the figures is that the ability of a grid to describe the cross-section varies between different PDFs. As the same Monte Carlo weights are used to produce the grid and the reference cross-sections, the discrepancy can be interpreted as the ability of the grid to accurately describe a given PDF. This varies between PDFs, as a grid is an approximation to a true PDF and the ability of the interpolation to estimate a PDF at a given point (x , Q^2) will depend on its functional form and parameters.

⁸A seedless cone algorithm was used to define jets from the final state.




Figure 4.17: Comparison of grid/standard for the inclusive jet-cross section at ATLAS ($0 < |\eta| < 1$). A standard cross-section is calculated using NLOJET⁺⁺ without an integration grid. (Weighted grid, $30 \times 30 \times 10$ bins in y_1 , y_2 , τ , and co-ordinate transform parameter $a=5$).




Figure 4.18: Comparison of grid/standard for the inclusive jet-cross section at ATLAS ($1 < |\eta| < 2$). A standard cross-section is calculated using NLOJET⁺⁺ without an integration grid. (Weighted grid, $30 \times 30 \times 10$ bins in y_1 , y_2 , τ , and co-ordinate transform parameter $a=5$).

4.4.1 Effect of Modifying Grid Size

The number of grid bins (in y and τ) can be modified. A greater number of bins within the same range of y and τ provides greater resolution in x and Q^2 but incurs a penalty in terms of computer memory and speed of processing. As a comparison reconstructing




Figure 4.19: Comparison of grid/standard for the inclusive jet-cross section at ATLAS ($2 < |\eta| < 3$). A standard cross-section is calculated using NLOJET⁺⁺ without an integration grid. (Weighted grid, $30 \times 30 \times 10$ bins in y_1 , y_2 , τ , and co-ordinate transform parameter $a=5$).

a cross-section with a $30 \times 30 \times 10$ (y_1 , y_2 , τ) grid takes approximately 0.7s as opposed to 0.1s for a $10 \times 10 \times 10$ grid (for 24 observable bins).

The effect of varying the number of y bins from 40 to 10 for 10 τ bins⁹ for a weighted grid is shown in Figures 4.20, 4.21 and 4.22 for the H12000ms PDF. As can be seen the smaller grid is sufficient at low rapidities $0 < |\eta| < 1$ but fails to describe the forward regions at high P_T . This is consistent with the co-ordinate transforms used in the grids equations 4.2 as the grids have better resolution at low x , at the expense of the high x region which is probed at high rapidity and P_T .

The effect of varying the number of τ bins (independently of the y -bins) from 5 to 30 for 30 y bins is shown in Figure 4.23. The accuracy of the grid is found to be less dependent on the resolution in τ bins than y bins and the improvement of grid accuracy with the number of τ bins appears to saturate above a value of 10.

As described earlier the ability of grids to describe different PDFs varies and it was found that the cross-section produced by a CTEQ6.1 best fit PDF is successfully described by a small ($10 \times 10 \times 10$) grid to within 0.1% over all rapidity ranges. The reason for this discrepancy could be that the CTEQ6.1 PDF is better smoothed by use of the weighting

⁹($40 \times 40 \times 10$) to ($10 \times 10 \times 10$) (y_1 , y_2 , τ).




Figure 4.20: Comparison of grid/standard for the inclusive jet cross section at ATLAS for different grid sizes ($0 < |\eta| < 1$). (Weighted grid, co-ordinate transform parameter $a=5$). A standard cross-section is calculated using NLOJET⁺⁺ without an integration grid.




Figure 4.21: Comparison of grid/standard for the inclusive jet cross section at ATLAS for different grid sizes ($1 < |\eta| < 2$). (Weighted grid, co-ordinate transform parameter $a=5$). A standard cross-section is calculated using NLOJET⁺⁺ without an integration grid.

functions than the older H1200ms used in Figures 4.20-4.23. Although the more recent CTEQ PDF provides a better description of the proton it is important that a grid-method is reliable for sub-optimal PDFs. The reason being that the use of grids in global PDF fits requires the accurate production of a cross-section with an assumed PDF which may not be an optimal description of the proton.




Figure 4.22: Comparison of grid/standard for the inclusive jet cross section at ATLAS for different grid sizes ($2 < |\eta| < 3$). (Weighted grid, co-ordinate transform parameter $a=5$). A standard cross-section is calculated using NLOJET⁺⁺ without an integration grid.




Figure 4.23: Comparison of grid/standard for the inclusive jet-cross section at ATLAS for different grid sizes ($2 < |\eta| < 3$). (Weighted grid, no. y bins=30, co-ordinate transform parameter $a=5$). A standard cross-section is calculated using NLOJET⁺⁺ without an integration grid.

4.4.2 The Effect of Interpolation Order

The interpolation function which allows the PDF to be determined at a general x , Q^2 from the discrete grid can be applied at different orders. Higher order interpolation draws on a greater number of discrete grid points to estimate the value of a PDF and can provide a more accurate description if the PDF is a well-behaved, smooth function of x and

Q^2 . The effect of increasing the interpolation order is shown in Figure 4.24. Increasing the interpolation order improves the accuracy of the grids, however it slightly lengthens their production time as a greater number of grid points have to be updated for each Monte Carlo weight that is produced. A choice of interpolation order between 3 and 5 (inclusive) was found to provide sufficient accuracy without adversely affecting the speed of grid generation.




Figure 4.24: Comparison of grid/standard for the inclusive jet-cross section at ATLAS for different interpolation orders ($(0 < |\eta| < 1)$). Weighted grid ($30 \times 30 \times 10$), co-ordinate transform parameter $a=5$). A standard cross-section is calculated using NLOJET⁺⁺ without an integration grid.

From Figure 4.24, the use of lower order interpolation (e.g. order 1) leads to increased error at low-mid P_T . This is possibly due to the gradient of the weighted PDFs (Figure 4.14) being higher at low x values ($x \sim 0.015$). The use of a low order interpolation may fail to adequately describe the PDFs curvature despite the increase in grid resolution afforded by the x co-ordinate transformation (Eqn 4.2).

4.4.3 The Effect of PDF Weighting

The effect of turning off the PDF weighting is shown in Figure 4.25 for a small $10 \times 10 \times 10$ grid. As can be seen, without weighting the small grid fails to describe the cross-section at low P_T . This is consistent with the PDF weighting being most effective at low x (and

hence Q^2) where the PDF curvature is greatest.




Figure 4.25: Comparison of grid/standard for the inclusive jet-cross section at ATLAS for weighted and unweighted ($10 \times 10 \times 10$) grids. ($(0 < |\eta| < 1)$, co-ordinate transform parameter $a=5$). A standard cross-section is calculated using NLOJET⁺⁺ without an integration grid.

It was found that larger $30 \times 30 \times 10$ grids did provide reasonable accuracy (to within $\sim 2\%$) over the P_T and η ranges tested¹⁰, without weighting. The unweighted grids, however, showed a greater sensitivity to a reduction in the size of the grid and to the choice of interpolation order. The best results overall were obtained by combining the large grid with PDF weighting.

4.4.4 The Effect of the Co-ordinate Transformation

The effect of the co-ordinate transform is shown in Figure 4.26 for small ($10 \times 10 \times 10$) grids. The free parameter in the co-ordinate transform was changed from values of 0 to 10. For a weighted grid, a choice of $a = 5$ was shown to give the best overall performance although it only differs from a choice of $a = 3$ by a negligible amount $O(10^{-4})$. At low values of ‘ a ’ the grids fail to describe the high P_T jets, with a considerable error being seen for choice $a = 0$. Higher ‘ a ’ values lead to an increased resolution of the grids at high x and thus better describe high P_T jet events. At very high values though ($a = 10$), the accuracy

¹⁰i.e. the ranges shown in Figures 4.17-4.19

for low P_T jets begins to suffer. Larger weighted and unweighted grids ($30 \times 30 \times 10$) are immune to the effects of the co-ordinate transformation $0 < a < 10$ to $O(10^{-3})$.




Figure 4.26: Comparison of grid/standard for the inclusive jet-cross section at ATLAS for different grid co-ordinate transforms. ($(2 < |\eta| < 3)$, $10 \times 10 \times 10$ weighted grids). A standard cross-section is calculated using NLOJET⁺⁺ without an integration grid.

4.4.5 Adjusting the Renormalisation and Factorisation Scales

The hardness scale (renormalisation/factorisation) can be changed for a grid after generation. Figure 4.27 shows that the scale may be changed freely without reducing the ability of a grid to recreate the standard cross-section.

4.5 Grids in Global PDF Fits

As described earlier, an advantage of using integration grids for the calculation of NLO cross-sections is that it allows hadron collider jet data to be included into global PDF analyses. Hadron collider data is useful in particular for constraining the gluon PDF at high x which is poorly determined by deep-inelastic scattering (DIS) experiments.

Integration grids were generated using NLOJET⁺⁺ for three regions of pseudorapidity $0 < \eta < 1$, $1 < \eta < 2$ and $2 < \eta < 3$ for the inclusive jet cross-section at ATLAS¹¹.

¹¹The standard $30 \times 30 \times 10$ weighted grids were used with $a = 5$ and 5th order interpolation




Figure 4.27: Comparison of grid/standard for the inclusive jet-cross section at ATLAS for different grid renormalisation/factorisation scales. ($(0 < |\eta| < 1)$, $30 \times 30 \times 10$ weighted grids). A standard cross-section is calculated using NLOJET⁺⁺ without an integration grid.

Experimental pseudo-data for ATLAS was simulated using the JETRAD [17] generator and then combined in a global PDF fit with DIS data from ZEUS [29].

Global PDF fits consider the experimental error on a measurement. If a measurement's error is large then its ability to constrain PDFs is correspondingly reduced. It was found that correlated systematic errors like the jet energy scale (JES), dominate the effectiveness of ATLAS data in constraining the gluon PDF¹². The effect of including ATLAS jet data on the uncertainty in the gluon distribution is shown in figures 4.28, 4.29 and 4.30 (fits carried out by Claire Gwenlan (Oxford/UCL)). The three figures show the reduction in the gluon uncertainty when ATLAS jet data is included for correlated systematic (JES) uncertainties of 0%, 1% and 3% respectively. The different bands of each figure represent increasing information being fed into the global fit. ZEUS-O is based solely on neutral and charged current interactions at ZEUS, whilst ZEUS-JETS also considers jet cross-sections from direct-photoproduction and QCD inclusive jets. The ATLAS-JETS takes all the information from ZEUS-JETS combined with the inclusive jet cross-section at ATLAS.

As can be seen, the ATLAS data is useful for constraining the gluon PDF at high x if

¹²The uncorrelated systematic and statistical errors play a far smaller role

the JES can be constrained to 1% (or better) for jets with $P_T > 1\text{TeV}$. The problem of determining the JES at high jet P_T and a possible solution will be described in Chapter 6.




Figure 4.28: Gluon PDF uncertainty from global fit with inclusion of ATLAS pseudo-data, without JES uncertainty. Results for 10fb^{-1} with a 5% uncorrelated systematic error. ZEUS-O considers data from neutral/charged current interactions from ZEUS whilst ZEUS-JETS includes data from ZEUS jet cross-sections. ATLAS-JETS in addition includes the ATLAS inclusive jet cross-section. (plot by Claire Gwenlan (Oxford/UCL)).




Figure 4.29: Gluon PDF uncertainty from global fit with inclusion of ATLAS pseudo-data, with 1% JES uncertainty. Results for 10fb^{-1} with a 5% uncorrelated systematic error. ZEUS-O considers data from neutral/charged current interactions from ZEUS whilst ZEUS-JETS includes data from ZEUS jet cross-sections. ATLAS-JETS in addition includes the ATLAS inclusive jet cross-section. (plot by Claire Gwenlan (Oxford/UCL)).




Figure 4.30: Gluon PDF uncertainty from global fit with inclusion of ATLAS pseudo-data, with 3% JES uncertainty. Results for 10fb^{-1} with a 5% uncorrelated systematic error. ZEUS-O considers data from neutral/charged current interactions from ZEUS whilst ZEUS-JETS includes data from ZEUS jet cross-sections. ATLAS-JETS in addition includes the ATLAS inclusive jet cross-section. (plot by Claire Gwenlan (Oxford/UCL)).

Chapter 5

The Inclusive Jet Cross-Section

5.1 Introduction

The inclusive jet cross-section describes the probability of obtaining a jet with a given transverse momentum. The cross-section considers all events which contain jets within acceptance¹ and as a result has contributions from many different processes at the LHC. A simplification can be made in theoretical predictions, however, by considering only basic QCD production (as shown in Section 3.2) as this is the dominant source of jets by several orders of magnitude.

As was demonstrated earlier, the inclusive jet cross-section at hadron colliders provides an implicit test of hadron structure as PDFs are needed to create theoretical predictions to compare with experimental data. A poorly determined PDF may lead to a discrepancy between the theoretical prediction and experiment. In addition the running of the strong coupling constant with Q^2 (see Section 3.1.1) can be measured which provides a non-trivial test of QCD. Such a measurement was carried out at the Tevatron [30] and a similar measurement may be made at ATLAS over a larger Q^2 range.

The relative Q^2 reach of the Tevatron and the LHC is shown in Figure 5.1. The principal cause of the difference between the colliders is the centre of mass energy $\sqrt{s}$ which

¹The cross-section was encountered in previous chapters (see Chapter 3, Figure 3.6, and tests of the integration grid in Chapter 4)

is 1.96TeV in the case of the Tevatron and 14TeV for the LHC. The greater $\sqrt{s}$ allows for higher Q^2 interactions and hence the cross-section extends to greater P_T . It should be noted in addition that ATLAS will eventually collect an integrated luminosity ~ 2 orders of magnitude higher than corresponding Tevatron detectors, which also increases its effective Q^2 reach.




Figure 5.1: Inclusive Jet Cross-Section for the LHC and Tevatron (Run II), (NLOJET, CTEQ6.1, $\mu_r, \mu_f = P_T(max)$). P_T is effectively a measure of Q^2 .

Considerable interest in the inclusive jet cross-section is directed toward new physics searches. High P_T jet production involves the highest Q^2 interactions available at a hadron collider and thus is a probe of interactions at the smallest obtainable length scales². By searching for deviations between experimental data and Standard Model predictions; sensitivity to unanticipated phenomena may be obtained. This is in contrast to specific searches (e.g. SUSY) which look for particular final-state signatures predicted by possible extensions to the Standard Model.

One extension to the Standard Model to which the inclusive jets cross-section is sensitive is quark compositeness. In the Standard Model, quarks are considered to be point-like fundamental particles. However, it is conceivable that quarks are composite objects comprised of multiple constituents bound together by a force of greater strength than QCD.

²A Q^2 of 1TeV relates to a probing of a length-scale $O(10^{-18}m)$, based on the De-Broglie wavelength of an exchange particle in the Breit frame.

This compositeness could have evaded detection in previous experiments in the same way that the compositeness of the atom and the proton escaped detection until experiments were created to probe the structure at sufficient Q^2 . Quark compositeness is a worthwhile study at the LHC as the collider is capable of obtaining higher Q^2 interactions than previously achieved and could discover or set a new limit on a compositeness scale. Compositeness manifests itself experimentally as an excess in the observed inclusive cross-section (when compared to SM theory) at high jet P_T .

Whether carrying out a specific search for compositeness, or a more general new physics search the errors on both the theoretical prediction and experimental measurement of the inclusive jet cross-section have to be quantified accurately. Failure to do this can mask or indeed lead to false signals of new physics. In this section an overview of the major sources of theoretical uncertainty are outlined, followed by an outline of the experimental errors.

5.2 Theoretical Uncertainties

5.2.1 Fixed-Order Uncertainties

The theoretical prediction for the inclusive jet cross-sections is carried out by a calculation similar to that outlined in Section 3.2. These calculations are referred to as fixed order as they only consider contributions (i.e. Feynman diagrams) up to a given complexity (i.e. power in the coupling constant). The remaining terms are neglected in the interests of calculational simplicity. This approximation is valid in a situation where the coupling constant involved is small (i.e. at high Q^2 for QCD) which suppresses the contribution from discarded terms to the overall cross-section.

The highest order predictions for inclusive jet cross-sections currently available are at next-to-leading order (NLO) and have been implemented as parton-level Monte Carlo programs (e.g. NLOJET⁺⁺ and JETRAD [16, 17]). These programs do not include a parton shower and hence produce a maximum of three partons in the final state from

which jets are inferred by applying a jet algorithm. Despite the simplicity of the final state, sufficient accuracy is obtained (after corrections) to adequately describe experimental data. Alternative parton shower programs such as PYTHIA [8] provide a realistic multi-particle final state for individual jets³ but fail to describe events with multiple hard jets. The use of parton shower programs is however helpful in gaining an understanding of the development of an outgoing parton into a jet.

Despite the inclusion of NLO corrections a considerable number of terms are neglected in the calculation and these confer an uncertainty to the theoretical prediction. This uncertainty is difficult to precisely quantify without a higher-order calculation which is absent due to calculational difficulty. It is anticipated that the perturbative series should converge as higher order corrections become less and less significant, hence the difference between the LO and NLO predictions could be taken as an upper limit on this error. The ratio of these predictions is given in Figure 5.2⁴.




Figure 5.2: K-Factors for inclusive jets in three regions of pseudorapidity, (NLOJET, CTEQ6.1, $\mu_r, \mu_f = P_T(max)$)

From Figure 5.2 it is clear that the differences between LO and NLO calculations are considerable, particularly in the forward regions at high P_T . In some respects the disagreement of the calculations in this region, is a result of the arbitrary scale choice (in

³The accuracy of a parton shower is typically at leading log (LLA)

⁴The seedless cone algorithm is used for all NLOJET⁺⁺ predictions.

this case $\mu_r, \mu_f = P_T(max)$, where $P_T(max)$ is the P_T of the leading jet). An alternative choice could be used which improved the agreement between LO and NLO predictions for a given point in P_T and η at the expense of other regions.

The unphysical renormalisation and factorisation scales both appear due to the calculation being carried out to finite order. The renormalisation scale affects the strength of the strong coupling constant, whilst the factorisation scale is an arbitrary scale defining at what point soft QCD radiation becomes attributed to the PDF in the calculation. An all-orders calculation would be independent of such choices and it can be shown that a NLO calculation has a reduced scale-dependence in comparison with LO [23]. Owing to the differing sensitivities to scale, K-Factors as shown Figure 5.2 are sensitive to the choice of scale.

The approach of comparing LO with NLO predictions to estimate the errors due to fixed-order calculations could be considered pessimistic as the perturbative series is expected to converge. An alternative approach is to simply vary the hard scales used in the NLO-QCD calculation (i.e. the renormalisation and factorisation scales). The sensitivity of the prediction to these scales can then be interpreted as a measure of their uncertainty.

The standard approach is to vary the renormalisation and factorisation scales (μ_r, μ_f) between two or more arbitrary choices and observe the resulting variation in the predicted cross-section. Generally, the choice of these scales is designed to reflect the hardness (Q^2) of the event concerned. Such a choice makes the strong coupling (α_s) representative of that in the hard scattering and allows soft processes to be subsumed into the PDF. This minimises the variation of the predicted cross-section with scale, increasing its accuracy.

As the choice of the hard scale is arbitrary, there exist a number of conventions for its selection. Often the scale is equated with either a multiple of the P_T of each jet in the event [23] (i.e. $Q^2 \sim P_T^2$) or the highest P_T jet (P_T^{max}) in the event [31]. The choice of multiple is again arbitrary, however, many calculations choose to multiply the jet P_T by a constant factor between 0.5 and 2.0.

An equally subtle decision has to be made when deciding how much the scales should be varied in order to estimate the error associated with the fixed-order calculation. The standard approach is to recalculate the cross-section using a scale that is half and double that used in the best approximation and to use the resulting predictions as an error band. This choice is a convention rather than being well motivated theoretically.

The use of integration grid methods can help in studies of this nature as they allow the renormalisation and factorisation scales to be varied, without the need to recalculate the NLO cross-section (See Section 4.4.5). An integration grid originally used for PDF fitting was used to look at the variation in the cross-section with scale (μ_r, μ_f) between $0.5P_T^{max}$ and $2.0P_T^{max}$ in comparison with a default value of P_T^{max} . Two methods were applied, the first considered the maximum and minimum cross-sections (for a given bin) found within the range of scale choices. The second method just considered the cross-section at the extreme choices of scale ($0.5P_T^{max}$ and $2.0P_T^{max}$) in comparison with the default. The plots are shown in Figures 5.3, 5.4 and 5.5.




Figure 5.3: Scale error for inclusive jets within $0 < |\eta| < 1$. The scale was varied from $0.5P_T^{max}$ to $2.0P_T^{max}$. The max-min series relate to the largest deviations of the cross-section within this range, whereas std-fix considers the values of the cross-section at the extreme choices of scale.(NLOJET, CTEQ6.1)

The two approaches for determining the fixed-order error are equivalent except in high P_T , forward regions. The reason for this can be seen by looking at the variation for the




Figure 5.4: Scale error for inclusive jets within $1 < |\eta| < 2$. The scale was varied from $0.5P_T^{max}$ to $2.0P_T^{max}$. The max-min series relate to the largest deviations of the cross-section within this range, whereas std-fix considers the values of the cross-section at the extreme choices of scale. (NLOJET, CTEQ6.1)




Figure 5.5: Scale error for inclusive jets within $2 < |\eta| < 3$. The scale was varied from $0.5P_T^{max}$ to $2.0P_T^{max}$. The max-min series relate to the largest deviations of the cross-section within this range, whereas std-fix considers the values of the cross-section at the extreme choices of scale. (NLOJET, CTEQ6.1)

cross-section in a particular P_T bin with scale. This is shown for two examples in Figures 5.6 and 5.7.

As can be seen in these figures at high P_T in forward regions, there is a rapid decrease in the cross-section with decreasing scale. The local maximum in Figure 5.7, is a result of terms that control scaling violations within the NLO calculation. Typically an increase in




Figure 5.6: Variation of the cross-section with scale for a P_T bin, $100 < P_T < 150 \text{ GeV}$, $2 < |\eta| < 3$. (NLOJET,CTEQ6.1)




Figure 5.7: Variation of cross-section with scale for a P_T bin $600 < P_T < 700 \text{ GeV}$, $2 < |\eta| < 3$. (NLOJET,CTEQ6.1)

scale reduces the cross-section, as the coupling constant and the PDFs decrease. However, in certain forward, high P_T regions the rise in the PDFs from high- x valence quarks leads to a change in the sign of the scaling violation terms.

The errors estimated by Figures 5.3, 5.4 and 5.5 due to the fixed order calculation are far smaller than the differences between the NLO and LO calculation previously shown in Figure 5.2 and only reach significant values ($\sim 15\%$) at very high P_T where experimental errors (including statistics) are likely to dominate.

5.2.2 PDF Errors

The PDFs which describe the parton content of the proton also contribute an uncertainty to the theoretical cross-section. As mentioned previously, PDFs are a mandatory component of any cross-section prediction at a hadron collider but cannot be determined directly from first principles. Instead the PDFs are parameterised and constrained in a global fit.

