

Displaced Vertex Search For Heavy Neutral Leptons Using The ATLAS Detector

by

Dominique Trischuk

B.Sc., The University of Toronto, 2017

A DISSERTATION SUBMITTED IN PARTIAL FULFILLMENT OF
THE REQUIREMENTS FOR THE DEGREE OF

DOCTOR OF PHILOSOPHY

in

THE FACULTY OF GRADUATE AND POSTDOCTORAL STUDIES

(Physics)

THE UNIVERSITY OF BRITISH COLUMBIA

(Vancouver)

February 2022

© Dominique Trischuk 2022



The following individuals certify that they have read, and recommend to the Faculty of Graduate and Postdoctoral Studies for acceptance, the dissertation entitled:

**Displaced Vertex Search For Heavy Neutral Leptons Using The ATLAS
Detector**

submitted by	Dominique Trischuk	in partial fulfillment of the requirements for
the degree of	Doctor of Philosophy	
in	Physics	

Examining Committee:

Colin Gay, Physics & Astronomy, UBC

Supervisor

Alison Lister, Physics & Astronomy, UBC

Co-Supervisor

Douglas Bryman, Physics & Astronomy, UBC

University Examiner

Roman Krems, Chemistry, UBC

University Examiner

Additional Committee Members:

Gary Hinshaw, Physics & Astronomy, UBC

Supervisory Committee Member

Ariel Zhitnitsky, Physics & Astronomy, UBC

Supervisory Committee Member

Abstract

This dissertation presents a search for long-lived heavy neutral leptons (HNLs) in proton-proton collisions at the Large Hadron Collider (LHC). The Standard Model (SM) of particle physics is an extremely successful theory and many of its major predictions have been precisely confirmed. However, the existence of neutrinos, with small nonzero masses, suggests that the SM is incomplete. Introducing HNLs into the SM is a natural way to generate the light neutrino masses through a seesaw mechanism. Theories that postulate the existence of such particles can also explain the asymmetry between matter and anti-matter in our universe and models with at least three HNLs provide a dark matter candidate. This experimental search uses ATLAS data collected between 2015 and 2018 at a centre-of-mass energy of 13 TeV. A non-standard technique is used to search for a displaced vertex from particle trajectories produced in the HNL decay to leptons. The dominant background from uncorrelated leptons crossing in the ATLAS detector is estimated using an object shuffling method. The reconstructed HNL mass is used to discriminate between signal and background. No excess of events is observed and constraints on the strength of the interactions between HNLs and neutrinos are imposed in various scenarios.

This dissertation also presents new methods to study the readout system and performance of a silicon strip tracking detector. The LHC is currently undergoing upgrades that will enable it to produce more than ten times the data that has already been collected. To meet the requirements of this challenging new environment, an all-silicon particle tracking system will be installed in ATLAS.

Lay Summary

The Standard Model (SM) of particle physics describes the fundamental building blocks of matter and their interactions. While many of its major predictions have been precisely confirmed, the SM is incomplete. For instance, the SM predicts that neutrinos, a type of fundamental particle, are massless. Thus, the observation of nonzero neutrino masses suggest that the SM is incomplete. This dissertation uses data from the ATLAS Experiment at the Large Hadron Collider to search for a particle that explains the very tiny neutrino masses. The analysis is designed to search for clues left by this theoretical particle after it travels a long distance through the ATLAS detectors. No evidence for the existence of a new particle is found. Various signal hypotheses are ruled out and this analysis thereby provides guidance to future searches about where new physics that explains fundamental characteristics about our universe might be found.

Preface

The research in this dissertation is based on ATLAS experiment data. ATLAS is a global collaboration of physicists, engineers, technicians, students and support staff. The collective effort of over 5500 members make the operation of the experiment, including the data collection, possible. For example, collaboration members are responsible for the development and maintenance of reconstruction software for physics objects such as muons and electrons. Other collaboration members are responsible for the proton collision simulation software. This simulation also relies on input from theory and generator experts outside of the collaboration. This dissertation would not have been possible without the collective effort of all members.

