

Exploring the Standard Model and beyond with jets from proton-proton collisions at $\sqrt{s}=13$ TeV with the ATLAS Experiment

A DISSERTATION PRESENTED
BY
JENNIFER KATHRYN ROLOFF
TO
THE DEPARTMENT OF PHYSICS

IN PARTIAL FULFILLMENT OF THE REQUIREMENTS
FOR THE DEGREE OF
DOCTOR OF PHILOSOPHY
IN THE SUBJECT OF
PHYSICS

HARVARD UNIVERSITY
CAMBRIDGE, MASSACHUSETTS
MAY 2019

©2019 – JENNIFER KATHRYN ROLOFF
ALL RIGHTS RESERVED.

Exploring the Standard Model and beyond with jets from proton-proton collisions at $\sqrt{s}=13$ TeV with the ATLAS Experiment

ABSTRACT

This thesis documents work from two analysis performed using $\sqrt{s} = 13$ TeV data taken with the ATLAS detector. The first portion of this thesis presents a measurement of the jet mass. This is the first measurement of a substructure observable that can be compared to analytical predictions beyond leading logarithmic accuracy. The measurement demonstrates the ability to precisely measure jet substructure observables, and the analytical calculations are shown to be in good agreement with the unfolded data distribution, paving the way towards a broader program of jet substructure measurements at the Large Hadron Collider.

The second part of this thesis lays the foundation for a search for R-parity violating supersymmetry in states which involve jets and displaced vertices. This search is sensitive to models that have long-lived particles decaying within the tracking detector. This thesis explores the sensitivity of this search, and provides preliminary background estimation techniques that will be used for this analysis.

Contents

0	INTRODUCTION	1
1	THE STANDARD MODEL	4
1.1	The Electroweak Sector	5
1.2	Electroweak Symmetry Breaking and the Higgs Sector	7
1.3	The Quantum Chromodynamics Sector	10
2	JETS AND JET ALGORITHMS	12
2.1	Parton Showers and Hadronization	13
2.2	Jet Algorithms	14
2.3	Jet Grooming	15
3	THE LARGE HADRON COLLIDER	22
4	THE ATLAS DETECTOR	25
4.1	The Inner Detector	26
4.2	The Calorimeters	30
4.3	The Muon Spectrometer	35
4.4	Trigger and Data Acquisition	36
4.5	Data Quality Monitoring	37
5	OBJECT RECONSTRUCTION	38
5.1	Track Reconstruction	39
5.2	Vertex Reconstruction	42
5.3	Jets	47
6	CONSTITUENT-LEVEL PILEUP MITIGATION	56
6.1	Pileup Mitigation Techniques	57
6.2	Impact of Pileup Mitigation on Small-R Jets	64
6.3	Impact of Pileup Mitigation on Large-R Jets	67
6.4	Summary	74
I	Measurements of Jet Substructure Observables	77
7	THEORETICAL BACKGROUND AND ANALYTICAL PREDICTIONS	79
8	EVENT SELECTION AND SAMPLES	81
8.1	Data Samples	81
8.2	Monte Carlo Samples	82
8.3	Event Selection	82

9	UNFOLDING	84
9.1	Iterative Bayesian Regularized Unfolding	88
9.2	Considerations in Unfolding	89
10	SYSTEMATIC UNCERTAINTIES	92
10.1	Statistical Uncertainties	93
10.2	Systematic Uncertainties	93
10.3	Summary of Uncertainties	98
11	MEASUREMENTS OF JET SUBSTRUCTURE OBSERVABLES FOR SOFT DROPPED JETS	103
11.1	Comparison of Unfolded Data to Monte Carlo Predictions	105
11.2	Comparison of Unfolded Data to Analytical Predictions	105
12	SUMMARY OF THE JET MASS MEASUREMENT	109
II	Searching for Long-Lived Particles using Displaced Vertices	111
13	SUPERSYMMETRY	113
13.1	The Basics of Supersymmetry	114
13.2	Long-Lived Particles	116
13.3	R-hadrons	120
13.4	Signal Model for the Search	123
14	EVENT AND VERTEX SELECTION	125
14.1	Trigger and DRAW Requirements	126
14.2	Event and Vertex Selection	129
14.3	Prospects for the Search	132
14.4	Backgrounds	133
15	BACKGROUND ESTIMATION FOR HADRONIC INTERACTIONS	139
15.1	Construction of a Monte Carlo Template	140
15.2	Validation in Data	141
16	MERGED VERTEX BACKGROUND ESTIMATION	158
16.1	Template Creation	159
16.2	Towards a background estimation	161
17	ADDITIONAL STUDIES ON A BACKGROUND ESTIMATION	165
18	SUMMARY OF THE SEARCH FOR DISPLACED VERTICES PLUS JETS	170
	APPENDIX A FULL PLOTS FOR SINGLE PERIOD BACKGROUND ESTIMATE	171
	APPENDIX B CORRECTIONS FOR TRUTH JETS ASSOCIATED TO LONG-LIVED PARTICLES	182
	APPENDIX C DIGITRUTH	187
	C.1 Signal modeling	188
	C.2 DigiTruth	190
	REFERENCES	199

Acknowledgments

I am immensely grateful to the many people who have supported me on the path to my Ph.D. I have been fortunate to have many wonderful friends, colleagues, mentors, and teachers throughout the years. While there is not enough space here to mention everyone who has helped me along the way, a few people have been particularly important. First, I would like to thank my advisor, John Huth, who has helped guide my work throughout the course of my studies, and who gave me the freedom to explore a wide variety of topics as a student.

I doubt that I would have ever entered high energy physics had I not been a CERN summer student with the University of Michigan. There are a number of wonderful people from that program, and in particular, I would like to thank Junjie Zhu, who has always been incredibly supportive and has helped me get to where I am today. I would also like to thank Homer Neal, whose insight and thoughtfulness are deeply missed.

I owe a great deal of thanks to David Lopez Mateos, who taught me an extraordinary amount of things about jets and being a physicist. He helped guide my work and pushed me in directions that I am still following now, and I am incredibly fortunate that I was able to work with him. Since that time, I have been lucky to work with Larry Lee, who constantly has new ideas for how to approach problems, and who has taught me a lot about long-lived particles.

I have been lucky to work with a number of amazing people on a variety of projects during my time at CERN. The grooming scan team was the first large group that I worked with, and the entire team has inspired me with their enthusiasm and insights. I have been especially fortunate to co-lead the effort with Joe Taenzer, whose skills I appreciate the more that I work with him. Within the Standard Model group, I have also been lucky enough to work on a couple of small teams of dedicated people. Ben Nachman has been a constant source of enthusiasm and ideas. He and Max Swiatlowski patiently taught me many of the things I know about performing measurements, and I have greatly enjoyed the continuation of this work with him and Matt LeBlanc. TJ Khoo has been a constant source of knowledge, both as co-leader of the pileup forum, as well as for a variety of technical issues, and he has

been kind enough to help me solve a number of problems that few others would have been able to.

The entire JetEtMiss group has been so welcoming and supportive, and I am glad that I have been able to work with many people from the group. In particular, Kate Pachal, Zach Marshall, and Steven Schramm have provided a lot of encouragement, advice, and assistance, and I am thankful for their kindness. My thanks as well to the many others who have helped me along the way, making my time at CERN enjoyable and interesting.

Finally, there is no doubt that I would not be here without my parents, who have continuously encouraged me to study a wide variety of things. I am grateful for the support they have given, and am fortunate to have such loving parents.

0

Introduction

The Large Hadron Collider (LHC) has successfully completed two full runs of data-taking. This period has been marked by a number of successes, including the discovery of the higgs boson, precise measurements of parameters of the Standard Model, and limits on new particles extending into the TeV range. In spite of this, no new physics beyond the Standard Model has been discovered. This does not mean that there is no new physics that can be discovered at the LHC, but it does mean that signs of new physics may be difficult to discover. Even though many aspects the Standard Model are well-understood, there are still many features that are difficult to describe analytically, and rely heavily on Monte Carlo simulation. In particular, quantum chromodynamics (QCD) is challenging to model due to the behavior of the strong force, which results in confinement, meaning that individual quarks can-

not exist on their own, and must be contained in color-neutral hadrons. Instead of observing quarks and gluons, collimated sprays of hadrons (*jets*) are reconstructed. Jets are important for understanding QCD, but also are challenging to model.

Recently, several advances have been made in the theoretical understanding of the inner structure of jets, allowing the first calculations of jet substructure observables beyond leading logarithmic accuracy. This thesis presents the first comparison of a measurement of a jet substructure observable to a theoretical prediction of this accuracy. This paves the way towards future measurements of jet substructure observables, which are important for creating better models of the behavior of jets, which is crucial at a hadron collider.

In addition to improving our understanding of the Standard Model, it is important to explore the full variety of searches for new physics which are possible at the LHC. Long-lived particles exist in a variety of models, but searches for long-lived particles often require specialized reconstruction and background estimation, and many searches are not able to set strong limits on these models. As stronger limits are being set on models with promptly decaying particles, it is important to also search for long-lived particles. This thesis presents the foundations for a search for long-lived particles which are produced in association with jets by searching for final states which have a displaced vertex and several jets. The studies shown here include preliminary techniques for the background estimation as well as studies of the sensitivity of this search.

The first few chapters provide the fundamentals for this search, starting with an overview of the Standard Model in Chapter 1, and an overview of jets in Chapter 2. The theoretical background is followed by a brief overview of the Large Hadron Collider in Chapter 3, and an overview of the ATLAS detector in Chapter 4. Then, a basic description of object reconstruction is provided in Chapter 5, and a few studies of jet reconstruction with constituent-level pileup mitigation in Chapter 6.

Part I describes the measurement of the jet mass, and is divided into several chapters for simplicity. Chapter 7 describes the theoretical background and analytical predictions relevant to the measurement. Then, the event selection for the measurement are described in Chapter 8. A description of unfolding is provided in Chapter 9, and the various systematic and statistical uncertainties are described in Chapter 10. Finally, the full measurement is

provided in Chapter 11, and a summary of the measurement and its implications are provided in Chapter 12.

Part II lays the foundations for a search for displaced vertices plus jets. Chapter 13 provides an overview of supersymmetry, and the theoretical models relevant for this search. The event selection for the search is provided in Chapter 14. Then, the background estimation techniques for hadronic interactions and merged vertices are described in Chapters 15 and 16, and a preliminary exploration of alternative background estimation technique is shown in Chapter 17. Finally, a summary of the status of the search is given in Chapter 18.

1

The Standard Model

The Standard Model (SM) of particle physics is a relativistic quantum field theory composed of force-mediating bosonic fields and fermionic matter fields, which is able to describe elementary particles and their interactions, with the exception of the gravitational force. It is the result of many years of experimental and theoretical developments, and was only recently completed in 2012 with the discovery of the Higgs boson by the ATLAS and CMS experiments^{5 54}.

The Standard Model is described by an $SU(3) \times SU(2) \times U(1)$ gauge symmetry, and contains 26 free parameters that must be determined experimentally. Within this framework, particles may be described by their spin, mass, and charges under the different symmetries. Fermions are categorized as *quarks*, which have color charge under the

SU(3) symmetry, and color-neutral *leptons*. These are then further subdivided into up-type and down-type quarks, where up-type quarks have electric charge of 2/3, and down-type quarks have electric charge of -1/3, while leptons are divided into electron-type, which have electric charge of -1, and neutrinos, which have no electric charge. It has been observed that each of these subcategories is composed of three *generations*, which are only distinguishable by the mass of the particles. A table showing all known particles and their properties may be seen in Table 1.1¹⁰⁴. While these particles are the basis for understanding the Standard Model, their interactions may be written more succinctly using the Standard Model Lagrangian, which may be written as

$$\mathcal{L}_{\text{SM}} = -\frac{1}{4}F_{\mu\nu}F^{\mu\nu} + i\psi^\dagger \not{D}\psi + \psi_i y_{ij} \psi_j \phi + h.c. + |D_\mu \phi|^2 - V(\phi). \quad (1.1)$$

In order to simplify the explanation of this equation, it is possible to break the Lagrangian into several components. The SU(2)×U(1) sector describes both electroweak (EW) interactions and the higgs sector, while the SU(3) sector describes quantum chromodynamics (QCD). In short, it may all be written as

$$\mathcal{L}_{\text{SM}} = \mathcal{L}_{\text{QCD}} + \mathcal{L}_{\text{EW}} + \mathcal{L}_{\text{Higgs}}. \quad (1.2)$$

Each of these sectors is described in detail below.

1.1 THE ELECTROWEAK SECTOR

The Electroweak Lagrangian describes the SU(2)×U(1) sector, where the U(1) symmetry is the gauge group for *hypercharge*, and the SU(2) symmetry is weak isospin. It may be further broken into several components in order to simplify the understanding of it.

$$\mathcal{L}_{\text{EW}} = \mathcal{L}_{\text{gauge}} + \mathcal{L}_{\text{fermions}} + \mathcal{L}_{\text{Higgs}} + \mathcal{L}_{\text{Yukawa}}. \quad (1.3)$$

The gauge fields for the U(1) and SU(2) symmetries may be written as B_μ and W_μ^a respectively, where $a \in [1, 3]$ since there are three generators for SU(2). Then, their field strength tensors may be written as $B_{\mu\nu}$ and $W_{\mu\nu}^a$,

Table 1.1: A table of the Standard Model particles and their properties

classification	name	symbol	spin	electric charge	mass
up-type quarks	up	u	1/2	+2/3	2.2 MeV
	charm	c	1/2	+2/3	1.3 GeV
	top	t	1/2	+2/3	173 GeV
down-type quarks	down	d	1/2	-1/3	4.7 MeV
	strange	s	1/2	-1/3	95 MeV
	bottom	b	1/2	-1/3	4.2 GeV
electron	electron	e	1/2	-1	0.51 MeV
	muon	μ	1/2	-1	105.6 MeV
	tau	τ	1/2	-1	1776 MeV
neutrino	electron neutrino	ν_e	1/2	0	< 2 eV
	muon neutrino	ν_μ	1/2	0	< 2 eV
	tau neutrino	ν_τ	1/2	0	< 2 eV
gauge boson	photon	γ	1	0	$< 1 \times 10^{-18}$ eV
	W^+	W^+	1	+1	80.3 GeV
	W^-	W^-	1	-1	80.3 GeV
	Z	Z^0	1	0	91.2 GeV
	gluon	g	1	0	0
Higgs boson	Higgs	H	0	0	125.2 GeV

where

$$B_{\mu\nu} = \partial_\mu B_\nu - \partial_\nu B_\mu,$$

$$W_{\mu\nu}^a = \partial_\mu W_\nu^a - \partial_\nu W_\mu^a - g\epsilon_{abc}W_\mu^b W_\nu^c,$$
(1.4)

Putting this together, and adding in fermionic Dirac spinors ψ for the matter fields, the electroweak Lagrangian may be written as

$$\mathcal{L}_{\text{gauge}} = -\frac{1}{4}B_{\mu\nu}B^{\mu\nu} - \frac{1}{4}W_{\mu\nu}^a W^{a\mu\nu},$$
(1.5)

where g, g' are the coupling constants for the U(1) and SU(2) symmetries respectively, and ϵ_{abc} is the totally asymmetric tensor.

Then, the fermionic term may be included using the Lagrangian

$$\mathcal{L}_f = \sum_\psi \bar{\psi} \gamma^\mu (i\partial_\mu - g'Y_W B_\mu - g\tau_{L_a} W_\mu^a) \psi$$

$$= \sum_{m=1}^F (\bar{q}_{mL}^0 iD\gamma^\mu q_{mL}^0 + \bar{l}_{mL}^0 iD\gamma^\mu l_{mL}^0 + \bar{u}_{mR}^0 iD\gamma^\mu u_{mR}^0 + \bar{d}_{mR}^0 iD\gamma^\mu d_{mR}^0 + \bar{e}_{mR}^0 iD\gamma^\mu e_{mR}^0 + \bar{\nu}_{mR}^0 iD\gamma^\mu \nu_{mR}^0),$$
(1.6)

where m sums over F , the number of families, and L and R refer to the chirality of the fields as left- or right-handed, and D are the covariant derivatives for each of the fields. Left-handed fields transform as SU(2) doublets,

$$\begin{aligned} D_\mu q_{mL}^0 &= (\partial_\mu + \frac{ig}{2}\tau\dot{W}_\mu + \frac{ig'}{6}B_\mu)q_{mL}^0 \\ D_\mu l_{mL}^0 &= (\partial_\mu + \frac{ig}{2}\tau\dot{W}_\mu + \frac{ig'}{6}B_\mu)l_{mL}^0, \end{aligned} \tag{1.7}$$

where τ_L are the generators of the SU(2) symmetry, and g, g' are the coupling constants for the U(1) and SU(2) symmetries respectively. Similarly, for the right-handed fields, which are SU(2) singlets, the covariant derivative may be defined as

$$\begin{aligned} D_\mu u_{mR}^0 &= (\partial_\mu + \frac{ig'}{3}B_\mu)u_{mR}^0 \\ D_\mu d_{mR}^0 &= (\partial_\mu + \frac{ig'}{3}B_\mu)d_{mR}^0 \\ D_\mu e_{mR}^0 &= (\partial_\mu + \frac{ig'}{3}B_\mu)e_{mR}^0 \\ D_\mu \nu_{mR}^0 &= \partial_\mu \nu_{mR}^0, \end{aligned} \tag{1.8}$$

Reading off from the Lagrangian, there are both kinetic terms for the bosonic and fermionic fields, as well as interactions between the fermions and bosons. Notably missing from this equation are mass terms for the fermions, which cannot be added directly since $m\bar{\psi}\psi$ is not SU(2) invariant. Fermions in the Standard Model are massive, so there is evidently still a missing piece to this equation. This is what motivated the inclusion of a Higgs boson, which was theorized in 1964.

1.2 ELECTROWEAK SYMMETRY BREAKING AND THE HIGGS SECTOR

The physics described in the previous two sections is sufficient to explain much of the interactions we see, but still cannot explain certain things such as the masses of gauge bosons. However, this can be explained if we include the Higgs sector, which is written as

$$\mathcal{L}_{\text{Higgs}} = (D_\mu \phi)^2 + \mu^2 \phi^\dagger \phi - \lambda(\phi^\dagger \phi)^2, \tag{1.9}$$

where ϕ is a complex doublet with hypercharge $\frac{1}{2}$, and is often referred to as the Higgs multiplet, and its covariant derivative $D_\mu\phi$ may be written as

$$D_\mu\phi = \partial_\mu\phi - igW_\mu^a\tau^a\phi - \frac{1}{2}ig'B_\mu H. \quad (1.10)$$

The inclusions of this term has remarkable consequences. The first observation is that $\mu^2\phi^\dagger\phi - \lambda(\phi^\dagger\phi)^2$, or the Higgs potential, has local minimum at $\sqrt{\frac{2\mu^2}{\lambda}}$. In order to be in the ground state, the symmetry of the potential is spontaneously broken. In order to minimize the potential, the symmetry is spontaneously broken. Solving for H , and making a choice of gauge that simplifies the result, it may be written as

$$H = \begin{pmatrix} 0 \\ \frac{v}{\sqrt{2}} + \frac{h}{\sqrt{2}} \end{pmatrix}, \quad (1.11)$$

where $v = \mu/\sqrt{\lambda}$, and h is a real scalar field.

This term gives mass to the gauge bosons, which can be seen from the term

$$|D_\mu H|^2 = g^2 \frac{v^2}{8} [(W_\mu^1)^2 + (W_\mu^2)^2 + (\frac{g'}{g}B_\mu - W_\mu^3)^2]. \quad (1.12)$$

In this basis, the different bosons are not mass eigenstates. To counter this effect, the W^3 and B bosons are typically rewritten as a linear combination that avoids this problem, with

$$\begin{aligned} Z_\mu &= \cos(\theta_w)W_\mu^3 - \sin(\theta_w)B_\mu \\ A_\mu &= \sin(\theta_w)W_\mu^3 - \cos(\theta_w)B_\mu, \end{aligned} \quad (1.13)$$

where $\tan(\theta_w) = \frac{g'}{g}$. In order to avoid other terms that couple to W^3 , the W^1 and W^2 terms are also transformed such that they couple instead to A , writing them as

$$W_\mu^\pm = \frac{1}{\sqrt{2}}(W_\mu^1 \mp W_\mu^2). \quad (1.14)$$

This results in three massive bosons $W^\pm$ and Z , and one massless boson A , which correspond to the bosons typ-

ically described in the Standard Model. Based on this result, it is clear that the masses of the gauge bosons are related, where

$$\begin{aligned} m_W &= \frac{v}{2}g \\ m_Z &= \frac{v}{2}\sqrt{g^2 + g'^2} = \frac{m_W}{\cos(\theta_W)}. \end{aligned} \tag{1.15}$$

Similarly, the inclusion of the Higgs boson solves the problem of fermion masses. It is possible to add a Yukawa coupling term to the Lagrangian,

$$\mathcal{L}_{\text{Yukawa}} = \sum_{m,n}^F (\bar{u}_{mL}^0 \Gamma_{mn}^u H u_{mR}^0 + \bar{d}_{mL}^0 \Gamma_{mn}^u H d_{mR}^0 + \bar{e}_{mL}^0 \Gamma_{mn}^u H e_{mR}^0 + \bar{\nu}_{mL}^0 \Gamma_{mn}^u H \nu_{mR}^0 + h.c.). \tag{1.16}$$

After spontaneous symmetry breaking, these terms have two parts: one for the Yukawa interaction, and one that resembles a mass term. For instance, this includes the term

$$\mathcal{L}_f = \sum_{m,n}^F (\bar{u}_{mL}^0 \Gamma_{mn}^u \frac{\nu}{\sqrt{2}} u_{mR}^0 + \bar{u}_{mL}^0 \Gamma_{mn}^u \frac{g M^\mu}{2M_W} u_{mR}^0), \tag{1.17}$$

where M^μ is the fermion mass matrix, meaning that spontaneous symmetry breaking also restores mass terms for fermions.

For the quark fields, there is an additional complication. It is impossible to choose a basis that simultaneously diagonalizes the up and down mass bases. However, it is possible to create mass eigenstates for the quarks, at the cost of including couplings that break the SU(2) symmetry. These terms are typically described using the Cabibbo-Kobayashi-Maskawa (CKM) matrix, which describes this misalignment. The end result of this is that W bosons are able to couple up-type and down-type quarks.

1.3 THE QUANTUM CHROMODYNAMICS SECTOR

The strong force is mediated by a spin-1 boson called the *gluon*, which interacts with itself and spin-1/2 particles called *quarks*. The Lagrangian dictating its dynamics may be written as

$$\mathcal{L}_{\text{QCD}} = -\frac{1}{4}F_{\mu\nu}^a F^{a\mu\nu} + \sum_r \bar{\psi}_{r\alpha} i D_\beta^\alpha \gamma^\mu \psi^{r\beta}, \quad (1.18)$$

where

$$F_{\mu\nu}^i = \partial_\mu G_\nu^i - \partial_\nu G_\mu^i - g_s f_{ijk} G_\mu^j G_\nu^k \quad (1.19)$$

is the field strength tensor, and g_s is the QCD gauge coupling constant, and

$$D_{\mu\beta}^\alpha = \partial_\mu \delta_{\alpha\beta} + i g_s G_\mu^i L_{\alpha\beta}^i. \quad (1.20)$$

Since there are eight generators for an SU(3) symmetry, i runs from 1 to 8, where each G^i represents one of the eight gluon fields. The $\alpha, \beta \in [1, 2, 3]$ indices represent the color, which are commonly referred to as *red*, *green*, and *blue*, and r is the quark flavor. The field strength tensor contains both the kinetic terms for the gluon field, as well as self-interaction terms, providing both a 3-gluon and 4-gluon interaction ⁸².

In a quantum field theory, the mass and coupling terms are not static, and instead *run* with the energy scale of the process. One interesting feature of QCD is the running of the coupling g_s

$$\beta(\alpha_s) = Q^2 \frac{\partial \alpha_s}{\partial Q^2} = -\left(\frac{\alpha_s}{4\pi}\right)^2 \sum_{n=0} \left(\frac{\alpha_s}{4\pi}\right)^n \beta_n. \quad (1.21)$$

The β function describes how the effective strength of α_s changes as a function of Q^2 , where Q is the momentum transfer. The different β_n terms of this equation may be solved for, and describe the running of the coupling.

In particular,

$$\beta_0(\alpha_s) = -\frac{\alpha_s^2}{2\pi} \left(\frac{11}{3} n_{\text{colors}} - \frac{4}{3} n_{\text{flavors}} \right). \quad (1.22)$$

Ignoring the higher order terms, this can be solved to give the strength of α_s as a function of Q and the QCD scale parameter Λ ,

$$\alpha_s(Q^2) = \frac{4\pi}{\beta_0 \ln(Q^2/\Lambda^2)}. \quad (1.23)$$

In this equation, there is a competition between the quarks screening the color charge, and the gluon anti-screening which occurs because of the gluon self-coupling⁶¹. In Standard Model QCD, there are 3 colors and 3 flavors, meaning that the β function has a negative value. This results in *asymptotic freedom*, where the strength of the strong force drops to zero at high energies or short distances. However, the converse is also true, where at large distances, the strength of QCD approaches infinity. One result of this, for which there is ample evidence, but still no theoretical proof, is *color confinement*, where colored objects will form color-neutral bound states with other objects, and will not be observed directly on their own. This means that instead of observing quarks directly, we observe hadrons, which are color-neutral bound states of quarks. This includes baryons like the proton and neutron, which are composed of any set of three quarks (or three antiquarks), and mesons, which are composed of quark-antiquark pairs. The experimental and theoretical implications of this are discussed more in the next chapter.

1.3.1 PARTON DISTRIBUTION FUNCTIONS

One important consequence of confinement for proton collider experiments is in the structure of a proton, which determines the probability of various hard-scatter processes. While a proton is approximately a bound state of two up quarks and a down quark, its inner structure is in fact much more complicated. Instead of just being composed of these valence quarks, it also includes virtual quarks and antiquarks, sometimes referred to as *sea* quarks. In order to predict the probability of a hard-scatter process, it is necessary to have *parton distribution functions* (PDFs), which are measurements of the probability of an interaction with a particular particle inside the proton. This is typically characterized as a function of x , the fraction of momentum carried by that particle, and Q^2 , the momentum scale of the proton.

2

Jets and Jet Algorithms

Unlike other forces in the Standard Model, the strong force increases in strength over long distances. In fact, bare quarks and gluons have never been observed, which is believed to be a result of *color confinement*, where as colored objects become separated, they have sufficient energy to create quark-antiquark pairs, which are eventually turned into color-neutral hadrons. This means that instead of detecting individual quarks and gluons, collider experiments instead detect *jets*, which are collimated sprays of particles that roughly correspond to the radiation from an individual parton. There are several consequences of this in high energy experimental physics. There are several scales associated to any process involving quarks, since in addition to the parton interactions, there is some fragmentation process that turns these into color-neutral objects. Second, this also implies that a quark or gluon does

not necessarily result in the production of a single hadron, and most likely will produce several hadrons during the fragmentation. In order to probe interactions between colored objects, it is critical to have a meaningful proxy for these objects that can be observed experimentally. This is typically done using *jets*, which may be defined in several different ways, and which cluster groups of color-neutral objects together to act as a proxy for the initiating parton. The processes that contribute to jet formation as well as the different algorithms that may be used to define a jet are the subject of this chapter.

2.1 PARTON SHOWERS AND HADRONIZATION

The first stage of jet formation is the *parton shower*, where quarks and gluons from the hard-scatter event can radiate gluons, and gluons can split into a quark-antiquark pair. The cross-section for the radiation of a gluon diverges for very soft or collinear gluons. While this behavior is not experimentally observable, it does require additional treatment in calculations and models of the parton shower.

In Monte Carlo simulations, the behavior of the parton shower may be modeled using an ordered shower, which can avoid the problems of divergences as well as issues in double-counting emissions. Typically, parton showers are either angular ordered, or transverse-momentum ordered. In angular ordering, the widest emissions occur first, and as the shower progresses, the emissions tend towards smaller angles. In transverse-momentum ordering, the highest p_T emissions occur first.

The parton shower does not produce color-neutral particles, and the evolution from these particles to color-neutral hadrons occurs during the hadronization stage. This evolution is non-perturbative, meaning that it cannot currently be described using perturbative QCD. Instead, several different algorithms are used to model its behavior. Two main models exist to describe this behavior: the Lund string model and cluster hadronization. Both of these models do a reasonable job of modeling the dynamics of jets, though they come with different benefits.

One of the most popular fragmentation models is the Lund string model, which is used by Pythia. In this model, a flux tube exists between the quarks in an event. Since the gluon couples to itself, the field between the quarks will be localized in space. As they move apart, the energy in this string increases with the distance between the ob-

jects. At any given point, the string has some probability to break at a *vertex* into a quark-antiquark pair, where the strength of the field increases linearly with time. This process continues until only on-shell hadrons exist¹¹.

An alternative approach to hadronization is cluster hadronization, which is used by the Herwig event generator. This model is based on the idea of preconfinement, in which color singlets clusters are formed that decay into the final state hadrons. These clusters are formed from the quarks and antiquarks produced in the parton shower (after all gluons have decayed to quark-antiquark pairs), and tend to be fairly localized and of similar momenta. The mass distribution of these clusters tends to be fairly small and is relatively insensitive to the scale of the originating hard process. The clusters are then *fragmented* into the final-state hadrons based on several assumptions. Low-mass clusters undergo isotropic two-body decay, with equal probability for all quark flavors, and the decay products are tested against the phase space available for this decay. Clusters with mass above some scale M_f first undergo fission into quark-antiquark pairs, which are then decayed. In the case of a cluster with a heavy quark, the quark is decayed until it no heavy flavors remain. Any unstable resonances are then decayed based on information in the Particle Data Tables¹⁰⁷.

2.2 JET ALGORITHMS

A *jet* is roughly defined to be the collimated spray of partons produced by a particle during its fragmentation. However, even the definition of what constitutes a particle is ambiguous. In the soft collinear limit, such as a quark emitting a soft collinear limit, the result is experimentally the same as a single particle. In the hard limit, such as a hard gluon emission at a wide angle, is not fundamentally different than when this process happens at tree-level, and so it should be considered as two separate particles. There is a whole realm in between, which makes it impossible to have a universal definition of a jet.

It is, however, still possible to create a definition of a jet. Based on the work from the *Snowmass accord*⁷⁸, in order to be theoretically robust and experimentally tenable, an algorithm used to define a jet must satisfy several criteria. These properties include the ability to implement it easily, both experimentally and theoretically, and also require a few theoretical properties. This includes that the jet definition must be defined at any order of perturba-

tion theory and also yield a finite cross-section at any order of perturbation theory. Finally, it must be relatively insensitive to hadronization, in order to avoid large corrections from a theoretical calculation to an experimental observable⁷⁸. These criterion imply that a jet definition must be infrared and collinear safe (IRC), meaning that a collinear splitting or addition of a soft emission must not change the set of hard jets that are created in the event.

Most jet definitions used by modern experiments are part of a class of sequential recombination algorithms. These algorithms use a distance metric d_{ij} to recursively combine particles together, where there are several different distance metrics defined below. The sequential recombination algorithm proceeds as follows:

1. Find the distance d_{ij} between all pairs of particles i and j .
2. Find the pair of particles i and j which minimize this distance.
3. If this minimum distance $d_{ij,min}$ is less than a threshold cut d_{cut} , then combine i and j into a single *pseudo-jet*.
4. If no pairs of particles meet this criterion, the algorithm terminates.

The three main sequential algorithms use the distance metric

$$d_{ij} = \min(p_{T,i}^{2p}, p_{T,j}^{2p}) \frac{\Delta R_{ij}^2}{R^2}, \quad (2.1)$$

where R is the jet radius, and ΔR_{ij} is the angular distance between the different particles, and the value of p determines the algorithm type. The k_t algorithm uses $p = 1$, meaning that the soft particles are clustered together first, while the anti- k_t algorithm uses $p = -1$, meaning the hard radiation is clustered together first. Similarly, the Cambridge-Aachen algorithm takes $p = 0$, meaning that only the angular distance between particles is considered in the clustering algorithm. The anti- k_t algorithm has the benefit of creating relatively circular jets, making it experimentally convenient for calibration. Theoretically, the C/A and k_t algorithms provide a nicer definition, as they are more closely related to the showering¹⁰¹.