The global fits utilise a large amount of data from a variety of experiments and perform a chi-squared fit to achieve the best-fit PDFs. The fits take into account both the statistical (uncorrelated) and systematic (correlated) errors of the individual experiments concerned. Following the CTEQ analysis [23], the chi-squared function is given by:

$$\chi'^2 = \sum_e \chi_e^2(a, r) \quad (5.1)$$

$$\chi_e^2(a, r) = \sum_i \frac{[D_i - \sum_k r_k \beta_{ki} - T_i(a)]^2}{\alpha_i^2} + \sum_k r_k^2 \quad (5.2)$$

where ‘ e ’ labels an experimental data set, ‘ i ’ labels a particular data point within that set and ‘ k ’ labels a particular source of systematic error. The parameters ‘ a ’ are those of the PDF being fitted and ‘ r ’ represents a set of shift variables designed to account for systematic experimental errors within the global fit. ‘ D_i ’ and ‘ T_i ’ represent the data and theory values respectively, with ‘ α_i ’ and ‘ β_{ki} ’ as the uncorrelated and correlated errors⁵. A global fit minimises ‘ χ'^2 ’ with respect to both the PDF parameters and the shifts ‘ r ’ in order to provide the best fit to the data.

The best fit to the data is, however, not the only acceptable solution. It would be expected that the parameter space in the immediate vicinity of the best-fit choice would also provide acceptable fits. This neighbourhood in parameter space can be described using the Hessian matrix given by:

$$H_{ij} = \frac{1}{2} \frac{d^2 \hat{\chi}^2}{da_i da_j} \quad (5.3)$$

⁵The magnitude of these errors are generally provided by the experimental collaboration from which the data originated.

where:

$$\hat{\chi}^2(a) \equiv \chi'^2(a, \hat{r}(a)) \quad (5.4)$$

and $\hat{r}$ represents the optimal shifts. In the Hessian approach to errors, the eigenvectors of the matrix are first found and a set of error PDFs produced by considering positive and negative excursions along each one. The number of eigenvectors is equal to the number of free parameters in the PDF parameterisation and these span the parameter space around the best-fit value. The magnitude of the excursions about the best-fit are defined by points at a chi-squared of $\hat{\chi}^2 = \hat{\chi}_0^2 + T^2$ where $\hat{\chi}_0$ is the best fit value and T is referred to as the tolerance. The tolerance effectively parameterises the maximum magnitude of an excursion for which the error PDF can still be considered an acceptable fit to data. The choice of the tolerance parameter is non-trivial and the CTEQ collaboration decided on a rather large value of 10 as being representative of a 90% confidence limit⁶ [32].

An estimate of the PDF error (δX) on an observable (X) may be found by looking for the maximum and minimum values of an observable defined with parameters within tolerance (T) of the best-fit. This can be carried out using the error sets described earlier by use of the equation (from [23]):

$$(\delta X_{\pm})^2 = \sum_{i=1}^n [X(a_i^{\pm}) - X(a_0)]^2 \quad (5.5)$$

The a_i^+ is defined as an error PDF set which causes an increase in the observable in comparison with the best fit and a_i^- is the same but for a corresponding decrease. The overall error bands on the observable are given as $X - \delta X_- < X < X + \delta X_+$.

The error on the predicted inclusive jet cross-section was calculated using the above method for the 40 CTEQ6.1M PDF error sets. The process involves calculating the inclusive jet cross-section for each of the error PDFs and is hence suited to integration grid techniques. The PDF error was hence estimated using the previous integration grid the results being shown in Figures 5.8, 5.9 and 5.10.

⁶This is not a rigorous statistical definition given the complexity of the analysis, but rather the result of comparing predictions with the experimental data in the analysis.




Figure 5.8: Uncertainty on inclusive jet cross-section from the 40 CTEQ6.1 error sets, $0 < |\eta| < 1$ (NLOJET, CTEQ6.1).




Figure 5.9: Uncertainty on inclusive jet cross-section from the 40 CTEQ6.1 error sets, $1 < |\eta| < 2$ (NLOJET, CTEQ6.1).

As can be seen the PDF uncertainty increases with both jet P_T and rapidity. The principal source of PDF uncertainty at high Q^2 for the inclusive jets cross-section arises from the poorly constrained high- x gluon⁷. Indeed, the majority of the error at high jet P_T as determined from equation 5.5 can be attributed to one eigenvector (PDF error sets 29 and 30) which describes uncertainty in the high- x gluon.

PDF errors dominate the theoretical uncertainty at high P_T on the inclusive jet cross-

⁷A large proportion of data that enters into global fits is from DIS experiments which are typically insensitive to the gluon as it does not carry electrical charge.




Figure 5.10: Uncertainty on inclusive jet cross-section from the 40 CTEQ6.1 error sets, $2 < |\eta| < 3$ (NLOJET, CTEQ6.1).

section and in previous experiments at the Tevatron inconsistencies between experimental data and theory have been attributed to PDFs. This sensitivity allows the jet data to be potentially useful in PDF fits (by use of integration grids) but in turn reduces sensitivity to new physics such as compositeness.

It should be noted that the PDF errors mentioned before, pertain to those concerning the global fitting procedure. Other errors can be introduced by a change in the underlying theoretical framework (such as the heavy quark treatment) which concern the evolution of PDFs. Recently, such a change was carried out between the CTEQ6.1 and CTEQ6.2 sets and resulted in a change in the PDFs in excess of the error band of the older PDFs at mid to low x , however this was not considered for this thesis.

5.3 Jet Errors - Experimental

5.3.1 Issues with Jet Reconstruction

Experimental errors on the inclusive jet cross-section arise from many sources. A jet can be defined in a theoretical calculation by applying an algorithm over a collection of final state particles of known momentum. Experimentally the complete set of final

state particles is unknown and the algorithm must instead be applied over a collection of segmented energy deposits from a range of detector sub-systems. Calorimetry plays a key role in jet measurements by measuring the energy and position of electrons, photons and hadrons⁸.

Jet reconstruction is complicated by the need to calibrate the raw calorimeter data which is simply information of the energy deposited in cells (a small area in $\eta - \phi$ space) over several longitudinal samplings. A raw measurement in a given cell is not a direct measurement of an incident particle's energy and instead the energy of incident species has to be reconstructed from the calorimetry information as a whole. This is a complicated process and without appropriate calibration, detector effects can lead to systematic errors on the energy of reconstructed jets.

ATLAS has a sampling calorimeter which measures a particle shower at discrete points (samplings) along its longitudinal development. As the number of samplings is finite and the energy deposited by a shower is a stochastic process, a statistical error is introduced to the total deposited energy. This is complicated further by a response which is dependent on the P_T spectrum of particles which make up a jet. For instance, the calorimeter will respond differently to a 200GeV jet comprised of 20GeV particles than to a 200GeV jet comprised of 10GeV particles. In addition the uniformity of the response over $\eta - \phi$ may vary due to differing amounts of dead-material and gaps between detector systems (e.g. barrel and end-caps). As referred to previously the ATLAS calorimetry is non-compensating which means that EM energy is more efficiently reconstructed than its hadronic counterpart. This is an issue as the EM fraction of jets fluctuate and represents a considerable proportion of a jet's energy⁹.

Besides the calorimeter response, other factors contribute to jet errors. The solenoidal magnetic field prevents low P_T particles ($< 1\text{GeV}$) from reaching the calorimeter and hence reduces the total measured jet energy. Another factor that decreases the measured energy of a jet is the presence of muons or neutrinos (from heavy-flavor decays) which do

⁸A description of how this is achieved is given in Section 2.2.2.

⁹Approximately one third of a hadronic jet's energy is EM in nature as a result of pion decay ($\pi^0 \rightarrow \gamma\gamma$).

not interact in the calorimeter. In a contrary fashion, the underlying event (created by soft parton-parton interactions) tends to increase the jet energy, as soft scattered particles are swept up into a jet by the jet reconstruction algorithm and contribute to its total energy.

The overall effect of these potential sources of error can be summarised in terms of a statistical and a systematic error. The statistical error leads to a symmetric smearing of a jet's energy about its true value, whereas the systematic error leads to an offset of the mean of the true and experimentally measured energy distributions. These uncertainties are referred to as the jet energy resolution and jet energy scale (JES) error respectively. Corrections can be applied for JES only to within the uncertainty to which it is known, the remaining uncertainty leads to a systematic error.

Further complications arise when considering jet cross-sections in contrast to isolated jet events. To measure an experimental cross-section a knowledge of the integrated luminosity and trigger efficiency is required to normalise the measurement. In addition, the knowledge of the jet energy resolution can become a factor. The reason for this comes from the steeply falling cross-sections (with P_T) that arise in jet measurements. For an experimentally measured jet of a given P_T it is more likely that it is a result of a low P_T jet that has been shifted up in P_T due to resolution smearing than a high P_T jet that has shifted down. This leads to a systematic error that tends to increase the jet cross-section in a given bin. This effect can be corrected for by use of an unfolding procedure but the accuracy is limited by the knowledge of the jet resolution.

5.3.2 Offline Jet Reconstruction and Calibration at ATLAS

The 'Benchmark procedure' of jet reconstruction used by ATLAS is a cone algorithm ($\Delta R = 0.7$). The longitudinal samplings are summed (at the EM scale) to form calorimeter towers which relate to an area of granularity 0.1×0.1 in $\eta - \phi$ space. The highest P_T tower is then used as a seed for the cone algorithm. In this manner a number of jets are defined from the calorimeter information made up of a number of energy deposits or towers. The next step is to calibrate the energy deposits to arrive at a better approximation

of the overall jet energy.

In order to calibrate the jet energy it is first parameterised as a function of the calorimeter cells that comprise it:

$$E^k = \sum_l f(a_l; \overline{\epsilon_{i,j}}) \quad (5.6)$$

where E^k is the energy of jet ‘ k ’, $\bar{\epsilon}$ is a cell’s energy at the EM scale (denoted by the overline symbol) and a_l are a set of parameters. The i and j index the longitudinal sampling and position (pseudorapidity) of the cell, and l is an index of the parameters used. The jet recombination (i.e. the summation of calorimeter cells that comprise the jet) is implicit in the function ‘ f ’ of equation 5.6.

The a_l are chosen to minimise:

$$\chi = \sum_{k=1}^K (E^k - E_{kin}^k)^2 + \beta \sum_{k=1}^K (E^k - E_{kin}^k) \quad (5.7)$$

where E_{kin}^k is the true or reference energy of jet ‘ k ’. The Lagrange multiplier (β) is designed to converge the reconstructed energy to the reference one. The minimisation of Eqn 5.7 is carried out with respect to the a_l parameters and β simultaneously. It should be noted that in general the reference (or true) energy of the jet E_{kin}^k is unknown, so the a_l are in turn parameterised as smooth functions of E_{kin}^k with further parameters b_n .

$$a_l = g_l(b_n; E_{kin}) \quad (5.8)$$

The b_n are established by looking at discrete points where the reference energy of a jet is known and fitting the smooth functions g_l to these points. After this has been carried out Eqn 5.8 can be used to obtain a values of a_l at any E_{kin} as the smooth functions provide the interpolation or extrapolation.

Once the smooth calibration functions (equation 5.8) are known from reference data, the process of reconstructing jet energy can then be carried out in the general case. For a given jet, a starting point is defined for its reconstructed energy using the EM scale calibration. This energy can then be used to determine the values of a_l from equation

5.8. These a_l are then in turn used to reconstruct the energy of the jet using equation 5.6. These steps are iterated until the change in jet energy is less than 1 MeV.

The simplest application of the technique above is in a sampling based calibration. In this method the energy obtained in a given longitudinal sampling is multiplied by a free parameter a_l :

$$E^K = a_1 \overline{E_1^k} + a_2 \overline{E_2^k} + a_3 \overline{E_3^k} + a_4 \overline{E_4^k} \quad (5.9)$$

where the subscripts (1, 2, 3, 4) relate to the presampler, EM calorimeter, hadronic calorimeter and the deposited energy in the cryostat. More sophisticated methods such as the H1 method consider corrections to the energy in individual cells of the calorimeter.

An alternative approach to jet calibration is also being pursued which involves identifying and correcting clusters of energy in the calorimeter before reconstructing them into objects such as jets etc. This local hadronic calibration relies on differentiating between electromagnetic and hadronic energy based on the topological properties of calorimeter clusters and applying corrections accordingly. This approach is more ambitious than the standard benchmark procedure but has the advantage that the calibration is independent of the jet finding algorithm.

5.3.3 Standard Candles for Calibration

The offline correction procedure for jets described above requires data where the true energy of a jet is known (from an independent source) in order to setup the calibration. Test beam runs allow the response of the detector to individual charged hadrons (e.g. from pions) of known energy to be determined. However, a jet cannot be simulated in a controlled way using a test beam. Despite this, by appealing to certain classes and topologies of event it is sometimes possible to infer the jet energy independently of the hadronic calibration.

One such class of events is the hadronic decay of W bosons, which provides a method of low P_T jet calibration. Experimentally measured jets resulting from W decay have an

invariant mass that can be compared to the known mass of the W in order to provide a calibration constraint. At the LHC, $t\bar{t}$ production provides a reliable source of such decays (as the top quark decays predominantly to a W-boson and a b-quark) but these must be separated from background processes. This separation is achieved by demanding high jet multiplicity, b-tagging and an isolated lepton. The procedure allows the jet energy scale to be fixed to within $\sim 1\%$ for jets with P_T between 50GeV and 200GeV. At lower and higher jet P_T the constraint is limited by the ability of Monte Carlo simulations to model out of cone corrections/FSR and jet overlap respectively.

Another approach to jet energy calibration is to connect it to the better known electromagnetic scale. The idea is to demand P_T balance in events with a single jet and entities well understood at the EM scale such as electrons or photons. A good source of these events arise from Z+jet events where the Z decays into charged leptons ($Z \rightarrow e^+e^-$). This procedure allows the jet energy scale to be again constrained to within $\sim 1\%$ for a jet of P_T between 50GeV and a few hundred GeV. At low P_T , however, the P_T of additional soft jets from initial state radiation becomes significant and affects the accuracy of the balancing of the Z decay products with the leading jet. At high P_T the cross-section for the process is too low to provide a sufficient number of events for an accurate calibration.

The high precision ($\leq 1\%$) methods of constraining the jet energy scale are only viable at low jet P_T ($P_T < 500\text{GeV}$) and for many analyses this is sufficient. However, the LHC will produce a considerable number of jets at very high P_T as predicted by the theoretical inclusive jet cross-section (see Figure 3.6). An accurate jet energy scale is important for these high P_T jets if they are to be used in compositeness searches or to constrain PDFs (see Section 4.5).

The offline calibration provides a correction for a jet at any P_T as the correction parameters (a_i) are expressed as smooth functions of the true jet energy (see Equation 5.8). This smooth function provides an extrapolation of the calibration to very high jet P_T . However, as the function is only constrained to high accuracy for jets with P_T below 500GeV it is questionable whether the extrapolation is reliable for jets of 1TeV and above.

5.3.4 Experimental Errors on Jet Measurements

The effect of experimental errors on jet cross-sections can be estimated by using Monte Carlo simulations to model the imperfect response of the detector. Complex computer models attempt to model the interaction of particles with the material of the detector (generated using a parton shower Monte Carlo e.g. PYTHIA) and then apply the same reconstruction algorithms as used in the experiment. Such models can be validated using test beam data to ensure that they provide a reasonable description of detector effects. It is also possible to compare reconstructed events with the ‘truth’ data which can be obtained directly from the Monte Carlo¹⁰.

The ATLAS simulation software (Athena 11.0.4), was used to investigate the effects of reconstruction on jet cross-sections. A comparison was made between reconstructed and ‘truth’ data for Cone R=0.7 jets and is shown in Figure 5.11. The leading (highest P_T) jet was selected from the ‘truth’ and reconstructed events and were required to have $\Delta R < 0.7$ ¹¹.

The data was taken from centrally produced samples¹² of inclusive jet production generated using PYTHIA and reconstructed with the full ATLAS simulation software.

The results show a good agreement in both P_T ranges with a systematic offset of $\sim 1\%$. The standard deviation of the fluctuation is also small at $\sim 4\%$. As a full detector simulation was used, a method of calibration was implemented to obtain jet properties (e.g P_T) from the simulated detector response as described in Section 5.3.2. The calibration constants in these samples were derived by comparing truth and reconstructed jets, for the R=0.7 cone algorithm. With this consideration it is perhaps unsurprising that the agreement between ‘truth’ and reconstructed jets is so good. As actual calibration constants must be obtained using ‘standard candles’ these results should not be interpreted as an overall accuracy of jet reconstruction. However the results do show that the effects of the act of reconstruction itself on jet measurements are expected to be small.

¹⁰Truth data is (of course) not available for actual experimentally recorded collisions in the detector.

¹¹ $\Delta R = \sqrt{\Delta\phi^2 + \Delta\eta^2}$

¹²The sample used were J5 and J6 from the CSC production using Athena version 11.0.4.




Figure 5.11: Comparison of the leading jet in ‘Truth’ and reconstructed inclusive jets in ATLAS (Athena version 11.0.4) for two ranges in P_T . Cone $R=0.7$, jets from calorimeter towers.

In addition a comparison was made of the different jet algorithms as shown in Figure 5.12. The differences between the algorithms are quite considerable $\sim 20\%$, however this does not imply a corresponding uncertainty in the cross-section, as a comparison between data and theory is generally made using the same algorithm.




Figure 5.12: Comparison of the Cone $R=0.4$ and KT (tower) algorithms to Cone $R=0.7$ in ATLAS (Athena version 11.0.4).

5.3.5 Experimental Errors on the Inclusive Jet Cross-Section

As mentioned earlier, in the case of high P_T jets the overall measurement error can be separated into a jet's resolution and scale uncertainty. A simplified approach for modelling these uncertainties is to apply assumed values for the jet energy resolution and scale to theoretical predictions. In this manner the distortion of the measured cross-section as a result of these factors can be assessed.

Once the effect of measurement errors on the cross-section are modelled, the experimental data can be corrected to take this into account. This correction is, however, limited by uncertainty in the input parameters of the model (i.e. jet energy resolution and scale) which have to be ascertained from a knowledge of the detector, its simulation and its calibration. The uncertainty in the jet resolution and scale leads therefore to a systematic error in a jet cross-section.

5.3.6 Statistical Errors

Before considering the effect of jet errors on cross-sections it is instructive to put these into context by estimating the statistical error. This error arises as only a finite number of jet events are recorded to form an experimental cross-section. As a result, the number of jets in a given histogram bin¹³ can suffer statistical fluctuations. By using a theoretical prediction of a cross-section (or experimental data) the error can be estimated using Poissonian statistics¹⁴. This is shown in Figures 5.13, 5.14 and 5.15.

As expected, the statistical error varies inversely with the jet cross-section. Statistical errors increase with jet P_T and rapidity as the cross-section falls (see Figure 3.6) and sampling errors in histogram bins increase. With reference to Figure 5.8, it is clear that even at low integrated luminosities ($0.1fb^{-1}$), theoretical errors resulting from PDFs are likely to dominate for jets with $P_T \sim 2TeV$ and above in the central $0 < |\eta| < 1$ region.

The plots do not take into account the effect of the trigger and prescales which tend to reduce the number of recorded events containing jets with low P_T . However, at high jet P_T

¹³Typically a P_T bin in the case of the inclusive jet cross-section.

¹⁴For a given bin containing ' n ' jets the error is estimated as $\sigma = \sqrt{n}$




Figure 5.13: Statistical Errors for inclusive jets at LHC for $0 < |\eta| < 1$ for three choices of integrated luminosity, (NLOJET, CTEQ6.1 from integration grids, $\mu_r, \mu_f = P_T(max)$)




Figure 5.14: Statistical Errors for inclusive jets at LHC for $1 < |\eta| < 2$ for three choices of integrated luminosity, (NLOJET, CTEQ6.1 from integration grids, $\mu_r, \mu_f = P_T(max)$)

where statistical errors begin to dominate; nearly all events are selected by the single jet trigger and are not pre-scaled. At ATLAS the current planned single jet trigger threshold (without pre-scales) is set to $P_T > 400 GeV$ [33]. Below $P_T = 400 GeV$ a set of pre-scales reduces the number of jet events recorded, however, the increase in statistical error this creates is negligible in comparison with other sources of theoretical and experimental error.

There are discontinuities in the distributions seen in figures 5.13, 5.14 and 5.15 where




Figure 5.15: Statistical Errors for inclusive jets at LHC for $2 < |\eta| < 3$ for three choices of integrated luminosity, (NLOJET, CTEQ6.1 from integration grids, $\mu_r, \mu_f = P_T(max)$)

the P_T bin width changes. This is expected as a larger bin size collects a greater number of events and hence has a reduced statistical error. This could be used as a method to reduce statistical errors at high jet P_T , however, the exponential fall in the inclusive jet cross-section with jet P_T tends to render this an ineffective approach.

Statistical errors affect the ability of experimental data to identify new physics and make precision measurements, but are comparatively easy to quantify. As it is impossible to reduce statistical uncertainty without increasing the size of the underlying sample; statistical errors serve as a benchmark from which the severity of other error sources can be judged.

5.3.7 Jet Energy Resolution Errors

The effect of jet energy resolution uncertainty on the inclusive jet cross-section can be estimated using a theoretical prediction produced via Monte Carlo. The form of the jet energy resolution can be parameterised as shown below:

$$\frac{\sigma}{E} = \frac{a}{\sqrt{E}} \oplus b \quad (5.10)$$

Where E is in GeV, ' a ' is $\sim 50(\%/GeV^{0.5})$ and ' b ' is $\sim 3(\%)$ from the ATLAS calorime-

try specifications (See Table 2.4). The σ relates to the width of a Gaussian distribution which smears the measured energy symmetrically about the true energy. The effect of this smearing can be corrected if the parameters a, b are precisely known but residual uncertainties lead to a mis-correction and an associated systematic error on the experimental cross-section.

This error can be estimated by applying the smearing given in Equation 5.10 to the theoretically calculated cross-section (on a jet by jet basis) and looking at the change in the final cross-section for different choices of a and b . The initial a and b parameter values were chosen to be 50 and 3 respectively (based on the expected performance of the barrel calorimeter) and then varied by 10% and 20%. The resulting change in the cross-section is shown in Figures 5.16, 5.17 and 5.18.




Figure 5.16: Jet Energy Resolution Errors for inclusive jets at LHC for $0 < |\eta| < 1$ for a shift of 10% and 20% of the sampling and constant term parameters (a, b), (NLO-JET,CTEQ6.1 at Born level)

The fractional uncertainty (or error) in the plots are given by:

$$\epsilon = \frac{shifted - central}{central} = \frac{shifted}{unshifted} - 1 \quad (5.11)$$

Where the ‘*central*’ cross section is given by that found using the initial parameters for a, b and the ‘*shifted*’ cross-section is that found using the shifted parameters e.g.