The ATLAS collaboration releases both public results in peer-reviewed journals and preliminary results that are not externally reviewed. However, all public and preliminary results are extensively reviewed within the collaboration. All of the nearly 3000 qualified authors are listed alphabetically on each publication or public result from the ATLAS collaboration. Figures presented in this text have labels that indicate the level to which they have been reviewed. Figures taken from a peer-reviewed publication are labelled with “ATLAS” or “ATLAS Simulation”. Figures taken from preliminary results are labelled with “ATLAS Preliminary”. Figures that are not approved by the collaboration are labelled with “ATLAS Work in Progress” or do not have a label are all prepared by me, with the exception of statistical analysis results in section 5.11, which were prepared by C. Appelt. The results are all undergoing the final stages of the internal ATLAS review process, after which they will be made public. All other figures not prepared by me have their references indicated in the captions.

Abbreviations are used in this dissertation to refer to a variety of objects and concepts ranging from hardware components to physics object reconstruction algorithms. Discipline-specific terms are also used throughout the text. These abbreviations and terms are defined in the glossary and accompanied by a short description. Throughout the text you will find references to the terms listed in the glossary.

I wrote all of the text for this dissertation. Chapters 5 and 6 are written by me, but parts of the text also appear in internal ATLAS documentations. Chapters 1 to 4 present material that serves as an introduction for the reader. Chapters 5 and 6 summarize my work. This has led to two journal publications, one of which (at the time of writing) is undergoing the final stages of the internal ATLAS review process. The remainder of this preface describes

my relevant contributions in each area. The search for displaced heavy neutral leptons is presented in chapter 5. This search is performed using 139 fb^{-1} of proton-proton collision data collected with the ATLAS detector. The analysis is approved internally by the ATLAS “Exotics” group and a journal publication is undergoing the final stages of the internal ATLAS review process. My work was performed under the supervision of A. Lister with guidance from M. Danninger. The analysis team consists of 16 members and I am one of three main analyzers in this search. Collaborators who helped me in this analysis are C. Appelt, J. Burzynski, D. Casper, S. Galantzan, C. Gay, M. Lutz, S. Oda, R. Newhouse, O. Ruchayskiy, J. Tastet and M. Wielers. This list below outlines my main contributions to the analysis. In addition, I provided support and feedback to others, including in group meetings.

- I was responsible for validating and coordinating the production of all signal samples used in the analysis.
- In collaboration with theory associates J. Tastet and O. Ruchayskiy, I developed the method used to interpret the results for the eight different signal models (see section 5.3). I also implemented and validated the method in our analysis framework.
- In collaboration with J. Tastet, I developed the method used to take into account spin correlations in the signal model (see section 5.4.3). I also implemented and validated the method in our analysis framework. M. Wielers provided assistance with the production of signal samples for the validation.
- I optimized the secondary vertexing algorithm described in section 5.5.2. D. Casper provided guidance.
- I developed the HNL mass reconstruction described in section 5.5.3 following guidance from C. Gay.
- I developed a custom derivation and coordinated the processing of all the data and simulated samples for the analysis. A derivation step is used in the data analysis framework to skim away information that is not required for the analysis (see section 5.7.1).
- I developed the analysis event selection code (see fig. 5.18). I also contributed to the development of the code that creates custom data structures (ntuples) for the relevant information in the search. R. Newhouse and J. Burzynski are the main developers of this ntuple maker code. C. Appelt developed the statistical analysis code.
- I designed and optimized the event selection (see section 5.7). S. Oda and R. Newhouse developed the filter and filter emulator respectively. I also optimized the additional analysis regions defined in section 5.7.4.

- I performed the studies of the main backgrounds and optimized the selections that remove non-random background events (see sections 5.8.2 to 5.8.4).
- I developed the novel technique to estimate the probability of uncorrelated particles crossing in the ATLAS inner detector described in section 5.8.6. I also developed the methods to estimate the background uncertainties.
- In collaboration with S. Galantzan, I implemented another random crossing estimate using a track shuffling method (see appendix A.4). I integrated the shuffling algorithm into the main analysis framework, performed the validation studies and performed the small-scale data test.
- I am the main author and editor of the internal analysis documentation. C. Appelt, M. Lutz, R. Newhouse and M. Danninger also contributed to the document.