2.3 JET GROOMING

This method of jet reconstruction is sufficient for reconstructing a jet that is at a similar energy scale as the initiating parton. For this type of reconstruction, typically a fairly small jet radius R , such as $R = 0.4$ is used, which con-

tains the majority of high-energy radiation, while not large enough to contain large amounts of other uncorrelated energy from things such as the underlying event or pileup. However, the internal structure of a jet is important for several things, in which case a larger jet radius is important for containing more of the shower. For hadronic decays of heavy boosted objects, the decay products can be reconstructed into a single large- R jet, where the necessary size of the radius depends on the Lorentz boost of the decaying particle. For these, the internal structure of the jet can be used to separate these jets from the more common jets initiated by light quarks and gluons. Additionally, the internal structure of the jet can be used to obtain a more detailed understanding of jet evolution, providing a more fundamental understanding of QCD as well as for tuning Monte Carlo generators.

Large- R jets cluster more uncorrelated energy into the jet, which decreases the power of using substructure. This uncorrelated energy includes pileup, which adds energy roughly uniformly throughout the event, as well as the underlying event, which has a similar effect. In addition, the hard boundary of a jet creates sensitivity to non-global logarithms, or radiation that leaves, and then re-enters a jet, which complicates the theoretical interpretation of a jet. In order to account for these effects, several *grooming* algorithms have been developed, which preferentially remove radiation that does not originate from the initiating particle. Each algorithm accounts for these effects in a different way, and as a result, the optimal configuration depends on the type of physics being explored. Several of these algorithms are described below, and a comparison of these algorithms will be shown in Section 6.3, and a measurement using one of the Soft Drop grooming algorithm is one of the main topics of this thesis.

2.3.1 TRIMMING

The trimming algorithm⁸⁰ identifies the regions of a jet that contain the bulk of the energy, and groom away constituents that are outside of these regions. It achieves this by reclustering the constituents of a large- R jet into subjets of radius $R_{\text{sub}} < R_{\text{jet}}$ using the k_t algorithm. and discarding any subjet with p_T less than $f_{\text{cut}} \times p_{T,\text{jet}}^{\text{ungroomed}}$. This gives two parameters to the algorithm: R_{sub} and f_{cut} , which can be tuned for optimal performance, and Figure 2.1 demonstrates an example of its behavior. This is an effective grooming algorithm, since it is able to retain most of the activity clustered around hard radiation within the jet, even for jets with multi-pronged decays, but

also reduces the catchment area of the jet, reducing the sensitivity to pileup contamination and the underlying event. This grooming algorithm with $R_{\text{sub}} = 0.2$, $f_{\text{cut}} = 5\%$ is the standard for large- R jets on ATLAS, based on an optimization done during Run 1²⁰, and has been found to improve the pileup stability and increase the mass resolution for large- R jet reconstruction in ATLAS.




Figure 2.1: A cartoon demonstrating the trimming algorithm²⁰

2.3.2 PRUNING

The pruning algorithm⁶⁵ takes an alternative approach, where it instead modifies the jet clustering sequence by discarding soft and wide-angle constituents, which are more likely to be from pileup or the underlying event. This is done by reclustering the constituents of a large- R jet using a modified version of the Cambridge-Aachen algorithm⁶², where at each step, if the clustering is either too soft or too wide-angled, the cluster with the lower p_T is discarded. The Cambridge-Aachen algorithm is used in order to achieve an angular-ordered cluster sequence, which is more meaningful than the anti- k_t algorithm used to determine the set of constituents in a jet. There are two pruning conditions that must be satisfied,

1. Wide Angle : $\Delta R_{ij} \geq R_{\text{cut}} \times 2 \frac{M_{ij}}{p_{T,ij}}$
2. Soft : $\frac{\min(p_{T,i}, p_{T,j})}{p_{T,i} + p_{T,j}} \leq z_{\text{cut}}$

where ΔR_{ij} , M_{ij} , and $p_{T,ij}$ are respectively the angular distance, mass, and transverse momentum of the cluster combination. An example of this algorithm applied to constituents is shown in 2.2, which demonstrates the clustering history as pairs of constituents either pass or fail the pruning algorithm.




Figure 2.2: An example of the pruning algorithm, where each square represents a constituent with p_T proportional to its size. (a) The unclustered constituents, (b) the first pair of constituents, which fail the pruning algorithm due to the soft condition, (c) the new set of constituents after discarding the softer constituent in the first splitting, (d) a pair of constituents that pass the pruning conditions, (e) another pair of constituents that pass the pruning conditions, (f) a pair of constituents that fail the wide-angle condition, and (g) the final jet with its clustering history.

2.3.3 SOFT DROP AND THE MODIFIED MASS DROP TAGGER

The Soft Drop grooming algorithm⁸³ acts on the clustering history of a sequential recombination jet algorithm⁶⁴.

In this algorithm, the constituents of the large- R jet are reclustered using the Cambridge-Aachen algorithm, creating an angular ordering of the clustering history. Next, the last step of the clustering is undone, and it is checked

whether the two proto-jets satisfy the following condition:

$$\frac{\min(p_{T,j_1}, p_{T,j_2})}{p_{T,j_1} + p_{T,j_2}} > z_{\text{cut}} \left(\frac{\Delta R_{12}}{R} \right)^\beta, \quad (2.2)$$

where p_T is the transverse momentum of the jet transverse, z_{cut} and β are algorithm parameters, and ΔR_{12} is the angular distance between the proto-jets. If the soft-drop condition in Eq. (2.2) is not satisfied, then the branch with the smaller p_T is removed. The procedure is then iterated on the remaining branch. Once the condition has been satisfied, the algorithm terminates, and the remaining constituents form the jet. Diagrams showing the different stages of the grooming algorithm are shown in Figure 2.3. The parameter z_{cut} sets the scale of the energy removed by the algorithm, while β tunes the sensitivity of the algorithm to wide-angle radiation. As β increases, the fraction of branches where the condition is satisfied increases, reducing the amount of radiation removed from the jet, and in the limit $\beta \rightarrow \infty$, the original jet is unmodified. It is noted that Soft Drop with $\beta = 0$ reduces to the modified Mass Drop Tagger algorithm, which is an algorithm that preceded the Soft Drop algorithm. However, for simplicity, this will still be referred to as Soft Drop in this thesis. The Soft Drop algorithm has several nice theoretical properties, which are discussed at greater length in Section 7, motivating its use in measurements of jet substructure observables.

2.3.4 RECURSIVE SOFT DROP

The standard Soft Drop algorithm aims to find the first hard splitting in the jet clustering history in order to define a groomed jet. In the case of a multi-pronged decay, this condition may not be sufficient to remove much of the uncorrelated energy in the jet, since the condition may be satisfied before removing all of this energy. A recent extension of the Soft Drop algorithm has been proposed, named *Recursive Soft Drop* (RSD), in which the algorithm continues until it has found N hard splittings in the clustering history. In the case of $N=1$, this reduces to the standard Soft Drop algorithm, while for larger values of N , this allows for more of the jet to be removed by grooming.




Figure 2.3: The stages of the soft drop grooming algorithm. (a) A jet clustered with the anti- k_t algorithm. (b) The same constituents reclustered with the C/A algorithm. (c) The first iteration of the soft drop algorithm, which removes the j_2 branch. (d) The second iteration of the soft drop algorithm, which finds a splitting that satisfies the soft drop condition.

2.3.5 BOTTOM-UP SOFT DROP

Finally, the bottom-up Soft Drop (BUSD) algorithm uses the Soft Drop condition, but applies it during jet clustering, in a way similar to the pruning algorithm. This algorithm may be applied to either the full event or to in-

dividual jets, but results shown in this thesis only consider a jet-by-jet grooming. In this algorithm, jets are reconstructed with the anti- k_t algorithm, and then reclustered using a modified version of the Cambridge-Aachen algorithm, where particles i and j with the smallest distance $d_{ij} = \Delta R_{ij}/R_0$ are combined to create a new pseudojet given by

$$p_{ij} = \begin{cases} \max(p_i, p_j), & \text{if the Soft Drop condition fails,} \\ p_i + p_j & , \text{otherwise.} \end{cases} \quad (2.3)$$

This means that if the Soft Drop condition is not met, then only the hardest constituent is used for the clustering history.

3

The Large Hadron Collider

The Large Hadron Collider (LHC)^{43,44,39} is located at the CERN laboratory, and is the largest existing proton-proton synchrotron, designed to accelerate protons up to 7 TeV. Four experiments are located at the collision points that are located around the 27 kilometer ring. Two of these detectors, ATLAS¹⁴ and CMS⁵⁵, are general-purpose detectors, whose different designs enable them to provide a cross-check of results. One of the other detectors, ALICE¹, is designed for the purpose of studying heavy-ion collisions. The last experiment, LHCb⁸, is designed to study CP-violation by using mesons that decay to b-quarks. The studies presented in this thesis are done using data taken with the ATLAS experiment, which is described in detail in the next chapter.

The protons are accelerated in several steps. The protons are taken from hydrogen atoms that have been ionized

in an electric field, and subsequently accelerated to 50 MeV in a linear accelerator, the LINAC 2. These protons are fed into the BOOSTER, where they are accelerated to 1 GeV. These are accelerated to 26 GeV in the Proton Synchrotron, and then to 450 GeV in the Super Proton Synchrotron, which is sufficient energy for stable beams in the LHC, where they are accelerated to their final energy of 6.5 TeV. Figure 3.1 shows the different accelerators involved in this process, as well as other accelerators present at CERN that are used for other experiments¹⁰².




Figure 3.1: Diagram of the LHC injection chain, as well as some of the relevant experiments at CERN

The protons being collided at the LHC are separated into *bunches*, which are kept together using the radio frequency acceleration. There are a 35640 buckets that can be filled with protons at the LHC, with a spacing of around 25 ns, and a maximum of 2808 regions may be filled⁴³. Typically, the bunches are organized into *bunch trains*, which are groups of these regions that are all filled, with some empty buckets between each bunch train. The exact configuration of these trains varies with the run.

The rate of data taking, or *instantaneous luminosity* can be calculated by the simple equation

$$L = \frac{f_{\text{rev}} b N^2}{A}, \quad (3.1)$$

where f_{rev} is the accelerator frequency, b is the number of colliding bunches, N is the number protons per bunch, and A is the transverse cross section of the overlap between the colliding beams. Thus, in order to increase the rate of data taking, with a fixed bunch configuration and revolution rate, either the number of protons per bunch may be increased, or the beam size can be decreased. In order to know the absolute luminosity of data taken, it is critical to have a measurement of the instantaneous luminosity, requiring measurements of the beam area and the number of protons. Typically, the beam area is measured using Van der Meer scans, where the rate of collisions is measured as a function of the distance between beams.

4

The ATLAS Detector

The ATLAS experiment is an general-purpose detector designed to be able to detect quasi-stable particles produced in collisions from the LHC, with the exception of neutrinos. This includes electrons, muons, taus, photons, and hadrons. Its design follows that of modern general-purpose detectors, with several concentric layers of specialized sub-detectors. The ATLAS experiment is composed of a *barrel* region, composed of a set of cylindrical detector systems, which is sandwiched by the *endcap* detectors, or a set of flat, circular detector systems, as depicted in Figure 4.1. The innermost layers are the tracking detectors, which measure the trajectory of charged particles without greatly affecting their momenta. The tracking system is encased in a 2 Tesla solenoid magnet in order to be able to measure their charge and momentum. Outside of this are the calorimeters, which measures the energy

of all particles (except muons), and are sufficiently dense to stop nearly all particles. Finally, the outermost layer is composed of toroidal magnets and the muon spectrometer, which are designed for precise measurements of muons, whose interactions are small enough that they can travel through the calorimeter. This design allows for the measurement of a wide range of particles, making it useful for everything from measurements of the Standard Model to searches for particles beyond the Standard Model. Each of the individual components of the detector will be described briefly in the following sections, with a focus on the detector systems most relevant for the specific analyses in this thesis.

There are several different coordinate systems used to describe the ATLAS detector. In general, ATLAS uses a Cartesian right-handed coordinate system, where the origin is at the nominal interaction point, the z-axis points along the beampipe, the y-axis points towards the surface of the earth, and the x-axis points towards the center of the LHC ring. Based on the cylindrical nature of the detector, a cylindrical coordinate system (r, ϕ, z) is frequently used, where r is the radial distance from the z-axis, ϕ is the azimuthal angle around the beampipe. Finally, a coordinate called the pseudorapidity (η) is frequently used, where $\eta = -\ln(\tan(\theta/2))$, and θ is the polar angle. Pseudorapidity has several convenient properties for collider experiments. For massless particles, differences in pseudorapidity are invariant under longitudinal boosts. Additionally, the production of pions is approximately flat as a function of η , meaning that in high luminosity, the energy density will be roughly flat as a function of η . Each of these different coordinate systems will be used throughout this thesis.

4.1 THE INNER DETECTOR

The inner detector (ID) of ATLAS is designed to reconstruct tracks produced from charged particles passing through the detector. It consists of three layers: the innermost pixel layers, the middle SemiConductor Tracker (SCT), and the outermost Transition Radiation Tracker (TRT), which are depicted in Figure 4.2. In order to be able to measure the momentum and charge of the particles, the tracking system is surrounded by a strong, solenoidal magnetic field, which bends the trajectory of the charged particles. These detectors provide tracking information for particles within $|\eta| < 2.5$ for tracks with p_T down to 0.5 GeV. To balance accurate track reconstruction with the cost of the




Figure 4.1: Diagram of the ATLAS detector

detectors, the innermost detector systems have a higher granularity than those that are farther from the beamline. Each of these detectors is described in detail in the following sections.

4.1.1 THE PIXEL LAYERS

The innermost layer of the tracking system is the pixel detector, a type of semiconducting silicon detector, covering the region $|\eta| < 2.5$. The pixel detector is designed to be a high-resolution, radiation-hard detector that is able to aid in primary vertex reconstruction, precise transverse impact parameter reconstruction, and b-tagging. In order to be able to perform these tasks, the original pixel detector was designed with three detector layers in the barrel region, and three disk layers in the endcap, as shown in Figure 4.2. Each layer is composed of several pixel modules, each of which contains 46,080 silicon sensors of size $50 \mu\text{m} \times 400 - 600 \mu\text{m}$ in order to provide precise position information. In order to provide full coverage in ϕ , the barrel modules are tilted with respect to the cylinder and overlap slightly with the next, while the the endcap modules are tilted with respect to the plane perpendicular to the beam. Signals in the sensors are amplified and compared to a threshold, and the time-over-threshold is stored for possible readout is the event is triggered. The thresholds may be tuned as the radiation does increases. In ad-




Figure 4.2: Diagram of the layout of the inner detector

dition, the pixel sensors were designed to be radiation tolerant, with n^+ implants on the readout side, and a p^+ - n junction on the back. As the sensors undergo radiation, the bulk type inverts, and still allows for some sensitivity⁶⁶. In order to improve the overall precision of tracking and vertexing, and to protect against failure of the first layer of the pixel detector, another layer was added in 2014, named the Insertable B-Layer (IBL), which is located between the beampipe and the first layer of the original pixel detector⁵⁰.

4.1.2 THE SEMICONDUCTOR TRACKER

The SemiConductor Tracker (SCT) is located outside of the pixel detector, meaning that it needs to fulfill less stringent requirements, both in terms of resolution as well as radiation sensitivity. Instead of using expensive rectangular pixel detectors, the SCT uses strips detectors, providing a compromise between accuracy and cost. Each SCT module consists of two sets of two silicon-strip sensors which are glued together with a stereo angle of 40 mrad between them in order to gain 3-dimensional hit information. In the barrel region, these are positioned approximately parallel to the magnetic field and beampipe. In order to be robust against radiation damage, it uses

the same type of doping as in the pixel layers. There are four radial layers of SCT modules in the barrel region and nine layers per endcap, as shown in Figure 4.2, making a total of nearly 16,000 sensors⁵⁶.

4.1.3 THE TRT

The outermost layer of the tracker is the Transition Radiation Tracker (TRT), where the name comes from its ability to identify electrons using their transition radiation in the form of x-ray photons, or radiation from charged particles passing through different materials. In order to do this, it was designed to be made with a xenon-based gas mixture (Xe(70%)CO₂(27%)O₂(3%)), whose high density gives it a long absorption length¹², which is ideal for detecting this type of radiation. However, due to large gas leaks in parts of the detector, some parts of the TRT are instead filled with an argon mixture⁹⁶.

The TRT is composed of around 300,000 straw tubes, each with a diameter of 4 mm, with a 31 μm gold-plated tungsten wire in the center. In the barrel region, the tubes are parallel to the beamline, where they cover up to $|z| < 720\text{mm}$. In the endcaps, the wires are perpendicular to the beamline, and cover $827\text{ mm} < z < 2774\text{ mm}$, and $617 < R < 1106\text{ mm}$ ⁹⁶. Since each tube spans a fairly large region, the occupancy becomes high under high pileup conditions, and in some cases, the occupancy can be greater than 50%¹⁰⁶.

4.1.4 THE SOLENOID

The tracking detectors are surrounded by a solenoid, which provides the 2 T magnetic field enabling the measurement of the transverse momentum of particles. A couple of considerations were taken into account when designing the solenoid. Since the magnet is in front of the calorimeters, it is important to minimize the amount of material in order to get the best possible energy measurement. Additionally, stronger magnetic fields are critical, since track momentum resolution increases proportionally to the strength of the field. With these constraints, the solenoid was designed to provide a 2 T axial field, with material that contributes around 0.66 radiation lengths¹⁴.

4.2 THE CALORIMETERS

The ATLAS calorimeters are designed to measure the energy of particles passing through them, with the exception of muons and neutrinos. This is done by stopping particles within the calorimeter, and measuring the energy of their showers. This needs to be possible for both the electromagnetic (EM) showers produced from photons and electrons, as well as the hadronic showers produced from hadrons. Since these showers differ in shape, composition, and depth, the ATLAS detector has two calorimeters: an inner electromagnetic calorimeter designed to measure EM showers, which are absorbed over a shorter scale, and an outer hadronic calorimeter designed to absorb hadronic showers.

The basics of EM and hadronic showers, as well as the fundamental principles of calorimetry are important to understanding the design of the ATLAS calorimeters. Electromagnetic showers are formed by photons and electrons when they interact with material. Generally, photons interacting with material can produce an electron-positron pair, while electrons may undergo bremsstrahlung and emit a photon. The typical scale associated to this showering is characterized by the radiation length. For electrons, the radiation length is the average distance it must travel such that it has lost all but $1/e$ of its initial energy to brehmstrahlung, while for photons, it is $7/9$ of the average distance it travels before pair production.

Hadronic showers are more complicated than EM showers. Not only do they involve more types of particles, but the types of interactions occurring are also more complicated. Typically, they result from an incoming hadron interacting with a nucleus, which creates a shower of hadrons. These may go on to produce cascading showers, or in the case of $\pi^0 \rightarrow \gamma\gamma$, they may produce EM showers. In addition to the complication of multiple types of particles, effects such as nuclear binding energy can result in *invisible energy*. With all of these factors at play, hadronic showers have much larger variations in their shape than EM showers. Since they rely on nuclear interactions, they also typically develop over a longer length than EM showers.

Both the EM and hadronic calorimeters on ATLAS are sampling calorimeters. The basic principle of these is that they are composed of alternating layers of a dense material to produce showering, and an active material to measure the showers. There are four separate calorimeters on ATLAS, which are described below.




Figure 4.4: A diagram of the signal shape in the LAr calorimeters

proximately cancel on average, since the downward fluctuations from nearby bunch crossings (out-of-time pileup) will tend to cancel the upward fluctuations from pileup within that bunch crossing (in-time pileup). The ADC counts from 4-5 samples S_{sample} at different time slices N are translated into a calibrated energy using the formula

$$E_{\text{cell}} = \frac{1}{f_{I/E}} \times F_{\text{gain}} \left(\sum_{\text{sample}=1}^N OF_{\text{sample,gain}} (S_{\text{sample}} - P_{\text{gain}}) \right). \quad (4.1)$$

With this equation, the pedestal for each gain P_{gain} is subtracted from the ADC sample, and these values are multiplied by the optimal filtering coefficients (OFCs) $OF_{\text{sample,gain}}$, which have been determined based on the pulse shape. This is then converted into a current in μA using F_{gain} , which is then converted into an energy with $f_{I/E}$. The number of samples changed from 5 to 4 in the middle of Run-2 of data taking, though up to 32 samples may be read out for calibration purposes⁸⁷.

4.2.2 THE HADRONIC BARREL (TILE) CALORIMETER

The Hadronic Barrel (Tile) calorimeter is composed of the barrel region, covering the region $|\eta| < 1.0$, with two extended barrels covering the region $0.8 < |\eta| < 1.7$. It is composed of cells of size 0.1×0.1 ($\Delta\eta, \Delta\phi$) in the barrel region, and 0.2×0.1 in the extended barrel region, with each cell being composed of alternating layers of iron plates and scintillating tiles, as shown in Figure 4.5. The tiles measure the scintillation light produced by the showers. Like the EM calorimeter, the tile calorimeter provides several radial measurements, and is designed to have a




Figure 4.5: Diagram of a tile calorimeter cell¹⁴.

geometry to have approximately even coverage in slices of η , as seen in 4.6.

4.2.3 THE HADRONIC END-CAP CALORIMETERS

The Hadronic end-cap calorimeters (HEC) are designed to detect hadronic showers for $1.5 < |\eta| < 3.2$. Similarly to the central EM calorimeter, the active material is liquid argon, while the absorber is copper. Each HEC is composed of two wheels, where each wheel has longitudinal segmentation, with two longitudinal read-outs per wheel. Its material covers approximately 12 interaction lengths, which was chosen to fully contain jets produced at 14 TeV. Its granularity in η, ϕ is approximately $0.1 \times 2\pi/64$ for the region $|\eta| < 2.5$, and $0.2 \times 2\pi/64$ ⁷³, and the overall layout may be seen in Figure 4.7.

4.2.4 THE FORWARD CALORIMETERS

The Forward Calorimeter (FCal) covers the region $3.1 < |\eta| < 4.9$, and therefore needs to be able to cope with a high particle flux, since it is close to the beampipe. It is important to have some longitudinal information for the showers, and so it is built in three sections, as shown in Figure 4.8. The module closest to the beam is built with




Figure 4.6: Diagram of a tile calorimeter segmentation¹⁴.




Figure 4.7: Diagram of the layout of the hadronic end-cap calorimeters⁷³.




Figure 4.8: Diagram of the layout of the forward calorimeters.




Figure 4.9: Electrode structure of the forward calorimeter.

electrode rods and calorimeter matrix made of copper for the absorber, whose conductivity avoids the problem of overheating the LAr. The second two modules have tungsten electrode rods as absorbers with a matrix built of tungsten slugs, which are dense¹⁴. The structure of the electrodes may be seen in Figure 4.9.

4.3 THE MUON SPECTROMETER

The Muon Spectrometer (MS) is the outermost layer of the ATLAS detector, built primarily in order to detect muons in the central region. As the muon system is not used in the analyses presented in this thesis, it is described

briefly. It is able to measure particles with $|\eta| < 2.7$, and trigger on them for $|\eta| < 2.4$. Unlike the other systems described above, the muon system uses separate detectors for triggering and measuring muons. The triggering system is composed of Resistive Plate chambers in the barrel region, and Thin Gap Chambers in the endcap region, while more precise measurements are performed with the Monitored Drift Tube Chambers in the barrel region, and Cathode Strip Chambers in the endcaps. In addition, the muon spectrometer contains the toroid magnets, which consist of the barrel toroid and the endcap toroids. The barrel toroid is composed of eight coils, each 25.3 m long, and 10.7 m tall. Each endcap contains a toroid, in order to be able to measure muons in the forward region. These toroids are composed of square coils and wedges, which are glued together to produce a single rigid structure¹⁴.

4.4 TRIGGER AND DATA ACQUISITION

The ATLAS detector takes an enormous amount of data, with collisions every 25 ns, and not all of these can be fully processed or stored. Most of these events are soft collisions which are not of interest to the average physics analysis, and so collisions are filtered using the trigger system, which selects events various criterion. Since ATLAS is a multi-purpose detector, there are a wide variety of triggers available in order to be able to analyze as much relevant data as possible. There are three stages of the trigger: the Level 1 (L1) trigger which is implemented in the hardware for fast processing, and the Level-2 (L2) trigger, and the event filter (EF)¹⁴.

The L1 trigger is implemented in hardware, and must make decisions quickly in order determine which events should be buffered for testing by the high level triggers (HLT). As such, it only takes information from the calorimeter and muon spectrometer, for which it is possible to apply a meaningful reconstruction at this rate. Instead of reconstructing the complete objects that are used for electron, photon, jet, and MET reconstruction, the L1 calorimeter trigger relies on towers in $\eta \times \phi$, or even more granular at higher η , which are used as proxies for the inputs to these objects, while the muon spectrometer relies on having at least 3 hits in coincidence. In order to trigger an event with a jet, a region of interest (RoI) is identified based on 0.2×0.2 towers which have a high sum of E_T , and which are a local maximum of sum E_T . The size of the RoI ranges from 0.4×0.4 to 0.8×0.8 , depending on the relevant jet multiplicity, with higher multiplicity corresponding to a smaller RoI to provide higher granularity⁶.




Figure 4.10: The total integrated luminosity for 2015-2018.

The data passing the L1 trigger is designed to have a rate of around 75 kHz, which must then be reduced to around 200 Hz by the HLT¹⁴. The HLT is composed of the Level-2 (L2) trigger and event filter (EF), but they are often referred to together due to similarities in how they are run. Events are first passed through the L2 trigger, which reads out information seeded by the RoIs from L1, but with more complete information available. Events passing the L2 trigger proceed to the event filter, which is able to reconstruct the full event, albeit with some modifications to the various reconstruction algorithms. This stage is designed to be as close to the offline reconstruction as possible within the limitations of the processing speed. Events that pass the offline trigger are then stored permanently for full offline reconstruction.

4.5 DATA QUALITY MONITORING

Data-taking with the ATLAS detector has mostly proceeded smoothly, but the quality of the data for each detector subsystem is carefully monitored throughout data-taking. Data-taking runs where detectors are problematic for too much of the run are marked as bad, and are not used in standard physics analysis. This information is compiled to create a *good runs list* that lists all of the runs which are of sufficiently high quality to be used in analysis. The total integrated luminosity, both before any selection is applied, after the trigger, and for high-quality physics data is shown in Figure 4.10, where it can be seen that from 2015-2018, ATLAS has taken around 139 fb^{-1} of data which may be used in analysis.

5

Object Reconstruction

In order to use information from the detector for physics analysis, it is necessary to process detector signals into meaningful physics objects. The standard types of physics objects used on ATLAS represent electrons, photons, muons, jets, tracks, and missing transverse momentum. In general, while these objects are standardized, there are often multiple supported definitions for each object, which may be chosen in order to best fit the needs of an analysis. The two analyses covered by this thesis only rely on jets, tracks, and vertices, and the reconstruction of these objects is described below.

5.1 TRACK RECONSTRUCTION

Tracks are reconstructed in order to generically represent charged particles traveling through the detector, and are reconstructed using information from the inner detector layers described in Section 4.1. Charged particles follow a helical trajectory, and so tracks may be parameterized by their four-momentum, and several additional observables, as seen in Figure 5.1. The d_0 parameter, also known as the transverse impact parameter, describes the point of closest approach to the z-axis that run along the beamline. Similarly, the z_0 parameter, or the longitudinal impact parameter, describes the distance along the z-axis to the x-y plane at the same point of closest approach. Then, θ is given by the angle that the four-momentum of the track makes with respect to the z_0 axis at the point of closest approach, while ϕ is the angle that the track makes with respect to the x-y plane at the point of closest approach. Generally, particles produced near the interaction point will have small values for both z_0 and d_0 . In addition, particles emitted collinearly from these will also have small impact parameters, since they will still point approximately toward the same origin.




Figure 5.1: The parameterization of a track on ATLAS.

The tracking algorithms on ATLAS rely on the three tracking detectors described in Section 4.1. Assuming 100% efficiency, a generic charged particle passing through the barrel region of the inner detector will leave hits in 3 pixel clusters, 8 SCT clusters, and upwards of 30 TRT straws. More realistically, there will be some detector inefficiencies from hits below threshold, gaps in the detectors, and incorrectly assigned hits. The tracking algorithms attempts to balance these effects with the rate of fake track reconstruction in the tracking algorithm.

Standard track reconstruction is performed in two stages: first with an *inside-out* approach, and then with an *outside-in* approach. The *inside-out* approach uses hits from the inner silicon detectors to seed tracks, which are then extended outwards, while the *outside-in* approach uses the TRT to seed tracks, which are then propagated inwards. The inside-out tracking algorithm is first seeded by the two silicon detectors (the pixel and SCT detectors), where a seed requires three space-points from three separate layers of the silicon detectors¹⁶. Since the size of each pixel is sufficiently small, a single hit in the pixel detector qualifies as a space point, while the SCT uses a combination of the two strips that are glued at 40 mrad with respect to each other, as described in Section 4.1. This seed allows for a rough estimate of the track parameters described above, under the assumption helical tracks. These seeds are then used as inputs to a window search, which creates a road along the trajectory of the seed. Hits along this road are then be included in the track candidate using a Kalman filter¹⁶. The quality of these track candidates is then assessed based on the number of clusters, holes, track fit quality, and track momentum²⁹.

After these tracks have been created, the outside-in algorithm is run. The electromagnetic calorimeter is used to determine regions of interest in the TRT, and segments are created out of TRT hits. These segments are then extended into the silicon detectors, where unassociated hits may be added to the track candidate. These are added to the standard set of tracks. If no associated silicon hits are found, these tracks are still used for photon conversion reconstruction³⁵.

5.1.1 LARGE RADIUS TRACKING

The standard tracking definition is designed to reconstruct tracks originating from a primary vertex. However, the cuts applied to these tracks reduce efficiency for tracks originating from long-lived particles, which decay far from

the interaction point. This motivated the development of the large-radius tracking (LRT) algorithm³⁵, which is able to reconstruct tracks with loosened parameter choices. The overall algorithm for large radius tracking is nearly the same as that of the standard inside-out tracking algorithm, but relaxes several requirements, and is only run on hits that have not been used by the other tracking algorithms. Several parameters are changed, including the maximum d_0 and z_0 , which are loosened from 300 mm to 10 mm, and from 250 mm to 1500 mm respectively. Two silicon modules are allowed to be shared instead of one, and only five unshared silicon hits are required instead of six. Finally, the algorithm uses a sequential Kalman filter instead of a combinatorial one, which avoids the creation of multiple track candidates. This algorithm improves the reconstruction efficiency for displaced hadrons significantly for particles with a production radius of greater than around 20 mm, as shown in Figure 5.2, which shows the reconstruction efficiency in Monte Carlo for long-lived gluinos in a split SUSY model that decay hadronically. The overall reconstruction efficiency is defined as ratio of truth particles from the gluino decay that have an associated reconstructed track to the total number of truth signal particles.




Figure 5.2: Inclusive track reconstruction efficiency for displaced hadrons as a function of the radius of production r_{prod} ³⁵.

It is clear that the large-radius tracking algorithm significantly increases the reconstruction efficiency for tracks from long-lived particles. Especially for very long lifetimes, some inefficiency is still observed, with only around 40% of tracks being reconstructed for particles with a radius of production above 250 mm. This inefficiency is a combination of several effects, including the fiducial volume, as well as the actual inefficiency of the algorithm. In order

to study the performance of the algorithm for displaced hadrons that fall within the fiducial volume, the *technical efficiency* may be defined. This determines the efficiency for truth particles associated to these displaced hadrons that would be expected to be reconstructed based on the detector layers that they pass through. This eliminates the effect of particles that would not pass through enough silicon layers to be reconstructed. In order to pass the basic requirements, a truth particle must have $p_T > 1$ GeV and be produced with $r_{prod} < 300$ mm, and $|\eta| < 5$, and it must deposit energy in at least 7 silicon layers. The technical efficiency is then the ratio of truth particles that are associated to a reconstructed object to all truth particles that pass this selection. The result of this may be seen in Figure 5.3, which demonstrates an efficiency of at least 80% for particles that are produced up to 300 mm.