Figure 5.17: Jet Energy Resolution Errors for inclusive jets at LHC for $1 < |\eta| < 2$ for a shift of 10% and 20% of the sampling and constant term parameters (a, b), (NLO-JET,CTEQ6.1 at Born level)




Figure 5.18: Jet Energy Resolution Errors for inclusive jets at LHC for $2 < |\eta| < 3$ for a shift of 10% and 20% of the sampling and constant term parameters (a, b), (NLO-JET,CTEQ6.1 at Born level)

$a(1 + 0.1)$. As can be seen from the figures, the effect of a lack of knowledge of the jet energy resolution does not lead to large uncertainties on the inclusive jet cross-section. The effects are most pronounced at mid rapidities $1 < |\eta| < 2$ where jets may have both large P_T and energy. However, the effect of resolution errors are within anticipated statistical errors at high P_T (even for high integrated luminosities $\sim 1000 fb^{-1}$) and otherwise constrained to within 1 – 2% at low P_T , where PDF and scale errors are expected to

dominate.

5.3.8 Jet Energy Scale Errors

In a similar manner, a jet energy scale error can be applied to jets in the theoretical prediction in order to determine the effect on the overall cross-section. This was carried out for jet energy scale errors of 1%, 5% and 10% and the results are shown in Figures 5.19, 5.20 and 5.21.




Figure 5.19: Jet energy scale Errors for inclusive jets at LHC for $0 < |\eta| < 1$ at 1%, 5% and 10%, (NLOJET, CTEQ6.1), $\mu_r, \mu_f = P_T(max)$

The errors in the plots are calculated in a similar manner to those of the jet resolution. In the case of a large downward shift in the cross-section (i.e. *shifted* $\ll$ *unshifted*) Equation 5.10 tends to -1 , which is observed in the high P_T limit of some of the plots.

As the inclusive jet cross-section is strongly dependent on jet P_T it is also highly sensitive to uncertainties on the jet energy scale. This source of error tends to dominate the experimental uncertainty and unless constrained to low values ($< 1\%$) can exceed the combined theoretical contributions from PDFs and the choice of the hard scale.

The error on the inclusive jet cross-section is seen to grow with increasing jet P_T . This is expected as a 1% error on the jet energy constitutes a larger absolute error on a high P_T jet than a low P_T jet leading to an increased uncertainty on the inclusive jet




Figure 5.20: Jet energy scale Errors for inclusive jets at LHC for $1 < |\eta| < 2$ at 1%, 5% and 10%, (NLOJET, CTEQ6.1), $\mu_r, \mu_f = P_T(max)$




Figure 5.21: Jet energy scale Errors for inclusive jets at LHC for $2 < |\eta| < 3$ at 1%, 5% and 10%, (NLOJET, CTEQ6.1), $\mu_r, \mu_f = P_T(max)$

cross-section. The magnitude of the errors vary with pseudorapidity due to the differing gradients ($d\sigma/dP_T$) of the inclusive jet cross-section in these regions.

In global PDF fits the jet energy scale is a key factor in determining the ability of collider data to constrain distributions. Unfortunately, the jet energy scale errors have a very similar form to those of the PDFs (increasing in P_T and in forward regions) making them difficult to disentangle.

In summary, we have reviewed the main sources of error on inclusive jet measurements

at the ATLAS experiment. For a high P_T (1 TeV), centrally produced jet the most significant theoretical error arises from PDFs ($\sim 12\%$), whilst the dominant experimental error originates in uncertainty on the JES ($\sim 30\%$ for a 5% scale error).

Chapter 6

Bootstrapping the Jet Energy Calibration

6.1 Calibration of high P_T jets

As mentioned earlier, difficulties are encountered calibrating jets at high P_T owing to the low cross-section of processes which allow the hadronic scale to be matched to the better understood electromagnetic scale (see Section 5.3.3). As a result, there is a reliance on the extrapolation of the jet energy scale calibration from the low P_T regime ($< 500\text{GeV}$) to several TeV. The reliability of this extrapolation is likely to decrease at high P_T in the absence of nearby in-situ calibration points and hence any method to provide a cross-check of the JES in this region is valuable for estimating its error.

One such method is to appeal to transverse momentum (P_T) balance in multi-jet events. The premise is that in a multi-jet event the transverse momentum of a single high P_T jet can be inferred from multiple low P_T jets that are well calibrated from in-situ techniques. Once the calibration of the higher P_T jets has been achieved they can subsequently be used to calibrate even higher P_T jets in a similar manner. Using this approach the jet calibration at low P_T from in-situ techniques can be ‘bootstrapped’ to high P_T .

6.2 Jet Balancing Cuts

A set of basic cuts were used to select multi-jet events and to construct a balance between the P_T of the highest P_T jet present (the leading jet) and the other softer jets. The procedure was as follows¹.

- Require events to have at least 3 jets.
- Require 3rd leading jet to have $P_T > 40\text{GeV}$.
- Identify the leading jet.
- Calculate $\mathbf{P}_T^{\text{bal}} = \sum_{i \neq \text{lead}} \mathbf{P}_T^i$ (sum over jets - except leading).
- Require that $|\phi_{\text{lead}} - \phi_{\text{bal}}| > 160^\circ$.
- Calculate the difference $P_T^{\text{dif}} = |\mathbf{P}_T^{\text{lead}}| - |\mathbf{P}_T^{\text{bal}}|$.

The aim of these cuts is to ensure that a multi-jet event is selected where the third jet carries a non-negligible P_T and that there is not a large component of missing P_T transverse to the leading jet axis. The difference in P_T between the leading jet and the summed balance (P_T^{dif}) was recorded in one of a set of histograms depending on the leading jet P_T . These histograms were then fitted with Gaussians and their mean values plotted against the leading jet P_T . If a jet energy scale error were present it is expected to manifest itself as a consistent deviation from zero of P_T^{dif} with P_T^{lead} .

6.3 Jet Balancing using the ATLAS simulation

The cuts were applied to J3-J6 jet events generated with PYTHIA [8] as part of the CSC production using ATLAS software releases 11.0.41 and 11.0.42. The ‘J’ samples relate to inclusive QCD jet production (at LLA) within different kinematic constraints placed on the initial hard-scattering. The cuts are designed to allow production of reasonable

¹Many of these cuts were first used in work by M.Kaneda (ATLAS Jet-Etmiss working group)

numbers of jet events for a range of jet P_T and are shown in Table 6.1. The results are shown in Figure 6.1 for cone $R = 0.7$, truth jets.

Sample	CKIN(3) GeV	CKIN(4) GeV
J1	17	35
J2	35	70
J3	70	140
J4	140	280
J5	280	560
J6	560	1120
J7	1120	2240
J8	2240	∞

Table 6.1: The range of allowed P_T values of the initial hard-scattering in PYTHIA for the ATLAS ‘J’ samples. CKIN(3) and CKIN(4) are the minimum and maximum values of permitted P_T allowed in GeV.




Figure 6.1: Means of $P_T^{dif} = P_T^{lead} - P_T^{bal}$ for J3-J6 Truth events, Cone 0.7

From Figure 6.1, it can be seen that P_T^{dif} is fairly stable for truth jets with increasing P_T^{lead} . However, it is apparent that a consistent positive bias in P_T^{dif} exists, suggesting that either the leading jet is being systematically over-estimated in P_T or the balance jets are being underestimated.

The origin of this bias can be attributed to a number of sources. The first is a selection bias owing to the use of leading jets in the analysis. Jets can contain a component of

invisible energy due to the presence of neutrinos produced by heavy flavour decays. It is hence possible that a leading jet is mis-tagged as sub-leading due to the presence of neutrinos, contributing a positive bias to P_T^{dif} .

Another effect is due to the preference for soft jets to be found in the opposite hemisphere to the leading jet. If a cut on the minimum P_T of jets (or indeed particles) that are recorded is present this can also produce a bias. In addition the underlying event can be swept into large jet cones and create a bias for a hemisphere containing a large number of jets.

It is observed in Figure 6.1 that P_T^{dif} varies in a complex manner with the leading jet P_T . This is unexpected, as there is no compelling theoretical motivation for the apparent discontinuities seen for leading jet P_T of $\sim 400\text{GeV}$ and $\sim 900\text{GeV}$. They may, however, be an artefact of the kinematic cuts used in the samples which were used and hence not expected in real data. At the upper P_T boundary of a ‘J’ sample, owing to the limits in the P_T of the hard-interaction, secondary emissions (which lead to multi-jets) are suppressed. Conversely, at the lower boundary of a ‘J’ sample secondary emissions are encouraged as contributions may arise from the entire generated range. This is likely to explain the instabilities in P_T^{dif} with increasing P_T , which occur. In order to (temporarily) avoid these instabilities and to consider the effect of basic cuts a study was carried out with PYTHIA and a standalone jet algorithm.

6.4 PYTHIA Generator-Level Jet Balancing

The PYTHIA study used a standard CDF Run I cone algorithm (JETCLU) with a default jet-seed of 2.0GeV and cone size of $R = 0.7^2$. The PYTHIA events produced were of the same type found in the ‘J’ samples, however, they were created in a single set between $300\text{--}1000\text{GeV}$ (CKIN 3-4) to avoid problems with the suppression and exacerbation of soft radiation at the boundaries between the ‘J’ samples.

²These are the ATLAS software defaults for full simulation. The default ATLAS cone algorithm is currently based on Tevatron implementations [34].

No cut was applied on the minimum P_T of particles that could comprise a jet or on the jets themselves which are considered in the balancing. This was done to avoid potential problems associated with P_T becoming lost. Furthermore, in an attempt to isolate different sources of the bias in P_T^{dif} , neutrinos were also included in the jet reconstruction. The analysis was carried out as before whilst the cone size, P_T cut on particles and jets and jet seed threshold were individually altered to examine the effects on P_T^{dif} . The effect of changing the cone size is shown in Figure 6.2.




Figure 6.2: Means of $P_T^{dif} = P_T^{lead} - P_T^{bal}$ for PYTHIA + jet-finder (JETCLU), Jet-Seed=2.0GeV

From Figure 6.2, it can be seen that changing the jet cone size has a strong effect on P_T^{dif} . As the cone size increases, P_T^{dif} decreases and becomes negative. Larger cones collect more energy particularly from the underlying event and due to the greater multiplicity of jets in the hemisphere away from the leading jets this affects P_T^{dif} .

The P_T^{dif} was found to be less sensitive to the other parameters, introducing a 10GeV Jet P_T cut created a small ~ 0.2 GeV (negative) bias; a 0.5 GeV particle P_T cut and lowering the jet seed threshold produced no discernible effect. Neglecting neutrinos from the jet finding produced a similarly small $\sim 0.2 - 0.3$ GeV positive bias. However, Figure 6.2 suggests that the bias on the jet balancing is strongly dependent on the jet algorithm used and increased sensitivity to all parameters (except neutrinos) was found when the

study was repeated with a smaller cone size of 0.4 and seed threshold of 1.5 GeV (i.e. the ATLFAST defaults).

The advantages of a choice of $R = 0.7$ for the cone radius arises by accident from the kinematics of the multi-jet events. The fragmentation (i.e. out of cone corrections) losses are less dependent on the size of the cone than the underlying event and act in an opposite sense [24]. Hence a cone size can be tuned to balance these effects³.

6.5 Determining Jet Energy Scale Uncertainty

A possible JES non-linearity was simulated by introducing a systematic JES error which increases uniformly with P_T , e.g. of the form of a 1% increase in JES over a jet P_T increase of 500GeV. As the lower P_T jets are expected to be well calibrated from in-situ methods this non-linearity was introduced for jets with $P_T > 500\text{GeV}$. The effect of introducing this error for the CSC truth data is shown below in Figure 6.3.




Figure 6.3: Means of $P_T^{dif} = P_T^{lead} - P_T^{bal}$ for ATLAS-CSC, J2-J6 truth jets, Cone 0.7 with a JES introduced for jet $P_T > 500\text{GeV}$ of 10% over 500GeV and 1% over 500GeV

From Figure 6.3, using the truth data a 10% JES error that starts within 500GeV would be clearly visible, however a 1% error over the same range could prove difficult to

³And any additional (constant) systematic error in P_T^{dif} .

detect. In practice, truth data is unavailable and the practice of jet balancing becomes more complicated for reconstructed events in experiments.

The first difficulty is that due to the smearing of jet P_T (from jet energy resolution), the transverse momentum of jets does not, in general, balance on an event by event basis. In addition, fluctuations in the P_T of jets due to resolution can also lead to a mis-identification of the leading jet in the balancing procedure which creates a positive bias in P_T^{dif} . Furthermore, a jet that has its P_T shifted up by a resolution fluctuation is more likely to be binned in a histogram at higher P_T . As the population of jets at low P_T far exceed those at high P_T this can lead to a ‘landslide’ effect, introducing a further positive bias to P_T^{dif} .

The practicalities of real measurements introduce numerous implicit cuts on particle P_T due to the solenoidal magnetic field, and the clustering over calorimeter cells as opposed to the truth particles. Figure 6.4 shows P_T^{dif} for reconstructed events in a manner similar to Figure 6.3.




Figure 6.4: Means of $P_T^{dif} = P_T^{lead} - P_T^{bal}$ for ATLAS-CSC, J2-J6 reconstructed events, Cone 0.7 Tower Jets with a JES introduced for jet $P_T > 500\text{GeV}$ of 10% over 500GeV and 1% over 500GeV

From Figure 6.4, it is clear that a 10% JES error would be clearly identified, however the difference in bias between the series for no JES and the truth (Figure 6.3) are considerable. For the reconstructed jets there is now a trend of increasing bias in P_T^{dif} with

leading jet P_T . If the reconstructed jet P_T balance was taken naively this trend could be mistaken for a JES error.

The study suggests that by looking for strong deviations in the $P_T^{dif} = P_T^{lead} - P_T^{bal}$ distributions that large (10%) deviations which build up gradually over a P_T of $\sim 500 GeV$ should be readily observable. However, precision jet balancing is complicated by the dependency on the jet algorithm chosen and the effects of reconstruction, especially resolution smearing. If the resolution effects can be well understood at high P_T then it may be possible to use Monte Carlo techniques to model and then correct for the bias they contribute. The limits on the accuracy of the JES that could be obtained from such a method probably result from the confidence in such a correction.

With this premise, a conservative estimate for the JES accuracy of the bootstrap method can be obtained by examining the bias in Figure 6.4 for the series where no jet energy scale error is applied. With reference to Figure 6.3 it can be determined what non-linearity applied at truth level leads to a similar bias at the reconstructed level (with no JES error) in the region $P_T > 500 GeV$. In order to quantify this bias, the gradient of a linear fit to the data (in Figures 6.3 and 6.4) in this region was considered and is shown in Figure 6.5. The process leads to an estimate of the accuracy of $\sim 1.5\%/500 GeV$, however this only relates to the ‘bootstrapping’ and does not consider errors that might have been propagated from the low ($P_T < 500 GeV$) region.




Figure 6.5: Means of $P_T^{dif} = P_T^{lead} - P_T^{bal}$ for reconstructed jets with no JES error and truth jets with a 1.5%/500GeV JES error. The thick continuous line is a linear fit of the reconstructed data. The truth data have been shifted in the y co-ordinate for ease of comparison. ATLAS-CSC, J2-J6 reconstructed events.

Chapter 7

Quark Compositeness at the LHC

7.1 Sensitivity to New-Physics

The inclusive jet cross-section is sensitive to ‘new physics’ at high P_T , and one possible scenario which can create an observable signal is quark compositeness. In the Standard Model, quarks are treated as fundamental particles, however, it is possible that quarks are comprised of a number of particles (preons) bound together by a new type of interaction referred to as hypercolor. The quark compositeness has an associated scale Λ , which defines a length scale for the quark and an energy at which hypercolour effects become large¹. If the compositeness scale (Λ) is sufficiently high then quark substructure would have escaped detection at previous experiments but may be discernible with the high centre of mass energies available at the LHC.

The motivation for quark compositeness stems from the desire to explain the existence of the three generations of the Standard Model (See Table 1.3) in a unified manner. If the ‘fundamental’ particles of the Standard Model were composite, then the mass-hierarchy of the generations could be explained as the result of excited states of bound hypercolor systems. Furthermore, the connection between quarks and leptons (regarding the hierarchy) would follow naturally if they were both bound states of fundamental

¹This may be viewed in analogy with $\Lambda_{QCD} \sim 200 - 300 MeV$ which determines a length scale for hadrons and an energy at which QCD effects become non-perturbative.

preons [4].

Compositeness itself, relates to a class of theories as the premise does not lead to a precise description of the underlying hypercolor interaction. It is possible, however, by appealing to general theoretical arguments to predict the likely experimental consequences. At scales well below Λ , it is thought that compositeness can be modelled as a contact interaction (generated by constituent exchange) in the same manner that the weak interaction may be below the electroweak scale ($O(100GeV)$).

The value of the compositeness scale (Λ) is not predicted by theory, but is expected to be above the electroweak scale and a lower limit exists from the Tevatron experiment which demands $\Lambda > 2.4TeV$ (at 95% confidence level)[35]. It is interesting to note that such a scale is considerably higher than the known mass of the quark. This is unusual as it would normally be expected that the substructure scale sets the mass of the composite object. For example, in the case of hadrons, the overall mass is governed by the QCD scale as opposed to the constituent mass of its partons. This difficulty can however be overcome in a natural way (according to t'Hooft) by appealing to symmetry constraints [36], [37].

A hypercolor interaction may permit flavour changing transitions, although such a feature would demand a very high compositeness scale ($> 100TeV$) due to experimental constraints [38]. It is possible, however, to construct a hypercolour interaction which permits only flavour conserving processes, with the much lower constraint ($> 2.4TeV$) provided by the Tevatron.

Such a flavour conserving process is given in equation 7.1 below. For simplicity the compositeness is considered only in left handed quarks and mediated by a colour singlet, isoscalar exchange [39]:

$$\mathcal{L}_{qq} = \eta_0 \frac{g^2}{2\Lambda^2} \bar{q}_L \gamma_\mu q_L \bar{q}_L \gamma^\mu q_L \quad (7.1)$$

Where g is an effective coupling constant and $\eta_0 = \pm 1$. This contact interaction formulation is valid in the case when the experimentally probed scale $\hat{s}$ is much smaller

than the compositeness scale Λ^2 .

The (effective) Lagrangian above is combined with that of the Standard Model and used to calculate cross-sections. This combination leads to two types of term not originally present in the Standard Model from ‘interference’ and ‘pure-contact’. The interference terms arise from QCD and compositeness processes sharing final states and may decrease or increase the magnitude of the Standard Model amplitude. In this model, constructive interference occurs for $\eta_0 = -1$ and destructive for $\eta_0 = +1$. The ‘pure-contact’ interaction is independent of the choice of η and can be thought of as the compositeness process that would take place regardless of the Standard Model.

For the inclusive jet cross section the effects of Eqn 7.1 are anticipated in events with high P_T jets, as these are dominated by high Q^2 quark scattering. As a result of the ‘pure-contact’ terms compositeness leads to an excess in the inclusive jet cross-section at high P_T regardless of the choice of η .

7.2 Compositeness Searches in Inclusive Jets

The compositeness scenario defined in equation 7.1 was investigated using the PYTHIA 6.4 event generator. The inclusive jet cross-section was generated first using standard QCD interactions and then including compositeness processes defined at a scale Λ to determine the change in the cross-section. This is shown in Figures 7.1, 7.2 and 7.3 where all quarks are considered composite and the compositeness effects are constructive.

As anticipated the effects of compositeness are most prominent at high P_T in central regions for a low Λ . It is noticeable that the strength of the compositeness signal appears to depend predominantly on the P_T of the jet only and is to a large degree independent of the rapidity region chosen. However, the central rapidity region offers the highest P_T (and hence Q^2) events and for a given P_T the errors on the inclusive jet cross-section tend to be greater in forward regions. As a result the highest sensitivity to compositeness for

²This may be considered in analogy to the Fermi contact interaction which describes the Weak interaction at scales well below M_W .




Figure 7.1: Change in inclusive jets cross-section at LHC for $0 < |\eta| < 1$ for four choices of compositeness scale, (constructive interference), PYTHIA 6.4




Figure 7.2: Change in inclusive jets cross-section at LHC for $1 < |\eta| < 2$ for four choices of compositeness scale, (constructive interference), PYTHIA 6.4

inclusive jets is found in the central region.

The figures shown are for constructive interference ($\eta_0 = -1$), however similar results are in the destructive case although the signal strengths are typically diminished. In the case of inclusive jets an excess in the cross-section is observed at high P_T regardless of the interference sign (η_0) as a result of the ‘pure-contact’ terms.

When considering the significance of a compositeness signal, the uncertainty on the inclusive jets cross-section has to be considered. In particular as the compositeness signal




Figure 7.3: Change in inclusive jets cross-section at LHC for $2 < |\eta| < 3$ for four choices of compositeness scale, (constructive interference), PYTHIA 6.4

is manifest at high jet P_T , where the cross-section is small, the effect of statistical uncertainties must be taken into account. 250 million PYTHIA events were used to generate figures 7.1, 7.2 and 7.3, corresponding to an integrated luminosity of $\sim 370 fb^{-1}$ ³. At this luminosity, corresponding to many years of high luminosity data-taking the statistical fluctuations at high P_T are very low and the systematic effects dominate.

The jet energy scale uncertainty limits the significance of a compositeness signal. As observed in earlier plots (Figures 5.19, 5.20 and 5.21), jet energy scale uncertainty can lead to an increase in the cross-section with jet P_T in a similar manner to a compositeness signal. The JES uncertainty is problematic as the exact form of a non-linearity may not be known *a priori*. For instance, it is expected that the jet energy scale will be well understood at low P_T but may deviate with increasing jet P_T in an unknown manner beyond the final in-situ calibration point.

The uncertainty on a potential compositeness signal caused by JES and PDF errors could be reduced by appealing to characteristics of the process. Most notable of these is the near independence of a compositeness signal to rapidity (unlike JES and PDF errors). Hence, if a consistent excess in the inclusive jet cross-section were observed occurring at approximately the same P_T in all rapidity regions this would strengthen the evidence for

³A cut was placed on the minimum P_T in the hard-scattering of 500GeV

new physics (assuming uniformity of the calorimeter).

However, as the inclusive jet cross-section is far smaller in forward than in central regions, it is possible that a compositeness signal could be observed at high P_T (3 TeV) in the central region without an available analogue in the forward regions due to a lack of statistics. The forward rapidity information is nevertheless useful in ascertaining whether an observed excess in the cross-section could be a result of PDF errors.

PDF errors arise from uncertainties in the description of the proton's parton content and are manifest in all pseudorapidity regions. However, the forward region is particularly sensitive to the poorly constrained high- x gluon PDF at high P_T as a result of the kinematics of an asymmetric collision. The forward ($2 < |\eta| < 3$) inclusive jet cross-section at modest P_T can hence provide a constraint on the high- x gluon content and a cross-check that an excess in the central high P_T cross-section is not a result of PDF uncertainty. The advantage of using the forward region to constrain PDF errors is that the modest P_T values present reduce the risk of contamination of new physics (e.g. compositeness).