To be listed as an author on ATLAS publication, collaboration members must perform at least 80 days technical work. My qualification work is described in section 6.2. I developed a temperature stress test for the ATLAS Inner Tracker (ITk) silicon strip detectors. This quality assurance test provides valuable information about the design and construction of the ITk-Strip detectors that will be included in future design reviews. My supervisor for this work was N. Hessey and my collaborators F. Guescini and T. Stack provided guidance and support. This work used a large-scale detector object (petal) that was designed and built by members of the ITk-Strip collaboration.

- I worked with industrial partners to select a climate chamber that performed the temperature stress tests.
- I set up a monitoring system to study the environment in the chamber during the temperature cycling in collaboration with F. Guescini.
- I built a small-scale prototype detector object (petalet) in collaboration with F. Guescini.
- I developed the temperature cycling program for the stress test.
- I wrote the measurement programs that evaluated the dimensional stability of the detector components.
- I performed all the temperature cycling tests and analyzed the results.
- I wrote the internal documentation.

My contribution to the ITk-Strip readout is described in section 6.3. My work was performed under the supervision of C. Gay. To study the demands of the detector readout:

- I developed a firmware emulator that mimics both the detector and the readout system (see section 6.3.1). My collaborators W. Fedorko, M. Trovato and M. Warren also provided guidance and support.
- I was part of a small team who set up a system that, for the first time, read out information from a new generation of on-detector electronic chips. This work was done in collaboration with B. Gallop, K. Krizka, S. Stucci, J.J Teoh and W. Wu.
- I built one of only a handful of readout systems in our lab at UBC that uses many of the final components that will be part of the ITk-Strip readout system at CERN.

This readout effort led to the following journal publication:

ATLAS ITk Collaboration. The ABC130 barrel module prototyping programme for the ATLAS strip tracker. [*Journal of Instrumentation*, 15:P09004](#), 2020.

In addition to the contributions listed above, I have presented my work, and other's work on their behalf, in a number of internal ATLAS collaboration meetings and regional conferences:

D. Trischuk. Recent Results from UEH Searches, ATLAS Physics and Performance Week: Exotics Plenary, online, 2021.

D. Trischuk. Semi-Electrical Petal Thermocycling, ATLAS ITk Week, online, 2021.

D. Trischuk. Secondary Vertexing with Displaced Leptons, ATLAS Week: Early-Career-Scientists Session, online, 2021.

D. Trischuk. Search for Displaced Heavy Neutral Leptons, Virtual Physics Conference, online, 2020.

B. Gallop, C. Meyer, C. Gay, D. Trischuk[†], E. Zhivun, J. Pinzino, J. Beacham, J.J. Teoh, O. Arnaez, P. Phillips, S. Stucci, T. Calvet and W. Wu. ITk Strips FELIX Demonstrators, ATLAS ITk Week, Geneva, Switzerland, 2019.

D. Trischuk. ATLAS Strip Detector Upgrade for the Inner Tracker: Tackling Big Data Readout Systems, Winter Nuclear & Particle Physics Conference, Banff, Canada, 2019.

D. Trischuk[†], W. Fedorko and C. Gay. ITK Strips HCCStar Emulator, ATLAS ITk Week, Geneva, Switzerland, 2018.

[†]Indicates the speaker.

Table of Contents

Abstract	iii
Lay Summary	iv
Preface	v
Table of Contents	ix
List of Tables	xii
List of Figures	xv
Glossary	.xxvi
Acknowledgements	xxxv
1 The Standard Model	1
1.1 Particles in the Standard Model	1
1.2 Chiral Fermions	3
1.3 Structure of the Standard Model	3
1.4 Electroweak Symmetry Breaking	5
1.5 Successes and Limitations of the Standard Model	6
2 Heavy Neutral Leptons	8
2.1 Neutrino Masses From Right-Handed Neutrinos	8
2.2 Interactions with the Standard Model	10
2.3 HNL Decays	12
2.4 Lepton Number Violation	12
2.5 HNL Cross Section	13
3 The ATLAS Detector at the Large Hadron Collider	17
3.1 The Large Hadron Collider	17
3.1.1 LHC Filling Scheme	18
3.1.2 Luminosity	18