Due to the relaxed requirements, large-radius tracking requires a large amount of computing power. As a result, these objects are not available by default, and this algorithm is only run on selected events where long-lived particles are likely to be relevant. The implications of this are discussed in greater detail in Section 14.1.




Figure 5.3: The technical reconstruction efficiency of signal particles as a function of the radius of production r_{prod} ³⁵.

5.2 VERTEX RECONSTRUCTION

Reconstructed tracks are the basis of vertex reconstruction, where vertices link groups of tracks with a common origin. Two types of vertices exist: primary vertices, which correspond to hard-scatter events, and secondary vertices,

which correspond to interactions and decays of particles produced in the hard-scatter events. The goals of both of these types of vertices is very different, and consequently, the algorithms used to reconstruct them also differ, as described in the following sections.

5.2.1 PRIMARY VERTEX RECONSTRUCTION

In order to determine the location of the hard-scatter collision of interest, as well as pileup interactions, it is critical to be able to reconstruct *primary vertices* (PVs). In general, the aim is to reconstruct one vertex per interaction, meaning that the amount of pileup may be characterized by the number of reconstructed primary vertices (N_{PV}). These PVs are reconstructed out of tracks fulfilling the following quality requirements³⁰:

- $p_T > 400$ MeV
- $|\eta| < 2.5$
- $N_{SIhits} > 9$ (if $|\eta| \leq 1.65$), or 11 if $|\eta| > 1.65$
- $N_{IBL+B-layerhits} \geq 1$
- At most one shared module
- No pixel holes
- At most 1 SCT hole

Tracks fulfilling these requirements are then used as inputs to PV reconstruction. A set of primary vertices is found by seeding the vertex position based on the global maximum of the z-position of all the input tracks. Then, the vertex position is determined using an adaptive fitting algorithm, which minimizes the χ^2 of the fit, and progressively down-weights poorly-matched tracks in each step. After the position has been fit, any tracks that are incompatible with the reconstructed vertex by more than 7σ are removed and used to create a new vertex using this same method. This algorithm continues until all tracks are included in vertices, or no new vertices may be found. Each vertex must have at least two associated tracks to be considered. After creating the set of primary vertices, their positions are refit using the original reconstructed position and the width of the beamspot¹⁵.

Using this set of vertices, it is possible to determine the interaction that a track is associated to. The vertex associated to the hard-scatter event is typically the one with the most momentum associated to it, while the other ones are typically from pileup interactions. The primary vertex association of tracks enables the removal of tracks originating from a pileup interaction.

5.2.2 SECONDARY VERTEX RECONSTRUCTION

In order to identify decays of long-lived particles, it is possible to reconstruct secondary vertices. These follow a different reconstruction algorithm than the primary vertex reconstruction, since the constraints are much different.

The set of secondary vertices is reconstructed out of the set of tracks passing the following preselection:

- The track has not been used for primary vertex reconstruction
- $p_T > 1 \text{ GeV}$
- $N_{\text{hits,SCT}} > 2$

Secondary vertex reconstruction is then applied using a multi-stage approach. First, the set of all compatible two-track vertices are found, where each track may be included in multiple vertices. In this stage, all pairs of tracks are formed into vertices, with the vertex position approximated by the track parameters. If the $d0$ or $z0$ of the tracks are more than 50 or 100 mm from the estimated vertex position, the vertex seed is thrown away. Then compatible pairs of tracks have their position estimated more precisely using a full track extrapolation. Vertex seeds are kept if this fit succeeds, $\chi_{vtx}^2/n_{dof} < 5$ and $r_{vtx} < 563 \text{ mm}$. These seeds are further filtered based on the hit pattern of the associated tracks. Since tracks typically travel outwards from the vertex, tracks are required to have no hits in tracker layers with smaller radii than the vertex position. For vertices that are very close to the outside of a tracker layer, this requirement is relaxed by one tracking layer. Additionally, tracks are required to have a hit in the first tracking layer beyond where the vertex is located. In the case of disabled pixel modules, which do not produce any hits, this requirement is automatically fulfilled to avoid inefficiencies.

The two-track vertex seeds are then used to create n-track vertices, using a graph that determines compatibility of each pair of tracks. Each n-track vertex is formed out of tracks where each pair of tracks in the n-track vertex is one of the 2-track vertex seeds. This may still leave some tracks associated to multiple vertices. In these cases, the goodness-of-fit is evaluated for this track in each vertex, and it is included the best-fitting vertex.

After the creation of the preliminary set of n-track vertices, additional steps are run to improve the reconstruction by merging compatible vertices together, and associating additional tracks that were not originally included.

The initial algorithm may reconstruct a single decay as multiple vertices. This includes the reconstruction of two nearby vertices instead of one, or a misreconstruction of the vertex position from tracks that cross multiple

times. Additionally, some tracks may not be included in the original vertex, due to effects like a backscatter that fails the hit pattern requirement, or a track not passing the original preselection. To account for these effects, several algorithms are run, first to attempt to merge vertices together, and second to associate additional tracks to these vertices. In order to determine that vertices to consider merging, the *distance significance* S is defined, where

$$S = (\vec{v}_2 - \vec{v}_1)(C_1 + C_2)(\vec{v}_2 - \vec{v}_1)^T, \quad (5.1)$$

where $\vec{v}_i$ is the position of vertex i , and C_i is the covariance matrix for vertex i . The merging algorithms are run in several stages on vertices whose distance significance is less than 100, where The merging algorithms are run sequentially, starting with vertices with the smallest number of tracks, which are most likely to be mismeasured. If a given merging algorithm succeeds, no other vertexing algorithms are tested for that vertex.

The first merging algorithm, *merge by shuffling*, tests if the position of the original vertex V_{original} was slightly misreconstructed. Looping through all other vertices in the event, starting with those with the most associated tracks, the algorithm refits the vertex V_{original} with the position of another vertex V_{test} , resulting in a new position V'_{original} . If the distance between the suggested position and the refitted position is less than 10, then the vertices are merged together.

Next, the *magnetic merging* algorithm is tested. Again, looping through all other vertices in the event, a track is borrowed from a target vertex and added to the original vertex. The original vertex is then refit with the position of the vertex where the track has been taken from. If the distance significance between the refit vertex and original vertex is the minimum, and is less than 10, then the vertices are merged together.

Finally, if the other strategies have not resulted in a merged vertex, the algorithm attempts to force merging between vertices with the *wild merging* algorithm. In this, all tracks are combined into a vertex, and the algorithm attempts to fit these with the position of the original vertex. The pair of vertices with the minimum distance significance between the new vertex and the original vertex is found. If this distance is less than the 10, then the vertices are merged together.

Finally, any vertices whose distance is less than 2mm are merged automatically. If any of the merging steps

have succeeded, the vertices are merged together, and the merged vertex is run through a function to improve the χ^2 . In this step, the vertex position is first refit, seeded by its original position. Then, while its χ^2 is above some threshold and $N_{\text{tracks}} > 2$, the algorithm attempts to remove single tracks from the vertex. If this improves the overall fit quality of the vertex, then the track is removed. This proceeds until the vertex has two tracks or until the fit quality is sufficiently high.

After this is completed, some tracks may still not be associated to a vertex, and an algorithm is run that allows for the association of these tracks to existing vertices. Vertices may have a maximum of one track associated to them by this algorithm. This process is only done for good quality tracks whose χ^2 per degree of freedom is less than 5 and $p_T > 1$ GeV. Additionally, the maximum d_0 and z_0 significance must both be less than 5. In order to avoid contamination from primary vertices, the association is only done for vertices that are at least 0.5 mm away from any primary vertex. The track is not required to pass the same hit pattern requirements that are used to create the vertex, in order to include back-scattered tracks. The association algorithm attempts to add each track to a vertex, starting with the highest n-track vertices. If the association succeeds, the track is added to the vertex and removed from the set of tracks to be considered. Once all vertices are created, merged, and associated, the vertex position is refit.

VERTEX CLEANING CUTS

In the initial processing of the data, an asymmetry in ϕ for vertex production was observed, which can be seen in Figure 5.4. Generally, a ϕ asymmetry indicates a problem with the reconstruction, since the detector is approximately symmetric in ϕ . Further studies demonstrated that this asymmetry only existed in vertices with an associated track, and that these tracks tended to have a d_0 value that corresponds to one of the pixel layers. The problem has now been understood to be an issue with the track extrapolation, but this solution was not implemented in time for these studies. Instead, these vertices are preferentially removed using cuts that target the problematic tracks. Specifically, vertices are tagged as bad if they have an associated track whose d_0 is between (26,39), (46, 56.5), (83.5, 93.5), or (117.5, 128). The results of this cut may be seen in Figure 5.4, which shows that although

this does not entirely remove the ϕ asymmetry, it does lead to a reduction in the behavior.




Figure 5.4: The track d_0 distribution and vertex ϕ before and after the cleaning cuts.

5.3 JETS

As mentioned in Section 2.2, there is no single way of defining a jet, and any set of 4-vectors may be used as inputs to jet reconstruction. In practice, it is recommended to have as close of a correlation between the reconstructed object and truth-level particles as possible. On ATLAS, there are several different inputs to jet reconstruction that have been developed through the years, which rely on some combination of tracking and calorimetry information. While a track is designed to correspond to a single charged particle, this is more complicated to achieve in the calorimeter. However, particles tend to create localized energy deposits that can be clustered together to approximately correspond to deposits from a single particle. The information from these two systems may be used in a

variety of ways in order to improve the performance of jet reconstruction for particular regions of phase space. The main types of jet inputs used on ATLAS are topological calorimeter-cell clusters, particle flow objects, and track-caloclusters, each of which is described in the sections below, as well as developments towards further improving these results for future use.

5.3.1 TOPOLOGICAL CALORIMETER CLUSTERS

Topological calorimeter clusters (topoclusters) are jet inputs created purely out of calorimeter deposits. These objects are designed to make use of the 3-dimensional calorimeter information while simultaneously suppressing the noise. The algorithm creates the inputs as follows:

1. Clusters are seeded by cells with $|E_{\text{cell}}| > 4\sigma_{\text{noise}}$, where σ_{noise} is the expected amount of noise for that cell, and E_{cell} may be negative to prevent a bias in the reconstructed energy.
2. Any cells that neighbor the seed that have $|E_{\text{cell}}| > 2\sigma_{\text{noise}}$ are added to that cluster. This continues until there are no neighboring cells with $|E_{\text{cell}}| > 2\sigma_{\text{noise}}$.
3. All neighboring cells are added to the cluster, regardless of energy

This method is called the ‘4-2-0’ algorithm, based on the seeding of the clusters. Once the clusters are created, they can then be split into multiple clusters, which happens when there are two local maxima within the cluster*. Once the algorithm has found the local maxima, the energy of the cells located between the maxima are divided between the two maxima based on the relative distance to each maxima. Then, the cluster is given an energy of the sum of the energy of its cells, η and ϕ of the energy-weighted average of the η and ϕ of the cells included in that cluster, and a mass of zero.

As mentioned in Section 4.2, hadronic and electromagnetic showers are fundamentally different. The ATLAS calorimeters are non-compensating, meaning that they do not account for the invisible energy present in hadronic showers. This means that the reconstructed signal for a hadron will be less than that for an electron or photon that deposits the same energy. In order to account for this, a weighting algorithm was developed that uses local shower information to calibrate the energy back to the hadronic scale, called “LCW” for locally calibrated

*The local maxima must be within the EM sampling layers EMB2, EMB3, EME2, and EME3, or FCAL0, and a local maxima is defined as an excess of at least 500 MeV as compared to at least four neighboring cells. Additionally, cluster splitting may occur in EMB1 or EME1, or can be in any of the hadronic calorimeter sampling layers.

weight. The LCW algorithm aims to correct signal inefficiencies in hadronic showers due to non-compensation of the calorimeter, signal losses in clustering, and signal losses from inactive material. It first determines the probability that a cluster is from an EM or hadronic shower. It then applies a correction for each of these effects, where correction depends on this probability. For instance,

$$w_{cell}^{calib} = P_{clus}^{EM} \times w^{emcalib} + (1 - P_{clus}^{EM}) \times w^{hadcalib}, \quad (5.2)$$

where P_{clus}^{EM} is the probability that a cluster is from an EM-shower, and $w^{emcalib}$ and $w^{hadcalib}$ represent the calibration weights derived for EM and hadronic showers respectively. Generally, LCW weights are applied to clusters used in analyses that are concerned with the substructure of a jet. Due to the noise introduced from pileup, these weights worsen the energy resolution of jets at low- p_T , and are not used for small-radius jets, but are important for large- R jets where the substructure of the jet is relevant.

5.3.2 PARTICLE FLOW

The tracking detector provides complementary information to the calorimeter, allowing for the combination of information from the two detectors in order to provide more precise inputs to jet reconstruction. The momentum resolution of the tracking detector is inversely proportional to the track momentum, while the calorimeter energy resolution scales inversely proportional to the square root of the energy of the particle. Generally, the angular resolution of the tracker is better than that of the calorimeter, since it is designed to measure the direction of a particle instead of its energy. The Particle Flow (PFlow) algorithm combined tracking and calorimetric information with the goal of improving the resolution and pileup stability of reconstructed jets, compared to jets created from topoclusters alone³. In this algorithm, tracks are matched to EM-scale topoclusters, whose creation is described above. Then, the expected amount of deposited energy is calculated, based on the track momentum and position of the topoclusters, using measured value of the ratio of E/p for single particles. Based on this ratio, if it is likely that the track has deposited energy into multiple topoclusters, more topoclusters may be included as well. Then, the expected energy is subtracted cell-by-cell from all of the matched topoclusters until either all of the

expected energy has been subtracted or no more cells remain. If cells remain after the subtraction, the algorithm checks if this energy is consistent with the expected shower fluctuations of a single particle. If it is, this energy is simply removed, and otherwise, it forms a (modified) cluster with no associated track. Any topoclusters that are not matched to tracks are assumed to originate from neutral particle energy depositions and are unmodified. After this algorithm has been completed, the set of particle flow objects (PFOs) consists of tracks whose energy has been scaled, modified neutral clusters, and unmodified neutral clusters. Typically, PFOs that are not associated to a pileup vertex are removed, which increases the pileup stability of the PFlow algorithm over unmodified topoclusters.

In some cases, it is difficult to determine an accurate estimate of the expected energy deposit for a track. For tracks that are nearby other tracks, it is difficult to determine the contributions from an individual track to a particular cluster, making the subtraction ineffective. Similarly, for high- p_T tracks, where the momentum resolution is sufficiently degraded, the expected shower cannot be calculated accurately enough to make a meaningful subtraction. With this in mind, the PFlow algorithm turns off gradually for tracks between 30 and 60 GeV, and no subtraction is done for tracks above 60 GeV. Similarly, tracks in dense environments are not used for subtraction.

5.3.3 TRACK-CALOCCLUSERS

A second algorithm exists to combine tracking and calorimetric information, which is called Track-CaloClusters (TCCs)³³. This algorithm was designed to improve the substructure resolution of highly boosted objects, by using tracks to split clusters, based on the increased angular resolution of the tracks. Since TCCs are geared towards jet substructure, this algorithm uses the LC-scale of the topoclusters instead of the EM scale used in PFlow. Unlike PFlow, the TCC algorithm does not attempt to do an accurate estimate of the energy contributions of a particular track, and instead assumes that each track matched to a cluster deposits energy proportional to its p_T . The algorithm works as follows. A mapping is created between tracks and clusters based on an extrapolation of each track to the center of the cluster. Then, for each cluster, the total momentum of associated tracks is computed. Then, one TCC object is created per track, where the energy of that object is determined based on fraction of momentum

that a given track contributes to a given cluster.

More precisely, the TCC splitting algorithm takes place in several steps. First, it is noted that each track τ seeds a TCC object. Then, for each cluster, the set of all clusters matched to the track τ is given by C_τ , and similarly, the set of all tracks matched to a cluster c is given by T_c .

The contribution of each cluster c in C_τ will contribute proportionally to the ratio of its p_T to the p_T of the 4-momentum sum of all associated clusters. This fraction f_τ^c can then be written as

$$f_\tau^c = \frac{p_T^c}{p_T[\sum_{c \in C_\tau} p_T^c]}. \quad (5.3)$$

Then, the amount of momentum contributed from a given cluster is determined based on all associated tracks T_c , with the momentum of each track weighted by the relative weight of the cluster for that track. This fraction is then given by F_c^τ , defined to be

$$F_c^\tau = \frac{p_T^\tau}{p_T[\sum_{\tau' \in T_c} p_T^\tau f_{\tau'}^c]}. \quad (5.4)$$

Finally, the total momentum sharing for a track τ is taken to be

$$M_\tau = \sum_{c \in C_\tau} p_T^c f_\tau^c F_c^\tau, \quad (5.5)$$

such that the 4-momentum of a TCC t seeded by track τ is given as

$$p_t = (p_T[M_\tau], \eta^\tau, \phi^{tau}, m[M_\tau]). \quad (5.6)$$

While this is a complicated equation, a couple examples may be used to demonstrate its behavior. For a track matched to two clusters, each with no other associated tracks, the energy and mass of the associated TCC is given by the sum of the four-momentum of the two clusters. For a cluster with energy E_c that has two tracks associated to it, with transverse momentum p_T^1 and p_T^2 , two TCC objects are created, with the energy of $(E_c \times p_T^1)/(p_T^1 + p_T^2)$ and $(E_c \times p_T^2)/(p_T^1 + p_T^2)$ respectively. Tracks that are not matched to clusters are not included in the final set

of TCC objects, and unmatched topoclusters are also added to the collection. Finally, any TCC objects that are matched to a pileup vertex are removed from the list of objects.

One feature of this algorithm is its pileup dependence. As pileup increases, the number of tracks increases, and the number of tracks matched to each cluster increases on average. Since the TCC algorithm does not consider the expected shower deposition of a particle, and only considers the geometrical matching of tracks to clusters, as pileup increases, the energy remaining from the hard-scatter event will decrease, since more energy will be removed by pileup tracks matched to clusters created by the hard-scatter event.

5.3.4 UNIFIED FLOW OBJECTS

PFlow tends to perform well for low- p_T jets, where the algorithm can take advantage of the better tracking resolution, while TCCs tend to perform well for calculating the substructure of highly boosted objects where the granularity of clusters is not sufficient to resolve the different prongs of a jet. In order to improve the overall performance, it is possible to combine these algorithms sequentially into an object called *unified flow objects* (UFOs), with a few small modifications. Starting with PFlow objects, the set of tracks associated to the hard-scatter vertex who are not used for an E/p subtraction are found. This includes tracks that were not sufficiently well-matched to a cluster, as well as tracks in dense environments or high- p_T tracks. Then, these tracks are used as inputs to the TCC algorithm, where they are able to split apart clusters that have not been well-matched to tracks in the PFlow algorithm. This provides improvements in a couple of ways. It enables the cluster-splitting for highly boosted objects, which is not possible with the current PFlow algorithm. However, it does this in a way that avoids unnecessary splitting of clusters with pileup tracks, since these are not included in the cluster-splitting algorithm, and retains the PFlow E/p subtraction where relevant. Since PFlow is done at the EM-scale currently, this algorithm is performed at the EM scale as well.

A comparison of the W-tagging performance of EM-scale PFlow, TCCs, and UFOs is shown in Figure 5.5. These ROC curves show the background rejection as a function of the signal efficiency using a simple 2-variable tagger. In this tagger, the smallest mass window that contains 67% of the signal W-jets is determined, and a 1-sided cut

of d_2 is scanned over to obtain different signal efficiencies. This demonstrates that UFOs consistently perform as well as or better than either TCCs or PFlow across a wide variety of p_T scales, motivating further studies of these objects.




Figure 5.5: A comparison of the tagging performance of EMPFlow, TCCs, and UFOs at several different p_T scales using a 2-variable tagger with the mass and d_2 .

5.3.5 JET RECONSTRUCTION

Once the jet inputs are created, jet reconstruction algorithms may be applied to these inputs. Jets on ATLAS are created using the anti- k_t algorithm, typically using radius parameter of either 0.4 or 1.0, which are known as small-

R and large- R jets respectively. The anti- k_t algorithm typically produces circular jets, which simplify their experimental interpretation. Most analyses use the standard small- R jet collection, since they capture most of the energy from the shower, and correspond roughly to a single parton. Analyses that are interested in the substructure of a jet, either to identify hadronic decays of heavy boosted objects or to study the inner structure of a jet typically use jets with a larger radius parameter (large- R jets), which captures more of the detailed shower information. The energy scale of reconstructed jets differs from the energy scale of a corresponding truth jet, due to effects from input reconstruction as well as additional radiation such as pileup. In order to account for this, a multi-stage calibration procedure is applied in order to bring the scale of the reconstructed jet closer to the truth scale.

The first stage of the calibration procedure is designed to correct effects from pileup. Pileup adds energy roughly uniformly throughout the event, meaning that it increases the scale of a jet. Since the amount of pileup changes on an event-by-event basis, it is important to have a more meaningful subtraction of this pileup energy so that a calibration factor applied to a jet in a low-pileup event will also be applicable to one applied to a jet in a high-pileup event. For a generic hard-scatter event, the energy density will be peaked in a few places, resulting in a median energy density of zero. However, assuming a uniform distribution of pileup, the median energy density in an event can be used to approximate the amount of pileup that has been included in a jet. The median energy density is calculated using the median p_T density of $R=0.4$ k_t jets with $|\eta| < 2$, where the k_t algorithm is used due to its sensitivity to soft radiation in the event, and then ρ is given by p_T/A_{jet} ². After the median energy density has been calculated, $\rho \times A$ is subtracted from each jet in order to bring its scale closer to the scale without pileup.

After the pileup subtraction has been applied, there is still some remaining pileup dependence of the calibrated jet response. In order to account for this, a two-stage residual correction is applied. It is observed that the p_T dependence of the jet response is fairly linear in both N_{PV} and μ . Fitting these slopes enables a further subtraction. In total, the pileup-corrected p_T of a jet is given by

$$p_T^{corr} = p_T^{reco} - \rho \times A - \alpha \times (N_{PV} - 1) - \beta \times \mu, \quad (5.7)$$

where α and β are determined by linear fits in several different η regions.

Once the pileup correction has been applied, a calibration is derived to bring the energy closer to that of the truth scale. The first stage of this is a calibration produced from Monte Carlo events, and is referred to as the Monte Carlo jet energy scale calibration, or MCJES. To derive this calibration, reconstructed jets are matched to truth jets if $\Delta R < 0.3$. This is done on isolated truth jets in order to avoid closeby effects. The jet energy scale is calibrated using the numerical inversion method²⁶. These calibration factors are then able to be used for all jets. After the MC calibration has been applied, a global sequential calibration (GSC) is derived in order to account for differences in the calorimeter response on the jet composition and energy distribution. This is designed to keep the same energy response, but to improve the overall resolution.

Finally, the last stage of the jet calibration accounts for differences in the jet response in data and Monte Carlo, which can arise from imperfect detector descriptions, as well as imperfect modeling of the detector response, or differences in the simulation of the hard-scatter or pileup event. These differences are measured by comparing the p_T of a jet to that of a different well-measured reference object, including $Z \rightarrow \mu\mu + \text{jet}$ events, $\gamma + \text{jet}$ events, and multijet events. A full description of this calibration may be found in ².

6

Constituent-Level Pileup Mitigation

Pileup is one of the major challenges impacting object reconstruction on ATLAS. Higher instantaneous luminosities enable huge amounts of data-taking, but also result in multiple (soft) interactions per bunch crossing. At the beginning of Run 2, each event typically had around 20 reconstructed primary vertices per bunch crossing, while this number has increased to nearly 60 by the end of Run 2. These additional interactions result in a number of effects, including an increased amount of energy not associated to the hard-scatter (HS) event of interest, higher fake rates in object reconstruction, decreased reconstruction efficiency, worsened resolution of reconstructed objects, and cut efficiencies which depend on PU. These effects are present in some combination for all different objects which are reconstructed, but are particularly challenging for jets, which have a comparatively large catchment area.

6.1 PILEUP MITIGATION TECHNIQUES

Historically, small- R jet reconstruction on ATLAS has relied on a two-fold approach to eliminate pileup. The first step corrects the expected PU contribution¹³, and later, jets are tagged as PU or HS based on their associated tracks. The PU correction is done using the jet-area subtraction⁴⁶, which calculates the area of a jet, and subtracts the expected amount of PU energy in that region based on the median energy density of the event. Pileup jet tagging is done with the jet vertex tagger (JVT)^{13,25}, which determines the likelihood of a jet being dominated by PU based on the percentage of associated tracks originating from the HS primary vertex. Large- R jets rely on grooming algorithms to remove pileup dependence. Since grooming algorithms remove soft and wide-angle radiation from jets, this also works to remove pileup from the jet. Since large- R jets are typically used at high p_T , a pileup jet tagger is not necessary.

While these techniques have worked effectively for early Run 2 conditions, they begin to fail in higher pileup conditions. In light of this, there has been a range of algorithms developed to improve these techniques, in particular by applying them to individual constituents instead of individual jets. Several of these algorithms are described below, and their performance in ATLAS simulation is evaluated.

6.1.1 CONSTITUENT SUBTRACTION

Constituent subtraction⁴⁰ offers an alternative way of defining area subtraction for constituents using ghost association. For an event with median energy density ρ , the expected amount of pileup energy in a region with area $A = \delta\eta \times \delta\phi$ is $\rho \times A$. With this in mind, massless, low- p_T virtual particles called *ghosts* are added uniformly to each event with a spacing of $l = 0.1$ in $\delta\eta$ and $\delta\phi$, such that they each cover an area (A_g) of 0.01. Each ghost is given the same p_T of $A_g \times \rho$, such that the energy density of the ghosts matches the median energy density.

After the ghosts have been added, the distance, $\Delta R_{i,k}$ between each possible pair of particles, i , and ghosts, k is calculated. Then, starting with the closest ghost-particle pairs, the p_T of each ghost and particle is modified by

The algorithm iterates until all pairs with $\Delta R_{i,k} > \Delta R_{max}$, where ΔR_{max} is one of the free parameters of the algorithm.

$$\begin{aligned}
\text{If } p_{T_i} \geq p_{T_k}^g: \quad & p_{T_i} \rightarrow p_{T_i} - p_{T_k}^g, \\
& p_{T_k}^g \rightarrow 0; \\
\text{otherwise:} \quad & p_{T_i} \rightarrow 0, \\
& p_{T_k}^g \rightarrow p_{T_i}.
\end{aligned}$$

6.1.2 VORONOI SUBTRACTION

Voronoi subtraction provides an alternative approach to applying area subtraction, using the geometrical area of each constituent given by its Voronoi area in $\eta - \phi$ space*. The cluster momentum is then corrected by

$$p_T^{\text{corr}} = p_T - \rho * A_{\text{Voronoi}}. \quad (6.1)$$

After this subtraction, many clusters are left with an unphysical negative transverse momentum, which cannot be used for jet finding. Two algorithms can be used to deal with this: negative suppression and spreading. For *negative suppression* (Vor Supp), negative p_T^{corr} clusters are simply removed from the event, and jet clustering is performed on the remaining clusters. Since the pileup energy density fluctuates throughout a given event, the median energy density may still be non-zero after applying negative suppression. An alternative approach called *Voronoi spreading* accounts for some local fluctuations by distributing negative p_T^{corr} to other nearby (positive p_T^{corr}) clusters. The algorithm proceeds as follows.

For each cluster i with $p_{T,i}^{\text{corr}} < 0$:

1. Find the set of clusters J , where each cluster $j \in J$ is within a cone of $\Delta R_{ij} < 0.4$, and has $p_{T,j}^{\text{corr}} > 0$
2. For each cluster $j \in J$, calculate a weight:

$$w_{ij} = \frac{1}{\Delta R_{ij}^2} \times \frac{1}{\sum_{\Delta R_{ij}^2} \frac{1}{\Delta R_{ij}^2}}, \quad (6.2)$$

such that $\sum_j w_{ij} = 1$.

3. Rescale the p_T of each cluster $j \in J$ to be $p_{T,j}^{\text{spread}} = \max(p_{T,j}^{\text{corr}} + w_{ij} \times p_{T,i}^{\text{corr}}, 0)$.
4. Set $p_{T,i}^{\text{spread}} = 0$

This algorithm ensures that $p_T^{\text{spread}} \geq 0$, and creates a valid set of inputs to jet clustering. Figure 6.1 shows the

*The Voronoi area of a constituent is the area of all points on a plane which are closer to that constituent than any other constituent

application of the Voronoi algorithms, starting from the original event, applying the subtraction, and then with both suppression and spreading.

6.1.3 SOFTKILLER

Unlike the Voronoi and CS algorithms, which rescale the p_T of each cluster, the SoftKiller (SK) algorithm⁴⁹ applies a p_T cut to constituents. This cut is determined on an event-by-event basis, and is designed to make the median p_T flow density (ρ) approximately zero. This is done by dividing the detector into an $\eta - \phi$ grid, with each grid cell containing all constituents in that region. Then, a p_T cut is chosen such that the half of the grid spaces are empty, which eliminates all constituents from 50% of the cells as well as some constituents from the remaining cells. The only free parameter of the SK algorithm is the size of the grid, and should be chosen to be large enough that many cells contain constituents before SK, but small enough that not all constituents are eliminated.

When not applied at particle level, SK is designed to be applied after some form of area subtraction such as Voronoi subtraction or CS, since clusters may be a combination of HS and PU energy. An example of applying SK to an event may be seen in Figure 6.2, which demonstrates the process of starting from an unsubtracted event, applying Voronoi subtraction, and then applying SK with two different grid sizes. This event display demonstrates the balance that must be struck in choosing a SK grid size: a grid length of 0.4 leaves many clusters which are not associated to truth jets, while using a grid length of 0.6 begins to cut clusters which are likely associated to the HS event.

6.1.4 PILEUP PER PARTICLE IDENTIFICATION (PUPPI)

While the first previous algorithms only rely on the 4-momentum of each constituent, it is also possible to use tracking information to improve the classification of particles. One such algorithm is called “pileup per particle identification”, or PUPPI⁴¹. PUPPI relies on particle flow information, which provides a classification of charged particles as either HS or PU. In this algorithm, each constituent i is assigned a likelihood that it is pileup. To determine this likelihood, first, the quantity α_i is calculated, where

$$\alpha_i = \log\left(\sum_j \frac{p_{T,j}}{\Delta R_{ij}} \times \Theta(R_{\min} \leq \Delta R_{ij} \leq R_0)\right),$$

where j are all of the charged constituents associated to the Primary Vertex (PV), R_0 is the maximum radial distance at which constituents are associated to each other, and $R_{\min}$ is the minimum radial distance of association. $R_{\min}$ is generally taken to be very small, and is set to 0.001 for these studies. The idea behind this metric is that neutral activity nearby high- p_T charged activity is more likely to also be pileup than if it is far away from any other HS activity. Once α has been calculated for all constituents, then the following quantity is determined:

$$\chi_i^2 = \Theta(\alpha_i - \bar{\alpha}_{\text{PU}}) \times \frac{(\alpha_i - \bar{\alpha}_{\text{PU}})^2}{\sigma_{\text{PU}}^2}, \quad (6.3)$$

where $\bar{\alpha}_{\text{PU}}$ is the mean value of α for all charged PU constituents in the event, and σ_{PU} is the RMS of that same distribution. Then, the 4-momentum of each neutral constituent i is weighted by

$$w_i = F_{\chi^2, \text{NDF}=1}(\chi_i^2), \quad (6.4)$$

where F_{χ^2} is the cumulative distribution of the χ^2 distribution. This eliminates all neutral constituents i whose calculated value of α_i is less than $\bar{\alpha}_{\text{PU}}$. Finally, a p_T cut is applied to the remaining reweighted constituents in order to further suppress noise. This is N_{PV} dependent, and is determined by

$$p_{T, \text{cut}} = a + b \times N_{\text{PV}}. \quad (6.5)$$

While PUPPI can technically be applied to topoclusters, the principles of the algorithm depend strongly on the association of neutral constituents to nearby charged HS constituents, and is therefore much more effective for particle-flow type algorithms. In the forward region of the detector, where charged information is not available, the algorithm needs to be optimized separately, and is done using α_i calculated from *all* constituents instead of charged HS constituents. For these studies, we take $R_0 = 0.3$, $a = 200$ MeV and $b = 14$ MeV, based on an optimization done for small- R jets made from PFOs.