The jet energy scale is difficult to constrain in the manner described above. As JES is a mis-calibration error it is possible (and indeed likely) that it will be a function of a jet's P_T . A non-linearity could hence occur for high P_T jets (3TeV-4TeV) and fake a compositeness signal in the central region, without affecting the forward regions at lower P_T . The mid- P_T inclusive jet cross-section ($1 < |\eta| < 2$), however, has access to multi-TeV jets, allowing for comparisons with jets at high P_T in central regions. A JES error would be manifest as a much larger deviation from the Standard Model at higher rapidity than at central rapidity. This approach is limited by the maximum P_T of jets available at mid-rapidities and is hence not able to assist compositeness searches for very high Λ .

A possible approach to determining whether an excess in the cross-section at high P_T is a result of compositeness would be to use the Standard Model to model the mid and forward rapidity regions in order to assess the JES and PDF errors. The error present in the central region could then be better predicted and used to interpret an excess in the cross-section.

7.2.1 Analysing the Inclusive Jets Cross-Section

A method of hypothesis testing can be applied in order to determine whether a compositeness signal has been found. It is, however, simpler to assess whether a given set of experimental data is in agreement with the Standard Model (the null hypothesis). This approach is reasonable given that an excess in the inclusive jet cross-sections may relate to new physics but not necessarily to compositeness.

A standard approach to hypothesis testing is to apply a chi-squared test, which takes into account theoretical predictions, experimental data and their statistical (finite sample) errors over many bins in P_T :

$$\chi^2 = \sum_i^n \frac{(T_i - y_i)^2}{\delta y_i^2} \quad (7.2)$$

Where the sum is over the bins in the inclusive jet cross-section, T_i and y_i are theoretical predictions and experimental data for the bin ‘ i ’, whilst δy_i is the experimental data’s statistical error. The value of χ^2 may then be compared with standard statistical tables to assess whether a significant deviation from anticipated results has been found.

This approach neglects the systematic errors on jet measurements and predictions which tend to dominate the statistical (finite sample) errors up to very high P_T . In order to allow for the effects of systematic errors, the CDF collaboration modified equation 7.2 to include sources of systematic error which had well defined shapes in terms of P_T [21]. The sources of systematic error were then allowed to vary by a multiplier determined in a multi-parameter minimisation of the overall χ^2 . This approach gives the systematic errors the opportunity to decrease the significance of any non-standard model signal if they are able to do so.

The CDF approach requires an accurate determination of how the uncertainty in the cross-section varies as a function of P_T for each source of systematic error. This is in practice difficult to ascertain and requires a good understanding of the detector, which is currently unavailable for ATLAS as the detector has not yet been fully commissioned⁴. A

⁴Typically the jet energy scale error is broken down into multiple sources such as high and low P_T

simplified approach which allows for an estimate of the sensitivity of inclusive jet cross-sections to compositeness signals can be made by considering the overall error bands of individual sources of systematic errors. In this approach, the overall magnitude of systematic errors on a P_T bin are calculated by a simple summation. The bin is only considered in the analysis if the discrepancy from the standard model $|T_i - y_i|$ is greater than the sum of systematic errors⁵. The analysis of the bins with a definite signal are carried out using the basic chi-square equation (7.2), where the experimental data is corrected by the systematic errors to minimise the discrepancy $|T_i - y_i|$.

The overall effect of this simplified approach is to allow the systematic errors unlimited freedom to follow the form of the compositeness signal and diminish its statistical significance. It can hence be considered as a conservative estimate of the sensitivity of the ATLAS inclusive jets data to potential compositeness signals. The resulting χ^2 and degrees of freedom for 3 choices of jet energy scale errors with PDF errors (CTEQ6.1) included are shown in Tables 7.1 and 7.2 for constructive and destructive interference respectively. The degrees of freedom vary as a bin was only considered if there was an identifiable signal above systematic errors. The significance or probability that the χ^2 is the value shown or above is given in Table 7.3. The value reflects the probability of the null hypothesis, in this case the Standard Model.

$\Lambda(\text{GeV})$	Jet Energy Scale Uncertainty					
	1%		5%		10%	
	χ^2	dof	χ^2	dof	χ^2	dof
5000	40911	22	26865	20	17552	19
10000	1013	16	478	13	209	9
15000	55	9	11	6	0.2	1
20000	2	4				

Table 7.1: Chi-Squared and degrees of freedom for inclusive jet cross-section at ATLAS for three compositeness scale choices (constructive interference) considering JES and PDF errors, $0 < |\eta| < 1$, $L = 10\text{fb}^{-1}$.

From Table 7.3 it can be seen that the inclusive jet cross-section is sensitive to a

hadron response, stability etc

⁵A minimum bin-occupancy of 10 was also required

$\Lambda(\text{GeV})$	Jet Energy Scale Uncertainty					
	1%		5%		10%	
	χ^2	dof	χ^2	dof	χ^2	dof
5000	20809	20	14568	19	10216	17
10000	200	11	105	9	37	6
15000	0.83	2				
20000						

Table 7.2: Chi-Squared and degrees of freedom for inclusive jet cross-section at ATLAS for three compositeness scale choices (destructive interference) considering JES and PDF errors, $0 < |\eta| < 1$, $L = 10fb^{-1}$.

$\Lambda(\text{GeV})$	Jet Energy Scale Uncertainty					
	1%		5%		10%	
	cons	dest	cons	dest	cons	dest
5000	$< 10^{-15}$	$< 10^{-15}$	$< 10^{-15}$	$< 10^{-15}$	$< 10^{-15}$	$< 10^{-15}$
10000	$< 10^{-15}$	$< 10^{-15}$	$< 10^{-15}$	$< 10^{-15}$	$< 10^{-14}$	2×10^{-6}
15000	1.4×10^{-8}	0.66	0.10		0.66	
20000	0.31					

Table 7.3: Significance of Chi-Squared (probability of null-hypothesis) for the inclusive jet cross-section at ATLAS for three compositeness scale choices for both constructive and destructive interference, considering JES and PDF errors, $0 < |\eta| < 1$, $L = 10fb^{-1}$.

compositeness signal for $\Lambda \leq 10\text{TeV}$. At $\Lambda = 15\text{TeV}$ it might be possible to discern a compositeness scenario, provided that its effects were constructive and that JES uncertainty could be constrained to within 1%. Elsewhere and at higher scales Λ , however, the signal could be interpreted as a systematic JES or PDF error. The significances shown relate to $10fb^{-1}$. A higher integrated luminosity would increase the significance of any signal and allow for an increase in the P_T range studied. However, there is uncertainty in the magnitude of systematic errors and as a result small signal strengths would probably be attributed to this in the first instance.

7.3 Di-Jet Angular Distributions

An alternative approach to compositeness searches is to appeal to the geometry of jet events. The advantage of this method is that only the trajectory of jets needs to be

measured as opposed to their precise P_T , granting it immunity to jet energy scale errors. A di-jet angular separation variable is defined as

$$\chi = e^{|\eta_1 - \eta_2|} \quad (7.3)$$

Where η_i is the pseudorapidity of jet i . The variable χ is equal to 1 for jets of the same rapidity and becomes larger as their η -separation increases. The χ variable can be related directly to the scattering angle θ^* in the centre of mass frame.

$$\chi = \frac{1 + |\cos\theta^*|}{1 - |\cos\theta^*|} \quad (7.4)$$

The relationship is shown in Figures 7.4 and 7.5.




Figure 7.4: Relationship of the variable χ with the scattering angle in the centre of mass frame θ^*

QCD processes tend to produce an approximately flat χ distribution as a result of a preference for forward production. Compositeness processes are more isotropic in nature leading to an excess at low χ . The effect of a compositeness signal on the di-jet angular distribution was investigated using the same PYTHIA model as before with the χ variable being measured in four bins of invariant jet mass (for the di-jet system). Although a jet energy measurement would be required to obtain the invariant di-jet mass, it would not have to be known to high precision (unlike the inclusive jet cross-section) to place




Figure 7.5: Relationship of the scattering angle in the centre of mass frame θ^* with the variable χ

events into wide bins. Four bins in invariant mass were chosen following the ATLAS TDR analysis [5], of $2000 < M_{jj} < 2300 \text{ GeV}$, $2300 < M_{jj} < 2800 \text{ GeV}$, $2800 < M_{jj} < 3400 \text{ GeV}$ and $M_{jj} > 3400 \text{ GeV}$. Jets were only considered with $|\eta| < 3$ and the leading jet was required to have $P_T > 400 \text{ GeV}$. The results are shown in Figures 7.6, 7.7, 7.8 and 7.9 for constructive interference.




Figure 7.6: Change in di-jet cross-section (w.r.t SM) for dijets with invariant mass $2000 < M_{jj} < 2300 \text{ GeV}$ for four choices of the compositeness scale Λ (constructive interference), cross-sections normalised to SM.

The deviation from the Standard Model is most prominent at low χ for high invariant




Figure 7.7: Change in di-jet cross-section (w.r.t SM) for dijets with invariant mass $2300 < M_{jj} < 2800 \text{ GeV}$ for four choices of the compositeness scale Λ (constructive interference), cross-sections normalised to SM.




Figure 7.8: Change in di-jet cross-section (w.r.t SM) for dijets with invariant mass $2800 < M_{jj} < 3400 \text{ GeV}$ for four choices of the compositeness scale Λ (constructive interference), cross-sections normalised to SM.

mass di-jets (which relate to high Q^2 interactions) as expected. In order to assess the significance with which this signal can be determined, the appropriate backgrounds have to be considered.




Figure 7.9: Change in di-jet cross-section (w.r.t SM) for dijets with invariant mass $M_{jj} > 3400 \text{ GeV}$ for four choices of the compositeness scale Λ (constructive interference), cross-sections normalised to SM.

7.3.1 Theoretical Uncertainties

Theoretical uncertainties on the predicted di-jet angular cross-section arise from PDFs and the use of a fixed order calculation (scale errors) in a similar manner to that of inclusive jets. These errors can be efficiently estimated by the use of an integration grid method, eliminating the need to repeat the time consuming Monte Carlo calculation. A choice must, however, be made as to which MC generator to use.

The signal significance for compositeness (in Figures 7.6, 7.7, 7.8 and 7.9) was generated using the PYTHIA Monte Carlo generator, which is effectively a leading-log approximation. Such generators are by construction likely to underestimate the effect of a hard secondary emission (i.e. 3-jet events) as a full treatment of NLO effects is absent. It would be expected that the absence of NLO contributions would not dramatically affect the significance of a compositeness signal as shown in the earlier figures as they would largely cancel out in the ratio of cross-sections.

However, in a practical compositeness analysis a comparison is made between experimental data and the theoretical prediction for the Standard Model. If there is uncertainty in the shape of the SM predictions then this will limit the confidence in any detected signal. The di-jet angular distribution was hence calculated for the Standard Model with

PYTHIA (at LO and LLA) and NLOJET⁺⁺ (at LO and NLO) in order to assess the effect of different contributions to the shape of the cross-section.




Figure 7.10: The Standard Model di-jet cross-section for di-jets with invariant mass $2000 < M_{jj} < 2300 \text{ GeV}$ from PYTHIA and NLOJET⁺⁺ at leading order for CTEQ5l




Figure 7.11: The Standard Model di-jet cross-section for di-jets with invariant mass $M_{jj} > 3400 \text{ GeV}$ from PYTHIA and NLOJET⁺⁺ at leading order for CTEQ5l

Figures 7.10 and 7.11 show the agreement of PYTHIA and NLOJET⁺⁺ when the calculation is carried out at leading order⁶. The NLOJET⁺⁺ prediction was made using an integration grid of dimension $(10 \times 10 \times 10)$, and the CTEQ5l PDF in order to agree with

⁶This involves turning off both initial and final state radiation parton shower in PYTHIA, and neglecting NLO corrections in NLOJET⁺⁺

the PYTHIA default. As only changes in shape are important regarding the compositeness search the overall NLOJET⁺⁺ cross-section over the range shown was normalised to the PYTHIA prediction.

As may be seen from Figures 7.10 and 7.11, at leading order the PYTHIA and NLOJET⁺⁺ predictions are in agreement in terms of shape for all χ at both low and high di-jet invariant mass. This is perhaps unsurprising, however, as the calculations are effectively identical.

A change in shape can be seen however between the LO and NLO predictions for the di-jet cross-section from NLOJET⁺⁺. The ratio of these two predictions is shown in Figures 7.12 and 7.13. The NLO prediction has an excess at low χ in a similar manner to a compositeness signal. As a consequence it is possible that the absence of higher order Standard Model contributions in theoretical calculations could be incorrectly interpreted as a compositeness signal.




Figure 7.12: The ratio of the Standard Model di-jet cross-section prediction for di-jets with invariant mass $2000 < M_{jj} < 2300 \text{ GeV}$ from NLOJET⁺⁺ at leading and next-to-leading order.

Another difference is seen when comparing predictions at LO and LLA. This is shown for PYTHIA in Figures 7.14 and 7.15. The leading-log approximation leads to lower predictions of the cross-sections as the parton shower tends to degrade the energy of a jet (in favour of forming multi-jet systems) and hence reduces the value of the di-jet mass.




Figure 7.13: The ratio of the Standard Model di-jet cross-section prediction for di-jets with invariant mass $M_{jj} > 3400\text{GeV}$ from NLOJET⁺⁺ at leading and next-to-leading order.

This reduction in cross-section is stronger at high χ which relates to events with highly separated jets in forward regions. For a given di-jet invariant mass bin, final-state partons in forward regions (high χ events) have a greater energy than those produced centrally. This higher energy of partons leads to a longer ‘evolution’ of partons to the hadronization scale in the parton shower and hence a greater probability of additional radiation. Accordingly the effects of LLA corrections (in comparison with LO) are more prominent at high χ . In addition the parton shower effects are more prominent for di-jets with low invariant mass as secondary emissions due to the parton shower are usually soft and hence have a larger (proportional) effect on less energetic jets.

Considering the shape of the LLA distribution (i.e. irrespective of its normalisation); the region at low χ which is sensitive to the compositeness signal remains quite flat and hence neglecting leading log contributions is unlikely to lead to a false ‘discovery’ of compositeness.

As a failure to consider leading log corrections is less likely to fake a compositeness signal than neglecting next-to-leading order terms, it is preferable to use a NLO prediction to compare with experimental data. This is a reasonable choice as NLO predictions are considered to be more accurate than LO or LLA when considering the configurations of




Figure 7.14: The ratio of the Standard Model cross-section prediction for di-jets with invariant mass $2000 < M_{jj} < 2300 \text{ GeV}$ from PYTHIA at leading order and leading-log approximation.




Figure 7.15: The ratio of the Standard Model cross-section prediction for di-jets with invariant mass $M_{jj} > 3400 \text{ GeV}$ from PYTHIA at leading order and leading-log approximation.

hard jets where ATLAS has sensitivity to compositeness. The accuracy of these NLO predictions have to be taken into account in the same way as for the inclusive jets earlier.

An integration grid was made for the χ angular distribution for di-jets in the four invariant mass ranges used earlier⁷. This allowed the PDF errors to be calculated using equation 5.5 and for the scale errors to be determined using the same method as in section

⁷A $(10 \times 10 \times 10)$ grid was found to provide sufficient accuracy to recreate the standard NLO produced cross-section for CTEQ6.1

5.2.1. This is shown in Figures 7.16, 7.17, 7.18 and 7.19 for the invariant mass boundaries used earlier.




Figure 7.16: PDF and Scale errors on SM di-jet cross-section for di-jets with invariant mass $2000 < M_{jj} < 2300 \text{ GeV}$ from an integration grid (CTEQ6.1 error sets used)




Figure 7.17: PDF and Scale errors on SM di-jet cross-section for di-jets with invariant mass $2300 < M_{jj} < 2800 \text{ GeV}$ from an integration grid (CTEQ6.1 error sets used)

The PDF errors in the plots are approximately flat in terms of χ although the overall magnitude of this uncertainty increases with increased di-jet invariant mass. This means that PDF errors are unlikely to present problems in terms of faking a compositeness signal as the shape of the cross-section is largely unaffected. The scale errors lead to an increase




Figure 7.18: PDF and Scale errors on SM di-jet cross-section for di-jets with invariant mass $2800 < M_{jj} < 3400 \text{ GeV}$ from an integration grid (CTEQ6.1 error sets used)




Figure 7.19: PDF and Scale errors on SM di-jet cross-section for di-jets with invariant mass $M_{jj} > 3400 \text{ GeV}$ from an integration grid (CTEQ6.1 error sets used)

in the cross-section at low χ (and hence a shape change) but this is less than 5% within the range $1 < \chi < 10$ over all invariant masses shown.

The scale errors are noticeably smaller than the difference between the LO and NLO predictions found with NLOJET⁺⁺ earlier. This may be viewed in analogy with the scale dependence of the NLO inclusive jet cross-section predictions and its ‘k’ factors as shown in Section 5.2.1.

7.3.2 Experimental Uncertainties

In addition to theoretical uncertainties it is possible that the reconstruction of experimental data will also lead to errors being introduced to the χ cross-section.

In order to investigate this possibility a sample of jets produced by the ATLAS software was used to determine the difference between ‘truth’ and ‘reconstructed’ events⁸. The sample chosen was the ATLAS J6 CSC data (see Chapter 6.1, table 6.1). The result is shown in Figure 7.20.




Figure 7.20: Experimental error based on the ratio of ‘truth’ to reconstructed data for ATLAS CSC J6 sample (212,000 events). The acceptance in the di-jet mass has been widened to $M > 2000 \text{ GeV}$ owing to the small sample size available.

As can be seen from the plot there is no apparent bias in χ as a result of the reconstruction. It should be noted that the ATLAS J6 sample contains jets with $560 < P_T < 1120 \text{ GeV}$ and as a result Figure 7.20 does not contain all contributions for the di-jet invariant mass range $m > 2000 \text{ GeV}$ ⁹. With regard to reconstruction errors, however, the results can be expected to provide a reasonable prediction, insofar as the simulation software provides an accurate description of the real detector.

⁸Truth results arrive from applying the jet algorithms on the kinematics of the final state hadrons taken directly from the generator (e.g. PYTHIA). In the case of reconstructed events, the effect of the detector’s response to the generated particles is considered.

⁹This is due to the absence of higher P_T , centrally produced di-jets.

7.3.3 Analysis Of Di-Jet Angular Distributions

In the same manner as the inclusive jet cross-section, a method of hypothesis testing is required in order to determine whether a potential compositeness signal can be explained adequately by a combination of statistical and systematic errors. It can be seen that the PDF and to an extent the fixed scale errors are flat in χ . This suggests that using a variable which considers a change in shape (irrespective of the overall normalisation) could provide a powerful way of separating a signal from these systematic errors.

Such a variable R_χ ¹⁰ considers the ratio of number of events below a given value of χ (ϵ_0) to that above this value within a defined range of χ .

$$R_\chi = \frac{N(\chi < \epsilon_0)}{N(\chi > \epsilon_0)} \quad (7.5)$$

The value of R_χ is calculated for $1 \leq \chi \leq 10$ and ϵ_0 is a free parameter which may be optimised to maximise the compositeness signal's significance¹¹. From previous plots (Figures 7.6, 7.7, 7.8 and 7.9) a compositeness signal is expected to be manifest as an increase in R_χ above the Standard Model prediction.

The benefits of switching to this new variable can be seen in the estimated errors arising from PDFs and fixed order calculations as show in Tables 7.4 and 7.5 for $\epsilon_0 = 5$. In comparison with the error bands as shown in Figures 7.16-7.19 the errors are much smaller. Furthermore, the scale errors now dominate those from PDFs, as only a change in shape of the χ distribution can contribute to uncertainty on R_χ as opposed to a change in its overall normalisation.

In a similar approach to the inclusive jets analysis, the strength of a compositeness signal can be assessed by the ratio of $R_\chi^\Lambda / R_\chi^{SM}$ where R_χ^Λ and R_χ^{SM} relate to the R_χ (c.f. Eqn 7.5) found for a compositeness signal and the Standard Model respectively. The results are shown in Tables 7.6 and 7.7 for the case of constructive and destructive interference respectively. The results are given with the statistical error expected for an

¹⁰Similar to that used in the TDR analysis [5].

¹¹An earlier analysis chose a value of $\epsilon_0 = 5$ [5].

$m(\text{GeV})$	PDF Uncertainty %	
	Up	Down
2000 $\rightarrow$ 2300	0.34	0.26
2300 $\rightarrow$ 2800	0.33	0.25
2800 $\rightarrow$ 3400	0.35	0.26
> 3400	0.51	0.38

Table 7.4: PDF errors (%) on R_χ as determined from CTEQ6.1 error sets (using integration grids).

$m(\text{GeV})$	$\mu_{r/f}$ Scale Uncertainty %	
	Up	Down
2000 $\rightarrow$ 2300	1.81	0.71
2300 $\rightarrow$ 2800	1.88	0.74
2800 $\rightarrow$ 3400	2.12	0.86
> 3400	2.57	1.01

Table 7.5: $\mu_{r/f}$ scale uncertainty (%) on R_χ as determined from integration grids by varying scales between $0.5 \rightarrow 2.0 P_T$ of the leading jet and considering maximum changes in cross-section.

integrated luminosity of $10fb^{-1}$ and a combined systematic error based on the PDF and scale uncertainty¹².

$m(GeV)$	Compositeness Signal R_x^Λ/R_x^{SM}								Systematic Error
	$\Lambda = 5TeV$		$\Lambda = 10TeV$		$\Lambda = 15TeV$		$\Lambda = 20TeV$		
	Signal	$(\pm)$	Signal	$(\pm)$	Signal	$(\pm)$	Signal	$(\pm)$	
$2000 \rightarrow 2300$	1.25	0.01	1.03	0.01					0.02
$2300 \rightarrow 2800$	1.53	0.01	1.06	0.01	1.02	0.01			0.02
$2800 \rightarrow 3400$	2.20	0.03	1.14	0.02	1.04	0.01	1.02	0.01	0.02
> 3400	3.83	0.07	1.49	0.03	1.14	0.03	1.06	0.02	0.03

Table 7.6: Compositeness signal (constructive interference) $R_\chi^\Lambda/R_\chi^{SM}$ and statistical error (for $10fb^{-1}$, $\epsilon_0 = 5$), only signals with strength above 1% shown. The final column is a combined (positive) systematic error on the signal from PDFs and scale choice.

In the case of constructive interference (Table 7.6), the choice of the R_χ variable allows sensitivity for compositeness signals at $\Lambda = 15TeV$ and below. At higher scales the effect of systematic errors becomes important. For example, the 6% signal at $\Lambda = 20TeV$ and $m > 3400GeV$ has associated values of $\sim 2.6\%$ and $\sim 0.5\%$ for scale and PDF errors

¹²These were added in quadrature for a positive (up) fluctuation.