3.1.3	Dataset	19
3.1.4	Pileup	20
3.1.5	High-Luminosity LHC	20
3.2	The ATLAS Detector	21
3.2.1	Coordinate System	21
3.2.2	Detector Overview	21
3.2.3	Inner Detector	22
3.2.4	Calorimeter System	26
3.2.5	Muon Spectrometer	27
3.2.6	Trigger and Data Acquisition	29
3.2.7	ATLAS Detector Simulation	31
4	Object Reconstruction	32
4.1	Standard Reconstruction	32
4.1.1	Overview	32
4.1.2	Tracks, Primary Vertices and Energy Clusters	33
4.1.3	Electrons	36
4.1.4	Muons	40
4.2	Non-Standard Reconstruction	42
4.2.1	Displaced Tracks	42
4.2.2	Displaced Leptons	44
4.2.3	Displaced Vertices	45
5	Search for Heavy Neutral Leptons	50
5.1	Previous Searches for HNLs	50
5.2	Search Overview	53
5.3	HNL Models	54
5.4	Data and Signal Samples	55
5.4.1	Data Samples	55
5.4.2	Signal Samples	55
5.4.3	Modelling Spin Correlations	58
5.4.4	Signal Kinematics	61
5.5	Event Reconstruction	64
5.5.1	Displaced Vertex Reconstruction for the HNL Signal	64
5.5.2	Custom Displaced Vertex Configuration	66
5.5.3	Reconstructed HNL Mass	74
5.6	Object Definitions	78
5.7	Event Selection	79

Table of Contents

5.7.1	Analysis Framework	80
5.7.2	Pre-Selections	81
5.7.3	Signal Region Selection	84
5.7.4	Additional Regions	87
5.8	Backgrounds	88
5.8.1	Random Track Crossings	89
5.8.2	Material Interactions	89
5.8.3	Metastable Particle Decays	92
5.8.4	Leptonic Z Boson Decays	96
5.8.5	Muons from Cosmic-Ray Interactions	97
5.8.6	Data-Driven Estimate of Random Track Crossings Using a Prompt Lepton Shuffling Method	98
5.9	Signal Systematic Uncertainties	107
5.10	Statistical Analysis	110
5.11	Results	112
6	ATLAS Inner Tracker Upgrade	120
6.1	ATLAS Inner Tracker	120
6.1.1	Inner Tracker Strip Detector	121
6.2	Quality Assurance for ITk-Strip Petals	122
6.2.1	Thermocycling Program	123
6.2.2	Thermocycling the Petalet	125
6.2.3	Thermocycling the Semi-Electrical Petal	128
6.3	ITk-Strip Readout	138
6.3.1	HCCStar Emulator	139
7	Conclusion	142
	Bibliography	144

Appendices

A	Data-driven Estimate of Random Track Crossings Using a Track Shuffling Method	152
A.1	Shuffling Analysis Framework	153
A.2	Track Shuffling Problem	154
A.3	Validation of the Track Shuffling Method	155
A.4	Small-Scale Test of the Track Shuffling Method	157

List of Tables

1.1	The four fundamental forces with their relative strengths and mediators. Table from [1].	2
1.2	The Standard Model (SM) fermion fields and their electric charges Q , third components of the weak isospin T^3 and hypercharges Y	4
2.1	The relative strength of the mixing, $x_\alpha = U_\alpha ^2/ U ^2$, with the three neutrino flavours for the inverted (IH) and normal (NH) hierarchy benchmark and single-flavour mixing scenarios.	11
4.1	Main differences between standard and large radius tracking (LRT).	43
5.1	A summary of the eight models used in the interpretation of the analysis results. Models are defined according to the number of HNLs (1SFH or 2QDH), the fraction of LNC and LNV decays and the relative strength of the mixing between the lepton generations. The inverted (IH) and normal (NH) hierarchy benchmarks are used for 2QDH models only, while the 1SFH models have single-flavour mixing.	54
5.2	The cross section for LNC and LNV decays for one Majorana HNL, one Dirac HNL, a quasi-Dirac HNL pair with “Majorana-like” decays and a quasi-Dirac HNL pair with “Dirac-like” decays. The total cross section is the sum of the LNC and LNV contributions.	55
5.3	The main differences in the track selection between the default VSI configuration and the lepton-only selection used in the custom VSI configuration in this analysis.	67
5.4	Single muon and electron triggers from 2015-2018 that are used to select events of interest.	82
5.5	The HNL filter and overall filter rate for a set of test events in each of the four data taking years (2015-2018).	83