6.1.5 CLUSTER VERTEX FRACTION

The cluster vertex fraction (CVF) algorithm is a simple form of particle flow modelled after the jet vertex tagging (JVT) method. Tracks are associated to a cluster when $\Delta R_{\text{cluster, track}} < 0.1$,[†] and CVF is then calculated as

$$\text{CVF} = \frac{\sum_{\text{tracks,HS}} p_T^{\text{track}}}{\sum_{\text{tracks,HS}} p_T^{\text{track}} + \sum_{\text{tracks,PU}} p_T^{\text{track}}}, \quad (6.6)$$

where $\sum_{\text{tracks,HS}}$ and $\sum_{\text{tracks,PU}}$ indicate the sum over all associated tracks from the hard scatter primary vertex and the pileup vertices, respectively. If there are no tracks found within $\Delta R < 0.1$ of the cluster then it is assigned a CVF value of -1. Values of CVF close to 1 indicate hard scatter activity, while $\text{CVF} = 0$ indicates pileup activity, and $\text{CVF} = -1$ indicates a neutral cluster with no associated tracks. Since pileup tends to be low- p_T , and since a single nearby PU track can skew the value of CVF, the CVF algorithm removes clusters based on both their calculated value of CVF and their p_T . Clusters with $\text{CVF}=0$ and $p_T > 2$ GeV are considered PU and removed from the event. This algorithm is essentially a rudimentary form of particle flow which is able to be tested easily without creating the full PFO collection. CVF may be used alone or in addition to other pileup mitigation techniques, and is found to perform better in conjunction with Voronoi Spreading than when applied on its own. Based on an optimization for small- R jets, Voronoi spreading is applied to the constituents, and CVF is applied afterwards.

[†]The effect of the magnetic field on (low momentum) tracks is taken into account by extrapolating the track to the first layer of the hadronic calorimeter using a simple model based on p_T , charge and η of the track.




Figure 6.1: Event displays at $\mu=20$ showing the variations on Voronoi subtraction. Each polygonal region is the Voronoi area associated to a cluster which is located at the center of the polygon. The color of the region shows magnitude of the p_T of each cluster after each step, and the black dotted circles indicate the truth jets over 20 GeV. (a) The event before Voronoi subtraction, (b) The event after Voronoi subtraction, (c), the event after Voronoi suppression, and (d), the event after Voronoi spreading




Figure 6.2: Example event display (a) before pileup removal, (b) after Voronoi subtraction is applied to constituents, (c) after VS and SoftKiller 0.4 are applied, and (d) after VS and SK 0.6 are applied. The size of the constituent is proportional to its p_T , and truth jets with $p_T > 20$ GeV are also shown in dashed lines.

6.2 IMPACT OF PILEUP MITIGATION ON SMALL-R JETS

The most common type of jets used on ATLAS are small- R jets. Currently, ATLAS supports these jets down to a p_T of around 20 GeV, but producing a calibration for these jets becomes more difficult as pileup decreases the correlation between the truth and reconstructed jets. The effects of pileup may be seen in both the response of the jet (p_T^{reco}/p_T^{true}) as well as the overall spread of this distribution (resolution). The effects of pileup on the jet p_T response can be corrected on average using a calibration, but even after such a calibration, the resolution degrades with increased pileup. However, particle-level studies indicate that constituent-level pileup mitigation algorithms are able to improve the stability of the jet response distribution as a function of pileup.

In order to perform a preliminary comparison of the different constituent-level pileup mitigation algorithms on ATLAS, their behavior is studied in low- p_T jets using the full ATLAS simulation. These studies use a multijet sample produced in PYTHIA 8¹⁰³, where the p_T of the leading $R=0.6$ truth jet is between 20 and 60 GeV. Three samples are produced with different amounts of pileup, in order to study the performance in early Run 2 conditions ($\langle\mu\rangle = 20$), Run 3 conditions ($\langle\mu\rangle = 80$), and in preparation for the high-luminosity LHC ($\langle\mu\rangle = 200$). The $\langle\mu\rangle = 20$ sample uses the A14 tune together with the NNPDF2.3LO PDF set⁹¹, while the other samples use the AU2 tune with the CT10 PDF set. The pileup in these samples is included by overlaying minimum-bias samples, which are generated with PYTHIA8 using tune A2¹⁹ and the MSTW2008LO PDF set. In all samples the detector response is simulated using Geant4⁷¹ and the ATLAS simulation framework¹⁷. Simulation of the detector response to pile-up proceeds using the same software as for the hard-scatter process.


Once the samples are produced, anti- k_t ^{47,45} jets with a radius parameter 0.4 are produced using FastJet⁴⁸. The inputs to the reconstructed jets are LCW topoclusters, which may be modified by a pileup mitigation algorithm, while the inputs for truth jets are stable particles from the hard scattering event with a lifetime of over 10 ps, excluding muons and neutrinos. Once the truth and reconstructed jet collections are created, a ΔR matching is performed, where a truth and reconstructed jet are matched if $\Delta R_{\text{truth, reco}} < 0.3$. These studies require truth jets to be isolated from all other truth jets whose p_T is greater than 5 GeV by at least $\Delta R \geq 0.6$, which minimizes close-by effects in these studies.

The width of the jet response is impacted by the calibration of jets, and so it is important to perform a rough calibration of jets in order to compare their resolution. This is done using the numerical inversion method²⁶ to calibrate reconstructed jets to the stable-particle-scale. The calibration is derived in bins of the number of primary vertices N_{PV} . First, in each N_{PV} bin, truth jets matched to reconstructed jets are binned in p_T^{true} to find the mode of the corresponding response distribution in that bin:


$$R_{p_T} = \frac{p_T^{reco}}{p_T^{true}} \quad (6.7)$$

This distribution is approximately Gaussian, and the jet response is given by the position of a Gaussian distribution fit to the response distribution. The jet response is binned in $\langle p_T^{true} \rangle$, and then fit in order to derive a functional form for $R(p_T^{true})$. Finally, with the definition $f(p_T^{true}) = p_T^{true} \times R(p_T^{true})$, the calibrated p_T value for a given reconstructed jet is defined as $p_T^{calib} = f^{-1}(p_T^{reco})$. Application of the calibration procedure in bins of N_{PV} has a similar effect to the residual correction described in Ref.³⁴.


The jet p_T resolution is then used to compare the performance of the different algorithms, where the resolution is defined as the width of a Gaussian fit of the calibrated response distribution. The results may be seen in Figure 6.3 for several different values of μ . These results demonstrate that constituent-level pileup mitigation is able to improve the jet resolution compared to both uncorrected and area subtracted jets. As expected, the improvement increases under higher pileup conditions. with up to a 20-30% improvement over standard techniques in HL-LHC conditions. The improvements are most obvious for low- p_T jets, since the relative pileup contributions are most significant at low- p_T , but some improvements may be seen up to 60 GeV. Overall, the CVF algorithm performs the best, which is expected considering its use of tracking information. Of the algorithms which only use calorimeter information, CS+SK is the most performant choice.



(a)



(b)



(c)

Figure 6.3: Resolution of low- p_T jets reconstructed with anti- k_t with $R = 0.4$ versus true jet p_T . An average of 20, 80 or 200 pile-up events is superposed on the multi-jet signal. The constituent-level pileup mitigation methods are explained in the text. Monte Carlo statistical uncertainties are negligible, and the estimated jet energy resolution uncertainty is 2-3 percent.

6.3 IMPACT OF PILEUP MITIGATION ON LARGE- R JETS

For analyses interested in the substructure of a jet, either to tag hadronic decays of heavy boosted objects, or to measure properties of the jet shape, large- R jets are typically used, since they include a larger portion of the shower. These analyses are typically interested in more than just the p_T and position (η, ϕ) of the jet, making use of jet substructure observables which are sensitive to the soft radiation within the jet. These observables differ from observables like the p_T , which is affected by the total amount of pileup remaining, not the distribution of pileup. This means that jet substructure observables are more strongly affected by pileup, since soft radiation is able to impact the calculation of these observables. Jet grooming algorithms are able to remove some of this pileup dependence, but some residual dependence may still exist after grooming. This means that constituent-level pileup mitigation algorithms are also applicable to large- R jets.

In order to study this performance, several different Monte Carlo samples are produced. For these studies, a sample of jets originating from light quarks and gluons (“QCD jets”) is generated using high- p_T multijet events, which are generated in slices of jet p_T in order to retain high statistics at high p_T . These events are then weighted to produce a smoothly falling p_T distribution predicted by PYTHIA8. Signal W jets are taken from the beyond the Standard Model spin-1 $W' \rightarrow WZ \rightarrow qqqq$ process, where a maximum of one jet per event is selected based on its matching to the truth W particle, with a maximum of $\Delta R < 0.75$ between the particle and the jet axis. An array of W' signal samples are generated, with W' masses ranging from 400-5000 GeV, in order to obtain a broad kinematic range of W jets. In order to remove the dependence on the signal model, jet-by-jet weights are applied based on the anti- k_t truth jet, such that the p_T spectrum of signal W jets matches that of the QCD jets.

In all cases, the samples are generated the PYTHIA 8.186 generator¹⁰³ using the A14 tune¹⁹ with the NNPDF2.3LO PDF set³⁷. After simulation of the parton shower and hadronization, events are passed through the full GEANT4⁷¹ simulation of the ATLAS detector¹⁷. Hadronic showers are simulated using the FTFP BERT model²³. Pileup is simulated by overlaying additional minimum-bias events generated with PYTHIA8. These events are overlaid onto the hard scattering events to match the average number of pp interactions per bunch crossing during Run 2 of the LHC. This overlay is performed in such a way that each sub-detector considers the effect of adjacent bunch cross-

ings as well as the current hard scatter, thereby accounting for both in-time and out-of-time pileup as described in Reference³¹.

Using these samples, jets are reconstructed with the anti- k_t algorithm with $R=1.0$ using a variety of inputs, including LCW topoclusters, PFOs, and TCCs (possibly with pileup mitigation techniques applied). Truth jets are reconstructed with the anti- k_t algorithm with $R=1.0$, and an exclusive ΔR matching in (η, ϕ) is performed between the truth jet and reconstructed jets for each type of input. The kinematics of this ungroomed truth jet are used to characterize the p_T scale of the hadronic energy flow entering the detector, which ensures that in a given kinematic region, the comparison between reconstruction algorithms is performed using the same set of jets.

In the case of the QCD jet background sample, the leading truth jet is always kept assuming that it has p_T greater than 300 GeV. The ungroomed reconstructed jet closest to the leading truth jet, within $\Delta R < 0.75$, is kept, where this ΔR matching requirement is performed independently for jets reconstructed from each type of input. Grooming is then applied to these truth and reconstructed jets, ensuring that the same population of jets studied for each type of grooming algorithm. In the case of the signal W -jet sample, rather than keeping the leading truth jet, the truth jet originating from the W boson is kept. The truth jet originating from the W boson is identified via the criteria $\Delta R(J_{\text{Truth}}, W) < 0.75$, where J_{Truth} is the truth jet nearest to the W boson. After determining which truth jet to keep, the procedure for deciding which ungroomed reconstructed jets to keep for each input type proceeds in the same way as for the QCD jet background sample.

In order to ensure full containment of the hadronic W boson decay, jets are only studied if they are matched to a truth jet with $p_T > 300$ GeV. Additionally, to ensure that the catchment area of the jet is largely contained within the region of the detector where charged-particle tracks may be reconstructed, an additional cut of $|\eta^{\text{true}}| < 1.2$ is applied.

6.3.1 PILEUP STABILITY OF THE W MASS DISTRIBUTION

For analyses looking to tag hadronic decays of boosted heavy objects like W bosons or top quarks, the jet mass is one of the best discriminating variables. Since the jet mass is a jet substructure observable, the mass distribution

shifts as a function of the number of reconstructed primary vertices (N_{PV}), as shown in Figure 6.4, which shows the mass distribution of W -boson jets in low and high pileup conditions. As expected, the jet mass increases with pileup, since more uncorrelated energy is included within the jet. Similarly, the jet mass scale and resolution are impacted, but this is challenging to compare without a calibration. Since the pileup dependence of a jet is affected by both the jet inputs as well as the grooming algorithm, it is difficult to study these effects directly without deriving a calibration for each combination of jet input and grooming algorithm. Instead, it is possible to compare the pileup stability of uncalibrated jets using metrics which are largely unaffected by the calibration. One such possibility is to use the reconstructed W boson peak central value and width. These values are much less affected by the calibration, and may also be used to gauge the relative dependence of the mass scale and resolution as a function of pileup.

To quantify these effects, the W boson peak reconstructed mass spectrum is iteratively fit with a Gaussian in exclusive bins of N_{PV} to extract the central value, and the size of the full 68% window is used as a measure of the width of the peak. These values are plotted as a function of N_{PV} and fit with a line, as shown in Figure 6.5, and the slope of this line is a measure of the pileup stability of the jet definition, which can be used to compare a large number of jet definitions. This figure also demonstrates how constituent-level pileup mitigation algorithms are able to be used to improve the pileup stability of the mass distribution.

The values of the slopes for both the central value and the width may be seen for all constituent and grooming configurations in Figure 6.6. The constituent-level pileup mitigation techniques significantly outperform uncorrected LCTopo and PFlow inputs, while it worsens the performance for TCCs, as seen in Figure 6.6. The behavior of TCCs is also notable in that the position of the W mass peak decreases as a function of pileup. As pileup increases, more tracks are associated to each cluster, and those tracks are removed from the event. As such, the TCC algorithm essentially overcorrects the effects of pileup, and this performance is further degraded by adding pileup mitigation algorithms.

While the variations in grooming algorithms are less drastic, a few trends may be observed. For Soft Drop grooming, smaller values of β are more stable under pileup, since they remove soft radiation more aggressively. Com-

pared to the other grooming algorithms, trimming without pileup mitigation is more robust against pileup effects at low p_T , though pileup mitigation techniques further stabilize the performance of the algorithm.




Figure 6.4: A representative comparison of reconstructed W boson jet mass in exclusive bins of the number of primary vertices (N_{PV}), illustrating the shift to higher masses for increased pileup.




Figure 6.5: (a) A representative comparison of reconstructed W boson jet mass in exclusive bins of the number of primary vertices (N_{PV}), illustrating the shift to higher masses for increased pileup. (b) The mean value of a Gaussian fit to the reconstructed W boson mass peak as a function of the number of reconstructed primary vertices (N_{PV}), presented here for the LCTopo, CS, and CS+SK for the Soft Drop grooming algorithm.





Figure 6.6: Pileup dependence of the (a) value and (b) width of the fitted W boson mass peak


6.3.2 PILEUP STABILITY OF D_2

In addition to the jet mass, other jet substructure observables are frequently used to tag W jets. In particular, the D_2 observable characterizes the radiation within a jet by using correlations between *all* pairs and triplets of jet constituents^{84,85,86}. This observable is defined in terms of the n -point energy correlation functions e_2 and e_3

$$\begin{aligned} e_2^\beta &= \frac{1}{p_{T,j}^2} \sum_{1 \leq i < j \leq n_j} p_{T,i} p_{T,j} R_{ij}^\beta, \\ e_3^\beta &= \frac{1}{p_{T,j}^3} \sum_{1 \leq i < j < k \leq n_j} p_{T,i} p_{T,j} p_{T,k} R_{ij}^\beta R_{ik}^\beta R_{jk}^\beta, \end{aligned} \tag{6.8}$$

with $D_2 = e_3/e_2^3$. There are several characteristics of the n -point correlation function which make it theoretically useful for discriminating between different radiation patterns within jets. The observable e_n^β is infrared and collinear safe, which is sensitive to the n -prong structure in a jet. For any value of n , e_n^β is insensitive to recoil, meaning that it does not depend on the angular distance between the hard jet core and the jet momentum axis, which can occur with soft, wide-angle radiation. Additionally, in the soft or collinear limit, $e_n^\beta \rightarrow 0$ for all configurations and for any number of particles within the jet. Unlike the n -subjettiness¹⁰⁵ observable, this D_2 observable does not rely on hard subjet axes to derive its discrimination power. For this reason, it can achieve a higher level of discrimination power.

However, since D_2 is dependent on the fine structure of the jet, this also leads to an increased pileup dependence. The D_2 distribution for a few choices of jet inputs and grooming algorithms may be seen in Figure 6.7. To quantify the effects of pileup on D_2 , the 50% signal efficiency cut is determined for W jets with $m_{jet} > 50$ GeV in a low pileup region ($0 < N_{PV} < 15$). This cut is then applied for a range of N_{PV} bins, and the efficiency of this cut is calculated in each bin, as shown in Figure 6.8. The efficiencies are parametrized with a linear fit as a function of N_{PV} , where the slope of this line characterizes the pileup dependence. Figure 6.9 shows a summary of the slope of the D_2 tagging efficiency as a function of pileup, for all different input and grooming combinations. The conclusions from the D_2 efficiency metric are consistent with those drawn from the W boson mass peak. Figure 6.9 shows the dependence of the background efficiency using the same D_2 threshold as for the signal.




(a) LCTopo, Soft Drop($z_{\text{cut}} = 0.10$, $\beta = 0$)

(b) CS+SK, Soft Drop($z_{\text{cut}} = 0.10$, $\beta = 0$)

Figure 6.7: The D_2 distribution of W boson jets in events with low and high pileup for jets formed from LCTopo clusters (a) and LCTopo+CS+SK corrected clusters (b) and groomed with the Soft Drop algorithm with $z_{\text{cut}} = 0.1$ and $\beta = 0$.

6.4 SUMMARY


As the instantaneous luminosity increases, constituent-level pileup mitigation techniques are becoming critical for jet reconstruction, both to improve low- p_T small- R jet reconstruction, as well as for analyses relying on jet sub-structure observables. While these results look promising, they are only the first set of steps towards using this for jet reconstruction on ATLAS. An alternative method of studying the impact of pileup on object reconstruction is shown in Appendix C, which outlines the basics of the DigiTruth method. Since it represents a large change from the current area-subtraction technique, a more complete set of studies is necessary before migrating to these inputs, particularly for small- R jets, which only studied the effects in a narrow p_T and η region. A more complete set of studies is ongoing, with the aim of preparing for Run 3. For large- R jets, these techniques are beginning to be




(a) Soft Drop($z_{\text{cut}} = 0.15, \beta = 0$)

Figure 6.8: The dependence of the signal W jet identification efficiency for a D_2 selection, chosen at low N_{PV} to provide 50% signal efficiency for W boson jets, as a function of the number of reconstructed primary vertices for a low- p_T kinematic selection of jets, as determined by the ungroomed truth jet p_T . Presented here for the LCTopo, CS, and CS+SK jet constituent inputs are a Soft Drop (a) configuration.

studied for use as the recommended jet definition. While they are currently not used by default, constituent-level pileup mitigation techniques show a promising route of study for improving jet reconstruction in high pileup conditions.



(a)



(b)

Figure 6.9: Pileup dependence of a D_2 cut on the (a) W boson jet selection efficiency and (b) QCD jet selection.

Part I

Measurements of Jet Substructure

Observables

The field of jet substructure exists to utilize and study the internal structure of jets. Historically, the focus has been on improving the ability to classify jets originating from light quarks or gluons from those originating from hadronic decays of boosted heavy objects, such as W bosons and top quarks. These taggers typically rely on jet substructure observables, such as the jet mass, which depend on the distribution of energy within the jet. Recently, efforts have been made to create a solid theoretical understanding of the behavior of jets. This work has resulted in some of the first predictions of the differential cross-sections of jet substructure observables. A comparison of these predictions of the jet mass to an unfolded measurement is the focus of the first part of this thesis, and is covered by Sections 7- 11.

This thesis documents a measurement of the Soft Drop jet mass using 32.9 fb^{-1} of $\sqrt{s} = 13 \text{ TeV}$ pp data collected in 2016 by the ATLAS detector. The studies shown here document the first direct comparison between collision data and calculations beyond the LL accuracy of Parton Shower (PS) MC programs⁶⁸. The jet mass measurement demonstrates both a strong level of control over our theoretical understanding of QCD as well as our ability to perform precise measurements of JSS observables. This measurement paves the way for a wide program of precision jet substructure, with which we may be able to extract fundamental parameters of the SM such as the strong coupling constant¹⁰ or the top quark mass⁷⁷.

The next few chapters document the details of this analysis. It starts in Section 7 by providing motivation for this measurement and describing the relevant details of the corresponding predictions. Section 8 summarizes the event and jet selection for the analysis along with the motivations for the particular cuts that were applied. Section 9 describes the basics of unfolding and how it applies to this analysis, followed by Section 10, which describes the various systematic and statistical uncertainties relevant for this analysis. Section 11 shows the results of the analysis, and Section 12 summarizes the results and describes various paths for future work related to this measurement.

7

Theoretical Background and Analytical Predictions

A full calculation of the jet mass distribution is beyond the scope of this thesis, and has been covered in depth in other papers. Even so, there are several features of the jet mass and its calculation that are relevant to understanding the results and their significance. This chapter describes some of these features, and provides the relevant information about the analytical predictions used for this measurement.

One of the most powerful jet substructure observables is the jet mass, which is defined as the norm of the four-momentum sum of all jet constituents. For hadronically-decaying, Lorentz-boosted massive particles, such as top quarks or W bosons, the mass of a jet is approximately the same as the mass of the originating particle, and typically receives small corrections from fragmentation and hadronization. However, for jets originating from light

quarks and gluons, the picture is more complicated. The mass of these particles is negligible, and they instead gain their mass from the fragmentation of highly virtual partons¹⁰¹, where a quark emits a gluon, or a gluon splits into two quarks.

While the fragmentation of partons is calculable, the difficulty arises because of the boundaries of the jet, which make the calculation of its mass sensitive to non-global logarithms (NGLs)⁵⁹. These NGLs are resummation terms arising from particles radiating out of, and then back into a jet. These terms are present at next-to-leading-logarithm accuracy, and are challenging to include in calculations. This has prevented a meaningful prediction of the jet mass distribution, until the creation of the modified Mass Drop, and later the Soft Drop grooming algorithms. These algorithms formally remove non-global logarithms from jets, enabling a prediction of the jet mass.

This measurement is compared to two analytical predictions. The first of these is a LO+NNLL calculation^{69 70}, and the second is a calculation at next-to-leading order (NLO) with NLL^{93 94}. For both of these, the mass distribution has been derived using the same p_T and ρ binning used in this measurement.

8

Event Selection and Samples

This analysis is focused on di-jet events, which enable comparisons to the analytical predictions of the jet mass.

The LHC produces copious amounts of di-jet events, so only a simple event selection is required to produce a pure sample of di-jet events across a wide range of jet p_T . The Monte Carlo samples, data samples, and event selection are all described below.

8.1 DATA SAMPLES

The entire 2016 dataset is used for this analysis. Since this analysis is not limited by statistics, this was simpler than combining with other years, where the trigger thresholds change.

8.2 MONTE CARLO SAMPLES

Several Monte Carlo simulations are used for this analysis. The unfolding matrices described in Section 9 are determined from di-jet events generated at leading order using PYTHIA¹⁰³ 8.186, with the $2 \rightarrow 2$ matrix element (ME) convolved with the NNPDF2.3LO parton distribution function (PDF) set⁵¹, and using the A14²⁴ set of tuned PS and underlying-event model parameters. Additional radiation beyond the ME was simulated in PYTHIA 8 using the LL approximation for the p_T -ordered PS⁵⁸.

In order to provide an uncertainty based on the modeling of the data, as well as to provide several comparisons for the unfolded data, two additional dijet samples were simulated. SHERPA 2.1.1⁷⁵ generates events using multi-leg $2 \rightarrow 3$ matrix elements, which are matched to the PS following the CKKW prescription⁵². These SHERPA events were simulated using the CT10 PDF set⁸¹ and the default SHERPA event tune. HERWIG++ 2.7.1^{36,57} events were generated with the $2 \rightarrow 2$ matrix element, convolved with the CTEQ6L1 PDF set¹⁰⁰ and configured with the UE-EE-5 tune⁷². Both of these samples differ from the PYTHIA samples, in that they use an angular ordering for the parton shower, and a cluster model for the hadronization¹⁰⁷, providing an interesting comparison of these effects.

For all of these samples, pileup is added to the events, by overlaying minimum bias events in order to match the pileup conditions in 2016 data taking. The minimum bias events are generated using PYTHIA 8, with the MSTW2008LO PDF set⁹¹ and A2 tune¹⁹. All samples are processed using the full ATLAS detector simulation¹⁷ based on GEANT 4⁷¹, in order to have good modeling of the inner structure of the jets.

8.3 EVENT SELECTION

Jets are clustered from LCW topoclusters whose four-momentum has been corrected to point towards the PV in order to improve the rapidity resolution. These topoclusters are then formed into jets using the anti- k_t jet algorithm with radius parameter $R = 0.8$ with FastJet⁴⁸. Jets are then calibrated using a jet energy calibration derived for $R = 0.8$ jets derived in Monte Carlo, which includes corrections of the scale as well as a residual correction to cor-

rect for pileup effects. Events are triggered on the lowest unrescaled small- R jet trigger, whose threshold is 400 GeV. Since the jets used in this analysis use a larger radius parameter (0.8 instead of 0.4), events are required to have at least one jet with $p_T > 600$ GeV, in order to have nearly 100% trigger efficiency. Jets are required to have $|\eta| < 1.5$, in order to be able to apply uncertainties to the topoclusters. In order to select dijet events, the jets are required to be fairly balanced in p_T by imposing that the leading two p_T -ordered jets satisfy $p_{T,1}/p_{T,2} < 1.5$. This produces a fairly pure sample of dijet events, and so no other event selection criterion are applied.

After the event selection has been applied, the Soft Drop algorithm is run on both jets in the event, using three different values of β : 0, 1, and 2, which enable comparisons of the different angular effects of radiation. The z_{cut} parameter is fixed at 0.1, which is small enough that $\log(z_{\text{cut}})$ resummation is negligible⁹³. Since the Soft Drop jet mass is highly correlated with p_T for light quark and gluon jets, we instead measure the observable $\log_{10}(\rho^2)$, where

$$\rho = m^{\text{soft drop}}/p_T^{\text{ungroomed}}. \quad (8.1)$$

This allows us to measure a dimensionless quantity ρ , which makes the measurement much less dependent on the p_T distribution of the jets*. This measurement is focused on understanding the region where resummation effects are dominant over the non-perturbative effects in the low-mass region, or the fixed-order calculation which dominates the high-mass region. The resummation region approximately covers $-3.7 < \log_{10}(\rho^2) < -1.7$, which motivates the measurement of $\log_{10}(\rho^2)$ instead of ρ .

*The ungroomed jet p_T is used because the groomed version is collinear unsafe when $\beta = 0$ ⁹³

9

Unfolding

In experimental physics, reconstructed observables differ from the theoretical ones in a number of ways, including the reconstruction efficiency, the fake rate, the scale of the observable, and the resolution of an observable. In order to compare a theoretical prediction to a measured quantity, the measured observable can be “unfolded”, which corrects for all of these effects, bringing the unfolded measured result close to the theoretical result. Since this accounts for detector effects, this also allows for comparisons of measurements across different experiments.

For this measurement, the particle-level selection is chosen to be as close as possible to the detector-level selection, minimizing the unfolding corrections. Particle-level jets are reconstructed with the same process as detector-level jets, using simulated particles with mean lifetime $\tau > 30$ ps, and excluding muons and neutrinos, which do

not interact significantly with the detector. These jets must pass the same event selection and dijet requirement as the detector-level jets. The detector-level and truth distributions for the ρ observables are shown in Figure 9.1, where some large differences may be seen between the two distributions, demonstrating the necessity of correcting for detector-level effects.

Unfolding uses Monte Carlo simulations of both the truth and reconstructed observable. If t is the truth value of a particular observable O , the distribution of t can be described as some function or histogram $f(t)$. Similarly, if m is the reconstructed value of the same observable, its distribution may be described by some different function or histogram $g(m)$. Finally, the probability of reconstructing a particular value m , given a truth value t can be described using the function $p(m|t)$.

For any measurement, since there are not infinite Monte Carlo statistics, the result will be measured discretely, with N_t (N_m) truth (reconstructed) bins, each with a bin width of Δt_j (Δm_i), and bin center t_j (m_i). We can then construct the response matrix R_{ij} , which describes the probability of reconstructing a value that falls in bin i , given a truth value in bin j , or written more precisely,

$$R_{ij} = p(m \in \text{bin}_i^{\text{reco}} | t \in \text{bin}_j^{\text{truth}}). \quad (9.1)$$

This matrix can then be used to determine the reconstruction efficiency and the fake rate. The reconstruction efficiency is the probability of a truth value being reconstructed at all within the fiducial cuts of the measurement, and is given by

$$\epsilon_j = \sum_{i=1}^{N_m} R_{ij} \quad (9.2)$$

Similarly, we can define the fake factor, or the probability that a reconstructed value corresponds to any truth value within the space of the measurement, as

$$\epsilon_i = \sum_{j=1}^{N_t} R_{ij} \quad (9.3)$$




Figure 9.1: Distributions of $\log_{10}(\rho^2)$ in data compared to reconstructed detector-level (Reco.) PYTHIA, SHERPA, and HERWIG++, and particle-level (Truth) PYTHIA simulations for $\beta = 0$ (left), $\beta = 1$ (right), and $\beta = 2$ (bottom). The ratio of the three detector-level MC predictions to the data is shown in the middle panel, and the size of the detector $\rightarrow$ particle-level corrections for PYTHIA is shown as the ratio in the bottom panel. The error bars on the data points and in the first ratio include the experimental systematic uncertainties in the cluster energy, angular resolution, and efficiency. The distributions are normalized to the integrated cross-section, σ_{resum} , measured in the resummation region, $-3.7 < \log_{10}(\rho^2) < -1.7$.

The fake and efficiency factors can be seen in Figure 9.2, where these values are close to one for most of the measurement, meaning that the efficiency is very high, and that very few fakes are present. In the lowest p_T bin,

the efficiency is lower, since a measured value of ρ may not pass the p_T requirements for the measurement, and will be counted as an inefficiency, and the fake factor is lower for similar reasons. In addition, the fake and efficiency factors are low for the lowest bin in ρ , since the mass value may fall outside the fiducial cuts.




Figure 9.2: The fake and efficiency correction factors

9.1 ITERATIVE BAYESIAN REGULARIZED UNFOLDING

The goal of unfolding is to find the unfolding matrix Θ_{ij} , which gives the joint probability of finding an truth value in bin j , given a measured value in bin i , such that

$$t_j = \sum_i i P(T_i|M_j) m_i = \sum_i \Theta_{ij} m_i. \quad (9.4)$$

Assuming that the response matrix given by Monte Carlo is not affected by statistical fluctuations and is a complete representation of the detector effects, the inverse of the response matrix could be used to provide this, since

$$t_j = \sum_i R_{ij}^{-1} m_i. \quad (9.5)$$

However, while this will be true for the Monte Carlo samples used to create these matrices, this will bias the results when this is applied to data. The studies shown here use iterative Bayesian unfolding in order to reduce the effects of statistical fluctuations on the unfolded result. Instead of inverting the response matrix, several steps are taken. First, the unfolding matrix is rewritten as

$$\Theta_{ij} = P(T_j|M_i) = \frac{P(M_i|T_j) \cdot P(T_j)}{\sum_j P(M_i|T_j) \cdot P(T_j)} = \frac{R_{ij} \cdot P(T_j)}{\sum_j R_{ij} \cdot P(T_j)}, \quad (9.6)$$

where R_{ij} are the elements of the response matrix, and $P(T_j)$ is the probability of finding an event in bin j of the truth distribution. $P(T_j)$ is referred to as the prior. Note that this equation does not yet use the Monte Carlo distributions for these matrices, and is simply a mathematical expression of the idealized unfolding matrix. The prior can then be estimated by the particle-level (truth) distribution t_j taken from a large Monte Carlo sample. In order to avoid the bias from using the Monte Carlo truth distribution, the unfolding is applied iteratively, each time using the output of the unfolding as the input to the prior for the next iteration. This means, for the n -th iteration of the unfolding,

$$\Theta_{ij}^n = \frac{R_{ij} \cdot P(t_{j,n-1})}{\sum_j R_{ij} \cdot P(t_{j,n-1})}, \quad (9.7)$$

where $t_{i,0} = t_i$, determined directly from Monte Carlo. As the number of iterations increases, the Monte Carlo bias decreases, but the statistical fluctuations become more important. The number of iterations should be chosen to balance these effects, and is typically small. For these results, the number of iterations is 4, and is implemented in the RooUnfold framework⁷.