$m(GeV)$	Compositeness Signal $R_\chi^\Lambda/R_\chi^{SM}$								Systematic Error
	$\Lambda = 5TeV$		$\Lambda = 10TeV$		$\Lambda = 15TeV$		$\Lambda = 20TeV$		
	Signal	$(\pm)$	Signal	$(\pm)$	Signal	$(\pm)$	Signal	$(\pm)$	
2000 $\rightarrow$ 2300	1.15	0.01							0.02
2300 $\rightarrow$ 2800	1.37	0.01							0.02
2800 $\rightarrow$ 3400	2.04	0.02	1.03	0.01					0.02
> 3400	4.05	0.08	1.25	0.03	1.01	0.02			0.03

Table 7.7: Compositeness signal (destructive interference) $R_\chi^\Lambda/R_\chi^{SM}$ and statistical error (for $10fb^{-1}$, $\epsilon_0 = 5$), only signals with strength above 1% shown. The final column is a combined (positive) systematic error on the signal from PDFs and scale choice.

respectively (which were added in quadrature to arrive at the 3% error quoted in the table). This would suggest a discernible signal, however, uncertainty on the magnitude of the systematic errors needs to be considered. For instance, the method of varying the ‘hard’ scale to estimate uncertainty in the fixed-order QCD calculation and the definition of ‘tolerance’ in PDF analyses are to an extent arbitrary. As a result, small signal strengths would reasonably be attributed to uncertainty on the magnitude of systematic errors.

In the case of destructive interference (Table 7.7), the sensitivity is reduced at high compositeness scales Λ in a similar manner as with inclusive jets. In this instance sensitivity is found for $\Lambda = 10TeV$ and below, the tentative signal at $\Lambda = 15TeV$ being insignificant in comparison with systematic uncertainties.

Varying the ϵ_0 parameter

The previous results above all relate to a choice of $\epsilon_0 = 5$, but this parameter can be optimised to maximise the significance¹³ of a compositeness signal. The effects of varying ϵ_0 on the signal strength and its significance for the high-invariant mass bin ($m > 3400GeV$) are shown in Figures 7.21 and 7.22 for constructive interference.

As can be seen, the signal strength increases as ϵ_0 is decreased and the low χ excess is better isolated. However, from Figure 7.22, this increase in signal strength eventually leads to a reduction in significance as the value of R_χ suffers increased statistical uncertainty.

¹³The significance in this instance is defined as the signal strength divided by the expected statistical uncertainty i.e. $(\frac{R_\chi^\Lambda}{R_\chi^{SM}} - 1)/\sigma_{stat}$.




Figure 7.21: Compositeness signal as a function of ϵ_0 , $M_{jj} > 3400 \text{ GeV}$, 10 fb^{-1} . No systematic errors considered (Constructive interference).




Figure 7.22: Significance of compositeness signal as a function of ϵ_0 , $M_{jj} > 3400 \text{ GeV}$, 10 fb^{-1} . No systematic errors considered (Constructive interference).

In order to assess any advantages to varying ϵ_0 , the effects on the systematic errors must also be considered, these are shown in Figure 7.23.

A decrease of ϵ_0 is seen to lead to an increase in the systematic errors of both PDFs and scale choice. The effects are, however, mild and are typically a fraction of a percent. The choice of ϵ_0 in reality will be a compromise between obtaining the maximum possible statistical significance and the production of a signal of sufficient strength to be unambiguous in terms of systematic errors. These criteria favour a choice of ϵ_0 between 2




Figure 7.23: PDF and scale errors on compositeness signal as a function of ϵ_0 , $M_{jj} > 3400\text{GeV}$, 10fb^{-1} .

and 4 (inclusive), however, it should be noted that for high compositeness scales (Λ) the advantages of optimising ϵ_0 become smaller.

In summary the di-jet R_x variable offers a similar sensitivity to compositeness as the inclusive jet cross-section for 10fb^{-1} of ATLAS data. A discernible signal is found for compositeness scales with $\Lambda \leq 10\text{TeV}$ regardless of the sign of interference terms and for $\Lambda \leq 15\text{TeV}$ in the constructive case only. The limitations of each method are however rather different.

In the case of inclusive jets the major background arises from uncertainty in the jet energy scale. Increased integrated luminosity will effectively allow a greater P_T reach of the cross-section but this benefit will probably be undermined by a greater JES uncertainty in this range (c.f. Figure 5.19).

In the case of the di-jet analysis the results are immune to JES uncertainty as only the spatial configuration of jets is considered. Indeed, the angular resolution of jets tends to improve in high Q^2 events as they become increasingly collimated. As a result, higher integrated luminosities are likely to produce improved sensitivity with this method as higher cuts on di-jet mass can be considered. The systematic errors have been shown to increase with increasing di-jet mass, however such increases appear modest. The sensitivity of this method is hence governed by the ability to choose a di-jet mass bin which has

sufficient signal to overcome the theoretical scale errors (which dominate the PDFs). As a result both greater integrated luminosity and an improved theoretical prediction lead to greater sensitivity of this method to compositeness.

Chapter 8

The ATLAS Semiconductor Tracker

8.1 Introduction to The SCT

A brief description of the ATLAS semiconductor tracker was given in Chapter 2. The detector, which is part of the inner detector subsystem is positioned between the pixel detector (which is the detector closest to the interaction point) and the TRT. The SCT is comprised of a barrel region consisting of 4 concentric cylinders and two end-cap regions consisting of 9 disks each (see Table 2.1).

The goal of the inner detector is to measure the path and momentum of charged particles. This is achieved by reconstructing a charged particle's path in the solenoidal magnetic field from discrete space points measured by detectors at different (radial) distances from the interaction point. The specific aim of the SCT is hence to simply measure where a charged particle traverses it to high precision.

8.2 The p-n junction

The basic building block of the SCT's barrel and end-cap is the silicon module. The module is effectively a large number of reverse-biased p-n junctions created at the surface of a silicon wafer. In a semiconductor, the valence and conduction bands are separated by a small amount of energy $\sim 1\text{eV}$. When a charged particle traverses a semiconductor,

the energy deposited creates a large number of electron-hole pairs as electrons are excited from the valence to the conduction band. This charge can then be read out by applying a potential difference across the semiconductor that sweeps out the electrons and holes in opposite directions. In pure silicon however, there exists a large number of free-carriers which allow the electron and hole pairs to recombine before a signal is produced.

In order to circumvent this difficulty a p-type and n-type semiconductor are joined to form a p-n junction. The difference between a 'p' and 'n' type semiconductor is the majority (most abundant) charge carrier, which are holes and electrons respectively¹. When a junction is formed between the two semiconductor types the majority carriers diffuse across the interface and undergo recombination. As both the p and n type semiconductors begin electrically neutral, this migration leads to the p-type semiconductor becoming negatively charged and the n-type becoming positively charged. This continues until sufficient charge has migrated that an electric field is produced which acts to balance the diffusion gradient. The interface hence becomes a 'depletion region' essentially free of charge-carriers, see Figure 8.1.




Figure 8.1: A p-n junction in thermal equilibrium.

By applying a reverse bias to the p-n junction, the depletion region increases in size.

¹These are produced in pure silicon by doping with group III elements (to produce acceptors - p type) and group V (to produce donors - n type)

If an electron-hole is produced in the depletion region (by a traversing ionising particle), the high electric field sweeps the electron and hole in opposite directions (preventing recombination) to the ‘n’ and ‘p’ sides respectively where it can be measured. In a silicon strip detector, a large number of p-n junctions are made along narrow parallel strips. Each strip is read out individually and is in effect an independent detector. The intrinsic resolution of the detector is hence determined by the size or pitch of the strip².

8.3 Module Composition and Operation

A barrel module is comprised of four $63.6\text{mm} \times 64.0\text{mm}$ silicon wafers each containing 768 parallel strips of $80\mu\text{m}$ pitch. The wafers are arranged as a double layer of two-modules placed end-to-end with a wrap-around hybrid which serves to readout the strips (see Figure 8.2). The strips in the first layer are aligned parallel to the beam line in order to provide an accurate determination of the ϕ of a traversing particle (needed to reconstruct a particles momentum). The presence of a second layer helps to improve the accuracy of the ϕ measurement and is offset with respect to the first by a stereo rotation of 40mrad . This angular offset allows a coarse measurement of a particles z co-ordinate to be made, which is needed to assist with the determination of the event’s primary vertex. The modules are also mounted at 10° to the tangent of the ‘barrel’ in order to reduce signal spread (caused by the 2T solenoidal field).

8.4 Reading Out Modules

As a strip on a module forms a p-n junction, charge is collected upon it when a particle traverses the depletion region. The p-type strip which collects charge is separated from the readout chip by an insulating layer. The charge induces a charge in the chip (via the capacitance effect) which is then measured (AC-coupling). This approach is only sensitive to changes in charge rather than its absolute value and helps to reduce noise. The signal

²The method of read-out e.g binary or analog also affects the spatial resolution.




Figure 8.2: An SCT barrel module from [40]

itself is measured by a chip on the module referred to as the ASICs (Application Specific Integrated Circuits) or an ABCD3TA. There are 12 chips per module (6 on either side) each of which reads out a total of 128 strips.

The charge measured on a given strip is a continuous quantity, however, the goal of a silicon strip detector is simply to determine whether a particle has traversed it. The role of chip is therefore to decide given a charge reading on a strip whether it has been hit or not and to pass this information on to the RODs (Read-Out Drivers) if a level 1 trigger signal is received.

A simplified diagram of the structure of an ABCD3TA chip is given in Figure 8.3. The front end of each chip processes the analog signal from each of the strips by performing charge integration, pulse shaping and amplitude discrimination. The discriminator registers a hit if the charge collected passes a threshold which can be set for the chip via a programmable DAC (Digital-Analogue-Converter). A TrimDAC allows an offset to this threshold to be applied from strip to strip. The decision is passed into a pipeline whilst a level-1 trigger signal is awaited. If a trigger signal is received, the values of every strip for the event in question (including the ones occurring directly before and after) are sent to the read-out-buffer before being compressed and sent via optical link to a ROD crate.

Only two chips in a module (one on each side) are attached to read-out links. The




Figure 8.3: A schematic of an ABCD3TA module from [41]

other modules pass their information to the read-out chip on their side in series. In order to maintain fault tolerance, a malfunctioning chip can be bypassed and in the event that a read-out chip becomes disabled the other chips on its side can be read out via the link on the opposite side. The four cylinders of the SCT barrel are comprised of a total of 2112 modules and hence comprises $\sim 3.2 \times 10^6$ channels.

8.5 Calibration of Modules

8.5.1 Setting the Overall Threshold

The calibration of modules is vital to ensure that optimum performance is achieved. As the readout of modules is binary, the ABCD3TA chip has to make a decision whether the charge it reads constitutes a true hit or not. This is complicated in a real detector by the presence of noise. Noise arises from various sources (e.g. thermal excitation and leakage currents)³ and leads to a build up of charge on a strip which can be mistaken for a hit. A threshold is hence introduced above which a strip is considered to have been hit by a traversing particle. As this threshold is increased the purity of the hits is higher (i.e.

³Inter-strip capacitance is the dominant source of noise in silicon strip detectors.

there are less noise hits), however, the efficiency of the detector decreases as real hits may be rejected. This is shown in Figure 8.4. A key reference point as regards calibration is referred to as t_{50} , which is the threshold at which an efficiency of 50% is obtained.




Figure 8.4: A diagram showing the relationship of efficiency (thick black line) with threshold for a silicon detector. The dashed line represents the noise distribution, the pale black line is the 'improved Landau distribution' from [42]

The background noise has a charge probability distribution of Gaussian form with width $\sim 0.15fC$. A minimally ionising particle (or MIP), is characterised with a Landau distribution with a width of $\sim 0.17fC$ and a most probable value of $\sim 3fC$. The overall charge probability distribution is hence a convolution of these two contributions to form the 'improved' Landau distribution as shown in Figure 8.4. The choice of threshold is determined from calibration and made to conform to the SCT specifications of a 99% efficiency of detecting hits and noise occupancy below 5×10^{-4} . The noise occupancy is defined as the average number of noise hits in a given strip per event and has a tendency to increase with prolonged running of the detector due to radiation damage.

Radiation damage causes a change in the effective doping of the bulk n-type silicon to p-type which leads to a need for a greater bias to maintain full depletion. As the decision of whether a hit has taken place is carried out by the ABCD3TA chip, noise occupancy implicitly takes into account amplification and pulse shaping in the front-end processing. A measure of noise with a more 'fundamental' character is provided by the 'equivalent

noise charge' (ENC), which describes a charge which if deposited on the strip would lead to the same output. In effect this value parameterises an approximate relationship between strip threshold and noise occupancy:

$$N \propto \exp\left(-\frac{t^2}{2\sigma^2}\right) \quad (8.1)$$

Where N is the noise occupancy, t is the threshold and σ is the ENC.

8.5.2 The Trim Range Test

Each channel has an associated $100fF$ capacitor which can be used to inject a known charge Q_i on to its input for the purpose of calibration. Repeated tests are made of injecting a known charge whilst the threshold is varied, allowing t_{50} to be determined for a given Q_i . It is expected that for an injected charge of $Q_i = 0$ that the corresponding t_{50} should also be zero. Any difference from zero is referred to as the 'threshold offset'.

In reality, the threshold offset varies between strips. As there is only one overall threshold for a given chip (set by the DAC) these offsets have to be brought into line in order to obtain uniform efficiency across a chip (or module). This is achieved by the TrimDAC which allows the threshold for a given strip to vary in relation to the global chip threshold. There are 16 possible TrimDAC settings for a given strip and 4 possible trim ranges for a chip. The trim ranges effectively set the scale of the TrimDAC units (i.e. the difference between the first TrimDAC setting with respect to the second). The need for a number of ranges is foreseen as radiation damage will lead to the need for greater trimming of strips.

A trim range test consists of injecting $Q_i = 1fC$ into strips and performing a threshold scan for all four of the possible trim ranges. This involves using all 16 possible TrimDAC settings for the first range (to check that all TrimDAC ranges can be set) and 4 in other ranges in order to save time. The mean t_{50} value is then obtained. The linear relationship between the TrimDAC setting and the mean t_{50} value can then be used to determine whether a TrimDAC setting exists in a given trim range for a given strip which allows

it to be set to an arbitrary t_{50} value referred to as a ‘trim target’. If a channel can be brought to the ‘trim target’ it is referred to as ‘trimmable’.

A ‘trim target’ is hence chosen and the list of ‘trimmable’ channels is recorded for each trim range. The TrimDAC and trim range settings are then chosen in such a way to minimise the spread in the resulting mean t_{50} of the channels. This aim is achieved whilst also looking to maximise the number of ‘trimmable’ channels and to use low (i.e. fine) trim ranges⁴.

8.5.3 Three point gain test

The trim range tests brings the threshold offsets into line, however, its absolute value (i.e. the t_{50} values for $Q_i = 0$) and the relationship of t_{50} to injected charge Q_i is not obtained. In order to ascertain these values a test is carried out by performing threshold scans for three injected charges for example: $Q_i = 1.5, 2.0$ and $2.5 fC$. The presence of noise leads to a smearing of the t_{50} value about its mean. The RMS of this error function can be considered as a measure of the output noise of a strip. The t_{50} ’s dependence on Q_i meanwhile provides an estimate for the gain of a channel (See Eqn 8.2).

$$Gain = \frac{dt_{50}}{dQ_i} \quad (8.2)$$

As mentioned earlier the output noise can be considered to be the result of an initial input noise (e.g. on the strip as a result of a leakage current) that is then amplified and adjusted through the front-end chip. The gain, by providing the correspondence between injected charge and measured (charge) voltage, can be used to estimate the input noise. The input noise is determined by simply dividing the output noise by the gain. The extrapolation of the t_{50} vs Q_i graph can also be used to estimate the threshold offset.

⁴Modules which have their TrimDAC and trim range settings determined by this method are referred to as running with ‘trim target settings’

8.5.4 Response Curve Test

The response curve test is, in effect, a lengthened version of the three point gain test with a greater range of injected test charges chosen ($Q_i = 0.5, 0.75, 1, 1.25, 1.5, 2, 3, 4, 6$ and $8fC$). The aim of the test is to better determine the relationship between the measured voltage at the chip with the charge injected onto the strip. Instead of a straight line fit of t_{50} vs Q_i as used in the three point gain test, the data is fit with an empirical function:

$$V = a + \frac{b}{1 + e^{-Q/c}} \quad (8.3)$$

Where a , b and c are fit parameters, V is the measured voltage on the chip and Q is the injected charge. As with the three-point gain test, the threshold offset and input/output noise can be ascertained. The TrimDAC settings can be adjusted to account for the measured threshold offset and if this is carried out the strips are referred to as running with their ‘response curve settings’.

8.5.5 Noise Occupancy Test

A noise occupancy test is a threshold scan performed without any charge injection. The aim of this test is to determine the relationship between the noise occupancy of a strip (or module) and its threshold. This is an important measurement as the SCT requirements demand that the noise occupancy of strips is below 5×10^{-4} . In addition, this method allows the noise component of the overall detected signal to be estimated. However, as mentioned earlier, output noise is also a function of gain (owing to the effect of the front-end processing). The noise occupancy test allows the ENC (Equivalent Noise Charge) to be determined.

8.6 SCT Barrel Tests in the pit - May 2007

8.6.1 Aim

The aim of this study was to identify ‘noisy’ modules from an SCT run carried out in the ATLAS pit from data obtained using the monitoring software and directly from the RODs. This information was then cross-checked with conditions data and earlier calibration tests on modules. Modules can become noisy for a number of reasons and it does not imply a permanent physical fault. For instance a module may appear noisy if the temperature of modules is not kept stable or a detector’s threshold settings are not correctly calibrated.

8.6.2 Sources of Data

The data used to identify noisy modules in this study is from run 7977, which tested the SCT barrel in the pit in ‘physics’ mode. This mode is designed to simulate running with colliding beams, however, the trigger is provided by a signal generator and all ‘hits’ registered in the detector arise from noise. The offline monitoring software was run over the first 1000 events of this raw-data (ByteStream) sample⁵ and compared with ROD data⁶ and the conditions database⁷. The calibration data was taken from tests (as described in Section 8.5) carried out in the pit and found on the web-page display⁸.

8.6.3 Noise Occupancy from Monitoring, RODs and Calibration

Noise occupancy is calculated as the average number of hits in a given strip per event and is automatically determined by the monitoring software. In this study, the noise occupancy is considered as a module average (i.e. over the 1536 strips) and individual modules were considered noisy if this exceeded 5×10^{-4} .

⁵[/castor/cern.ch/user/s/sofia/tmp/daqROSEventBuilder7977](http://castor.cern.ch/user/s/sofia/tmp/daqROSEventBuilder7977)

⁶The ROD level data was obtained in the form of a ROOT ntuple directly from Sergey Burdin

⁷The conditions data used comes directly from the conditions database from a query carried out by Saverio D’Auria.

⁸<http://pcphsctr02.cern.ch/cgi-bin/datadisplay.cgi?locn=SR1&testIndex=1¶meterIndex=2&dataSet=23&barrel=0&chipIndex=0&cellfillindex=0&compIndex=0>

Of the 2112 barrel modules of the SCT, 1452 were present in run 7977 of which 23 were found to be noisy in both the monitoring and ROD level data. The calibration data does not identify any noisy modules, however, in the calibration tests noisy strips are automatically suppressed. A table of the location of the noisy modules and their noise occupancy is given in Table 8.1. For comparison, the average noise occupancy and RMS for well behaved modules are of order $\sim 5 \times 10^{-5}$ as can be seen in Figure 8.5.

From Table 8.1, it can be seen that there exists a correlation between the noise detected from the monitoring and the ROD level data, however, the calibration data show considerably lower noise occupancies. At the time of the tests it was not possible to mask all noisy channels from the ‘offline’ analysis. As a result some modules could appear noisy as a result of a few individual strips that should have been neglected (masked). A change in running conditions such as high voltage bias or temperature might also differ between the calibration and ‘physics’ runs and affect the measured noise occupancies.

The ratio of the monitoring and RODs data for all modules and for noisy modules is given in Figures 8.6 and 8.7. As can be seen there appears to be little bias between the sources in either case with the exception of one noisy module which is an outlier (barrel=3, eta=-3, phi=22). This would suggest that for the main part that there is little distinction between ROD and ‘physics run’ data regarding noisy modules.




Figure 8.5: Distribution of noise occupancies (below 50×10^{-5}) from the monitoring data of the first 1000 events of run 7977 for 1452 barrel modules.

Module Position			Noise Occupancy $\times 10^{-5}$		
barrel	eta	phi	Monitoring	Rods	Calibration
0	-6	28	301.69	297.46	7.9
0	-5	28	186.59	186.44	7.36
0	-4	28	167.58	172.06	9.43
0	-2	28	141.15	143.97	7.34
0	-1	28	229.04	224.55	7.32
0	3	24	7074	7082.8	3.35
0	6	10	314.65	318.38	5.28
1	-4	20	264.58	265.81	7.62
1	-4	39	494.6	498.54	11.5
1	-3	20	367.25	363.54	6.38
1	-3	26	6502.7	6469.6	No data
1	-3	39	114.91	111.81	6.17
1	-2	39	184.9	188.16	7.71
1	2	14	1909.1	1899	11.1
2	-3	33	66.537	66.956	2.21
2	5	21	67.839	67.972	4.61
3	-6	54	67.643	67.585	4.59
3	-4	41	76.302	74.378	11.1
3	-3	22	1141.2	380.81	5.41
3	1	38	99349	99349	No data
3	2	17	51.172	52.272	18.9
3	2	32	70.117	70.176	7.06
3	4	43	70.703	70.881	8.18

Table 8.1: Noisy Modules from run 7977 (SCT Barrel in pit)




Figure 8.6: Ratio of monitoring results to RODs data from run 7977 (all modules with noise occupancy below 50×10^{-5}).




Figure 8.7: Ratio of monitoring results to RODs data from run 7977 (for the modules with noise occupancy $> 50 \times 10^{-5}$). The noisy module at barrel=3, eta=-3 and phi=22 is not plotted and has an exceptionally high ratio $monitoring/rods \sim 3$.

8.6.4 Strip-Level Information

The monitoring software and the ROD data contain information on the modules at the strip level. This is shown in Figures 8.8 and 8.9 and, as can be seen, in some modules the high noise can be attributed to one or a small number of strips. In these cases the agreement between monitoring data and that from RODs is very good with both clearly identifying the problem channel. In other modules, all the strips are noisy and in

Module Position			Channel Error
barrel	eta	phi	
2	-3	33	STUCKON
2	5	21	DEAD
2	-6	54	DEAD
3	-4	41	STUCKON
3	1	38	DEAD
3	2	32	STUCKON
3	4	43	STUCKON

Table 8.2: Noisy modules with strip defects identified in calibration that match noisy channels found in physics run 7977 from [44].

these cases it is possible this results from running conditions (e.g. calibration settings). There appears sometimes to be an offset between monitoring and ROD level data in these cases (e.g. barrel=0, eta=3, phi=24). However, both correctly identify channels as noisy and the noise occupancies involved are much higher than expected for a normal working module. It is also observed in the ROD data that occasionally 12 spikes (in noise) occur evenly spaced across the strips. These spikes, however, are the product of a known problem in decoding the first channel in each chip at the ROD level and are therefore not seen in the monitoring data [43].

