
Search for long-lived Supersymmetric particles using displaced vertices with the ATLAS detector at the LHC

Emily Anne Thompson



Fakultät für Mathematik und Physik
Albert-Ludwigs-Universität Freiburg



Search for long-lived Supersymmetric particles using displaced vertices with the ATLAS detector at the LHC

Dissertation
zur
Erlangung des Doktorgrades (Dr. rer. nat.)
der
Fakultät für Mathematik und Physik
der
Albert-Ludwigs-Universität Freiburg

vorgelegt von
Emily Anne Thompson

April 2022

Dekan

Prof. Dr. Michael Thoss

Betreuer der Arbeit

Prof. Dr. Beate Heinemann, Dr. Federico Meloni

Datum der mündlichen Prüfung: 15.07.2022

Prüfer

Prof. Dr. Gregor Hertel

Prof. Dr. Beate Heinemann

JProf. Dr. Stefan Vogl

Acknowledgements

Working on this thesis as a student at DESY and in the ATLAS collaboration has been an immensely rewarding experience. Though I certainly enjoy searching for new physics - I also know that this experience wouldn't have been half as enjoyable if it weren't for the amazing people with whom I've had the pleasure to work with.

To Federico: Thank you for advising me these past few years. I really appreciate every time you've given me feedback on physics studies, sat through a practice presentation, reviewed one of my applications, read through this thesis, and answered my panicked messages even if they were on a Sunday. Thanks for believing in me and for all of your really great advice about everything. I hope one day to be as excellent of an advisor as you are.

To Beate: thank you for your invaluable advice these past few years. I learned something from each one of our conversations. Thank you for answering my cold email inquiring about a PhD position and for giving me a chance. Thank you for your professional advice and for being an excellent role model.

To the SUSY DV team: thanks for being the best colleagues anyone could ask for. Christian, Larry, and Simone - thank you for your sincere interest and invaluable feedback at every meeting and for spreading your infectious passion for this analysis. Giulia, David, Atsushi, Moe, Filip, Evan, Oleg, Alex, Stefan, Dave, Marija, Jennifer, and Rebecca - thanks for making this analysis fun.

To Marco, Sonia, Tadej, and the entire SUSY/LL team at DESY: Thanks for your thoughtful questions and insightful conversations every Wednesday morning. Thanks for always being available to answer my questions.

To Stefan, Jan, David, Max, Tova, Ben and all of my FTK colleagues: Thank you for answering my naive questions as a new PhD student and for helping me get through my authorship qualification task.

To my RPV-meets-RPC colleagues: Carlos - it was a pleasure to work with you and learn together how to "fake it til' we make it". Lorenzo - thank you for joining us and picking up the momentum in moving in this analysis forward.

To the DESY ATLAS group - thank you for creating an immensely supportive and welcoming environment where I could grow as a physicist. Thank you especially to the group leaders - Ingrid, Klaus, and Krisztian - and to the group secretaries - Sabine, Carolin, and Ramona - for getting the group through the pandemic in one piece and for taking special care of us international students.

To my mentor Shan and my fellow Dynament mentees: I learned so much from our interactions and

always left our conversations feeling more inspired and motivated to continue this journey. Thank you for that!

To my ECSB family - Ben, Flavia, Giordon, Sukanya, Sebastien, Michele, Dimitrii, Lesya, and Reina: Thank you for inspiring me with your passion and ambition to make the ATLAS collaboration an even better working environment for early career scientists. I hope our paths cross again and we can all meet in-person one day.

To my fellow PhD students at DESY, especially my officemates - Craig, Tom, Jordi, Fang-Ying, and Hamish: thank you for all of the great lunch and coffee breaks and wonderful memories. Per request, I would like to give a special thanks to Alessandro Guida.

To my family - mom, dad, and Denise: thank you for your unwavering support during these years. I couldn't have done this without you!

Abstract

The existence of long-lived particles (LLPs) is a common feature in many theories beyond the Standard Model. For example, models with small couplings (i.e. R-parity-violating supersymmetry) and models allowing for decays via highly virtual intermediate states (i.e. Split supersymmetry) predict the presence of LLPs. With lifetimes ranging from picoseconds to nanoseconds, massive LLPs could decay to several electrically charged particles in the inner tracking volume of the ATLAS detector at the Large Hadron Collider, resulting in the reconstruction of a displaced secondary vertex. This thesis presents the first search for high mass, multi-track displaced vertices and multiple energetic jets using the full Run-2 dataset collected by the ATLAS detector, corresponding to an integrated luminosity of 139 fb^{-1} at $\sqrt{s} = 13 \text{ TeV}$. The observed event yields are compatible with those expected from background processes. The results are presented as exclusion limits at 95% confidence level on supersymmetric scenarios with non-zero R-parity violating couplings.

This thesis also presents the ATLAS FastTracKer (FTK), which is a hardware-based track finder envisioned to be integrated into the ATLAS trigger system. The trigger system is responsible for selecting interesting events with a high purity at a rate of only 1 kHz. The FTK was designed to provide tracking information of all charged particles with $p_T > 1 \text{ GeV}$ in the ATLAS inner detector, which is a powerful discriminant for selecting interesting events. The training of the FTK system to achieve a high track-finding efficiency and the adaptability of the system to new detector conditions is presented. Furthermore, the applicability of the system to select candidate events containing long-lived particles is described.

Zusammenfassung

Die Existenz langlebiger Teilchen (LLPs) ist ein gemeinsames Merkmal vieler Theorien jenseits des Standardmodells. Zum Beispiel sagen Modelle mit kleinen Kopplungen (R-Paritätsverletzende Supersymmetrie) und Modelle, die Zerfälle über hochvirtuelle Zwischenzustände zulassen (d.h. geteilte Supersymmetrie), das Vorhandensein von LLPs voraus. Mit Lebensdauern im Bereich von Pikosekunden bis Nanosekunden könnten massive LLPs im inneren Spurdetektor des ATLAS-Detektors am Large Hadron Collider in mehrere elektrisch geladene Teilchen zerfallen, was zur Rekonstruktion eines verschobenen sekundären Vertex führt. In dieser Arbeit wird die erste Suche nach schweren, mehrspurigen verschobenen Vertices in Verbindung mit mehreren energetischen Jets vorgestellt, wobei der gesamte vom ATLAS-Detektor gesammelte Run-2-Datensatz verwendet wird, der einer integrierten Luminosität von 139 fb^{-1} bei $\sqrt{s} = 13 \text{ TeV}$ entspricht. Die beobachtete Anzahl von Ereignissen ist mit der durch Hintergrundprozesse erwarteten Anzahl kompatibel. Die Ergebnisse werden als Ausschlussgrenzen auf einem 95% Vertrauensniveau für supersymmetrische Szenarien mit R-Paritätsverletzungs-Kopplungen ungleich Null dargestellt.

In dieser Arbeit wird auch der ATLAS FastTracKer (FTK) vorgestellt, ein Hardware-basierter Spurensucher, der in das ATLAS-Trigger-System integriert werden soll. Das Triggersystem ist für die Auswahl interessanter Ereignisse mit einer hohen Reinheit bei einer Rate von nur 1 kHz verantwortlich. Das FTK-System wurde so konzipiert, dass es Informationen über die Spuren aller geladenen Teilchen mit $p_T > 1 \text{ GeV}$ im inneren ATLAS-Detektor liefert, welche starke Unterscheidungsmerkmale für die Auswahl interessanter Ereignisse liefern. Das Training des FTK-Systems, das durchgeführt wurde, um eine hohe Effizienz bei der Spurensuche zu erreichen, und die Anpassungsfähigkeit des Systems an neue Detektorbedingungen werden vorgestellt. Darüber hinaus wird die Anwendbarkeit des Systems zur Auswahl von Ereigniskandidaten, die langlebige Teilchen enthalten, beschrieben.

Contents

1	Introduction	1
2	Theory and motivation	3
2.1	Elementary particles and fundamental forces	3
2.2	Mathematical formalism of the Standard Model	5
2.3	Unanswered Questions in Particle Physics	11
2.4	Supersymmetry	12
2.5	Long-lived particles in Supersymmetry	16
2.6	Long-lived particles in non-SUSY extensions of the Standard Model	22
3	Experimental Setup	25
3.1	The Large Hadron Collider	25
3.2	Phenomenology of proton-proton collisions	26
3.3	The ATLAS Detector	32
3.4	Trigger	44
4	Event reconstruction	47
4.1	Standard tracks	47
4.2	Large Radius tracks	49
4.3	Primary vertices	51
4.4	Secondary Vertices	53
4.5	Jets	57
4.6	Electrons	61
4.7	Muons	62
4.8	Missing transverse momentum	62
5	The ATLAS Fast Tracker	63
5.1	Motivation for tracking in the trigger	64
5.2	Track finding with the FTK	64

5.3	Configuring the FTK	67
5.4	Production of pattern bank for commissioning of 128-PU system	72
5.5	Adaptation of the FTK to new detector conditions	75
6	Search for long-lived particles with displaced vertices in multi-jet events	81
6.1	Motivation	82
6.2	Analysis strategy	83
6.3	Data and simulated event samples	84
6.4	Event selection	85
6.5	Background estimation	93
6.6	Alternative background estimation	104
6.7	Uncertainties	123
6.8	Results	126
6.9	Statistical analysis	128
6.10	Interpretation	130
7	Conclusion	135
	Bibliography	139
	List of Figures	149
	List of Tables	153

Introduction

Exploration is fundamental to human nature. To this day, humans are exploring the depths of the ocean, the vastness of space, and even the blueprints of life itself. The work presented in this thesis is an exploration of the fabric of the universe. Particle physics is the study of the nature and interactions of the smallest building blocks of matter, called elementary particles. The Large Hadron Collider (LHC), which accelerates and collides protons at unprecedented energies, allows for the exploration of elementary particles at the energy frontier. The proton collisions are recorded by large particle detectors, one of which is the ATLAS detector. A great achievement of the LHC was the discovery of the Higgs Boson in 2012 by both the CMS and ATLAS experiments. This discovery completed the Standard Model - an impressive theory which describes the nature of all of the known elementary particles and how they interact with three of the four fundamental forces. But there are several missing pieces in our understanding of the universe which the Standard Model fails to provide an answer for. What is the nature of the 85% of matter which is 'dark'? What is the origin of neutrino masses? How is the mass of the Higgs Boson stabilized? Only through exploration into the unknown may we find answers to these questions.

There is no shortage of theories predicting new physics beyond the Standard Model which could answer these open questions. One such theory is Supersymmetry, an extension of the Standard Model with the ability to stabilize the Higgs mass, provide a candidate dark matter particle, and unify the fundamental forces at higher energy scales. Supersymmetric particles have long been sought after at the LHC, starting in Run-1 in 2009. But on the date of writing, there have been no obvious signs that supersymmetric particles are being produced at the LHC. This motivates the effort to search for supersymmetry in more challenging regions of parameter space. Are supersymmetric particles being produced in LHC collisions but evading detection in analysis so far?

If the new particles which are being produced in LHC collisions have a long ($\gtrsim 0.01$ ns) lifetime, it

is possible that they could have evaded detection until now. Particles with long lifetimes are present in a wide range of theories predicting new physics beyond the Standard Model, including Supersymmetry. New, long-lived particles being produced in LHC collisions could traverse a macroscopic distance before decaying within the detector leading to unconventional signatures. Searches for long-lived particles therefore require dedicated techniques, as the detectors and reconstruction algorithms used at the LHC were designed with promptly-decaying particles in mind.

The work presented in this document focuses on a search for one specific signature of long-lived particles, called displaced vertices. A displaced vertex may be reconstructed when a long-lived particle travels a discernible distance ($\gtrsim 4$ mm) into the inner tracker before decaying into multiple charged particles. Searches for displaced vertices with a high mass and track multiplicity are unique for several reasons. First, there are no long-lived particles in the Standard Model which produce high mass ($m_{\text{DV}} > 10$ GeV) displaced vertices. These searches therefore benefit from having very low background. Secondly, searches for displaced vertices are agnostic to many of the properties of the long-lived particle itself. This means that these searches may have sensitivity to a wide range of theories predicting new, massive long-lived particles. In this thesis I present the first search for high mass, multi-track displaced vertices and multiple energetic jets with the ATLAS detector using the full Run-2 dataset. I contributed to nearly all aspects of this search, including the event selection, background estimation, study of systematic uncertainties, and the interpretation of the results.

The very first step of exploration at the LHC begins with the trigger - the system which decides which proton-proton collisions to read off the detector for further analysis. With bunches of protons colliding every 25 ns, the trigger system must make its decisions very quickly. The tracking of charged particle trajectories is a powerful discriminant for selecting interesting events. It is also a computationally expensive task which cannot be performed for each event in the triggering process. In this thesis I will present the ATLAS Fast TracKer (FTK) system, which is a hardware-based track finder envisioned to be integrated into the ATLAS trigger system. I contributed to the project by ensuring that the system was ready to be adapted quickly to new detector conditions during data-taking. Furthermore, I developed the system to meet its full design size specifications during commissioning at the end of Run-2 data-taking and trained it to achieve a high track-finding efficiency. I will also discuss how this system could be used to trigger directly on signature of long-lived particles.

This document begins with a theoretical introduction to particle physics and motivation to search for long-lived particles in Chapter 2. Chapter 3 describes the setup of the LHC and the ATLAS detector. Chapter 4 describes how the signals from proton-proton collisions in the detector are reconstructed for use in higher-level analysis. In Chapter 5, the ATLAS FastTracKer is presented, with a focus on the areas of the project to which I contributed. Chapter 6 presents the search for displaced vertices in multijet events using the full Run-2 dataset with the ATLAS detector. Chapter 7 provides concluding remarks and some insight into future directions for research.

Theory and motivation

In this chapter, the search for long-lived particles using displaced vertices is motivated. Sections 2.1 and 2.2 introduce the elementary particles which have been observed in nature and the theory which is used to describe them, called the Standard Model. In Section 2.3, several questions about the universe which the Standard Model leaves unanswered are discussed. These open questions motivate the introduction of Supersymmetry (SUSY) in Section 2.4. Long-lived particles are predicted in many supersymmetric extensions of the Standard Model, a few of which are outlined in Section 2.5. Lastly, Section 2.6 introduces a few non-supersymmetric extensions of the Standard Model which predict the existence of long-lived particles, demonstrating their importance beyond SUSY scenarios.

2.1 Elementary particles and fundamental forces

Elementary particles are the basic building blocks of the universe. They are characterized by their observable quantities, for example mass, charge, and spin, which is an intrinsic form of angular momentum. The particles which make up the visible matter of the universe - called fermions - have spin $1/2$. Particles which communicate the fundamental forces between the matter particles - called bosons - have integer spin values. There are four fundamental forces which have been observed to date: the strong, electromagnetic, weak, and gravitational forces. The elementary particles which have been observed in nature are shown in Figure 2.1.

Fermions are characterized according to which fundamental forces they carry charge under. Fermions that participate in the strong interaction are called quarks. There are six types of quarks, each with three varieties of color charge. Each quark has an electric charge of either $+2/3$ (up-type), or $-1/3$ (down-type). Fermions that do not carry color charge are called leptons. There are six leptons, three of which carry electric charge, and three of which do not, called neutrinos. Both quarks and




Figure 2.1: Elementary particles which have been observed in nature [1].

leptons are organized into generations, indicated by the columns in Figure 2.1. The first generation contains the lightest of each particle type and comprises most of the stable and ordinary matter in the observed universe. The masses of the charged leptons and quarks increase with each generation.

Each fermion in the SM has an antimatter partner particle. These antimatter particles have exactly the same properties as their particle-counterparts, but with opposite charges. This means that, for example, an up-type quark with charge $+2/3$ and “red” color charge has an antimatter counterpart with charge $-2/3$ and “anti-red” color charge.

Another important characteristic of elementary particles is their chirality. Like spin, chirality is a purely quantum mechanical phenomena without a classical explanation. The chirality is closely related to a particle’s helicity, which indicates the orientation of the particle’s spin with respect to the direction of its momentum. A particle has right-handed helicity if the spin is in the same direction as its momentum, and vice versa for particles with left-handed helicity. For a particle moving with velocity less than the speed of light ($v < c$), it is possible to boost to another reference frame where the helicity of the particle changes. Therefore, helicity is not a Lorentz invariant quantity. Chirality is a mathematically more abstract quantity which is Lorentz invariant. The chirality of fermions is important because left-handed and right-handed fermions interact differently with some of the bosons.

2.2 Mathematical formalism of the Standard Model

The elementary particles introduced in the previous section are described mathematically with a relativistic quantum field theory, called the Standard Model (SM). In the SM, elementary particles are described by fields. The evolution of these fields is described by the SM Lagrangian, from which the phenomenology of all SM particles can be predicted. The SM Lagrangian is built so that it is Lorentz invariant, gauge symmetric (invariant under transformations of the SM gauge group), renormalizable (provides finite predictions for physical observables), and minimal (uses the fewest terms needed to describe observed particles and interactions). The SM gauge group is $SU(3)_c \times SU(2)_L \times U(1)_Y$, and its Lagrangian can be split into several parts:

$$\mathcal{L}_{\text{SM}} = \mathcal{L}_{\text{EW}} + \mathcal{L}_{\text{QCD}} + \mathcal{L}_{\text{Higgs}} + \mathcal{L}_{\text{Yukawa}}, \quad (2.1)$$

both of which will be explained in the following sections.

The electroweak sector

The electroweak sector of the SM is described by the $SU(2)_L \times U(1)_Y$ gauge group. It describes the unified electromagnetic and weak forces. The construction of this gauge group begins with the quantum field theoretical treatment of the electromagnetic force with a $U(1)_{\text{em}}$ gauge theory, called Quantum Electrodynamics (QED). In QED, interactions between elementary particles carrying an electric charge are mediated by the exchange of a photon. Photons are massless and electrically neutral spin-1 particles.

The weak force is responsible for processes such as β decay, for example the decay of a neutron into a proton, electron, and antineutrino. An important difference from QED is that the weak interaction is a short-range interaction with massive mediator bosons: the $W^\pm$ and the Z^0 . Furthermore, the weak charged-current interactions (those mediated by $W^\pm$) only couple to left-handed fermions and right-handed anti-fermions.

The Glashow-Salam-Weinberg model describes how QED and the weak force are different features of one unified theory [2, 3], called electroweak interactions. This model extends the $U(1)_{\text{em}}$ gauge theory to the electroweak (EW) gauge group $SU(2)_L \times U(1)_Y$, which contains one $U(1)$ gauge boson, B_μ , and three $SU(2)$ gauge bosons, W_μ^a ($a = 1, 2, 3$). The lagrangian is given by

$$\mathcal{L}_{\text{EW}} = \sum_i \bar{\psi}_i (i\gamma^\mu D_\mu) \psi_i - \frac{1}{4} B_{\mu\nu} B^{\mu\nu} - \frac{1}{4} W_{\mu\nu}^a W^{a\mu\nu} \quad (2.2)$$

where the sum i is over all fermion fields ψ . The first term contains the covariant derivative, D_μ ,

which is given by

$$D_\mu = \partial_\mu + igT_a W_\mu^a + ig'YB_\mu, \quad (2.3)$$

where g and g' are gauge couplings. The generators of the $U(1)$ and $SU(2)$ gauge groups are given by the hypercharge, Y , and a set of three 2×2 complex matrices, T_a , respectively. The generators T_a for right-handed fermions are 0, reflecting the fact that the W_μ bosons do not interact with the right-handed fermions. The eigenvalues of the non-zero T_a for left-handed fermions are called the weak isospin. The $U(1)_Y$ symmetry describing electroweak interactions differs from the $U(1)_{\text{em}}$ symmetry describing the electromagnetic force, but their charges are related via

$$Q_{\text{em}} = \frac{1}{2}Y + T_3, \quad (2.4)$$

where T_3 is the third component of the weak isospin.

The remaining two terms in $\mathcal{L}_{\text{EW}}$ are called the kinetic terms, and are given by

$$B_{\mu\nu} = \partial_\mu B_\nu - \partial_\nu B_\mu, \quad (2.5)$$

and

$$W_{\mu\nu}^a = \partial_\mu W_\nu^a - \partial_\nu W_\mu^a + gf^{abc}W_\mu^b W_\nu^c, \quad (2.6)$$

where f^{abc} is the structure constant of the group, defined by the commutator $[T_a, T_b] = if^{abc}T_c$. The extra term in $W_{\mu\nu}^a$ arises because the $SU(2)$ gauge symmetry is non-abelian, meaning the generators of the group do not commute. This gives rise to new interactions between the gauge bosons W_μ^a of the group.

Quantum Chromodynamics

The SM gauge group is completed with the addition of an $SU(3)_c$ gauge group, which is the mathematical framework of Quantum Chromodynamics (QCD). QCD describes the strong force, which acts on all particles that carry color charge. The QCD part of the Lagrangian, $\mathcal{L}_{\text{QCD}}$, is written down in a form analogous to $\mathcal{L}_{\text{EW}}$. The generators of the $SU(3)$ group are eight 3×3 complex matrices. Thus, QCD has eight massless gauge bosons, called gluons, and three conserved charges, called color. The $SU(3)$ group is non-abelian, meaning that the gluons are self interacting.

An important feature of QCD is color confinement, which states that elementary particles which carry color charge are always confined to color singlet states. Quarks are always confined to colorless hadrons and are never observed on their own in nature. There is no analytical proof for color confinement, however it is thought to arise due to the self interaction of gluons [4].

Running of the coupling

The coupling constant is a dimensionless quantity which determines the strength of a force exerted in an interaction. It is dependent on the energy scale at which the coupling is being probed. Conceptually this arises because of the screening (or enhancement) of a force due the presence of quantum fluctuations of virtual particles. The amount of these fluctuations grow with higher energies.

The QED coupling is small at low energies and increases gradually at higher energies. The QCD coupling (α_s) behaves quite differently, due to the self interaction of the gluons. In interactions with low momentum transfer ($q = 1 \text{ GeV}$), the QCD coupling is large ($\alpha_s \approx 1$) and therefore QCD is non-calculable with perturbative measures. This energy regime describes the physics of bound hadronic states. The QCD coupling decreases logarithmically as the energy increases, which is called asymptotic freedom. At energies relevant for particle collisions ($q > 100 \text{ GeV}$), α_s is sufficiently small so that perturbation theory can be used.

Spontaneous electroweak symmetry breaking

The fermions observed in nature have been measured to have a non-zero mass. However including a mass term of the form $m\psi\bar{\psi}$ for chiral fermions in the SM Lagrangian is not allowed, because it would violate gauge invariance. Furthermore, writing down mass terms for the gauge bosons of the form $\frac{1}{2}m^2 G_\mu G^\mu$ would also violate gauge invariance. This is not a problem for the massless gauge bosons of QED and QCD, but the weak force has been observed to be mediated by massive gauge bosons. The generation of masses for the fermions, $W^\pm$ and Z^0 bosons are explained by the Brout-Englert-Higgs mechanism ("Higgs mechanism") and spontaneous electroweak symmetry breaking.

The Higgs mechanism introduces a new scalar field, Φ , with a potential of the form:

$$V(\Phi) = \mu^2 \Phi^\dagger \Phi + \frac{\lambda^2}{2} (\Phi^\dagger \Phi)^2, \quad \text{where } \Phi \equiv \frac{1}{\sqrt{2}} \begin{pmatrix} \sqrt{2}\phi^+ \\ \phi^0 + ia^0 \end{pmatrix}. \quad (2.7)$$

The field components of Φ are split according to their charge and CP-symmetry, which describes how they transform under a charge-parity transformation. ϕ^0 and a^0 are the CP-even and CP-odd neutral components of Φ , respectively, and ϕ^+ is the complex charged component [5]. If $\mu^2 > 0$, the potential has a shape similar to what is shown in Figure 2.2(a), where the function is symmetric around its minimum. If $\mu^2 < 0$, the potential changes shape such that the minima no longer correspond to positions around which the potential is symmetric (see Figure 2.2(b)). In this case, the field Φ develops a vacuum expectation value at its minimum when spontaneous symmetry breaking occurs:

$$SU(2)_L \times U(1)_Y \rightarrow U(1)_{\text{em}}.$$




Figure 2.2: Schematic of Higgs potential in two dimensions. The shape of the potential in (b) causes spontaneous symmetry breaking to occur, because the potential is no longer symmetric around its minimum, as it is in (a).

Through the mechanism of spontaneous symmetry breaking, the W_μ^3 and B_μ fields mix to generate the massive Z^0 gauge boson and the massless photon A_μ :

$$A_\mu = \cos \theta_W B_\mu + \sin \theta_W W_\mu^3 \quad (2.8)$$

$$Z_\mu = -\sin \theta_W B_\mu + \cos \theta_W W_\mu^3, \quad (2.9)$$

where θ_W is the weak mixing angle, defined as

$$\tan \theta_W = \frac{g'}{g}. \quad (2.10)$$

The massive $W^\pm$ are linear combinations of the W_μ^1 and W_μ^2 gauge bosons:

$$W^\pm = \frac{1}{\sqrt{2}}(W^1 \mp iW^2). \quad (2.11)$$

The Lagrangian for the new scalar field is given by

$$\mathcal{L}_{\text{Higgs}} = |D_\mu \phi|^2 - V(\Phi), \quad (2.12)$$

where D_μ is defined as in Eq. 2.3. The first term generates the masses of the gauge bosons, which are

given by

$$M_{W^\pm} = \frac{1}{2}g\nu \quad M_{Z^0} = \frac{\nu}{2}\sqrt{g^2 + g'^2},$$

where ν is the vacuum expectation value, measured to be $\nu \approx 246$ GeV.

Fermions acquire mass via interactions with the Higgs field directly, called Yukawa couplings.¹ The Yukawa terms combine both left-handed and right-handed fermion fields and are of the form

$$\mathcal{L}_{\text{Yukawa}} \supset Y_i \bar{\psi}_L \phi \psi_{i,R}. \quad (2.13)$$

The Yukawa coupling strengths, Y_i , which are proportional to the fermion's observed mass via $\sqrt{2}m_f/\nu$, are free parameters of the Standard Model.

The Higgs mass is also a free parameter of the Standard Model and therefore cannot be predicted directly. In 2012, the ATLAS and CMS experiments at the Large Hadron Collider announced the observation of a new particle with mass $m \approx 125$ GeV [7, 8]. This particle's properties are so far consistent with those of the Standard Model Higgs boson.

2.2.1 Long-lived particles

The proper lifetime of a particle, τ , is inversely proportional to its decay width, Γ , which is the probability per unit time that the particle will decay. It is given by

$$\Gamma = \tau^{-1} = \frac{1}{2m_\chi} \int d\Pi_f |\mathcal{M}(m_\chi \rightarrow \{p_f\})|^2. \quad (2.14)$$

Here, m_χ is the mass of the particle, $\mathcal{M}$ is the matrix element for the particle's decay into its decay products, $\{p_f\}$, and $d\Pi_f$ is the phase space for its decay. A particle will have a long lifetime if it has a small matrix element or a small phase space for the decay. This can occur if:

1. The coupling of the particle to its decay products are small
2. The decay of the particle is mediated by a higher-dimensional operator, so that the coupling constant is suppressed by the scale at which the decay is mediated
3. The available phase space for the particle's decay is small

Figure 2.3 displays the mass and proper lifetime of some particles in the Standard Model. The plot is split into three separate regions. The first, labeled “Detector-Prompt”, refers to particles which

¹ Neutrinos are an exception. In the Standard Model, neutrinos are massless. However, the observation of neutrino oscillations between different flavors indicates that neutrinos have a small but non-zero mass [6]. Introducing a mechanism by which neutrinos obtain mass in the Standard Model would require the extension of its current particle content, or the inclusion of lepton violating processes which have yet to be observed. Therefore, the origins of neutrino mass are thought to come from some mechanism beyond the Standard Model.

have a short enough lifetime such that the typical distance between their production and decay points is shorter than the spatial resolution of the detector. The intrinsic spatial resolution of the ATLAS pixel detector is $10\text{ }\mu\text{m}$ and $115\text{ }\mu\text{m}$ in the $R\phi$ and z directions, respectively [9]. The second region, labeled “Detector-Stable”, refers to those particles with lifetimes long enough that they typically travel through and decay outside of the ATLAS detector. Thus, for purposes of data analysis, this particle is considered stable. The third and most interesting region for this thesis is the un-shaded region in the center. Here, particles have lifetimes such that they have a high chance of decaying in the detector but are long-lived enough to be distinguished from promptly decaying particles.




Figure 2.3: The masses and proper lifetimes of some particles in the Standard Model. The shaded regions on the left and right represent regions where the particles are considered prompt and stable, respectively. The un-shaded region in the middle contains particles which have a significant probability of decaying within the ATLAS detector [10].

The probability that a particle survives for time t before decaying is described by an exponential decay distribution, whose time constant depends on the particle’s proper lifetime τ and Lorentz factor, γ ,

$$P(t) = e^{-t/(\gamma\tau)}. \quad (2.15)$$

This has important implications for the signatures of long lived particles. For particles with a proper lifetime corresponding to the un-shaded region in Figure 2.3, a certain fraction of these particles will decay promptly, and a certain fraction will pass through the detector without decaying. The exact amount of these fractions depends on the particle’s lifetime and boost factor when it is produced. This is important because it implies that searches for promptly-decaying particles also have some sensitivity to long-lived particles, and vice versa.

2.3 Unanswered Questions in Particle Physics

The Standard Model has triumphed with the prediction and subsequent discoveries of the gauge bosons and the Higgs boson, and has been demonstrated to accurately describe particle interactions across a vast range of energies. However, there remain some observations of the universe which the Standard Model is unable to explain. Two of these open questions are summarized in the following paragraphs. They concern the stabilization of the Higgs mass and the large portion of the universe known as ‘dark matter’ which has yet to be observed.

The Hierarchy problem addresses the fact that the Higgs mass is on the order of the electroweak scale, even though it is unstable to quantum corrections. The Higgs mass is subject to quantum corrections from loop diagrams with fermions and bosons, as shown in Figure 2.4. The correction to




Figure 2.4: Feynman diagrams for (a) bare Higgs mass, quantum loop correction with (b) top quark loop, (c) $W^\pm$ or Z^0 boson loop, and (d) Higgs boson loop.

the Higgs mass, m_H^2 , for a loop diagram containing a fermion, as shown in Figure 2.4(a), is given by,

$$\Delta m_H^2 = -\frac{|Y_f|^2}{8\pi^2} \Lambda_{UV}^2 + \dots, \quad (2.16)$$

where Y_f is the Yukawa coupling to the fermion f , and terms not shown are proportional to m_f^2 and grow at most logarithmically with Λ_{UV} . Here, Λ_{UV} is the ultraviolet momentum cutoff, a very large energy scale, above which new physics may alter the high-energy behavior of the theory. The dependence of Δm_H^2 on $|Y_f|^2$ indicates that the more massive a particle is, the larger its contribution to the Higgs mass correction.

Without these large quantum contributions being fine-tuned to cancel out, the Higgs mass would be expected to lie around the Planck mass (10^{19} GeV). While there is nothing in the Standard Model which forbids fine-tuning, this observation has motivated a search for an elegant solution to stabilize the Higgs mass.

Dark matter is another feature of the universe which the Standard Model does not provide an explanation for. The existence of dark matter was proposed in the 1930's, when it was observed that luminous objects in the universe, such as stars and even entire galaxies, move much faster than what is

expected from the gravitational pull of other luminous objects in their vicinity [11]. This and further astrophysical observations, for example the cosmic microwave background and the structure formation of the universe, strongly indicate that a significant portion of the matter in the universe ($\sim 80\%$) is ‘dark’, i.e. no more than weakly interacting with electromagnetic radiation.

The composition of dark matter is unknown. A candidate dark matter particle needs only to interact with visible matter predominantly through the gravitational interaction. Thus, in principle, dark matter could be comprised of baryonic matter in the SM, such as protons or neutrons. However, the study of gravitational microlensing, the cosmic microwave background, and the lack of abundance of heavy elements in the universe do not support this scenario [12]. Instead, dark matter is thought to be composed of a particle that is stable on cosmological timescales, was non-relativistic during galaxy formation, and which does not interact strongly with electromagnetic radiation. In other words, dark matter is likely to be composed of a new particle which is not already present in the Standard Model.

2.4 Supersymmetry

Supersymmetry (SUSY) is one BSM theory which can address some of these unanswered questions. SUSY extends the Standard Model to include a symmetry relating fermions and bosons. A supersymmetric transformation turns a boson into a fermion, and vice versa, leaving all other quantum numbers unchanged. This additional symmetry implies the existence of additional elementary particles called *superpartners*, which differ from their Standard Model partner particles only by $1/2$ unit of spin. Extending the Standard Model with this symmetry has the potential to solve the Hierarchy problem in the form of cancellations to the Higgs mass corrections between SM particles and their superpartners. Many supersymmetric models also provide a dark matter candidate particle. Furthermore, some SUSY models predict that at some high energy scale, the electroweak and strong forces unify under a single gauge group. For these reasons, Supersymmetry has been extensively studied and searched for at the Large Hadron Collider.

The following paragraphs describe the Minimal Supersymmetric Standard Model (MSSM), i.e. the Standard Model with the absolute minimum necessary requirements to add this additional symmetry relating fermions and bosons. Then, the mechanism of supersymmetry breaking and R-parity conservation will be discussed.

2.4.1 The Minimal Supersymmetric Standard Model

The single-particle states of supersymmetry are organized into *supermultiplets*, each containing both a fermionic and a bosonic state (the SM particle and its superpartner). The *chiral supermultiplets* contain the SM fermions and their bosonic superpartners, which are listed in Table 2.1 [13]. Each SM fermion has two scalar superpartners, one partner to the left-handed fermions, and the other to the

right-handed fermions. The superpartners in chiral supermultiplets are named by adding the prefix ‘s-’ to the corresponding SM fermion’s name (i.e. bottom to sbottom).

The massless SM gauge bosons (prior to electroweak symmetry breaking) and their fermionic superpartners, called gauginos, reside in *gauge supermultiplets*. These are shown in Table 2.2. Gauginos are named by adding the suffix ‘-ino’ to the SM gauge boson’s name (i.e. gluon to gluino).

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

Table 2.1: Chiral supermultiplets in the Minimal Supersymmetric Standard Model [13]. The superpartners of SM particles are denoted with a tilde ($\sim$).

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

Table 2.2: Gauge supermultiplets in the Minimal Supersymmetric Standard Model [13]. The superpartners of SM gauge bosons are denoted with a tilde ($\sim$).

Though not included yet in the Standard Model, the elementary particle responsible for gravity, the spin-2 graviton, also has a superpartner particle. This superpartner, called the gravitino, has a spin of 3/2.

In order to generate the Yukawa couplings for both the up-type and down-type quarks in a supersymmetric theory, there must be two Higgs chiral supermultiplets, H_u and H_d . Through Yukawa couplings, the H_u generates mass for the up-type quarks, while the H_d generates mass for the down-type quarks and charged leptons. Electroweak symmetry breaking proceeds in an analogous way to the Standard Model (Section 2.2) and each Higgs doublet acquires a vacuum expectation value, i.e.

$$\langle H_u \rangle = \frac{1}{\sqrt{2}} \begin{pmatrix} 0 \\ v_u \end{pmatrix}, \quad \langle H_d \rangle = \frac{1}{\sqrt{2}} \begin{pmatrix} v_d \\ 0 \end{pmatrix}. \quad (2.17)$$

The vacuum expectation values are related to the mass of the Z^0 boson (m_Z) and the electroweak gauge couplings via the relation:

$$v_u^2 + v_d^2 = v^2 = 2m_Z^2/(g^2 + g'^2) \approx (174 \text{ GeV})^2. \quad (2.18)$$

The ratio of these vacuum expectation values is a free parameter of the theory, defined with β as:

$$\tan \beta \equiv v_u/v_d. \quad (2.19)$$

After electroweak symmetry breaking, there are a total of five physical Higgs bosons, the lightest of which corresponds to the Standard Model Higgs boson with $m_h = 125 \text{ GeV}$.