Figure 9.3 shows the unfolding matrices for each choice of β , inclusive in p_T . Because of the large migrations between detector- and particle-level distributions, as seen in Figure 9.1, there are large off-diagonal terms in the unfolding matrix. This happens particularly for very negative values of $\log_{10}(\rho^2)$, where the mass is smallest, and so the granularity of the detector is not always sufficient to resolve the mass. This effect becomes more obvious for larger values of β , which grooms the jet less aggressively, and therefore is affected by softer collinear splittings that may be difficult for the detector to resolve.

9.2 CONSIDERATIONS IN UNFOLDING

In order to determine the efficacy of the unfolding, several tests are performed to verify that it has worked as intended. Typically, the more diagonal the unfolding matrix, the more accurate the unfolding is, since the necessary corrections are smaller. This means that the binning used for the measurement should have the on-diagonal elements with values larger than around 0.5. For the smallest values of $\log_{10}(\rho^2)$, the diagonality degrades significantly below this, as seen in Figure 9.3, but since the region of interest is the resummation region, which is a bit higher than this, the binning is considered sufficient for these studies.

Once the binning has been chosen, the technical closure of the unfolding is tested, which checks that the unfolding matrix determined in Monte Carlo is able to unfold the same Monte Carlo detector-level measurement to its truth-level measurement. This is checked to avoid pathologies in the unfolding, such as a low correlation between the truth and detector-level distributions that is insufficient for making an unfolding matrix. The results of this are shown for $\beta = 0$ in Figure 9.4, and as expected, the unfolded distribution matches the particle-level distribution.




Figure 9.3: The unfolding matrices for each β parameter

An additional test is also run to determine the sensitivity of the unfolding to the shape of the underlying distribution, which is done using the data-driven closure test. In this test, the shape of the particle-level Monte Carlo distribution is reweighted smoothly, and the detector-level distribution is altered based on folding the truth distribution with the response matrix. The reweighting is chosen such that the detector-level distributions for data and MC match more closely after the reweighting. Then, both the data and the reweighted MC distribution are un-




Figure 9.4: The technical closure for $\beta=0$

folded with the original unfolding matrix. These distributions should then match, since the underlying shape of the detector-level distribution is the same.

Once it is determined that the unfolding procedure is able to correct for detector effects, this can be used to determine the statistical and systematic uncertainties, and produce the final measurements. These steps are described in the following chapters.

10

Systematic Uncertainties

Both statistical and systematic uncertainties must be accounted for in the measurement. The goal of these is to determine the sensitivity of the measurement to various effects, including the amount of data and Monte Carlo that is available, the modeling of the detector simulation, and the modeling of the underlying physics processes. All of the different uncertainties included in this measurement are described below, with a summary of the impact of each uncertainty on the measurement at the end of this chapter.

10.1 STATISTICAL UNCERTAINTIES

The statistical uncertainties must be estimated for both the data as well as the MC. This can be done by creating a number of pseudo-experiments based on the MC prediction, where each bin in the measured distribution is fluctuated based on a Poisson distribution. These fluctuations are propagated to the response matrix, and the result is unfolded. The statistical uncertainty for each bin is then determined based on the inter-quartile range of the unfolded results for that bin.

Additionally, there is a data-driven nonclosure uncertainty in order to estimate the potential bias from a given choice of prior and the number of iterations in the iterative Bayesian method⁹⁰. In this method, the inverse of the response matrix is applied to the particle-level spectrum, which has been reweighted until the folded spectrum agrees with data. This modified detector-level distribution is then unfolded with the nominal response matrix, and the difference between this unfolded result and the reweighted particle-level spectrum is taken as the nonclosure uncertainty.

10.2 SYSTEMATIC UNCERTAINTIES

Various systematic uncertainties contribute to the overall uncertainty of this measurement. This come from effects such as the ability of the Monte Carlo to model the jet fragmentation and the detector response to these effects. The different systematics that are considered for this measurement are described below.

10.2.1 THEORETICAL MODELING UNCERTAINTIES

The theoretical modeling uncertainty arises from the jet fragmentation (QCD Modeling), and is one of the dominant uncertainties. To estimate the size of this uncertainty, the PYTHIA generator is used for the nominal sample, and comparisons are made with SHERPA and HERWIG++. To be specific, the SHERPA and HERWIG++ reconstructed distributions are unfolded using the PYTHIA unfolding matrix, and the results are compared to the particle-level distributions for SHERPA and HERWIG++ respectively. The full difference between the unfolded re-

sult and the original particle-level distribution is taken to be the uncertainty. The results from both SHERPA and HERWIG++ are compatible, so only the variation with SHERPA is used to compute the systematic uncertainty.

10.2.2 CLUSTER UNCERTAINTIES

The experimental uncertainties on the shape of the mass distribution are given by propagating uncertainties on the topoclusters through to the jet mass distribution. This includes uncertainties on the cluster energy scale and resolution, position, and their reconstruction efficiency.

CLUSTER ENERGY SCALE AND RESOLUTION UNCERTAINTIES

The uncertainty is determined as follow: isolated clusters are matched to tracks, giving a distribution of the cluster energy to track momentum ratio (E/p). This is done in bins of η and p , in order to account for effects of how the detector responds to different particles at different momenta as well as the geometry of the detector. An attempt is made to fit this E/p distribution with a Gaussian distribution. For low momentum bins, this fit may fail, in which case the mean and standard deviation of the overall distribution are used instead. This distribution will naturally have a peak at zero, due to effects from the cluster reconstruction efficiency. Since this is accounted for elsewhere, the value of this bin is set to the average of the bin above and below. It is noted that in some cases, E/p may fall below one, due to the shaping of the LAr calorimeter response, as described in Section 4.2.1.

An example of this fit in one bin of η and p is shown in Figure 10.1. The mean and standard deviation of these distributions are calculated in data and Monte Carlo, and the results are smoothed across momentum, where it is expected to change smoothly. The data and Monte Carlo values for the smoothed mean value of E/p are shown in Figure 10.2.

For any clusters that have $|\eta| > 0.6$ and $p > 30$ GeV, and all clusters with $p > 350$ GeV, a flat 10% uncertainty is used for both the energy scale and resolution, motivated by earlier studies³². Then, the systematic uncertainty at a given η and p bin is taken by the ratio of the data and MC distributions, as shown in Figure 10.3. It is difficult to obtain a pure sample of isolated cluster-track pairs at high momentum. In order to augment this informa-

tion, results from the combined test beam are used for particles with $30 < p < 350$ GeV. Above this, a flat 10% uncertainty is applied for both the scale and resolution of the clusters.

These uncertainties are then propagated to the jet level as follows. For the cluster energy scale, all clusters in an event are shifted up or down based on the data-MC ratio shown in Figure 10.3, while for the cluster energy resolution, the energy of all clusters is smeared based on the ratio of the width of the distributions in data and MC.




Figure 10.1: The E/p distribution for central $|\eta|$ and at high p in log (right) and linear (left) scale.

CLUSTER ANGULAR RESOLUTION UNCERTAINTY

The systematic uncertainty on the cluster angular resolution is derived using $Z \rightarrow \mu\mu$ events, where $p_T^Z > 30$ GeV, and where a large fraction of clusters are isolated. Then, non-muon tracks are selected that have at most one cluster within a $\Delta R < 0.15$ cone around their position, extrapolated to the second layer of the calorimeter. The $\Delta R_{\text{cluster,track}}$ is plotted in data and MC, as shown in Figure 10.4. This distribution has two peaks, particularly




Figure 10.2: The smoothed mean E/p for data (left) and MC (right).

for low-momentum particles, but only the first peak is relevant to the single particle uncertainty, while the second peak arises from the cluster isolation requirement. In order to quantify the effects on the single particle angular resolution, only the region with $\Delta R_{\text{cluster, track}} < 0.075$ is used for the data/MC comparisons. An example of the extracted values of the resolution in the barrel region for data and Monte Carlo are shown in Figure 10.5, where the differences are typically less than 1 mrad. This uncertainty was measured at $\sqrt{s} = 7$ TeV instead of $\sqrt{s} = 13$ TeV, but the effect of this is expected to be small, based on a comparison of results at $\sqrt{s} = 7$ TeV and $\sqrt{s} = 8$ TeV. In order to be conservative, a 5 mrad smearing is used for the measurements at $\sqrt{s} = 13$ TeV.

CLUSTER RECONSTRUCTION EFFICIENCY UNCERTAINTY

For low- p_T particles, the amount of detector material in front of the calorimeter may be sufficient prevent the particle from having enough energy to seed a cluster in the calorimeter. In order to account for this, an uncertainty




Figure 10.3: The ratio of the E/p (a) means, and (b) widths with smoothing.

is included that compares the cluster reconstruction efficiencies in data and MC. The efficiency is determined in single hadrons by measuring the percentage of tracks with an associated cluster. An example of the efficiency as a function of particle momentum is shown in Figure 10.6.

Then, the difference between data and MC are fit with the function $ae^{b \times p} + be^{-c \times p^2}$, with the requirement that this goes to zero at 30 GeV, since all particles above this momentum should have sufficient energy to seed a cluster. Examples of these fits are shown in Figure 10.7. Based on these fits, the uncertainty is estimated by randomly dropping clusters based on the value of these functions for a given bin of η and p .

10.2.3 PILEUP UNCERTAINTY

Since the Soft Drop algorithm is designed to remove soft and wide-angle radiation from a jet, it is not expected for pileup to have a large effect on the measurement. Additionally, since this measurement is done using the 2016




Figure 10.4: Left (Right): The ΔR between isolated low (high) momentum tracks and clusters in $Z \rightarrow \mu\mu$ events in the barrel of the detector

dataset, it is not as sensitive to pileup as one that uses the full Run 2 dataset, where the pileup is higher on average. Nevertheless, it is still important to measure the uncertainty from pileup, since it could have some effect on the measurement. The sensitivity of the unfolding procedure to pileup is determined by reweighting events to vary the distribution of the number of interactions in the MC simulation by 10%. The difference between these results and the nominal is taken to be the uncertainty, which is very small.

10.3 SUMMARY OF UNCERTAINTIES

The uncertainties for the inclusive p_T measurement are shown for all values of β in Figure 10.8. They are dominated by QCD modeling and the cluster energy scale. The modeling uncertainties are largest ($\lesssim 20\%$) at low masses where the non-perturbative effects are largest. In the resummation region and fixed-order region, the modeling uncertainties are less than 10%. In the low mass region, the cluster energy scale uncertainty is relatively large,




Figure 10.5: Left (Right): The RMS of the $\Delta\eta$ ($\Delta\phi$) between isolated single particle tracks and clusters for tracks extrapolated to the second layer of the calorimeter in the barrel of the detector.

since these jets typically only have a few clusters, making them sensitive to the scale of each cluster. In the high mass region, the cluster energy scale uncertainties dominate. In this region, the mass of the jet is driven by the energy of the hard prongs, instead of their opening angle. As expected, this measurement is not sensitive to either pileup effects, or the statistical uncertainties on either data or Monte Carlo.




Figure 10.6: The probability for finding a calorimeter cell cluster matched to a track as a function of the track momentum in various bins of η .




Figure 10.7: The cluster reconstruction efficiency uncertainty (data/MC difference) as a function of the track momentum in various bins of η .




Figure 10.8: The breakdown of systematic and statistical uncertainties as a function of $\log_{10}(\rho^2)$ for $\beta = 0$.

11

Measurements of Jet Substructure Observables for Soft Dropped Jets

With the event selection and unfolding method finalized, and with the determination of the relevant uncertainties, it is possible to unfold the detector-level distribution for $\log_{10}(\rho^2)$. As described in Section 8, the $\log_{10}(\rho^2)$ and p_T distributions are simultaneously unfolded for three different values of the Soft Drop β parameter. In each case, the unfolded distribution in the resummation region, between $-3.7 < \log_{10}(\rho^2) < -1.7$. The unfolded distribution for each p_T bin is shown in Figure 11.1, where the different p_T bins have been normalized by a different factor indicated in the legend, in order to compare them all on the same plot. The shape of the $\log_{10}(\rho^2)$ distribution changes

very slowly with p_T , since this observable was chosen over measuring the jet mass directly for this behavior. Because of this, it is meaningful to consider the inclusive distribution instead of individual p_T bins. This is done by unfolding the data simultaneously in $\log_{10}(\rho^2)$ and p_T , summing the results, and then normalizing the inclusive distribution. The unfolding is still done differentially in order to obtain the most accurate results.




Figure 11.1: The unfolded $\log_{10}(\rho^2)$ distribution in data for each of the p_T bins in the analysis, overlaid with a multiplicative factor indicated in the legend.

Before looking at the results, it is worth reminding of the different regions of the distribution, where each region is dominated by different physics. Generally, *non-perturbative effects* are large for more negative values of $\log_{10}(\rho^2)$ (lower mass), which is where the small-angle and soft-gluon emissions dominate. This region is approx-

imately covered by $\log_{10}(\rho^2) < -3.7$, though this value depends somewhat on both the value of β and the jet p_T . Above this is the *resummation region*, which is the focus of this measurement, which lies approximately in the range $-3.7 < \log_{10}(\rho^2) < -1.7$. Above this, for $\log_{10}(\rho^2) > -1.7$, *fixed-order* effects dominate, based on large-angle gluon emissions.

11.1 COMPARISON OF UNFOLDED DATA TO MONTE CARLO PREDICTIONS


The unfolded data, inclusive in p_T , and normalized in the resummation region are shown in Fig. 11.2, where they are compared to MC predictions from PYTHIA, SHERPA, and HERWIG++ generators, as described in Section 8.

Generally, the Monte Carlo generators agree well with the unfolded data, which is expected since the Soft Drop algorithm removes a large amount of the non-perturbative effects that are hard to model. As β increases, the fraction of radiation removed by soft-drop grooming decreases and the impact of non-perturbative effects grows larger^{70 69}, resulting in worsened agreement between data and MC for higher values of β in the low-mass region. For $\beta = 1, 2$, the data fall below the MC predictions in the more negative values of $\log_{10}(\rho^2)$. PYTHIA models the data well throughout the entire region, but agreement between the data and all the MC generators generally remains within uncertainties for all values of β .

11.2 COMPARISON OF UNFOLDED DATA TO ANALYTICAL PREDICTIONS


Finally, the unfolded data may be compared to the analytical predictions, and these results are shown in Figure 11.3.

Two of the lines are from the NLO+NLL prediction from Ref.^{93,94}, where NLO+NLL shows the pure calculation, while NLO+NLL+NP includes non-perturbative corrections. These corrections are determined using an average of change in the distribution in various MC models from turning non-perturbative effects off and on, and the uncertainty is given by the envelope of these effects⁹³. The LO+NNLL prediction from Refs.^{69 70} is shown with two sets of markers. The closed markers are drawn in the regions where the calculation is expected to accurate, while the open markers are drawn where the non-perturbative effects are expected to be large. Both the NLL and NNLL calculations use NLOJet++^{53,97} (MG5_aMC⁹) with the CT14nlo⁶³ (MSTW2008LO) PDF set for matrix ele-



(a) $\beta = 0$

(b) $\beta = 1$



(c) $\beta = 2$

Figure 11.2: The unfolded $\log_{10}(\rho^2)$ distribution for anti- k_t $R = 0.8$ jets with $p_T^{\text{lead}} > 600$ GeV, after the soft drop algorithm is applied for $\beta \in \{0, 1, 2\}$, in data compared to PYTHIA, SHERPA, and HERWIG++ particle-level, and NLO+NLL+NP⁹³ and LO+NNLL^{70 69} theory predictions. The LO+NNLL calculation does not have non-perturbative (NP) corrections; the region where these are expected to be large is shown in a open marker, while regions where they are expected to be small are shown with a filled marker. All uncertainties described in the text are shown on the data; the uncertainties from the calculations are shown on each one. The distributions are normalized to the integrated cross section, σ_{resum} , measured in the resummation region, $-3.7 < \log_{10}(\rho^2) < -1.7$. The NLO+NLL+NP cross-section in this resummation regime is 0.14, 0.19, and 0.21 nb for $\beta = 0, 1, 2$, respectively⁹³.

ment calculations. The uncertainties due to the analytical calculation come from independently varying each of the renormalization, factorization, and resummation scales by factors of 2 and 1/2.

The analytical calculations are expected to be accurate in different regions of $\log_{10}(\rho^2)$ ^{93 70 69}. For all values of β , the measured and predicted shapes agree well in the resummation region, demonstrating the accuracy of the

predictions. The resummation region, should have the most reliable prediction for the LO+NNLL analytical calculation, while the NLO+NLL calculation should also be accurate for the fixed-order region. This effect is seen slightly in the high-mass region of these distributions, particularly for $\beta = 0$. At more negative values of $\log_{10}(\rho^2)$, non-perturbative effects lead to distinctly different predictions between the calculations without NP corrections; the data fall below the predictions for all β values. Interestingly, the NNLL calculation is not exclusively a better model of the data than the NLL calculation in the resummation regime and NP effects can also be comparable to the higher order resummation corrections in this regime. Therefore, improved precision for the future will require a deeper understanding of different perturbative choices in the calculations as well as NP effects beyond PS MC.




Figure 11.3: The unfolded $\log_{10}(\rho^2)$ distribution for anti- k_t $R = 0.8$ jets with $p_T^{\text{lead}} > 600$ GeV, after the soft drop algorithm is applied for $\beta \in \{0, 1, 2\}$, in data compared to PYTHIA, SHERPA, and HERWIG++ particle-level, and NLO+NLL+NP⁹³ and LO+NNLL^{70 69} theory predictions. The LO+NNLL calculation does not have non-perturbative (NP) corrections; the region where these are expected to be large is shown in an open marker, while regions where they are expected to be small are shown with a filled marker. All uncertainties described in the text are shown on the data; the uncertainties from the calculations are shown on each one. The distributions are normalized to the integrated cross section, σ_{resum} , measured in the resummation region, $-3.7 < \log_{10}(\rho^2) < -1.7$. The NLO+NLL+NP cross-section in this resummation regime is 0.14, 0.19, and 0.21 nb for $\beta = 0, 1, 2$, respectively⁹³.

12

Summary of the Jet Mass Measurement

This measurement of the jet mass enabled the first comparisons of unfolded data to analytical predictions for jet substructure observables beyond leading logarithmic accuracy. The analytical calculations are able to model the behavior observed in data for the region of the measurement where they are expected to be accurate, giving confidence in this type of prediction. This type of prediction was not expected to be possible when the field of jet substructure was just beginning, 10 years ago, and indicates both a deep understanding of quantum chromodynamics as well as an experimental understanding of the ATLAS detector.

These reasons make a measurement of the jet mass interesting, but the implications for this extend far beyond measuring single jet substructure observables. Confidence in the analytical methods for these calculations enables

other calculations of more fundamental parameters, and the experimental precision means that these can be measured rigorously. There are two measurements that are particularly interesting. First, there has been significant work towards a robust calculation of the top mass distribution in the boosted regime, where the top decays hadronically⁷⁷. In this prediction, large- R jets are groomed using the Soft Drop algorithm, which allows this measurement to be factorized. Current measurements of the top mass extract the mass using Monte Carlo-based kinematics, and are able to achieve precision of less than 1 GeV. However, this method requires an additional uncertainty based on the top mass scheme uncertainty, making the interpretation of these measurements challenging. The Soft Drop measurement would avoid this uncertainty, and if the top mass could be measured with precision of 2-3 GeV, it could provide an alternative measurement of the top mass which is theoretically robust and simpler to interpret. The other fundamental parameter of the Standard Model which may be possible to measure using jet substructure is the strong coupling constant α_s ¹⁰, where differences in the mass distribution for quarks and gluons may make it possible to measure this quantity, possibly to 5-10% precision. Progress still needs to be made both towards a sufficiently accurate prediction as well as improving experimental uncertainties, but it is an interesting possibility for the future.

Part II

Searching for Long-Lived Particles using Displaced Vertices

The second part of this thesis documents the steps towards searching for supersymmetry in events with jets and a displaced vertex. While this search is not yet completed, this section documents the progress made towards the search, including theoretical motivation, sensitivity studies, and preliminary background estimates. A search for displaced vertices plus jets has not been done on ATLAS since Run 1 of data-taking, and there are many compelling reasons to revisit this search. Most obviously, the LHC is running at higher energies, meaning that a new search would be able to constrain a wider array of models. Similarly, Run 2 of the LHC has taken significantly more data than in Run 1, increasing the power of the search. On the theoretical side, there have also been developments in the modeling of long-lived R-hadrons, which has now been included in the ATLAS simulation, making the interpretation of the results more robust.

Long-lived particles typically require special object reconstruction, and recently, there has been significant effort towards improving the reconstruction of displaced vertices. These developments improve both the acceptance of signal vertices as well as the rejection of fake vertices, increasing the sensitivity of the search. However, these changes also have a significant effect on the background estimation strategy, since the current strategy is highly dependent on object reconstruction. This means that the background estimation techniques need to be rethought and validated, which is particularly challenging given the unusual nature of these backgrounds. This, in addition to the simulation of R-hadrons, is the focus of this chapter. While this search is not complete, these studies are important steps towards creating a robust analysis.

13

Supersymmetry

As discussed in Chapter 1, the Standard model provides a robust framework for describing the known particles and their interactions, with the exception of gravity, which has a negligible effect at collider energy-scales. However, it is not a complete theory of everything. At the Planck scale, quantum gravity will need to be included within the model, though there is still no known way of doing this. Since the Planck scale is on the order of 10^{19} GeV, while current colliders are able to probe energies on the order of 10^4 GeV, on its own, this is not immediately concerning. However, there are still a number of other reasons to believe that the Standard Model is incomplete, and that there is more physics yet to be discovered at scales below the Planck scale.

One of the motivations for physics beyond the Standard Model comes from the structure of the Standard Model

itself in the form of the *hierarchy problem*. In a quantum field theory, particle masses and couplings *run* with energy, meaning that their strengths change as a function of the energy scales. This running can be calculated for each mass or coupling separately, but the running of the Higgs boson mass are particularly noteworthy. Given some UV cutoff Λ_{UV} , the Higgs mass receives corrections in the form of

$$\Delta m_H^2 = -\frac{|\lambda|^2}{8\pi^2} \Lambda_{\text{UV}}^2 + \dots, \quad (13.1)$$

where the other corrections grow at most logarithmically with Λ_{UV} . If there are no other corrections up to the Planck scale (meaning that $\Lambda_{\text{UV}} = M_{\text{Planck}}$), then the corrections to the Higgs mass are 30 orders of magnitude larger than the Higgs mass itself, implying that the Higgs mass has been fine-tuned.

In addition to the theoretical motivations for physics beyond the Standard Model, there is also cosmological evidence of physics that is not included in the Standard Model. The evidence for this ranges from things such as galactic rotations to the anisotropies of the cosmic microwave background, which are most easily explained by the presence of some sort of particle referred to as *dark matter*. However, while these observations do indicate a new type of matter, they do not provide a straightforward theory for what the physics of it is.

All of this evidence indicates that the Standard Model is incomplete, and a number of theories have been developed in order to explain various combinations of these open questions. One theory that has gained a lot of traction is that of supersymmetry, based in part on its ability to explain all of these features. Supersymmetry encompasses a broad array of models, and a full review of its theory is not covered here, and the rest of this chapter will focus on its basic principles as well as a few specific supersymmetric models that are particularly relevant to searches for long-lived particles.

13.1 THE BASICS OF SUPERSYMMETRY

The basic principle of supersymmetry is that there is a symmetry that links fermionic and bosonic states. This results in an operator Q that transforms a bosonic state into a fermionic state, and vice versa, which changes only the spin state of the particle. Under this assumption, its theory may be described more simply using *supermulti-*

plets, which are sets of bosons and fermions linked by this symmetry. Since Q only changes the spin-state of the particle, each supermultiplet must contain the same number of degrees of freedom for the bosonic and fermionic states. In order to satisfy this, a Weyl fermion must be in a supermultiplet with two real scalars, which will be referred to here as a *matter* supermultiplet. while a spin-1 boson will be in a supermultiplet with a single Weyl fermion, and is referred to as a *gauge* supermultiplet. Additionally, if a spin-2 graviton is included, it will be in a supermultiplet with a spin-3/2 particle. Finally, the complex doublet scalar fields corresponding to the Higgs boson must be included in a matter supermultiplet. A direct inclusion of a single Higgs supermultiplet results in a gauge anomaly, and so any supersymmetric extension of the Standard Model must include at least two Higgs supermultiplets.

Any supersymmetric extension of the Standard Model requires adding a significant number of particles to the theory. In order to ease interpretation, a simple naming scheme is used by relating the Standard Model particles to their superpartners. The superpartners of Standard Model fermions are typically referred to by adding an *s* for *scalar* to the beginning of the particle name, such as the *stop*, which is the supersymmetric partner of the top quark, and will sometimes be written with a tilde, such as $\tilde{t}$. Similarly, the supersymmetric partners of gauge bosons will sometimes be referred to as *gauginos*, and the individual particles will be referred to by adding *ino* as a suffix to the end, such as *gluino* or *wino*.

Supersymmetry is a broad theory, and there are many models that can be written following its basic principles. The Minimal Supersymmetric Standard Model (MSSM) aims to simplify this by writing a theory in which the fewest interactions and particles can be consistently added to the Standard Model. The interactions allowed by the MSSM may be summarized by its superpotential,

$$W_{MSSM} = \bar{u}y_uQH_u - \bar{d}y_dQH_d - \bar{e}y_eLH_d + \mu H_uH_d, \quad (13.2)$$

where $\bar{u}$, H_u , $\bar{d}$, Q , H_d , $\bar{e}$, L , and H_d are all matter superfields, while y_u , y_d , and y_e are the Yukawa coupling parameters, and μ provides a mass term for the Higgs fields. It is noted that superpotential assumes the conservation of R-parity, which is discussed more in Section 13.2.1.

One interesting consequence of supersymmetry is the solution to the problem of the Higgs mass corrections. Fermionic and bosonic corrections to the Higgs mass enter with an opposite sign, the corrections from the superpartners will exactly cancel the quadratic terms from the Standard Model particles, leaving only terms which are logarithmic in Λ_{UV} . In addition, if the number of supersymmetric particles is conserved using R-parity, a concept discussed more in Section 13.2.1, the lightest supersymmetric particle will not be able to decay, and is a possible candidate for dark matter. Even so, there are significant problems with the theory. Most obvious is the fact that it does not exist as an unbroken theory, since the superpartners of Standard Model particles would have the same masses as the Standard Model particles, which is not been observed. On its own, this does not exclude the possibility of supersymmetry, but it does mean that if it is an exact symmetry, it must be broken spontaneously such that the vacuum state of the model is not supersymmetric. Many such models exist that preserve some of the appealing features of unbroken supersymmetry, and supersymmetry breaking can also enable a mechanism for electroweak symmetry breaking.

13.2 LONG-LIVED PARTICLES

Many searches on ATLAS focus on detecting final states with objects which originate from the interaction point, including objects such as jets, leptons, photons, or missing transverse momentum. While these searches are important, they are also broadly insensitive to a broad range of models which include long-lived particles. Frequently, the object definitions include cuts which have reduced efficiency for long-lived particles or their decay products, such as cuts on tracks which require that they point roughly towards the interaction point. These cuts help reduce noise from fakes and pileup, but also result in losing sensitivity to a large number of models which include particles with a noticeable lifetime. In order to regain sensitivity to these models, it is important to have dedicated searches and reconstruction techniques.

The next few chapters focus on a search for long-lived particles decaying within the inner detector of the ATLAS experiment. In this context, the term *long-lived particles* generically refers to particles with a lifetime that is observable with the ATLAS detector. This may mean anything from particles decaying shortly after the hard-

scatter interaction, whose decay products can be used to point to an origin that is not the interaction point, to particles that have a long enough lifetime to pass through the ATLAS detector without decaying, or particles that are stopped within the detector and decay after the window for detection. A number of searches exist at the ATLAS experiment that are focused on long-lived particles. Some of these look for the particles directly by looking for unusual signatures within the detector, such as highly ionizing particles¹⁸, or disappearing tracks that start at the interaction, but disappear partway through the tracker due to a decay²¹. Other searches look for the decay products of long-lived particles, which include searches for non-pointing photons²², jets with unusual energy deposits²⁷, searches like this for secondary vertices^{4,28}, and other similar signatures.

There are a number of ways in which long-lived particles can occur in a model, and in fact, several examples exist within the Standard Model itself. Typically, in order to have a significant lifetime, a particle must have either small couplings to other particles, small mass splittings with other particles, or be connected to a softly broken symmetry. These features can exist generically, but there are a few particular SUSY models that commonly result in long-lived particles due to constraints of the theory. These models include R-parity violating SUSY, split SUSY, and gauge-mediated supersymmetry breaking, each of which is discussed briefly below.

13.2.1 R-PARITY VIOLATING SUSY

The most generic versions of a supersymmetric theory include terms in the form of

$$W_{\Delta L=1} = \frac{1}{2}\lambda^{ijk}L_iL_jE_k^c + \lambda'^{ijk}L_iQ_jD_k^c + \mu^iL_iH_u, W_{\Delta B=1} = \frac{1}{2}\lambda''^{ijk}U_i^cD_j^cD_k^c, \quad (13.3)$$

where L_i and Q_i are respectively the lepton and quark superfields, E_i^c , U_i^c , and D_i^c are respectively the charged lepton, up quark, and down quark superfields, and H_u is the higgs superfield. The couplings λ_{ijk} , λ'_{ijk} , and μ_i are responsible for L-violating couplings, while λ''_{ijk} is responsible for the B-violating couplings⁹². These terms respectively allow for the lepton and baryon number to be changed by one, which are both quantities conserved by the Standard Model.

However, in many models, these couplings are required to be zero, which was motivated by both experimental

observation and theoretical assumptions. Experimentally, non-zero values of these terms can lead to various phenomena that are already tightly constrained. For instance, if both λ and λ'' or λ' and λ'' are non-zero, this can result in the proton decay, where a squark could mediate the decay of a proton into a lepton and π^0 ⁹². Additional constraints exist on other combinations of couplings based on similar reasoning⁴². This leads to the assumption in SUSY models of R-parity, which is the conservation of

$$RP = (-1)^{(3B+L+2S)}, \quad (13.4)$$

where B is the baryon number, L is the lepton number, and S is the spin. Under this symmetry, all SM particles have even values of R -parity, while all sparticles have odd values, meaning that sparticles must be produced in pairs, and sparticles cannot decay entirely to SM particles. Theoretically, if sparticles are able to decay to exclusively SM particles, this would eliminate the possibility of a stable sparticle that could be a dark matter candidate, with the exception of a super-light gravitino which has not yet been excluded. For these reasons, these terms are required to be zero in many SUSY models.

While there are some constraints on the choices of R-parity violating (RPV) couplings that may exist, the couplings may be non-zero within the existing constraints, and it is possible to build R-parity violating models that are able to answer some of the open questions about the Standard Model. A couple of observations are possible. First, many SUSY searches rely on searching for large values of missing transverse momentum based on the assumptions about the existence of a lightest supersymmetric particle (LSP). However, for RPV SUSY, this signature is not as ubiquitous, and so existing RPC SUSY searches do not set strong limits on RPV SUSY models. For instance, many RPV SUSY models include signatures with many jets, from the hadronic decay of the lightest supersymmetric particle. In addition, since the R-parity violating couplings are required to be small, this may also result in sparticles with a measurable lifetime.