5.6	Requirements of the four HNL analysis filters. Definitions of the lepton selections (e.g. ID or isolation requirements) used in the filter can be found in the descriptions of the lepton reconstruction algorithms in section 4.1.3 and section 4.1.4.	83
5.7	Signal region selections that are applied in each channel.	86
5.8	Description of the selection criteria used to defined the SR. All criteria are further explained within this chapter.	86
5.9	Definition of the inverted-prompt-lepton VR. A prompt lepton is defined as any muon or electron object with <i>Loose</i> or better quality that has $ d_0 < 3$ mm and $ z_0 \sin(\theta) < 0.5$ mm.	88
5.10	Description of the heavy-flavour decay veto.	96
5.11	Description of the Z mass veto.	97
5.12	The relative systematic uncertainties on the background estimate, estimated by creating shuffled events from SS DVs or from prompt leptons taken from events with different DV channels. The total uncertainty is the quadrature sum of the two contributions. The uncertainties are treated as uncorrelated among the different channels.	101
5.13	Estimated background predictions with statistical and systematic uncertainties in the signal region and the control region. The observed yields are shown for OS DVs and SS DVs.	102
5.14	The KS test probabilities obtained when comparing the shuffled and SS background estimate HNL mass (m_{HNL}) and tri-lepton mass (m_{lll}) distributions. The KS test is done separately for each channel. Two distributions are considered incompatible if the KS probability is much less than one.	107
5.15	Table summarizing the systematic uncertainties that impact the signal prediction. Systematic uncertainties in this table are the impacts on the signal yields in the SR.	107
5.16	Post-fit background predictions, observed event-yields and scale factors μ_b with statistical and systematic uncertainties in the control and signal regions. The results are shown for the 2QDH model with inverted-hierarchy mixing in the “Majorana-limit”. The results are also shown for all final states and for the 1SFH Dirac-limit, LNC configuration $l_\alpha^\pm - l_\alpha^\mp l_\gamma^\pm$	113
6.1	Measurement results for flatness of Side A and B, as well as the bow for the sensor on side A.	127
6.2	Measurement uncertainties for the features on the semi-electrical petal.	132
6.3	Overview of the custom made ASICs for the ITk Strip Detector [91].	139

A.1	The different regions used in the random track crossing background estimation.	153
-----	--	-----

List of Figures

1.1	The elementary particles that make up the Standard Model. Figure adapted from [2].	2
1.2	The potential $V(\Phi) = \mu^2 \Phi ^2 + \lambda \Phi ^4$ in a theory without (a) and with (b) local symmetry breaking around a minimum of the potential.	5
1.3	Summary of several Standard Model cross-section measurements. The measurements are corrected for branching fractions, compared to the corresponding theoretical expectations. Figure taken from [8].	7
2.1	The elementary particles that make up the Standard Model with the addition of $\mathcal{N} = 3$ right-handed sterile neutrinos. Figure adapted from [2].	9
2.2	A plot that shows the allowed combinations of $ U_\alpha ^2/ U ^2$ that are consistent with the recent NuFIT [21] fit to neutrino oscillation data. The allowed regions under normal (blue) and inverted (green) hierarchy are shown. The red cross markers represent two benchmark scenarios that are defined for the two different hierarchy assumptions. Figure adapted from [20].	11
2.3	Feynman diagrams for HNL production from the decay of a W boson and the HNL decay into leptonic and semi-leptonic final states.	12
2.4	Lepton number conserving (LNC) and lepton number violating (LNV) diagrams for the charge-current leptonic HNL decay.	13
2.5	Branching ratio of $N \rightarrow ll\nu$ with a sum of all neutrino flavours as a function of m_N for muon-only (a), electron-only (b), NH (c) and IH (d) mixing scenarios. In each mixing scenario, the branching ratio is shown for $N \rightarrow ee\nu$ (blue), $N \rightarrow e\mu\nu$ (orange), $N \rightarrow \mu e\nu$ (pink) and $N \rightarrow \mu\mu\nu$ (purple) decays.	14
2.6	The total mixing $ U ^2$ as a function of HNL mass m_N for the muon-only mixing scenario. Each $ U ^2$ is computed for m_N between 2 and 20 GeV with proper decay lengths $c\tau_N = 1, 10$ and 100 mm.	15