Due to electroweak symmetry breaking, the neutral higgsinos and neutral electroweak gauginos mix together to form four mass eigenstates, called neutralinos. The neutralino mass term in the supersymmetric Lagrangian is given by

$$\mathcal{L}_{\text{neutralino mass}} = -\frac{1}{2}(\psi^0)^T \mathbf{M}_{\tilde{\chi}^0} \psi^0 + \text{charge conjugate}, \quad (2.20)$$

where $\psi^0 = (\tilde{B}, \tilde{W}^0, \tilde{H}_d^0, \tilde{H}_u^0)$ and $\mathbf{M}_{\tilde{\chi}^0}$ is the neutralino mass matrix given by

$$\mathbf{M}_{\tilde{\chi}^0} = \begin{pmatrix} M_1 & 0 & -\cos \beta \sin \theta_W m_Z & \sin \beta \sin \theta_W m_Z \\ 0 & M_2 & \cos \beta \cos \theta_W m_Z & -\sin \beta \cos \theta_W m_Z \\ -\cos \beta \sin \theta_W m_Z & \cos \beta \cos \theta_W m_Z & 0 & -\mu \\ \sin \beta \sin \theta_W m_Z & -\sin \beta \cos \theta_W m_Z & -\mu & 0 \end{pmatrix}, \quad (2.21)$$

where M_2 , M_1 , and μ are the wino, bino and higgsino mass terms in the supersymmetric Lagrangian. The eigenvalues of this matrix correspond to the masses of the four physical neutralinos ($\tilde{\chi}_i^0$). The subscript i indicates the mass eigenstate, with $i = 1$ being the lightest. Similarly, the charged higgsinos and winos mix to form two mass eigenstates with differing charge, called charginos ($\tilde{\chi}_i^\pm$). The chargino mass term in the supersymmetric Lagrangian is given by

$$\mathcal{L}_{\text{chargino mass}} = -\frac{1}{2}(\psi^\pm)^T \mathbf{M}_{\tilde{\chi}^\pm} \psi^\pm + \text{charge conjugate}, \quad (2.22)$$

where $\psi^\pm = (\tilde{W}^+, \tilde{H}_u^+, \tilde{W}^-, \tilde{H}_d^-)$ and the chargino mass matrix is given by

$$\mathbf{M}_{\tilde{\chi}^\pm} = \begin{pmatrix} \mathbf{0} & \mathbf{X}^T \\ \mathbf{X} & \mathbf{0} \end{pmatrix}, \text{ with } \mathbf{X} = \begin{pmatrix} M_2 & \sqrt{2} \sin \beta m_W \\ \sqrt{2} \sin \beta m_W & \mu \end{pmatrix}, \quad (2.23)$$

and the eigenvalues of $\mathbf{M}_{\tilde{\chi}^\pm}^\dagger \mathbf{M}_{\tilde{\chi}^\pm}$ give the masses of the two physical charginos.

2.4.2 Supersymmetry breaking

If SUSY were realized in nature, superpartners would have already been discovered at collider experiments, with the same mass as their Standard Model counterpart. Given that no superpartners have been discovered yet, SUSY must be a spontaneously broken symmetry, yielding supersymmetric particles which are much heavier than their Standard Model counterparts. In order to generate spontaneous symmetry breaking, the MSSM needs to be extended to include new particles and interactions at higher mass scales. The exact mechanism of spontaneous symmetry breaking is unknown, however two scenarios and their phenomenological consequences on long-lived SUSY particles will be discussed in Sections 2.5.1 and 2.5.3. Instead of assuming a single spontaneously broken scenario, supersymmetry is broken explicitly in the MSSM by adding to its effective Lagrangian *soft supersymmetry breaking* terms, which do not affect the UV behavior of the theory. Thus, one can consider the MSSM effective Lagrangian of the form

$$\mathcal{L} = \mathcal{L}_{\text{SUSY}} + \mathcal{L}_{\text{soft}} . \quad (2.24)$$

Here, the $\mathcal{L}_{\text{SUSY}}$ term contains gauge and Yukawa interactions and the $\mathcal{L}_{\text{soft}}$ contains mass terms and positive dimension coupling parameters [13]. Supersymmetry breaking occurs exclusively in the $\mathcal{L}_{\text{soft}}$ term, which ensures that the act of supersymmetry breaking leaves the important cancellations to the Higgs mass intact.

2.4.3 R-parity

The most general gauge-invariant and renormalizable superpotential includes the following terms,

$$W_{\Delta L=1} = \frac{1}{2} \lambda^{ijk} L_i L_j \bar{e}_k + \lambda'^{ijk} L_i Q_j \bar{d}_k + \mu'^i L_i H_u \quad (2.25)$$

$$W_{\Delta B=1} = \frac{1}{2} \lambda''^{ijk} \bar{u}_i \bar{d}_j \bar{d}_k \quad (2.26)$$

which violate lepton number and baryon number, respectively. Allowing for these terms in the MSSM superpotential would result in un-discovered processes such as proton decay, which would occur if both the λ' and λ'' couplings are non-zero. In addition, each of the above coupling terms is constrained by fundamental tests of lepton number and baryon number conservation, as described in [14].

To avoid these terms, many SUSY models impose R-parity conservation, defined as

$$P_R = (-1)^{3(B-L)+2s} , \quad (2.27)$$

where B, L and s stand for baryon-, lepton-, and spin- quantum numbers, respectively. All Standard Model particles have an R-parity of +1, while supersymmetric particles have an R-parity of -1. R-parity

conservation implies that all SUSY particles will be produced in pairs at colliders, will decay eventually to the lightest supersymmetric particle (LSP), and that the LSP is stable [15]. Stable, neutral LSP's provide a viable candidate for WIMP dark matter. In the MSSM, the lightest neutralino is usually assumed to be the LSP, because it is the favored candidate for dark matter. However, models with sneutrino and gravitino LSPs exist as well.

2.5 Long-lived particles in Supersymmetry

On the date of writing, no evidence of superpartner particles has been observed at the LHC. The majority of these searches have focused on promptly-decaying SUSY particles. However, many SUSY scenarios involve long-lived particles as well. Due to the challenges with reconstructing signatures of long-lived particles in the ATLAS detector, SUSY scenarios with long-lived particles have been less explored, and there are still many viable regions of parameter space which are accessible at LHC energies. In the following sections, various ways in which long-lived particles arise in supersymmetric theories are discussed and the current status of searches for these particles at the LHC is provided.

2.5.1 Gauge-mediated SUSY breaking

In Gauge-mediated SUSY breaking (GMSB), the supersymmetric SM gauge group is extended with an additional secluded sector, which is responsible for symmetry breaking [16, 17]. The SUSY breaking is communicated to the MSSM via heavy messenger states, which are charged under the SM gauge group. As a result, SUSY breaking is only communicated to the MSSM via its interactions with SM gauge fields. In this class of models, the gravitino is the LSP, which has a mass scaling with

$$m_{3/2} \sim \frac{F}{M_{pl}} . \quad (2.28)$$

Here, F is the energy at which SUSY breaking occurs and the suppression from the Planck mass, M_{pl} (10^{19} GeV), results in a very light gravitino. The next-lightest supersymmetric particle (NLSP) decays to the gravitino and a SM particle via higher dimensional operators which are suppressed by $1/F$. The most natural choices for the NLSP are the lightest neutralino or the stau. The decay rate of the stau NLSP is given by

$$\Gamma(\tilde{\tau} \rightarrow \tau \tilde{G}) = k^2 \left(\frac{m_{\tilde{\tau}}}{100 \text{ GeV}} \right)^5 \left(\frac{100 \text{ TeV}}{\sqrt{F}} \right)^4 2 \times 10^{-3} \text{ eV} , \quad (2.29)$$

where $k < 1$ is a model dependent coefficient. The decay rate of the neutralino LSP to a photon and the gravitino is given by an analogous formula depending instead on $m_{\chi_1^0}^5$ and with an additional branching factor ratio to take into account the (likely suppressed) decays to a gravitino and a Z or

neutral Higgs boson instead [16]. In the range of low-scale to intermediate-scale SUSY breaking ($F \sim 10^5$ GeV to $F \sim 10^{10}$ GeV), the NLSP lifetime varies from prompt to macroscopic.

Searches for charged leptons or photons at the LHC which do not point back to the interaction point or which have delayed arrival times to detector components provide sensitivity to GMSB models with a long-lived NLSP. The exclusion limits on the mass of the NLSP as a function of its lifetime in two simplified GMSB models are shown in Figure 2.5. In Figure 2.5(a), the NLSP is assumed to be either a single stau or a selectron, smuon, and stau as co-NLSPs. Stau NLSPs with a lifetime of 0.1 ns are excluded with masses up to 340 GeV [18]. A similar search for displaced leptons in the CMS experiment excludes staus NLSPs with a lifetime of 0.06 ns up to with masses of 400 GeV [19]. In Figure 2.5(b), the NLSP is assumed to be a neutralino, which then decays dominantly to a photon at the gravitino. Neutralinos with proper decay lengths of 1 m are excluded with masses up to 525 GeV from a dedicated search for delayed photons in the CMS experiment [20].




Figure 2.5: In (a), the exclusion limits on the NLSP lifetime as a function of the mass of the NLSP are shown for a simplified GMSB model with either the stau as the single NLSP or a selectron, smuon and stau as co-NLSPs [18]. The shaded regions are excluded. In (b), the exclusion limits on the $c\tau(\text{NLSP})$ as a function of the NLSP mass are shown for a simplified GMSB model with a neutralino NLSP [20]. The regions to the left of the curves are excluded.

2.5.2 Mini-split SUSY

The lack of any evidence for sparticles at the LHC to date make it difficult for viable supersymmetric theories to solve the hierarchy problem without introducing more fine-tuning itself. Mini-split SUSY [21] focuses on the other appealing aspects of SUSY theories, for example the existence of a viable dark matter candidate and coupling constant unification. Mini-split SUSY has received notable interest recently because it correctly predicts the SM Higgs mass to lie within $\sim 120 - 150$ GeV [22].

Mini-split SUSY proposes that the masses of scalar sfermions (squarks and sleptons) are at the SUSY breaking scale, and thus inaccessible at LHC energies. However, the masses of the gauginos and higgsinos are $O(1 \text{ TeV})$, which successfully produces gauge coupling unification, regardless of the scalar sparticle masses. This mass hierarchy results in the suppression of gaugino decays to the LSP.

Unlike electroweak gauginos, the gluinos do not have decay channels which can circumvent the high sfermion scale. The gluino must decay via a highly virtual squark into a quark, antiquark and the LSP. This results in a long gluino lifetime given approximately by

$$\tau \simeq 8 \left(\frac{m_S}{10^9 \text{ GeV}} \right)^4 \left(\frac{1 \text{ TeV}}{m_{\tilde{g}}} \right)^5 \text{ s}, \quad (2.30)$$

where m_S is the squark mass [23].

Gluinos carry color charge and thus will hadronize into *R-hadrons* when they have a lifetime above the hadronization timescale (10^{-23} s). An R-hadron is a composite state of a strongly-interacting sparticle with Standard Model quarks and/or gluons. R-hadrons may be electrically neutral, charged, or even doubly charged and may change their internal state as they interact with material when traversing the detector [24]. Searches for R-hadrons at the LHC rely on their direct interaction with detector material or the detection of their decay products. Figure 2.6 shows the exclusion limits on the gluino mass as a function of its lifetime in a split-supersymmetry scenario. These limits were calculated using the results of searches for prompt and long-lived particles using several different signatures in the ATLAS detector. Similar searches have also been performed in the CMS experiment [25–27]. In the model shown in Figure 2.6, the LSP is assumed to be the lightest neutralino with a mass of 100 GeV. R-hadron masses below 2.0 TeV are excluded across the entire lifetime range. It should be noted that these exclusion limits only apply to the specific simplified model, described in Ref. [28]. Other models predicting R-hadrons do not necessarily have such stringent exclusion limits.

2.5.3 Mass-degeneracies in the MSSM

In addition to mass hierarchies, as shown in Mini-Split SUSY, mass degeneracies can also lead to long-lived particles in the MSSM. The nature of the lightest neutralino determines the mass splitting between the lightest chargino and neutralino. If this mass splitting becomes small, the lightest chargino will have a long lifetime.

As discussed in Section 2.4.1, neutralinos and charginos are the mass eigenstates resulting from the mixing of the electroweak gauginos and higgsinos. This mixing and the resulting masses of the neutralinos and charginos are determined by the nature of soft supersymmetry breaking, the details of which are still unknown. The correct electroweak symmetry breaking is easily achieved if the LSP is assumed to be “bino-like”, i.e. $\tilde{\chi}_1^0 \approx \tilde{B}$, resulting in a sizeable mass difference between the lightest chargino and the LSP [13]. However, this is not a requirement. The LSP can also be wino- or higgsino-




Figure 2.6: The exclusion limits on the gluino mass as a function of its lifetime in a split-supersymmetry scenario [29]. The regions below the curves are excluded. The dots represent results for which the particle is assumed to be stable. In this context, stable means escaping the ATLAS detector.

like, which lead to scenarios with compressed mass spectra.

Wino-like LSPs, which arise in models with *anomaly-mediated* SUSY breaking, result in a compressed mass hierarchy between the lightest chargino and neutralino. In Section 2.5.1, the GMSB was introduced, which communicates SUSY breaking to the MSSM via gauge interactions. In a similar manner, gravity-mediated SUSY breaking communicates the symmetry breaking via gravity. If the invisible sector does not have the necessary structure for either gauge-mediated or gravity-mediated SUSY breaking, then communication of SUSY breaking via anomalies becomes dominant. This is called anomaly-mediated SUSY breaking (AMSB) [30]. A defining characteristic of the AMSB are the gaugino masses given by

$$M_i = \frac{\beta(g_i^2)}{2g_i^2} m_{\tilde{G}} , \quad (2.31)$$

where g_i is the gauge coupling constant, $\beta(g_i^2)$ describes the evolution of the coupling with energy scale, and $m_{\tilde{G}}$ is the gravitino mass. The indices $i = 1, 2, 3$ correspond to the gauge groups $U(1)$, $SU(2)$, and $SU(3)$, respectively. The masses of the bino, wino, and gluino then occur with an approximate mass ratio of $M_1 : M_2 : M_3 \approx 3 : 1 : 7$, resulting in a wino-like LSP. The mass splitting between the charginos and LSP is then given by

$$m_{\chi^\pm} - m_{\chi^0} = \frac{M_W^4}{\mu^3} \sin 2\beta + \frac{M_W^4 \tan^2 \theta_W}{(M_1 - M_2)\mu^2} \sin^2 2\beta , \quad (2.32)$$

where μ is the higgsino mass and $\tan \beta$ is the ratio of vacuum expectation values acquired by the up- and down-type Higgs doublets [31]. In the case where the higgsino mass is large, the mass splitting between the chargino and the lightest neutralino is small. This gives rise to a chargino with a long lifetime, which then decays to the LSP and a SM particle, typically a soft charged pion. Searches for short charged particle tracks and for slow-moving, highly ionizing particles in the ATLAS and CMS detectors provide sensitivity to AMSB models with long-lived chargino [25, 32–34]. The exclusion limits on the mass of the chargino as a function of its lifetime in an AMSB model using results from the ATLAS experiment are shown in Figure 2.7(a). The mass limit is much weaker for charginos with a shorter lifetime, due to experimental challenges with reconstructing tracks from short-lived charged particles [35].

Higgsino-like LSPs are motivated by naturalness, which suggests that the higgsino mass should lie near the weak scale, while the bino and wino mass parameters, M_1 and M_2 can be much larger [36]. This results in the three lightest electroweakino states, $\tilde{\chi}_1^0$, $\tilde{\chi}_1^\pm$, $\tilde{\chi}_2^0$, respectively, being dominated by the higgsino component. The decay length of the chargino is given by

$$c\tau \simeq 0.7 \text{ cm} \times \left[\left(\frac{\Delta m}{340 \text{ MeV}} \right)^3 \sqrt{1 - \frac{m_\pi^2}{\Delta m^2}} \right]^{-1}, \quad (2.33)$$

where m_π is the charged pion mass and $\Delta m = \tilde{\chi}_1^\pm - \tilde{\chi}_1^0$ [37]. For chargino masses ranging from 150 GeV to 1000 GeV, the mass splitting ranges from approximately 280 MeV to 350 MeV, leading to charginos with $c\tau \approx 7 - 14$ mm. The exclusion limits on higgsino pair production at the LHC as a function of the chargino and neutralino masses is shown in Figure 2.7(b). For a pure higgsino scenario, chargino masses up to 210 GeV are excluded using the results of a search for short charged particle tracks in the ATLAS experiment [35]. A similar search in the CMS experiment excludes chargino masses up to 160 GeV in a pure higgsino scenario [34].

2.5.4 R-parity violating SUSY

In Section 2.4.3, R-parity conservation was introduced as a way to prevent lepton- and baryon- number violating terms in the superpotential, as shown in Equations 2.25 and 2.26. However, R-parity conservation is not a necessary requirement in supersymmetric theories. SUSY scenarios with R-parity violating (RPV) couplings can still evade experimental constraints, if only one of the λ' or λ'' couplings are non-zero, and if all non-zero RPV couplings are small.

Allowing for RPV couplings, albeit small, has a drastic consequence on the phenomenology of SUSY signatures at colliders. First, one could have single sparticle production, in addition to the usual sparticle pair production processes. Furthermore, the LSP will decay via RPV couplings into SM particles. The number and type of SM particles which it decays to depends on the RPV coupling mediating the decay. In the case of a $\tilde{\chi}_1^0$ LSP, decays via the λ coupling would result in two oppositely




Figure 2.7: In (a), the constraints on the chargino mass-vs-lifetime plane for an AMSB model with $\tan(\beta) = 5$ and $\mu > 0$ are shown. The regions below the curves are excluded. In (b), the exclusion limits of higgsino pair production as a function of the chargino and neutralino masses are shown. The shaded regions are excluded. Taken from Ref. [29].

charged and possibly different-flavor leptons and a neutrino. Decays mediated by the λ' coupling would result in a pair of jets and either a charged lepton or neutrino. Lastly, decays mediated by the baryon-number violating λ'' coupling will result in final states with multiple jets.

The LSP lifetime depends both on the RPV coupling strength and on the masses of the sfermions involved in the decay. In the case of a bino-like neutralino LSP which decays via the λ'' coupling, the mean decay length is given by

$$L(\text{cm}) \approx \frac{0.9\beta\gamma}{\lambda''^2} \left(\frac{m(\tilde{q})}{100 \text{ GeV}} \right)^4 \left(\frac{1 \text{ GeV}}{m(\tilde{\chi}_1^0)} \right)^5, \quad (2.34)$$

where $m(\tilde{q})$ is the scale of the squark masses [28]. A range of long-lived particle lifetimes can be generated with RPV couplings. In the limit that RPV couplings are infinitely small, R-parity is essentially conserved and the LSP is stable. Very small RPV couplings lead to very long LSP lifetimes, with decay lengths much larger than the dimension of the ATLAS detector. As the RPV couplings increase, the lifetime of the LSP decreases, until its decay is prompt. Increasing the RPV coupling further allows for the direct decay of non-LSP particles to SM particles via the RPV coupling.

Exclusion limits on the strength of the λ'' RPV coupling are shown in Figure 2.8. In Figure 2.8(a), the gluino decays to SM quarks and the LSP, which then decays via the λ''_{112} coupling to light-flavor Standard Model quarks. The LSP is assumed to be a bino-like neutralino with $m(\tilde{\chi}^0) = 200 \text{ GeV}$. These limits were set using the results of searches for events with a high multiplicity of energetic jets in the ATLAS experiment. Similar searches have been performed with the CMS experiment [38, 39]. Though not included in this figure, gluinos are excluded up to masses of $\approx 1.3 \text{ TeV}$ in the

high- λ''_{112} region by a dedicated search using multijet events in the ATLAS experiment [40]. There is an uncovered region where the neutralino decays at a displaced location with respect to the interaction point but still within the detector.

In Figure 2.8(b), the stop decays to a top quark and the LSP, which then decays via the λ''_{323} RPV coupling to two down-type Standard Model quarks. The LSP is assumed to be a bino-like neutralino with $m(\tilde{\chi}^0) = 200$ GeV. These limits were set using searches for events with highly energetic jets originating from heavy flavor quarks and for events with exactly two highly energetic jets in the ATLAS experiment. Similar searches have been performed in the CMS experiment [41, 42]. There is an under-covered region for values of the λ''_{323} coupling corresponding to χ_1^0 lifetimes of ~ 1 ns. This corresponds to a proper decay length within the ATLAS inner tracker ($R \lesssim 300$ mm).




Figure 2.8: In (a), the exclusion limits on the mass of the gluino as a function of the RPV coupling λ''_{112} is shown. In (b), the exclusion limits on $m(\tilde{t})$ as a function of the RPV coupling λ''_{323} is shown. The limits are calculated from the results of searches for promptly-decaying supersymmetric particles in the ATLAS detector. Taken from Ref. [28].

2.6 Long-lived particles in non-SUSY extensions of the Standard Model

Long-lived particles are present in a wide variety of non-SUSY extensions of the Standard Model as well. A comprehensive discussion of these theories can be found in Ref. [10, 43, 44]. Several examples are described in the following paragraphs.

Most simplified models of dark matter (DM) include a single DM particle which interacts weakly with the Standard Model via a mediator particle. At early times in the universe, the DM and the SM particles were in chemical equilibrium. At some point when the universe has cooled sufficiently, freeze-out occurs, where the number-changing processes cease to occur. At this point, the DM density approaches a constant value, called the relic density.

In s-channel setups, the mediator particle couples to either a pair of DM particles, or a pair of SM

particles, which always ensures a sizeable decay width and prompt decay. If, however, a t-channel approach is taken instead, then the mediator couples to one DM particle and one SM particle. This produces two avenues for a long-lived mediator: if the coupling is small or if there is a compressed mass splitting between the mediator and the DM [45].

A particularly interesting scenario is that when the coupling is small, such that the common assumption that DM and its co-annihilation partner are in chemical equilibrium during freeze-out is no longer valid. Omitting this assumption opens up a new region of parameter space where the relic density can be achieved with a small mass splitting, called *conversion-driven freeze-out*. The cosmologically viable parameter space for conversion-driven freeze-out is shown in Figure 2.9 [46]. The shaded regions are excluded by existing searches for R-hadrons at the LHC and from Big Bang nucleosynthesis. The un-covered region of parameter space corresponds to long-lived mediators with decay lengths on the order of centimeters.




Figure 2.9: The cosmologically viable parameter space in the conversion-driven freeze-out region (below solid black line) is shown as a function of the mass splitting Δm between the t-channel mediator and the dark matter particle, and the mass of the dark matter particle (m_χ). Constant mediator lifetimes are shown as gray dotted curves from 10^{-3} s to 10^3 s in steps of an order of magnitude. The excluded regions from R-hadron searches at the 8 TeV and 13 TeV LHC are shown in dark and light blue, respectively. The red shaded region is excluded from Big Bang nucleosynthesis constraints. Taken from Ref. [45].

Long-lived particles are also abundant in “hidden-sector” models [47]. This general class of models extends the SM gauge group by an additional gauge interaction, under which all SM particles are neutral. Interactions between this additional group and the SM gauge group are possible via higher dimension operators, thus rendering long lifetimes of mediators which are charged under both the SM and the hidden sector. This general model is applied in many BSM theories, including the GMSB, Twin Higgs models [48], milli-charged DM particles [49], and even string theory [50].

Experimental Setup

The Large Hadron Collider (LHC) is the world's largest and most powerful particle collider [51]. The ATLAS detector records the particles resulting from these collisions, allowing for precision tests of the Standard Model and searches for new physics beyond the Standard Model [9]. In this chapter, the setup of the LHC and the ATLAS detector is presented and the phenomenology of proton-proton collisions at the LHC is discussed.

3.1 The Large Hadron Collider

The LHC is a 27 km circular accelerator, located approximately 100 m underground near Geneva, Switzerland. Two beams consisting of bunches of protons are accelerated in opposite directions around the ring. The LHC collides the bunches of protons head-on at a rate of 40 MHz. From 2015-2018 (Run-2), the LHC accelerated each beam of protons to 6.5 TeV in order to achieve a center of mass collision energy of 13 TeV. The LHC is also capable of colliding heavy ions.

3.1.1 Accelerator basics

Several different technologies are used to accelerate, steer, and focus the proton beams at the LHC. Protons are accelerated using radiofrequency cavities, which are metallic chambers containing an electromagnetic field. The radiofrequency cavities are tuned to resonate at a specific frequency, so that each bunch of protons gain some energy as they pass through the cavity. Several types of superconducting electromagnets operated at cryogenic temperatures are used to focus and bend the proton beams. As the energy of the beams increase, it becomes more and more difficult to bend the particles in a circular trajectory. The LHC achieves this with 1,232 powerful dipole magnets, each of which are 15 meters long and produce an 8.3 T field. The beams are squeezed to a size of a few

microns in order to increase the chance of collisions occurring with the incoming beam. This is done with quadrupole magnets which are arranged symmetrically around the beam. Corrector magnets with up to 12 poles are used to stabilize the beam against small field disturbances and field irregularities of the dipole and quadrupole magnets [52].

3.1.2 The CERN accelerator complex




Figure 3.1: Schematic of the CERN accelerator complex. Linac 4 replaced Linac 2 in 2020. Modified from Ref. [53].

A series of accelerators are used to generate and accelerate the proton beams, which are shown in Figure 3.1. The protons start in a simple bottle of hydrogen gas. An electric field strips the hydrogen atoms of their electrons, yielding protons. In Run-2, the first accelerator in the chain is Linac 2, which is a linear accelerator that produces a 50 GeV proton beam. The next step is the Proton Synchrotron Booster, which is composed of four superimposed synchrotron rings. The incoming beam is split and accelerated simultaneously to 1.4 GeV in each of the rings, then recombined and injected into the Proton Synchrotron. The Proton Synchrotron and Super Proton Synchrotron are the last two steps before the LHC. Both use conventional (room-temperature) magnets to accelerate the beams up to 25 GeV and 450 GeV, respectively. Following these steps, the beams are injected into the LHC ring, in opposite directions, where they are accelerated up to their final energy of 6.5 TeV.

3.2 Phenomenology of proton-proton collisions

In this section, the phenomenology of proton-proton collisions produced at the LHC is discussed.

3.2.1 Constituents of the proton

A proton is a bound state of two up quarks and one down quark. These so-called *valence quarks* interact with each other via the exchange of gluons. Virtual quark-antiquark pairs produced from the gluons are called *sea quarks*. Therefore there are three main constituents of protons, called partons: valence quarks, sea quarks, and gluons. In a 13 TeV proton-proton collision at the LHC, it is the constituents of the protons that are colliding rather than the protons themselves.

The parton distribution function (PDF), $f_a(x, \mu^2)$, describes the probability to find a parton of type a with a momentum fraction x of the proton's total momentum when probed at energy scale μ^2 . Due to the large QCD coupling constant at proton energy scales ($\alpha_s \sim O(1)$ at interactions with momentum transfer $|q| \approx 1$ GeV), parton distribution functions cannot be calculated directly with perturbation theory. Instead they are measured from global fits using a wide range of experimental data, for example from inelastic electron-proton interactions at HERA [54].

The PDFs with $\mu^2 = 10^4$ GeV² are shown in Figure 3.2. The PDFs of valence quarks are characterized by a peak below $x = 1/3$. The radiation of gluons with large momenta is suppressed by the $1/q^2$ gluon propagator, and therefore the PDFs of gluons and sea quarks is characterized by a steeply falling spectrum in x . The total fraction of the momentum of the proton carried by quarks and antiquarks is just over 50%. The remainder is carried by the gluons.




Figure 3.2: The parton distribution functions for partons within the proton [55]. The subscript v indicates a valence quark.

The PDFs depend somewhat on the energy scale at which the proton is being probed, μ^2 . At higher μ^2 , the proton is observed to have more low- x sea quarks and gluons and a softer x spectrum for

the valence quarks. Conceptually this arises because probing the proton with higher μ allows the interaction to "see" with higher resolution the partons within the proton. Thus, the number of resolved partons which share the proton's momentum increases with μ , resulting in an increased probability to find a quark or gluon at small x , and a decreased chance of finding a valence quark with high x . The evolution of the PDFs with μ^2 is calculated in perturbative QCD using the DGLAP equations [56].

3.2.2 Parts of a proton-proton collision

The LHC collides bunches of protons. One bunch crossing is referred to as an *event*. The vast majority of proton-proton collisions at the LHC are soft, meaning they produce potentially lots of new particles with low transverse momenta. Most of the physics studied at the LHC originates from a *hard interaction*, which is an interaction between partons where heavy objects are created or a large momentum transfer occurs. The hard interaction is only one part of the activity which occurs in a proton-proton collision. The full spectrum of activity which can occur in a proton-proton collision is depicted in Figure 3.3.

The particles entering and exiting the hard process are quarks and gluons. They can radiate gluons which can then produce quark-antiquark pairs. This results in a shower of partons originating from either the final or initial partonic state of the hard process. The radiation of a particle from the incoming (outgoing) particles participating in the hard interaction is called *initial- (final-)state radiation*.

Due to color confinement, the partons in these showers do not propagate freely. As the parton showers evolve to lower momentum scales, the partons become bound within colorless hadrons - a process called hadronization. The result is a spray of colorless hadrons traveling in the same direction as the original parton - called a jet. As QCD processes are dominant in LHC collisions, there are typically many jets in the events.

Many of the hadrons produced in the hadronization process are unstable resonances. Thus after the hadronization process, there will be the decays of some hadrons into lighter particles, some promptly and others after traveling a macroscopic distance in the detector.

The underlying event encapsulates all activity in an event which does not come directly from the primary hard scattering. *Multiple parton interactions* occur when additional partons from each proton interact with each other, resulting in a soft, secondary collision and low-momentum jets. Interactions between protons in the bunch which are not involved in the hard process are called *pileup*. Pileup and multiple parton interactions make analyzing LHC data challenging, because they make detector full of activity which is not directly related to the hard interaction of interest.

3.2.3 Luminosity and cross sections

The LHC is a counting experiment; events with certain characteristics are counted and compared to the theoretical predictions from the Standard Model or BSM theories. The number of events expected




Figure 3.3: Schematic of a proton-proton collision [57]. The two incoming protons are each depicted by three horizontal dark green arrows. Each color indicates a different type of interaction, which is indicated in the legend.

for a given process i is given by

$$N_i = \sigma_i \int \mathcal{L}(t) dt, \quad (3.1)$$

where σ_i is the *cross section* of the process and $\mathcal{L}$ is the *instantaneous luminosity*. The instantaneous luminosity is a process-independent quantity that indicates the conditions of the beams which have an impact on the probability of a collision happening. For head-on collisions of two beams, the luminosity of a single bunch crossing is given by

$$\mathcal{L}_b = f \frac{n_1 n_2}{2\pi \Sigma_x \Sigma_y}, \quad (3.2)$$

where f is the bunch crossing frequency, $n_{1,2}$ are the number of particles in each colliding bunch, and $\sigma_{x,y}$ are the convolved horizontal and vertical beam sizes. The sum of $\mathcal{L}_b$ over all colliding bunch pairs gives the total instantaneous luminosity. The luminosity is measured using van der Meer scans

and luminosity-sensitive detectors [58]. The LHC achieved a peak luminosity of $21.0 \times 10^{33} \text{ cm}^{-2} \text{ s}^{-1}$ in 2018, as shown in Figure 3.4(a).

The luminosity integrated over time between 2015 and 2018 is shown in Figure 3.4(b). The luminosity which is recorded by the ATLAS detector is less than the luminosity delivered by the LHC. This is due to inefficiencies in the data acquisition system and in the time needed for the tracking detectors to prepare for data-taking after stable beams in the LHC have been reached. Data which is recorded by ATLAS when one or more of the subdetectors or magnet systems is experiencing integrity issues is not used for physics analysis. This may occur as a result of a surge in noise from the detector electronics or from a high-voltage trip in a detector module. In the full Run-2 dataset, approximately 95.6% of the data recorded by the ATLAS Detector achieved a high quality and was considered good for use in physics analyses [59].




Figure 3.4: In (a), the peak luminosity during each run is shown for several months in 2018. In (b), the overall luminosity delivered by the LHC and collected by ATLAS in Run-2 is shown [60].

The cross section σ_i shown in Eq. 3.1 is a measure of the quantum mechanical probability for the interaction i to take place. This is calculated in QCD by factorizing the cross section into a low-energy and high-energy part. The low-energy part describes the behavior of bound partons within the proton with PDFs and the high-energy part describes the hard interaction. Therefore the cross section for two protons colliding and producing final state X is given by

$$\sigma_{p_1 p_2 \rightarrow X} = \sum_{i,j \in \{q,g\}} \int dx_1 f_i(x_1, \mu^2) \int dx_2 f_j(x_2, \mu^2) \hat{\sigma}_{ij \rightarrow X}(\hat{s}, \mu^2), \quad (3.3)$$

where i and j sum over all partons in the protons, f is the PDF, $\hat{\sigma}$ is the cross section of the hard interaction with center of mass energy $\hat{s} = x_1 x_2 s$. The cross section of the hard interaction is calculated using perturbative QCD. Quarks and gluons may be treated as quasi-free particles rather than strongly-bound within the proton due to asymptotic freedom.

Even though α_s at high μ^2 is sufficiently small for perturbation theory, it is not so small that higher-order corrections can be neglected. Therefore, cross sections involving QCD are often calculated including higher-order diagrams involving loops, at the level of next to leading order (NLO) or next to next to leading order (NNLO).

3.2.4 Event simulation

Proton-proton collisions are simulated using Monte-Carlo techniques [61]. The simulated events aim to reproduce as close as possible what would be measured in the detector with a real proton-proton collision. Simulated events are essential to extract a signal of new physics from the Standard Model background. They are used to design analyses for maximal sensitivity to the physics process of interest, to understand and compensate for the effects of the detector on measurements, and to provide useful input for the design of new detectors, new triggers, and new analysis techniques.

All of the activity in proton-proton collisions described in Section 3.2.2 needs to be simulated. QCD is involved in several steps of the simulation at very different scales and not all processes are calculable using first principles. An important theoretical point which simplifies event generation is the QCD factorization property, which states that the treatment of hard and soft scales are separable. The cross section and event structure depend on the hard process while parton showers and hadronization happen at softer scales. Assuming hard and soft scales are separable, the soft-scale physics can be added to the event without changing the cross section.

The event generation is built around the hard process - one specifies the hard process of interest, and event generators produce partonic events (a list of outgoing partons and their momenta) according to the PDFs, matrix element and phase space. This is performed with a dedicated matrix-element and phase-space generator, such as MadGraph [62].

The generation of the hard process using lowest-order matrix elements describes the momenta of the outgoing jets well. However higher-order calculations are needed to accurately model the parton showers, which quickly become cumbersome using the matrix element method. Instead, the effect of all higher orders is simulated through a *parton shower algorithm*. General-purpose event generators, like Pythia [63], take the partonic events from dedicated matrix-element generators and perform the parton showering.

To be useful, event generators must replicate the particles which deposit energy in the detector. These are the hadrons which are stable on detector length scales, not the parton showers. Thus the parton showers are evolved to lower momentum scales, until it reaches a region (1 GeV) in which QCD becomes strongly interacting, and perturbation theory breaks down. It is at this scale that a hadronization model is used to describe the confinement of the system of colored partons into colorless hadrons. Hadronization models use fragmentation functions, which give the probability that a given observed final state of hadrons originates from a given quark or gluon. Generators then simulate the

decays of unstable hadrons, which is based on a combination of experimental results and theoretically motivated assumptions. Furthermore, other soft processes in the underlying event are simulated with models with tunable parameters which are constrained by data.

The passage of particles through the ATLAS detector and the detector response are simulated using the GEANT4 or ATLFAST-2 toolkits [64]. The GEANT4 simulation toolkit uses detailed detector geometry and simulation techniques to provide an accurate model of detector response as particles traverse the detectors. This detailed simulation is especially important for accurately modeling the response of the detector to unconventional particle signatures, for example those produced by massive, long-lived particles. The ATLFAST-2 simulation toolkit uses a dedicated fast tracking and calorimeters simulation algorithms to perform simulation 100 times faster than the full GEANT4 simulation. It is used to provide large samples of simulated events to supplement full simulation studies.

After this production chain, simulated events are in the same format as data collected with the ATLAS detector, with the exception that they have a ‘truth’ record stored as well - a list of the true kinematics and decay chains of all particles present in the event.

3.3 The ATLAS Detector

The ATLAS (A Toroidal LHC ApparatuS) detector is a general purpose detector at the LHC. It aims to detect particles originating from the proton-proton collisions. Detecting interesting events at the LHC poses a unique challenge, because the inelastic proton-proton collision cross section (approximately 80 mb [65]) is very large compared to rare processes. A further complication arises from the fact that quarks and gluons participate in the strong force. This results in a high jet production cross section, which also dominates over rare processes. The ATLAS detector is designed to overcome these challenges and exploit the full discovery potential of the LHC.

3.3.1 Coordinate system

ATLAS uses a right-handed coordinate system, with the origin at the nominal interaction point at the center of the detector and the beam line along the z-axis. The x-y plane is referred to as the transverse plane. The azimuthal angle ϕ is measured around the beam axis, and the polar angle θ is measured from the beam axis.

The incoming beams nominally travel in the z-direction only. The quarks and gluons within a proton carry varying fractions of the proton’s momentum. Thus, the initial conditions along the z-axis of a collision between partons is unknown. However, due to conservation of momentum, the momenta of particles in the transverse plane, $p_T = |\mathbf{p}| \sin(\theta)$, is expected to be close to 0. The missing transverse momentum is therefore an important quantity. It is defined as the negative vector sum of the transverse momenta of all reconstructed objects in the ATLAS detector.