13.2.2 SPLIT SUPERSYMMETRY

Split supersymmetry relaxes the requirement of naturalness, and assumes that the SUSY scalars are significantly heavier than the SUSY fermions, with the exception of one Higgs, which is finely-tuned to be light. While these heavy scalars will be inaccessible to collider searches, the gaugino masses are still allowed to be light due to a chiral symmetry, and are taken to have masses near the weak scale⁷⁶. This results in very different phenomenological models, since more standard theories result in cascading decays of sparticles mediated by the scalar sparticles. Instead, decays to electroweak gauginos are suppressed based on the virtuality of the squark. With the reduced phase space for decay, this model frequently results in a long-lived gluino, and is approximately given by

$$\tau_{gluino} = 8 \left(\frac{m_S}{10^9 \text{GeV}} \right)^4 \left(\frac{1 \text{TeV}}{m_{\tilde{g}}} \right)^5 s, \quad (13.5)$$

where m_S is the scalar mass. Depending on this value, the lifetime can range from barely measurable to the lifetime of the universe, meaning that there are a broad range of relevant signatures⁷⁶.

13.2.3 GAUGE-MEDIATED SUPERSYMMETRY BREAKING

Finally, in gauge-mediated supersymmetry breaking (GMSB), where the supersymmetry breaking terms arise from gauge interactions⁷⁹ instead of gravity. This is accomplished by adding chiral supermultiplets that couple to the (s)fermions and higgs(inos). In this model, the next-to-lightest supersymmetric particle (NLSP) is the lightest neutralino χ_1^0 , the lightest stau, or very rarely, the lightest sneutrino. Then, the average distance traveled by the NLSP before decaying is given by

$$L = \frac{1}{K_\gamma} \left(\frac{100 \text{GeV}}{m} \right)^5 \left(\frac{\sqrt{F/k}}{100 \text{TeV}} \right)^4 \sqrt{\frac{E^2}{m^2} - 1} \times 10^{-2} \text{cm}, \quad (13.6)$$

where

$$K_{\gamma, \text{neutralinos}} = |N_{11} \cos(\theta_W) + N_{12} \sin(\theta_W)| K_{\gamma, \text{staus}} = 1, \quad (13.7)$$

where m and E is the mass and energy of the NLSP, θ_W is the Weinberg angle, and F is the scale of supersymmetry breaking inside supermultiplets, and $k = F/F_0$, where F_0 is the fundamental scale of supersymmetry breaking. Depending on the value of $\sqrt{F/k}$, the NLSP may have a measurable lifetime relevant for long-lived particle searches^{74 95}.

13.3 R-HADRONS

Much like their Standard Model counterparts, squarks and gluinos are not color singlets. In the case where they decay promptly, this does not result in any modifications to their behavior, but when their lifetime is sufficiently long, they will form color-singlet objects that are referred to as *R-hadrons*. This adds several complications to modeling their behavior, which needs to be accounted for properly in order to obtain meaningful limits on a particular SUSY model. These complications are described below, with a focus on the considerations that are able to impact this analysis.

13.3.1 R-HADRON MASS SPECTRUM

One interesting aspect of the R-hadron simulation is the mass-spectrum of the R-hadrons being produced. Unlike for Standard Model hadrons, there is no sample of R-hadrons that can be used to determine the spectrum of R-hadrons. The mass spectrum currently being used in Pythia8 is based on the approximation that the mass of a (generic) hadron may be estimated by

$$m_{\text{hadron}} = \sum_i m_i - k \sum_{i \neq j} \frac{(F_i \cdot F_j)(S_i \cdot S_j)}{m_i m_j}, \quad (13.8)$$

where m_i are the masses of the constituent partons, F_i represents the SU(3) color matrix, S_i is the SU(2) spin matrix, and k is chosen to be 0.043 GeV³ for mesons, and 0.026 GeV³ for baryons, based on measurements of Standard Model hadron masses⁸⁸. There are newer calculations for the R-hadron mass spectrum, based on lattice calculations⁶⁷, but these are currently not included in the simulation. However, the search presented here is not heavily dependent on the mass spectrum, so the Pythia8 version is considered sufficient for these studies.

13.3.2 INTERACTION OF R-HADRONS WITH THE DETECTOR

Since R-hadrons have a non-zero lifetime, it is important to consider the simulation of how they will interact with the detector as they travel. Unlike for the standard particles observed at the LHC, there is no dataset that can be used to measure their interactions. Although this means that there will be some uncertainty on the specifics of their interactions, it is possible to produce a rough picture of how they will interact, based on our understanding from Standard Model hadrons. Two general models are considered for this: the generic and triple Regge models, which are described below. In the generic model, the scattering cross-section is assumed to be constant per light quark, meaning that the cross-section of interactions does not depend on the Lorentz-boost of the R-hadron.

The triple Regge model was originally formulated to describe the dynamics of generic heavy hadrons⁶⁰, and has since been included in GEANT4 as a possible interaction for R-hadrons. Since the cross-section of a heavy parton and a target nucleon is proportional to $\frac{1}{m^2}$, the dynamics of the R-hadron will be dominated by the light quark system, which also implies that R-baryons are almost twice as likely to interact as R-mesons, since there are twice as many light quarks⁸⁹. In this model, it is also possible for an R-meson to become an R-baryon, with a probability of around 10%. The opposite transformation is not allowed in interactions with ordinary detector material due to conservation of baryon number⁸⁸. This is because the transformation of an R-meson into an R-baryon happens through quark exchange with a proton, but the opposite process cannot occur without mesons, which are not ordinarily present. Additionally, this may result in the oscillation of R-hadrons to their antiparticle. The model implemented in GEANT4 uses measurements of low-energy hadron-hadron to predict this behavior for R-hadrons⁸⁹. Figure 13.1 shows the cross-sections for interactions for gluino baryons and mesons in both the generic and triple Regge models as a function of their Lorentz factor γ .

13.3.3 R-HADRON DECAY

The final complication of R-hadrons is their decay. There are two main components of this decay. First, the sparticle itself will decay, which is determined by the SUSY model being considered. Once this is done, the remaining partons need to be treated properly, including their showering and fragmentation. Several Monte Carlo genera-




Figure 13.1: The cross-section for interactions for gluino R-hadrons in the generic (dashed lines) and triple Regge (solid lines) models⁸⁹.

tors exist that can account for showering and hadronization, but Pythia8 is used for these studies, since it already includes the necessary information to account for the treatment of R-hadrons, and so the same generator may be used for R-hadron production and decay.

13.3.4 SUMMARY OF R-HADRON SIMULATION

With these considerations in mind, the process of producing Monte Carlo samples for this analysis proceeds as follows. Events are produced with Madgraph plus Pythia8, which contain detector-stable particles, including R-hadrons. These particles are then propagated through the detector using GEANT4, until the decay of the R-hadron. At this point, the R-hadron is decayed with Pythia8, using the same parameters as in the generation step. The decay products are hadronized with Pythia8, and the stable particles are passed back to GEANT4, where they are propagated through the detector.

13.4 SIGNAL MODEL FOR THE SEARCH

As discussed towards the beginning of the chapter, there are a wealth of models that are able to produce long-lived particles in association with jets. For the purposes of this analysis, the search will be done in the context of only one of these models: R-parity violating SUSY, where the neutralino is long-lived, which is shown in Figure 13.2. In this model, there are three relevant parameters that are varied: the mass of the long-lived gluino, the mass of the neutralino, and the lifetime of the neutralino. Monte Carlo samples are produced with four different neutralino lifetimes of 0.01 ns, 0.1 ns, 1 ns, and 10 ns. The gluino mass ranges from 1600 to 2600 GeV, while the neutralino mass ranges from 10 to 2000 GeV, and Figure 13.3 shows the grid of signal samples produced for different gluino and neutralino masses. The mass splitting between the gluino and the neutralino is related to the prompt jets produced in association to the gluino decay, and the mass of the neutralino affects the energy of the quarks that it decays to.




Figure 13.2: A diagram showing the signal model for this search. The model is an RPV SUSY model, where each gluino decays promptly to two quarks plus a neutralino, and the neutralino has a long lifetime based on the small λ'' coupling.

There are currently plans to include additional models, as well as provide material for reinterpretation

of the results of this search. With this in mind, the analysis cuts are made as general as possible within the constraints of the ATLAS data and background estimation techniques.




Figure 13.3: The different grid points produced for the signal samples. The gluino masses range from 1600 to 2600 GeV, while the neutralino masses range from 10 to 2500 GeV. Since the gluino decays to two quarks plus a neutralino, the mass of the neutralino is required to be smaller than the gluon mass.

14

Event and Vertex Selection

This analysis uses the full Run-2 dataset, which has a combined luminosity of around 139 fb^{-1} , meaning that it will be able to set stronger constraints on models than the Run 1 analysis²⁸ both because of the higher energy of collisions as well as the increased amount of data. There are a number of complications in processing the data for this analysis. As mentioned in Section 5.1.1, large-radius tracking requires too much computing power to be run on every event. Instead, for data, the algorithm is run on a set of events that pass a pre-selection in order to filter fully reconstructed offline events, similar to applying another trigger selection, but with access to the fully reconstructed event. This selection is designed to apply to several analyses, including the one described in this thesis, and so it includes several filters on a variety of reconstructed objects, including jets, muons and missing transverse

Table 14.1: The lowest unscaled trigger thresholds by data-taking year

Year	3-jet threshold	4-jet threshold	5-jet threshold	6-jet threshold	7-jet threshold
2015	175 GeV	85 GeV	55 GeV	45 GeV	45 GeV
2016	200 GeV	100 GeV	70 GeV	60 GeV	45 GeV
2017	200 GeV	120 GeV	85 GeV	70 GeV	45 GeV
2018	200 GeV	120 GeV	85 GeV	70 GeV	45 GeV

momentum. The jet-related filters are described here in more detail for clarity.

14.1 TRIGGER AND DRAW REQUIREMENTS

While displaced vertices may occur in a number of different event topologies, this search is focused on displaced vertices that are produced in association to jets. This motivates triggering on events that pass a multi-jet trigger. In all cases, events are required to pass the lowest unscaled trigger. The thresholds for these triggers change each year based on the different beam conditions, and the different thresholds may be seen in Table 14.1. In order to understand the impact of the trigger selection on the relevant signals, we test the efficiency of signal events passing the lowest unscaled triggers from 2017. This result is shown in Figure 14.1, which shows the efficiency of passing any trigger as a function of several different parameters of the model. As can be seen from these plots, the trigger efficiency is very high for all parameter choices, though some efficiency begins to be lost for very large neutralino lifetimes. In these cases, the lifetime may be large enough that the LLP does not decay in time to produce a large energy deposit in the calorimeter, lowering the trigger efficiency.

As mentioned in Section 5.1.1, large-radius tracking requires a significant amount of computing resources such that running it on all events is unfeasible. In order to deal with this, events are filtered prior to reconstruction. This filter is called the *derived RAW filter*, or *DRAW filter*, based on the stage in event reconstruction where this filter is applied. The goal of the filter is to remove as many background events while still retaining a high signal efficiency. However, this is challenging for a search for a multi-jet signature at a hadron collider. A two-step approach is taken in order to reduce rates but also retain a high signal efficiency.

Events may pass the DRAW filter either by passing an offline jet requirement in addition to a trackless jet re-




Figure 14.1: The efficiency of passing any unrescaled multi-jet trigger for several different neutralino lifetimes, gluino masses, and neutralino masses

quirement, or by passing a stricter jet requirement. The offline jet requirement combined with the trackless jet requirement is referred to as the *DV multi-jet trackless filter*, and is based on the idea that jets originating from long-lived particles have fewer tracks associated to the primary vertex. A trackless jet is defined as a jet where $\sum_{tracks} p_T^{track}$ is less than 5 GeV, and $|\eta| < 2.5$, and events must have either one trackless jet with p_T above 70 GeV, or two trackless jets with p_T above 50 GeV. In addition, these events are required to pass a multijet filter,

Table 14.2: The jet thresholds for the trackless multi-jet requirement and high- p_T multi-jet requirement

Jet multiplicity	Trackless multi-jet	High- p_T multi-jet
2-jet	220 GeV	500 GeV
3-jet	120 GeV	180 GeV
4-jet	100 GeV	220 GeV
5-jet	75 GeV	170 GeV
6-jet	45 GeV	50 GeV
7-jet	45 GeV	45 GeV

which requires at least N jets each with p_T above some threshold X .

$$\text{DVmulti-jet trackless filter} = ((1j_{\text{trackless}}||2j_{\text{trackless}})\&\&((2j\&\&3j)||4j||5j||6j||7j)\&\&(\text{multi-jet trigger})). \quad (14.1)$$

The multi-jet trackless jet filter is not fully efficient for all signal samples. In particular, it loses some sensitivity for small mass splittings between the gluino and neutralino, as well as for long lifetimes. To counteract this, events may also pass the DRAW filter by passing a pure multi-jet filter. In order to keep the event rate low enough, the multi-jet p_T thresholds are higher than for the trackless jet filter, but are still low enough to significantly increase the signal efficiency of events passing the DRAW filter.

$$\text{DV multi-jet filter} = ((2j_h\&\&3j_h)||4j_h||5j_h||6j_h||7j_h), \quad (14.2)$$

where j_h refers to the higher jet p_T thresholds. Putting this together, events must pass either the DV multi-jet trackless filter or the DV multi-jet filter.

The exact multi-jet thresholds depend on the data-taking year, but a sample of these may be seen in Table 14.2. In addition, the filter efficiency for signal events may be seen in Figure 14.2, which shows a high efficiency across all different parameters for the RPV SUSY model of interest.




Figure 14.2: The efficiency of passing the DRAW filter for several different neutralino lifetimes, gluino masses, and neutralino masses

14.2 EVENT AND VERTEX SELECTION

14.2.1 THE MATERIAL MAP

In order to remove most hadronic interactions, a map of dense material regions is created. This map is then used to veto vertices that are located within these dense material regions, and are therefore likely to be an interaction with the material instead of the decay of a long-lived particle. In order to create this map, secondary vertices are reconstructed in data. To remove contributions from K_{SS} , vertices in the mass range of $0.45 \text{ GeV} < m_{DV} < 0.55$

GeV are removed. In addition, vertices are required to have a mass less than 2.5 GeV in order to reduce the effects of random track crossings as well as to avoid sensitivity to the signal region.

In the region of $R_{xy} < 150$ mm, the number of reconstructed vertices is sufficient to construct the map purely from data. The fact the material in the detector is roughly periodic in ϕ can be used to improve the precision of the map by folding in ϕ . This is done for each tracking layer individually, and the layers are first corrected for small distortions of the detector that cause small translations in R_{xy} . The IBL is folded 7 times, while the pixel layers are folded 11, 19, and 13 times respectively. After folding, the results of the map are smoothed, and a threshold of number of reconstructed vertices is applied to determine which regions are considered dense material. In the outer region of the detector ($R_{xy} > 150$ mm), the material map is constructed from Monte Carlo and validated in data.

The distributions of displaced vertices in (x, y) is shown in Figure 14.3 before and after the material veto, and in (z, R_{xy}) in Figure 14.4.




Figure 14.3: Distributions of displaced vertices in data in the x and y plane. Data are required to have at least one baseline muon with $|d_0| > 1.0$ mm. (a) Before applying a material veto. (b) after applying a material veto




Figure 14.4: Distributions of displaced vertices in data in the x and y plane. Data are required to have at least one baseline muon with $|d_0| > 1.0$ mm. (a) Before applying a material veto. (b) after applying a material veto

14.2.2 EVENT AND VERTEX SELECTION

The complete event and vertex selection for this search has not yet been finalized, the studies shown here use a vertex selection based on the Run 2 search for displaced vertices plus missing transverse momentum⁴ in combination with basic trigger efficiency studies. The standard quality cuts will not be changed, though the precise offline jet thresholds, in addition to the selection of displaced vertices may be slightly affected. The offline jet threshold cuts are based the requirement of the jet trigger being fully efficient at those cuts, and should not be significantly affected. The displaced vertex requirements are a result of the current background estimation techniques, and cannot be loosened significantly using the current strategy.

All data used in these studies must pass a *good runs list*, which is designed to filter out data-taking periods with problematic detector conditions. Events must pass the lowest un-prescaled trigger for that year. Events must also pass jet cleaning requirements, which are designed to filter out events with beam induced backgrounds. These cleaning requirements are designed specifically for long-lived particles, since the standard cuts apply timing requirements that cut away the signal. Events must pass the offline jet selection described in the previous section.

For the vertices, two different selections that are applied. The first is the *pre-selection*, which is used to study the overall behavior of displaced vertices, while the full selection is called the *post-selection*, which is the selection applied for the signal region. Since the expected backgrounds are very small, both selections are necessary in order to estimate the background contributions. The vertex pre-selection is designed to be a loose selection, and includes the following requirements:

1. The vertex must satisfy $R_{xy} < 300$ mm and $|z_{DV}| < 300$ mm.
2. The vertex fit must be of good quality, such that $\chi^2/N_{\text{DOF}} < 5$
3. The vertex must be at least 4 mm away from any reconstructed primary vertex in the transverse plane
4. The vertex must pass the cuts designed to eliminate wrongly associated tracks, as described in Section 5.1.1
5. No value of the covariance matrix may be larger than 1000, which indicates a poor knowledge of the vertex position
6. The vertex must not be within a material-dense region, as determined by the material map

If the vertex passes these requirements, then it is considered for studies of pre-selected events. In a few cases, vertices that pass all of these requirements except for the material map veto will be used, but this is explicitly stated when this is done. For the final selection of the signal region, each vertex must pass the pre-selection, and additionally have at least five associated tracks, and a reconstructed invariant mass of at least 10 GeV.

14.3 PROSPECTS FOR THE SEARCH

Before discussing the details of the background estimation, it is useful to do a short prospective study on the sensitivity of this search. Typically, searches for DV+X are done in signal regions with approximately zero background, so we make the assumption that it is possible to exclude regions of phase space where at least 3 events are expected. The sensitivity may be studied using production cross sections, which have been calculated at NNLO_{approx} + NNLL accuracy in the context of the MSSM³⁸. These cross sections can then be used, in combination with the amount of data, and the expected efficiencies of the signal region in order to estimate the sensitivity of the search. During Run 2, the LHC has delivered 156 fb^{-1} of data, of which 139 fb^{-1} is usable for ATLAS physics results, and this is the number that will be used for these studies.

A few basic selections are required in order to study the sensitivity. First, the signal efficiencies are calculated for events that pass the *pre-selection* requirements, which require the event to pass the jet trigger, the DRAW filter, and the basic DV selections. Then, the efficiency of passing the *post-selection* is also studied, which requires the event to pass both the pre-selection, as well as additional requirements on the displaced vertex. These efficiencies are studied for a few different values of lifetime of the neutralino. By studying the sensitivity of both of these requirements separately, it is possible to see how much could be gained in these models by loosening the requirements on the mass and number of tracks associated to the DV.

Figures 14.5 and 14.5 shows the expected number of events after various selections for a sample of different particle lifetimes. As can be seen in these figures, the DV post-selection has by far the greatest effect on the expected sensitivity of this search. This is particularly true for models with a small neutralino mass, but is true as well for higher neutralino masses.

Loosening the post-selection requirements could significantly improve the limits for certain signal models. In particular for smaller neutralino masses, loosening both the N-track and vertex mass requirement could increase the limits on the gluino mass by hundreds of GeV as seen in Figure 14.7, while this effect decreases for larger neutralino masses (Figure 14.8 where the displaced vertices will be both more massive and have more associated tracks. In both cases, loosening the N-track requirement tends to have a larger effect than loosening the mass requirement.

14.4 BACKGROUNDS

There are three main backgrounds that must be accounted for in this search. The first of these, *hadronic interactions*, is a result of interactions of hadrons produced in the event with the detector material. This background is dramatically reduced by vetoing regions of the detector with dense material, though the efficiency of this is not 100%. The other two backgrounds, *merged vertices* and *accidental crossings*, are the result of the algorithms used to reconstruct displaced vertices. Figure 14.9 shows a cartoon of each of these types of backgrounds for clarity. In the current plans for this analysis, each of these backgrounds is estimated separately. The following chapters will document the technique used to estimate the background for hadronic interactions as well as preliminary inves-






Figure 14.5: The expected number of events after various selections for several different gluino lifetimes, and a small neutralino mass, with a red line at 3 events.


tigations towards estimating the merged vertex background, while the accidental crossings background will not be covered. It is noted that both of these are significant alterations of the background estimations for the Run-1 DV+jets search, due to considerations in the underlying physics for the hadronic interactions, and changes in the vertex merging for the merged vertex background. In addition, a few preliminary studies of alternative background estimation techniques are shown in order to test the possibility of moving away from the strategy of creating three separate background estimations.




(a) $\lambda_{\chi_1} = .01 \text{ ns}$



(b) $\lambda_{\chi_1} = .1 \text{ ns}$



(c) $\lambda_{\chi_1} = 1 \text{ ns}$



(d) $\lambda_{\chi_1} = 10 \text{ ns}$

Figure 14.6: The expected number of events after various selections for several different gluino lifetimes, and a large neutralino mass, with a red line at 3 events.







Figure 14.7: The expected number of events after various selections for several different gluino lifetimes, and a small neutralino mass, with a red line at 3 events




(a) $\lambda_{\chi_1} = .01 \text{ ns}$



(b) $\lambda_{\chi_1} = .1 \text{ ns}$



(c) $\lambda_{\chi_1} = 1 \text{ ns}$



(d) $\lambda_{\chi_1} = 10 \text{ ns}$

Figure 14.8: The expected number of events after various selections for several different gluino lifetimes, and a small neutralino mass, with a red line at 3 events




Figure 14.9: A cartoon demonstrating the three types of expected backgrounds, where (a) shows an accidental crossing, (b) shows a hadronic interaction, and (c) shows a merged vertex.

15

Background Estimation for Hadronic Interactions

One of the major sources of background for long-lived particles is interactions of hadrons with detector material, which can produce particles whose origin is displaced from the hard-scatter interaction. While these are typically low energy, with 96% of hadrons having momentum less than 10 GeV and an average of 4 GeV, the sheer number of hadrons produced at the ATLAS experiment means that these will still populate the signal region. Since these are the result of interactions with the detector material, the bulk of these may be removed by discarding vertices that occur in dense regions of the detector, which is the purpose of the material map described in Section 14.2.1. Even with this veto, there can still be interactions in the less dense regions of the detector, inefficiencies in the material map, or vertices not vetoed due to the resolution of the vertex position. The number of vertices that are not

vetoed by the map are estimated using a template fit derived from Monte Carlo events.

In previous results, this background had been estimated using an exponential fit of the vertex mass distribution, and taking the integral of the exponential above the signal region mass. However, there was not a solid reason to believe that the hadronic interaction mass distribution should fit an exponential, and large uncertainties were associated to this method. For some of the results shown in this chapter, the results of the template are compared to that of the exponential fit in order to demonstrate the improvement over the previous technique.

15.1 CONSTRUCTION OF A MONTE CARLO TEMPLATE

To estimate the contributions of the remaining vertices, Monte Carlo simulations may be used to construct a template of the shape of the mass distribution for hadronic interactions. For this template, it is important to have a pure sample of displaced vertices that originate from a hadronic interaction, and that do not include accidental crossings. This can be done using the truth associations of tracks to determine whether the tracks are correctly associated. Since truth particles from pileup interactions are not stored in Monte Carlo simulations, three classes of displaced vertices are defined: pure hadronic interactions, hadronic interactions with one or more accidental crossings, and unknown. To classify a vertex, each of its associated tracks is looped over as follows:

1. A track is matched to its truth particle. If there is no truth particle link, then this is a fake track, and is an accidental crossing.
2. If this track is matched to a truth particle, its truth particle link is followed. If the link is unavailable, the track is matched to a pileup truth particle, and the track is classified as pileup.
3. If the truth particle is available, then the production point of the truth particle is compared to the origin of the displaced vertex. If these are more than X mm away from each other, the track is considered an accidental crossing, where X is a parameter of the method.

Any vertices with no pileup tracks and no accidental crossings is considered a pure hadronic interaction. Any vertices without pileup tracks but with any number of hadronic interactions are considered hadronic interactions with accidental crossings. All vertices with any number of pileup tracks are considered to be of unknown origin, and are not used for either template. Two templates are reconstructed: one for hadronic interactions, and one for all vertices that do not contain pileup. Only the first is used for the background estimation, but both are used to validate the method. For clarity, a few examples of this classification are shown in a cartoon in Figure 15.1.

The definition of X can be tuned to give a conservative estimate of which tracks are accidental crossings, in order to provide a pure template of hadronic interactions. Since this technique is not being used to construct an accidental crossing template, a conservative definition is suitable for these studies. In order to determine the value of X to use, several different choices are compared, and the resulting vertex mass distribution for pure hadronic interactions is compared in Figure 15.2. Based on these distributions, it is clear that values of X less than around 2 mm do not significantly change the shape of the template, while very large values of X such as 10 mm, more vertices begin to be included in the tails, indicating the inclusions of high-mass vertices from accidental crossings. Based on these results, a value of 0.5 is chosen as a conservative definition of an accidental crossing.

15.2 VALIDATION IN DATA

Once the hadronic interaction template has been created, it is necessary to validate its performance in data. The templates are created in bins of R_{xy} in order to avoid correlations between the radius of production and the mass distribution, with bins approximately corresponding to regions of the detector. In particular, there are 13 total layers at [4, 22, 25, 29, 38, 46, 73, 84, 111, 120, 145, 180, 300], where the first layer is the region before any dense material, the second corresponds to the beampipe, and the next layers alternate between material-sparse regions and layers of the tracking detector. Validating the templates requires a fairly pure sample of hadronic interactions, which is difficult to acquire since the tails of the vertex mass distribution tend to be populated by vertices with accidental crossings. This becomes even more problematic in high pileup conditions, where the greater number of reconstructed tracks makes it more likely that an accidental crossing is included in a given vertex. This effect may be seen in Figure 15.3, where the shape of the vertex mass distribution changes dramatically with the number of reconstructed primary vertices. Even in zero pileup conditions, some accidental crossings will be present, but the effect is much smaller.

Requiring $N_{PV} < 10$ minimizes the pileup effects while still leaving sufficient statistics for comparing to Monte Carlo. There will still be some residual pileup effects, but it is enough to determine whether the templates are reasonable. In addition, since there are still accidental crossings when no pileup is present, the hadronic interaction

template will still not match the vertex mass distribution in data. Instead, we can use the template with all vertices without pileup tracks in order to determine whether the Monte Carlo template models the data correctly, and then use the hadronic interaction template to obtain the background estimate.

Before estimating the background, it is important to determine the region in which to scale the templates to the data. Ideally, this should be done in a regime where accidental crossings do not have a significant impact on the shape of the results. Based on the shape of the mass distributions in different bins of pileup, it is clear that this value depends in part on how many tracks are in the vertex; vertices with fewer tracks tend to be affected at lower masses than those with more tracks. However, normalizing in the region $0.5 \text{ GeV} < m_{DV} < N \text{ GeV}$, where N is the number of tracks, should avoid significant biases for all vertices.

Several checks are performed in order to verify that the template derived in Monte Carlo is valid for the data. It is important to avoid unblinding the results, and so a multi-step procedure is used. First, the template is compared to data in events that fail the material veto. This region will be dominated by the effects of material interactions, and so they are not sensitive to small fluctuations from a signal. The results of the background estimate in this region are compared to the observed number of events in the signal region to test the overall performance of the method. In addition, to test that the templates are also valid outside of dense material regions, the results are also compared to data from a single period, which is a small enough sample to avoid problems with unblinding. The results for these studies are described in the following sections.







Figure 15.1: Cartoons demonstrating the definition of accidental crossings. (a) is defined as a pure hadronic interaction from the hard-scatter event, and is used in the template. (b) One truth particle has an origin that is displaced from the vertex position, and is classified as an accidental crossing. (c) A hadronic interaction, but from pileup tracks, which is not able to be used in the template. (d) A hadronic interaction with an accidental crossing, where one track is associated to a truth pileup particle, and is therefore unable to be used in the accidental crossing template.




(a) $N_{\text{track}} = 3$



(b) $N_{\text{track}} = 4$




(c) $N_{\text{track}} = 5$




(d) $N_{\text{track}} > 5$


Figure 15.2: The hadronic interaction template for several different distances X .




(a) $N_{\text{track}} = 3$



(b) $N_{\text{track}} = 4$




(c) $N_{\text{track}} = 5$




(d) $N_{\text{track}} > 5$


Figure 15.3: The NPV dependence of the vertex mass in data




(a) $N_{\text{track}} = 3$



(b) $N_{\text{track}} = 4$



(c) $N_{\text{track}} = 5$




(d) $N_{\text{track}} > 5$


Figure 15.4: The cumulative distribution function of the vertex mass distribution in different bins of N_{PV} in data

15.2.1 VALIDATION IN MATERIAL REGIONS


In order to obtain a large sample of hadronic interactions without unblinding the results, we can use the vertex mass distribution in material-dense regions. In these regions, the effects of a signal would not be visible, since hadronic interactions are numerous. The results of this can be seen in Figure 15.5, where the pure mass distribution is shown, and Figure 15.8, where the cumulative distribution of the mass distribution is shown. In order to test the method, it is first useful to create the background estimate for the low-pileup region of N_{PV} less than 10. As discussed before, this distribution should correspond fairly well to the template that includes accidental crossings from the hard-scatter event. These distributions may be seen in Figure 15.5 for a material-dense region, using a subset of the 2016 dataset. As shown in this figure, the templates model the data fairly accurately, and tables comparing the estimated number of vertices to the observed number of vertices are shown in Tables 15.1- 15.4. The templates predict a number of vertices that is on the same order of magnitude as what is actually observed. Especially since the data still contains some pileup, which will increase the number of observed vertices, this indicates the basic functionality of the templates. Since the additional vertices in the tails of the mass distribution are from accidental crossings or merged vertices, the fact that the template that includes all hard-scatter DVs is consistently lower than the number of observed vertices is expected when pileup is present.




(a) $N_{\text{track}} = 3$



(b) $N_{\text{track}} = 4$




(c) $N_{\text{track}} = 5$




(d) $N_{\text{track}} > 5$


Figure 15.5: The vertex mass distribution fit with the two templates. The red lines show the region of normalization.




(a) $N_{\text{track}} = 3$



(b) $N_{\text{track}} = 4$



(c) $N_{\text{track}} = 5$



(d) $N_{\text{track}} > 5$

Figure 15.6: The vertex mass distribution fit with the two templates. The red lines show the region of normalization.