The module at barrel=3, eta=2 and phi=17 does not fit well into either category and shows low noise except for the last ~ 100 channels. This suggests a possible problem in one of the readout chips (each chip reads out 128 strips) or a damaged edge of a sensor.

8.6.5 Known problem strips from calibration

The calibration tests as outlined earlier can provide information on defective strips. Eventually, problem strips will be masked in order to suppress their output. To take account of this, a list of known problem strips from an earlier ‘three point gain’ test was cross-referenced with the strip-level hit maps. Seven of the noisy modules can be attributed to isolated defective strips identified in calibration see Table 8.2.

The actual description of the channel error from calibration is, in some cases, inconsistent with the observed problem of a noisy channel. For example the problem (noisy)




Figure 8.8: Strip level noise occupancies for modules with noise above 50×10^{-5} . Black represents the monitoring software results, the red the ROD level information (run 7977)




Figure 8.9: Strip level noise occupancies for modules with noise above 50×10^{-5} . Black represents the monitoring software results, the red the ROD level information (run 7977)

channels in some modules are described as ‘dead’ from calibration and hence would be expected to give no measurement whatsoever. The correlation between the known strip problem and the noisy channel is, however, very strong suggesting that the same underlying problem is the cause. This can be seen in Figure 8.10 as an example.




Figure 8.10: STUCKON strip errors from calibration for module at barrel=2, $\eta = -3$, $\phi = 33$. The problem strip coincides exactly with the ‘noisy’ strip from the ‘physics run’ 7977.

8.6.6 Conditions Data

A record is kept of the running conditions of the SCT in the conditions database. A number of parameters are recorded including the module temperature and the bias of the high-voltage supply, which are factors in module noise. Currently there is no record of luminosity block or run-number within the conditions database and instead the exact time of a run must be known for conditions data to be retrieved. The time of run 7977 analysed in this note was found to be May 13th 18:07 and was established from the time-stamp on the bytestream file as recorded on a DAQ PC at Point 1. However, this time is recorded in CERN local time (Central European Summer Time) which is two hours ahead of UTC (Coordinated Universal Time) used by the conditions database.

The conditions database records parameters every twenty seconds if that parameter has changed more than a predefined amount since the previous reading. For this reason a

search of the database was made for the maximum and minimum values of temperature (Tm0 and Tm1) and high voltage supply (HVchVolt) between the times of 16:05 and 16:09 on May 13th (2007). In all cases the high voltage supply was found to be stable (i.e. unchanged) across this time period. The value of HVchVolt and the maximum and minimum average temperature are given in Table 8.3.

Unfortunately, the recording of conditions data was compromised in the May runs, leading to doubt over the accuracy of the reported values. The problem appeared to stem from a mass operation (e.g switching on/off a large number of modules) which led to the database being inundated and data on changes of state being lost. If a parameter was stable then no data-point is recorded⁹. This meant that the values listed often related to those taken many hours previously. This increased the risk that data loss could lead to a false recorded value of a parameter¹⁰.

As can be seen from Table 8.3, some of the noisy modules show a very low high-voltage bias (nominal working bias is $\sim 150V$, however, the temperature remains fairly stable around $25^\circ C$ ¹¹. Table 8.4 shows the HVchVolt for a list of modules which showed a large proportion of noisy channels. There appears to be a correlation between the modules which show multiple channels being noisy and a low HVchVolt bias although there is an exception in the case of the module at barrel=-1, eta=-3, phi=26.

8.6.7 Summary and Postscript

An SCT barrel noise run in the pit (7977) was investigated by running the offline monitoring software over the bytestream data. The results were compared with data at the ROD level, calibration data and the conditions database. Twenty-three noisy modules (with average noise occupancies greater than 50×10^{-5}) were identified in the monitoring (and ROD) data. No noisy modules were found in the calibration data, however, in this

⁹After upgrades ‘confidence’ values are now taken at least every 30 minutes.

¹⁰All the measurements of HVchVOLT from the conditions database listed are from times before the run was taken.

¹¹The SCT modules could not be cooled to their design running temperature of $-7^\circ C$ due to the risk of condensation.

Module Position			Noise Occupancy $\times 10^{-5}$	Conditions Data		
barrel	eta	phi	Monitoring	HVchVolt	Min Temp	Max Temp
0	-6	28	301.69	150	26.68	26.68
0	-5	28	186.59	150.12	26.29	26.29
0	-4	28	167.58	150.12	26.898	26.898
0	-2	28	141.15	150.12	26.637	26.637
0	-1	28	229.04	150	26.637	26.637
0	3	24	7074	0	25.824	26.077
0	6	10	314.65	10.25	26.126	26.126
1	-4	20	264.58	150.12	26.249	26.249
1	-4	39	494.6	150.12	26.679	26.679
1	-3	20	367.25	150.12	26.204	26.204
1	-3	26	6502.7	150	25.203	25.532
1	-3	39	114.91	150	25.908	25.908
1	-2	39	184.9	150	25.657	25.657
1	2	14	1909.1	5	26.376	26.376
2	-3	33	66.537	50	24.617	24.617
2	5	21	67.839	150	25.534	25.534
3	-6	54	67.643	150.12	26.65	26.65
3	-4	41	76.302	50	24.325	24.325
3	-3	22	1141.2	4.875	26.853	26.853
3	1	38	99349	4.875	21.707	21.707
3	2	17	51.172	50	27.238	27.685
3	2	32	70.117	150.12	25.954	25.954
3	4	43	70.703	150.12	27.165	27.165

Table 8.3: Noisy Modules in SCT with conditions data for temperature and high voltage bias for run 7977 - the data may have been compromised due to the conditions data becoming overloaded

case noisy strips are automatically suppressed. Data loss in the recording of conditions data for this run make it difficult to be certain of the exact running conditions of the modules during the run.

A later ‘physics’ run (42325) carried out as part of the M6 commissioning runs in March 2008 identified 13 modules with a noise occupancy above 50×10^{-5} . Of these 13 modules, only 5 were found to be present in the list of noisy modules of run 7977 (studied here). These modules were also all to be found in the list of modules with known defects (Table 8.2). This suggests that running conditions account for a considerable proportion of ‘noisy’ modules and that most ‘noisy’ modules are not the result of permanent defects.

Module Position			High Voltage bias
barrel	eta	phi	HVchVolt
0	3	24	0
0	6	10	10.25
1	-3	26	150.0
1	2	14	5.0
3	-3	22	4.875
3	1	38	4.875

Table 8.4: A list of noisy modules with all channels noisy and their HVchVolt from the conditions database

8.7 Online Monitoring of the SCT and DCS

The purpose of monitoring is to ensure confidence in the data produced by a detector. In a detector sub-system as complicated as the SCT, constant monitoring is vital to ensure the quality of the data is maintained and to guarantee that its operation remains within safe limits. The monitoring can be separated into the ‘online’ which deals with the real-time running of the ATLAS detector and ‘offline’ which considers recorded events on disk.

The online monitoring samples completed events after the event builder at the SFI (sub-farm input), before the event filter. The sampled events (which represent only a small fraction of the total) are then processed, the results being published on Information Service servers (IS) for external viewing. The SCT online monitoring includes a number of tools which display occupancies, tracking quality and generic DAQ errors (e.g associated with timing mis-calibration or the failure of read out links).

Another important function of the online monitoring is to determine the noise occupancy of modules. This is complicated in normal runs (as opposed to a noise occupancy runs) by the presence of real hits associated with traversing particles which should be neglected. The noise monitoring tool first receives the collection of hits associated with tracks in a given event. These hits are referred to as ‘reconstruction input objects’ (RIOs) and are stored in a 1D histogram whose bins represent the strips for a given module. When an event is read out from the SCT, it contains complete information on the given bunch crossing and in addition, those which occur immediately before and after (three time bins). The noise tool considers hits occurring in all three of these time bins, demanding that the central time bin has a hit (X1X run mode) to create an analogous histogram to that of the RIOs. The noise occupancy is then obtained by simply subtracting the two histograms and dividing by the number of events. In this manner true hits are prevented from being considered as noise.

The ‘Detector Control System’ monitoring or DCS monitoring is also vital for assuring the reliability of data from the experiment, but has a slightly different focus. DCS monitoring does not examine the event data that passes through the DAQ but rather monitors

the status and environment of the detector as a whole. This includes for instance the monitoring of temperatures, bias voltages and currents etc. As such monitoring is desirable whenever the detector is active, regardless of data taking; the system can be run independently of the DAQ. The scope of the DCS itself goes beyond passive monitoring and forms the basis of control of all hardware elements that comprise the detector, including its configuration. The DCS conditions data is needed in both online and offline worlds. The need for online access to DCS information is self-evident and is obtained through an ‘Information Service’ (IS) server to which the data is published and may be retrieved in real time. Offline access to the DCS information is also important to ensure the reliability of data (as was carried out earlier see Section 8.6) and is obtained through the conditions database.

8.7.1 Online DCS monitoring tool

The DCS and SCT online monitoring are by design independent systems. However, it is advantageous, particularly in the event of a module exhibiting high levels of noise, to be able to assess key DCS information on that module. To allow this to take place with a common interface, a dedicated DCS monitoring tool (written by the author) was added to the SCT online monitoring software. The tool does not replace the role of the DCS as regards modules, as the DCS must be able to run independently of the DAQ. However, the tool does provide the online monitoring team with more information in a consistent style with which to infer the likely cause of errant modules.

The DCS monitoring tool is designed to work within the framework of the SCT online monitoring and as such is hence executed in the same manner as the noise tool, i.e. after a sampled event is directed to it from the SFI. Unlike the other monitoring tools, the DCS tool does not require any event data to function and instead reads conditions data directly from an IS server when triggered.

The tool then interprets the IS data and provided a mapping between a module’s location and the crate-channel numbers given by the DCS in order to publish histograms

alongside those containing event data from the other monitoring tools.

8.7.2 Testing in the M5 commissioning runs

In order to test the DCS monitoring and the wider online SCT monitoring package it was included in the M5 commissioning runs carried out in October 2007. These runs were intended to test the integration of various systems in the DAQ including monitoring in a realistic setting to assess the compatibility of sub-systems and the demands on online computing resources.

As it was not possible for the M5 runs to use the SCT with the DAQ (due to an absence of module cooling), a test box¹² was set up with four working modules in order to provide a pseudo-realistic interface and source of data for the purposes of the test. An example histogram produced by the DCS monitoring tool during the M5 runs is shown in Figure 8.11, the active modules are visible at $\phi=21$ and $\eta=-6$ to -3 .




Figure 8.11: HVchVolt for the four modules of the H8 test box in the M5 commissioning runs from the DCS monitoring tool. The x-axis relates to ‘eta’ and the y-axis to ‘phi’. The plot is taken from a ‘screen-shot’ of a histogram from the OH (online histogram display).

The HVchVolt values from the DCS online monitoring tool (Figure 8.11) agreed with those obtained directly from the DCS. In addition tests showed that the regular mass-reading of information from the IS server does not put undue stress in terms of CPU and

¹²This test box is referred to as the ‘H8’ test box (named after the test beam).

memory on the online computing system.

A more thorough test of the DCS monitoring tool was carried out in the M6 commissioning runs in March 2008. These allowed for measurements to be taken from the entire SCT barrel (which was cooled) and hence a more realistic test of the tool's performance. Figure 8.12 shows a map of the HV bias for the inner layer of the SCT modules. Additional functionality was added for the M6 runs which allowed the state of the high voltage and low voltage supplies to be displayed and this is shown in Figure 8.13. The state is coded as given in Table 8.5.

State	Code (hex)
OFF	0
ON	1
STB	2
MANUAL	3
MASK_OFF	4
MASK_ON	5
HARD_RESET	6
DISABLED	7
RAMPING	8
INTERLOCKED	9
TRIP_HW	A
TRIP_SW	B
LVCARD_LATCH	C
NO_MATCH	D
UNKNOWN	E
ANY	F

Table 8.5: Table of HV and LV state codes.

The time taken to access the complete IS server information for the SCT was found to be 1-2 seconds. In order to prevent the server from being inundated with access requests, a safeguard was added to the DCS monitoring tool which prevents a request being made at intervals of less than 20 seconds.




Figure 8.12: HVchVolt for layer 0 of the SCT Barrel in the M6 commissioning runs from the DCS monitoring tool. The x-axis relates to ‘eta’ and the y-axis to ‘phi’. The plot is taken from a ‘screen-shot’ of a histogram from the OH (online histogram display).




Figure 8.13: HVState for layer 0 of the SCT Barrel in the M6 commissioning runs from the DCS monitoring tool. The x-axis relates to ‘eta’ and the y-axis to ‘phi’. The plot is taken from a ‘screen-shot’ of a histogram from the OH (online histogram display).

Chapter 9

Tracking Efficiency in Minimum-Bias Events

The goal of the inner detector is to measure the trajectory of charged particles emerging from the interaction point (IP). However, the individual detector sub-systems measure discrete space-points where the particle traverses a layer, as opposed to a continuous track. As a consequence, a tracking algorithm is needed to reconstruct tracks from the discrete points. This is a non-trivial problem due to the wide range in transverse-momentum (P_T) of particles, the strong solenoidal magnetic field and the high occupancy of detectors, particularly near the IP.

The tracking (or pattern recognition) algorithms typically begin by defining ‘roads’ or track candidates. In an attempt to reduce ambiguities, the seeds that define the initial track candidates are often taken from a detector at a radius from the IP (e.g. the TRT) where the track density is lower (the histogramming method). The role of the road is to define a subset of hits from the detectors which can be considered when fitting a helix to define a track.

The information on the hits from inner layers is included by a process of pattern recognition which creates a geometric fit to a track trajectory. These algorithms are often rather sophisticated, for example the Global Chi-Squared approach¹ and the Kalman

¹A parameterised track hypothesis is tested and those with the minimum χ^2 values are accepted.

filter smoother formalism which allows the track candidate to be refined during the fitting process². In order to resolve ambiguities, pattern recognition algorithms decide whether a hit belongs to a track candidate, as opposed to whether a particular track should pass through the location of a given hit. This strategy is motivated experimentally as it is easier to discern a good quality track than a hit.

Owing in part to their complexity, characterising the behaviour of tracking algorithms is important in order to correctly interpret experimental results. The track reconstruction efficiency, the acceptance of the quality cuts (used to define tracks) and fake rates all need to be measured. The performance of a tracking algorithm, however, is hard to define without recourse to Monte Carlo simulation. In a simulated event the flavour of primary charged particles and their kinematics are known from ‘truth information’. Reconstructed tracks can thus be matched to truth primaries to allow a determination of acceptance, fake rates and efficiencies.

It is desirable, however, to assess the tracking efficiency directly from data. One such approach is to use track insertion. The premise is to take an event (either from simulation or experimental data) and to insert a single simulated track (from Monte Carlo) of known type and kinematics. The merged event is then reconstructed as normal and compared to the reconstructed unmerged event. The presence (or absence) of additional tracks in the merged event (with respect to the unmerged) can then be used to estimate the reconstruction efficiency. The presence of ‘spectator’ hits in the merged event allow an appraisal of the track reconstruction in a realistic environment, where ambiguities and high track density could reduce efficiency.

²The filter is run from the outermost layers of the ID inward taking into account multiple-interactions and bremsstrahlung radiation. In this process the track candidate is continually updated. After this a smoothing procedure runs from the inner layers out which uses the complete information obtained from the filtering.

9.1 Criteria for Successful Track Reconstruction

In determining whether an inserted track has been reconstructed, a broad range of criteria can be applied. These may include the requirement that the trajectory and kinematics of the reconstructed track agree with the generation level values within a predefined error margin. If the matching criteria are too tight however, the resolution of the detector simulation may begin to affect the measured efficiency.

For this reason, an approach was chosen to apply loose cuts to determine whether an inserted track had been successfully reconstructed and then to demonstrate that the use of loose cuts did not introduce a bias between generated and reconstructed values of P_T , η and ϕ .

A cone of radius $R = 0.7$ in $\eta - \phi$ space was defined; centred on the trajectory of the track in the generated event. This cone was then projected on to the original minimum-bias event (without the inserted track) and the number of reconstructed tracks contained within the cone counted. The same technique was carried out for the merged event. The inserted track was said to have been successfully reconstructed if the merged event contained one or more tracks (in the cone) in addition to those of the original event.

This method has the advantage that it is unnecessary to explicitly determine the parameters of the overlaid track in the merged event to perform the efficiency calculation. However, in order to ascertain whether the cuts applied produce systematic biases (in P_T , η and ϕ) the reconstructed inserted track has to be identified in order to directly compare it to the generated level. This was achieved by selecting the reconstructed track in the merged event closest (in terms of $\eta - \phi$) to the known generated track trajectory.

9.2 Minimum-Bias Events

One area in which a knowledge of track reconstruction efficiency is important is the study of minimum-bias events. A minimum-bias event does not directly refer to a specific physics channel, but rather to a class of events expected in the absence of selection

by an experimental trigger. By construction, the observed processes are those which dominate the proton-proton cross-section (elastic, diffractive and non-diffractive inelastic interactions)³. The study of minimum-bias events provides an insight and constraint on models of low Q^2 QCD interactions such as hadronization, which are described by complex Monte Carlo generators. Minimum-bias studies are also important for detector commissioning and the estimation of backgrounds to hard processes (referred to as the underlying event).

Minimum bias events tend not to contain high P_T leptons or jets and instead studies tend to concern themselves with inclusive charged particle distributions (as a function of P_T and η) and distributions such as the KNO scaling variable $n/\langle n \rangle$ (where ‘ n ’ is the particle multiplicity in an event). The track reconstruction efficiency is hence an important consideration when interpreting minimum-bias data.

A set of 28500 non-diffractive inelastic scattering (minimum-bias) events were generated by PYTHIA and simulated using the ATLAS software (version 12.0.31). This simulated data set reflects which parts of the detector registered hits during the event (i.e. before reconstruction) and are often referred to as ‘DIGI’ files. The simulated data was then reconstructed using the ATLAS software (version 12.0.5) to produce the ‘unmerged’ sample.

An equal number of simulated single-track muon events were then produced using the single particle gun tool⁴. The kinematics of the overlaid track were chosen by sampling global distributions of track P_T , η and ϕ obtained from reconstruction of the minimum-bias events. In addition the Z-coordinate (longitudinal) of the track origin was set to match that of the reconstructed vertex of a corresponding event in the minimum-bias sample. The simulated (DIGI) files of the corresponding minimum-bias and single track events were then overlaid using the InDetOverlay tool⁵ to form the merged event which was in turn reconstructed. The single-track simulation files were also reconstructed in

³The definitions of minimum bias events vary in the literature and are sometimes associated with purely non-diffractive inelastic processes.

⁴Muons were chosen as a benchmark as they are stable (over the distance scales of the inner detector) and undergo less pair-production or bremsstrahlung radiation than electrons.

⁵Written by Andrei Gaponenko

isolation (i.e. without merging).

9.2.1 Effect Of Selection Cuts

Before an estimate of track reconstruction efficiency is made it is necessary to examine whether the selection cuts create a bias in the kinematics of reconstructed merged tracks. This is carried out by comparing the nearest track (in $\eta - \phi$) in the merged event with the known generated track trajectory for events that pass the selection cuts. The difference in η , ϕ and P_T between the generated and merged track are shown in Figures 9.1, 9.2 and 9.3. At the time of the study the track reconstruction algorithm (in 12.0.5) was not robust for tracks with $P_T \leq 0.5 \text{ GeV}$. As a result only tracks with $P_T > 0.5 \text{ GeV}$ are considered.




Figure 9.1: Pseudorapidity of closest track to generated track trajectory in merged event minus pseudorapidity of generated track (For tracks with $P_T > 0.5 \text{ GeV}/c$). The mean of the distribution is zero (within errors).

In all cases it can be shown that the loose cuts applied to determining whether an inserted track has been found, do not lead to a systematic bias between the kinematics of the generated and reconstructed track. The plots also show that the cone radius of 0.7 (in $\eta - \phi$ space) used to compare tracks in the original and merged events is much greater than the resolution of the detector to tracks.

Another potential source of bias could arise from the manner in which tracks are inserted into the original event. The model as outlined so far assumes there is no spatial




Figure 9.2: ϕ of closest track to generated track trajectory in merged event minus ϕ of generated track (For selected events, $P_T > 0.5 \text{ GeV}/c$). The mean of the distribution is zero (within errors).




Figure 9.3: P_T of closest track to generated track trajectory in merged event minus P_T of generated track (For selected events, $P_T > 0.5 \text{ GeV}/c$). The mean of the distribution is zero (within errors).

correlation between tracks in a minimum bias event and that a track may be inserted by sampling global distributions (of P_T , η and ϕ). If, however, tracks in a given minimum-bias event tend to cluster, then inserting a track according to global distributions may not provide a realistic test of reconstruction efficiency.

This concern was addressed by looking at distributions of all possible combinations of ΔR (defined as $\Delta R = \sqrt{\Delta\phi^2 + \Delta\eta^2}$) between tracks in a standard minimum bias event

and a toy MC event generated from global distributions (of track P_T , η and ϕ). The toy MC event can be considered as representing the non-correlated tracks used in the insertion process. The results are shown for events containing 5, 10 and 15 tracks in Figures 9.4, 9.5 and 9.6.




Figure 9.4: Distribution of ΔR of all tracks in event, for events with 5 tracks. Distributions are from PYTHIA (used to create ATLAS simulated data) and Toy MC (equivalent to track insertion technique). For 1697 events




Figure 9.5: Distribution of ΔR of all tracks in event, for events with 10 tracks. Distributions are from PYTHIA (used to create ATLAS simulated data) and Toy MC (equivalent to track insertion technique). For 926 events

The figures show that there is a mild-clustering effect in minimum-bias events that is more pronounced as the number of tracks in an event decreases. This effect has to




Figure 9.6: Distribution of ΔR of all tracks in event, for events with 15 tracks. Distributions are from PYTHIA (used to create ATLAS simulated data) and Toy MC (equivalent to track insertion technique). For 545 events

be considered when assessing track-reconstruction efficiency by the insertion method as pattern recognition is typically more difficult when tracks are closer to each other than well separated.

9.2.2 Reconstruction Efficiency

The track reconstruction efficiency was assessed using the criteria as outlined earlier. The results as a function of inserted track η , P_T and the number of tracks in the unmerged minimum bias event are shown in Figures 9.7, 9.8 and 9.9.

As can be seen, within errors, the effect of track η , the number of tracks in an event and the track's P_T (for $P_T > 0.5 \text{ GeV}/c$) on reconstruction efficiency is small. The efficiency values are close to those found previously ($\sim 97\%$ ⁶) by a simple comparison of the properties of generated and reconstructed muons [7]. The values can vary between published results, however, due to differing cuts used in truth-matching.