Rapidity is a relativistic measure of longitudinal velocity, which has a desirable property that differences in rapidity are invariant under Lorentz transformations along the beam direction. Pseudorapidity is typically used in collider experiments, as it is easier to construct than rapidity and a valid estimate of rapidity in the $\beta \rightarrow 1$ limit. It is defined as

$$\eta = -\ln \left[\tan \left(\frac{\theta}{2} \right) \right].$$

Particles traveling only in the transverse plane have $\eta = 0$. Pseudorapidity is used in the following sections to indicate the placement of sub-detectors in ATLAS.

3.3.2 General design concepts

The ATLAS detector spans 25 m in height and 44 m in length. It was designed with several concepts in mind [9]. These are

1. Excellent electromagnetic calorimetry for electron and photon identification measurements.
2. Large coverage of hadron calorimetry to reconstruct jets and missing transverse momentum.
3. Efficient charged particle tracking for lepton momentum measurements and the tagging of electrons, photons, taus, and heavy flavor decays.
4. Precise measurements of muon momentum and charge.
5. Large acceptance in η coverage.
6. Triggering and measurements of particles with low p_T .

To achieve these design concepts, ATLAS is comprised of cylindrical layers of various segmented sub-detectors, each with distinct detection mechanisms. It has a high detector granularity, a large acceptance in the solid angle around the interaction point, and radiation hard electronics and subdetectors. It is situated with a superconducting solenoid magnet around the inner detector region and large superconducting air-core toroid magnets in the outer region. Figure 3.5 depicts a schematic of the ATLAS detector. The various subdetectors are labeled and are described in the following sections.

Having various sub-detectors with different particle detecting mechanisms is crucial to identifying particles and detecting a wide range of interesting physics processes. The unique signatures left by several types of particles in the ATLAS detector is shown in Figure 3.6. The inner detector tracks the trajectory of charged particles. It is immersed in a magnetic field so that the trajectory of charged particles is helical and the bending of the track allows for a measurement of the particle's momentum. The calorimeters detect particles via a destructive process. The particle's energy is converted into a detectable signal until the particle is absorbed. The calorimeters nominally absorb all particles in the




Figure 3.5: Schematic of the ATLAS detector, with labeled sub-detectors [9].

Detector component	Required resolution	η coverage	
		Measurement	Trigger
Tracking	$\sigma_{p_T}/p_T = 0.05\% p_T \oplus 1\%$	± 2.5	
EM calorimetry	$\sigma_E/E = 10\%/\sqrt{E} \oplus 0.7\%$	± 3.2	± 2.5
Hadronic calorimetry (jets)			
barrel and end-cap	$\sigma_E/E = 50\%/\sqrt{E} \oplus 3\%$	± 3.2	± 3.2
forward	$\sigma_E/E = 100\%/\sqrt{E} \oplus 10\%$	$3.1 < \eta < 4.9$	$3.1 < \eta < 4.9$
Muon spectrometer	$\sigma_{p_T}/p_T = 10\% \text{ at } p_T = 1 \text{ TeV}$	± 2.7	± 2.4

Table 3.1: General performance goals of the ATLAS detector [9]. Units for energy (E) and transverse momentum (p_T) are in GeV.

Standard Model with the exception of muons and neutrinos. Muons traverse through the calorimeters mostly unimpeded, depositing only little of their energy. Therefore, the muon spectrometer is situated outside of the calorimeters, in which the trajectory of muons in a magnetic field is measured using a tracker. Neutrinos interact only weakly and therefore travel through the ATLAS detector without leaving any detectable signal. However, the existence of neutrinos (or any weakly-interacting BSM particle) can be inferred from the E_T^{miss} .

3.3.3 Energy and momentum resolution

The design goals of the ATLAS detector systems in terms of energy and momentum resolution and coverage are listed in Table 3.1.




Figure 3.6: Schematic of the signatures of different particles in the various sub-detectors of ATLAS [66].

For a large number of track measurements (N) of a particle with high transverse momentum p_T , the resolution of the track's momentum is approximated by the Gluckstern Formula:

$$\frac{\sigma_{p_T}}{p_T} \approx \sqrt{\frac{720}{N+4}} \left(\frac{\sigma_x p_T}{0.3BL^2} \right), \quad (3.4)$$

where σ_x is the intrinsic resolution of each measurement, B is the strength of the magnetic field, and L is the distance between the first and last track measurement in the bending plane. It is beneficial to build a large tracking detector with a strong magnetic field which allows for many precise track measurements in order to have the highest momentum resolution. The resolution decreases for particles with higher momenta since the curvature of their trajectory through the tracker will be smaller. At lower momenta, an additional non-negligible term arises from the multiple scattering (small angle deflections due to interaction with material detector) of the particle as it traverses the tracker. For this reason, the amount of material in the tracker is minimized.

In contrast to momentum measurements using magnetic deflection, the energy resolution of

calorimeters improves with energy. The energy resolution of a calorimeter is given by

$$\frac{\sigma}{E} = \frac{a}{\sqrt{E}} \oplus \frac{b}{E} \oplus c, \quad (3.5)$$

where E is the incident energy of the particle. The first term parameterizes the resolution from intrinsic fluctuations in the shower development. The second term parameterizes the noise, i.e. from the readout chain, and depends on the type of detection mechanism used in the calorimeter. This source can be reduced with calibration and is typically negligible at high energies. The constant term c includes any contributions which do not depend on energy, i.e. response non-uniformity due to the detector geometry.

For sampling calorimeters, in which active material layers are alternated with absorber material, the dominant source of uncertainty in the energy measurement arises from fluctuations in the number of charged particles traversing each active layer. The energy resolution of sampling calorimeters increases with thinner absorber layers, a higher number of sampling layers, and more particles detected in each sampling layer.

The calorimeter resolution also degrades with increased leakage, i.e. showers that are not contained within the detector. The resolution also degrades when particles lose energy before reaching the calorimeter. For these reasons it is beneficial to build a large calorimeter in which there is minimal material situated between the particle's production point and the start of the calorimeter.

3.3.4 Magnet system

The ATLAS magnet system is important because it bends the trajectory of charged particles as they pass through the detector, allowing for a measurement of their momentum. ATLAS utilizes four large superconducting magnets spanning an area of 22 m in diameter and 26 m in length. A schematic of the magnet system is shown in Figure 3.7. The central solenoid is situated in front of the calorimeter system and provides a central field of 2T in the inner detector. The barrel toroid is situated around the calorimeter volume and provides a central field of 0.5 T. Two end-cap toroids are inserted into the barrel toroid at each end and line up with the central solenoid, providing a central field of 1.0 T. The toroids provide bending power in the muon spectrometer.




Figure 3.7: A schematic of the magnet system in ATLAS [9]. The barrel and endcap toroid coils are shown in orange. The solenoid, situated within the calorimeter volume, is not visible in this schematic.

3.3.5 Inner detector

The inner detector is situated closest to the interaction point and tracks the trajectory of charged particles. A schematic of the inner detector is shown in Figure 3.8. The inner detector uses two techniques to identify particles: silicon semiconductor tracking and a straw tube tracker. It consists of three parts: The Pixel detector, the Semiconductor tracker (SCT), and the Transition Radiation Tracker (TRT). The precision trackers (pixel and SCT) cover the region $|\eta| < 2.5$ while the TRT covers the region $|\eta| < 2.0$.

The Pixel detector is located closest to the interaction point and spans radii of 33 to 123 mm. It contains four concentric cylindrical layers of silicon pixels, with three disks of silicon pixels situated perpendicular to the beam axis at each end. The pixel layers are segmented in ϕ and z (R) in the barrel (endcap). When a charged particle passes through the silicon, charge is generated from the formation of electron-hole pairs. The length of time during which the signal generated from the drifting charges is above a predefined threshold (Time-over-Threshold) is recorded for further analysis. The Time-over-Threshold is proportional to the amount of charge deposited in the pixel. A calibration procedure is performed to infer the deposited charge from the recorded Time-over-Threshold.

The pixel detector has a high granularity, resulting in a large number of readout channels (~ 80.4 million) and very high resolution tracking. The three outermost pixel layers achieve a hit resolution of about $10 \mu\text{m}$ in $R\phi$ and $115 \mu\text{m}$ in z (R) in the barrel (endcap). The innermost barrel layer of the Pixel detector, called the insertable B-layer (IBL) [68], has an intrinsic hit resolution of $10 \mu\text{m}$ in $R\phi$ and $75 \mu\text{m}$ in z . The IBL is situated close to the interaction point to allow for a precise determination of the point of closest approach of each track to the interaction point. This is important




Figure 3.8: A schematic of the ATLAS Inner Detector [67].

to resolve tracks and vertices which are displaced from the primary interaction, for example those which arise from B-hadrons and tau leptons.

The Semiconductor tracker (SCT) surrounds the pixel detector, spanning radii of 299 mm to 514 mm. The SCT works according to the same physics principles as the pixel detector, but with strips of silicon instead of pixels. The SCT utilizes binary read out, which means that only the occurrence of a hit is recorded rather than the Time-over-Threshold. While using strips reduces the tracking resolution, they allow the SCT to track particles in a much larger volume in a cost-affordable way. The SCT contains four cylindrical layers of silicon strips and nine disks of silicon strips at each end. Each cylindrical layer contains two sub-layers of strips. One set is parallel to the beam direction, providing an $R\phi$ measurement. The other set is situated at a small angle (40 mrad) with respect to the first set, allowing for a measurement of the z coordinate as well. The strips in the end cap are arranged in a similar way, with one set situated radially and the other offset by a 40 mrad angle. The intrinsic resolution is $17 \mu\text{m}$ in $R\phi$ and $580 \mu\text{m}$ in z (R) in the barrel (endcap).

The outermost layer of the ID is the Transition Radiation Tracker (TRT). The TRT spans radii of 554 to 1082 mm from the interaction point and is situated with both barrel and end-cap layers. It allows for tracking at larger radii from the interaction point in the $|\eta| < 2.0$ region in a cost-effective

way and with much less material disturbing the particles as they pass through. It is composed of about 50,000 gas-filled straw tubes which are 4 mm in diameter, each surrounded by material which induces transition radiation. The transition radiation is emitted when a charge particle passes through inhomogeneous material, and is enhanced when the charged particle is ultrarelativistic. Thus, high energy electrons passing through the TRT leave a strong signal in the straw tubes, distinguishing them from other charged particles. The readout has two independent thresholds, allowing for the distinction between tracking hits (passing the lower threshold) and transition radiation hits (passing the higher threshold).

Each straw in the TRT provides a drift-time measurement, allowing for a measuring of $R\phi$ with an intrinsic resolution of 170 μm . It does not provide a z measurement. Though less precise than the precision silicon trackers, the TRT contributes significantly to the momentum measurement of charged particles because of the large number of hits at larger radii. Charged particles travelling through the ID will typically leave 4 hits in the pixel detector, 8 hits in the SCT, and 36 hits in the TRT.

3.3.6 Calorimeter

Following the ID are the electromagnetic and hadron calorimeters, which measure the energy of particles and jets. Calorimetry is powerful because it can detect both charged and neutral particles, its performance improves with energy, and signals are fast, so can be used for triggering. ATLAS utilizes sampling calorimeters, where layers of absorber and active materials are alternated. The absorber material initiates particles showers and the active material allows for the shower to be measured. The various parts of the ATLAS calorimeter system are listed in Table 3.2 and shown in Figure 3.9.




Figure 3.9: A schematic of the ATLAS calorimeters [9].

Calorimeter	Material	Pseudorapidity coverage	Thickness
LAr electromagnetic barrel	lead-LAr	$ \eta < 1.475$	$24 X_0$
LAr electromagnetic end-cap	lead-LAr	$1.374 < \eta < 3.2$	$26 X_0$
Tile barrel	iron-scintillator	$ \eta < 1.0$	11λ
Tile extended barrel	iron-scintillator	$0.8 < \eta < 1.7$	11λ
LAr hadronic end-cap	copper-LAr	$1.5 < \eta < 3.2$	12λ
LAr forward	copper/tungsten-LAr	$3.1 < \eta < 4.9$	9λ

Table 3.2: The various parts of the ATLAS calorimeter system. The materials listed indicate those used in the absorber-active layers. The thickness of each calorimeter part includes the material which is situated in between the calorimeter and the interaction point.

Electromagnetic calorimeter

Immediately following the solenoid surrounding the inner detector is the electromagnetic calorimeter (ECAL). It is split into a barrel and two end-cap components. The granularity of the ECAL ranges from $\Delta\eta = 0.003 - 0.1$ and $\Delta\phi = 0.025 - 0.1$. The ECAL has the highest granularity in the η region covered by the inner detector in order to provide precision measurements of electrons and photons.

The ECAL is an accordion-shaped sampling calorimeter with alternating layers of liquid argon (LAr) and lead. The lead material induces electromagnetic showers, where electrons and photons undergo bremsstrahlung and pair production, respectively. The shower stops when photons no longer have enough energy to undergo pair production, and other forms of energy loss besides bremsstrahlung (ionization / excitation) begin to dominate for electrons. The showers are sampled with the layers of liquid argon, where charged particles passing through create a signal via ionization. In electromagnetic calorimeters, the shower development is governed by a radiation length (X_0), which is the distance traveled through material after which one electron has only $(1/e)$ (37%) of its incident energy.

$$E = E_0 e^{-x/X_0} \quad (3.6)$$

In the barrel region, the ECAL extends from 1.15 m to 2.25 m ($24 X_0$). This thickness is chosen so that most electromagnetic showers are fully contained within the ECAL.

The ECAL is preceded by a presampler detector in the $|\eta| < 1.8$ region, which is used to correct for the energy lost in the material upstream of the calorimeter (inner detector, cryostat walls, solenoid coil). The presampler is a thin LAr layer of thickness 1.1 cm (0.5 cm) in the barrel (endcap) region.

Hadron calorimeter

Following the ECAL is the hadronic calorimeter (HCAL). The HCAL aims to measure the energy of hadrons. It is split into a tile calorimeter and two end-cap calorimeters. The tile calorimeter contains steel and scintillating tiles while the end-cap calorimeters utilize copper and liquid argon.

The granularity of the HCAL ranges from $\Delta\eta = 0.1 - 0.2$ and $\Delta\phi = 0.1 - 0.2$.

Particle showers are produced in the HCAL, similar to the ECAL. However, due to the strong interaction with detector material, these showers are more complex than electromagnetic showers. Hadrons that pass through the HCAL interact with the absorber nuclei, resulting in intranuclear cascades, spallation, nuclei deexcitations, and neutron capture. Energy deposited in the calorimeter via these mechanisms is referred to as hadronic energy. Some components of hadronic energy are ‘invisible’, meaning their signals cannot be detected at all, or only with a significantly reduced efficiency. Sources of invisible energy include delayed photons, soft neutrons, and the binding energy of nucleons released from nuclei. The invisible energy may represent up to 40% of the total non-em energy and can have large event-to-event fluctuations [69].

Particles which decay electromagnetically, i.e. π^0 , η , are also produced in hadronic showers. These particles will initiate electromagnetic showers, thereby transferring energy from the hadronic part of the shower to an electromagnetic part. The average fraction of the total shower energy which is transferred to the electromagnetic part increases non-linearly with the energy.

The *response* of a calorimeter is defined as the conversion efficiency from deposited energy to generated signal. Due in large part to the invisible component of the hadronic energy, the response of a calorimeter to the electromagnetic and hadronic components of a shower, e and h , are usually not the same. A calorimeter with a value of e/h different from 1, called a *non-compensating* calorimeter, will have deviation from linearity in the hadronic response versus energy, resulting in a worse energy resolution and non-gaussian tails in the calorimeter response. The value of the e/h ratio depends on a number of factors such as the thickness of the passive layers, the thickness of the active layers, and the sampling fraction [70]. The (e/h) of the tile calorimeter, for example, is 1.36 ± 0.11 [71]. Event-by-event fluctuations in the electromagnetic fraction of the hadronic shower therefore contribute significantly to the energy resolution of the calorimeter. A careful calibration procedure of the energy of particles measured in the calorimeter is needed, in order to set an accurate overall energy scale and to obtain as uniform of a response as possible.

Hadronic shower development is dictated by the interaction length (λ), defined as the mean free path of a particle between two nuclear collisions in the material. In terms of cm, λ is typically much larger than X_0 . The HCAL of ATLAS is therefore larger than the ECAL, spanning 2.28 m to 4.25 meters in the barrel region. This thickness is sufficient to keep the amount of hadronic showers which are not fully contained within the HCAL below an acceptable level.

The Forward Calorimeter (FCAL) is a hadronic calorimeter situated in the very forward direction with respect to the beam direction. It allows for the energy measurement of forward jets and helps to shield the muon spectrometer from contamination from other particles. The FCAL consists of metal (copper and tungsten) matrices with regularly spaced longitudinal channels filled with concentric rods and tubes, with a gap filled with LAr in between. The granularity of the FCAL endcap ranges from $\Delta\eta = 0.2$ and $\Delta\phi = 0.2$.

Sub-detector	Purpose	Location	$ \eta $ coverage	$\sigma(z/R)$	$\sigma(\phi)$	timing
MDT	Tracking	barrel, endcap	$ \eta < 2.7$	$35 \mu\text{m} (z)$	-	-
CSC	Tracking	endcap layer 1	$2.0 < \eta < 2.7$	$40 \mu\text{m} (R)$	5 mm	7 ns
RPC	Trigger	barrel	$ \eta < 1.0$	10 mm (z)	10 mm	1.5 ns
TGC	Trigger	endcap	$1.0 < \eta < 2.4$	2-6 mm	3-7 mm	4 ns

Table 3.3: The purpose, location, and resolution of various sub-detectors within the ATLAS Muon spectrometer.

3.3.7 Muon Spectrometer

The muon spectrometer (MS) occupies the outermost layer of ATLAS and measures the momentum of muons. It provides tracking and triggering of muons in the $|\eta| < 2.7$ and $|\eta| < 2.4$ range, respectively. A schematic of the MS is depicted in Figure 3.10. The MS consists of large, superconducting air-core




Figure 3.10: A schematic of the ATLAS Muon Spectrometer [9].

toroid magnets and four sub-detector systems, which are listed in Table 3.3 and shown in Figure 3.11. Monitored Drift Tubes (MDT) and Cathode Strip Chambers (CSC) provide precision tracking while Resistive Plate Chambers (RPC) and Thin Gap Chambers (TGC) utilize fast timing to provide trigger capabilities.

Precision tracking chambers

The monitored drift-tube chambers (MDTs) provide a precision measurement of the track coordinates over most of the $|\eta|$ range. They consist of aluminum tubes of 30 mm diameters with a central wire, with lengths ranging from 70 cm to 630 cm. Precision tracking chambers are arranged in three concentric cylindrical shells around the beam axis at radii of approximately 5 m, 7.5 m, and 10 m. In




Figure 3.11: A schematic of the various muon sub-detectors [9]. The MDTs (not labeled) are shown in green and turquoise. The HCAL is shown in red.

the endcap, they are arranged as large wheels, located at distances ranging from $|z| = 7.4 - 21.5$ m from the interaction point. Each chamber consists of 3-8 layers of drift tubes.

The background rate is highest at large pseudorapidity and close to the interaction point. Therefore, MDTs are replaced by CSCs in the $|\eta| > 2$ region of the first endcap layer. The CSCs consist of multiwire proportional chambers with cathode strips situated in orthogonal directions. Compared to the MDTs, the CSCs have higher granularity which allows them to withstand the higher flux of particles in the high $|\eta|$ region. The CSCs are segmented into eight chambers in ϕ , each chamber containing four CSC planes, resulting in four independent measurements of η and ϕ for each track.

Trigger chambers

The trigger chambers provide fast information on muon tracks which is used to quickly decide if the event should be kept for further analysis. They are designed to measure the muon p_T , provide bunch-crossing identification, and perform tracking to be used in the trigger. Triggering on muons is uniquely challenging in the endcap region. To achieve the same p_T resolution as in the barrel, the endcap must be constructed with higher granularity. Furthermore, the endcap region has much higher levels of background compared to the barrel. For these reasons, two different technologies are used for triggering in the barrel and endcap regions.

In the barrel, RPCs are used to provide good spatial and time resolution. The RPCs consist of a narrow gas gap formed by two parallel resistive plates (2 mm), separated by insulating spacers (2 mm).

Charged particles passing through the gap produce ionization and the primary ionisation electrons are multiplied into avalanches with a high, uniform electric field of around about 4.9 kV/mm. The RPCs are situated in three planes, located near the middle and outer MDT barrel layers. Muons passing through all three stations will produce a track with 6 measurements in η and ϕ .

In the endcap region, TGCs are used. The TGCs are multiwire proportional chambers with a wire-to-cathode distance which is smaller than the wire-to-wire distance. This configuration allows for a good time resolution for charged particles passing through. They are situated in four planes, three of which are located close to the middle MDT layer, while the fourth is located near the innermost MDT layer. In addition to providing fast trigger signal, they are also used to complement the track measurements of the MDT in the radial direction.

3.4 Trigger

The trigger system is tasked with selecting events to be recorded for offline analysis. The data acquisition system is not able to read out all proton-proton collision events at a rate of 40 MHz. Furthermore, the majority of these events do not contain a hard scatter, which is of interest to most ATLAS analyses. Therefore the trigger system aims to choose only the most interesting events to be read offline for analysis. In Run-2, ATLAS used a two level trigger system: a hardware Level-1 trigger (L1) and a software-based high level trigger (HLT). The L1 trigger uses information from the calorimeters and triggering chambers in the muon spectrometer to reduce the event rate to 100 kHz. If an event passes the L1 trigger, further information from all sub-detectors of ATLAS is read out and the HLT decides if the event should be kept for offline analysis. The Level-1 and HLT triggers reduce the event rate from 40 MHz to approximately 1 kHz, on a timescale of 2.5 μ s and 200 ms, respectively [72].

3.4.1 Level-1 trigger

The L1 trigger is the first system which analyzes each proton-proton collision. It is a hardware-based trigger which uses reduced granularity information from the calorimeter and muon detectors to decide on each event. The L1 trigger selects events using event-level quantities (i.e. MET), the multiplicity of objects above certain p_T thresholds, and topological considerations (i.e. angular distances between objects). The L1 trigger also identifies *Regions of Interest* (RoI) which are investigated further in the HLT. Signals from the inner detector are not used in the L1 trigger decision, because tracking is too expensive in terms of processing time per event and computing cores needed.

The L1 Calorimeter trigger (L1Calo) [73] takes as input the signals in the electromagnetic and hadronic calorimeters. Rather than using the full granularity of the calorimeter system, the L1Calo uses coarser regions called *trigger towers*. The size of the trigger towers are $\Delta\eta \times \Delta\phi = 0.1 \times 0.1$

in the central region and vary slightly in other areas of the detector. The L1Calo system identifies electron, photon, and tau-lepton, and jet candidates and produces global sums of total and missing transverse momentum using the calorimeter signals in trigger towers.

The L1 Muon trigger (L1Muon) [74] processes input from the RPCs and TGCs in the muon spectrometer. Muon candidates are identified by coincident signals on different planes which are consistent with the trajectory of a muon from the interaction point. The typical dimensions of RoIs around L1Muon-triggered objects is $\Delta\eta \times \Delta\phi = 0.1 \times 0.1$ (0.3×0.3) in the RPCs (TGCs).

The L1 Topological trigger [75] analyzes the angular and energy information of each trigger object in the event. It receives as input each trigger object identified by the L1Calo and L1Muon systems and forms trigger decisions based on the topological angular and kinematic selections. This topological information is critical to keeping low p_T trigger thresholds, by improving background rejection with minimal signal loss.

3.4.2 High level trigger

If an event is accepted by the L1 trigger, data is read out from the remaining detector subsystems and sent to the HLT. The HLT uses finer-granularity calorimeter information and precision measurements from the muon spectrometer. In addition, tracking information from the inner detector is integrated into the trigger at this step.

Tracking is performed within RoIs identified by the L1 trigger. Tracking in the HLT takes an average 2.5 ms per RoI, with a tail at the per-mille level at around 60 ms [76]. Tracking for hadronic tau and b-jet triggers occur in large RoIs. Consequently, tracking in these triggers takes longer, on average 4.8 ms per RoI. Tracking in the HLT is over 99% efficient for electrons with $p_T > 20$ GeV and muons with $p_T > 6$ GeV. For a small subset of events, tracking over the full detector may be performed, but this is limited due to CPU requirements.

To control the rate of accepted events, a prescale n can be applied, for which each event passing that trigger has a probability of $1/n$ to be accepted.

Event reconstruction

The search for displaced vertices in multijet events makes use of both standard and special reconstruction algorithms. The standard reconstruction algorithms are optimized for reconstructing *prompt* objects, which are those that originate from or near the interaction point. These algorithms are not designed to reconstruct the signatures of particles which originate from a *displaced* ($R \gtrsim 4$ mm) location in the detector. The importance of displaced signatures for searches for long-lived particles has motivated the development of dedicated reconstruction algorithms for displaced objects. This search utilizes two of these algorithms, which are designed to reconstruct the tracks of charged particles which do not necessarily point back to the interaction point, and to find secondary vertices using those tracks which are significantly displaced from the interaction point. With these algorithms, the ATLAS experiment is able to provide competitive results when searching for long-lived SUSY particles with displaced vertices. In this section, both standard and special reconstruction algorithms are described.

4.1 Standard tracks

A track describes the trajectory of a charged particle as it passes through the detector. Standard tracks are built from the location of charge deposits in the silicon tracking detectors and in the transition radiation tracker. Tracks are described by five parameters and a reference point, using a perigee representation, as listed in Table 4.1 and shown in Figure 4.1.

The *standard tracking algorithm* [77] aims to reconstruct the tracks of charged particles which originate from or near the interaction point. The algorithm begins by clustering together the charge deposited on neighboring pixels and strip sensors. The clusters are converted to space-points, which represent the three-dimensional point where the charged particle traversed through the detector layer. Each pixel hit equates to one space-point, while two hits on each sub-layer of strips are combined to




Figure 4.1: Schematic of the track parameters with a perigee representation.

Parameter	Description
q/p	the ratio of the charge of the reconstructed track divided by the magnitude of its momentum
ϕ	the azimuthal angle of the track's momentum in the transverse plane, measured from $[-\pi, \pi]$
θ	the polar angle of the track's momentum in the longitudinal plane, measured from $[0, 2\pi]$
d_0	the transverse impact parameter, defined as the point of closest approach of the track to the reference point in the transverse plane
z_0	the longitudinal impact parameter, defined analogously to d_0

Table 4.1: A description of the five track helix parameters.

create a space-point in the SCT.

Tracks are seeded from sets of three space-points. Track reconstruction first follows an inside-out approach, in which track seeds formed from space points in the Pixel and SCT layers which pass p_T and impact parameter requirements are used to build track candidates. A combinatorial Kalman filter [78] identifies space-points which are compatible with the track seed and adds them to the track. Tracks reconstructed in the inside-out step are called *Si-seeded tracks*.

In this process, every combination of realistic space-points is made, producing many track candidates. Clusters may be shared between multiple tracks. This is expected if two particles pass through the same detector element close together, so that the charge deposited by each are merged into the same cluster. A neural network uses quantities like the measured charge and relative position of pixels in the

cluster to identify merged clusters. The ambiguity solving step is then performed to limit the number of clusters which are shared between tracks, but which are not identified as merged.

The ambiguity solver begins by ordering tracks according to their quality. The quality of the track depends on several quantities, for example it is smaller for tracks which are missing an expected hit along its trajectory (called a *hole*) and larger for tracks with a good χ^2 from the fit. The ambiguity solver then assigns shared clusters to tracks, giving priority to tracks of higher quality. After the ambiguity solver, all tracks have at most one shared hit, and all clusters are shared by no more than two tracks.

After the ambiguity solver, track candidates are extrapolated to the TRT. The TRT timing information is turned into calibrated drift circles and this information is used to assign TRT hits to the track extrapolation. TRT hits which are compatible with the extrapolation are assigned to the track candidate [79]. This completes the inside-out tracking step.

It is possible that charged particles originate far from the interaction point, but have small impact parameters, for example from the decay of K_S^0 . In this case, the charged particle would not leave many hits in the pixel and SCT layers, but only in the TRT layers. Therefore an outside-in tracking pass is performed, where TRT track segments are extended back to the silicon layers, and any compatible hits in the silicon layers which have not already been associated to a track candidate are added. Tracks reconstructed in this pass are called *TRT-seeded tracks*. The final collection of tracks from both the inside-out and outside-in tracking steps are called *standard tracks*.

The efficiency of reconstructing tracks from isolated, prompt particles (i.e. muons) is greater than 99% [77]. The reconstruction of tracks in regions with high charged particle density (i.e. near a jet) is more challenging, as particles are more likely to produce hits which are close together or overlapping. The reconstruction efficiency of prompt charged particles in simulated dijet events is shown in Figure 4.2. The efficiency is calculated as a function of the angular distance to a jet, where the angular distance is defined as $\Delta R = \sqrt{(\Delta\eta)^2 + (\Delta\phi)^2}$. The efficiency reduces at smaller angular distances to the jet axis, and with higher jet p_T , which is due to the higher density of charged particles in the center of high- p_T jets.

4.2 Large Radius tracks

Tracks from the decay products of long-lived particles may have large impact parameters ($d_0 \gtrsim 10$ mm) with respect to the interaction point, which results in inefficiencies in the standard tracking algorithm. This is demonstrated in Figure 4.3, where the reconstruction efficiency of the standard tracking algorithm reduces significantly if the long-lived particle decays beyond a radius of 10 mm.

To increase the efficiency of reconstructing displaced tracks, a dedicated algorithm known as *Large Radius Tracking* is used [80]. Large Radius Tracking is run after the standard tracking algorithm. It follows a similar method as the inside-out tracking, with loosened selection criteria on the tracks.




Figure 4.2: The efficiency to reconstruct prompt particles in jets with $|\eta| < 1.2$ is shown as a function of the angular distance of the particle from the jet axis for various p_T . Taken from [77].




Figure 4.3: Track reconstruction efficiency for displaced charged hadrons produced in the decay of long-lived particles. Taken from Ref. [80].

In particular, the Large Radius Tracking algorithm loosens the impact parameter requirement from $|d_0| < 10$ mm to $|d_0| < 300$ mm and $|z_0| < 250$ mm to $|z_0| < 1500$ mm. The selection criteria on the track p_T and hit requirements are also adjusted, as listed in Table 4.2.

Large Radius Tracking begins by identifying spacepoints from hits which have not been assigned to a standard track. It follows a similar procedure as the inside-out standard tracking approach, by first creating track seeds from the space points, from which track candidates are built. Large Radius Tracking uses a sequential Kalman filter rather than a combinatorial one (as in Standard tracking) to extend the track seeds into track candidates. This prevents multiple track candidates from being created from different possible combinations of points.

The standard and large radius track collections are combined to form a final track collection. The impact of including Large Radius Tracking on the track reconstruction efficiency from long-lived particle decays is shown in Figure 4.3. The combined track collection is able to retain a much higher tracking efficiency than the standard track collection alone for charged particles originating at $R > 10$ mm.

Selection	Standard Tracking	Large Radius Tracking
Maximum $ d_0 $ (mm)	10	300
Maximum $ z_0 $ (mm)	250	1500
Maximum $ \eta $	2.7	5
Maximum shared silicon modules	1	2
Minimum unshared silicon hits	6	5
Minimum silicon hits	7	7
Seed extension	Combinatorial	Sequential

Table 4.2: Selection applied to tracks in the different tracking algorithms. Modules are the individual silicon sensors. A shared pixel module has at least one shared pixel hit and an SCT module is shared if it has at least two shared SCT hits.

The accuracy of track parameters reconstructed with the large radius tracking algorithm is measured in a sample of simulated events. The simulated events contain pair-produced long-lived R-hadrons with $m = 1.4$ TeV which decay into two Standard Model quarks and a neutralino with $m = 1.3$ TeV. These events produce displaced hadrons with a relatively soft p_T spectrum ranging from a few GeV to a few tens of GeV. The resolutions of the d_0 and the q/p track parameters, shown in Figure 4.4, are measured as the difference of the reconstructed parameter of the tracks relative to the signal particles in the truth record to which they are matched. Large radius tracks have a worse resolution for both the d_0 and q/p parameters when compared to standard tracks. This is expected as the large radius tracks are originating from particles produced at larger radii compared to standard tracks and therefore the extrapolation of the track parameters from the first hit on the track trajectory to the interaction point will result in larger uncertainties.

4.3 Primary vertices

Primary vertex (PV) reconstruction aims to identify the locations of proton-proton interactions and the tracks which originate from them. The reconstruction of primary vertices is important to identify the location of the hard scatter interaction and the position of the beamspot.

The PV reconstruction [81] takes as input the reconstructed standard tracks which satisfy several stringent requirements, including selection on the number of silicon hits (N_{si}) and the number of holes in the pixel or SCT layers (N^{hole}). These selection are applied to reduce contamination from poorly




Figure 4.4: The residual of the (a) d_0 and (b) q/p parameters for standard and large radius tracks. Taken from [80].

measured tracks or tracks originating from secondary interactions. They are

- $p_T > 500 \text{ MeV}$, $|\eta| < 2.5$
- $N_{\text{si}} \geq 9$ if $|\eta| \leq 1.65$, $N_{\text{si}} \geq 11$ if $|\eta| > 1.65$
- At least one hit on the IBL or B-layer
- No more than 1 shared module (1 shared pixel hit or 2 shared SCT hits)
- $N_{\text{pix}}^{\text{hole}} = 0$, $N_{\text{SCT}}^{\text{hole}} \leq 1$

Tracks satisfying the above selection criteria are used to find a seed position for the primary vertex. The seed position of the vertex fit in the transverse plane is determined from the position of the beamspot in the transverse plane. In the longitudinal plane, the seed vertex position is chosen as the mode of the z-coordinates of tracks at their points of closest approach to the center of the luminous region.

The tracks and seed vertex position are used as input to a vertex fit. The vertex fitting proceeds iteratively, each time down-weighting tracks which are less compatible with the updated vertex position. The *beamspot constraint* removes tracks which are outside of the beamspot from the fitting procedure. This constraint improves the transverse position resolution of vertices reconstructed with few tracks. Tracks with sufficiently low compatibility (7σ) with that vertex are removed. This procedure is repeated, using as input the remaining tracks in the event which have not yet been associated to a primary vertex.

All vertices with at least two associated tracks are considered as primary vertex candidates. The primary vertex (PV) is defined as the collision vertex with the largest $\sum p_T^2$ of its associated tracks [81, 82]. The remaining reconstructed collision vertices are referred to as pileup.

The seed vertex position and beamspot constraint require knowledge of the beamspot position and size. The beamspot is determined using a subset of data, collected roughly every 10 minutes throughout the data-taking period. In each event, the primary vertex with the highest $\sum p_T^2$ is identified using the primary vertex reconstruction technique, without a beamspot constraint. Primary vertices with at least 5 tracks and good fit quality are selected. These requirements ensure that the PVs have a resolution in the transverse plane which is comparable to the beamspot size. These PVs are fit to determine the centroid position of the luminous region, the tilt angles in the x-z and y-z planes, and the width of the beamspot in all directions. In Run-2, the luminous region ranged from 7-14 μm in the transverse planes and 30-50 mm in the longitudinal direction.

4.4 Secondary Vertices

A dedicated algorithm is used to reconstruct vertices from the decays of long-lived particles which are significantly displaced from the interaction point [83]. The algorithm takes as input all standard and large radius tracks which satisfy $p_T > 1$ GeV. The first step of the algorithm is two-track seed vertex finding, where all pairs of tracks which are roughly compatible with the decay of a long-lived particle are identified. Due to the high multiplicity of tracks in LHC collisions, a preselection is applied to the tracks which are considered for two-track seed vertex finding. These so-called *core tracks* must not be associated to any primary vertex in the event. Furthermore, they must satisfy several selection based on the number of hits in the pixel (N_{PIX}), SCT (N_{SCT}), and TRT (N_{TRT}) subdetectors. These requirements are:

- If $N_{\text{PIX}} = 0$, then the track must have $N_{\text{SCT}} \geq 6$
- If $N_{\text{PIX}} < 2$, then the track must have $N_{\text{TRT}} \geq 1$
- If $p_T < 25$ GeV, then the track must have $N_{\text{SCT}} \geq 7$
- If $p_T < 25$ GeV and $|\eta| < 1.7$, then $N_{\text{TRT}} \geq 20$

All possible pairs of tracks satisfying the above conditions are used as inputs to a vertex fit. The vertex fitting takes as input a list of tracks and a best guess vertex position. It aims to estimate the position of the interaction point, determine the momentum vectors of the outgoing tracks at this point, and determine their goodness of fit to that position vertex. The fitting is based on the Kalman filtering method [84].