Table 15.1: Background estimate for material-dense regions using low-pileup data for $nTrack = 3$

	Data	Had Int Template	Accidental Template
$22 \leq R_{xy} \leq 25$	1.000 ± 1.000	0.403 ± 0.152	1.921 ± 0.325
$29 \leq R_{xy} \leq 38$	11.000 ± 3.317	0.956 ± 0.276	3.090 ± 0.483
$46 \leq R_{xy} \leq 73$	15.000 ± 3.873	0.500 ± 0.204	4.076 ± 0.565
$84 \leq R_{xy} \leq 111$	9.000 ± 3.000	0.203 ± 0.143	2.389 ± 0.469
$120 \leq R_{xy} \leq 145$	20.000 ± 4.472	0.000 ± 0.000	2.134 ± 1.067
$180 \leq R_{xy} \leq 300$	47.000 ± 6.856	0.000 ± 0.000	1.576 ± 1.115

Table 15.2: Background estimate for material-dense regions using low-pileup data for $nTrack = 4$

	Data	Had Int Template	Accidental Template
$22 \leq R_{xy} \leq 25$	4.000 ± 2.000	0.488 ± 0.147	1.414 ± 0.239
$29 \leq R_{xy} \leq 38$	2.000 ± 1.414	0.738 ± 0.213	3.335 ± 0.431
$46 \leq R_{xy} \leq 73$	8.000 ± 2.828	0.466 ± 0.165	3.044 ± 0.400
$84 \leq R_{xy} \leq 111$	4.000 ± 2.000	0.371 ± 0.166	1.171 ± 0.269
$120 \leq R_{xy} \leq 145$	5.000 ± 2.236	0.000 ± 0.000	1.966 ± 0.879
$180 \leq R_{xy} \leq 300$	3.000 ± 1.732	0.000 ± 0.000	1.238 ± 0.875

Table 15.3: Background estimate for material-dense regions using low-pileup data for $nTrack = 5$

	Data	Had Int Template	Accidental Template
$22 \leq R_{xy} \leq 25$	2.000 ± 1.414	0.405 ± 0.122	1.456 ± 0.215
$29 \leq R_{xy} \leq 38$	5.000 ± 2.236	1.057 ± 0.225	2.096 ± 0.293
$46 \leq R_{xy} \leq 73$	6.000 ± 2.449	0.615 ± 0.159	1.558 ± 0.235
$84 \leq R_{xy} \leq 111$	5.000 ± 2.236	0.247 ± 0.123	0.841 ± 0.198
$120 \leq R_{xy} \leq 145$	3.000 ± 1.732	0.000 ± 0.000	1.015 ± 0.508

Table 15.4: Background estimate for material-dense regions using low-pileup data for $nTrack \geq 6$

	Data	Had Int Template	Accidental Template
$22 \leq R_{xy} \leq 25$	6.000 ± 2.449	4.979 ± 0.339	9.615 ± 0.421
$29 \leq R_{xy} \leq 38$	7.000 ± 2.646	4.680 ± 0.383	8.101 ± 0.442
$46 \leq R_{xy} \leq 73$	6.000 ± 2.449	2.472 ± 0.256	4.986 ± 0.323
$84 \leq R_{xy} \leq 111$	0.000 ± 0.000	0.232 ± 0.095	0.957 ± 0.162
$120 \leq R_{xy} \leq 145$	2.000 ± 1.414	0.355 ± 0.355	1.197 ± 0.489

This exercise may be repeated without the restriction on N_{PV} , which is shown in Figure 15.7. The template with all non-pileup vertices models the data well, though some residual nonclosure is observed in the high mass region. The template works better for vertices with more tracks, likely due to the fact that pileup has a smaller

effect in this region. Tables 15.5- 15.8 show the estimated number of events from each template, compared to the observed number of events. Since more pileup exists in these samples, the observed number of events is consistently larger than the predicted number of events, which is expected since these high-mass vertices originate from one of the other backgrounds, and should not be included in this estimate. However, the shape of the distributions match fairly well in the normalization region, indicating that this method is still reasonable.







Figure 15.7: The vertex mass distribution fit with the two templates. The red lines show the region of normalization.




(a) $N_{\text{track}} = 3$



(b) $N_{\text{track}} = 4$



(c) $N_{\text{track}} = 5$



(d) $N_{\text{track}} > 5$

Figure 15.8: The CDF of the vertex mass distribution fit with the two templates. The red lines show the region of normalization.

Table 15.5: Background estimate for material-dense regions for $n\text{Track} = 3$

	Data	Had Int Template	Accidental Template
$22 \leq R_{xy} \leq 25$	140.000+/-11.832	5.430+/-2.052	25.881+/-4.375
$29 \leq R_{xy} \leq 38$	470.000+/-21.679	12.042+/-3.476	38.903+/-6.076
$46 \leq R_{xy} \leq 73$	783.000+/-27.982	5.707+/-2.330	46.507+/-6.449
$84 \leq R_{xy} \leq 111$	1014.000+/-31.843	2.086+/-1.475	24.592+/-4.823
$120 \leq R_{xy} \leq 145$	2559.000+/-50.587	0.000+/-0.000	24.531+/-12.266
$180 \leq R_{xy} \leq 300$	17811.000+/-133.458	0.000+/-0.000	42.545+/-30.084

Table 15.6: Background estimate for material-dense regions for $nTrack = 4$

	Data	Had Int Template	Accidental Template
$22 \leq R_{xy} \leq 25$	54.000+/-7.348	5.905+/-1.780	17.113+/-2.893
$29 \leq R_{xy} \leq 38$	157.000+/-12.530	8.185+/-2.363	37.002+/-4.777
$46 \leq R_{xy} \leq 73$	289.000+/-17.000	4.958+/-1.753	32.357+/-4.249
$84 \leq R_{xy} \leq 111$	216.000+/-14.697	3.875+/-1.733	12.239+/-2.808
$120 \leq R_{xy} \leq 145$	463.000+/-21.517	0.000+/-0.000	22.844+/-10.216
$180 \leq R_{xy} \leq 300$	3306.000+/-57.498	0.000+/-0.000	37.897+/-26.797

Table 15.7: Background estimate for material-dense regions for $nTrack = 5$

	Data	Had Int Template	Accidental Template
$22 \leq R_{xy} \leq 25$	31.000+/-5.568	4.778+/-1.441	17.191+/-2.535
$29 \leq R_{xy} \leq 38$	66.000+/-8.124	10.896+/-2.323	21.606+/-3.025
$46 \leq R_{xy} \leq 73$	84.000+/-9.165	6.617+/-1.708	16.755+/-2.526
$84 \leq R_{xy} \leq 111$	63.000+/-7.937	2.486+/-1.243	8.462+/-1.994
$120 \leq R_{xy} \leq 145$	93.000+/-9.644	0.000+/-0.000	11.731+/-5.865

Table 15.8: Background estimate for material-dense regions for $nTrack \geq 6$

	Data	Had Int Template	Accidental Template
$22 \leq R_{xy} \leq 25$	73.000+/-8.544	61.412+/-4.179	118.590+/-5.196
$29 \leq R_{xy} \leq 38$	76.000+/-8.718	50.809+/-4.162	87.946+/-4.798
$46 \leq R_{xy} \leq 73$	85.000+/-9.220	26.666+/-2.765	53.788+/-3.479
$84 \leq R_{xy} \leq 111$	31.000+/-5.568	2.372+/-0.968	9.777+/-1.653
$120 \leq R_{xy} \leq 145$	41.000+/-6.403	3.169+/-3.169	10.686+/-4.363

15.2.2 VALIDATION WITH SINGLE PERIOD

The first set of plots were done for material-dense regions, where there are no significant concerns of the results being sensitive to new signals. Similarly, it is possible to look outside of these regions, as long as only a small portion of the data is considered. With this in mind, one way of validating these results is to use data from a single run period to test if the results from the background estimate are consistently smaller than the number observed in data. For these studies, data is taken from Period D of 2015. Two examples of the template fits are seen in Figure 15.9 and 15.10, where it can be seen that the templates do a good job of matching the data in the low mass region, even after the material veto has been applied. As seen in Tables 15.9- 15.12, the number of hadronic interactions predicted by the templates is consistently lower than the number of observed vertices. However, since contribution

from hadronic interactions is expected to be subdominant to both the accidental crossing and merged vertex backgrounds in the high mass region, this is the expected behavior.







Figure 15.9: The vertex mass distribution in data, and the template fit, using Period D of 2015. The red lines show the region of normalization.




(a) $N_{\text{track}} = 3$



(b) $N_{\text{track}} = 4$



(c) $N_{\text{track}} = 5$



(d) $N_{\text{track}} > 5$

Figure 15.10: The vertex mass distribution in data, and the template fit, using Period D of 2015. The red lines show the region of normalization.

Table 15.9: Background estimate for 2015 Period D for nTrack = 3

	Data	Had Int Template	Accidental Template
$4 \leq R_{xy} \leq 22$	209.000+/-14.457	0.000+/-0.000	4.697+/-2.348
$22 \leq R_{xy} \leq 25$	1.000+/-1.000	0.015+/-0.006	0.076+/-0.013
$25 \leq R_{xy} \leq 29$	7.000+/-2.646	0.000+/-0.000	0.640+/-0.320
$29 \leq R_{xy} \leq 38$	2.000+/-1.414	0.025+/-0.007	0.083+/-0.013
$38 \leq R_{xy} \leq 46$	8.000+/-2.828	0.038+/-0.022	0.184+/-0.047
$46 \leq R_{xy} \leq 73$	24.000+/-4.899	0.050+/-0.020	0.422+/-0.057
$73 \leq R_{xy} \leq 84$	19.000+/-4.359	0.000+/-0.000	5.435+/-2.431
$84 \leq R_{xy} \leq 111$	19.000+/-4.359	0.054+/-0.038	0.591+/-0.121
$111 \leq R_{xy} \leq 120$	14.000+/-3.742	0.000+/-0.000	1.174+/-1.174
$120 \leq R_{xy} \leq 145$	33.000+/-5.745	0.000+/-0.000	1.065+/-0.532
$145 \leq R_{xy} \leq 180$	93.000+/-9.644	0.000+/-0.000	0.000+/-0.000
$180 \leq R_{xy} \leq 300$	216.000+/-14.697	0.000+/-0.000	1.874+/-1.325

Table 15.10: Background estimate for 2015 Period D for nTrack = 4

	Data	Had Int Template	Accidental Template
$4 \leq R_{xy} \leq 22$	60.000+/-7.746	0.000+/-0.000	2.651+/-1.875
$22 \leq R_{xy} \leq 25$	0.000+/-0.000	0.016+/-0.005	0.050+/-0.008
$25 \leq R_{xy} \leq 29$	0.000+/-0.000	0.000+/-0.000	0.241+/-0.241
$29 \leq R_{xy} \leq 38$	0.000+/-0.000	0.007+/-0.002	0.036+/-0.004
$38 \leq R_{xy} \leq 46$	0.000+/-0.000	0.013+/-0.013	0.036+/-0.021
$46 \leq R_{xy} \leq 73$	3.000+/-1.732	0.039+/-0.014	0.263+/-0.034
$73 \leq R_{xy} \leq 84$	4.000+/-2.000	0.000+/-0.000	1.176+/-1.176
$84 \leq R_{xy} \leq 111$	4.000+/-2.000	0.081+/-0.036	0.255+/-0.059
$111 \leq R_{xy} \leq 120$	1.000+/-1.000	0.000+/-0.000	0.000+/-0.000
$120 \leq R_{xy} \leq 145$	2.000+/-1.414	0.000+/-0.000	0.840+/-0.376
$145 \leq R_{xy} \leq 180$	5.000+/-2.236	0.000+/-0.000	0.000+/-0.000
$180 \leq R_{xy} \leq 300$	12.000+/-3.464	0.000+/-0.000	0.823+/-0.582

Table 15.11: Background estimate for 2015 Period D for nTrack = 5

	Data	Had Int Template	Accidental Template
$4 \leq R_{xy} \leq 22$	27.000+/-5.196	0.000+/-0.000	3.541+/-2.504
$22 \leq R_{xy} \leq 25$	0.000+/-0.000	0.034+/-0.010	0.144+/-0.019
$25 \leq R_{xy} \leq 29$	0.000+/-0.000	0.000+/-0.000	0.000+/-0.000
$29 \leq R_{xy} \leq 38$	0.000+/-0.000	0.010+/-0.002	0.021+/-0.003
$38 \leq R_{xy} \leq 46$	2.000+/-1.414	0.000+/-0.000	0.011+/-0.006
$46 \leq R_{xy} \leq 73$	0.000+/-0.000	0.031+/-0.008	0.088+/-0.013
$73 \leq R_{xy} \leq 84$	0.000+/-0.000	0.000+/-0.000	0.000+/-0.000
$84 \leq R_{xy} \leq 111$	1.000+/-1.000	0.078+/-0.039	0.266+/-0.063
$111 \leq R_{xy} \leq 120$	3.000+/-1.732	0.000+/-0.000	0.348+/-0.348
$120 \leq R_{xy} \leq 145$	1.000+/-1.000	0.000+/-0.000	0.427+/-0.213
$145 \leq R_{xy} \leq 180$	1.000+/-1.000	0.000+/-0.000	0.000+/-0.000

Table 15.12: Background estimate for 2015 Period D for $nTrack \geq 6$

	Data	Had Int Template	Accidental Template
$4 \leq R_{xy} \leq 22$	19.000+/-4.359	0.000+/-0.000	1.523+/-1.523
$22 \leq R_{xy} \leq 25$	0.000+/-0.000	0.102+/-0.007	0.200+/-0.009
$25 \leq R_{xy} \leq 29$	1.000+/-1.000	1.000+/-0.577	1.421+/-0.580
$29 \leq R_{xy} \leq 38$	0.000+/-0.000	0.092+/-0.008	0.163+/-0.009
$38 \leq R_{xy} \leq 46$	0.000+/-0.000	0.036+/-0.010	0.070+/-0.013
$46 \leq R_{xy} \leq 73$	2.000+/-1.414	0.138+/-0.015	0.297+/-0.019
$73 \leq R_{xy} \leq 84$	0.000+/-0.000	0.000+/-0.000	0.000+/-0.000
$84 \leq R_{xy} \leq 111$	1.000+/-1.000	0.048+/-0.020	0.214+/-0.035
$111 \leq R_{xy} \leq 120$	0.000+/-0.000	0.000+/-0.000	0.000+/-0.000
$120 \leq R_{xy} \leq 145$	2.000+/-1.414	0.142+/-0.142	0.473+/-0.193
$145 \leq R_{xy} \leq 180$	1.000+/-1.000	0.000+/-0.000	1.000+/-1.000
$180 \leq R_{xy} \leq 300$	1.000+/-1.000	0.000+/-0.000	0.000+/-0.000

16

Merged Vertex Background Estimation

One of the backgrounds for long-lived particles is *merged vertices*, or distinct sets of vertices that have been merged together based on the vertexing algorithm. As described in Section 5.2.2, there are several ways that vertices may be merged together, which is done to improve the signal efficiency of the algorithm. However, this also results in the merging of some unrelated vertices. These studies represent a preliminary investigation into a background estimation technique, and are not intended to provide a complete background estimate. However, since the technique used in Run 1 is not viable due to changes in the vertex merging algorithm, these are important steps towards having a method to estimate this background.

Given a random distribution of displaced vertices in an event, then the distribution of distances between vertices


in the same event should match the distribution for different vertices taken from different events as well. However, if uncorrelated vertices are merged together, then this will result in a deficit of nearby vertices in the same event. While these assumptions are imperfect, it is possible to test the basic premise by comparing these distributions.

This hypothesis may be easily tested in data. As usual, the standard event selection is applied, and vertices are required to pass the standard pre-selection. The distribution of distance significance is plotted for pairs of vertices both with 2 tracks (2+2), pairs where one has 2 tracks and the other has 3 (2+3), and pairs where both vertices have 3 tracks (3+3). The result of this is shown in Figure 16.1 for same-event and mixed-event pairs of vertices. As shown, the two distributions match fairly well above distance significance of 100, where they should match if the assumptions are correct, since vertex merging only occurs below distance significance of 100, as described in Section 5.2.2. However, at significance below 100, the two distributions begin to disagree, indicating that vertices are being merged within the same event. This demonstrates that it may be possible to use the ratio of these two distributions to estimate the number of vertices that have been merged, and this ratio will be referred to as a template.

In addition, it is noted that the templates for (2+2), (2+3), and (3+3) track pairs are all remarkably similar, as seen in Figure 16.2. While it is still important to study the dependence of these distributions on various vertex properties, these distributions motivate further study of both the distributions and the background estimation technique.


16.1 TEMPLATE CREATION

In order to obtain a more realistic estimate of the merging, the templates should take into account the factors that make a vertex more or less likely to be merged with another. One natural consideration is the proximity of a DV to a jet. There are likely to be more displaced vertices where there is more jet activity, meaning that vertices close to jets should be more likely to be merged. The ratio of (normalized) same-event to merged-event pairs is shown as a function of $\min(\Delta R_{DV,jet})$ for (2, 2)-track and (2, 3)-track events in Figure 16.3, where each horizontal row is normalized separately. Based on this, it is clear that this dependence changes slowly as a function of ΔR , with a cutoff around 0.7. This motivates creating the templates in a few bins of ΔR in order to account for this depen-



(a)

(b)



(c)

Figure 16.1: The distance significance between same-event and mixed-event vertices.

dence.

Another parameter that can impact frequency of merging is the radial displacement of the vertex (R_{xy}). More vertices tend to be produced towards the center of the detector, making it more likely for vertices to be merged. We can study this dependence by plotting the same ratio as a function of R_{xy} using the same bins as for the hadronic interaction templates, where two points are plotted for each pair of vertices (one for each vertex). These results may be seen in Figure 16.4. As seen in these figures, aside from very large values of R_{xy} where merging is highly unlikely, no strong dependence is visible.




Figure 16.2: The templates for (2+2), (2+3), and (3+3) track pairs

Based on these results, the templates are created as a function of several variables. In order to avoid problems with statistics, results are currently binned inclusively in R_{xy} , and in two bins of ΔR , based on whether either vertex is within $\Delta R < 0.6$ of a jet or not. This is likely to be further optimized for the full background estimate, but is sufficient for basic tests of the method.

16.2 TOWARDS A BACKGROUND ESTIMATION

Once the templates are created, it is possible to test the closure of this method. The basic idea is as follows.

1. Obtain a template of the ratio of (2, 2)-track vertices in same-event vs. mixed events.
2. Scale the number of pairs of vertices in mixed events with (n, m)-tracks to be the same as for same-event pairs in the region $S > 100$.
3. Multiply the distribution for (n, m)-tracks by the template obtained in step 1. to obtain the predicted number of same-event pairs.
4. Subtract this distribution from the number of (scaled) mixed-event pairs to obtain the predicted number of merged events.

The predicted number of same-event (n, m)-track pairs can be compared to the actual number of same-event pairs to test the closure of the method. For these studies, in order to retain high statistics, only the DRAW filter is applied to events, but no selection specific to DV+jets is made.

An example of this may be seen in Figure 16.5, which shows the number of 2+2 pairs observed in one R_{xy} , $\min(\Delta R)$ bin, the number of mixed-event pairs, scaled to the same-event in the high distance significance region,




Figure 16.3: The ratio of distance significance for same-event to mixed-event vertices as a function of the minimum angular distance to a jet.

and the difference between them, which is the expected number of merged vertices. For the full background estimation, the 2+2 template of the ratios of same-event to mixed-event pairs will be applied to the 2+3 and 3+3 pairs. For these, a simple closure test may be performed in order to confirm that the number of events predicted by the template is the same as the difference between mixed-event and same-event pairs. This may be seen for 2+3 pairs in Figure 16.6, which demonstrates that that using the 2+2 template provides fairly consistent results to the observed difference for 2+3 pairs. While these results demonstrate that it is possible to predict shape differences between same-event and mixed-event pairs of vertices, this does not show whether this difference can be fully ex-




Figure 16.4: The ratio of distance significance for same-event to mixed-event vertices as a function of $R_{xy,DV}$.

plained by vertices that are merged. This question remains to be studied in the future but will not be covered here.




Figure 16.5: An example of the number of predicted vertices for 2+2-track events




Figure 16.6: An example of the number of predicted vertices for 2+3-track events

17

Additional Studies on a Background Estimation

While historically, each of the three backgrounds have been estimated separately, there is reason to believe that this is not the optimal approach. If instead, all of the backgrounds could be estimated simultaneously, this would significantly ease the background estimation strategy. In order to do this, it is first relevant to consider why displaced vertices are reconstructed, and which features they are correlated to. This chapter documents some of the first studies towards creating such a background estimation, though as seen here, there are still a number of unresolved issues that will need to be understood before an full application of this strategy.

One potentially interesting way of estimating the background is by parameterizing the likelihood of finding a displaced vertex, given various event- or jet-level quantities. In particular, since displaced vertex production is cor-

related with jet production, this section will study the possibility of parameterizing the probability of producing a DV as a function of various jet properties, such as the p_T , η , and other relevant features. The studies shown in this chapter are done with dijet Monte Carlo events, which allow some correlation between truth and reconstructed vertices and tracks.

To test the assumption that DV production is correlated to jet production, we can plot the minimum radial distance ΔR from each DV to any jet with p_T above 15 GeV in the event, as shown in Figure 17.1, which shows this for jets constructed from EM-scale topoclusters (EMTopo), and jets constructed from tracks. For these preliminary studies, no calibration is applied to the jets, as this is simply attempting to understand the basic correlation between displaced vertices and jets. As seen, most vertices fall within a 0.4 cone of the jet, but there is a long tail of DVs that appear to be uncorrelated to any jet, which is particularly seen when the matching is done with EMTopo jets. To understand if this is simply a factor of the jet p_T threshold, this same distribution is plotted as a function of the jet p_T threshold, as seen in Figure 17.2. As shown, this distribution is fairly stable for thresholds above 10 GeV. While lowering the EMTopo jet threshold is potentially interesting, current jet calibrations do not include recommendations for jets below 20 GeV due to difficulties in calibration, and track jets are a more likely route. However, for completeness, the studies shown in this chapter will be shown with both cluster and track jets, with an uncalibrated p_T threshold of 15 GeV for both.

In order to understand the dependence on the features of DV, this same distribution is shown as a function of the mass of the DV, the number of tracks in the DV, and the R_{xy} of the vertex in Figures 17.3, 17.4, and 17.5. The shape of the distribution does not appear to be strongly dependent on the vertex mass or N_{track} , and overall, there is not a strongly dependence on R_{xy} . However, at small values of R_{xy} , the ΔR distribution is more uniform, which is likely due to the presence of a large number of fakes.

These figures demonstrate that displaced vertex production does seem to be correlated with jet activity. For using this information to estimate the background, it is important to use the properties of jets to predict the number of expected displaced vertices. This means that it is necessary to study the probability of producing a DV given a jet with certain properties. The most basic parameterization is to look at the average number of DVs produced in




Figure 17.1: The minimum radial distance to a jet in PYTHIA di-jet events for vertices passing the pre-selection (without a material veto). This is shown for anti- k_t $R=0.4$ jets created from (a) EM-scale topoclusters, and (b) tracks.




Figure 17.2: The minimum radial distance to a jet in PYTHIA di-jet events for vertices passing the pre-selection (without a material veto) as a function of the jet p_T threshold that is applied. This is shown for anti- k_t $R=0.4$ jets created from (a) EM-scale topoclusters, and (b) tracks.

an $\Delta R = 0.4$ cone of a jet as a function of the jet p_T and η , which is shown in Figure 17.6. In all cases, it seems that the likelihood of finding a DV within a jet is very low. However, there does seem to be a slight correlation with the p_T of the jet and the probability of it containing a displaced vertex, where higher p_T jets are more likely to contain a DV.




Figure 17.3: The minimum radial distance to a jet in PYTHIA di-jet events for vertices passing the pre-selection (without a material veto) as a function of the mass of the DV. This is shown for anti- k_t $R=0.4$ jets created from (a) EM-scale topoclusters, and (b) tracks.




Figure 17.4: The minimum radial distance to a jet in PYTHIA di-jet events for vertices passing the pre-selection (without a material veto) as a function of the number of tracks in the DV. This is shown for anti- k_t $R=0.4$ jets created from (a) EM-scale topoclusters, and (b) tracks.

While this is by no means a complete set of studies sufficient for a background estimate, this does demonstrate that parameterizing the probability of creating a displaced vertex in terms of various track jet features could be a viable way to rethink the current background estimation techniques used for this type of analysis. Additional studies still need to be done to understand features such as the vertices that appear to be outside of the jet cone,




Figure 17.5: The minimum radial distance to a jet in PYTHIA di-jet events for vertices passing the pre-selection (without a material veto) as a function of the R_{xy} of the DV. This is shown for anti- k_t $R=0.4$ jets created from (a) EM-scale topoclusters, and (b) tracks.




Figure 17.6: The average number of DVs per jet as a function of the jet p_T and η .

how to correctly treat the probabilities for jets that are closeby to each other, and the full set of parameters that are necessary to characterize the track jets. However, given the possible gains in expanding the signal region, this is a promising direction that should be followed.

18

Summary of the Search for Displaced Vertices Plus Jets

The search for displaced vertices in association with jets is important to constrain models with long-lived particles that would not be constrained otherwise in searches for prompt jets. This search has not been done since Run 1, and so there is significant room for constraining parameters and models that are important to explore. Such a search is not without its challenges, including specialized object reconstruction, non-standard backgrounds, and even data processing considerations. The studies documented here demonstrate the potential reach of such a search, and additionally form the basis for much of the background estimation that will be critical to this search, and will hopefully provide a basis for this search to proceed smoothly in order to be able to explore a wide array of models.



Full Plots for Single Period Background Estimate

The first set of plots were done for material-dense regions, where there are no significant concerns of the results being sensitive to new signals. Similarly, it is possible to look outside of these regions, as long as only a small portion of the data is considered. With this in mind, one way of validating these results is to use data from a single run period to test if the results. For these studies, data is taken from Period D of 2015.




Figure A.1: The vertex mass distribution in data, and the template fit, using the full dataset




Figure A.2: The vertex mass distribution in data, and the template fit, using the full dataset




Figure A.3: The vertex mass distribution in data, and the template fit, using the full dataset




Figure A.4: The vertex mass distribution in data, and the template fit, using the full dataset




Figure A.5: The vertex mass distribution in data, and the template fit, using the full dataset




Figure A.6: The vertex mass distribution in data, and the template fit, using the full dataset




Figure A.7: The vertex mass distribution in data, and the template fit, using the full dataset




Figure A.8: The vertex mass distribution in data, and the template fit, using the full dataset




Figure A.9: The vertex mass distribution in data, and the template fit, using the full dataset




Figure A.10: The vertex mass distribution in data, and the template fit, using the full dataset

B

Corrections for Truth Jets associated to long-lived particles

In Monte Carlo simulations, truth particles are stored with a 4-momentum which describes their momentum in a local coordinate system. This means that a truth particle whose origin is translated in space will have the same 4-momentum as before translation, even though this is not what would be observed globally. For most Monte Carlo simulations, this is a reasonable way of modeling the behavior of particles. However, for decay products of long-lived particles, this can result in difficulties when trying to match truth particles to reconstructed objects, which is often important for studies of detector response, resolution, and efficiency. Typically, matching between truth

particles and reconstructed objects is done using a ΔR matching, but as seen in Figure B.1, this may not work when the decay product is not produced collinearly to the long-lived particle. Instead, a correction is necessary in order to be able to study the behavior of such jets.




Figure B.1: A diagram demonstrating the difference between a local and global coordinate system for a decay product of a long-lived particle.

It is possible to calculate a simple correction which may be applied to truth jets which are matched to the decay products of long-lived particles. The simplest correction ignores the effects of a magnetic field, and simply extrapolates particles linearly. First, the distance traveled by the decay product between its creation and the calorimeter can be calculated as

$$d_{\text{traveled}} = \frac{-b + \sqrt{b^2 - 4 \times a \times c}}{2 \times a}, \quad (\text{B.1})$$

where

$$\begin{aligned} a &= \sin(\theta_{\text{decay}})^2 \\ b &= 2 \times r_{\text{decay}} \times \sin(\theta_{\text{LLP}}) \times \sin(\theta_{\text{decay}}) \times (\cos(\phi_{\text{LLP}}) \times \cos(\phi_{\text{decay}}) + \sin(\phi_{\text{LLP}}) \times \sin(\phi_{\text{decay}})) \\ c &= (r_{\text{decay}} \times \theta_{\text{LLP}})^2 - R_{\text{calorimeter}}^2, \end{aligned} \quad (\text{B.2})$$

where θ_{LLP} and ϕ_{LLP} are the θ and ϕ direction of the long-lived particle respectively, and θ_{decay} and ϕ_{decay} are the θ and ϕ direction of the decay product, r_{decay} is the distance traveled by the long-lived particle before its decay, and $R_{\text{calorimeter}}$ is the longitudinal distance between the beampipe and the calorimeter

Then, the distance and direction of the decay product, as well as the position of the calorimeter, this can be

used to determine the 3-dimensional position where the decay product hits the calorimeter.

$$\begin{aligned}
x_{\text{cal}} &= r_{\text{decay}} \times \sin(\theta_{\text{LLP}}) \times \cos(\phi_{\text{LLP}}) + d_{\text{traveled}} \times \sin(\theta_{\text{decay}}) \times \cos(\phi_{\text{decay}}), \\
y_{\text{cal}} &= r_{\text{decay}} \times \sin(\theta_{\text{LLP}}) \times \sin(\phi_{\text{LLP}}) + d_{\text{traveled}} \times \sin(\theta_{\text{decay}}) \times \sin(\phi_{\text{decay}}), \\
z_{\text{cal}} &= r_{\text{decay}} \times \cos(\theta_{\text{LLP}}) + d_{\text{traveled}} \times \cos(\theta_{\text{decay}}),
\end{aligned} \tag{B.3}$$

Finally, all of this can be put together to obtain the corrected direction for the decay product,

$$\begin{aligned}
\theta_{\text{new}} &= \text{acos}\left(\frac{z_{\text{cal}}}{\sqrt{x_{\text{cal}}^2 + y_{\text{cal}}^2 + z_{\text{cal}}^2}}\right), \\
\phi_{\text{new}} &= \text{atan2}(y_{\text{cal}}, x_{\text{cal}}).
\end{aligned} \tag{B.4}$$

Ideally, an extrapolation could be performed for all truth particles, additionally including effects such as their bending from the magnetic field. However, the implementation of this is unnecessary for the most basic studies. Instead, the decay products of a long-lived particles are matched to truth jets, and the correction is applied to these truth jets. This means that these jets may have particles clustered in from prompt decays, but the effects of this are typically small. The effect of this correction may be seen in Figure B.2, which shows the minimum ΔR between a truth jet and reconstructed jet, where the truth jet has been matched to one of the decay products of the neutralino using a ΔR matching. For small lifetimes, the effect of the correction is observable, but most truth jets are still matched to a reconstructed jet. As the lifetime increases, the truth and reconstructed jet direction become less correlated, as seen by the flattened ΔR distribution, and the correction becomes essential for studying the performance of jets produced from decay products of long-lived particles.

An example of this transformation may be seen in Figures B.3- B.4, which shows the reconstructed jets, prompt truth jets, and the correction applied to the truth jets matched to decay products of long-lived particles. As can be seen, the corrections can be insignificant, while in other cases, it enables the truth jet to be matched to a reconstructed jet.




Figure B.2: The ΔR between truth and reconstructed jets before and after the long-lived particle correction is applied.




Figure B.3: An event display showing the correction applied to truth jets matched to decay products of long-lived particles.




Figure B.4: An event display showing the correction applied to truth jets matched to decay products of long-lived particles.



DigiTruth

As ATLAS goes to higher instantaneous luminosity, pileup is playing an increasingly significant role in jet reconstruction. In recent years, many algorithms have been designed to mitigate these effects ^{49 40 41}. These algorithms have been studied in-depth at the particle level, where information exists about which particles are hard-scatter and which are from pileup. For studies on ATLAS, these particles must be run through the GEANT4 simulation of the detector, in order to simulate the full reconstruction process. Since multiple particles may deposit energy in the same cell, the information is no longer as simple as at particle-level, making it harder to test and understand the effects of pileup.

Currently, we have few and inefficient methods to try to understand these effects. While we can produce zero-

pileup events, these only allow us to understand the effects on average, and not for individual events. Calibration hits are another option, but these require a lot of memory space due to the large number of hits, and also only give the true energy depositions, not what would have been reconstructed after running through the full simulation. Finally, the general method of looking at NPV and μ dependence only tells how things scale with NPV, not whether we are actually correcting the right effects.

Having a handle on the per-cluster or per-cell contributions from pileup would greatly improve our understanding of the effects of pileup. Not only could we better understand why certain algorithms perform better, but we could also design new algorithms. With this in mind, we have designed the DigiTruth method. This tool gives information on the per-cell level about how much of the reconstructed energy was contributed by HS. This information may be propagated to the cluster-level to give a more concise picture, or studied in-depth at the cell-level.