The track reconstruction efficiency determined by track insertion can be thought of as being comprised of two parts. The first is the efficiency of the single generated track being reconstructed in isolation of any other tracks. The second component (given the first) is

⁶This result was for muons with $P_T = 1 \text{ GeV}$.




Figure 9.7: Track reconstruction efficiency as determined by insertion into a minimum-bias (merged) and in isolation (isolated) as a function of η in event for min-bias events ($P_T > 0.5 \text{ GeV}/c$). The three bins relate to the barrel, transition and end-cap regions of the inner detector.




Figure 9.8: Track reconstruction efficiency as determined by insertion into a min-bias (merged) and in isolation (isolated) as a function of track P_T for minimum-bias events in three regions of pseudorapidity ($P_T > 0.5 \text{ GeV}/c$).

whether this track can also be reconstructed in the track environment of a minimum-bias event. The separation of these two elements is seen in Figures 9.7 and 9.8 where the efficiency is given when the generated track is merged with a minimum-bias event and when it is reconstructed in isolation. As can be seen, the efficiency is always slightly higher when the track is reconstructed in isolation (in comparison to when merged), but




Figure 9.9: Track reconstruction efficiency as a function of the number of tracks in the original minimum-bias event ($P_T > 0.5 \text{ GeV}/c$).

the effect is mild ($1 - 2\%$). This difference can be thought of as the loss in efficiency due to the difficulty of pattern recognition in the multi-track environment of minimum-bias events. As this loss is small it suggests that in minimum-bias events the dominant factor determining efficiency is the fundamental track reconstruction as opposed to the nature of the track environment.

As was mentioned earlier there is a mild-clustering effect in minimum-bias events with low track multiplicity. The effect this has on reconstruction efficiency can be determined by plotting track reconstruction efficiency as a function of the distance between the inserted track trajectory and the nearest track (in terms of ΔR in the original event. This is shown in Figure 9.10. As can be seen the track reconstruction efficiency seems to be insensitive to the distance of neighbouring tracks and as a result it can be assumed that the failure to take into account the mild clustering effects in minimum bias events should not affect the calculated efficiency.




Figure 9.10: Track reconstruction efficiency as a function of ΔR from inserted track trajectory to nearest track in the original minimum-bias event ($P_T > 0.5 \text{ GeV}/c$).

Chapter 10

Conclusions

A method of using integration grids to allow generation of NLO jet cross-sections independent of PDFs was implemented. After optimisation the grids are found to reconstruct the standard Monte Carlo result to within 1% accuracy for events with a broad range of Q^2 . This accuracy allows jet data from colliders to be included into global PDF fits in a manner consistent with other sources (e.g. from deep-inelastic scattering). Early studies suggest that ATLAS data will be able to constrain the high x gluon PDF if the jet energy scale (JES) can be constrained to 1%.

The effect of uncertainties within PDFs and fixed order calculations on theoretical predictions of the inclusive jet cross-section was investigated. For a $P_T \sim 1$ TeV centrally produced ($0 < |\eta| < 1$) jet, there is an uncertainty of approximately $\sim 12\%$ and $\sim 7\%$ arising from PDFs and hard-scale choice respectively. Corresponding experimental errors from jet energy resolution and scale (JES of 5%) were also considered and lead to uncertainties of $< 1\%$ and $\sim 30\%$ respectively. Uncertainty in the jet energy scale dominates the experimental error and is difficult to constrain at high P_T , owing to an absence of processes (with sufficient cross-section) that allow the hadronic scale to be connected to the better known electromagnetic scale. An approach to cross-check the JES, by ‘bootstrapping’ the calibration for low P_T jets using multi-jet events, was tested. The technique suggests that a non-linearity in energy scale can be detected at $\sim 1.5\%/500\text{GeV}$ but a more precise constraint of $1\%/500\text{GeV}$ is likely to be challenging.

The sensitivity of ATLAS jet data to potential compositeness signals (left-left isoscalar model) was also studied. The inclusive jet cross-section was found to be sensitive to scales of $\Lambda = 10\text{TeV}$ for JES uncertainties as high as 10% regardless of whether the compositeness effects are constructive or destructive for 10fb^{-1} of data. The sensitivity to compositeness at the higher scale of $\Lambda = 15\text{TeV}$ is limited to constructive interference and the requirement that the JES can be constrained to 1%. A similar sensitivity is found from dijet-angular methods, although in this case there is no constraint from JES uncertainties and the signal is limited by statistics (for 10fb^{-1} of data).

The noise occupancy of silicon modules that form the SCT detector at ATLAS were studied from a run in ‘physics’ mode in the pit (May 2007). The data suggested 23 modules were ‘noisy’, 7 of which could be attributed to known problems with single-channels which would be suppressed in later runs. A later run taken in March 2008 found 13 noisy modules of which only 5 were in common with those of the earlier run. These common modules are all associated with known problems. This suggests that most ‘noisy’ modules are the result of running conditions (e.g. calibration settings) and not the result of permanent defects.

The track reconstruction efficiency in minimum-bias events was investigated by track insertion. The results suggest that the efficiency is high ($\sim 95\%$), independent of P_T above 0.5GeV and is not greatly affected by ambiguities in pattern recognition associated with the track environment in minimum-bias events.

Whilst this thesis investigates the effects of experimental errors on jet measurements as regards cross-sections, a prediction of their magnitude will naturally have to be reviewed in the light of experimental data. The commissioning process will allow a better understanding of the detector and how this will affect the accuracy of cross-sections and sensitivity to new physics. The LHC will offer higher centre of mass energies than any previous accelerator and allow for a thorough exploration of the electroweak sector in addition to potential new physics processes such as supersymmetry. As unexpected phenomena may be encountered the need for a reliable and well understood detector, in

addition to theoretical uncertainties, is vital.

List of Figures

2.1	Accelerators at CERN	9
2.2	The ATLAS detector [3]	11
2.3	The ATLAS inner detector, incorporating the pixel detector, SCT and TRT[3] .	13
2.4	The ATLAS EM and hadronic calorimetry [3]	16
2.5	The ATLAS Muon-Chambers [7]	19
2.6	The multi-level trigger [7]	20
3.1	Confinement and particle production due to QCD. The curly lines represent gluons.	26
3.2	Dominant QCD amplitudes for jet production at LO (t-channel), curly lines represent gluons, plain lines are quarks or antiquarks.	29
3.3	Some QCD amplitudes for Jet production at NLO, curly lines represent gluons, plain lines are quarks or antiquarks.	31
3.4	Summary of $\alpha_s(M_Z)$ from Particle Data Group (2006) [19]. The errors shown include theoretical uncertainties and an average gives $\alpha_s(M_Z) = 0.1176 \pm 0.0009$.	35
3.5	PDFs for the up, down, anti-up quarks and the gluon from CTEQ6.1. $Q_0 =$ 10GeV and the gluon distribution has been scaled by 0.1 for ease of comparison [20]	36
3.6	Inclusive jet cross-section for a proton-proton collider with $\sqrt{s} = 14TeV$ (LHC),CTEQ6.1 PDFs for three regions in pseudorapidity, generated by NLOJET ⁺⁺	37
4.1	A PDF described by interpolation on a grid.	45

4.2	A one-dimensional grid with $N_y = 4$, $\delta y = 1$ and $y_{min} = 3$ to describe the function $f(y) = 2y$	46
4.3	Combining a weight-grid with a PDF in a fast step.	49
4.4	Example weight-grid (NLOJET ⁺⁺ , inclusive jets), as a function of y_1 and y_2 projected over $2.90 < \tau < 2.96$	53
4.5	Co-ordinate transformation of the x-variable (Eqn 4.2).	54
4.6	Co-ordinate transformation of the Q^2 -variable, $\Lambda = 0.25$ (Eqn 4.3).	54
4.7	Variable co-ordinate transformation of the x -variable Eqn 4.18.	55
4.8	Sum of the up, down and gluon PDFs at $Q=5\text{GeV}$, 50GeV , 500GeV and 5000GeV for CTEQ6.1 best fit PDFs.	56
4.9	Gradient of the sum of up, down and gluon PDFs at $Q=5\text{GeV}$, 50GeV , 500GeV and 5000GeV for CTEQ6.1 best fit PDFs.	57
4.10	Parameter a_1 found by fitting Eqn 4.19 to the sum of the up, down and gluon PDFs for the 40 error PDFs of CTEQ6.1 at $Q=5\text{GeV}$, 50GeV , 500GeV and 5000GeV	58
4.11	Parameter a_2 found by fitting Eqn 4.19 to the sum of the up, down and gluon PDFs for the 40 error PDFs of CTEQ6.1 at $Q=5\text{GeV}$, 50GeV , 500GeV and 5000GeV	58
4.12	Parameter a_1 found by fitting Eqn 4.19 to the sum of the up, down and gluon PDFs for 31 best fit PDFs (See Table 4.1) at $Q=5\text{GeV}$, 50GeV , 500GeV and 5000GeV	59
4.13	Parameter a_2 found by fitting Eqn 4.19 to the sum of the up, down and gluon PDFs for 31 best fit PDFs (See Table 4.1) at $Q=5\text{GeV}$, 50GeV , 500GeV and 5000GeV	59
4.14	Weighted PDF(CTEQ6.1) (black) for $a_1=-1.5$, $a_2=3$ and unweighted PDFs (red) at $Q=5\text{GeV}$, 50GeV , 500GeV and 5000GeV . Function $g(x) = x$ in the ‘unweighted’ case and $g(x) = 1/\text{weight}'$ in the ‘weighted’ case, where weight is defined by equation 4.19.	60

4.15	Weighted PDF grid (CTEQ6.1), $a_1=-1.5$ and $a_2=3$ for $F^{(0)}$ as a function of y_1, y_2 , projected over $2.90 < \tau < 2.96$	61
4.16	Weighted PDF grid (CTEQ6.1), $a_1=-1.5$ and $a_2=3$ for $F^{(0)}$ as a function of y_1 and τ , projected over $0 < y_2 < 6$	61
4.17	Comparison of grid/standard for the inclusive jet-cross section at ATLAS ($0 < \eta < 1$). A standard cross-section is calculated using NLOJET ⁺⁺ without an integration grid. (Weighted grid, $30 \times 30 \times 10$ bins in y_1, y_2, τ , and co-ordinate transform parameter $a=5$).	63
4.18	Comparison of grid/standard for the inclusive jet-cross section at ATLAS ($1 < \eta < 2$). A standard cross-section is calculated using NLOJET ⁺⁺ without an integration grid. (Weighted grid, $30 \times 30 \times 10$ bins in y_1, y_2, τ , and co-ordinate transform parameter $a=5$).	63
4.19	Comparison of grid/standard for the inclusive jet-cross section at ATLAS ($2 < \eta < 3$). A standard cross-section is calculated using NLOJET ⁺⁺ without an integration grid. (Weighted grid, $30 \times 30 \times 10$ bins in y_1, y_2, τ , and co-ordinate transform parameter $a=5$).	64
4.20	Comparison of grid/standard for the inclusive jet cross section at ATLAS for different grid sizes ($0 < \eta < 1$). (Weighted grid, co-ordinate transform parameter $a=5$). A standard cross-section is calculated using NLOJET ⁺⁺ without an integration grid.	65
4.21	Comparison of grid/standard for the inclusive jet cross section at ATLAS for different grid sizes ($1 < \eta < 2$). (Weighted grid, co-ordinate transform parameter $a=5$). A standard cross-section is calculated using NLOJET ⁺⁺ without an integration grid.	65
4.22	Comparison of grid/standard for the inclusive jet cross section at ATLAS for different grid sizes ($2 < \eta < 3$). (Weighted grid, co-ordinate transform parameter $a=5$). A standard cross-section is calculated using NLOJET ⁺⁺ without an integration grid.	66

4.23	Comparison of grid/standard for the inclusive jet-cross section at ATLAS for different grid sizes ($2 < \eta < 3$). (Weighted grid, no. y bins=30, co-ordinate transform parameter $a=5$). A standard cross-section is calculated using NLOJET ⁺⁺ without an integration grid.	66
4.24	Comparison of grid/standard for the inclusive jet-cross section at ATLAS for different interpolation orders ($(0 < \eta < 1)$). Weighted grid ($30 \times 30 \times 10$), co-ordinate transform parameter $a=5$). A standard cross-section is calculated using NLOJET ⁺⁺ without an integration grid.	67
4.25	Comparison of grid/standard for the inclusive jet-cross section at ATLAS for weighted and unweighted ($10 \times 10 \times 10$) grids. ($(0 < \eta < 1)$, co-ordinate transform parameter $a=5$). A standard cross-section is calculated using NLOJET ⁺⁺ without an integration grid.	68
4.26	Comparison of grid/standard for the inclusive jet-cross section at ATLAS for different grid co-ordinate transforms. ($(2 < \eta < 3)$, $10 \times 10 \times 10$ weighted grids). A standard cross-section is calculated using NLOJET ⁺⁺ without an integration grid.	69
4.27	Comparison of grid/standard for the inclusive jet-cross section at ATLAS for different grid renormalisation/factorisation scales. ($(0 < \eta < 1)$, $30 \times 30 \times 10$ weighted grids). A standard cross-section is calculated using NLOJET ⁺⁺ without an integration grid.	70
4.28	Gluon PDF uncertainty from global fit with inclusion of ATLAS pseudo-data, without JES uncertainty. Results for $10 fb^{-1}$ with a 5% uncorrelated systematic error. ZEUS-O considers data from neutral/charged current interactions from ZEUS whilst ZEUS-JETS includes data from ZEUS jet cross-sections. ATLAS-JETS in addition includes the ATLAS inclusive jet cross-section. (plot by Claire Gwenlan (Oxford/UCL)).	71

4.29	Gluon PDF uncertainty from global fit with inclusion of ATLAS pseudo-data, with 1% JES uncertainty. Results for $10fb^{-1}$ with a 5% uncorrelated systematic error. ZEUS-O considers data from neutral/charged current interactions from ZEUS whilst ZEUS-JETS includes data from ZEUS jet cross-sections. ATLAS-JETS in addition includes the ATLAS inclusive jet cross-section. (plot by Claire Gwenlan (Oxford/UCL)).	72
4.30	Gluon PDF uncertainty from global fit with inclusion of ATLAS pseudo-data, with 3% JES uncertainty. Results for $10fb^{-1}$ with a 5% uncorrelated systematic error. ZEUS-O considers data from neutral/charged current interactions from ZEUS whilst ZEUS-JETS includes data from ZEUS jet cross-sections. ATLAS-JETS in addition includes the ATLAS inclusive jet cross-section. (plot by Claire Gwenlan (Oxford/UCL)).	72
5.1	Inclusive Jet Cross-Section for the LHC and Tevatron (Run II), (NLOJET, CTEQ6.1, $\mu_r, \mu_f = P_T(max)$). P_T is effectively a measure of Q^2	74
5.2	K-Factors for inclusive jets in three regions of pseudorapidity, (NLOJET, CTEQ6.1, $\mu_r, \mu_f = P_T(max)$)	76
5.3	Scale error for inclusive jets within $0 < \eta < 1$. The scale was varied from $0.5P_T^{max}$ to $2.0P_T^{max}$. The max-min series relate to the largest deviations of the cross-section within this range, whereas std-fix considers the values of the cross-section at the extreme choices of scale.(NLOJET, CTEQ6.1) .	78
5.4	Scale error for inclusive jets within $1 < \eta < 2$. The scale was varied from $0.5P_T^{max}$ to $2.0P_T^{max}$. The max-min series relate to the largest deviations of the cross-section within this range, whereas std-fix considers the values of the cross-section at the extreme choices of scale. (NLOJET, CTEQ6.1) .	79
5.5	Scale error for inclusive jets within $2 < \eta < 3$. The scale was varied from $0.5P_T^{max}$ to $2.0P_T^{max}$. The max-min series relate to the largest deviations of the cross-section within this range, whereas std-fix considers the values of the cross-section at the extreme choices of scale. (NLOJET, CTEQ6.1) .	79

5.6	Variation of the cross-section with scale for a P_T bin, $100 < P_T < 150\text{GeV}$, $2 < \eta < 3$.(NLOJET,CTEQ6.1)	80
5.7	Variation of cross-section with scale for a P_T bin $600 < P_T < 700\text{GeV}$, $2 < \eta < 3$.(NLOJET,CTEQ6.1)	80
5.8	Uncertainty on inclusive jet cross-section from the 40 CTEQ6.1 error sets, $0 < \eta < 1$ (NLOJET, CTEQ6.1).	83
5.9	Uncertainty on inclusive jet cross-section from the 40 CTEQ6.1 error sets, $1 < \eta < 2$ (NLOJET, CTEQ6.1).	83
5.10	Uncertainty on inclusive jet cross-section from the 40 CTEQ6.1 error sets, $2 < \eta < 3$ (NLOJET, CTEQ6.1).	84
5.11	Comparison of the leading jet in ‘Truth’ and reconstructed inclusive jets in ATLAS (Athena version 11.0.4) for two ranges in P_T . Cone R=0.7, jets from calorimeter towers.	91
5.12	Comparison of the Cone R=0.4 and KT (tower) algorithms to Cone R=0.7 in ATLAS (Athena version 11.0.4).	91
5.13	Statistical Errors for inclusive jets at LHC for $0 < \eta < 1$ for three choices of integrated luminosity, (NLOJET, CTEQ6.1 from integration grids, $\mu_r, \mu_f = P_T(max)$)	93
5.14	Statistical Errors for inclusive jets at LHC for $1 < \eta < 2$ for three choices of integrated luminosity, (NLOJET, CTEQ6.1 from integration grids, $\mu_r, \mu_f = P_T(max)$)	93
5.15	Statistical Errors for inclusive jets at LHC for $2 < \eta < 3$ for three choices of integrated luminosity, (NLOJET, CTEQ6.1 from integration grids, $\mu_r, \mu_f = P_T(max)$)	94
5.16	Jet Energy Resolution Errors for inclusive jets at LHC for $0 < \eta < 1$ for a shift of 10% and 20% of the sampling and constant term parameters (a, b), (NLOJET,CTEQ6.1 at Born level)	95

5.17	Jet Energy Resolution Errors for inclusive jets at LHC for $1 < \eta < 2$ for a shift of 10% and 20% of the sampling and constant term parameters (a, b), (NLOJET, CTEQ6.1 at Born level)	96
5.18	Jet Energy Resolution Errors for inclusive jets at LHC for $2 < \eta < 3$ for a shift of 10% and 20% of the sampling and constant term parameters (a, b), (NLOJET, CTEQ6.1 at Born level)	96
5.19	Jet energy scale Errors for inclusive jets at LHC for $0 < \eta < 1$ at 1%, 5% and 10%, (NLOJET, CTEQ6.1), $\mu_r, \mu_f = P_T(max)$)	97
5.20	Jet energy scale Errors for inclusive jets at LHC for $1 < \eta < 2$ at 1%, 5% and 10%, (NLOJET, CTEQ6.1), $\mu_r, \mu_f = P_T(max)$)	98
5.21	Jet energy scale Errors for inclusive jets at LHC for $2 < \eta < 3$ at 1%, 5% and 10%, (NLOJET, CTEQ6.1), $\mu_r, \mu_f = P_T(max)$)	98
6.1	Means of $P_T^{dif} = P_T^{lead} - P_T^{bal}$ for J3-J6 Truth events, Cone 0.7	102
6.2	Means of $P_T^{dif} = P_T^{lead} - P_T^{bal}$ for PYTHIA + jet-finder (JETCLU), Jet-Seed=2.0GeV	104
6.3	Means of $P_T^{dif} = P_T^{lead} - P_T^{bal}$ for ATLAS-CSC, J2-J6 truth jets, Cone 0.7 with a JES introduced for jet $P_T > 500 GeV$ of 10% over 500GeV and 1% over 500GeV	105
6.4	Means of $P_T^{dif} = P_T^{lead} - P_T^{bal}$ for ATLAS-CSC, J2-J6 reconstructed events, Cone 0.7 Tower Jets with a JES introduced for jet $P_T > 500 GeV$ of 10% over 500GeV and 1% over 500GeV	106
6.5	Means of $P_T^{dif} = P_T^{lead} - P_T^{bal}$ for reconstructed jets with no JES error and truth jets with a 1.5%/500GeV JES error. The thick continuous line is a linear fit of the reconstructed data. The truth data have been shifted in the y co-ordinate for ease of comparison. ATLAS-CSC, J2-J6 reconstructed events.	108
7.1	Change in inclusive jets cross-section at LHC for $0 < \eta < 1$ for four choices of compositeness scale, (constructive interference), PYTHIA 6.4	112
7.2	Change in inclusive jets cross-section at LHC for $1 < \eta < 2$ for four choices of compositeness scale, (constructive interference), PYTHIA 6.4	112

7.3	Change in inclusive jets cross-section at LHC for $2 < \eta < 3$ for four choices of compositeness scale, (constructive interference), PYTHIA 6.4 . . .	113
7.4	Relationship of the variable χ with the scattering angle in the centre of mass frame θ^*	118
7.5	Relationship of the scattering angle in the centre of mass frame θ^* with the variable χ	119
7.6	Change in di-jet cross-section (w.r.t SM) for dijets with invariant mass $2000 < M_{jj} < 2300 GeV$ for four choices of the compositeness scale Λ (constructive interference), cross-sections normalised to SM.	119
7.7	Change in di-jet cross-section (w.r.t SM) for dijets with invariant mass $2300 < M_{jj} < 2800 GeV$ for four choices of the compositeness scale Λ (constructive interference), cross-sections normalised to SM.	120
7.8	Change in di-jet cross-section (w.r.t SM) for dijets with invariant mass $2800 < M_{jj} < 3400 GeV$ for four choices of the compositeness scale Λ (constructive interference), cross-sections normalised to SM.	120
7.9	Change in di-jet cross-section (w.r.t SM) for dijets with invariant mass $M_{jj} > 3400 GeV$ for four choices of the compositeness scale Λ (constructive interference), cross-sections normalised to SM.	121
7.10	The Standard Model di-jet cross-section for di-jets with invariant mass $2000 < M_{jj} < 2300 GeV$ from PYTHIA and NLOJET ⁺⁺ at leading order for CTEQ5l	122
7.11	The Standard Model di-jet cross-section for di-jets with invariant mass $M_{jj} > 3400 GeV$ from PYTHIA and NLOJET ⁺⁺ at leading order for CTEQ5l	122
7.12	The ratio of the Standard Model di-jet cross-section prediction for di-jets with invariant mass $2000 < M_{jj} < 2300 GeV$ from NLOJET ⁺⁺ at leading and next-to-leading order.	123