Pairs of tracks are considered *compatible* if the vertex fit has $\chi^2/N_{\text{dof}} < 5$, the position of the fitted vertex is within the SCT/TRT boundary ($R_{xy} < 563$ mm), and both tracks satisfy hit pattern requirements. The hit pattern requirements, depicted in Figure 4.5, ensure that both tracks are consistent with having come from a particle originating from the vertex position. Tracks are required

to have hits in the next outermost tracking layer with respect to the vertex position, and not have hits on innermost tracking layers. This requirement is not enforced for tracks included in vertices which are reconstructed close to tracking layers, to avoid inefficiencies due to a mis-measured vertex position.




Figure 4.5: Hit pattern requirements placed on tracks in two-track seed vertices. Taken from Ref. [83].

The next step of the algorithm is to form N-track vertices. This step focuses on *incompatible* pairs of tracks which have failed one or more of the above compatibility selection. For each event, an *incompatibility graph* is built with nodes representing tracks and edges connecting incompatible nodes. Groups of nodes which are fully compatible with each other are extracted by removing irrelevant nodes from the incompatibility graph, until all remaining nodes are completely isolated (without edges). The group of nodes (tracks) remaining are fully compatible with each other and they are fit into a single N-track vertex.

After all combinations of fully compatible groups of tracks have been formed into N-track vertices, an ambiguity solver assigns tracks which are included in multiple N-track vertices to a single vertex. The goodness of association of the track to all of its shared vertices is compared, and the vertex with the worst goodness (χ^2 per track) is identified. If $\chi^2 < 6$, then that track is removed from that DV, and the sequence is repeated. Otherwise, the two closest vertices that share that same track are merged.

The next step is to merge N-track vertices which are compatible with each other. This is important to ensure that a single, massive long-lived particle decay with many charged decay products is not incorrectly reconstructed as multiple vertices with lower masses and track multiplicities. DVs which have a distance significance $S < 10 \sigma$ are considered for merging. The distance significance is calculated as

$$S = \sqrt{(\vec{v}_1 - \vec{v}_2)(C_1 + C_2)^{-1}(\vec{v}_1 - \vec{v}_2)^T}, \quad (4.1)$$

where $\vec{v}_1$ and $\vec{v}_2$ are the displaced vertex 3D positions and the quantities C_1 and C_2 are the covariance matrices of the DVs. Pairs of DVs with $S < 10 \sigma$ are merged if any of the following conditions are

satisfied:

1. *Re-assembling* - Vertices with a larger number of tracks have an improved position resolution compared to those with a smaller number of tracks. Therefore, constituent tracks of the lower N-track DV are extrapolated to the higher N-track DV. If all of the constituent tracks are pointing to the high N-track DV, then the DVs are merged.
2. *Suggested refitting* - The lower N-track DV is refit using the position of the higher N-track DV as the starting position. If $S < 4 \sigma$ after the refit, then they are merged.
3. *Magnet merging* - Each track from the higher N-track DV is iteratively associated to the lower N-track DV, and the lower N-track DV is refit. If $S < 4 \sigma$ after any of these refits, then the vertices are merged.
4. *Wild merging* - All tracks from both DVs undergo a vertex fit. If the refit DV position is within 4σ of the original position of the higher N-track DV, then the DVs are merged.
5. *Final merging* - All pairs of DVs within 1 mm of each other are merged.

The displaced vertex reconstruction efficiency is measured in samples of simulated events from several extensions of the Standard Model predicting the existence of massive long-lived particles. The long-lived particle decays in the truth record are matched to displaced vertices by comparing the p_T -weighted fraction of the reconstructed tracks included in the vertex which are matched to long-lived particle decay products in the truth record. The efficiency is calculated as the fraction of all long-lived particle decays in the truth record which are matched to a reconstructed vertex. It is calculated as a function of the radial decay position of the long-lived particle, as shown in Figure 4.6(a).

The efficiency is highest in the $4 < r < 300$ mm region. The efficiency drops off closer to the interaction point due to the $|d_0|$ selection on core tracks, and at larger radii due to there being fewer displaced tracks reconstructed which satisfy the hit requirements.

Several tracks from long-lived particle decays may fail the strict selection placed on core tracks, which are necessary to reduce the CPU of the two-track vertex finding step. This results in displaced vertices being reconstructed with lower masses and track multiplicities. To improve these vertex kinematic quantities, the *track attachment* step of the algorithm attaches lower quality tracks to reconstructed DVs.




Figure 4.6: Performance of the secondary vertexing algorithm. In (a), the total vertex reconstruction efficiency as a function of radial decay position of the long-lived particle is shown [83]. In (b), the efficiency for selecting DVs with $N_{\text{trk}} \geq 5$ and $m > 10 \text{ GeV}$ as a function of the radial decay position is shown. The simulated events used to calculate the efficiency shown in (b) contain massive ($m = 200 - 2000 \text{ GeV}$) long-lived particle decays to Standard Model quarks, and are described further in Section 6.3.

To be considered in the track attachment step, tracks must satisfy the following selection:

- Not already be included in any reconstructed displaced vertex
- $p_T > 1 \text{ GeV}$
- $\chi^2 < 20$
- Has a hit on the next outermost layer with respect to the DV position (hits on innermost layers are allowed, unlike for core tracks)
- $|d_0|/\sigma_{d_0} < 5$ and $|z_0|/\sigma_{z_0} < 5$ with respect to the DV

Tracks passing this selection are added sequentially to DVs (starting from the highest N-track DV) if a vertex fit with that track is successful. Each track is only allowed to be associated to one vertex. The track attachment step does not change the reconstruction efficiency of the algorithm. However, it does have an impact on the efficiency of selecting displaced vertices satisfying certain track multiplicity and mass requirements. The fraction of long-lived particle decays in simulated events which satisfy $n_{\text{Tracks}}^{\text{DV}} \geq 5$ and $m > 10 \text{ GeV}$ before and after the track attachment step is shown in Figure 4.6(b). The selection efficiency increases by 15%-30% for long-lived particle masses between 50 GeV and 2.5 TeV.

The vertex position resolution is calculated by taking the difference between the reconstructed vertex position and the true long-lived particle decay position, and taking the standard deviation of the resulting distribution. The vertex position resolution calculated as a function of the number of tracks in the vertex and on the radial position of the true long-lived particle decay are shown in Figure 4.7. The resolution significantly improves as the track multiplicity of the displaced vertex increases and as the

distance from the primary vertex decreases. The latter trend occurs because displaced vertices which are reconstructed further from the primary vertex will be composed of tracks which have fewer hits on the innermost detector layers. These tracks will therefore have a worse track-parameter resolution, which will result in a worse vertex position resolution. Furthermore, fewer tracks will be reconstructed for charged particles originating from significantly displaced locations (see Figure 4.3), which further reduces the position resolution at higher radii.




Figure 4.7: The σ_r , σ_z , and $r \times \sigma_r$ resolutions as a function of (a) the number of tracks and (b) the radial distance between the true long-lived particle decay position and the primary vertex. The position resolution is calculated in simulated events containing long-lived neutralinos with $m_{\tilde{\chi}_1^0} = 2050$ GeV. Taken from Ref. [83].

4.5 Jets

Jets are reconstructed from clustered energy deposits in the calorimeter [85]. The calorimeter cells are clustered into three-dimensional topological clusters (topo-clusters) using a nearest-neighbor algorithm [86]. An event-by-event correction is applied to the topo-clusters which takes into account the assumption that the jet originates from the location of the primary vertex. The topo-clusters are then formed into jets using the anti- k_t algorithm with a size parameter value of $R = 0.4$ [86].

The anti- k_t algorithm defines the distance between two objects as

$$d_{ij} = \min(p_{Ti}^{-2}, p_{Tj}^{-2}) \frac{\Delta_{ij}^2}{R^2}, \quad (4.2)$$

and the distance between an object and the beam (B) as

$$d_{iB} = p_{Ti}^{-2}, \quad (4.3)$$

where $\Delta_{ij}^2 = (y_i - y_j)^2 + (\phi_i - \phi_j)^2$ and p_{Ti} , y_i and ϕ_i are the transverse momentum, rapidity, and

azimuth of object i . The algorithm iteratively clusters pairs of objects together which have the smallest distance between them in the event. If the smallest distance is d_{iB} , then object i is called a jet and removed from the collection of objects. This procedure repeats until no objects are left. The anti- k_t algorithm is a powerful jet clustering algorithm because it is configured with only one parameter, can be run relatively fast, and it reconstructs jets which are collinear and infrared safe [87]. This means that the final jet collection is stable against additions of soft emissions or any splitting of objects into collinear pairs.

The four-momentum of the final jets are calculated from the sum of the energy and momenta of topo-clusters used to create them. The energy of these jets are defined at the electromagnetic (EM) scale, which means that they correctly account for energy depositions from electromagnetic showers. However, the energy of the final jets needs to be calibrated in order to better model the jet energy at the particle level. The jet calibration procedure aims to correct the measured energy of jets to account for pileup effects, geometrical effects, a non-compensating calorimeter response, and differences between data and simulated events.

Simulated events are used to determine the energy scale and resolution of jets by comparing the reconstructed jets with particle-level truth jets. The truth jets are reconstructed using stable final-state particles in the truth record, excluding particles which leave little or no energy in the calorimeters (muons and neutrinos) as well as any particles from pileup interactions. Truth jets are matched to reconstructed jets if they are within $\Delta R = 0.3$ of each other.

The jet energy scale calibration proceeds in the following steps:

1. Pileup correction: This correction remove the excess energy from additional proton-proton interactions occurring in the same and in nearby bunch crossings. The corrections are based on the jet area and transverse momentum density of the event, as well as the mean number of interactions per bunch crossing and the number of reconstructed primary vertices in the event.
2. Absolute JES calibration: This correction ensures that the reconstructed jets agree in energy and direction with truth jets reconstructed from simulated dijet events, accounting for effects like the non-compensating calorimeter response (see Section 3.3.6) and energy loss in passive material. The jet energy response, defined as the as the mean of a Gaussian fit to the core of the $E^{\text{reco}}/E^{\text{true}}$ distribution, is shown as a function of the jet- p_T in Figure 4.8(a). The energy reconstructed in the calorimeters is significantly lower than the true energy of the particle, and it depends significantly on the amount of passive material which the particles pass through, hence the dependence on η . The jet calibration factor in this step is taken as the inverse of the average energy response shown in Figure 4.8(a).
3. Global sequential calibration: Fluctuations in the jet response may still exist after the above calibrations, originating from differences in the flavor and energy distribution of constituent particles and on the jet development in the calorimeter. This correction aims to reduce effects

from these fluctuations in order to improve the jet resolution, without changing the average jet energy response.

4. Residual in situ calibration: A final correction is applied to correct for any remaining difference between jets in data and in simulated events, that arise from imperfect simulation of the hard scatter, underlying event, jet formation, pileup, and particle interactions with the detector.

The uncertainty on the jet energy scale as a function of the jet- p_T is shown in Figure 4.8(b). These uncertainties are derived from the *in situ* measurements, pile-up effects, flavour dependence, and some additional effects. The uncertainty is 1% for jets with $250 < p_T < 2000$ GeV and increases to 5% at very low jet- p_T and to 3.5% at very high jet- p_T [85]. The jet energy resolution is measured using dijet events in which the transverse momenta is well balanced between the two leading jets. Deviations from an exact balance in the transverse momenta allow for the measurement of the jet energy resolution. Particle level effects such as additional radiation and non-perturbative processes (i.e. hadronization, multi-parton interactions) also contribute to the measured asymmetry. These effects are subtracted out from the final result. The jet energy resolution and its systematic uncertainty as a function of the jet- p_T is shown in Figure 4.9. The jet energy resolution is measured in both simulated dijet events and in data. The resolution measured in simulated events tends to underestimate the data, and this non-closure is translated into a systematic uncertainty. The relative jet energy resolution ranges from 0.35-0.04 as a function of the jet- p_T .




Figure 4.8: In (a), the jet energy response as a function of the jet- p_T is shown. In (b), the fractional jet energy scale systematic uncertainty components for central jets as a function of jet p_T is shown. The total uncertainty (all components summed in quadrature) is shown as a filled region topped by a solid black line. Taken from Ref. [85].




Figure 4.9: The relative jet energy resolution as a function of the jet- p_T is shown. Taken from Ref. [85].

Jets from non-collision backgrounds

There are several sources of non-collision backgrounds [88] which may be misidentified as jets in the calorimeters. These are:

1. Beam-induced background: Protons may escape the beam upstream of the detector, producing secondary cascades and muons which may reach the ATLAS detector and be misidentified as a jet.
2. Cosmic-ray showers: Cosmic ray showers produced in the atmosphere above the ATLAS detector may overlap with collision events. The ATLAS detector is situated underground which provides some shielding, however high- p_T muons from cosmic ray showers may traverse the detector.

In both of the above cases, the muons traversing the detector may deposit radiative energy losses in the calorimeters (i.e. from bremsstrahlung), which could be misidentified as a jet. While the rate of misidentified jets from beam-induced background and cosmic ray showers are low, they are still an important background in searches for rare processes.

Noisy calorimeter cells may also be misidentified as a jet. Dedicated selection criteria, called jet cleaning, are designed to efficiently reject jets from background processes while keeping the highest efficiency selection for jets produced in proton-proton collision. This is achieved by checking for abnormal pulse-shapes in the signal from the LAr calorimeter, comparing the amount of energy deposited in different layers of the calorimeters, and by using track-based variables [89].

Track jets

A special collection of jets, called track jets, are built from tracks instead of energy calorimeter deposits. These track jets are used to estimate backgrounds in the search for displaced vertices in multijet events. The track jets are constructed with the anti- k_T algorithm with $R = 0.4$ using all standard tracks with $p_T > 1$ GeV, $|d_0| < 2$ mm, and which pass several loose selection criteria designed to remove fake reconstructed tracks. In order to take into account the effects arising from the presence of pile-up interactions, track jets are reconstructed for each collision vertex in the event. Prior to track jet finding, each track is associated to only one vertex according to the following procedure: first, tracks are assigned to a collision vertex if their longitudinal impact parameter significance $|z_0/\sigma_{z_0}| < 3$ with respect to that vertex. Second, any remaining unassociated tracks are assigned to a collision vertex if their longitudinal impact parameter satisfies $|z_0| < 0.5$ mm with respect to that vertex. If a track is found to be compatible with multiple collision vertices, then the vertex with the highest $\sum p_T^2$ is chosen. The four-momenta of the track jets are calculated from the sum of the four-momenta of its constituent tracks. Unlike jets formed from calorimeter deposits, which are subject to variations from the calorimeter response, track jets do not need to undergo a calibration procedure.

b-tagging

Jets originating from b-quarks are tagged with a dedicated algorithm [90]. The b-quark from which the jet originates hadronises into b-hadrons, which then decay into lighter hadrons. With a mean lifetime of approximately 1.6 ps [91], b-hadrons with sufficient momentum can decay at a measurable distance from the collision point, creating a secondary vertex. The b-tagging algorithm uses a multivariate approach to identify jets originating from b-quarks.

First, low-level tagging algorithms reconstruct the features of b-jets. Tracks are matched to jets using the ΔR between the track momenta and the jet axis. The low-level algorithms utilize the fact that tracks from the b-hadron decays can have large impact parameters, and multiple tracks from the same b-hadron decay can form a displaced vertex. The outputs of the low-level tagging algorithms are then combined into high-level algorithms consisting of multivariate classifiers.

4.6 Electrons

Electrons are reconstructed using tracks and clusters of energy deposits in the electromagnetic calorimeter. The clusters are formed from energy deposits in topologically connected calorimeter cells-[86]. The clusters are matched to inner detector tracks which are consistent with having originated from the beam interaction region. The tracks are refit to account for bremsstrahlung. The electron reconstruction efficiency in $Z \rightarrow e^+e^-$ events reaches approximately 95% for electrons with $p_T > 10$ GeV [92].

The energy of the electron is estimated using an area corresponding to 3×7 (5×5) second-layer cells in the barrel (endcap) region of the electromagnetic calorimeter. The energy scale is calibrated using simulation and $Z \rightarrow e^+ e^-$ and $J/\psi \rightarrow e^+ e^-$ decays in data [93].

4.7 Muons

Muons have a signature in the detector characterized by a track in the inner detector and muon spectrometer with few energy deposits in the calorimeters. Tracks are reconstructed in the muon spectrometer from short, straight-line track segments which are formed from hits in an individual MS tracking station [94]. Segments are combined which are compatible with having been produced from a particle originating from the direction of the interaction point. A track fit is performed which takes into account the trajectory of the muon through the magnetic field as well as effects of interactions with detector material. If there is a matching track reconstructed in the inner detector, then a combined fit using both inner detector and muon spectrometer hits is performed, taking into account energy loss in the calorimeters.

Additional selection on the number of hits in different detector subsystems and on the ratio of the charge and momentum (q/p) measured in the inner detector and muon spectrometer are used to identify good muon candidates. The identification efficiency of muons satisfying a ‘loose’ identification selection in simulated $t\bar{t}$ events is above 98% for muons with $p_T > 10$ GeV [94].

4.8 Missing transverse momentum

The missing transverse momentum vector, E_T^{miss} , is defined as the negative vector sum of the transverse momenta of all reconstructed and calibrated objects in the ATLAS detector, with the addition of a *soft term* [95]. The soft term is built from the momenta of tracks reconstructed in the inner detector which are associated to the primary vertex but are not associated to any of the calibrated objects. Momentum conservation in the plane transverse to the beam axis implies that the transverse momenta of all particles in the final state should sum to zero. Therefore, E_T^{miss} is measured as a way to infer the presence of weakly-interacting stable particles (i.e. neutrinos or neutralinos) which do not leave a signal in the ATLAS detector. E_T^{miss} may also be generated from the mis-measurement of momenta of objects in the event.

The ATLAS Fast Tracker

The ATLAS Fast Tracker (FTK) is a hardware-based track finder which was envisioned to be integrated into the ATLAS trigger system [96]. The FTK is designed to provide the HLT with tracking information over the full fiducial volume of the ATLAS inner detector at the full L1 trigger rate of 100 kHz. In this chapter, the FTK system as it was built for commissioning at the end of Run-2 data taking and during Long Shutdown 2 (LS2) is described. The FTK project was discontinued in 2019 and will not be integrated into the trigger system for Run-3. Instead, more tracking in the HLT will be performed for signatures that are expected to benefit the most (i.e. jets and E_T^{miss}). Integrating tracking information into the trigger at the earliest possible stage is a top priority of LHC experiments. For the high luminosity upgrade of the LHC (HL-LHC), the CMS experiment plans to perform hardware-based track finding [97] while the ATLAS experiment plans to perform extensive tracking in the trigger using CPU cores and accelerator hardware (GPUs, FPGAs, ...) [98].

Section 5.1 motivates the use of tracking information in the trigger. Section 5.2 provides a brief overview of the FTK's track-finding system and Section 5.3 describes the production of the FTK track-finding configuration. The configuration of the scaled-up FTK system for the purpose of commissioning in 2019 is discussed in Section 5.4, and the adaptability of the FTK system to new detector conditions is evaluated in Section 5.5.

I contributed to the FTK project in several ways. First, I produced and compared two possible configurations for the scaled-up FTK system (Section 5.4). Secondly, I increased the adaptability of the FTK system to changing beamspot and silicon module positions (Section 5.5). I demonstrated the ability to produce separate configurations to recover lost efficiency from a shifted beamspot position, and implemented an algorithm to correct the FTK system to take into account shifts in the IBL module positions.

5.1 Motivation for tracking in the trigger

Charged particle tracks are critical for determining which L1-triggered events should be kept for further processing. For example, tracks can be used for vertex finding, giving the HLT early access to event-level quantities, like the primary vertex positions and the total number of primary vertices in the event. Tracks can also be used to distinguish objects originating from the primary vertex and pileup vertices, thereby allowing for better pileup rejection [99] and the ability to keep p_T thresholds on triggered objects low even in high pileup conditions.

Due to limitations in computing power and trigger latency, software-based track finding in the HLT is limited to small areas of interest around L1-triggered objects, called Regions of Interest (ROI). The FTK removes these limitations by providing the HLT with all tracks reconstructed in the inner detector with $p_T > 1$ GeV for every event accepted by the L1 trigger. Without the need to perform software-based tracking in ROIs, the HLT would have more resources available for sophisticated trigger algorithms using FTK tracks.

Having global tracking information available in the HLT will allow for event and object selections which are more closely matched to the offline selections, thereby reducing inefficiencies caused by a mismatch between online and offline algorithms. FTK tracks can also be used in the triggers of track-based signatures, such as hadronic tau decays and b-jets. For example, b-tagging in the HLT could be improved in multiple ways with the use of FTK tracks; either the FTK tracks could be used directly as input to a custom b-tagging algorithm, or they could be used to select jets of interest (i.e. jets containing FTK tracks with high d_0), so that the HLT would have enough resources to use a near-offline quality b-tagger. Finally, the FTK would introduce a new capability of triggering directly on displaced tracks, which is a signature of long-lived particle decays.

5.2 Track finding with the FTK

Track-finding in the FTK proceeds in two stages. In the first stage, eight of the twelve silicon tracking layers are used to perform pattern recognition and initial track fitting. The pattern recognition uses Associative Memory to identify track candidates using pre-defined combinations of hits which are consistent with having come from a charged particle originating from the interaction point. The helix parameters of track candidates are calculated using a linear system of equations. The calculation utilizes a set of pre-defined fit constants, which take into account the detector geometry and alignment. In the second stage, track candidates are extrapolated to the four remaining silicon layers and the track helix parameters are evaluated. A list of all tracks in the event which pass several quality selection criteria and which have $p_T > 1$ GeV are sent to the HLT. In this section, each step of the track-finding method used in the FTK is described. A detailed description of the FTK hardware system can be found in Ref. [100].

5.2.1 System overview

Following a Level-1 accept, data from eight tracking layers of the inner detector are sent to the FTK Auxiliary Card (AUX). Three tracking layers from the Pixel and five tracking layers from the SCT are used, shown in blue on Figure 5.1. The AUX clusters the hits into coarse resolution hits, called super-strips, which are typically 128 strips or 30×76 pixels in size.




Figure 5.1: Illustration of the FTK track-finding system. Rounded boxes represent Superstrips. Taken from Ref. [100].

The super-strips are sent to the Associative Memory Board (AMB), where pattern matching is performed. A pattern is a set of eight super-strips from different layers, corresponding to the positions on each silicon layer where a charged particle originating from the interaction point could pass through. The hit super-strips sent from the AUX are compared to a set of pre-defined patterns in a pattern bank. The pattern bank is stored in content-addressable memory, the Associative Memory (AM), which allows for rapid pattern comparisons at a rate of ~ 12 petabytes/s on each AMB [100]. Patterns which have seven or eight layers hit are considered matched. A set of super-strips matched to a pattern is called a road. A schematic of an 8-layer (8L) road is shown in blue on Figure 5.1.

Next, the hits on each road are sent to the track fitter. The track fitter uses full resolution hits rather than super-strips. Every combination of hits on different layers within the road is considered an 8L track candidate. The χ^2 squared of each 8L track candidate is calculated and required to satisfy $\chi^2/\text{ndof} < 6$. If multiple tracks in the same road pass the χ^2 selection and have fewer than 6 non-matching hit coordinates, they are considered duplicates, and the track with the best χ^2 is retained.

Following the 8L track fit, track candidates are extrapolated to the four remaining detector layers

(the IBL and three SCT tracking layers), denoted by the purple super-strips on Figure 5.1. Nearby hits on each layer are identified and a full twelve layer (12L) track fit is performed. Resulting tracks are required to satisfy $\chi^2/\text{ndof} < 4$. Tracks on all roads with fewer than 6 non-matching hit coordinates are considered duplicates, and the track with the best χ^2 is retained. The track parameters are evaluated for each 12L track and tracks with $p_T > 1$ GeV are sent to the HLT. This p_T threshold was chosen to match the offline selection for most physics analyses.

Track fitting in the FTK does not utilize minimization. Instead, a system of linear equations with pre-defined fit constants is evaluated. Using this method, approximately 10^9 track candidates can be fit per second on the FPGAs. The 8L track fit in the AUX has eleven coordinates (five SCT layers and three pixel layers) and five helix parameters. Thus there are six degrees of freedom, each of which corresponds to a function of coordinates that should equal zero. The χ^2 is calculated as the sum of the squares of those functions,

$$\chi^2 = \sum_{i=1}^6 \left(\sum_{j=1}^{11} S_{ij} x_j + h_i \right)^2, \quad (5.1)$$

where x_j is a hit coordinate and S and h are arrays of pre-defined fit constants. Separate sets of fit constants are defined for different regions of the detector, called sectors. A sector is comprised of one silicon module from each tracking layer (a few centimeters in size). Though wider than patterns, sectors are narrow enough that the linearized track fit provides helix parameters with sufficiently high resolution.

The same underlying principle is also used for the 12L track fit. In addition to the χ^2 , the track parameters are also evaluated with the equation,

$$\tilde{p}_i = \sum_{l=1}^{16} C_{il} x_l + q_i, \quad (5.2)$$

where x_l is a hit coordinate and C and q are arrays of pre-defined fit constants. The array $\tilde{p}$ contains the five track helix parameters and 11 partial χ values, which are summed together as in Equation 5.1. In both the 8L and 12L track fit, one hit may be dropped from the fit in order to improve efficiency. If a hit is dropped, the best expected hit position on that layer is used in order to compute track parameters, which is called the guessed hit (shown in orange in Figure 5.1).

The FTK system must perform fast track-finding while handling a large data input rate. To handle this challenge, the FTK system is segmented in $4 (\eta) \times 16 (\phi) = 64$ overlapping regions, called towers. Each tower consists of two processing units (PU), each containing an AUX and AMB. Track-finding is performed in parallel on each tower.

A software simulation of the FTK system is a vital tool for the design and commissioning process. Offline simulation of the FTK (FTKsim) includes the full FTK data processing chain: hit clustering, data organization and distribution, pattern matching, track fitting, and final processing [101]. FTKsim

takes as input the silicon hits from either real data or from simulated events and outputs a list of tracks reconstructed with the FTK. The same sectors, constants, and patterns intended for use in the real FTK system can be used in FTKsim, which makes it an instrumental tool in optimizing the configuration of the FTK system.

5.3 Configuring the FTK

This section describes the method of producing sectors, constants, and patterns which are used to configure the FTK system.

5.3.1 Production of sectors and constants

Sectors and constants are determined offline and loaded into the FTK system before track-finding. In order to define a robust set of sectors, one needs to understand all possible sets of modules that a charged particle originating from the interaction point might pass through as it traverses the inner detector. The production of sectors (and constants) therefore begins with the simulation of muons traversing the inner detector. Muons are chosen because they provide a relatively clean track in the inner detector as they do not scatter and undergo hadronic interactions as much as electrons and charged pions do. Approximately 300 million simulated muons are needed to cover all possible combinations of modules which are compatible with tracks originating from the beam spot position. Muons (with both positive and negative charge in equal measure) are drawn from uniform distributions in the five track parameters: $d_0 \in [-2.2, 2.2]$ mm, $z_0 \in [-120, 120]$ mm, $q/p_T \in [-0.8, 0.8]$ GeV⁻¹, $|\eta| < 3$, and $\phi \in [-\pi, \pi]$. The muon tracks are reconstructed, and each set of modules hit by the muons is considered a sector.

A detailed simulation of the muons as they traverse the detector is critical to produce fit constants. However, the full ATLAS simulation does not fully reproduce small detector misalignments, which may impact the FTK's track-finding efficiency and track parameter accuracy. Therefore, the single-muon events are simulated with the ATLAS detector conditions and alignment which is observed in real data using the ATLAS 'data overlay' infrastructure [102]. The conditions and alignment of one run in data are used to generate the sectors and constants for that data-taking year.

A set of fit constants is calculated for each sector which has a sufficient number of muons passing through it (typically at least five). The fit constants are evaluated by minimizing the distance between the linear track parameters $\tilde{p}_i$ used in Equation 5.2 and the muon's true track parameters p_i :

$$\min \left(\left\langle (\tilde{p}_i - p_i)^2 \right\rangle \right) \forall i = 1, \dots, 5, . \quad (5.3)$$

This procedure is performed separately for 8L and 12L tracks, resulting in an average 13,000 8L and 28,000 12L sectors per tower.

5.3.2 Production of pattern banks

In addition to sectors and constants, pattern banks are also determined offline and loaded into the FTK system before track-finding. The process of producing a robust pattern bank is similar to that of sectors - the trajectories of charged particles passing through the inner detector are analyzed and the hit super-strips (rather than the hit modules) are recorded. In the first step of pattern bank production, a *thin-space pattern (TSP) bank* is produced using patterns of the highest available resolution, which is determined by the super-strip width. In the second step, "don't care" (DC) bits are used to produce a *DC pattern bank* containing patterns with variable width. Variable width patterns allows the bank to retain a high track-finding efficiency while reducing the overall number of patterns to fit within the hardware specifications of the AM chips. Lastly, patterns in the DC pattern bank are redistributed to flatten the FTK's track-finding efficiency as a function of track η .

Thin-space pattern bank production

The tracks used to produce the candidate patterns are drawn from the same uniform distributions in the five track parameters that are used in the production of sectors. Due to the fact that patterns are thinner than sectors, a much larger number of muons (approximately 400 billion) are needed to produce a pattern bank. It is not feasible to use the full ATLAS simulation to determine the corresponding hit positions of each muon. Instead, the previously determined sectors and fit constants are used to estimate the hit positions of each track. Equation 5.2 is inverted so that one can determine the approximate hit coordinates from the track parameters. The hit coordinates of each track are then transformed into a sequence of super-strips. Each pattern candidate corresponds to eight coarse resolution hit positions on different detector layers which have been traversed by a muon. The multiplicity of tracks producing the same pattern is referred to as the pattern's coverage. An example of the coverage for two towers in a pattern bank is shown in Figure 5.2. Patterns with a high coverage are most important for track finding. The higher a pattern's coverage is, the more likely a track passes through this pattern and the larger its volume is in the phase space of track parameters. Due to geometrical considerations, patterns in the endcap tend to have a higher coverage compared to patterns in the barrel region. The full set of unique pattern candidates, ordered by tower, sector, and coverage, is called the TSP bank.

Variable-width pattern bank production

The number of patterns in a pattern bank is limited by the available space in the AM chips. With the nominal FTK configuration containing 128 PU's (2 per tower), the FTK is able to store 2^{24} (16.8 million) patterns per tower. The TSP bank contains approximately one billion patterns per tower. Therefore the size of the TSP bank far exceeds the hardware capacity. This motivates the introduction of patterns with variable width, which allows for a reduction in the overall number of patterns while keeping the pattern bank's track-finding efficiency high.




Figure 5.2: The coverage of patterns in two towers of a TSP pattern bank. The coverage is defined as the number of tracks traversing each pattern.

The benefit of introducing variable-width patterns is illustrate in Figure 5.3. To fully cover the trajectory of three charged particles, one would need three high resolution patterns, as shown in (a). To reduce the number of patterns, one could increase their width uniformly as in (b), for example by increasing the super-strip width. However, because the pattern is wider, there will be many more spurious hits and fake tracks resulting from that pattern. Introducing variable-width patterns allows one to build a pattern as in (c), where all three particle trajectories are covered with one pattern, while the rate of fake tracks is minimized.




Figure 5.3: In this illustration, three charged particle trajectories are covered with (a) three high resolution patterns, (b) one wide pattern, and (c) one pattern with variable-width resolution. The colored (gray) boxes indicate super strips which are (are not) included in the pattern.

Defining patterns with a variable width is made possible because the AM exploits ternary logic, where each bit is identified by 0, 1, or X. Here, X means "don't care" (DC) in the hit-to-stored-value comparison. The structure of the AM chips allows for the use of X values to logically OR neighboring super-strips, thereby increasing the width of that pattern on one detector layer. The pattern is considered

matched, regardless of which of the neighboring super-strips is hit on that detector layer. This feature allows for the width of DC patterns to vary, layer by layer, while increasing the effective number of patterns stored in the available hardware. Setting a total number of N_X DC bits across all eight layers of a pattern increases the effective number of patterns stored in a single AM address to 2^{N_X} .

With more effective patterns stored in the AM chips with the use of DC bits, the FTK's track-finding efficiency will increase. However, patterns with many DC bits will suffer from a degraded spatial resolution and a higher rate of combinatorial background in the track fitting step. The latter will increase the FTK's dataflow, which is the number of matched patterns and track fits performed per event. The dataflow must be kept within hardware limits so that the FTK can operate at the full L1 trigger rate of 100 kHz. This is achieved by setting a maximum number of DC bits per pattern, $N_X \leq N_{X,\max}$. The maximum number of DC bits per pattern, summed across all layers, is defined separately for the barrel ($N_{X,\max}^B$) and endcap ($N_{X,\max}^{EC}$) regions. Typical values for the maximum number of DC bits per pattern are $N_{X,\max}^B = 7 - 9$ and $N_{X,\max}^{EC} = 4 - 6$.

The DC pattern bank is built by filling AM addresses with TSP patterns, in order of decreasing coverage. Where possible, patterns are merged with previously stored patterns using DC bits. This process continues until all AM addresses are filled and there is no room to include further pattern candidates.

The efficiency of a DC pattern bank, as a function of the number of patterns per tower, is shown in Figure 5.4. As the number of patterns in the bank increases, the track-finding efficiency quickly increases before reaching a plateau. The standard 128-PU FTK configuration allows for the storage of 16.8 million patterns, which results in a track-finding efficiency on the plateau of the efficiency curve.

The η -inclusive track-finding efficiency does not always fall between the efficiencies calculated in the barrel and endcap regions separately. This is because it includes the transition region ($1.2 < \eta < 1.6$) between barrel and endcap detector components, which has a lower efficiency.

Partitioning

The track-finding efficiency of a DC bank as a function of track η is shown in blue in Figure 5.5. Two drops in efficiency are observed at $\eta \approx \pm 1.2$, which is due to the change in geometry from the barrel to endcap detector layers in these regions. To optimize the performance of the pattern banks the number of patterns are redistributed to flatten the efficiency as a function of η . This process is called partitioning.

In the partitioning step, TSP patterns are reassigned to AM addresses, giving a larger budget of AM addresses to bins of η with low efficiency. Each sector is assigned to exactly one bin of the η distribution. The η bin is determined by averaging the minimum and maximum η of tracks in that sector. For each sector falling into an η bin with efficiency ϵ less than the target efficiency ϵ_0 , a weight




Figure 5.4: Simulated track-finding efficiency of the AM pattern bank as a function of the number of patterns per tower in the bank. The DC pattern banks are produced from the same TSP patterns and no partitioning is applied. The efficiency is calculated as the fraction ‘true’ muons which are reconstructed by the FTK in a single-muon MC sample. Tracks are required to have $|z_0| < 110$ mm and $|d_0| < 2.2$ mm. Taken from Ref. [100].

is applied to the number of patterns allowed in that sector:

$$w = \frac{-1}{\ln(1 - \epsilon/\epsilon_0)}, \quad (5.4)$$

The budget of AM addresses in η bins with a high efficiency are lowered accordingly, to keep the overall number of AM addresses constant. This iterative process is repeated until the track-finding efficiency is roughly flat as a function of η .

The track-finding efficiency of the DC bank after partitioning is shown in red in Figure 5.5. The remaining oscillations in the η distribution after partitioning is larger than the uncertainty due to statistical precision. This oscillation is likely due the crude estimation of the average η in each sector, when in reality each sector contains tracks with a distribution of η values.




Figure 5.5: FTK track-finding efficiency as a function of the track η , before and after applying partitioning. Tracks are taken from a single-muon Monte Carlo sample and are required to have $|z_0| < 110$ mm and $|d_0| < 2.2$ mm. The pattern bank is created for a 64-PU configuration with $N_{X,\max}^B = 9$ and $N_{X,\max}^{EC} = 6$. Taken from Ref. [100].

5.4 Production of pattern bank for commissioning of 128-PU system

During LS2, commissioning of the FTK was performed in order to prepare the system for Run-3 data-taking. In 2018, the FTK was commissioned with a pattern bank compatible with a 64-PU system, which contains half of the number of patterns which would be available with the full 128-PU FTK system. To facilitate the scaling up of the FTK system from a 64-PU to a 128-PU system, a new pattern bank is produced. More processing units means that there are more AM addresses for the storage of patterns in the pattern bank. Thus, a higher track finding efficiency can be achieved. However, the number of allowed DC bits needs to be adjusted in order to keep the dataflow within the hardware limits.