3.1	An illustration of various experimental facilities and beam lines at the CERN accelerator complex relevant to the LHC. The protons are accelerated through the injection chain and into the LHC. The path they follow is: LINAC2 $\rightarrow$ BOOSTER $\rightarrow$ Proton Synchrotron (PS) $\rightarrow$ Super Proton Synchrotron (SPS) $\rightarrow$ LHC. Figure adapted from [34].	18
3.2	Total integrated luminosity versus time (a) delivered to ATLAS (green), recorded by ATLAS (yellow), and certified to be good quality data (blue) during Run 2. Distribution of the mean number of interactions per bunch crossing (b) recorded by the ATLAS experiment in the years 2015-2018. Figures taken from [37]. . .	19
3.3	The LHC/HL-LHC timeline and run schedule as of Fall 2021. This diagram includes the planned shutdown and upgrade periods. Figure taken from [38]. . .	20
3.4	A cut-away view of the the ATLAS detector. Figure taken from [39].	22
3.5	A transverse view in the R - ϕ plane of the barrel region of the inner detector. Tracks from the proton collisions traverse outward in radius, with most of them passing through the four silicon pixel layers, four silicon-microstrip (SCT) layers and approximately 36 straw tubes (TRT). In this figure, the fourth innermost layer of the pixel detector (IBL) is not shown. Figure taken from [39].	23
3.6	A longitudinal view of the active region of the pixel detector. In this figure, the fourth innermost layer of the pixel detector (IBL) is not shown. Figure taken from [41].	24
3.7	An SCT module showing the 4 silicon detectors connected to a readout hybrid assembly. Figure taken from [40].	25
3.8	A cut-away view of the ATLAS calorimeter system. Figure taken from [39]. . .	26
3.9	A cut-away view of the ATLAS muon system. Figure taken from [39].	28
3.10	The ATLAS trigger and data acquisition system (TDAQ) system in Run 2 showing the various parts of the system relevant to triggering, detector readout and data flow. Figure adapted from [47].	30
4.1	A schematic of different particles interacting with the ATLAS detector in a section of the $x - y$ plane. Figure adapted from [52].	33
4.2	A schematic of the path of an electron through the ATLAS detector. The solid red trajectory shows the path of the electron traversing through the tracking system (pixel, SCT and TRT layers) and entering the electromagnetic calorimeter. The dashed red trajectory indicates the path of a photon produced by the interaction of the electron with the material in the tracking system. Figure taken from [59].	37