C.1 SIGNAL MODELING

C.1.1 SIGNAL MODELING IN THE LIQUID ARGON CALORIMETER

The LAr calorimeter signal processing is modeled in the GEANT4-based ATLAS detector simulation [8] following the actual online and offline reconstruction procedures as closely as possible. The output of the LAr simulation are *hits* which collect the energy deposits of particles in the active liquid argon (visible energy) of calorimeter readout cells, together with timing information. This is necessary to correctly convert these simulated deposits into the shaped LAr signal⁹⁸. Like in the readout of the experiment, the simulated signal is converted into a digital signal in four 25 ns wide time bins (*digitization*). The contributions from the actual electronics, like electronic noise and cross talk, as well as in-time and out-of-time pile-up are included in this step. Following this, and mimicking the digital filtering applied in the readout drivers (RODs) used in the experiment, energy and time are reconstructed from the digital signal and the quality of this reconstruction is measured and stored for each calorimeter cell.

CONVERTING HITS TO DIGITAL SIGNALS

During the detector simulation, the hits in each calorimeter cell are summed in bins of the time t_{hit} . To determine t_{hit} , the time-of-flight $t_{\text{flight}} = c|\vec{R}|$ of the particle travelling the distance $|\vec{R}|$ from the nominal collision vertex to the location of the energy deposit at speed of light c , is subtracted from the time t_{dep} of the energy deposit provided by GEANT4, $t_{\text{hit}} = t_{\text{dep}} - t_{\text{flight}}$. With this approach, the timing between the readout clock in the experiment and the actual time for the shower development is represented appropriately by t_{hit} . Variable binning is used for this energy collection, with narrow bins of order of 1 ns for small t_{hit} and wider bins of order 10 ns for larger t_{hit} . For each hit its deposited energy is added to the summed energy in the appropriate bin of t_{hit} . At the end of the simulation of one event, this procedure results in a list of energies $E_{\text{hit},j}$ summed up in discrete time intervals j of variable widths in each calorimeter cell.

The digital signal in the experiment consist of four digital readings $S_i = S(t_{\text{digi},i})$ defined in equidistant bins i of time. These bins are centered at $t_{\text{digi},i}$ and their width is $\Delta t_{\text{digi}} = 25\text{ns}$. The collected hit energy is projected into these time slots and converted into the S_i in units of ADC counts, according to the following calculation performed in each time sample i :

$$S_i = S_{\text{ped}} + S_{\text{noise}} + \frac{1}{f_{\text{sample}}} \times \frac{1}{c_{\text{ADC} \rightarrow \text{MeV}}} \times \quad (\text{C.1})$$

C.1.2 DIGITIZATION OF THE TILE CALORIMETER

As with the LAr simulation, the inputs to the Tile simulation⁹⁹ are hits from the GEANT4 simulation. For each cell, all hits within 1ns are merged, giving the time and energy of the merged deposits. Then, within a time window of ± 62.5 ns, the total energy is converted into the number of photoelectrons which are produced. This is calculated by inverting the sampling fraction to give a scaling factor from energy to number of photoelectrons.

This number is then used to convert to the number of photoelectrons actually measured, which is done with a random number from a Poisson distribution. All energies are rescaled by the ratio of measured photo-electrons to

produced photo-electrons, using the formula

$$E_{hit,rescaled} = \frac{E_{hit,orig}}{frac_{saml}} \times \frac{N_{pe,meas}}{N_{pe,prod}}. \quad (C.2)$$

To convert these energies into an actual signal, first they are multiplied by the sampling fraction, which converts an energy into an ADC. Then, 7 samples are determined from the pulse shape and the gain. If the high gain samples saturate the output, then the samples are determined again from the low gain.

Each of the (rescaled, merged within 1 ns) hits is converted into a pulse, which peaks at the time of the hit, and have a height of the rescaled energy times the MeV→ADC conversion factor. These pulses are summed together to create the actual signal seen. From this, 7 samples are taken, from -75ns to 75 ns. Noise is added to each sample using two Gaussians. The pedestal is added to each sample based on the channel number and the gain. Once these have been added, they are saved as a TileDigit.

Once we have these samples, reconstruction is run in the same way as for data. The optimal filtering coefficients are applied to obtain the measured amplitude, time, and pedestal. The two channels belonging to the same cell are merged to give the output of the tile cell.

C.1.3 DIGITRUTH IN THE TILE CALORIMETER

For the DigiTruth method, instead of merging all hits, only hits from the HS event are merged. Since the ratio $\frac{N_{pe,meas}}{N_{pe,prod}}$ is generated using a random number generator, extra care is taken to ensure that the same ratio is used for DigiTruth reconstruction as for standard reconstruction. Similarly, the noise is kept the same for both reconstruction processes.

C.2 DIGITRUTH

DigiTruth is a method to estimate the per-cell, and therefore the per-cluster contributions to the reconstructed energy from the hard-scatter event. Hits from the HS event are run through the digitization process without pileup hits, alongside with the standard reconstruction. This is slightly different than a zero-pileup simulation, since care

is taken to ensure that the same conditions for reconstruction are used. The gain in each cell for DigiTruth reconstruction is the same as for the standard reconstruction, and the optimal filtering coefficients are chosen based on pileup conditions. After finding the HS digits for each channel, DigiTruth digits are put through the same reconstruction process as the standard cells, resulting in a second cell collection with DigiTruth information.

While extensive studies using the DigiTruth method will not be shown here, we will show a few plots demonstrating its purpose and validating the method. One of the most basic applications of DigiTruth is to create a cluster-level ratio of HS energy to total reconstructed energy. This can be done by simply adding together the energies from each cell which is part of the cluster for both the DigiTruth cells and the standard cells.

From Figure C.1, we can see that there are two main populations of clusters, one with a ratio near 0 corresponding to pileup clusters, and one with a ratio near 1, corresponding to hard scatter. This shows that even at $\mu=40$, most clusters are either PU or HS, and not usually a noticeable combination. Second, it is also apparent that most pileup clusters have very low energies, as would be expected from soft radiation.

Clusters within $\Delta R < 0.2$ from a truth jet of $p_T > 20$ GeV, the cluster energy ratio tends to be closer to 1, while clusters which are not within a $\Delta R > 0.6$ cone of any truth jets tend to have an energy ratio of around 0, as seen in Figure C.2. This demonstrates the capabilities of this method in discriminating between hard-scatter and pileup clusters, and that this ratio is a meaningful measure of whether or not a cluster originates from hard-scatter.

Finally, this ratio is shown as a function of η in Figure C.3, which shows how this ratio changes with the detector. More statistics are needed in the forward region in order to draw any conclusions, but some detector features are evident. A full study of this is left for future work.




Figure C.1: Ratio of DigiTruth HS energy to total reconstructed energy in clusters (a) for all clusters, and (b) normalized in each slice of energy




Figure C.2: Ratio of DigiTruth HS energy to total reconstructed energy in clusters normalized in each slice of energy (a) for clusters associated to truth jets, (b) for clusters not associated to truth jets.




Figure C.3: Ratio of DigiTruth HS energy to total reconstructed energy as a function of η normalized in each slice of η

References

- [1] 413235 (1995). ALICE: Technical proposal for a large ion collider experiment at the CERN LHC.
- [2] Aaboud, M. et al. (2017a). Jet energy scale measurements and their systematic uncertainties in proton-proton collisions at $\sqrt{s} = 13$ TeV with the ATLAS detector. *Phys. Rev.*, D96(7), 072002.
- [3] Aaboud, M. et al. (2017b). Jet reconstruction and performance using particle flow with the ATLAS Detector. *Eur. Phys. J.*, C77(7), 466.
- [4] Aaboud, M. et al. (2018). Search for long-lived, massive particles in events with displaced vertices and missing transverse momentum in $\sqrt{s} = 13$ TeV pp collisions with the ATLAS detector. *Phys. Rev.*, D97(5), 052012.
- [5] Aad, G. et al. (2012). Observation of a new particle in the search for the Standard Model Higgs boson with the ATLAS detector at the LHC. *Phys. Lett.*, B716, 1–29.
- [6] Achenbach, R., Adragna, P., Andrei, V., Apostologlou, P., Åsman, B., Ay, C., Barnett, B. M., Bauss, B., Bendel, M., Böhm, C., Booth, J. R. A., Brawn, I. P., Thomas, P. B., Charlton, D. G., Collins, N. J., Curtis, C. J., Dahloff, A., Davis, A. O., Eckweiler, S., Edwards, J. P., Eisenhandler, E., Faulkner, P. J. W., Fleckner, J., Föhlich, F., Garvey, J., Gee, C. N. P., Gillman, A. R., Hanke, P., Hatley, R. P., Hellman, S., Hidvégi, A., Hillier, S. J., Jakobs, K., Johansen, M., Kluge, E. E., Landon, M., Lendermann, V., Lilley, J. N., Mahboubi, K., Mahout, G., Mass, A., Meier, K., Moa, T., Moyse, E., Müller, F., Neusiedl, A., Nöding, C., Oltmann, B., Pentney, J. M., Perera, V. J. O., Pfeiffer, U., Prieur, D. P. F., Qian, W., Rees, D. L., Rieke, S., Rühr, F., Sankey, D. P. C., Schäfer, U., Schmitt, K., Schultz-Coulon, H. C., Schumacher, C., Silverstein, S., Staley, R. J., Stamen, R., Stockton, M. C., Tapprogge, S., Thomas, J. P., Trefzger, T., Watkins, P. M., Watson, A., Weber, P., & Woehrling, E. E. (2008). The ATLAS Level-1 Calorimeter Trigger. *Journal of Instrumentation*, 3(03), P03001–P03001.
- [7] Adye, T. (2011). Unfolding algorithms and tests using RooUnfold.
- [8] Alves, Jr., A. A. et al. (2008). The LHCb Detector at the LHC. *JINST*, 3, S08005.
- [9] Alwall, J., Frederix, R., Frixione, S., Hirschi, V., Maltoni, F., Mattelaer, O., Shao, H. S., Stelzer, T., Torrielli, P., & Zaro, M. (2014). The automated computation of tree-level and next-to-leading order differential cross sections, and their matching to parton shower simulations. *JHEP*, 07, 079.
- [10] Andersen, J. R. et al., Eds. (2018). *Les Houches 2017: Physics at TeV Colliders Standard Model Working Group Report*.
- [11] Andersson, B., Gustafson, G., Ingelman, G., & Sjöstrand, T. (1983). Parton fragmentation and string dynamics. *Physics Reports*, 97(2), 31 – 145.
- [12] Andronic, A. & Wessels, J. (2012). Transition radiation detectors. *Nuclear Instruments and Methods in Physics Research Section A: Accelerators, Spectrometers, Detectors and Associated Equipment*, 666, 130 – 147. Advanced Instrumentation.

- [13] ATLAS (2016). Performance of pile-up mitigation techniques for jets in pp collisions with the ATLAS detector. *Nucl. Instrum. Meth.*, A824, 367–370.
- [14] ATLAS Collaboration (2008). The atlas experiment at the cern large hadron collider. 3(08), S08003–S08003.
- [15] ATLAS Collaboration (2010a). *Performance of primary vertex reconstruction in proton-proton collisions at $\sqrt{s}=7$ TeV in the ATLAS experiment*. Technical Report ATLAS-CONF-2010-069, CERN, Geneva.
- [16] ATLAS Collaboration (2010b). *Performance of the ATLAS Silicon Pattern Recognition Algorithm in Data and Simulation at $\sqrt{s}=7$ TeV*. Technical Report ATLAS-CONF-2010-072, CERN, Geneva.
- [17] ATLAS Collaboration (2010c). The ATLAS Simulation Infrastructure. *Eur. Phys. J. C*, 70, 823–874.
- [18] ATLAS Collaboration (2011). Search for massive long-lived highly ionising particles with the ATLAS detector at the LHC. *Phys. Lett. B*, 698, 353.
- [19] ATLAS Collaboration (2012). Summary of ATLAS Pythia 8 tunes. *ATL-PHYS-PUB-2012-003*.
- [20] ATLAS Collaboration (2013a). Performance of jet substructure techniques for large- R jets in proton-proton collisions at $\sqrt{s}=7$ TeV using the ATLAS detector. *JHEP*, 09, 076.
- [21] ATLAS Collaboration (2013b). Search for charginos nearly mass degenerate with the lightest neutralino based on a disappearing-track signature in pp collisions at $\sqrt{s}=8$ TeV with the ATLAS detector. *Phys. Rev. D*, 88, 112006.
- [22] ATLAS Collaboration (2013c). Search for nonpointing photons in the diphoton and E_T^{miss} final state in $\sqrt{s}=7$ TeV proton–proton collisions using the ATLAS detector. *Phys. Rev. D*, 88, 012001.
- [23] ATLAS Collaboration (2014a). *A measurement of single hadron response using data at $\sqrt{s}=8$ TeV with the ATLAS detector*. Technical Report ATL-PHYS-PUB-2014-002, CERN, Geneva.
- [24] ATLAS Collaboration (2014b). ATLAS Pythia8 tunes to 7 TeV data. *ATL-PHYS-PUB-2014-021*.
- [25] ATLAS Collaboration (2014c). Tagging and suppression of pileup jets with the ATLAS detector. ATLAS-CONF-2014-018.
- [26] ATLAS Collaboration (2015a). Jet energy measurement and its systematic uncertainty in proton-proton collisions at $\sqrt{s}=7$ TeV with the ATLAS detector. *Eur. Phys. J. C*, 75, 17.
- [27] ATLAS Collaboration (2015b). Search for long-lived, weakly interacting particles that decay to displaced hadronic jets in proton–proton collisions at $\sqrt{s}=8$ TeV with the ATLAS detector. *Phys. Rev. D*, 92, 012010.
- [28] ATLAS Collaboration (2015c). Search for massive, long-lived particles using multitrack displaced vertices or displaced lepton pairs in pp collisions at $\sqrt{s}=8$ TeV with the ATLAS detector. *Phys. Rev. D*, 92, 072004.
- [29] ATLAS Collaboration (2015d). *The Optimization of ATLAS Track Reconstruction in Dense Environments*. Technical Report ATL-PHYS-PUB-2015-006, CERN, Geneva.
- [30] ATLAS Collaboration (2015e). *Vertex Reconstruction Performance of the ATLAS Detector at $\sqrt{s}=13$ TeV*. Technical Report ATL-PHYS-PUB-2015-026, CERN, Geneva.
- [31] ATLAS Collaboration (2016). Performance of pile-up mitigation techniques for jets in pp collisions at $\sqrt{s}=8$ TeV using the ATLAS detector. *Eur. Phys. J. C*, 76, 581.
- [32] ATLAS Collaboration (2017a). A measurement of the calorimeter response to single hadrons and determination of the jet energy scale uncertainty using LHC Run-1 pp -collision data with the ATLAS detector. *Eur. Phys. J. C*, 77(1), 26.

- [33] ATLAS Collaboration (2017b). *Improving jet substructure performance in ATLAS using Track-CaloClusters*. Technical Report ATL-PHYS-PUB-2017-015, CERN, Geneva.
- [34] ATLAS Collaboration (2017c). Jet energy scale measurements and their systematic uncertainties in proton-proton collisions at $\sqrt{s} = 13$ TeV with the ATLAS detector. *Phys. Rev. D*, 96(7), 072002.
- [35] ATLAS Collaboration (2017d). *Performance of the reconstruction of large impact parameter tracks in the ATLAS inner detector*. Technical Report ATL-PHYS-PUB-2017-014, CERN, Geneva.
- [36] Bahr, M. et al. (2008). Herwig++ physics and manual. *Eur. Phys. J. C*, 58, 639–707.
- [37] Ball, R. D. et al. (2015). Parton distributions for the LHC Run II. *JHEP*, 04, 040.
- [38] Beenakker, W., Borschensky, C., Krämer, M., Kulesza, A., & Laenen, E. (2016). NNLL-fast: predictions for coloured supersymmetric particle production at the LHC with threshold and Coulomb resummation. *JHEP*, 12, 133.
- [39] Benedikt, M., Collier, P., Mertens, V., Poole, J., & Schindl, K. (2004). *LHC Design Report*. CERN Yellow Reports: Monographs. Geneva: CERN.
- [40] Berta, P., Spousta, M., Miller, D. W., & Leitner, R. (2014). Particle-level pileup subtraction for jets and jet shapes. *JHEP*, 06, 092.
- [41] Bertolini, D., Harris, P., Low, M., & Tran, N. (2014). Pileup Per Particle Identification. *JHEP*, 10, 059.
- [42] Bhattacharyya, G. (1997). A Brief review of R-parity violating couplings. In *Beyond the desert 1997: Accelerator and non-accelerator approaches. Proceedings, 1st International Conference on Particle Physics beyond the Standard Model, Tegernsee, Ringberg Castle, Germany, June 8-14, 1997* (pp. 194–201).
- [43] Brüning, O. S., Collier, P., Lebrun, P., Myers, S., Ostojic, R., Poole, J., & Proudlock, P. (2004a). *LHC Design Report*. CERN Yellow Reports: Monographs. Geneva: CERN.
- [44] Brüning, O. S., Collier, P., Lebrun, P., Myers, S., Ostojic, R., Poole, J., & Proudlock, P. (2004b). *LHC Design Report*. CERN Yellow Reports: Monographs. Geneva: CERN.
- [45] Cacciari, M. & Salam, G. P. (2006). Dispelling the N^3 myth for the k_t jet-finder. *Phys. Lett.*, B641, 57–61.
- [46] Cacciari, M. & Salam, G. P. (2008). Pileup subtraction using jet areas. *Phys. Lett.*, B659, 119–126.
- [47] Cacciari, M., Salam, G. P., & Soyez, G. (2008). The anti- k_t jet clustering algorithm. *JHEP*, 04, 063.
- [48] Cacciari, M., Salam, G. P., & Soyez, G. (2012). FastJet user manual. *Eur. Phys. J. C*, 72, 1896.
- [49] Cacciari, M., Salam, G. P., & Soyez, G. (2015). SoftKiller, a particle-level pileup removal method. *Eur. Phys. J.*, C75(2), 59.
- [50] Capeans, M., Darbo, G., Einsweiler, K., Elsing, M., Flick, T., Garcia-Sciveres, M., Gemme, C., Pernegger, H., Rohne, O., & Vuillermet, R. (2010). ATLAS Insertable B-Layer Technical Design Report.
- [51] Carrazza, S., Forte, S., & Rojo, J. (2013). Parton Distributions and Event Generators.
- [52] Catani, S. et al. (2001). Qcd matrix elements + parton showers. *JHEP*, 11, 063.
- [53] Catani, S. & Seymour, M. H. (1997). A General algorithm for calculating jet cross-sections in NLO QCD. *Nucl. Phys.*, B485, 291–419. [Erratum: Nucl. Phys.B510,503(1998)].
- [54] Chatrchyan, S. et al. (2012). Observation of a new boson at a mass of 125 GeV with the CMS experiment at the LHC. *Phys. Lett.*, B716, 30–61.

- [55] CMS Collaboration (2008). The CMS experiment at the CERN LHC. *JINST*, 3, S08004.
- [56] collaboration, T. A. (2014). Operation and performance of the ATLAS semiconductor tracker. *Journal of Instrumentation*, 9(08), P08009–P08009.
- [57] Corcella, G., Knowles, I. G., Marchesini, G., Moretti, S., Odagiri, K., Richardson, P., Seymour, M. H., & Webber, B. R. (2001). Herwig 6: An event generator for hadron emission reactions with interfering gluons (including supersymmetric processes). *JHEP*, 01, 010.
- [58] Corke, R. & Sjöstrand, T. (2010). Improved parton showers at large transverse momenta. *Eur. Phys. J. C*, 69, 1–18.
- [59] Dasgupta, M. & Salam, G. P. (2001). Resummation of nonglobal QCD observables. *Phys. Lett. B*, 512, 323–330.
- [60] de Boer, Y. R., Kaidalov, A. B., Milstead, D. A., & Piskounova, O. I. (2008). Interactions of Heavy Hadrons using Regge Phenomenology and the Quark Gluon String Model. *J. Phys.*, G35, 075009.
- [61] Deur, A., Brodsky, S. J., & de Teramond, G. F. (2016). The QCD Running Coupling. *Prog. Part. Nucl. Phys.*, 90, 1–74.
- [62] Dokshitzer, Y. L. et al. (1997). Better jet clustering algorithms. *JHEP*, 08, 001.
- [63] Dulat, S., Hou, T.-J., Gao, J., Guzzi, M., Huston, J., Nadolsky, P., Pumplin, J., Schmidt, C., Stump, D., & Yuan, C. P. (2016). New parton distribution functions from a global analysis of quantum chromodynamics. *Phys. Rev.*, D93(3), 033006.
- [64] Ellis, S. D. & Soper, D. E. (1993). Successive combination jet algorithm for hadron collisions. *Phys. Rev. D*, 48, 3160–3166.
- [65] Ellis, S. D., Vermilion, C. K., & Walsh, J. R. (2010). Recombination Algorithms and Jet Substructure: Pruning as a Tool for Heavy Particle Searches. *Phys. Rev.*, D81, 094023.
- [66] et al, G. A. (2008). ATLAS pixel detector electronics and sensors. *Journal of Instrumentation*, 3(07), P07007–P07007.
- [67] Farrar, G. R., Mackeprang, R., Milstead, D., & Roberts, J. P. (2011). Limit on the mass of a long-lived or stable gluino. *JHEP*, 02, 018.
- [68] For a recent review, see, Buckley, A., et al. (2011). General-purpose event generators for LHC physics. *Phys. Rept.*, 504, 145–233.
- [69] Frye, C. et al. (2016a). Factorization for groomed jet substructure beyond the next-to-leading logarithm. *JHEP*, 07, 064.
- [70] Frye, C. et al. (2016b). Precision physics with pile-up insensitive observables.
- [71] GEANT4 Collaboration (2003). GEANT4: a simulation toolkit. *Nucl. Instrum. Meth. A*, 506, 250–303.
- [72] Gieseke, S., Rohr, C., & Siodmok, A. (2012). Colour reconnections in herwig++. *Eur. Phys. J. C*, 72, 2225.
- [73] Gingrich, D. M. (2007). Construction, assembly and testing of the ATLAS hadronic end-cap calorimeter. *Journal of Instrumentation*, 2(05), P05005–P05005.
- [74] Giudice, G. F. & Rattazzi, R. (1999). Theories with gauge mediated supersymmetry breaking. *Phys. Rept.*, 322, 419–499.

- [75] Gleisberg, T., Hoeche, S., Krauss, F., Schonherr, M., Schumann, S., Siegert, F., & Winter, J. (2009). Event generation with sherpa 1.1. *JHEP*, 02, 007.
- [76] Hewett, J. L., Lillie, B., Masip, M., & Rizzo, T. G. (2004). Signatures of long-lived gluinos in split supersymmetry. *JHEP*, 09, 070.
- [77] Hoang, A. H., Mantry, S., Pathak, A., & Stewart, I. W. (2017). Extracting a Short Distance Top Mass with Light Grooming.
- [78] Huth, J. E. et al. (1990). Toward a standardization of jet definitions. In *1990 DPF Summer Study on High-energy Physics: Research Directions for the Decade (Snowmass 90) Snowmass, Colorado, June 25-July 13, 1990* (pp. 0134–136).
- [79] Kolda, C. F. (1998). Gauge mediated supersymmetry breaking: Introduction, review and update. *Nucl. Phys. Proc. Suppl.*, 62, 266–275. [266(1997)].
- [80] Krohn, D., Thaler, J., & Wang, L.-T. (2010). Jet Trimming. *JHEP*, 02, 084.
- [81] Lai, H.-L., Guzzi, M., Huston, J., Li, Z., Nadolsky, P. M., Pumplin, J., & Yuan, C. P. (2010). New parton distributions for collider physics. *Phys. Rev. D*, 82, 074024.
- [82] Langacker, P. (2010). Introduction to the Standard Model and Electroweak Physics. In *Proceedings of Theoretical Advanced Study Institute in Elementary Particle Physics on The dawn of the LHC era (TASI 2008): Boulder, USA, June 2-27, 2008* (pp. 3–48).
- [83] Larkoski, A. J., Marzani, S., Soyez, G., & Thaler, J. (2014a). Soft Drop. *JHEP*, 05, 146.
- [84] Larkoski, A. J., Moult, I., & Neill, D. (2014b). Power Counting to Better Jet Observables. *JHEP*, 12, 009.
- [85] Larkoski, A. J., Moult, I., & Neill, D. (2016). Analytic Boosted Boson Discrimination. *JHEP*, 05, 117.
- [86] Larkoski, A. J., Salam, G. P., & Thaler, J. (2013). Energy correlation functions for jet substructure. *JHEP*, 06, 108.
- [87] M. Aharrouché et al, A. E. B. C. C. (2006). Energy linearity and resolution of the atlas electromagnetic barrel calorimeter in an electron test-beam. *Nucl. Instrum. Meth. A*, 568.
- [88] Mackeprang, R. & Hansen, J. R. (2007). Stable Heavy Hadrons in ATLAS. Presented 05 Oct 2007.
- [89] Mackeprang, R. & Milstead, D. A. (2010). An updated description of heavy-hadron interactions in geant-4. *The European Physical Journal C*, 66(3), 493–501.
- [90] Malaescu, B. (2009). An Iterative, dynamically stabilized method of data unfolding.
- [91] Martin, A. D., Stirling, W. J., Thorne, R. S., & Watt, G. (2009). Parton distributions for the LHC. *Eur. Phys. J. C*, 63, 189.
- [92] Martin, S. P. (1997). A Supersymmetry primer. (pp. 1–98). [Adv. Ser. Direct. High Energy Phys.18,1(1998)].
- [93] Marzani, S., Schunk, L., & Soyez, G. (2017a). A study of jet mass distributions with grooming. *JHEP*, 07, 132.
- [94] Marzani, S., Schunk, L., & Soyez, G. (2017b). The jet mass distribution after Soft Drop.
- [95] Meade, P., Reece, M., & Shih, D. (2010). Long-Lived Neutralino NLSPs. *JHEP*, 10, 067.
- [96] Mindur, B. (2016). *ATLAS Transition Radiation Tracker (TRT): Straw tubes for tracking and particle identification at the Large Hadron Collider*. Technical Report ATL-INDET-PROC-2016-001, CERN, Geneva.

- [97] Nagy, Z. (2003). Next-to-leading order calculation of three jet observables in hadron hadron collision. *Phys. Rev.*, D68, 094002.
- [98] Oliveira Damazio, D. (2013). *Signal Processing for the ATLAS Liquid Argon Calorimeter : studies and implementation*. Technical Report ATL-LARG-PROC-2013-015, CERN, Geneva.
- [99] on behalf of the ATLAS Collaboration, S. K. (2014). Simulation and validation of the ATLAS Tile Calorimeter response.
- [100] Pumplin, J., Stump, D. R., Huston, J., Lai, H. L., Nadolsky, P. M., & Tung, W. K. (2002). New generation of parton distributions with uncertainties from global qcd analysis. *JHEP*, 07, 012.
- [101] Salam, G. P. (2010). Towards Jetography. *Eur. Phys. J.*, C67, 637–686.
- [102] Schindl, K. (1999). The Injector Chain for the LHC; rev. version. (CERN-PS-99-018-DI), 7 p.
- [103] Sjöstrand, T., Mrenna, S., & Skands, P. Z. (2008). A brief introduction to PYTHIA 8.1. *Comput. Phys. Commun.*, 178, 852–867.
- [104] Tanabashi, M., Hagiwara, K., Hikasa, K., Nakamura, K., Sumino, Y., Takahashi, F., Tanaka, J., Agashe, K., Aielli, G., Amsler, C., Antonelli, M., Asner, D. M., Baer, H., Banerjee, S., Barnett, R. M., Basaglia, T., Bauer, C. W., Beatty, J. J., Belousov, V. I., Beringer, J., Bethke, S., Bettini, A., Bichsel, H., Biebel, O., Black, K. M., Blucher, E., Buchmuller, O., Burkert, V., Bychkov, M. A., Cahn, R. N., Carena, M., Cacciucci, A., Cerri, A., Chakraborty, D., Chen, M.-C., Chivukula, R. S., Cowan, G., Dahl, O., D’Ambrosio, G., Damour, T., de Florian, D., de Gouvêa, A., DeGrand, T., de Jong, P., Dissertori, G., Dobrescu, B. A., D’Onofrio, M., Doser, M., Drees, M., Dreiner, H. K., Dwyer, D. A., Eerola, P., Eidelman, S., Ellis, J., Erler, J., Ezhela, V. V., Fetscher, W., Fields, B. D., Firestone, R., Foster, B., Freitas, A., Gallagher, H., Garren, L., Gerber, H.-J., Gerbier, G., Gershon, T., Gershtein, Y., Gherghetta, T., Godizov, A. A., Goodman, M., Grab, C., Gritsan, A. V., Grojean, C., Groom, D. E., Grünwald, M., Gurtu, A., Gutsche, T., Haber, H. E., Hanhart, C., Hashimoto, S., Hayato, Y., Hayes, K. G., Hebecker, A., Heinemeyer, S., Heltsley, B., Hernández-Rey, J. J., Hisano, J., Höcker, A., Holder, J., Holtkamp, A., Hyodo, T., Irwin, K. D., Johnson, K. F., Kado, M., Karliner, M., Katz, U. F., Klein, S. R., Klempt, E., Kowalewski, R. V., Krauss, F., Kreps, M., Krusche, B., Kuyanov, Y. V., Kwon, Y., Lahav, O., Laiho, J., Lesgourgues, J., Liddle, A., Ligeti, Z., Lin, C.-J., Lippmann, C., Liss, T. M., Littenberg, L., Lugovsky, K. S., Lugovsky, S. B., Lusiani, A., Makida, Y., Maltoni, F., Mannel, T., Manohar, A. V., Marciano, W. J., Martin, A. D., Masoni, A., Matthews, J., Meißner, U.-G., Milstead, D., Mitchell, R. E., Mönig, K., Molaro, P., Moortgat, F., Moskvic, M., Murayama, H., Narain, M., Nason, P., Navas, S., Neubert, M., Nevski, P., Nir, Y., Olive, K. A., Pagan Griso, S., Parsons, J., Patrignani, C., Peacock, J. A., Pennington, M., Petcov, S. T., Petrov, V. A., Pianori, E., Piepke, A., Pomarol, A., Quadt, A., Rademacker, J., Raffelt, G., Ratcliff, B. N., Richardson, P., Ringwald, A., Roesler, S., Rolli, S., Romaniouk, A., Rosenberg, L. J., Rosner, J. L., Rybka, G., Ryutin, R. A., Sachrajda, C. T., Sakai, Y., Salam, G. P., Sarkar, S., Sauli, F., Schneider, O., Scholberg, K., Schwartz, A. J., Scott, D., Sharma, V., Sharpe, S. R., Shutt, T., Silari, M., Sjöstrand, T., Skands, P., Skwarnicki, T., Smith, J. G., Smoot, G. F., Spanier, S., Spieler, H., Spiering, C., Stahl, A., Stone, S. L., Sumiyoshi, T., Syphers, M. J., Terashi, K., Terning, J., Thoma, U., Thorne, R. S., Tiator, L., Titov, M., Tkachenko, N. P., Törnqvist, N. A., Tovey, D. R., Valencia, G., Van de Water, R., Varelas, N., Venanzoni, G., Verde, L., Vinciter, M. G., Vogel, P., Vogt, A., Wakely, S. P., Walkowiak, W., Walter, C. W., Wands, D., Ward, D. R., Wascko, M. O., Weiglein, G., Weinberg, D. H., Weinberg, E. J., White, M., Wiencke, L. R., Willocq, S., Wohl, C. G., Womersley, J., Woody,

- C. L., Workman, R. L., Yao, W.-M., Zeller, G. P., Zenin, O. V., Zhu, R.-Y., Zhu, S.-L., Zimmermann, F., Zyla, P. A., Anderson, J., Fuller, L., Lugovsky, V. S., & Schaffner, P. (2018). Review of particle physics. *Phys. Rev. D*, 98, 030001.
- [105] Thaler, J. & Van Tilburg, K. (2011). Identifying Boosted Objects with N-subjettiness. *JHEP*, 03, 015.
- [106] Vogel, A. (2013). *ATLAS Transition Radiation Tracker (TRT): Straw Tube Gaseous Detectors at High Rates*. Technical Report ATL-INDET-PROC-2013-005, CERN, Geneva.
- [107] Webber, B. (1984). A qcd model for jet fragmentation including soft gluon interference. *Nuclear Physics B*, 238(3), 492 – 528.