Two 128-PU pattern banks are created with different DC bit configurations. The first configuration has fewer DC bits ($N_{X,\max}^B = 7$ and $N_{X,\max}^{EC} = 4$) and therefore results in lower dataflow, while the second configuration has more DC bits ($N_{X,\max}^B = 9$ and $N_{X,\max}^{EC} = 6$) resulting in a higher dataflow. These new pattern banks are compared to the 64-PU pattern bank used for commissioning in 2018, which has $N_{X,\max}^B = 9$ and $N_{X,\max}^{EC} = 6$. The efficiency of these pattern banks are compared in Table 5.1. The pattern bank with 128-PU and the most DC bits has the highest track-finding efficiency at 94.4%.

The track-finding efficiency of these pattern banks as a function of the five track parameters are shown in Figures 5.6-5.8. The track-finding efficiency peaks for tracks with low impact parameters (d_0 and z_0) and high p_T , and stays roughly constant as a function of track $|\eta|$. The 128-PU bank with $N_{X,\max}^B = 9$ and $N_{X,\max}^{EC} = 6$ results in the highest track-finding efficiency across the entire range of all track parameters, with the largest improvement at high $|d_0|$ and low p_T .

5.4 Production of pattern bank for commissioning of 128-PU system

Track-finding efficiency of various pattern-bank configurations		
64-PU, $N_{X,\max} = (9, 6)$	128-PU, $N_{X,\max} = (7, 4)$	128-PU, $N_{X,\max} = (9, 6)$
$90.0 \pm 0.1\%$	$91.2 \pm 0.1\%$	$94.4 \pm 0.1\%$

Table 5.1: Track-finding efficiency of pattern banks with respect to ‘truth record’ tracks in a single-muon Monte Carlo sample. Only tracks with $|\eta| < 2.5$, $|z_0| < 110$ mm and $|d_0| < 2.2$ mm are considered in the efficiency calculation. Each column indicates the DC bit configuration of the bank $N_{X,\max} = (N_{X,\max}^B, N_{X,\max}^{EC})$. The statistical uncertainties are shown.




Figure 5.6: FTK track-finding efficiency calculated with respect to the MC ‘truth record’ as a function of the (a) d_0 and (b) z_0 of tracks in a single-muon MC sample. Tracks are required to satisfy $|\eta| < 2.5$, $|z_0| < 110$ mm, and $|d_0| < 2$ mm. The efficiency is calculated for three pattern banks of varying size and DC bit configurations, $N_{X,\max} = (N_{X,\max}^B, N_{X,\max}^{EC})$, as indicated by the legend. Taken from Ref. [100].




Figure 5.7: FTK track-finding efficiency calculated with respect to the MC ‘truth record’ as a function of the (a) η and (b) p_T of tracks in a single-muon MC sample. Tracks are required to satisfy $|\eta| < 2.5$, $|z_0| < 110$ mm, and $|d_0| < 2$ mm. The efficiency is calculated for three pattern banks of varying size and DC bit configurations, $N_{X,\max} = (N_{X,\max}^B, N_{X,\max}^{EC})$, as indicated by the legend. Taken from Ref. [100].




Figure 5.8: FTK track-finding efficiency calculated with respect to the MC ‘truth record’ as a function of the track ϕ in a single-muon MC sample. Tracks are required to satisfy $|\eta| < 2.5$, $|z_0| < 110$ mm, and $|d_0| < 2$ mm. The efficiency is calculated for three pattern banks of varying size and DC bit configurations, $N_{X,\max} = (N_{X,\max}^B, N_{X,\max}^{EC})$, as indicated by the legend.

The gradual decrease in efficiency at high d_0 and z_0 and at low p_T is due to a lower coverage in those regions. More patterns are needed in those regions to achieve the same high efficiency as in the central (or high p_T) regions. The steeper drop off observed in the d_0 distribution is due to an edge effect; while all other track parameters are plotted within ranges smaller than the range used to select tracks in the pattern bank production, the d_0 is shown for the full range of d_0 used in pattern bank production.

The dataflow of each 128-PU bank is compared in Table 5.2. The dataflow is simulated using FTKsim with a subset of RAW events collected with 2017 data with a pileup of $\mu = 60$. The dataflow is measured by counting the number of roads, 8L track candidates, 12L track candidates, and 12L tracks which pass overlap removal and are sent to the HLT. Compared to the lower dataflow configuration, the high-dataflow bank results in about a factor of 2 increase in the number of roads, a 50% increase in the number of 8-layer and 12-layer track candidates, and a 10% increase in the number of final FTK 12-layer tracks after overlap removal.

Dataflow of 128-PU pattern banks	$N_{X,\max} = (7, 4)$		$N_{X,\max} = (9, 6)$	
	Barrel	Endcap	Barrel	Endcap
Roads	1700	2100	3800	4200
8L-track candidates	43	49	63	75
12L-track candidates	37	38	54	57
Final 12L-tracks	3.8	1.9	4.2	2.1

Table 5.2: Measures of dataflow for two 128-PU pattern banks with different DC bit configurations, $N_{X,\max} = (N_{X,\max}^B, N_{X,\max}^{EC})$, as indicated by the column header. The average number of roads, 8L and 12L track candidates, and final 12L tracks after overlap removal has been applied are listed per event and per tower.

For the purpose of commissioning, it was decided to push the FTK system to its performance limits. The high dataflow 128-PU pattern bank with $N_{X,\max}^B = 9$ and $N_{X,\max}^{EC} = 6$ was chosen for FTK commissioning. In August 2019, this pattern bank was loaded into the FTK hardware and the system successfully produced 12-layer tracks in one FTK tower which was being used for commissioning.

5.5 Adaptation of the FTK to new detector conditions

An FTK configuration, which is defined by sectors, constants, and patterns, is sensitive to the detector geometry and beamspot position. In the production of an FTK configuration, the nominal detector geometry and best guess of the beamspot position for the upcoming data-taking period is used. It is therefore important to understand how the FTK's track-finding performance changes with small shifts in silicon module or beamspot positions and to develop strategies to quickly adapt the FTK system to handle these changing conditions.

5.5.1 Shifting beamspot position

The beamspot is the region in the detector where the proton-proton collisions occur and where primary vertices are reconstructed (see Section 4.3). The beamspot is defined with respect to the ATLAS coordinate system, which is defined through a detector alignment procedure. Ideally, the beamspot is at the center of the detector. In reality, it is centered slightly off axis in the transverse plane (~ 0.5 mm). The beamspot can move slightly during data-taking, due to small changes in the LHC beams or from shifting detector components which change the alignment. The position of the beamspot in the x direction in all Run-2 data-taking periods is shown in Figure 5.9. The beamspot shifts during a run are much smaller, typically within 0.2 mm in the x and y directions.




Figure 5.9: The position of the beamspot during Run-2. The data points are the hourly average of results of a maximum likelihood fit to the spatial distribution of primary vertices collected over a two minute period.

The nominal beamspot position is used when generating sectors, constants, and patterns. Changes in the beamspot position with respect to the nominal position introduce a modulation of the track $|d_0|$ vs. ϕ . This modulation is well understood and can be corrected for in the offline tracking algorithm. However it is still important to study if there is any impact on the FTK track-finding efficiency. In order to study the effect that small beamspot displacements have on the FTK performance, the efficiency of the FTK at various stages is evaluated using single-muon MC events with varying beamspot positions between ± 1 mm. This value was chosen because it encompasses the typical beamspot shifts over a data-taking period. The results are shown in Figure 5.10(a).

In the ‘sectors as patterns’ stage, pattern matching is performed with sectors instead of patterns, without subsequent track fits. Sectors are much wider than patterns, and therefore this results in a much higher efficiency. The efficiency is roughly constant as a function of beamspot position, which indicates that the sectors are largely unaffected by small shifts in the beamspot position. The sectors are wide enough that tracks originating from a displaced beamspot position within ± 1 mm still pass through the sector.

In the ‘pattern’ recognition stage, pattern matching is performed with the pattern bank, without subsequent track fits. As the beamspot moves away from the nominal position, the efficiency decreases by at most 5% when the beamspot is shifted by $\Delta x = +1$ mm. This indicates that the patterns are sensitive to small beamspot shifts; some tracks originating from a displaced beamspot position are no longer passing through a pre-defined patterns.

The efficiency of the FTK after performing 8L (‘1st stage’) and 12L (‘2nd stage’) track fits result in a further drop in efficiency. This efficiency loss is constant as a function of the beamspot position when compared to the pattern matching stage. This demonstrates that the constants are largely unaffected by small beamspot shifts; the χ^2 of the tracks still pass the selection required after each track fit.

The performance of patterns is impacted by small beamspot shifts, while sectors and constants are not. Therefore a new pattern bank will be needed in case of small beamspot shifts. Early commissioning data at the start of each data-taking run could be used to determine the beamspot position, so that an optimised pattern bank could be made in time for physics data-taking. Additionally, multiple pattern banks can be produced prior to data-taking, each ‘trained’ with tracks originating from a different beamspot position around the nominal. If the beamspot position were to change during data-taking, the pattern bank stored in the FTK system could be quickly (within a few hours) changed accordingly.

The feasibility of this proposal is demonstrated in Figure 5.10(b), where the efficiency of five pattern banks trained with muons originating from various beamspot x-positions is shown. By utilizing a different pattern bank at various beamspot shifts, most of the FTK’s track-finding efficiency is retained.




Figure 5.10: In (a), the efficiency of the FTK system at various stages as a function of the beamspot position in the single-muon MC sample is shown. In (b), the efficiency of the full FTK system using pattern banks trained with shifted beamspot positions in single-muon MC sample is shown. Tracks are required to satisfy $|\eta| < 2.5$, $|z_0| < 110$ mm, and $|d_0| < 2$ mm.

5.5.2 Changing silicon module positions

The detector conditions and alignment observed in one run of data is used to generate sectors, constants, and patterns. The FTK is therefore sensitive to the detector geometry (i.e. location and orientation of detector elements). Ideally the detector conditions and alignment stays constant throughout the data-taking year. However in practice, small changes in detector alignments can occur due to changes in environmental conditions of the detector (i.e. from magnetic fields or power and cooling systems) and mechanical relaxations after sudden changes [104].

One example of this is the Insertable B-layer (IBL), whose staves bow in the azimuthal direction (module local- x direction) by up to hundreds of microns relative to the design geometry. These distortions are caused by the use of materials with different coefficients of thermal expansion, the severity of which is determined by the operating temperature [105]. The magnitude of the bowing is parameterized with a parabolic function with the maximum at the center of the stave. The magnitude of the distortions are parameterized by a single number, B_x , which indicates the level of distortion with respect to the IBL stave orientation at the nominal operating temperature, $T = -20$ °C. Small changes in temperature occur over time due to accumulated radiation dose and instantaneous luminosity.

To study the impact that the IBL bowing has on the FTK performance, simulated single-muon events are reconstructed with the detector alignment and conditions found in two separate runs with significantly different IBL distortions: Run 349014 with $B_x = -6$ μm and Run 364292 with $B_x = -20$ μm . These events are then used to test the performance of the FTK system with FTKsim, using a pattern bank with a 64-PU configuration with $N_{X,\text{max}}^B = 9$ and $N_{X,\text{max}}^{EC} = 6$.

Pattern matching does not utilize the IBL layer, and therefore the pattern bank is unaffected by IBL

shifts. Furthermore, the overall FTK track finding efficiency remains high ($\epsilon \approx 90\%$) in both MC samples, which indicates that extrapolation to the IBL layer and the 12L track fit still produce tracks with a high quality χ^2 , despite small shifts in the IBL module positions.

The accuracy of the FTK track parameters which are sent to the HLT and will be used in trigger decisions is also important. The residuals of the five track parameters reconstructed with the FTK with respect to the offline track reconstruction, are shown in Figure 5.11. The p_T , η and z_0 of the tracks are unaffected by the small shift in the IBL displacement. The d_0 residual peak shifts by $22 \mu\text{m}$ and the ϕ residual peak shifts by 0.4 mrad in the sample reconstructed with more severe IBL displacement. This indicates that the fit constants are sensitive to the bowing of the IBL.

As the production of fit constants is computationally expensive and takes several weeks, it is not feasible to remake them when a module shifts position. Thus an algorithm to update quickly the fit constants during data-taking is desired. A simplified approach is developed which works upon the following principle: The track parameters, $\tilde{p}_i$, are derived from module-local hit positions using Eq. 5.2. If only small shifts in the local ϕ and η directions are considered (Δ_j) and the correlation between coordinates is ignored, then a change in the module positions,

$$x_j \rightarrow x'_j = x_j - \Delta_j, \quad (5.5)$$

can be incorporated directly into Equation 5.2. This leads to the following relation:

$$\tilde{p}_i = \sum_{l=1}^N C_{il}(x'_l + \Delta_j) + q_i \quad (5.6)$$

$$= \sum_{l=1}^N C_{il}x'_l + \left(\sum_{l=1}^N C_{il}\Delta_l + q_i \right) \quad (5.7)$$

$$= \sum_{l=1}^N C_{il}x'_l + q'_i. \quad (5.8)$$

The corrected fit constants q'_i are calculated using the original fit constants and the shift of the module center position. The results of this correction are shown in blue on Figure 5.11. The algorithm successfully reduces the d_0 residual by 32% , and the ϕ residual by 54% .

To reduce the d_0 and ϕ residuals further, one may apply a correction to the average hit coordinates of muon tracks in each sector, which are used as input to calculate the fit constants. A track extrapolator can be used to measure the change in the average hit coordinates with the new detector geometry with respect to the nominal detector geometry. This algorithm would in principle be able to take into account correlations between hit coordinates and track parameters, and would be able to consider more complicated module motion beyond shifts in the ϕ and η directions. However, it is limited by the accuracy of the track extrapolator [106].




Figure 5.11: Residuals of the five track parameters in a single-muon MC sample reconstructed with the FTK, with respect to the offline reconstructed values. The legend indicates the run whose conditions were used to reconstruct the simulated muons. In blue, the fit constants have a correction applied to take into account the difference in IBL distortions between Run 349014 and Run 264292. The mean, μ , of each residual peak is shown in the legend.

Search for long-lived particles with displaced vertices in multi-jet events

In this chapter, a search for new physics processes producing a massive, multi-track displaced vertex and multiple jets in the ATLAS detector is presented. The search uses 139 fb^{-1} of proton-proton collision data collected by the ATLAS experiment at a center-of-mass energy of $\sqrt{s} = 13 \text{ TeV}$. The analysis does not require that the jets are originating from the displaced vertex in order to be inclusive of various new physics models. Signal regions are defined which aim to select signal-like events from the dataset. The event selection are designed to maximize the expected signal yield while reducing the expected background to approximately one event per signal region. The estimate of background yields in the signal region are derived using data-driven techniques.

This is the first search for high mass, multi-track displaced vertices and multiple energetic jets with the ATLAS detector using $\sqrt{s} = 13 \text{ TeV}$ pp collision data. This search benefits from many improvements with respect to a previous version of this analysis performed with Run-1 LHC data [107]. These improvements include the new vertexing algorithm described in Section 4.4, custom analysis-level selections to reduce backgrounds, and a novel data-driven background estimation technique using jet-DV correlations.

I have made significant contributions to several parts of the analysis, including the displaced vertex (DV) track cleaning (Section 6.4.5), the estimate of background from all sources inclusively (Section 6.5), the estimate of background from accidental crossings (Section 6.6.1) and merged vertices (Section 6.6.2), the systematic uncertainty due to non-standard tracking and vertexing (Section 6.7), and the interpretation of the results (Section 6.10).

6.1 Motivation

If a long-lived, massive particle decays into multiple charged particles in the inner detector of ATLAS, a secondary vertex may be reconstructed. This secondary vertex will be characterized by a high invariant mass and a high track multiplicity. A high multiplicity of jets can be present in the final state from the cascade decay of beyond Standard-Model particles or if the new long-lived particle decays into quarks. Searches for displaced vertices are unique because the signature does not directly depend upon the nature of the LLP itself. The ability to detect a LLP with a displaced vertex does not change whether the LLP is electrically charged or neutral, or if it has hadronized into an R-hadron. This allows the search for displaced vertices to have sensitivity to a wide range of models predicting LLPs.

RPV SUSY models are chosen as the focus of this analysis. Searches for displaced vertices are important in the context of RPV SUSY because they probe the region of parameter space in which the LSP is likely to decay after traveling a discernible distance in the detector. This region of parameter space is often left uncovered by conventional searches for prompt SUSY particles, as shown in Figure 2.8.

A schematic of the RPV models is shown in Figure 6.1. In both models, SUSY particles are pair-produced and decay either directly or through a cascade to Standard Model quarks via the λ'' RPV coupling. The LSP in both models is the neutralino, which obtains a long lifetime from the smallness of the RPV coupling. The first model, called the *Strong RPV* model, includes pair-produced gluinos which decay promptly to neutralinos. The final state includes prompt jets which originate from the interaction point and displaced jets which originate from some other location in the detector. The prompt jets originate from the gluino decay and the displaced jets originate from the long-lived neutralino decay. The level of displacement of the jets depends on the distance the neutralinos travel from the interaction point before decaying. The second model, called the *electroweak (EWK) RPV* model, includes pair-produced long-lived electroweakinos (charginos and/or neutralinos). Compared to the strongly produced model, this physics process is characterized by fewer jets in the final state, all of which may originate from a displaced position.




Figure 6.1: The benchmark models used to guide analysis decisions and interpret the search results.

6.2 Analysis strategy

The benchmark models shown in Figure 6.1 are used to guide the design of this analysis. The sensitive variables used to distinguish displaced vertices originating from signal processes and from Standard Model processes or from instrumental/algorithmic effects are the invariant mass and track multiplicity of the vertex. The invariant mass of a vertex is calculated from the sum of the four-vectors of all of its associated tracks, assuming that each track has the charged pion mass. A wide range of beyond Standard Model theories can give rise to high mass, multi-track displaced vertices in multi-jet events [23, 46, 47]. Therefore, care is taken not to optimize the analysis for these benchmark scenarios only, but to keep the analysis as model-independent as possible.

There are no Standard Model processes which give rise to high mass displaced vertices. The backgrounds in this search stem from instrumental and algorithmic effects. The three main sources of background are:

1. **Hadronic interaction (HI):** Particles traveling through the detector may interact with the nuclei of detector material. If multiple charged particles are produced in this interaction, then it may be reconstructed as a displaced vertex. Due to the kinematics of this process, the charged particles resulting from the interaction are typically collimated, giving rise to relatively low-mass displaced vertices. However, there is a non-negligible component which may result in a high-mass displaced vertex.
2. **Accidental crossings (AX):** A low-mass displaced vertex (hadronic interaction or Standard Model long-lived particle decay) may be crossed by an unrelated track. If that track is mistakenly included in the vertex, then the mass of the vertex can be promoted to the signal region.
3. **Merged vertices (MV):** If two low-mass displaced vertices (hadronic interactions or Standard Model long-lived particle decays) occur in the vicinity of each other, the vertexing algorithm may mistakenly include them in the same vertex. This may promote the mass of the displaced vertex to the signal region.

Two signal regions are designed which aim to select signal-like events from the dataset while reducing the expected number of events from background processes. One signal region targets signal processes with many high- p_T jets in the final state (i.e. the Strong RPV model) and the other signal region targets signal processes with fewer high- p_T jets but with displaced jets in the final state (i.e. EWK RPV model). The event selection criteria in both signal regions are designed to reduce the level of background from all three sources combined to $O(1)$ background event. The remaining component of background is estimated with two independent data-driven methods. The first method estimates all sources of background inclusively and is used as the main method to interpret the results of the search. The second method estimates all sources of background separately using different data-driven techniques. This method is used as a cross check of the inclusive method.

Similar searches have been performed previously by the ATLAS and CMS collaborations at $\sqrt{s} = 7$, 8, and 13 TeV [107–111]. This analysis is the first search for high mass, multi-track displaced vertices in multijet events with the ATLAS detector using $\sqrt{s} = 13$ TeV pp collision data. Novel techniques for the estimation of background processes that can produce a displaced vertex signal have been developed and sensitivity to RPV SUSY models predicting the production of long-lived electroweakinos is achieved for the first time.

6.3 Data and simulated event samples

This analysis uses 139 fb^{-1} of proton-proton collision data collected by the ATLAS experiment from 2015-2018.

Simulated events are used to estimate signal yields and to aid in the estimation of sources of background. SUSY signal samples were generated to model the production of gluinos and chargino/neutralino pairs. The matrix element calculation is performed to leading order precision with MadGraph5_aMC@NLO 2.6.2 [62]. The events are interfaced to Pythia 8.212 [63] for parton showering and hadronisation. The matrix element calculation includes the emission of up to two additional partons for all signal samples.

In the Strong RPV model, the proton-proton interaction produces a pair of gluinos, each of which promptly decay with a 100% branching ratio to a neutralino and a Standard Model $q\bar{q}$ pair. The neutralino then decays to three Standard Model quarks/antiquarks via the RPV λ'' coupling. All other RPV couplings are assumed to be zero. This ensures that the model does not allow for processes which have not been observed, like proton decay [14]. The neutralino decays are simulated for either light-flavor quarks (up, down, charm, or strange), heavy-flavor quarks (strange, top, or bottom), or a fully democratic decay in which all six quark flavors are possible. The nominal cross-sections and their uncertainties are calculated at NNLO precision and derived using the PDF4LHC15 PDF set, following the recommendations of Ref. [112]. The cross section ranges from 8.87 fb ($\pm 15.94\%$) at $m_{\tilde{g}} = 1 \text{ TeV}$ to 0.0462 fb ($\pm 30.31\%$) at $m_{\tilde{g}} = 2.6 \text{ TeV}$.

In the EWK RPV model, the proton-proton interaction produces charginos and/or neutralinos. The possible pairs are $\tilde{\chi}_1^\pm \tilde{\chi}_1^\pm$, $\tilde{\chi}_2^0 \tilde{\chi}_1^0$, $\tilde{\chi}_1^\pm \tilde{\chi}_1^0$, and $\tilde{\chi}_1^\pm \tilde{\chi}_2^0$. They are produced via a pure higgsino state. The charginos and neutralinos are all nearly mass-degenerate, with a mass difference of 1 GeV between the lightest neutralino and others. All particles ($\tilde{\chi}_1^\pm$, $\tilde{\chi}_2^0$, and $\tilde{\chi}_1^0$) decay via the λ'' RPV coupling into light-flavor Standard Model quarks. The nominal production cross sections and their uncertainties are computed at NLO precision in a limit of mass-degenerate higgsinos $\tilde{\chi}_1^\pm$, $\tilde{\chi}_2^0$, and $\tilde{\chi}_1^0$ with all the other sparticles assumed to be heavy and decoupled [113, 114]. The cross section ranges from 16.8 pb ($\pm 3.5\%$) at a common chargino/neutralino mass of 100 GeV to 0.021 fb ($\pm 8.4\%$) at 1.7 TeV.

To study the background estimation techniques, simulated dijet events are used. The simulated dijet events are generated using Pythia 8.230 and incorporates matrix elements for the leading order

processes of dijet production [115].

6.4 Event selection

In this section, the event selection is described. Events are collected from the detector using a series of triggers. Events which pass additional offline selection, called a *filter*, are reconstructed using dedicated LLP reconstruction algorithms. The *signal region* (SR) selection is designed to select candidate signal events, i.e. events with many highly energetic jets which contain a high-mass displaced vertex with a high track multiplicity. Furthermore *control regions* (CR) and *validation regions* (VR) are defined which select events which are not signal-like. The events in the CR are used to study and estimate the background in the SR, while the events in the VR are used to test the performance of the background estimate technique. The CR and VR selects events which either do not have many highly energetic jets, or which contain a displaced vertex with low mass or low track multiplicity. The following sections describe the signal region selection. The control and validation regions are described in Section 6.5.

6.4.1 Trigger and filter

Events are triggered using triggers which require the presence of multiple (4-7) jets. The jet- p_T thresholds at the HLT range from 45-120 GeV, depending on the jet multiplicity and data-taking year.

This search makes use of the Large Radius Tracking algorithm, which increases the time needed to reconstruct each event by a factor of roughly 2.5 with 2015/16 pileup conditions [80]. Therefore it is too computationally expensive to run on all events collected with the multi-jet triggers. Therefore a selection referred to as a *filter* selects a subset of events to enter a special data processing stream in which the Large Radius Tracking algorithm is run. This analysis utilizes two filters which apply selections on uncalibrated jets. The first filter, called the *Trackless jet filter*, selects events with four or more jets with p_T ranging from 45-100 GeV. The exact p_T requirements for each jet multiplicity are shown in the Low- p_T filter column of Table 6.1. Furthermore, events are required to have at least one trackless jet with $p_T > 70$ GeV and $|\eta| < 2.5$ or at least two trackless jets with $p_T > 50$ GeV and $|\eta| < 2.5$. A jet is considered trackless if the sum of the transverse momenta of all (standard) tracks associated to the jet is less than 5 GeV. All tracks with $p_T > 500$ MeV are used in the sum.

The second filter, called the *High- p_T jet filter*, selects events with four or more jets with p_T ranging from 75-220 GeV. There is no selection on the presence of trackless jets. The jet- p_T thresholds are chosen to limit the overall rate of filtered events. The rate of the Trackless and High- p_T filters combined was approximately 5-10 Hz in the Run-2 dataset.

Jet multiplicity	Low- p_T		High- p_T		Trackless jets	
	Filter	Offline	Filter	Offline	Filter	Offline
1	-	-	-	-	70	78
2	-	-	-	-	50	56
4	120	137	220	250	-	-
5	85	101	170	195	-	-
6	70	83	100	116	-	-
7	45	55	75	90	-	-

Table 6.1: Jet multiplicity and jet- p_T selection for the Trackless jet and High- p_T jet parts of the filter. All p_T values are in units of GeV.

6.4.2 Object definitions

After events are collected with the trigger and filter, physics objects are reconstructed and calibrated according to Chapter 4 and must pass additional selection.

Jets are required to have $p_T > 20$ GeV and $|\eta| < 2.8$. Jets are required to pass a dedicated cleaning which reduces contamination from noise and non-collision backgrounds, while maintaining a high efficiency for reconstructing jets from LLP decays.

Electrons are required to have $p_T > 10$ GeV, $|\eta| < 2.47$, and to satisfy the ‘LooseAndBLayer’ requirement [116], which is based on a likelihood using measurements of shower shapes in the calorimeter and track properties in the inner detector as input variables.

Muon candidates are required to have $p_T > 10$ GeV and satisfy the ‘medium’ identification requirements defined in Ref. [117]. These requirements apply selection on the number of hits in the different inner detector subdetectors and muon spectrometer tracking stations, and the track fit properties. Furthermore, a selection is applied on the ratio of the charge and momentum (q/p) measured in the inner detector and muon spectrometer subsystems, in order to ensure compatibility between the measurements made in these different subsystems.

A single particle or collection of particles may be reconstructed as multiple objects in the ATLAS detector. To avoid double-counting, a procedure to remove overlapping objects is applied. First, electron candidates are removed if they share an inner detector track with a muon candidate. Next, jets within an angular distance of $\Delta R = \sqrt{(\Delta\eta)^2 + (\Delta\phi)^2} = 0.2$ of an electron candidate are discarded, in order to remove jets which originate from the calorimeter shower of electrons.

Track jets are used to aid in the background estimation (see Section 6.5). Track jets are required to have at least two tracks and to satisfy $p_T > 20$ GeV to be used in the inclusive estimate. They are not used in the overlap removal procedure.

6.4.3 Baseline jet selection

The baseline jet- p_T selection is defined based on the multiplicity and p_T of calibrated jets. The jet- p_T may change during calibration. The baseline jet- p_T selection are chosen so that jets passing this selection have a high probability of having passed the filter-level p_T requirements prior to calibration. This threshold is calculated by studying the rate at which jets pass the filter-level p_T requirements, as a function of their calibrated p_T . The pass rate for the filter thresholds at 45 GeV and 500 GeV are shown in Figure 6.2. The selections on the calibrated jet- p_T are chosen so that the pass rate reaches 98%. The calibrated jet- p_T selection on each jet multiplicity are listed in the columns labeled *Offline* in Table 6.1.




Figure 6.2: In these plots, the fraction of jets which pass the (a) 45 GeV and (b) 120 GeV uncalibrated jet- p_T requirement as a function of the calibrated jet- p_T is shown. The dashed line indicates the uncalibrated jet- p_T selection, and the two dotted lines indicate the point at which the pass rate reaches 95% and 98%, respectively.

The calibrated jet- p_T selection are used to define event-level baseline jet- p_T requirements which target each of the filters:

$$\text{High-}p_T \text{ jet selection} = \{ 4j_h \parallel 5j_h \parallel 6j_h \parallel 7j_h \}$$

and

$$\text{Trackless jet selection} = \{ (4j_l \parallel 5j_l \parallel 6j_l \parallel 7j_l) \& (1j_t \parallel 2j_t) \},$$

where j_h , j_l and j_t are the offline jet- p_T selection for Low- p_T , High- p_T , and trackless jets which are listed in Table 6.1. Furthermore, events passing the Trackless jet selection are required to fail the High- p_T jet selection. This ensures that the two jet selections yield statistically independent sets of events.

The High- p_T jet selection is over 95% efficient for the majority of the Strong RPV model parameter space of interest. The efficiency drops to 60% only for low $m_{\tilde{g}}$, low $m_{\tilde{\chi}_1^0}$, and long $\tau_{\tilde{\chi}_1^0}$. In this corner

of parameter space, the prompt jets in the final state have a softer p_T spectrum that fail the selection requirements, and the displaced decay products from the boosted neutralino decay are more likely to fall out of the acceptance of the detector. The Trackless jet selection achieves a maximum of 15% efficiency for the Strong RPV model for samples with long $\tau_{\tilde{\chi}_1^0}$.

The efficiency of the High- p_T jet selection for the RPV EWK signal model is shown in Figure 6.3(a). The efficiency drops quickly for EWK RPV model with $m_{\tilde{\chi}_1^0} < 1.1$ TeV, where the jets in the final state are no longer energetic enough to pass the High- p_T jet requirements. The Trackless jet selection recovers some of this efficiency at low $m_{\tilde{\chi}_1^0}$, as shown in Figure 6.3(b).




Figure 6.3: The (a) High- p_T and (b) Trackless jet selection efficiency for the EWK RPV model.

6.4.4 Displaced vertex preselection

To pass the preselection, displaced vertices must satisfy several criteria which are listed in Table 6.2 and described in this section. Displaced vertices must be reconstructed within $r < 300$ mm and $|z| < 300$ mm, where the reconstruction efficiency of the vertexing algorithm is the highest. The displaced vertex position must be at least 4 mm from any primary vertex in the event, which reduces background displaced vertices from Standard Model processes like heavy flavor decays. Displaced vertices must also be reconstructed with a high goodness-of-fit to reject displaced vertices resulting from the accidental crossing of random tracks.

Displaced vertices are also required not to be reconstructed in regions of material. Particles traveling through detector material may undergo a hadronic interaction, producing displaced vertices with large track multiplicities, which are hard to distinguish from displaced vertices from massive long-lived particle decays. A material map veto is designed to remove the majority of hadronic interactions. The material map is parameterized as a function of the position in the detector (R , ϕ , and z). The positional granularity of the map differs in the inner ($R < 150$ mm) and outer ($R > 150$ mm) regions. The bin

DV Selection	
$r_{\text{DV}} < 300 \text{ mm}, z_{\text{DV}} < 300 \text{ mm}$	
$r_{\text{DV-PV}} > 4 \text{ mm}$	
$\chi^2/N_{\text{DoF}} < 5$	
Pass strict material map veto	
$m_{\text{DV}} > 10 \text{ GeV}, N_{\text{trk}} \geq 5, N_{\text{trk}}^{\text{sel}} \geq 2$	

Table 6.2: Summary of the displaced vertex selections. The first four selection constitute the DV preselection. The last selection on m_{DV} , N_{trk} , and $N_{\text{trk}}^{\text{sel}}$ constitute the DV full selection.

sizes in the inner material map are $\Delta R_{xy} = 0.2 \text{ mm}$, $\Delta\phi = \pi/423$, and $\Delta z = 1 \text{ mm}$ and in the outer material map they are $\Delta R_{xy} = 0.8 \text{ mm}$, $\Delta\phi = \pi/217$, and $\Delta z = 2 \text{ mm}$. The map is constructed from the measured positions of low-mass vertices in data and the known positions of detector elements. To be used in the map construction, low-mass vertices must satisfy several selection criteria on the invariant mass and angle of outgoing tracks in order to remove contamination from known long-lived particle decays (i.e. $K_S^0 \rightarrow \pi^+ \pi^-$) and photon conversions.

The material map veto used in previous searches vetoes DVs if their position corresponds to a bin in the material map which has material. This veto, called the *loose material map veto*, removes $\approx 42\%$ of the fiducial volume. The positions of reconstructed displaced vertices which pass the preselection but fail the loose material map veto are shown Figure 6.4.




Figure 6.4: The positions in the (a) x-y plane and (b) r-z plane of reconstructed displaced vertices in data that pass the DV preselection but fail the loose material map veto. Taken from Ref. [109].

To reduce backgrounds from hadronic interactions further, an altered version of this material map is used in this analysis. In addition to checking the bin corresponding to the position of the displaced vertex for material, the altered material map veto also checks all surrounding bins in the R , ϕ , and z

Name	Selection	Region	DV track type
Transverse momentum	$p_T > 2 \text{ GeV}$	all	both
	$p_T > 3 \text{ GeV}$	outside beampipe	attached only
	$p_T > 4 \text{ GeV}$	outside last pixel layer	attached only
d_0 -significance $= d_0 /\sigma(d_0)$	$d_0\text{-sig} > 10$	within beampipe	both
	$d_0\text{-sig} > 15$	within last pixel layer	attached only
	$d_0\text{-sig} > 10$	outside last pixel layer	core only
Small angle, low- p_T	$\Delta\alpha(\text{track}, r_{\text{DV-PV}}) > 0.2$ or $p_T > 4 \text{ GeV}$	outside beampipe	both
Backwards-going tracks	$\Delta\alpha(\text{track}, r_{\text{DV-PV}}) < \pi/2$	all	attached only
Upstream hit veto	Cannot have hits on tracking layers at $r < r_{\text{DV}}$	all	both

Table 6.3: Summary of DV-track cleaning.

directions. If any of those bins has material in them, the DV is vetoed. This is referred to as the *strict material map veto*. The strict material map veto removes $\approx 48\%$ of the fiducial volume.

6.4.5 Displaced vertex full selection

Displaced vertices passing the full selection must have at least five tracks, two of which must be core tracks (see Section 4.4), and a reconstructed invariant mass greater than 10 GeV. In order to reduce the background DVs passing this criteria, only certain DV-tracks passing several selection criteria are used in the DV track multiplicity counting and invariant mass calculation. These selection are defined separately for core tracks and attached tracks and depend on the radial position of the DV with respect to the position of the beam pipe ($R = 25 \text{ mm}$) and the outermost pixel barrel layer ($R = 145 \text{ mm}$). The selection are listed in Table 6.3 and motivated in the following paragraphs.

The transverse momentum and d_0 -significance of DV-tracks in background DVs and reconstructed χ_1^0 decays are shown in Figure 6.5. Tracks included in background DVs tend to have lower p_T and a smaller d_0 -significance. The p_T selections are tighter for attached DV-tracks in the outermost radial regions, where there is a larger rate of fake tracks attached to background DVs. The d_0 -significance also differs for attached and core tracks in different radial regions. These selection were chosen to remove the most background while maintaining a high signal efficiency.

The distribution of the DV-track p_T and the 3D angle between the DV-track and the PV-DV vector ($\Delta\alpha$) for background and signal DVs reconstructed outside of the beampipe is shown in Figure 6.6. Tracks in background DVs tend to have small $\Delta\alpha$ and a low p_T . On the contrary, tracks in DVs from reconstructed χ_1^0 decays have a larger spread in p_T and $\Delta\alpha$. This difference is exploited by rejecting DV-tracks with $\Delta\alpha < 0.2$ and $p_T < 4 \text{ GeV}$.




Figure 6.5: The (a) transverse momentum and (b) d_0 -significance of DV-tracks in background DVs and reconstructed χ_1^0 decays in all radial regions. The background DVs are studied from simulated dijet events. The signal models shown have $m_{\tilde{g}} = 2$ TeV and include all lifetimes. All DV-track cleaning are applied except for the transverse momentum selection in (a) and the d_0 -significance selection in (b).




Figure 6.6: The DV-track p_T vs. $\Delta\alpha$ in (a) background DVs taken from simulated dijet events and (b) simulated Strong RPV signal events with $m_{\tilde{g}} = 2.0$ TeV and $m_{\tilde{\chi}_1^0} = 200$ GeV. All DV-track cleaning are applied except for any angular selection. DVs are required to be reconstructed at radii larger than the position of the beampipe. The region inside the red box is rejected by the DV-track cleaning.