7.13	The ratio of the Standard Model di-jet cross-section prediction for di-jets with invariant mass $M_{jj} > 3400\text{GeV}$ from NLOJET ⁺⁺ at leading and next-to-leading order.	124
7.14	The ratio of the Standard Model cross-section prediction for di-jets with invariant mass $2000 < M_{jj} < 2300\text{GeV}$ from PYTHIA at leading order and leading-log approximation.	125
7.15	The ratio of the Standard Model cross-section prediction for di-jets with invariant mass $M_{jj} > 3400\text{GeV}$ from PYTHIA at leading order and leading-log approximation.	125
7.16	PDF and Scale errors on SM di-jet cross-section for di-jets with invariant mass $2000 < M_{jj} < 2300\text{GeV}$ from an integration grid (CTEQ6.1 error sets used)	126
7.17	PDF and Scale errors on SM di-jet cross-section for di-jets with invariant mass $2300 < M_{jj} < 2800\text{GeV}$ from an integration grid (CTEQ6.1 error sets used)	126
7.18	PDF and Scale errors on SM di-jet cross-section for di-jets with invariant mass $2800 < M_{jj} < 3400\text{GeV}$ from an integration grid (CTEQ6.1 error sets used)	127
7.19	PDF and Scale errors on SM di-jet cross-section for di-jets with invariant mass $M_{jj} > 3400\text{GeV}$ from an integration grid (CTEQ6.1 error sets used) .	127
7.20	Experimental error based on the ratio of ‘truth’ to reconstructed data for ATLAS CSC J6 sample (212,000 events). The acceptance in the di-jet mass has been widened to $M > 2000\text{GeV}$ owing to the small sample size available.	128
7.21	Compositeness signal as a function of ϵ_0 , $M_{jj} > 3400\text{GeV}$, $10fb^{-1}$. No systematic errors considered (Constructive interference).	132
7.22	Significance of compositeness signal as a function of ϵ_0 , $M_{jj} > 3400\text{GeV}$, $10fb^{-1}$. No systematic errors considered (Constructive interference).	132

7.23	PDF and scale errors on compositeness signal as a function of ϵ_0 , $M_{jj} > 3400\text{GeV}$, 10fb^{-1}	133
8.1	A p-n junction in thermal equilibrium.	136
8.2	An SCT barrel module from [40]	138
8.3	A schematic of an ABCD3TA module from [41]	139
8.4	A diagram showing the relationship of efficiency (thick black line) with threshold for a silicon detector. The dashed line represents the noise distribution, the pale black line is the 'improved Landau distribution' from [42]	140
8.5	Distribution of noise occupancies (below 50×10^{-5}) from the monitoring data of the first 1000 events of run 7977 for 1452 barrel modules.	145
8.6	Ratio of monitoring results to RODs data from run 7977 (all modules with noise occupancy below 50×10^{-5}).	147
8.7	Ratio of monitoring results to RODs data from run 7977 (for the modules with noise occupancy $> 50 \times 10^{-5}$). The noisy module at barrel=3, eta=-3 and phi=22 is not plotted and has an exceptionally high ratio $monitoring/rods \sim 3$	147
8.8	Strip level noise occupancies for modules with noise above 50×10^{-5} . Black represents the monitoring software results, the red the ROD level information (run 7977)	149
8.9	Strip level noise occupancies for modules with noise above 50×10^{-5} . Black represents the monitoring software results, the red the ROD level information (run 7977)	150
8.10	STUCKON strip errors from calibration for module at barrel=2, $\eta = -3$, $\phi = 33$. The problem strip coincides exactly with the 'noisy' strip from the 'physics run' 7977.	151

8.11	HVchVolt for the four modules of the H8 test box in the M5 commissioning runs from the DCS monitoring tool. The x-axis relates to ‘eta’ and the y-axis to ‘phi’. The plot is taken from a ‘screen-shot’ of a histogram from the OH (online histogram display).	157
8.12	HVchVolt for layer 0 of the SCT Barrel in the M6 commissioning runs from the DCS monitoring tool. The x-axis relates to ‘eta’ and the y-axis to ‘phi’. The plot is taken from a ‘screen-shot’ of a histogram from the OH (online histogram display).	159
8.13	HVState for layer 0 of the SCT Barrel in the M6 commissioning runs from the DCS monitoring tool. The x-axis relates to ‘eta’ and the y-axis to ‘phi’. The plot is taken from a ‘screen-shot’ of a histogram from the OH (online histogram display).	159
9.1	Pseudorapidity of closest track to generated track trajectory in merged event minus pseudorapidity of generated track (For tracks with $P_T > 0.5\text{GeV}/c$). The mean of the distribution is zero (within errors).	164
9.2	Phi of closest track to generated track trajectory in merged event minus phi of generated track (For selected events, $P_T > 0.5\text{GeV}/c$). The mean of the distribution is zero (within errors).	165
9.3	P_T of closest track to generated track trajectory in merged event minus P_T of generated track (For selected events, $P_T > 0.5\text{GeV}/c$). The mean of the distribution is zero (within errors).	165
9.4	Distribution of ΔR of all tracks in event, for events with 5 tracks. Distributions are from PYTHIA (used to create ATLAS simulated data) and Toy MC (equivalent to track insertion technique). For 1697 events	166
9.5	Distribution of ΔR of all tracks in event, for events with 10 tracks. Distributions are from PYTHIA (used to create ATLAS simulated data) and Toy MC (equivalent to track insertion technique). For 926 events	166

-
- 9.6 Distribution of ΔR of all tracks in event, for events with 15 tracks. Distributions are from PYTHIA (used to create ATLAS simulated data) and Toy MC (equivalent to track insertion technique). For 545 events 167
- 9.7 Track reconstruction efficiency as determined by insertion into a minimum-bias (merged) and in isolation (isolated) as a function of η in event for min-bias events ($P_T > 0.5 GeV/c$). The three bins relate to the barrel, transition and end-cap regions of the inner detector. 168
- 9.8 Track reconstruction efficiency as determined by insertion into a min-bias (merged) and in isolation (isolated) as a function of track P_T for minimum-bias events in three regions of pseudorapidity ($P_T > 0.5 GeV/c$). 168
- 9.9 Track reconstruction efficiency as a function of the number of tracks in the original minimum-bias event ($P_T > 0.5 GeV/c$). 169
- 9.10 Track reconstruction efficiency as a function of ΔR from inserted track trajectory to nearest track in the original minimum-bias event ($P_T > 0.5 GeV/c$). 170

List of Tables

1.1	First generation of particles in the Standard Model, all are fermions i.e. have spin = $1/2$	2
1.2	Interactions of the Standard Model. The relative couplings are defined as: strong= α_S (strong coupling), EM= α (fine-structure constant) and Weak= $G_F m_p^2$ (where G_F is the four fermion coupling and m_p is the proton mass). The true strength of the interactions vary with the momentum transfer (Q^2) between participating particles. Here the low Q^2 limit is considered.	3
1.3	The three generations of the Standard Model	4
2.1	Inner detector subsystems, resolution and acceptance (from [3])	13
2.2	Anticipated track-parameter resolutions (from [3]) for muons (momentum and angular) and pions (impact parameters). $\sigma_X(\infty)$ relates to the resolution at infinite track momentum and P_X relates to a momentum where contributions from multiple scattering are equal to the intrinsic detector resolution.	14
2.3	Calorimeter subsystems and acceptance for jets (from [3]).	17
2.4	Design specifications for resolution of calorimeter subsystems (from [6],[3])	18
4.1	Key for Figures 4.12 and 4.13, all PDFs from the LHAPDF 5.0 interface. .	60

6.1	The range of allowed P_T values of the initial hard-scattering in PYTHIA for the ATLAS ‘J’ samples. CKIN(3) and CKIN(4) are the minimum and maximum values of permitted P_T allowed in GeV.	102
7.1	Chi-Squared and degrees of freedom for inclusive jet cross-section at ATLAS for three compositeness scale choices (constructive interference) considering JES and PDF errors, $0 < \eta < 1$, $L = 10fb^{-1}$	116
7.2	Chi-Squared and degrees of freedom for inclusive jet cross-section at ATLAS for three compositeness scale choices (destructive interference) considering JES and PDF errors, $0 < \eta < 1$, $L = 10fb^{-1}$	117
7.3	Significance of Chi-Squared (probability of null-hypothesis) for the inclusive jet cross-section at ATLAS for three compositeness scale choices for both constructive and destructive interference, considering JES and PDF errors, $0 < \eta < 1$, $L = 10fb^{-1}$	117
7.4	PDF errors (%) on R_χ as determined from CTEQ6.1 error sets (using integration grids).	130
7.5	$\mu_{r/f}$ scale uncertainty (%) on R_χ as determined from integration grids by varying scales between $0.5 \rightarrow 2.0 P_T$ of the leading jet and considering maximum changes in cross-section.	130
7.6	Compositeness signal (constructive interference) $R_\chi^\Lambda/R_\chi^{SM}$ and statistical error (for $10fb^{-1}$, $\epsilon_0 = 5$), only signals with strength above 1% shown. The final column is a combined (positive) systematic error on the signal from PDFs and scale choice.	130
7.7	Compositeness signal (destructive interference) $R_\chi^\Lambda/R_\chi^{SM}$ and statistical error (for $10fb^{-1}$, $\epsilon_0 = 5$), only signals with strength above 1% shown. The final column is a combined (positive) systematic error on the signal from PDFs and scale choice.	131
8.1	Noisy Modules from run 7977 (SCT Barrel in pit)	146

8.2	Noisy modules with strip defects identified in calibration that match noisy channels found in physics run 7977 from [44].	148
8.3	Noisy Modules in SCT with conditions data for temperature and high voltage bias for run 7977 - the data may have been compromised due to the conditions data becoming overloaded	153
8.4	A list of noisy modules with all channels noisy and their HVchVolt from the conditions database	154
8.5	Table of HV and LV state codes.	158

Chapter 11

Acknowledgements

I¹ would like to thank the University of Glasgow for both their facilities and financial support over the last 4 years, which allowed me to pursue my studies and make numerous visits to CERN in addition to two international conferences.

I would also like to take this opportunity to thank my supervisor Craig Buttar for his support and advice over the course of my research and for encouraging me to pursue my own ideas.

I am indebted to Tancredi Carli, Gavin Salem, Claire Gwenlan and Amanda-Cooper Sarkar who I worked with on the integration grids project. I would also like to thank Martin White and Dave Robinson who introduced me to the online-software world and with whom I worked on SCT commissioning and monitoring.

From Glasgow I would like to thank William Bell, Stan Thompson, Graeme Stewart, Paul Millar, Arthur Moraes, Chris-Collins Tooth, Sarah Allwood, Samir Ferrag and Saverio D'aurio for their advice and help with technical issues.

I wish to thank my friends Michael John McCarthy² and Kevin O'Holleran for their continued support and insights over the last years. Finally I want to thank my parents and brothers (Matt and Tom) for always being there when needed.

¹I find breaking from the third person passive increasingly difficult and hence its use is restricted to this section and the 'Declaration'.

²Of McCarthy and Zoey Van Goey fame

Bibliography

- [1] Tancredi Carli, Gavin P. Salam, and Frank Siegert. A posteriori inclusion of PDFs in NLO QCD final-state calculations. *arXiv.org*, hep-ph/0510324, 2005.
- [2] Oliver Bruning and Paul Collier. Building a behemoth. *NATURE*, 448, 2007.
- [3] The ATLAS Collaboration. The ATLAS Experiment at the CERN Large Hadron Collider. Technical Specifications https://twiki.cern.ch/twiki/pub/Atlas/AtlasTechnicalPaper/Main_inst_submission.pdf, ATLAS, In press - Journal of Instrumentation (JINST) 2008.
- [4] W.B Rolnick. *The Fundamental Particles and Their Interactions*. Addison-Wesley, 1st edition, 1994.
- [5] ATLAS Collaboration. ATLAS, Detector and Physics Performance, Technical Design Report, Volume 2. Technical Design Report CERN/LHCC/99-15, ATLAS, May 1999.
- [6] ATLAS Collaboration. Atlas, technical proposal. Technical Design Report CERN/LHCC/94-43, ATLAS, December 1994.
- [7] ATLAS Collaboration. ATLAS, Detector and Physics Performance, Technical Design Report, Volume 1. Technical Design Report CERN/LHCC/99-14, ATLAS, May 1999.
- [8] T Sjostrand et al. High-Energy-Physics Event Generation with PYTHIA 6.1. *Computer Phys Commun* 135, p238, 2001.

-
- [9] Ketevi Assamagan et al. Report of the AOD format task force. ATLAS Public Note (CDS) ATL-SOFT-PUB-2006-011, The ATLAS Collaboration, December 2006.
- [10] Barrand G. et al. Gaudi - a software architecture and framework for building hep data processing applications. *Computer Physics Communications*, 140:45–55(11), October 2001.
- [11] J. Allison, K. Amako, J. Apostolakis, H. Araujo, P. Arce Dubois, M. Asai, G. Barrand, R. Capra, S. Chauvie, R. Chytrcek, G.A.P. Cirrone, G. Cooperman, G. Cosmo, G. Cuttone, G.G. Daquino, M. Donszelmann, M. Dressel, G. Folger, F. Foppiano, J. Generowicz, V. Grichine, S. Guatelli, P. Gumplinger, A. Heikkinen, I. Hrivnacova, A. Howard, S. Incerti, V. Ivanchenko, T. Johnson, F. Jones, T. Koi, R. Kokoulin, M. Kossov, H. Kurashige, V. Lara, S. Larsson, F. Lei, O. Link, F. Longo, M. Maire, A. Mantero, B. Mascialino, I. McLaren, P. Mendez Lorenzo, K. Minamimoto, K. Murakami, P. Nieminen, L. Pandola, S. Parlati, L. Peralta, J. Perl, A. Pfeiffer, M.G. Pia, A. Ribon, P. Rodrigues, G. Russo, S. Sadilov, G. Santin, T. Sasaki, D. Smith, N. Starkov, S. Tanaka, E. Tcherniaev, B. Tome, A. Trindade, P. Truscott, L. Urban, M. Verderi, A. Walkden, J.P. Wellisch, D.C. Williams, D. Wright, and H. Yoshida. Geant4 developments and applications. *Nuclear Science, IEEE Transactions on*, 53(1):270–278, Feb. 2006.
- [12] LCG Project. “LHC computing grid”. Technical Design Report CERN-LHCC-2005-024, LCG, June 2005.
- [13] A.J.G. Hey I.J.R. Aitchison. *Gauge Theories in Particle Physics*. IOP Publishing Ltd, 2nd edition, 1989.
- [14] M.E Peskin and D.V Schroeder. *An introduction to Quantum Field Theory*. Westview Press, 1st edition, 1995.
- [15] G. Corcella et al. HERWIG 6.5 release note. *arXiv.org*, hep-ph/0210213, 2002.

-
- [16] Zoltan Nagy. Next-to-leading order calculation of three-jet observables in hadron-hadron collisions. *Phys. Rev.*, D68:094002, 2003.
- [17] W. T. Giele, E. W. Nigel Glover, and David A. Kosower. Higher order corrections to jet cross-sections in hadron colliders. *Nucl. Phys.*, B403:633–670, 1993.
- [18] Stefano Frixione and Bryan R. Webber. Matching nlo qcd computations and parton shower simulations. *JHEP*, 06:029, 2002.
- [19] W.-M. Yao, C. Amsler, D. Asner, R.M. Barnett, J. Beringer, P.R. Burchat, C.D. Carone, C. Caso, O. Dahl, G. D’Ambrosio, A. DeGouvea, M. Doser, S. Eidelman, J.L. Feng, T. Gherghetta, M. Goodman, C. Grab, D.E. Groom, A. Gurtu, K. Hagiwara, K.G. Hayes, J.J. Hernández-Rey, K. Hikasa, H. Jawahery, C. Kolda, Kwon Y., M.L. Mangano, A.V. Manohar, A. Masoni, R. Miquel, K. Mönig, H. Murrayama, K. Nakamura, S. Navas, K.A. Olive, L. Pape, C. Patrignani, A. Piepke, G. Punzi, G. Raffelt, J.G. Smith, M. Tanabashi, J. Terning, N.A. Törnqvist, T.G. Trippe, P. Vogel, T. Watari, C.G. Wohl, R.L. Workman, P.A. Zyla, B. Armstrong, G. Harper, V.S. Lugovsky, P. Schaffner, M. Artuso, K.S. Babu, H.R. Band, E. Barberio, M. Battaglia, H. Bichsel, O. Biebel, P. Bloch, E. Blucher, R.N. Cahn, D. Casper, A. Cattai, A. Ceccucci, D. Chakraborty, R.S. Chivukula, G. Cowan, T. Damour, T. DeGrand, K. Desler, M.A. Dobbs, M. Drees, A. Edwards, D.A. Edwards, V.D. Elvira, J. Erler, V.V. Ezhela, W. Fetscher, B.D. Fields, B. Foster, D. Froidevaux, T.K. Gaisser, L. Garren, H.-J. Gerber, G. Gerbier, L. Gibbons, F.J. Gilman, G.F. Giudice, A.V. Gritsan, M. Grünewald, H.E. Haber, C. Hagmann, I. Hinchliffe, A. Höcker, P. Igo-Kemenes, J.D. Jackson, K.F. Johnson, D. Karlen, B. Kayser, D. Kirkby, S.R. Klein, K. Kleinknecht, I.G. Knowles, R.V. Kowalewski, P. Kreitz, B. Krusche, Yu.V. Kuyanov, O. Lahav, P. Langacker, A. Liddle, Z. Ligeti, T.M. Liss, L. Littenberg, L. Liu, K.S. Lugovsky, S.B. Lugovsky, T. Mannel, D.M. Manley, W.J. Marciano, A.D. Martin, D. Milstead, M. Narain, P. Nason, Y. Nir, J.A. Peacock, S.A. Prell, A. Quadt, S. Raby, B.N. Ratcliff, E.A. Razuvaev, B. Renk, P. Richardson, S. Roesler,

- G. Rolandi, M.T. Ronan, L.J. Rosenberg, C.T. Sachrajda, S. Sarkar, M. Schmitt, O. Schneider, D. Scott, T. Sjöstrand, G.F. Smoot, P. Sokolsky, S. Spanier, H. Spieler, A. Stahl, T. Stanev, R.E. Streitmatter, T. Sumiyoshi, N.P. Tkachenko, G.H. Trilling, G. Valencia, K. van Bibber, M.G. Vincter, D.R. Ward, B.R. Webber, J.D. Wells, M. Whalley, L. Wolfenstein, J. Womersley, C.L. Woody, A. Yamamoto, O.V. Zenin, J. Zhang, and R.-Y. Zhu. Review of Particle Physics. *Journal of Physics G*, 33:1+, 2006.
- [20] Diagram from: HEPDATA Durham Database Group. On-line parton distribution calculator with graphical display. <http://durpdg.dur.ac.uk/hepdata/pdf3.html>.
- [21] T Affolder and the CDF collaboration. Measurement of the inclusive jet cross-section in proton-antiproton collisions at $\sqrt{s} = 1.8\text{TeV}$. *Phys Rev D*, V64, 032001, 2001.
- [22] Matteo Cacciari and Gavin P Salam. Dispelling the n^3 myth for the k_t jet-finder. *Phys.Lett.B641:57-61*, 2006.
- [23] Daniel Stump et al. Inclusive jet production, parton distributions, and the search for new physics. *JHEP*, 10:046, 2003.
- [24] J W Huston J M Campbell and W J Stirling. Hard interactions of quarks and gluons: a primer for LHC physics. *Rep. Prog. Phys.*, 70:89–193, 2007.
- [25] D Clements. Prospects for the constraining of PDFs from ATLAS jet data. DIS 2006, World-Scientific, 2006.
- [26] T. Kluge, K. Rabbertz, and M. Wobisch. Fast pQCD calculations for PDF fits. *arXiv.org*, hep-ph/0609285, 2006.
- [27] F. James and M. Roos. Minuit - a system for function minimization and analysis of the parameter errors and correlations. *Computer Physics Communications*, 10(6):343–367, 1975.

-
- [28] W. Giele et al. The QCD/SM working group - Summary report. *arXiv.org*, hep-ph/0204316, 2002.
- [29] A. M. Cooper-Sarkar. An NLO QCD analysis of inclusive cross-section data and jet production data from the ZEUS experiment at HERA-I. *arXiv.org*, hep-ph/0407309, 2004.
- [30] CDF Collaboration. Measurement of the strong coupling constant from inclusive jet production at the tevatron $p\bar{p}$ collider. *Phys Rev Lett*, V88:042001, 2002.
- [31] T. Affolder, H. Akimoto, A. Akopian, M. G. Albrow, P. Amaral, S. R. Amendolia, D. Amidei, K. Anikeev, J. Antos, G. Apollinari, T. Arisawa, T. Asakawa, W. Ashmanskas, F. Azfar, P. Azzi-Bacchetta, N. Bacchetta, M. W. Bailey, S. Bailey, P. de Barbaro, A. Barbaro-Galtieri, V. E. Barnes, B. A. Barnett, S. Baroiant, M. Barone, G. Bauer, F. Bedeschi, and S. Belforte. Measurement of the inclusive jet cross section in $p\bar{p}$ collisions at $\sqrt{s} = 1.8$ TeV. *Phys. Rev. D*, 64(3):032001, Jun 2001.
- [32] Joey Huston. Private communication - 3rd august 2007. email.
- [33] T.Schorner-Sadenius and S.Taprogge. ATLAS Trigger Menus for the LHC Start-up Phase. ATLAS Public Note (CDS) ATL-DAQ-2003-004, The ATLAS Collaboration, July 2003.
- [34] Gerald C. Blazey, Jay R. Dittmann, Stephen D. Ellis, V. Daniel Elvira, K. Frame, S. Grinstein, Robert Hirosky, R. Piegaia, H. Schellman, R. Snihur, V. Sorin, and Dieter Zeppenfeld. Run II jet physics: Proceedings of the run II qcd and weak boson physics workshop. *arXiv.org*, hep-ex/0005012, 2000.
- [35] B. Abbott et al. The dijet mass spectrum and a search for quark compositeness in $p\bar{p}$ collisions at $\sqrt{s} = 1.8$ TeV. *Phys. Rev. Lett.*, 82:2457–2462, 1999.
- [36] 't Hooft. *Recent Developments in Gauge Theories*. Springer, 1st edition, 1980.

-
- [37] H Harari. Composite models for quarks and leptons. *Physics at very high-energies: Proceedings of the 10th SLAC summer institute on particle physics (SSI 82), Stanford, Calif.*, Aug 1982. SLAC-0259.
- [38] I Bars. Proceedings of the seventeenth rencontre de moriond. Edition Frontieres, 1982.
- [39] Estia J. Eichten, Kenneth D. Lane, and Michael E. Peskin. New tests for quark and lepton substructure. *Phys. Rev. Lett.*, 50(11):811–814, Mar 1983.
- [40] A Carter et al. Sct barrel module: Requirements and specification of barrel modules. Technical Specifications SCT-BM-FDR-4, ATLAS, March 2002.
- [41] A J Barr. *Studies of supersymmetry models for the ATLAS experiment at the Large Hadron Collider*. PhD thesis, Cambridge Univ., 2002. Presented on 30th Nov 2002.
- [42] M.J Palmer. *Studies of Extra-Dimensional Models With The ATLAS Detector*. PhD thesis, Cambridge Univ., 2005.
- [43] Sergey Burdin. Private communication - 29th may 2007. email.
- [44] Dave Robinson. Private communication - 26th june 2007. email.