Lastly, particles originating from real LLP decays are expected to travel in a similar direction as the LLP which they originated from (i.e. away from the interaction point). This is especially true for lighter, highly boosted LLPs. Two selection are applied to remove DV-tracks which are pointing backwards with respect to the DV position. The first removes DV-tracks which have a $\Delta\alpha < \pi/2$. This selection is only applied to attached tracks, as those tracks have a lower quality and were observed to more often be pointing backwards with respect to the background DV they were included in. The second requires that tracks cannot have hits on tracking layers which are situated at a smaller radii than the DV position.

Region	Selection
High-p_T signal region	Event passes the High- p_T jet selection and has at least one DV passing the full selection
Trackless signal region	Event passes the Trackless jet selection, fails the High- p_T jet selection, and has at least one DV passing the full selection

Table 6.4: A summary of signal region definitions.

The selection efficiency of the full DV selection for the EWK RPV model is shown in Figure 6.7. The selection achieves a maximum efficiency of 75% for samples with $\tau_{\tilde{\chi}_1^0} = 0.1$ ns. The loss of efficiency at shorter and longer lifetimes is due largely to the fiducial volume acceptance, which is also impacted by the boost (and therefore mass) of the $m_{\tilde{\chi}_1^0}$. A similar trend is observed for the Strong RPV model.




Figure 6.7: The selection efficiency of the full displaced vertex selection in the EWK RPV model. The efficiency is computed with respect to the total number of events passing the High- p_T or Trackless jet selections.

6.4.6 Signal region definitions

The High- p_T signal region selects events that pass the High- p_T jet selection and have at least one displaced vertex passing the full selection. The Trackless signal region selects events which fail the High- p_T jet selection, pass the Trackless jet selection, and have at least one displaced vertex passing the full selection. These selection are summarized in Table 6.4. The signal efficiency of each selection in the High- p_T SR and Trackless SR are shown in Figure 6.8.




Figure 6.8: The signal efficiency of several representative signal models in the (a) High- p_T and (b) Trackless signal regions. The selection are applied cumulatively.

6.5 Background estimation

In this chapter, the method used to estimate all sources of background inclusively is presented. This method is used to interpret the results of the search.

The production of high mass displaced vertices from background sources (hadronic interactions, merged vertices, accidental crossings) is correlated to the presence of track jets. One can utilise this correlation to estimate the expected number of events in a multi-jet final state which contain a high mass displaced vertex from a background source. This estimation technique works in two steps. First, the probability that a SR-like DV is produced in proximity to a track jet, called the *jet-DV probability*, is calculated using track jets and DVs in data events which contain a lower jet multiplicity than the signal region (called the control region). Then, the jet-DV probability is applied to track jets in events passing the baseline jet- p_T selection to estimate the total number of expected events with a background displaced vertex passing all signal region selection.

The control region (CR) is a collection of events which are used to measure the correlations between jets and DVs. It is populated with events which pass a single photon trigger with a $p_T > 140$ GeV threshold, fail both signal region baseline jet selections, and has at least three track jets with $p_T > 20$ GeV. The single photon trigger is selected because it is not expected to significantly alter the composition of jets in the control region with respect to the signal region.

6.5.1 Study of jet-DV correlations

The assumption that the production of track jets and background displaced vertices are correlated is checked by studying events in the control region. The track jet multiplicity, track jet- p_T , and number of tracks per jet in the CR and in events passing the baseline jet- p_T selections are shown in Figure 6.9. The control region by design has on average fewer track jets than the events passing the baseline jet

selection. The events passing the High- p_T jet selection have track jets with higher p_T and more tracks than those in the control region. Conversely, events passing the Trackless jet selection have track jets with slightly lower p_T and fewer tracks than those in the control region. The jet-DV probability is parameterized as a function of both the track jet- p_T and $N_{\text{trk}}^{\text{jet}}$ in order to take into account these differences.




Figure 6.9: A comparison of the track jet (a) multiplicity, (b) p_T , and (c) n_{Tracks} in the control region and in events passing the High- p_T and Trackless jet selection. All distributions utilize data from the full Run-2 dataset and are normalised to unity.

The fraction of events in the control region which contain a background DV as a function of track jet multiplicity is shown in Figure 6.10(a). There is a clear trend that events with more track jets are more likely to contain a SR-like DV.

The proximity of a DV and a track jet is calculated with

$$\Delta R_{\text{DV,jet}} = \sqrt{\Delta\eta^2 + \Delta\phi^2}, \quad (6.1)$$

where $\Delta\eta$ and $\Delta\phi$ are measured between the PV-DV vector and the jet axis. The average $\Delta R_{\text{DV,jet}}$ as a function of the DV track multiplicity is shown in Figure 6.10(b). DVs with a higher track multiplicity tend to be produced in closer proximity to track jets, which is an indication that DVs with a high track multiplicity have a stronger correlation to track jet activity.




Figure 6.10: In (a), the fraction of events in the CR with a preselected DV satisfying $m_{\text{DV}} > 3 \text{ GeV}$ and $n_{\text{Tracks}} > 2$ is shown. In (b), the average distance of preselected DVs with $m_{\text{DV}} > 2 \text{ GeV}$ to the closest track jet in the event is shown, as a function of n_{Tracks} .

The proximity of background DVs to track jets in events from the control region is also checked. Background DVs in the CR are matched to the closest jet in the event using the $\Delta R_{\text{DV,jet}}$ parameter. Each DV is matched to only one jet in each event, and the DV position is not required to be within the jet cone ($\Delta R = 0.4$). The ΔR between DVs and the closest track jet for various DV masses and track multiplicities are shown in Figure 6.11(a) and Figure 6.11(b), respectively. These distributions are characterized by a peak at low ΔR and a long tail extending to larger ΔR values. A large fraction of DVs (about 70-80%) are within the $\Delta R = 0.4$ cone of the nearest track jet, which further supports the assumption that the background DVs are correlated to the presence of track jets.

6.5.2 Calculation of jet-DV probability

The probability that a DV is produced given the presence of a jet is calculated in the CR as:

$$\text{Prob}(\text{DV} | \text{track jet}) = \frac{\text{Number of track jets with a matched DV}}{\text{Total number of track jets}}, \quad (6.2)$$

where the probability may be parameterized as a function of jet properties (i.e. p_{T} , η , $N_{\text{trk}}^{\text{jet}}$, etc). The jet-DV probability is shown as a function of the number of tracks in the jet and track jet- p_{T} in Figures 6.12-6.13. High mass DVs with many tracks are more likely to be produced in the vicinity of a track jet which has a high track multiplicity and moderate p_{T} . There are some shape differences




Figure 6.11: These figures display the ΔR between DVs and the nearest track jet in the CR using the full Run-2 dataset, as a function of various DV properties. In (a), the ΔR of DVs with $n_{\text{Tracks}} > 3$ and varied mass is shown. In (b), the ΔR of DVs with $m_{\text{DV}} > 2$ GeV and various track multiplicities are shown.

in the jet-DV probabilities for DVs with different masses and track multiplicities; this demonstrates the need to derive separate jet-DV probability distributions for DVs with different masses and track multiplicities. For example, in order to estimate the contribution of 4-track background DVs in events passing the SR baseline jet requirements, the correlation between jets and only 4-track DVs in the CR should be used.




Figure 6.12: Jet-DV probability as a function of the number of tracks in the jet, for DVs of various (a) masses and (b) track multiplicities, calculated using events in the CR with the full Run-2 dataset. Each distribution is normalised so that the shapes can be compared. The last bin includes the overflow. In (a), the $4 < m < 6$ GeV and $6 < m < 8$ GeV mass windows are merged in order to allow for finer binning in $N_{\text{trk}}^{\text{jet}}$.

The jet-DV probability used for the SR background estimate is parameterized as a function of track jet- p_T and $N_{\text{trk}}^{\text{jet}}$. The jet-DV probability is calculated using all DVs in the CR satisfying $n_{\text{Tracks}}^{\text{DV}} > 3$ and




Figure 6.13: Jet-DV probability as a function of the jet- p_T , for DVs of various (a) masses and (b) track multiplicities, calculated using events in the CR with the full Run-2 dataset. Each distribution is normalised so that the shapes can be compared. The last bin includes the overflow.

$m_{DV} > 5$ GeV. The DV mass and track multiplicity requirements in the CR are looser than the SR-DV requirements ($n_{Tracks}^{DV} > 4$ and $m_{DV} > 10$ GeV) so that there are adequate statistics in the CR to calculate the jet-DV probability. The number of DVs in the CR matched to jet is shown in Figure 6.14(a). The jet-DV probability shown in Figure 6.14(c) is calculated by dividing this distribution by the total number of jets in the CR, shown in Figure 6.14(b). The jet-DV probability is very small ($O(10^{-6})$), owing to the rarity of SR-like DVs and the high number of jets produced in the ATLAS detector.

6.5.3 Application of jet-DV probability

To estimate the number of background DVs expected in the High- p_T (Trackless) SR, all track jets in events passing the High- p_T (Trackless) jet selection are assigned a weight ($P(DV | jet)$) equal to the corresponding jet-DV probability in Figure 6.14(c). The sum of weighted jets gives the estimated number of events with a SR-like DV, i.e.

$$N_{Bkg} = f \cdot \sum_{i=0}^{N_{track\ jet}} P(DV | jet_i), \quad (6.3)$$

where f is an additional weight which takes into account the difference in the DV mass and track multiplicity requirements used to calculate the jet-DV probability and in the SR selection. It is calculated by taking the ratio of the number of events with a SR-like DV in the control region:

$$f = \frac{N_{Event}(m > 10 \text{ GeV}, N_{track} > 4)}{N_{Event}(m > 5 \text{ GeV}, N_{track} > 3)} \text{ in CR} = 3/75 = 0.040. \quad (6.4)$$

The jet-DV probabilities may be applied in this way because the probability that more than one




Figure 6.14: The jet-DV probability of SR-like DVs satisfying $n_{\text{Tracks}} > 3$ and $m_{\text{DV}} > 5$ GeV in the CR.

SR-like DV is produced in the same event is negligible.

The number of track jets in the same p_T and $N_{\text{track}}^{\text{jet}}$ bins used in the jet-DV probability distribution in events passing the High- p_T (Trackless) jet selection is shown in Figure 6.15(a) (6.16(a)). The sum of weighted jets in events passing the High- p_T (Trackless) jet selection is shown in Figure 6.15(b) (6.16(b)). The sum of the bin contents in Figures 6.15(b) and 6.16(b) give the estimated number of events with a background DV in the signal regions. This method predicts 0.49 (0.69) background events in the High- p_T (Trackless) signal region.




Figure 6.15: These plots include events passing the **High- p_T** jet selection. In (a), the number of track jets is shown, while in (b) the sum of weighted jets (see text for details) is shown. The sum over all bins in (b) gives the inclusive background estimate in the High- p_T SR.




Figure 6.16: These plots include events passing the **Trackless** jet selection. In (a), the number of track jets is shown, while in (b) the sum of weighted jets (see text for details) is shown. The sum over all bins in (b) gives the inclusive background estimate in the Trackless SR.

6.5.4 Uncertainties

Several sources of uncertainty are quantified for this background estimate, which are listed in Table 6.5. First, the statistical uncertainty on the number of events with a SR-like DV in the CR is propagated to the final estimate. A toy method is used which performs $N = 5000$ pseudo-experiments. The number of events in the CR with a SR-like DV (see Figure 6.14(a)) is varied bin-by-bin within its statistical uncertainty and used to calculate a new jet-DV probability distribution. The standard deviation of the background estimate from all statistically varied toy experiments is taken as an uncertainty on the final background estimate. The number of pseudo-experiments was chosen so that the resulting uncertainty

using this procedure varies by less than one percent.

Region	Estimated	Uncertainty				Total
		CR statistical	Non-linearity	Pileup	Non-closure	
High-pt SR	0.49 ± 0.28	58%	0.10 %	5.9%	4%	59%
Trackless SR	0.69 ± 0.42	58%	3.8 %	20%	-	61%

Table 6.5: Estimate of background in the SR with systematic uncertainties. The uncertainties are assumed to be uncorrelated and are added in quadrature to provide the overall uncertainty on the estimate in the second column.

An uncertainty is assigned to take into account any residual dependence of the jet-DV probability on the amount of jet activity in the event. One of the main assumptions of this method is that the jet-DV probability can be decoupled from the jet multiplicity in the event. This allows for the jet-DV probability to be calculated in events with fewer jets (the control region) and applied to events with more jets (those which pass the baseline jet- p_T selection). To test this assumption, the fraction of jets in the CR which are matched to a SR-like DV is calculated as a function of the number of track jets in the event. This distribution is shown in Figure 6.17(a). This distribution is multiplied bin-by-bin with the number of track jets in the CR and in events passing the baseline jet- p_T selection. The resulting histograms are integrated over all bins (N), and the difference between N_{CR} and $N_{\text{High-}p_T \text{ jet selection}}$ ($N_{\text{Trackless jet selection}}$) gives the uncertainty assigned to the background estimate in the High-pt SR (Trackless SR).

An uncertainty is applied to account for differences in the pileup distribution of events in the control region and events passing the High- p_T and Trackless jet selections. The pileup distributions are shown in Figure 6.17(b). To quantify this uncertainty, the events in the control region are reweighted so that their pileup distribution matches the distribution in events passing the High- p_T or Trackless jet selections. The weighted events are used to calculate new jet-DV probabilities, which are then used to estimate the background. The difference between the nominal and the reweighted background predictions is taken as an uncertainty. This procedure results in an additional 5.9% (20%) uncertainty on the background prediction in the High- p_T (Trackless) signal regions.

Lastly, an uncertainty is applied to account for the non-closure of the estimate in some of the validation regions. This method underestimates the number of DVs reconstructed inside of material in events which pass the High- p_T jet selection. The validation regions in which this non-closure is observed are populated with background DVs from hadronic interactions, which is expected to be a small component of the background in the signal region. To calculate an uncertainty to account for this, the non-closure in the validation regions is scaled by the expected contribution of hadronic interaction in the signal region. From this procedure, an uncertainty of 4% is assigned to the estimate of background in the High- p_T signal region, and no additional uncertainty is assigned to the Trackless signal region. The estimation of this uncertainty is described in Section 6.5.5.




Figure 6.17: In (a), the fraction of jets in the CR which are matched to a SR-like DV ($n_{\text{Tracks}} > 3$ and $m_{\text{DV}} > 5$ GeV) is shown. In (b), the pileup distributions in the control region and in event passing the High- p_T and Trackless jet selections is shown.

6.5.5 Validation

Validation regions (VR) are used to test the performance of the background estimation technique. The validation regions used in this analysis are:

1. Low- $n_{\text{Tracks}}^{\text{DV}}$ and low- m_{DV} sideband regions of the signal region in data
2. Inside of material regions in data
3. High- $n_{\text{Tracks}}^{\text{DV}}$ and high- m_{DV} regions in simulated dijet events

Validation in SR sideband regions:

First, the inclusive background estimation technique is tested in low- m_{DV} and low- $n_{\text{Tracks}}^{\text{DV}}$ sideband regions of the High- p_T SR and Trackless SR in data. These sideband validation regions are populated with events which satisfy all SR selection, except they contain a reconstructed DV with fewer than five tracks or a mass less than 10 GeV. In each validation region, a separate jet-DV probability is calculated from the CR. These jet-DV probabilities are used to estimate the background and uncertainties in each validation region separately.

There are few DVs in the CR available to calculate the jet-DV probabilities in the validation regions with $n_{\text{Tracks}}^{\text{DV}} = 4$, $m_{\text{DV}} > 10$ GeV and $n_{\text{Tracks}}^{\text{DV}} > 4$, $5 < m_{\text{DV}} < 10$ GeV. Therefore, the m_{DV} and/or $n_{\text{Tracks}}^{\text{DV}}$ requirements are slightly loosened to calculate the jet-DV probability, and the final background estimate is weighted down accordingly, similar to the procedure performed for the SR estimate.

The results of the validation test in the sideband regions of the High- p_T SR and Trackless SR are shown in Figure 6.18 and Figure 6.19, respectively. There is good agreement between the estimated

event counts and observed data in all validation regions.




Figure 6.18: Validation of background estimate in various sideband regions of the **High- p_T SR**. In (a), the regions contain events containing 4-track DVs with various mass requirements. In (b), the regions contain events with > 4 -track DVs with $m_{DV} = 5\text{-}10 \text{ GeV}$. The rightmost bin in (b) shows background estimate in the SR ($n_{\text{Tracks}} \geq 5$ and $m_{DV} > 10 \text{ GeV}$).




Figure 6.19: Validation of background estimate in various sideband regions of the **Trackless SR**. In (a), the regions contain events containing 4-trk DVs with various mass requirements. In (b), the regions contain events with > 4 -track DVs with $m_{DV} = 5\text{-}10 \text{ GeV}$. The rightmost bin in (b) shows background estimate in the SR ($n_{\text{Tracks}} \geq 5$ and $m_{DV} > 10 \text{ GeV}$).

Inside material regions:

Next, the inclusive estimate is used to estimate the number of background DVs which are reconstructed inside of material in data. The background inside of material is larger than outside of material, which allows for the use of higher- m_{DV} regions as validation regions. DVs are required to pass the full DV

selection, with the exception that they fail the loose material map veto. The jet-DV probability is recalculated in each $n_{\text{Tracks}}^{\text{DV}}$ and m_{DV} window using DVs in the control region which are reconstructed inside of material regions. The results of this validation test are shown in Figure 6.20. In events passing the High- p_T jet selection, there is a significant underestimation of the background in the lower m_{DV} windows for both 4-trk and >4-trk DVs. There is good agreement between the observed and expected background in the higher m_{DV} mass windows. The lower- m_{DV} mass windows inside material regions are populated with dominantly hadronic interactions. This validation test indicates that the inclusive background estimate does not predict as well the background originating from hadronic interactions. This is not an issue for the background prediction, because as will be shown in Section 6.6.4, the background from hadronic interactions is expected to constitute a very small fraction of the DVs in the signal regions.




Figure 6.20: Validation of the inclusive background estimate **inside of material** regions in data. DVs are required to pass the full DV selection with the exception that they fail the loose material map veto. Only the statistical uncertainties propagated from the CR are drawn.

An uncertainty is assigned to the estimate of background in the signal region to account for the

non-closure in the inside material validation regions. The uncertainty is calculated by taking the largest non-closure observed in a validation region inside of material in data and scaling it by the fraction of the total background coming from hadronic interactions which is expected in the signal region. This fraction is estimated from the separate data-driven background estimates, described in Section 6.6. In the High- p_T signal region, the largest non-closure is observed for 4-trk DVs with $m_{DV} = 10 - 15$ GeV inside of material (see Figure 6.20). In this validation region, the background is underestimated by $\approx 400\%$. The fraction of the total background coming from hadronic interactions in the High- p_T signal region is estimated to be $\approx 0.13\%$ (see Table 6.10). When varying the hadronic interactions estimate and total background estimate within their systematic uncertainties, the maximum fraction of hadronic interactions achieved is $\approx 0.35\%$. To produce a conservative uncertainty, this fraction is rounded up to 1% . Therefore, an uncertainty of $400 \times 0.01 = 4\%$ is assigned to the background estimate in the High- p_T SR. The non-closure in the inside material validation regions in events passing the Trackless jet selections is much smaller than in events passing the High- p_T jet selection. A similar uncertainty calculated for the Trackless SR is negligible and not included.

Closure test in simulated dijet events:

The inclusive estimate is also tested using simulated dijet events. The jet-DV probability is derived from simulated dijet events which fail all of the the multijet triggers used in this analysis. The jet-DV probabilities are used to estimate the background in events which pass at least one multijet trigger used in this analysis.

The results of this validation test are shown in Figure 6.21. There is good agreement between the expected and observed background in all $n_{Tracks}^{DV} = 4$ validation regions. There is a consistent overestimation of the background by 20-40% in the $m_{DV} = [5, 10]$ GeV validation regions. This behavior is not observed in the same m_{DV} and n_{Tracks}^{DV} regions in data and the underestimation is not significant in the high- m_{DV} regions. Therefore, an additional uncertainty is not applied to the estimate of background in the signal region.

6.6 Alternative background estimation

An additional method is used as a cross check of the inclusive background estimate. In this method, all three sources of background are estimated individually with dedicated data-driven techniques. Each estimation technique involves first building a mass template which describes the shape of m_{DV} from a specific background source. Then the mass templates are individually normalised before being combined into a complete background estimate.




Figure 6.21: Validation of background estimate in various regions in **simulated dijet** events, with a looser control and validation region event selection. Only the statistical uncertainties propagated from the CR are drawn.

6.6.1 Accidental crossings

A low-mass, displaced vertex can be promoted to a higher mass if an unrelated track is attached to it during reconstruction. If the unrelated track has a high p_T , or crosses the DV with a large angle with respect to the other tracks in the DV, then the mass of this DV with the unrelated track may be promoted to the signal region ($m_{DV} > 10$ GeV). This source of background is referred to as the *accidental crossing* background.

The background from accidental crossings is estimated in four steps:

1. **Produce track database:** The *track database* is a collection of tracks which have been incorrectly attached to DVs in data. Accidental crossing tracks are identified using displaced $K_S^0 \rightarrow \pi^+ \pi^-$. To identify accidental crossing tracks, 3-track DVs where two of the tracks can be combined to form a vertex consistent with the K_S^0 mass are identified. The third track in the vertex is considered to be an accidental crossing track. The properties and orientation of the accidental crossing track with respect to the DV position are saved and it is added to the track database.
2. **Build mass template:** The shape of the mass distribution of accidental crossing DVs is estimated separately for each DV track multiplicity. The $(n+1)$ -track mass distribution is generated by attaching one additional track from the database to n -track DVs in data. The $(n+1)$ -track DV mass is recalculated and used to build the mass template.
3. **Normalise mass template:** The *crossing factor* measures the probability that a DV is reconstructed with an unrelated track. It is used to normalise the mass templates. The crossing factor

Region name	Radial boundary
Inside BP	$R_{xy} < 25$ mm
Inside IBL	$25 \leq R_{xy} < 38$ mm
Inside L2	$38 \leq R_{xy} < 120$ mm
Inside SCT	$120 \leq R_{xy} < 300$ mm

Table 6.6: The estimate of accidental crossings is estimated separately in the above radial regions.

is calculated by measuring the fraction of all $K_S^0 \rightarrow \pi^+ \pi^-$ decays reconstructed in data which have an accidentally crossing track.

4. **Integrate:** The final estimate of accidental crossing DVs in the signal region is calculated by integrating all $(n+1)$ -track mass templates with $n \geq 4$ in the $m_{DV} > 10$ GeV region.

The behavior of the accidental crossings background is expected to differ depending on the position of the DV in the fiducial volume. This is due to multiple factors, for example the vertexing reconstruction efficiency and the larger incident angles of tracks crossing DVs at large radii. Therefore the above procedure is performed separately in different radial regions, which are determined by the position of the beampipe (BP), the IBL, the second pixel tracking layer (L2), and the innermost SCT tracking layer. The radial region definitions are listed in Table 6.6. This means that the full estimate of accidental crossings is comprised of four $(n+1)$ -track templates for each DV track multiplicity.

Track database

Tracks which have been incorrectly attached to displaced vertices are identified in data using reconstructed $K_S^0 \rightarrow \pi^+ \pi^-$ decays. The K_S^0 has an invariant mass of 498 MeV, a mean lifetime of 90 ps, and it decays with a 70% branching ratio via the $K_S^0 \rightarrow \pi^+ \pi^-$ channel. This decay channel is important because it is fully reconstructable, meaning that the K_S^0 decay can be reconstructed as a 2-track DV with an invariant mass close to the K_S^0 mass. If a random track is incorrectly attached to the K_S^0 decay vertex, then a 3-track DV is reconstructed and two of the tracks have an invariant mass close to the K_S^0 mass.

To produce the track database, all 3-track DVs in data which pass the DV preselection and which are present in events passing the baseline jet selection are identified. The DV-track cleaning is not applied when calculating the DV track multiplicity. The mass of each pair of tracks in the 3-track DV is calculated (m_{2-trk}). If m_{2-trk} of any pair of tracks is within 50 MeV of the K_S^0 mass, then the DV is considered to be a K_S^0 decay with an accidental crossing track. If this accidental crossing track passes the DV-track cleaning, it is added to the track database. If there is more than one accidental crossing track candidate in a single DV, then the pair of tracks with an invariant mass closest to the K_S^0 mass is identified and the third track is added to the database. The radial and z position of the DV which the AX track is attached to is also recorded in the track database.

Next, the properties of the accidental crossing tracks with respect to the DV they are attached to is calculated. The p_T , η and crossing angle $\Delta\phi$ of the track are stored together in the track database so that any correlations between them are retained. The crossing angle $\Delta\phi$ is defined as the angle between the accidental crossing track at the closest approach to the DV and the PV-DV pointing vector. This angle is important to preserve because the higher the crossing angle, the larger the invariant mass of the vertex with the addition of that track.

The properties of the accidental crossing tracks in all radial regions combined are shown in Figure 6.22. The accidental crossing tracks are characterized by a smoothly falling p_T spectrum and a wide $|\eta|$ distribution. The majority of AX tracks have a small opening angle with respect to the PV-DV pointing vector ($\Delta\phi$). The accidental crossing tracks with higher opening angles, $\Delta\phi \approx \pm\pi/2$, occur mostly in the outermost radial region.




Figure 6.22: Distributions of the accidental crossing track (a) p_T , (b) η and (c) $\Delta\phi$ in all radial regions. Separate track databases are built for events passing the High- p_T and Trackless jet selections.

Mass template

Templates of the mass distribution of DVs with an accidental crossing track are constructed from data. To construct the mass template for $(n+1)$ -track DVs, one random track from the track database is attached to each n -track DV in data. All n -track DVs in events passing the baseline jet selection and which pass the DV preselection are used in the construction of the template.

Tracks in the database are randomly sampled from the same radial and z region which the n -track DV is located. The fiducial region is divided into eight z regions for this purpose. When attaching an accidental crossing track to an n -track DV, the track η and p_T are taken directly from the database. The track ϕ is recalculated using the PV-DV pointing vector, $\phi_{\text{DV-PV}}$, so that the crossing angle is preserved, i.e:

$$\phi_{\text{new}} = \phi_{\text{DV-PV}} - \Delta\phi. \quad (6.5)$$

These three variables, along with the charged pion mass assumption, define the accidental crossing track's four-vector, which is used to calculate the mass of the $(n+1)$ -track DV.

DVs in the SR are required to have at least two core tracks which pass the DV-track cleaning. Therefore, if the n -track DV in data has only one core track, then a core track from the AX track database is attached to it. DVs without any core tracks are not used in the creation of the mass template.

The shapes of the 4-track accidental crossings mass templates for each radial region are compared in Figure 6.23. The shape of the template changes in different radial regions, which motivates the need for radial segmentation of the estimate. The mass templates shift to higher masses as the radial distance from the interaction point increases. This can be attributed to a higher incident angle of accidental crossing tracks as they cross n -track DVs at larger radii, which results in a higher $(n+1)$ -track DV mass.




Figure 6.23: Accidental crossing mass templates for 4-track DVs in the (a) High- p_T and (b) Trackless signal regions, split into radial regions. Each template is an average of $N = 4$ templates and is normalized to unity. Statistical error bars are drawn.

Crossing factor

The mass templates are normalised using the crossing factor, which is the fraction of $K_S^0 \rightarrow \pi^+ \pi^-$ decays which have been reconstructed with an unrelated track. $K_S^0 \rightarrow \pi^+ \pi^-$ decays are identified in a similar way as in the construction of the track database. To increase the number of K_S^0 decays available in this study, the crossing factor is calculated inclusively for both signal regions, using all events which pass any of the multi-jet triggers used in this analysis. Furthermore, both $K_S^0 \rightarrow \pi^+ \pi^-$ candidate tracks (in both 2-track and 3-track DVs) must be core tracks and at least one of them needs to satisfy $d_0 > 2$ mm. These selection ensure that only K_S^0 decays which would have been reconstructed without the accidental crossing track are used in this study.

The number of $K_S^0 \rightarrow \pi^+ \pi^-$ decays reconstructed without an accidental crossing track ($N_{\text{without-AX}}$) is measured by fitting the 2-track DV mass spectrum to a Gaussian function and a linear background in the $[0.4, 0.6]$ GeV mass window. The fitted background is subtracted from the mass distribution which is then integrated in the $[0.46, 0.54]$ GeV mass range. The number of $K_S^0 \rightarrow \pi^+ \pi^-$ decays reconstructed with an accidental crossing track ($N_{\text{with-AX}}$) is measured by fitting the mass distribution of all possible pairs of tracks in 3-track DVs with the same parameterisation. The mass of track pairs is only included in the distribution if the third track passes the DV-track cleaning.

In regions where the crossing factor is small, the K_S^0 mass peak in 3-track DVs is small compared to the level of combinatorial background. An example of this is shown in Figure 6.24(a) for DVs in the `insideBP` and `insideIBL` regions combined. This large background below the K_S^0 mass peak is not desirable, as it will lead to a larger uncertainty on the measurement of the crossing factor. To reduce the contamination from non- K_S^0 background, a constraint on the pointing vector of the candidate pairs of K_S^0 decay tracks is applied. The pointing vector is defined as $\cos(\Delta\phi(\vec{p}_{\text{tracks}}, \vec{r}_{\text{DV-PV}}))$, where $\vec{p}_{\text{tracks}}$ is the vectorial sum of four-momenta of pairs of tracks in the 2-track or 3-track DVs, and $\vec{r}_{\text{DV-PV}}$ is the PV-DV pointing vector. In Figure 6.25, the pointing vector of pairs of tracks with a mass consistent with the K_S^0 mass is compared with background-enriched distributions built from pairs of tracks with a mass above or below the K_S^0 mass. In 2-track DVs (Figure 6.25(a)), the K_S^0 enriched distribution peaks at higher $\cos(\Delta\phi(\vec{p}_{\text{tracks}}, \vec{r}_{\text{DV-PV}}))$ compared to the background enriched distribution. A similar trend is observed in 3-track DVs (Figure 6.25(b)), however the difference between the distributions is smaller, due to the larger contamination of background in the signal enriched distribution. To reduce the non- K_S^0 background, pairs of candidate K_S^0 decay tracks are required to satisfy

$$\cos(\Delta\phi(\vec{p}_{\text{tracks}}, \vec{r}_{\text{DV-PV}})) > 0.9999 . \quad (6.6)$$

After applying this selection, the K_S^0 mass peak is more clearly resolved, as shown in Figure 6.24(b).

The fitted mass distribution for 2-track and 3-track DVs in the outermost radial region is shown in Figure 6.26. The 2-track mass distribution is characterized by a distinct peak at the K_S^0 mass surrounded by a low level of background from non- K_S^0 decays. The kinematic cutoff at 280 MeV




 Figure 6.24: Impact of pointing vector selection on K_S^0 mass peak.


Figure 6.25: The pointing vector of K_S^0 decay candidate tracks in (a) 2-track and (b) 3-track DVs. The black distribution is built from pairs of tracks with a mass consistent with the K_S^0 mass, and the blue and red distributions are selected from pairs of tracks with a mass above or below the K_S^0 mass peak. The grey dashed line indicates the selection used to reduce the non- K_S^0 background.

arises due to the charged pion mass assumption used in the track reconstruction. The mass distribution of pairs of tracks in 3-track DVs is characterized by a peak at the K_S^0 mass surrounded by moderate levels of combinatorial background and other sources of 3-track DV background.

The crossing factor is calculated with:

$$\text{Crossing factor} = \frac{N_{\text{with-AX}}}{N_{\text{with-AX}} + N_{\text{without-AX}}} . \quad (6.7)$$

In each radial region, two crossing factors are calculated. The Core AX crossing factors are calculated using only K_S^0 decays reconstructed with an AX track that was in the core vertex. The All




Figure 6.26: In (a), the mass distribution of 2-track DVs in data is shown in the inside SCT region. In (b), the mass distribution of 2-track permutations of 3-track vertices is shown in the inside SCT region. The mass distributions are fit to a Gaussian (blue) and a linear background (red). The total number of entries in the peak after the fitted background is subtracted is printed on each histogram.

AX crossing factors do not require any additional selection on the accidental crossing track.

The crossing factors in each radial region are shown in Figure 6.27 and listed in Table 6.7. There are too few K_S^0 's reconstructed with an accidental crossing track in the insideBP and insideIBL regions to calculate the crossing factor there separately. Therefore the regions are merged for the crossing factor calculation. The crossing factor tends to increase with increasing radial position of the DV. The largest crossing factor occurs in the outermost radial region, which is attributed to the poorer vertex position resolution after the last pixel detector layer.




Figure 6.27: Crossing factors calculated in each radial region. The error bars are calculated from varying the subtraction of the non- K_S^0 background from the fitted peak results.

The crossing factors are used to weight each entry in the mass templates. The All AX (Core AX)

Radial region	Crossing Factors	
	All AX	Selected AX
Inside BP	$(1.284 \pm 0.095) \times 10^{-4}$	$(8.26 \pm 0.76) \times 10^{-5}$
Inside IBL	$(1.284 \pm 0.095) \times 10^{-4}$	$(8.26 \pm 0.76) \times 10^{-5}$
Inside L2	$(5.11 \pm 0.12) \times 10^{-4}$	$(4.29 \pm 0.11) \times 10^{-4}$
Inside SCT	$(2.910 \pm 0.049) \times 10^{-3}$	$(1.201 \pm 0.031) \times 10^{-3}$

Table 6.7: Crossing factors calculated in each radial region. The statistical errors are listed. The same crossing factor is used for the insideBP and insideIBL regions.

crossing factors are used to weight the mass of $n+1$ -track DVs when the n -track DV has at least (less than) two core tracks passing the DV-track cleaning. The weighted (3+1)-track, (4+1)-track, and ($\geq 5+1$)-track mass templates in all radial regions are shown in Figure 6.28. The estimate of background DVs in the signal region with an accidental crossing track is calculated by integrating these templates in the $m_{\text{DV}} > 10$ GeV region.

Uncertainties

To assign a systematic uncertainty to this background estimate, several assumptions are tested. One such assumption is that the properties and rate of accidental crossing tracks are independent of the properties of the DV that is being crossed. This allows for the rate and track database to be built from AX tracks crossing K_S^0 's, and applied to other types of background displaced vertices, for example hadronic interactions. To test this assumption, the crossing factor is calculated using $\Lambda^0 \rightarrow p\pi$ decays. The kinematics of $\Lambda^0 \rightarrow p\pi$ decays differ from $K_S^0 \rightarrow \pi^+\pi^-$ due to the fact that Λ^0 's are heavier, and they decay into particles with different masses. This makes Λ^0 's a good candidate to use as a cross-check that the AX estimate does not depend on the Standard Model LLP used to build it.

Displaced vertices originating from $\Lambda^0 \rightarrow p\pi$ decays in data are identified using the 2-track and 3-track DV mass spectrum in data in a similar procedure to $K_S^0 \rightarrow \pi^+\pi^-$ decays. The candidate Λ^0 decay track with the higher p_T is assigned the proton mass and the lower p_T track is assigned the charged pion mass. The crossing factor calculated with $\Lambda^0 \rightarrow p\pi$ decays is found to be 18% higher than the K_S^0 crossing factor in the $R = 20 - 120$ mm region, and 20% lower than the K_S^0 crossing factor in the $R = 120 - 300$ mm region. Therefore, a flat systematic uncertainty of 25% is assigned to the estimate of accidental crossings in both the High- p_T and Trackless signal regions. This value is chosen to be conservative and to cover other sub-dominant effects such as the dependence of the crossing factor on the DV-track cleaning (applied to candidate K_S^0 decay tracks), the event-level selection, and the radial and $|z|$ segmentation.




Figure 6.28: Final accidental crossing mass templates for various DV track multiplicities in both signal regions, as indicated in the legend of each plot. Each template is scaled by the appropriate crossing factor. Statistical error bars are drawn on the combined template only.

6.6.2 Merged vertices

Two low-mass Standard Model long-lived particle decays, hadronic interactions, or accidental crossings may be merged together in the vertexing algorithm, resulting in a high-mass DV that satisfies the signal region selection. This is called a merged vertex and its contribution to background in the signal region is estimated fully from data. The merged vertex background is estimated by comparing the distributions for the distance significance $S(V_1, V_2)$ for pairs of DVs taken from the same event and from distinct events. Merging can only occur between pairs of DVs in the same event which satisfy $S < 10$. Therefore the merging rate is measured in the $S < 10$ region using the deficit of DV pairs from the same event compared to DV pairs from distinct events. The steps to estimate the number of the merged vertices contribution in the signal region are as follows:

1. **Calculate merging rate:** The vertex merging rate is derived from the ratio of the distance significance distributions between preselected DV pairs taken from the same event and from different events. To increase the available statistics, only a subset of the DV-track cleaning criteria is applied in this step.
2. **Build mass template:** A mass template of merged vertices is built by merging pairs of DVs taken from different events which satisfy $S < 10$. The mass of a merged DV pair is the invariant mass calculated using the four-vectors of all DV tracks.
3. **Normalise mass template:** The merging rate is used to normalise the merged vertex mass templates. Each merged vertex mass in the template is assigned the same weight, so that the integral of the template is consistent with the measured merging rate.
4. **Apply full DV-track cleaning:** The full DV-track cleaning is applied to merged DV pairs in the mass templates.
5. **Integrate mass template:** The number of vertices with a merged mass in the SR is calculated by integrating the appropriate templates in the $m > 10$ GeV region.

The above steps are performed separately for each DV-track multiplicity. The details of these steps are described in the following sections.

Merging rate

To calculate the merging rate, the distance significance between pairs of preselected DVs in the same event is calculated. The distance significance distributions are built separately for DVs with different track multiplicities. If there are N possible pairs of DVs in the same event, then the distance significance between all pairs is calculated and they are added to the distribution with weight $1/N$.

Next, this procedure is repeated with preselected DVs taken from different events (called 'mixed' event). The distance significance between each DV and the DVs in $n = 10$ distinct events is calculated and added to the mixed-event distance significance distributions. The value of n is chosen to reduce the statistical uncertainty on the merging rate.

The distribution of pairs of preselected 2-track DVs taken from the same event and from distinct events are shown in Figure 6.29(a) for events which pass the High- p_T jet selection. The mixed-event distance significance distribution is normalised to match the same-event distance significance distribution in the $S > 10$ region, where merging cannot occur. The shape of the distributions match well in the $S > 10$ region. In the $S < 10$ region, there is a deficit of pairs of DVs in the same-event distance significance distribution, which indicates that merging has taken place. The ratio of these two distributions is called the merging rate. The absolute difference between these distributions in the $S < 10$ region is taken as the total number of events with a 4-track merged vertex with any mass in all events passing the High- p_T jet selection. This procedure is repeated for DVs with higher track multiplicities and in events which pass the Trackless jet selection.




Figure 6.29: In (a), the distance significance (squared) between pairs of preselected DVs taken from the same event and from distinct events passing the High- p_T jet selection is shown. The mixed event distribution is normalised to match the same event distribution in the $S^2 > 100$ region. In (b), the (2+2)-track merged vertex mass distribution is shown, before and after the full DV-track cleaning is applied.

The number of pairs of DVs drops quickly at small distance significance. In order to stabilize the merging rate against statistical fluctuations, the DV-track cleaning described in Section 6.4.5 is loosened before calculating the merging rate. It is important to only loosen DV-track cleaning selection which will not have an impact on the merging rate. To check this, the merging rate is calculated as a function of each of the variables included in the DV-track cleaning criteria (see Table 6.3). For example, the merging rate as a function of the maximum track- p_T and d_0 -significance in (2+2)-track pairs of DVs is shown in Figure 6.30. The merging rate is flat as a function of the maximum track- p_T ,

but it increases as the maximum d_0 -significance decreases. This check is repeated for the remaining DV-track cleaning criteria, and only the d_0 -significance and small angle, low- p_T selection have a significant impact on the merging rate. These two selection are therefore applied before calculating the merging rate.




Figure 6.30: Distribution of pairs of 2-track DVs with $S < 10$ as a function of the maximum p_T and maximum d_0 -significance of all tracks in the DV. The distribution of pairs of DVs taken from different events (mixed event) is normalised to match the number of pairs of DVs taken from the same event in the $S > 10$ region.

Mass template and normalisation

Once the total number of events with a merged vertex is determined, the next step is to determine how many of these merged vertices have a high mass ($m_{DV} > 10$ GeV) after requiring all DV-track cleaning. First, pairs of DVs from mixed events with $S < 10$ are merged without vertex refitting. The merged DV mass is calculated using the four-vectors of the tracks which constitute the DVs in the pairs. The mass of the merged DV using the loosened DV-track cleaning is added to the merged vertex mass template, as shown for (2+2)-track DVs in Figure 6.29(b). Each DV in the template is assigned a uniform weight so that the integral over the entire mass range is equal to the total number of expected merged vertices. Then, the full DV-track cleaning is applied to the DVs which have been added to the mass template. The m_{DV} and n_{Tracks}^{DV} is recalculated for each DV after the full DV-track cleaning and added to the appropriate mass template, preserving the weights from the previous step. This results in the template shown in Figure 6.29(b). Finally, the contribution of merged vertices to the SR is estimated by integrating the $n_{Tracks}^{DV} \geq 5$ templates in the $m_{DV} > 10$ GeV region.

The statistical uncertainty of the merged vertex template is 4% and 3% in the High- p_T and Trackless signal regions, respectively. The statistical uncertainty of the number of DV pairs with $S < 100$ in the same event and in different events is propagated to the final estimate of merged vertices. An uncertainty of 40% in events passing the High- p_T jet selection and from 10% in events passing the

Trackless jet selection is assigned to cover this.

Impact of DV-DV correlations

This background estimation method is sensitive to the presence of correlated DVs in the same event. Pairs of DVs are correlated if they originate from the same cascade hadronic interaction, for example. DVs which are correlated tend to be reconstructed close together, leading to an excess of DV pairs from the same event with a small distance significance between them. This is demonstrated when estimating the background from merged vertices inside of material regions in data, where there are more correlated DVs from hadronic interactions. The distance significance distributions of pairs of DVs which fail the loose material map veto are shown in Figure 6.31(a). The merging rate cannot be calculated in this case because there are more DV pairs from the same event than from different events in the $S < 10$ region.




Figure 6.31: The distance significance (squared) between pairs of preselected DVs failing the loose material map veto, taken from the same event and from distinct events passing the Trackless jet selection is shown. The mixed-event distribution is normalised to match the same-event distribution in the $S^2 > 100$ region. In (a), there is no additional selection on pairs of DVs, while in (b) DV pairs are required to satisfy $\Delta\phi > 0.2$.

Several variables are studied to identify and remove the excess of DV pairs in the same event with $S < 10$. The merged vertex mass of pairs of DVs with $S < 10$ is shown in Figure 6.32(a). The excess of DV pairs in the same event peaks at a low mass, which is consistent with what is expected from correlated DVs originating from the same hadronic interaction. The $\Delta\phi$ between the reconstructed position of pairs of DVs is shown in Figure 6.32(b). The excess of DV pairs in the same event tend to have a small $\Delta\phi$ between them. Therefore, to remove this component of correlated DVs inside material regions, DV pairs with $\Delta\phi < 0.2$ are excluded from the estimate. The distance significance distributions after applying the $\Delta\phi > 0.2$ selection are shown in Figure 6.31(b). After applying this selection, there is a deficit of DV pairs in the same event with $S < 10$ and a merging rate can be

calculated.




Figure 6.32: The $\Delta\phi$ between pairs of preselected DVs with $S < 10$, taken from the same event and from distinct events passing the Trackless jet selection is shown. The mixed-event distribution is normalised to match the same-event distribution in the $S > 10$ region.

Though the contribution from DV-DV correlations outside of material is small enough to enable a merging rate to be calculated without applying a $\Delta\phi$ selection, it is still possible that they have a non-negligible impact on the background estimate. To evaluate this impact, two sets of merged vertex mass templates for each track multiplicity of DVs outside of material are produced. The first set of templates is produced without a $\Delta\phi > 0.2$ selection, while the second set of templates is produced with the $\Delta\phi > 0.2$ selection. These templates are shown in Figure 6.33. The merged vertex background estimate of 4-trk DVs decreases after adding the $\Delta\phi > 0.2$, while it increases for DVs with higher track multiplicities. The nominal merged vertex background estimate is taken as the average of the two sets of templates. The difference between the two sets of templates and the nominal is taken as a systematic uncertainty. In the $m_{\text{DV}} > 10$ GeV region, this uncertainty is 7.4% (20%) for 4-trk DVs and 15% (28%) for ≥ 5 -trk DVs in events passing the Trackless (High- p_T) jet selection.




Figure 6.33: The merged vertex mass templates with and without a $\Delta\phi > 0.2$ selection for (a) 4-trk DVs and (b) ≥ 5 -trk DVs in events passing the Trackless jet selection is shown.

6.6.3 Hadronic interactions

Hadronic interactions originate from the interaction of particles from the proton-proton collision interacting with detector material. The material map veto removes the majority of this background component. The residual component which remains is due to hadronic interactions occurring in less dense regions of the detector or in regions which are missed by the material map veto due to the resolution of the vertex position.

Hadronic interactions dominate the low-mass region of the m_{DV} distribution. Therefore, to estimate this background in the signal region, the m_{DV} distribution is fit to a function in the $m_{DV} < 10$ GeV region and extrapolated to the signal region with $m_{DV} > 10$ GeV. A study of the truth record in simulated events indicates that the contamination from non-HI background displaced vertices in the $m_{DV} < 10$ GeV region is expected to be negligible. The hadronic interactions background appears to follow an exponentially falling distribution in the high mass region. At low mass, the number of DVs rises sharply before peaking, and this depends on the DV-track multiplicity. Therefore, using just an exponential function to fit the m_{DV} distribution results in a fit result which is heavily dependent on the fit range used. Instead, the entire m_{DV} spectrum in the $m_{DV} < 10$ GeV region is fit using a function which encapsulates both the sharp rise and the exponential decay of the hadronic interactions. This function is

$$N_{DV}(m_{DV}) = \frac{1}{\frac{1}{C(m_{DV}-b)} + \frac{1}{e^{-\frac{m_{DV}-B}{l}}}}, \quad (6.8)$$

where C dictates the slope of the initial straight line, b dictates when it begins to rise from 0, B and l describe the magnitude and lifetime of the exponential decay. A least-squares fit is performed in the $m_{DV} = [0, 10]$ GeV mass range with an m_{DV} bin size of 0.5 GeV. The value of the fitted function

at the center of each bin is taken as the estimate of hadronic interactions in that bin. The hadronic interactions estimate in bins where the fitted function is negative is set to 0.

The uncertainties on the fitted parameters are extrapolated to the estimate of background in the signal region. This uncertainty is around 200% in the High- p_T SR and 50% in the Trackless SR.

6.6.4 Combined estimate

The separate background estimates are combined to form a complete background estimate in the signal and validation regions. Each mass template is normalised according to their respective data-driven technique described in the previous sections.

Closure test in simulated dijet events

The performance of the combined background estimates is tested using simulated dijet events. The three separate background m_{DV} templates are built and normalised from events in the simulated dijet events, treating them exactly as data. The closure test is performed separately for DVs inside and outside of material. The results of the closure tests, both outside and inside material, are shown in Figure 6.34. The backgrounds are estimated in the high- m_{DV} regions by integrating the corresponding templates in the appropriate mass range. The results are listed in Table 6.8 (outside material) and Table 6.9 (inside material). There is a consistent underestimation of the background in the high- m_{DV} regions.

The non-closure in the high- m_{DV} regions is used to assign a systematic uncertainty on the combined accidental crossings and merged vertices background estimates. Ideally one would use the non-closure in the outside material regions (rightmost column of Table 6.8) to calculate the uncertainty. However, the non-closure in those regions is not significant due to the large statistical uncertainties. Therefore, the non-closure in the high- m_{DV} regions inside of material is used to assign the uncertainty instead

Region	MV	HI	AX	Comb.	Obs.	Non-clos.
4-trk, $m_{DV} > 10$ GeV	6.2 ± 2.5	$(1.5 \pm 1.0) \times 10^{-3}$	5.3 ± 1.3	11.5 ± 2.8	17	48%
≥ 5 -trk, $m_{DV} > 10$ GeV	1.23 ± 0.47	$(3.2 \pm 3.0) \times 10^{-3}$	0.83 ± 0.21	1.96 ± 0.52	4	104%

Table 6.8: The combined background estimates in simulated dijet events **outside of material**. The observed number of DVs is also listed. The non-closure is calculated as the percent difference between the observed and expected background.

Region	MV	HI	AX	Combined	Observed	Non-clos.
4-trk, $m_{DV} > 12$ GeV	8.5 ± 3.5	0.306 ± 0.026	44 ± 11	52 ± 11	89	71%
≥ 5 -trk, $m_{DV} > 15$ GeV	7.3 ± 3.0	3.34 ± 0.18	29.5 ± 7.4	40.1 ± 7.9	67	67%

Table 6.9: The combined background estimates in simulated dijet events **inside of material**. The observed number of DVs is also listed. The non-closure is calculated as the percent difference between the observed and expected background.




Figure 6.34: The combined background predictions in simulated dijet events for DVs inside and outside of material. The overflow is drawn in the last bin.

(rightmost column of Table 6.9). Although hadronic interactions are more prevalent inside of material, the high mass region ($m_{\text{DV}} > 12$ GeV for 4-trk DVs and $m_{\text{DV}} > 15$ GeV for ≥ 5 -trk DVs) is still dominated by accidental crossings and merged vertices. With this method, an additional uncertainty of 70% on the accidental crossings and merged vertices background estimates in the signal regions is assigned.

Background estimates in data

The normalised background predictions for 4-track and ≥ 5 -track DVs in events passing the High- p_T and Trackless jet selection are shown in Figures 6.35 and Figure 6.36, respectively. Good closure is achieved in the low- m_{DV} and low- $n_{\text{Tracks}}^{\text{DV}}$ sideband validation regions.

The estimate of background in the signal region is calculated by integrating all ≥ 5 -track mass templates in the $m_{\text{DV}} > 10$ GeV region. The uncertainties on the separate background estimates are added in quadrature to produce the overall uncertainty on the combined background estimate.

The resulting estimates in the High- p_T and Trackless signal regions are shown in Table 6.10. These estimates are consistent with those predicted by the inclusive background estimate, thereby providing a successful cross-check of the inclusive background estimate.




Figure 6.35: The combined background predictions in events passing the **High- p_T jet selection**. In (a), only the mass of 4-track DVs is shown. In (b), the mass of ≥ 5 -track DVs is shown. Events in the signal region ($m_{DV} > 10$ GeV) are not drawn.




Figure 6.36: The combined background predictions in events passing the **Trackless jet selection**. In (a), only the mass of 4-track DVs is shown. In (b), the mass of ≥ 5 -track DVs is shown. Events in the signal region ($m_{DV} > 10$ GeV) are not drawn.

Region	MV	HI	AX	Combined	Inclusive
High- p_T SR	0.79 ± 0.66	$(2.7 \pm 5.6) \times 10^{-4}$	0.28 ± 0.21	1.07 ± 0.83	0.49 ± 0.28
Trackless SR	1.5 ± 1.1	$(1.1 \pm 1.4) \times 10^{-3}$	0.32 ± 0.24	1.8 ± 1.3	0.69 ± 0.40

Table 6.10: The separate background estimates in the signal regions. The uncertainties on the total background estimate are calculated assuming no correlation between the separate estimates. The background estimate from the inclusive method is also listed.

6.7 Uncertainties

Simulated signal events are used to interpret the results of this search. The predicted yield of signal events depends on the signal efficiency, or the fraction of signal events which pass all event selection and fall into the signal region. These yields of simulated signal events are compared to the observed data. It is therefore important to understand and quantify differences between simulated events and real data that may impact the signal selection efficiencies. These sources are quantified as systematic uncertainties and are included in the fitting procedure as nuisance parameters.

The dominant uncertainty stems from differences in the performance of the Standard and Large Radius Tracking algorithms between data and simulation. The uncertainty of standard tracking is determined by the ATLAS tracking combined performance group to be around 1.7%. The uncertainty associated with the large radius tracks is estimated by comparing K_S^0 yields in data and simulated dijet events and is propagated to the simulated signal events. The procedure is described in Section 6.7.1. This uncertainty is treated as correlated between signal regions in the interpretation of the results.

The jet energy scale (JES) and jet energy resolution (JER) uncertainties quantify the differences in jet reconstruction in data and simulated events. The impact that these uncertainties have on the signal efficiency are less than 2% for the Strong RPV model. The impact on the EWK RPV model is less than 10% for the majority of the parameter space and reaches a maximum of 50% for samples with low $m_{\tilde{\chi}}$ and long $\tau_{\tilde{\chi}}$. The jet- p_T response was also studied as a function of the displacement of jets in the simulated signal samples. The standard jet calibration performs well also for the displaced jets in the analysis, and no additional uncertainty is assigned to the jet- p_T due to displacement.

All other sources of systematic uncertainties (luminosity of the dataset, trigger efficiency, pileup modeling, simulated event modeling, ...) are negligible with respect to these uncertainties.

6.7.1 Uncertainty due to non-standard tracking

To estimate this uncertainty, the number of reconstructed K_S^0 's in different radial regions is compared between data and the simulated dijet events. A deviation of the data/simulation ratio is used to derive an uncertainty on the track reconstruction efficiency in simulated events. This track reconstruction efficiency is then propagated to the signal efficiency.

Candidate K_S^0 's are required to pass the DV preselection, have exactly two tracks, the sum of which

must point back to the interaction point ($\cos(\Delta\phi) > 0.9999$) and have a mass within 50 MeV of the K_S^0 mass. The tracks originating from K_S^0 decays at small radii are typically reconstructed as standard tracks, which means that the reconstruction efficiency and its uncertainty are well understood in this region. Therefore, the K_S^0 yield in the simulated events is normalised to data in the $R_{xy} < 22$ mm region.

The number of K_S^0 's as a function of radial decay position is shown in Figure 6.37(a). In each radial bin, the 2-track DV mass spectrum is fit and the flat background is subtracted before counting the number of K_S^0 's. There is good agreement between simulation and data in the $R_{xy} < 50$ mm region, beyond which the simulation overestimates the K_S^0 yield in data. This is due to inefficiencies in data which are not fully described in the simulation, the main sources of which are:

1. Multiple scattering - a charged particle undergoes a large multiple scattering or hadronic interaction, and its track can no longer be followed. This is a source of inefficiency because the simulation does not perfectly model multiple scattering and hadronic interactions.
2. Disabled modules - a charged particle passes through a disabled module which is not perfectly accounted for in the simulation. In that case, the reconstruction algorithm expects their to be hit where the disabled module is, when in fact it registers a hole.

When propagating this inefficiency to the signal yields to calculate an uncertainty, it is important to understand if they should be treated as correlated or uncorrelated for tracks in the same DV. If the dominant source of inefficiency comes from multiple scattering, then the uncertainty will be uncorrelated between tracks reconstructed in the same DV. This because each particle has an independent probability that it will undergo multiple scattering. This type of inefficiency is dominant for standard tracks and has a gradual η and ϕ dependence corresponding to the position of material in the detector.

If the dominant source of inefficiency comes from disabled modules, then this will be correlated between tracks in the same DV, as they are likely to pass through the same modules. While this is a subdominant effect for standard tracks, this could have a larger impact on large radius tracks since they have fewer silicon hits and therefore the impact of losing one from a disabled module may be larger. This type of inefficiency would have a large η and ϕ dependence, corresponding to the position of the disabled modules. Furthermore, it could have a time dependence if modules become disabled or encounter problems during the data-taking. The data/simulation ratio of the K_S^0 yield as a function of η , ϕ , and run number were studied, and no obvious trends were observed. Therefore, the contribution to the inefficiency from disabled modules is assumed to be negligible, and the uncertainty between tracks in the same DV are treated as uncorrelated.

A per-track uncertainty is calculated using the deviation of the data/simulation ratio from unity. This per-track uncertainty is shown in Figure 6.37(b), where an uncertainty of 1.7% from the standard tracking prescription has been added in quadrature as a conservative measure. The tracking




Figure 6.37: In (a), the K_S^0 yield in different radial regions is compared between data and simulation. The simulated events are normalised to match the data in the innermost radial region (first bin). In (b), the tracking uncertainty is shown, assuming fully uncorrelated uncertainties for tracks in the same DV.

uncertainties are propagated to the signal efficiency using a track-killing method where DV tracks in the simulated signal events are randomly removed with a probability equal to the per-track uncertainty. Several track-killed samples are produced in this way, and the average efficiency of these track-killed samples is calculated. The difference of this efficiency with respect to the nominal efficiency is taken as the uncertainty on the signal yield, shown in Figure 6.38. The uncertainty on the signal efficiency reaches a maximum of 15% for signal samples with lower mass neutralinos and long lifetimes, while the efficiency is much lower for higher mass neutralinos. This is because the higher mass neutralinos are more likely to be reconstructed with higher track multiplicities, so that removing one track will not remove that DV from the SR.




Figure 6.38: The uncertainty on the signal efficiency as a function of neutralino lifetime.

Region	Expected background yield	Observed yields
High- p_T SR	0.49 ± 0.28	1
Trackless SR	0.69 ± 0.42	0

Table 6.11: Expected background yields and observed yields in data in both signal regions.

Selection	High- p_T SR		Trackless SR	
	Events	Rel. Eff. [%]	Events	Rel. Eff. [%]
Multijet trigger	122549332		122549332	
Filter	34476391	28.1	34476391	28.1
Jet requirements	2315962	6.7	4846972	14.1
Event has ≥ 1 DV passing:				
Fiducial volume	1628663	70.3	3389705	69.9
$r_{xy}(\text{PV-DV}) > 4$ mm	1613222	99.1	3357880	99.1
χ^2/n_{Dof}	1603256	99.4	3344324	99.6
Material veto	739963	46.2	1841992	55.1
$n_{trk} \geq 5$	712	0.1	1011	0.1
$n_{trk}^{sel} \geq 2$	702	98.6	993	98.2
$m_{DV} > 10$ GeV	1	0.1	0	0.0

Table 6.12: The yields in data after applying various selection in the High- p_T and Trackless signal regions. The selections are applied cumulatively and the relative efficiency of each selection is calculated with respect to the previous selection.

6.8 Results

The expected background and observed data in both signal regions is listed in Table 6.11. In both signal regions, the observed event yields are in good agreement with the background-only expectations. The one event passing all selections in the High- p_T SR was recorded in 2017. It contains a displaced vertex reconstructed within the beampipe ($R = 7.5$ mm), with 5 associated tracks and a mass of 32.6 GeV. This event is fully consistent with the background predictions.

The event yields in data are shown after various selections in Table 6.12. The yields are shown separately for the High- p_T and Trackless signal regions, which differ after the jet requirements are applied. The relative efficiency of each displaced vertex selection is similar for both signal regions. The largest variation occurs for the material map veto where the relative efficiency of the selection is 46% in events passing the High- p_T jet selection and 55% in events passing the Trackless jet selection. This variation is likely due to the fact that the relative fraction of the different sources of background vary between the two signal regions. The density of the material map veto is not uniform in the fiducial volume and therefore it is expected to have a different impact on the various background sources.

The event yields in data and in several signal models are shown after various selections in Figure 6.39. The relative contributions of data and signal demonstrate the importance of the signal region selection,

especially the m_{DV} and DV track multiplicity requirements, in increasing the sensitivity of the analysis to these signal models.




Figure 6.39: The event yields in data and several signal models after applying selection cumulatively in the (a) High- p_T and (b) Trackless signal regions.

The m_{DV} distributions of preselected DVs with at least five tracks are shown in Figure 6.40. The distributions are shown for data, the expected background, and several signal models. The data is in good agreement with the expected background over the entire mass range.




Figure 6.40: Distributions of invariant mass m_{DV} for preselected DVs with at least five tracks in events which pass the (a) High- p_T and (b) Trackless jet selection. The shaded histograms show the separate background estimates and the dashed lines show several different signal models. The error bars on the combined background estimate are the statistical and systematic errors added in quadrature. The black arrow denotes the mass window of the signal region. In events which have more than one preselected DV, only the most massive DV is included in the distribution.

6.9 Statistical analysis

A statistical analysis is performed on the results of the search in order to set an upper limit on the cross section of the benchmark signal models and to exclude regions of the parameter space with a certain level of confidence. This is performed using a hypothesis test, which quantifies the level of disagreement between the observed data and the expected outcomes of the experiment under the assumption of different signal models.

The likelihood function L quantifies the likelihood that the data originates from the hypothesis of signal strength μ . It is built from the product of Poisson probability distribution functions, one for each bin in the fit:

$$L(\text{data} | \mu, \theta) = \prod_{i=1}^{N_{\text{bins}}} \frac{(\mu s_i + b_i)^{n_i}}{n_i!} e^{-(\mu s_i + b_i)}, \quad (6.9)$$

where n_i is the number of observed events in each bin i , and N_{bins} is the total number of bins. The expected event yield in each bin is given by $\mu s_i + b_i$, where s_i and b_i are the expected number of signal and background events in bin i , respectively.

Systematic uncertainties (θ) are parametrised as nuisance parameters and constrained in the fit. They result in a loss of sensitivity of the analysis. The post-fit values of these nuisance parameters are not of interest. Therefore, a profile likelihood function is utilised, which "profiles out" the nuisance parameters by replacing θ with the conditional maximum likelihood estimators of the nuisance parameters, $\hat{\theta}$ (i.e. the value of θ that maximizes L for a specified μ). This reduces the dependency of the likelihood function to only the parameter of interest and allows for the construction of a one-dimensional confidence interval in the parameter of interest which is independent of the nuisance parameter values [118].

A *test statistic* is a quantity which reduces the observed and expected data to a single number that is sensitive to differences between the null ($\mu = 0$) and alternative hypotheses being tested. The test statistic $\tilde{q}_\mu$ used in the interpretation of these results is built from the profile likelihood ratio, defined as

$$\lambda(\mu) = \frac{L(\mu, \hat{\theta})}{L(\hat{\mu}, \hat{\theta})}.$$

Here, $\hat{\mu}$ and $\hat{\theta}$ are the unconditional maximum likelihood estimators of μ and θ , respectively. The denominator of the profile likelihood ratio is the maximised likelihood function, ensuring that $\lambda \leq 1$.

The test statistic is defined as:

$$\tilde{q}_\mu = \begin{cases} -2 \ln \frac{L(\mu, \hat{\theta})}{L(0, \hat{\theta})} & \hat{\mu} < 0 \\ -2 \ln \frac{L(\mu, \hat{\theta})}{L(\hat{\mu}, \hat{\theta})} & 0 \leq \hat{\mu} \leq \mu \\ 0 & \hat{\mu} > \mu \end{cases} \quad (6.10)$$

The form of the test statistic depends on the value of $\hat{\mu}$ to ensure that $\mu \geq 0$ (avoiding unphysical cases) and that the test statistic is not penalized in case of an upward fluctuation of $\hat{\mu}$. Lower values of $\tilde{q}_\mu$ indicate a better compatibility between the data and the hypothesized value of μ .

The test statistic is used to calculate a *p-value*:

$$p_\mu = \int_{\tilde{q}_{\mu, \text{obs}}}^{\infty} f(\tilde{q}_\mu | \mu) d\tilde{q}_\mu, \quad (6.11)$$

where $\tilde{q}_{\mu, \text{obs}}$ is the value of the test statistic observed in data. The p-value quantifies the level of disagreement between the outcome of the experiment and the hypothesis being tested. The distribution of the test statistic in different hypotheses (μ) is required in order to perform a hypothesis test. These distributions are obtained either using pseudoexperiments or using asymptotic formulae [119]. It is expected that the commonly used asymptotic formulae for the distributions of $\tilde{q}_\mu$ under various signal strength hypotheses are not applicable in this analysis, due to the very low expected background and signal yields. Therefore pseudoexperiments are instead used to determine the distributions of $\tilde{q}_\mu$ numerically.

The CL_s method uses the p-value to set exclusion limits. It is defined as

$$\text{CL}_s = \frac{p_{s+b}}{1 - p_b} \quad (6.12)$$

where p_{s+b} is the level of disagreement between the data and the signal plus background hypothesis and p_b is the level of disagreement between the data and the background-only hypothesis. If the CL_s value is less than 0.05 for a certain signal strength μ , then that signal hypothesis is excluded with 95% confidence level. The CL_s method [120] is used to set limits rather than p_{s+b} alone, because it avoids excluding models which the analysis does not have sensitivity to. It does this by penalizing the p_{s+b} with $1 - p_b$. If an analysis does not have sensitivity to a particular model, then $\tilde{q}_\mu$ for the signal plus background and background-only hypothesis will be small. This will lead to a small value of $1 - p_b$, which will increase the CL_s such that that signal model is not excluded.

Expected exclusion limits are quantified assuming the absence of any excesses in the data. It is used to quantify the sensitivity of an analysis. To calculate expected exclusion limits, the $\tilde{q}_{\mu, \text{obs}}$ in Eq. 6.11 is replaced with the median $\tilde{q}_\mu$ expected with data distributed according to $\mu = 0$.

6.10 Interpretation

The results are interpreted in the context of the Strong RPV and EWK RPV signal models described in Section 6.1. The expected number of Strong RPV signal events in the High- p_T signal region is shown in Figure 6.41(a) for a fixed gluino mass of 2.6 TeV and in Figure 6.41(b) with a fixed neutralino lifetime of 0.1 ns. The trackless SR does not add significant yields for this model. For the EWK RPV signal model, both signal regions are important, and the expected yields are shown in Figure 6.42.




Figure 6.41: The expected event yields of the Strong RPV signal model in the High- p_T SR with (a) fixed gluino mass of 2.6 TeV and (b) fixed neutralino lifetime of 0.1 ns.




Figure 6.42: The expected event yields of the EWK RPV signal model in the (a) High- p_T SR and (b) Trackless SR.

Exclusion limits on the Strong RPV and EWK RPV models are set using the CL $_s$ prescription described in the previous section. The exclusion limit setting is implemented in pyhf using 1 million pseudoexperiments [121]. It is implemented using a two-bin fit where the two bins correspond to the

	Uncertainty	Treatment
Expected background	CR statistical	Uncorrelated
	Non-linearity	Correlated
	Non-closure	-
	Pileup	Correlated
Expected signal	Cross section	Correlated
	JES/JER	Correlated
	Tracking	Correlated
	MC stats	Uncorrelated

Table 6.13: Treatment of correlation between uncertainties on the expected yields in the High- p_T and Trackless signal regions. The Non-closure uncertainty only applies to the expected background yield in the High- p_T signal region.

High- p_T and Trackless signal regions. The treatment of correlation between uncertainties in both signal regions are summarized in Table 6.13.

The 95% confidence level upper limits on the production cross section of the Strong RPV signal model for two gluino masses is shown in Figure 6.43. For $\tau = 0.1$ ns, the cross-section upper limits are set below 0.04 fb for neutralinos with $m_{\chi_1^0} > 250$ GeV. This analysis is most sensitive to models where the neutralino lifetimes is around 0.1 ns, so that the proper decay length falls within the inner tracker. The sensitivity decreases for both shorter and longer lifetimes due to acceptance and tracking/vertexing efficiency effects.




Figure 6.43: Upper limits on the production cross section of gluinos in the Strong RPV model as a function of m_{χ^0} for several values of τ_{χ^0} .

The expected and observed 95% confidence level exclusion limits on the mass of the neutralino as a function of its mean proper lifetime for $m_{\tilde{g}} = 2.0$ TeV and $m_{\tilde{g}} = 2.4$ TeV are shown in Figure 6.44. A larger portion of the parameter space with $m_{\tilde{g}} = 2.0$ TeV is excluded due to the higher cross section of this process at the LHC. Regions in the parameter space with a heavy neutralino with a short lifetime

are not excluded because those signal models fall out of the acceptance due to the small boost of the neutralino. Similarly, regions with a light neutralino mass with a long lifetime are also not excluded, as these tend to decay at larger radii from the interaction point, where the tracking efficiency falls off, or even outside of the tracker and therefore out of the acceptance of this analysis.




Figure 6.44: Exclusion limits in the $m_{\tilde{\chi}_1^0}$ vs. $\tau_{\tilde{\chi}_1^0}$ plane for the Strong RPV model and a gluino mass of (a) 2.0 TeV and (b) 2.5 TeV. The area between the two solid red lines (shaded) is the excluded parameter space.

The expected and observed 95% confidence level exclusion limits on the mass of the gluino as a function of the mean proper lifetime of the neutralino for $m_{\tilde{\chi}_1^0} = 50 \text{ GeV}$ and $m_{\tilde{\chi}_1^0} = 200 \text{ GeV}$ are shown in Figure 6.45. The excluded region is larger for more massive neutralinos, as their decays are more likely to produce displaced vertices which pass the high mass and track multiplicity requirements of the signal region. Figure 6.45(b) is particularly interesting, as it can be directly compared to the existing limits on the λ_{112}'' RPV coupling shown in Figure 2.8(a). For a fixed $m_{\tilde{\chi}_1^0} = 200 \text{ GeV}$ and $\tau_{\tilde{\chi}_1^0} = 0.1 \text{ ns}$, gluinos with masses up to 2550 GeV are excluded. This analysis therefore successfully probes a large portion of the uncovered regions of $m_{\tilde{g}}$ vs. $\tau_{\tilde{\chi}_1^0}$ parameter space in this model.

The 95% confidence level upper limits on the production cross section of the EWK RPV model as a function of the mass of the neutralino is shown in Figure 6.46(a). For $\tau = 0.1 \text{ ns}$, the cross-section upper limits are set below 0.05 fb for neutralinos with masses $m_{\tilde{\chi}_1^0} > 600 \text{ GeV}$. The 95% confidence level exclusion limits on the mass of the neutralino as a function of its lifetime are shown in Figure 6.46. For a mean proper lifetime of $\tau = 0.1 \text{ ns}$, neutralino masses below 1.6 TeV are excluded. For masses on the order of the DV mass requirement of 10 GeV and below, the sensitivity is expected to drastically decrease. For the wide range of mean proper lifetime values between $\tau = 0.01 \text{ ns}$ to $\tau = 10 \text{ ns}$, neutralino masses below 1 TeV are excluded.

These are the first search results for high mass, multi-track displaced vertices and multiple energetic jets with the ATLAS detector using $\sqrt{s} = 13 \text{ TeV}$ pp collision data. Furthermore, this is the first search in the ATLAS experiment which has set limits on an RPV SUSY model predicting the production




Figure 6.45: Exclusion limits in the $m_{\tilde{g}}$ vs. $\tau_{\chi_1^0}$ plane for the Strong RPV model and a neutralino mass of (a) 50 GeV and (b) 200 GeV. The area below the solid red lines (shaded) is the excluded parameter space.




Figure 6.46: In (a), the upper limit on the production cross section of the EWK RPV model is shown. In (b), the exclusion limits on $m_{\chi_1^0}$ as a function of $\tau_{\chi_1^0}$ are shown. The region below the solid red line (shaded) is excluded.

of long-lived electroweakinos. A similar search for displaced vertices was performed in the CMS experiment [27]. In this search, an RPV SUSY models with pair-produced, long-lived neutralinos which decay via an RPV coupling to heavy-flavor quarks (top, bottom, and strange) are considered. Pair-produced neutralinos with $\tau = 0.03 \text{ ns}$ are excluded up to 1.3 TeV. The limits from this search are slightly stronger, excluding pair-produced neutralinos with $\tau = 0.03 \text{ ns}$ up to 1.5 TeV.

Conclusion

This thesis focuses on a search for new, long-lived particles with the ATLAS experiment at the LHC. New, long-lived particles can leave unconventional signatures in the detector which are challenging to search for and therefore may have evaded detection until now. With the continuing absence of any obvious signs of new physics in LHC data, it is motivated to target unconventional signatures of new particles to ensure that new physics isn't hiding within the dataset.

Integrating tracking information into the trigger is critical for searches for new physics which rely on track-based signatures. In Chapter 5, the ATLAS FastTracKer (FTK) is presented which is a hardware-based tracking system designed to be integrated into the ATLAS trigger system. I produced the configuration of the system for the purpose of commissioning in 2019 and prepared the system to be able to adapt quickly to new beamspot and silicon module positions during data-taking. Though the FTK system will not be installed in Run-3, the integration of charged particle tracking information into the trigger at the earliest possible stage is still a top priority for the ATLAS experiment. In Run-3, the ATLAS triggering system will perform more CPU-based tracking in the HLT for signatures that benefit the most (i.e. jets and E_T^{miss}). Furthermore, in the HL-LHC, the ATLAS experiment plans to perform extensive tracking in the trigger using CPU cores and accelerator hardware. (GPUs, FPGAs, ...) [98].

In Chapter 6, a search for new, long-lived particles using displaced vertices (DV) is presented. I contributed to nearly all aspects of the search, most notably the event selections, background estimation, study of systematic uncertainties, and the interpretation of the results. The search uses the full Run-2 dataset corresponding to 139 fb^{-1} at $\sqrt{s} = 13 \text{ TeV}$. It is the first search for high mass, multi-track displaced vertices and multiple energetic jets with the ATLAS detector using $\sqrt{s} = 13 \text{ TeV}$ pp collision data. Two signal regions are utilized which require a high multiplicity of high- p_T jets or trackless jets and at least one reconstructed displaced vertex with a high mass and high track multiplicity. The backgrounds in this search stem from algorithmic and instrumental sources and are estimated using a

data-driven technique which exploits the correlation between background DVs and jets. With this technique, 0.49 ± 0.28 (0.69 ± 0.42) events are predicted in the High- p_T (Trackless) signal regions. The data agree with the expected background yields, with one event observed in the High- p_T SR and no events observed in the Trackless SR.

The results of the search are interpreted in the context of supersymmetric models with R-parity violating couplings. For $\tau = 0.1$ ns, the cross-section upper limits for the Strong (EWK) RPV models are set below 0.04 (0.05) fb for neutralinos with $m_{\tilde{\chi}_1^0} > 250$ (600) GeV. Gluinos in the Strong RPV SUSY model are excluded up to 2440 GeV for $m_{\tilde{\chi}_1^0} = 50$ GeV and $\tau = 0.01$ ns. Neutralinos in the EWK RPV SUSY model are excluded up to 1.6 TeV for $\tau = 0.1$ ns.

Compared to the Run-1 version of this search [107] and similar searches for displaced vertices in Run-2 [122], a much simplified background estimation is adopted which utilizes the correlation between jets and DVs. This simplified approach greatly reduces the complexity of the analysis and introduces fewer systematic uncertainties into the estimate of background in the signal region. Furthermore, the Run-1 version of this search only selected events which contain trackless jets for dedicated LRT reconstruction. The addition of the High- p_T SR in this analysis increased the sensitivity of the search to models with cascade decays of SUSY particles and to models with shorter LLP lifetimes where charged particles within the jets are more likely to be reconstructed as standard tracks.

There are several improvements which could be made to the analysis strategy which might improve the sensitivity of the search in the future. The reduction of backgrounds was a significant challenge in this search. A strict DV-track cleaning was applied to reduce the expected backgrounds to less than one event per signal region. However this also reduced the sensitivity to low mass LLPs. If the backgrounds could be reduced at an earlier stage, for example in the vertexing algorithm, then the DV-track cleaning could be loosened so that more low mass LLPs could enter the signal region while still keeping the level of background under control.

Merged vertices are the largest component of background in the high m_{DV} region. This background could be reduced by placing tighter constraints on the requirements for merging in the vertexing algorithm, i.e. by tightening the selection on the distance significance between pairs of DVs which are eligible for merging. The second largest component of background is accidental crossings. This background could be reduced by studying the shifts of the reconstructed DV positions when iteratively removing tracks from the DV. Accidental crossing tracks which originate from pileup vertices are likely to cause a large shift in the DV position along the beamline direction when removed from the DV. This approach is utilized by similar searches for displaced vertices in the CMS experiment [27].

Future iterations of this search will certainly benefit from the addition of data collected in Run-3, which will begin in 2022. Furthermore, a new displaced tracking algorithm which is less CPU-intensive and which reduces tracks reconstructed from random combinations of hits will be employed. This development will remove the requirement for filtering events prior to displaced vertex finding, giving more flexibility to event-level selections. Lastly, searches for displaced vertices may benefit from the

addition of unconventional tracking triggers which are in development for Run-3. One example is a hit-based displaced vertex trigger, which identifies events with a large number of hits that are not associated to prompt tracks after the first inner detector tracking layers. The ability to trigger directly on displaced tracks or displaced vertices in addition to multiple jets could allow for the reduction of the jet- p_T trigger thresholds, thereby increasing the efficiency of selecting new physics events with softer final state particles.

While no evidence of new, long-lived massive particles were uncovered in this search, the results presented in this thesis can help guide future searches at the LHC towards a discovery. This search was designed to be as model-independent as possible in order to maximize its sensitivity to a broad range of theories predicting massive long-lived particles. Though the results of the search were interpreted in the context of RPV SUSY models, it is also expected to have sensitivity to models like mini-split SUSY (R-hadrons), long-lived scalars from Higgs decays [123], and emerging jets from a dark sector [124]. This means that the results of this search can be used to identify still-uncovered regions of parameter space in a broad range of new physics models, thereby giving direction to future exploration in Run-3 and at the HL-LHC.

Bibliography

- [1] D. Galbraith, *Standard Model of the Standard Model*,
url: <http://davidgalbraith.org/portfolio/ux-standard-model-of-the-standard-model/>.
- [2] S. Glashow, *Partial Symmetries of Weak Interactions*, *Nucl. Phys.* **22** (1961) 579.
- [3] S. Weinberg, *A Model of Leptons*, *Phys. Rev. Lett.* **19** (21 1967) 1264.
- [4] J. Greensite, *An introduction to the confinement problem*, Springer, 2011.
- [5] J. F. Gunion et al., *The Higgs Hunter's Guide*, *Front. Phys.* **80** (2000) 1.
- [6] Q. Ahmad et al., *Measurement of the rate of $\nu_e + d \rightarrow p + p + e^-$ interactions produced by ^8B solar neutrinos at the Sudbury Neutrino Observatory*, *Phys. Rev. Lett.* **87** (2001) 071301, arXiv: [nucl-ex/0106015](#).
- [7] CMS Collaboration,
Observation of a new boson at a mass of 125 GeV with the CMS experiment at the LHC,
Phys. Lett. B **716** (2012) 30, arXiv: [1207.7235 \[hep-ex\]](#).
- [8] ATLAS Collaboration, *Observation of a new particle in the search for the Standard Model Higgs boson with the ATLAS detector at the LHC*, *Phys. Lett. B* **716** (2012) 1, arXiv: [1207.7214 \[hep-ex\]](#).
- [9] ATLAS Collaboration, *The ATLAS Experiment at the CERN Large Hadron Collider*, *JINST* **3** (2008) S08003.
- [10] L. Lee et al., *Collider Searches for Long-Lived Particles Beyond the Standard Model*, *Prog. Part. Nucl. Phys.* **106** (2019) 210, arXiv: [1810.12602 \[hep-ph\]](#).
- [11] A. H. G. Peter, *Dark Matter: A Brief Review*, Frank N. Bash Symposium (2011), arXiv: [1201.3942 \[astro-ph.CO\]](#).
- [12] L. Bergstrom, *Non-Baryonic Dark Matter - Observational Evidence and Detection Methods*, *Rep. Prog. Phys.* **63** (2000) 793.
- [13] S. P. Martin, *A Supersymmetry primer*, (1997) 1, [Adv. Ser. Direct. High Energy Phys.21,1(2010); Adv. Ser. Direct. High Energy Phys.18,1(1998)], arXiv: [hep-ph/9709356 \[hep-ph\]](#).

- [14] M. Chemtob,
Phenomenological constraints on broken R parity symmetry in supersymmetry models,
Prog. Part. Nucl. Phys. **54** (2005) 71, arXiv: [hep-ph/0406029](#) [[hep-ph](#)].
- [15] G. Bertone, D. Hooper and J. Silk, *Particle dark matter: evidence, candidates and constraints*,
Physics Reports **405** (2005) 279, ISSN: 0370-1573, URL:
<https://www.sciencedirect.com/science/article/pii/S0370157304003515>.
- [16] G. F. Giudice and R. Rattazzi, *Theories with gauge mediated supersymmetry breaking*,
Phys. Rept. **322** (1999) 419, arXiv: [hep-ph/9801271](#) [[hep-ph](#)].
- [17] S. Knapen and D. Redigolo, *Gauge mediation at the LHC: status and prospects*,
JHEP **01** (2017) 135, arXiv: [1606.07501](#) [[hep-ph](#)].
- [18] ATLAS Collaboration,
Search for displaced leptons in $\sqrt{s} = 13$ TeV pp collisions with the ATLAS detector,
Phys. Rev. Lett. **127** (2020) 051802, arXiv: [2011.07812](#) [[hep-ex](#)].
- [19] CMS Collaboration, *Search for long-lived particles decaying to displaced leptons in proton-proton collisions at $\sqrt{s} = 13$ TeV*, (2021).
- [20] CMS Collaboration, *Search for long-lived particles using delayed photons in proton-proton collisions at $\sqrt{s} = 13$ TeV*, *Phys. Rev. D* **100** (2019) 112003, arXiv: [1909.06166](#) [[hep-ex](#)].
- [21] A. Arvanitaki et al., *Mini-Split*, *JHEP* **02** (2013) 126, arXiv: [1210.0555](#) [[hep-ph](#)].
- [22] N. Arkani-Hamed and S. Dimopoulos, *Supersymmetric unification without low energy supersymmetry and signatures for fine-tuning at the LHC*, *JHEP* **06** (2005) 073, arXiv: [hep-th/0405159](#) [[hep-th](#)].
- [23] J. L. Hewett et al., *Signatures of long-lived gluinos in split supersymmetry*,
JHEP **09** (2004) 070, arXiv: [hep-ph/0408248](#).
- [24] ATLAS Collaboration, *Generation and Simulation of R -Hadrons in the ATLAS Experiment*,
ATL-PHYS-PUB-2019-019, 2019, URL: <https://cds.cern.ch/record/2676309>.
- [25] CMS Collaboration,
Search for long-lived charged particles in proton-proton collisions at $\sqrt{s} = 13$ TeV,
Phys. Rev. D **94** (2016) 112004, arXiv: [1609.08382](#) [[hep-ex](#)].
- [26] CMS Collaboration, *Search for decays of stopped exotic long-lived particles produced in proton-proton collisions at $\sqrt{s} = 13$ TeV*, *JHEP* **05** (2018) 127, arXiv: [1801.00359](#) [[hep-ex](#)].
- [27] CMS Collaboration, *Search for long-lived particles decaying to jets with displaced vertices in proton-proton collisions at $\sqrt{s} = 13$ TeV*, (2021), arXiv: [2104.13474](#) [[hep-ex](#)].

-
- [28] ATLAS Collaboration, *Reinterpretation of searches for supersymmetry in models with variable R -parity-violating coupling strength and long-lived R -hadrons*, ATLAS-CONF-2018-003, 2018, URL: <https://cds.cern.ch/record/2308391>.
- [29] ATLAS Collaboration, *SUSY March 2022 Summary Plot Update*, (2022), URL: <https://atlas.web.cern.ch/Atlas/GROUPS/PHYSICS/PUBNOTES/ATL-PHYS-PUB-2022-013/>.
- [30] T. Gherghetta, G. F. Giudice and J. D. Wells, *Phenomenological consequences of supersymmetry with anomaly induced masses*, Nucl. Phys. **B559** (1999) 27, arXiv: [hep-ph/9904378](https://arxiv.org/abs/hep-ph/9904378) [hep-ph].
- [31] G. F. Giudice et al., *Gaugino mass without singlets*, JHEP **12** (1998) 027, arXiv: [hep-ph/9810442](https://arxiv.org/abs/hep-ph/9810442) [hep-ph].
- [32] ATLAS Collaboration, *Search for metastable heavy charged particles with large ionization energy loss in pp collisions at $\sqrt{s} = 13$ TeV using the ATLAS experiment*, Phys. Rev. D **93** (2016) 112015, arXiv: [1604.04520](https://arxiv.org/abs/1604.04520) [hep-ex].
- [33] ATLAS Collaboration, *Search for long-lived charginos based on a disappearing-track signature using 136 fb^{-1} of pp collisions at $\sqrt{s} = 13$ TeV with the ATLAS detector*, (2022), arXiv: [2201.02472](https://arxiv.org/abs/2201.02472) [hep-ex].
- [34] CMS Collaboration, *Search for disappearing tracks in proton–proton collisions at $\sqrt{s} = 13$ TeV*, Phys. Lett. B **806** (2020) 135502, arXiv: [2004.05153](https://arxiv.org/abs/2004.05153) [hep-ex].
- [35] ATLAS Collaboration, *Search for long-lived charginos based on a disappearing-track signature using 136 fb^{-1} of pp collisions at $\sqrt{s} = 13$ TeV with the ATLAS detector*, ATLAS-CONF-2021-015, 2021, URL: <https://cds.cern.ch/record/2759676>.
- [36] H. Fukuda et al., *Higgsino dark matter or not: Role of disappearing track searches at the LHC and future colliders*, Physics Letters B **781** (2018) 306, URL: <https://doi.org/10.1016%2Fj.physletb.2018.03.088>.
- [37] H. Fukuda et al., *Higgsino Dark Matter or Not: Role of Disappearing Track Searches at the LHC and Future Colliders*, Phys. Lett. B **781** (2018) 306, arXiv: [1703.09675](https://arxiv.org/abs/1703.09675) [hep-ph].
- [38] CMS Collaboration, *Search for natural and split supersymmetry in proton–proton collisions at $\sqrt{s} = 13$ TeV in final states with jets and missing transverse momentum*, JHEP **05** (2018) 025, arXiv: [1802.02110](https://arxiv.org/abs/1802.02110) [hep-ex].
- [39] CMS Collaboration, *Search for new phenomena with the M_{T2} variable in the all-hadronic final state produced in proton–proton collisions at $\sqrt{s} = 13$ TeV*, Eur. Phys. J. C **77** (2017) 710, arXiv: [1705.04650](https://arxiv.org/abs/1705.04650) [hep-ex].

- [40] ATLAS Collaboration, *Search for R-parity-violating supersymmetric particles in multi-jet final states produced in pp collisions at $\sqrt{s} = 13$ TeV using the ATLAS detector at the LHC*, *Phys. Lett. B* **785** (2018) 136, arXiv: 1804.03568 [hep-ex].
- [41] CMS Collaboration, *Search for direct top squark pair production in events with one lepton, jets, and missing transverse momentum at 13 TeV with the CMS experiment*, *Journal of High Energy Physics* **2020** (2020),
URL: <https://doi.org/10.1007/2Fjhep05%282020%29032>.
- [42] CMS Collaboration,
Searches for pair production of third-generation squarks in $\sqrt{s} = 13$ TeV pp collisions, *Eur. Phys. J. C* **77** (2017) 327, arXiv: 1612.03877 [hep-ex].
- [43] D. Curtin et al., *Long-lived particles at the energy frontier: the MATHUSLA physics case*, *Reports on Progress in Physics* **82** (2019) 116201,
URL: <https://doi.org/10.1088/1361-6633/ab28d6>.
- [44] J. Alimena et al.,
Searching for long-lived particles beyond the Standard Model at the Large Hadron Collider, *Journal of Physics G: Nuclear and Particle Physics* **47** (2020) 090501,
URL: <https://doi.org/10.1088/1361-6471/ab4574>.
- [45] M. Garny et al., *Top-philic dark matter within and beyond the WIMP paradigm*, *Phys. Rev. D* **97** (2018) 075002, arXiv: 1802.00814 [hep-ph].
- [46] M. Garny et al., *Coannihilation without chemical equilibrium*, *Phys. Rev. D* **96** (2017) 103521, arXiv: 1705.09292 [hep-ph].
- [47] M. J. Strassler and K. M. Zurek, *Echoes of a hidden valley at hadron colliders*, *Phys. Lett. B* **651** (2007) 374, arXiv: hep-ph/0604261 [hep-ph].
- [48] Z. Chacko, H.-S. Goh and R. Harnik,
The Twin Higgs: Natural electroweak breaking from mirror symmetry, *Phys. Rev. Lett.* **96** (2006) 231802, arXiv: hep-ph/0506256 [hep-ph].
- [49] D. Feldman, Z. Liu and P. Nath, *The Stueckelberg Z-prime Extension with Kinetic Mixing and Milli-Charged Dark Matter From the Hidden Sector*, *Phys. Rev. D* **75** (2007) 115001, arXiv: hep-ph/0702123 [HEP-PH].
- [50] M. Cvetič, P. Langacker and G. Shiu,
Phenomenology of a three family standard like string model, *Phys. Rev. D* **66** (2002) 066004, arXiv: hep-ph/0205252 [hep-ph].
- [51] L. Evans and P. Bryant, *LHC Machine*, *Journal of Instrumentation* **3** (2008) S08001,
URL: <https://doi.org/10.1088/1748-0221/3/08/s08001>.

-
- [52] A. Ijspeert, M. Karppinen and J. Salminen, *Design of corrector magnets for the LHC*, (1999), URL: <https://cds.cern.ch/record/451608>.
- [53] E. Mobs, *The CERN accelerator complex - 2019. Complexe des accélérateurs du CERN - 2019*, (2019), General Photo, URL: <https://cds.cern.ch/record/2684277>.
- [54] F. D. Aaron et al., *Combined measurement and QCD analysis of the inclusive $e \pm p$ scattering cross sections at HERA*, *Journal of High Energy Physics* **2010** (2010), ISSN: 1029-8479, URL: [http://dx.doi.org/10.1007/JHEP01\(2010\)109](http://dx.doi.org/10.1007/JHEP01(2010)109).
- [55] R. D. Ball et al., *Parton distributions from high-precision collider data*, *The European Physical Journal C* **77** (2017), URL: <https://doi.org/10.1140%2Fepjc%2Fs10052-017-5199-5>.
- [56] V. N. Gribov and L. N. Lipatov, *Deep inelastic $e p$ scattering in perturbation theory*, *Sov. J. Nucl. Phys.* **15** (1972) 438.
- [57] T. Gleisberg et al., *Event generation with SHERPA 1.1*, *Journal of High Energy Physics* **2009** (2009) 007, URL: <https://doi.org/10.1088%2F1126-6708%2F2009%2F02%2F007>.
- [58] ATLAS Collaboration, *Luminosity determination in pp collisions at $\sqrt{s} = 13$ TeV using the ATLAS detector at the LHC*, ATLAS-CONF-2019-021, 2019, URL: <https://cds.cern.ch/record/2677054>.
- [59] ATLAS Collaboration, *ATLAS data quality operations and performance for 2015–2018 data-taking*, *JINST* **15** (2020) P04003, arXiv: 1911.04632 [physics.ins-det].
- [60] *ATLAS Luminosity Results in Run 2*, CERN, 2021, URL: <https://twiki.cern.ch/twiki/bin/view/AtlasPublic/LuminosityPublicResultsRun2>.
- [61] A. Buckley et al., *General-purpose event generators for LHC physics*, *Physics Reports* **504** (2011) 145, ISSN: 0370-1573, URL: <http://dx.doi.org/10.1016/j.physrep.2011.03.005>.
- [62] J. Alwall et al., *The automated computation of tree-level and next-to-leading order differential cross sections, and their matching to parton shower simulations*, *JHEP* **07** (2014) 079, arXiv: 1405.0301 [hep-ph].
- [63] T. Sjöstrand et al., *An introduction to PYTHIA 8.2*, *Comput. Phys. Commun.* **191** (2015) 159, arXiv: 1410.3012 [hep-ph].
- [64] ATLAS Collaboration, *The ATLAS Simulation Infrastructure*, *Eur. Phys. J. C* **70** (2010) 823, arXiv: 1005.4568 [physics.ins-det].

- [65] R. Bruce et al., *LHC Run 2: Results and challenges*, (2016) MOAM5P50. 7 p,
URL: <https://cds.cern.ch/record/2201447>.
- [66] L. R. F. Castillo, *The Search and Discovery of the Higgs Boson*, 2053-2571,
Morgan & Claypool Publishers, 2015, ISBN: 978-1-6817-4078-2,
URL: <http://dx.doi.org/10.1088/978-1-6817-4078-2>.
- [67] ATLAS Collaboration,
Study of the material of the ATLAS inner detector for Run 2 of the LHC,
JINST **12** (2017) P12009, arXiv: 1707.02826 [hep-ex].
- [68] ATLAS Collaboration, *ATLAS Insertable B-Layer Technical Design Report*,
ATLAS-TDR-19; CERN-LHCC-2010-013, 2010,
URL: <https://cds.cern.ch/record/1291633>, Addendum: ATLAS-TDR-19-ADD-1;
CERN-LHCC-2012-009, 2012, URL: <https://cds.cern.ch/record/1451888>.
- [69] M. Livan and R. Wigmans, *Misconceptions about Calorimetry*, 2017,
arXiv: 1704.00661 [physics.ins-det].
- [70] C. W. Fabjan and F. Gianotti, *Calorimetry for particle physics*,
Rev. Mod. Phys. **75** (2003) 1243.
- [71] ATLAS Collaboration, *ATLAS tile calorimeter: Technical design report*, (1996).
- [72] ATLAS Collaboration, *Operation of the ATLAS trigger system in Run 2*,
JINST **15** (2020) P10004, arXiv: 2007.12539 [hep-ex].
- [73] R. Achenbach et al., *The ATLAS Level-1 Calorimeter Trigger*, CERN, 2008,
URL: <https://cds.cern.ch/record/1080560>.
- [74] F. Anulli et al., *The Level-1 Trigger Muon Barrel System of the ATLAS experiment at CERN*,
Journal of Instrumentation **4** (2009) P04010,
URL: <https://doi.org/10.1088/1748-0221/4/04/p04010>.
- [75] ATLAS Collaboration, *Performance of the ATLAS Level-1 topological trigger in Run 2*,
(2021), arXiv: 2105.01416 [hep-ex].
- [76] ATLAS Collaboration, *Performance of the ATLAS trigger system in 2015*,
Eur. Phys. J. C **77** (2017) 317, arXiv: 1611.09661 [hep-ex].
- [77] ATLAS Collaboration, *Performance of the ATLAS track reconstruction algorithms in dense environments in LHC Run 2*, *Eur. Phys. J. C* **77** (2017) 673, arXiv: 1704.07983 [hep-ex].
- [78] R. Frühwirth, *Application of Kalman filtering to track and vertex fitting*,
Nuclear Instruments and Methods in Physics Research Section A: Accelerators, Spectrometers, Detectors and Associated Equipment **262** (1987) 444, ISSN: 0168-9002, URL: <https://www.sciencedirect.com/science/article/pii/0168900287908874>.

-
- [79] T. Cornelissen et al., *The new ATLAS track reconstruction (NEWT)*, *Journal of Physics: Conference Series* **119** (2008) 032014, URL: <https://doi.org/10.1088/1742-6596/119/3/032014>.
- [80] ATLAS Collaboration, *Performance of the reconstruction of large impact parameter tracks in the inner detector of ATLAS*, ATL-PHYS-PUB-2017-014, 2017, URL: <https://cds.cern.ch/record/2275635>.
- [81] ATLAS Collaboration, *Reconstruction of primary vertices at the ATLAS experiment in Run 1 proton–proton collisions at the LHC*, *Eur. Phys. J. C* **77** (2017) 332, arXiv: [1611.10235 \[hep-ex\]](#).
- [82] ATLAS Collaboration, *Vertex Reconstruction Performance of the ATLAS Detector at $\sqrt{s} = 13$ TeV*, ATL-PHYS-PUB-2015-026, 2015, URL: <https://cds.cern.ch/record/2037717>.
- [83] ATLAS Collaboration, *Performance of vertex reconstruction algorithms for detection of new long-lived particle decays within the ATLAS inner detector*, ATL-PHYS-PUB-2019-013, 2019, URL: <https://cds.cern.ch/record/2669425>.
- [84] V. Kostyukhin, *VKalVrt - package for vertex reconstruction in ATLAS*, ATL-PHYS-2003-031, 2003, URL: <https://cds.cern.ch/record/685551>.
- [85] ATLAS Collaboration, *Jet energy scale and resolution measured in proton–proton collisions at $\sqrt{s} = 13$ TeV with the ATLAS detector*, *Eur. Phys. J. C* **81** (2020) 689, arXiv: [2007.02645 \[hep-ex\]](#).
- [86] ATLAS Collaboration, *Topological cell clustering in the ATLAS calorimeters and its performance in LHC Run 1*, *Eur. Phys. J. C* **77** (2017) 490, arXiv: [1603.02934 \[hep-ex\]](#).
- [87] M. Cacciari, G. P. Salam and G. Soyez, *The anti- k_t jet clustering algorithm*, *JHEP* **04** (2008) 063, arXiv: [0802.1189 \[hep-ph\]](#).
- [88] ATLAS Collaboration, *Beam-induced and cosmic-ray backgrounds observed in the ATLAS detector during the LHC 2012 proton–proton running period*, *JINST* **11** (2016) P05013, arXiv: [1603.09202 \[hep-ex\]](#).
- [89] ATLAS Collaboration, *Selection of jets produced in 13 TeV proton–proton collisions with the ATLAS detector*, ATLAS-CONF-2015-029, 2015, URL: <https://cds.cern.ch/record/2037702>.
- [90] ATLAS Collaboration, *ATLAS b -jet identification performance and efficiency measurement with $t\bar{t}$ events in pp collisions at $\sqrt{s} = 13$ TeV*, *Eur. Phys. J. C* **79** (2019) 970, arXiv: [1907.05120 \[hep-ex\]](#).

- [91] C. Patrignani et al., *Review of Particle Physics*, *Chin. Phys.* **C40** (2016) 100001.
- [92] ATLAS Collaboration, *Electron and photon performance measurements with the ATLAS detector using the 2015–2017 LHC proton–proton collision data*, *JINST* **14** (2019) P12006, arXiv: [1908.00005 \[hep-ex\]](#).
- [93] ATLAS Collaboration, *Electron and photon energy calibration with the ATLAS detector using 2015–2016 LHC proton–proton collision data*, *JINST* **14** (2019) P03017, arXiv: [1812.03848 \[hep-ex\]](#).
- [94] ATLAS Collaboration, *Muon reconstruction and identification efficiency in ATLAS using the full Run 2 pp collision data set at $\sqrt{s} = 13$ TeV*, *Eur. Phys. J. C* **81** (2020) 578, arXiv: [2012.00578 \[hep-ex\]](#).
- [95] ATLAS Collaboration, *E_T^{miss} performance in the ATLAS detector using 2015–2016 LHC pp collisions*, ATLAS-CONF-2018-023, 2018, URL: <https://cds.cern.ch/record/2625233>.
- [96] ATLAS Collaboration, *ATLAS Fast TracKer (FTK): Technical Design Report*, ATLAS-TDR-021; CERN-LHCC-2013-007, 2013, URL: <https://cds.cern.ch/record/1552953>.
- [97] CMS Collaboration, *The Phase-2 Upgrade of the CMS L1 Trigger Interim Technical Design Report*, CERN, 2017, URL: <https://cds.cern.ch/record/2283192>.
- [98] ATLAS Collaboration, *Technical Design Report for the Phase-II Upgrade of the ATLAS Trigger and Data Acquisition System - Event Filter Tracking Amendment*, CERN, 2022, URL: <https://cds.cern.ch/record/2802799>.
- [99] ATLAS Collaboration, *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** (2016) 581, arXiv: [1510.03823 \[hep-ex\]](#).
- [100] ATLAS Collaboration, *The ATLAS Fast TracKer system*, *JINST* **16** (2021) P07006, arXiv: [2101.05078 \[hep-ex\]](#).
- [101] J. Adelman et al., *ATLAS FTK Challenge: Simulation of a Billion-fold Hardware Parallelism*, *Procedia Computer Science* **66** (2015) 540, 4th International Young Scientist Conference on Computational Science, ISSN: 1877-0509, URL: <https://www.sciencedirect.com/science/article/pii/S1877050915034109>.
- [102] A. Haas, *ATLAS Simulation using Real Data: Embedding and Overlay*, *Journal of Physics: Conference Series* **898** (2017) 042004, URL: <https://doi.org/10.1088/1742-6596/898/4/042004>.

-
- [103] ATLAS Collaboration, *FTK Public Results*,
URL: <https://twiki.cern.ch/twiki/bin/view/AtlasPublic/FTKPublicResults>.
- [104] ATLAS Collaboration, *Alignment of the ATLAS Inner Detector in Run-2*,
Eur. Phys. J. C **80** (2020) 1194, arXiv: 2007.07624 [hep-ex].
- [105] ATLAS Collaboration, *Study of the mechanical stability of the ATLAS Insertable B-Layer*,
ATL-INDET-PUB-2015-001, 2015, URL: <https://cds.cern.ch/record/2022587>.
- [106] A. Salzburger, *The ATLAS Track Extrapolation Package*, CERN, 2007,
URL: <https://cds.cern.ch/record/1038100>.
- [107] ATLAS Collaboration, *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** (2015) 072004, arXiv: 1504.05162 [hep-ex].
- [108] ATLAS Collaboration, *Search for displaced vertices arising from decays of new heavy particles in 7 TeV pp collisions at ATLAS*, *Phys. Lett. B* **707** (2012) 478,
arXiv: 1109.2242 [hep-ex].
- [109] ATLAS Collaboration,
Search for long-lived, massive particles in events with a displaced vertex and a muon with large impact parameter in pp collisions at $\sqrt{s} = 13$ TeV with the ATLAS detector,
Phys. Rev. D **102** (2020) 032006, arXiv: 2003.11956 [hep-ex].
- [110] CMS Collaboration, *Search for R-parity violating supersymmetry with displaced vertices in proton–proton collisions at $\sqrt{s} = 8$ TeV*, *Phys. Rev. D* **95** (2017) 012009,
arXiv: 1610.05133 [hep-ex].
- [111] CMS Collaboration, *Search for long-lived particles with displaced vertices in multijet events in proton–proton collisions at $\sqrt{s} = 13$ TeV*, *Phys. Rev. D* **98** (2018) 092011,
arXiv: 1808.03078 [hep-ex].
- [112] J. Butterworth et al., *PDF4LHC recommendations for LHC Run II*,
Journal of Physics G: Nuclear and Particle Physics **43** (2016) 023001, ISSN: 1361-6471,
URL: <http://dx.doi.org/10.1088/0954-3899/43/2/023001>.
- [113] B. Fuks et al.,
Gaugino production in proton-proton collisions at a center-of-mass energy of 8 TeV,
JHEP **10** (2012) 081, arXiv: 1207.2159 [hep-ph].
- [114] B. Fuks et al., *Precision predictions for electroweak superpartner production at hadron colliders with RESUMMINO*, *Eur. Phys. J. C* **73** (2013) 2480, arXiv: 1304.0790 [hep-ph].
- [115] ATLAS Collaboration, *Multijet simulation for 13 TeV ATLAS Analyses*,
ATL-PHYS-PUB-2019-017, 2019, URL: <https://cds.cern.ch/record/2672252>.

- [116] ATLAS Collaboration, *Electron reconstruction and identification in the ATLAS experiment using the 2015 and 2016 LHC proton–proton collision data at $\sqrt{s} = 13$ TeV*, *Eur. Phys. J. C* **79** (2019) 639, arXiv: 1902.04655 [hep-ex].
- [117] ATLAS Collaboration, *Muon reconstruction performance of the ATLAS detector in proton–proton collision data at $\sqrt{s} = 13$ TeV*, *Eur. Phys. J. C* **76** (2016) 292, arXiv: 1603.05598 [hep-ex].
- [118] R. D. Cousins, *Lectures on Statistics in Theory: Prelude to Statistics in Practice*, (2018), arXiv: 1807.05996 [physics.data-an].
- [119] G. Cowan et al., *Asymptotic formulae for likelihood-based tests of new physics*, *The European Physical Journal C* **71** (2011), ISSN: 1434-6052, URL: <http://dx.doi.org/10.1140/epjc/s10052-011-1554-0>.
- [120] A. L. Read, *Presentation of search results: the CL_s technique*, *Journal of Physics G: Nuclear and Particle Physics* **28** (2002) 2693, URL: <https://doi.org/10.1088/0954-3899/28/10/313>.
- [121] L. Heinrich et al., *pyhf: pure-Python implementation of HistFactory statistical models*, *Journal of Open Source Software* **6** (2021) 2823, URL: <https://doi.org/10.21105/joss.02823>.
- [122] ATLAS Collaboration, *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. D* **97** (2018) 052012, arXiv: 1710.04901 [hep-ex].
- [123] D. Curtin and C. B. Verhaaren, *Discovering uncolored naturalness in exotic Higgs decays*, *Journal of High Energy Physics* **2015** (2015) 1, URL: <https://doi.org/10.1007%2Fjhep12%282015%29072>.
- [124] P. Schwaller, D. Stolarski and A. Weiler, *Emerging jets*, *Journal of High Energy Physics* **2015** (2015), URL: <https://doi.org/10.1007%2Fjhep05%282015%29059>.

List of Figures

2.1	Particles in the Standard Model	4
2.2	Higgs Potential	8
2.3	SM particle lifetimes	10
2.4	Quantum loop corrections to Higgs mass	11
2.5	GMSB exclusion	17
2.6	Current mass-vs-lifetime constraints	19
2.7	Wino/Higgsino exclusion limits	21
2.8	Exclusion limits on RPV λ''	22
2.9	Conversion driven freeze-out	23
3.1	CERN accelerator complex	26
3.2	Parton distribution functions	27
3.3	Proton-proton collision	29
3.4	Luminosity collected by ATLAS	30
3.5	Schematic of ATLAS detector	34
3.6	Particle signatures in ATLAS	35
3.7	Magnet system	37
3.8	ATLAS Inner detector	38
3.9	ATLAS Calorimeters	39
3.10	ATLAS Muon spectrometer	42
3.11	Muon trigger chambers	43
4.1	Track parameters	48
4.2	Track reconstruction efficiency in jets	50
4.4	LRT track parameter residuals	52
4.5	Hit pattern requirements	54
4.6	Secondary vertexing algorithm performance	56
4.7	Secondary vertex position resolution	57
4.8	Jet performance	59

4.9	Jet energy resolution	60
5.1	FTK track-finding illustration	65
5.2	Coverage of pattern bank	69
5.3	FTK variable-width pattern illustration	69
5.4	FTK track-finding efficiency vs. number of patterns	71
5.5	FTK track-finding efficiency vs. η , before and after partitioning	72
5.6	FTK track-finding efficiency vs. d_0 and z_0	73
5.7	FTK track-finding efficiency vs. η and p_T	73
5.8	FTK track-finding efficiency vs. ϕ	74
5.9	Beamspot position	75
5.10	FTK track-finding efficiency vs. beamspot position	77
5.11	FTK-track residuals due to IBL distortion	79
6.1	Benchmark models	82
6.2	Filter turn-on curve	87
6.3	Baseline jet selection efficiency	88
6.4	Material map	89
6.5	p_T and d_0 -significance of DV-tracks	91
6.6	p_T vs. $\Delta\alpha$ of DV-tracks	91
6.7	DV selection efficiency	92
6.8	Signal cutflow	93
6.9	Track jet distributions in different regions	94
6.10	Test of jet-DV correlations	95
6.11	ΔR between DVs and closest track jet	96
6.12	Jet-DV probability vs. number of tracks in jet	96
6.13	Jet-DV probability vs. jet- p_T	97
6.14	Jet-DV probability used to estimate background in SR	98
6.15	Inclusive background estimate in High- p_T SR	99
6.16	Inclusive background estimate in Trackless SR	99
6.17	Uncertainty calculations on inclusive estimate	101
6.18	Validation of inclusive estimate in High- p_T SR sideband regions	102
6.19	Validation of inclusive estimate in Trackless SR sideband regions	102
6.20	Inclusive background estimate inside of material (data)	103
6.21	Inclusive background estimate in simulated dijet events	105
6.22	AX track database kinematics	107
6.23	Shape of AX mass templates	108
6.24	Impact of pointing vector selection on K_S^0 mass peak	110

6.25	Pointing vector of 2-track and 3-track DVs	110
6.26	K_S^0 mass peaks	111
6.27	Crossing factors	111
6.28	Single-AX mass templates	113
6.29	Merging rate	115
6.30	Merging rate	116
6.31	Distance significance distributions	117
6.32	Distribution of $\Delta\phi$ between pairs of DVs	118
6.33	Merged vertex mass templates	119
6.34	Closure test in simulated dijet events	121
6.35	Combined background estimate in High- p_T regions	122
6.36	Combined background estimate in Trackless regions	122
6.37	Difference in K_S^0 yield in data and simulated events	125
6.38	Uncertainty on signal efficiency due to LRT	125
6.39	Cutflow in data and signal	127
6.40	Unbinned m_{DV} distributions	127
6.41	Expected event yields (Strong RPV)	130
6.42	Expected event yields (EWK RPV)	130
6.43	Upper limits on cross section (Strong RPV model)	131
6.44	Excluded Strong RPV parameter space	132
6.45	Excluded Strong RPV parameter space	133
6.46	Exclusion limits (EWK RPV model)	133

List of Tables

2.1	SUSY particle content (1)	13
2.2	SUSY particle content (2)	13
3.1	Performance goals of the ATLAS Detector	34
3.2	Parts of the ATLAS calorimeters	40
3.3	Parts of the ATLAS muon spectrometer	42
4.1	Track parameters	48
4.2	Track selection in tracking algorithms	51
5.1	FTK track-finding efficiency for various pattern bank configurations	73
5.2	Dataflow of 128-PU pattern banks	74
6.1	Jet- p_T thresholds of DRAW filter and baseline jet selections	86
6.2	Selection on displaced vertices	89
6.3	DV-track cleaning selection	90
6.4	Summary of signal region definitions	92
6.5	Estimate of inclusive background in SR with uncertainties	100
6.6	Radial regions used in accidental crossings estimate	106
6.7	Crossing factors	112
6.8	Combined background estimates outside of material (simulation)	120
6.9	Combined background estimates inside of material (simulation)	120
6.10	Combined background estimates in the signal regions	123
6.11	Observed events in signal regions	126
6.12	Cutflow in data	126
6.13	Treatment of uncertainties in fit	131