4.3	The total reconstruction efficiency for simulated electrons in a single-electron sample as a function of the true transverse energy, E_T , for each step of the electron reconstruction: the energy cluster reconstruction (red triangles), track reconstruction using the Global χ^2 Track Fitter (blue open circles), cluster and track reconstruction using GSF tracking (yellow squares), and the final electron-candidate reconstruction, which includes the track-to-cluster matching (black circles). Since the cluster reconstruction requires that cluster seeds have $E_T > 2.5$ GeV, the total reconstruction efficiency is less than 60% below 4.5 GeV (dashed line). Figure taken from [59].	38
4.4	The electron identification efficiency in $Z \rightarrow ee$ events in data as a function of E_T . The efficiencies are obtained by applying data-to-simulation efficiency ratios measured in $J/\psi \rightarrow ee$ and $Z \rightarrow ee$ events to $Z \rightarrow ee$ simulation. The <i>LHMedium</i> operating point efficiency curve is plotted using red squares. Figure taken from [60].	39
4.5	The muon reconstruction efficiency for the <i>Medium</i> identification working point as a function of muon p_T in the region $0.1 < \eta < 2.5$. This reconstruction efficiency is obtained with $Z \rightarrow \mu\mu$ and $J/\Psi \rightarrow \mu\mu$ events. The error bars on the efficiencies indicate the statistical uncertainty. The panel at the bottom shows the ratio of the measured to predicted efficiencies, with statistical and systematic uncertainties. Figure taken from [64].	41
4.6	Track reconstruction efficiency for displaced charged particles produced by the decay of long-lived signal particles. The efficiency for standard tracking, LRT and the sum of the two is shown as a function of the radius of production of the displaced particles ($r_{\text{prod}}^2 = x^2 + y^2$). Displaced leptons are shown on the left and displaced hadrons on the right. Figure taken from [65].	44
4.7	The main steps in the VSI displaced vertexing algorithm.	45
4.8	A schematic illustration of the fake tracks rejection method. In example (a), the vertex is between the pixel B-Layer and Layer-1. The tracks of the reconstructed secondary vertex must not have hits in the layers behind the vertex (i.e. IBL and B-Layer), and must have hits in the closest layer ahead of the vertex (i.e. Layer-1). In example (b), the vertex is close to the Pixel B-Layer. The tracks in this vertex are not allowed to have hits in the IBL, but may have hits in the B-Layer, and must have hits on the Layer-1. Analogous requirements are made on vertices close to the other layers. Figure taken from [66].	47
4.9	An example of finding N -track vertices using an incompatibility graph. The layout of tracks and vertices in an event (a) and the incompatibility graph (b) that is formed by connecting two tracks when the two-track vertices are determined to be incompatible. Figure taken from [67].	48

4.10	An example of track sharing by multiple vertices. Figure taken from [67]. . . .	49
5.1	Observed 95% confidence level (CL) exclusion limits for HNL to muon neutrino mixing angle $ U_\mu ^2$ versus HNL mass m_N for the prompt signature (the region above the black line is excluded) and the displaced signature (the region enclosed by the red line is excluded). The solid lines show limits assuming lepton-number violation (LNV) for half of the decays and lepton-number conservation (LNC) for the other half. The long-dashed line shows the limits assuming LNC for all decays. The green and yellow bands indicate the ranges of expected limits obtained within ± 1 and ± 2 standard deviations of the median limit reflecting uncertainties in the signal and background yields. Figure taken from [70]. . . .	51
5.2	Sensitivity to HNLs with dominant mixing to first, second and third lepton generations. Current bounds on the mixing (filled regions) and prospects for new results in the next 5-15 years are shown as a function of the HNL mass m_N . Figures taken from [71].	52
5.3	Feynman diagrams included in the combinations for the muon-only mixing ($\mu-\mu\mu$, $\mu-\mu e$ and $\mu-ee$) and electron-only mixing ($e-e\mu$, $e-ee$ and $e-\mu\mu$) scenarios. Neutral current decays are also possible for the $\mu-\mu\mu$ and $e-ee$ channels.	56
5.4	Additional Feynman diagrams for signal processes that contribute to the signal strength if both $ U_\mu ^2$ and $ U_e ^2$ are nonzero. Diagrams are shown for the charged-current interaction only.	58
5.5	Comparison of the lepton p_T distributions for LNC decays between MADGRAPH samples (purple) and the weighted (pink) and unweighted (gray) PYTHIA8 samples for a $\mu-\mu\mu$ signal using the NC+CC event weights. Samples are generated with $m_N = 10$ GeV and $c\tau_N = 10$ mm signal sample.	60
5.6	Comparison of the lepton p_T distributions for LNC decays between MADGRAPH samples (blue) and the weighted (orange) and unweighted (gray) PYTHIA8 samples for a $\mu-\mu e$ signal using the CC event weights. Samples are generated with $m_N = 10$ GeV and $c\tau_N = 10$ mm signal sample.	61
5.7	Simulated lepton p_T distributions for the three charged leptons in the HNL event and the invariant mass of the tri-lepton system (m_{lll}). The leading displaced lepton is defined as the lepton in the DV with the largest p_T , while the subleading displaced lepton has the smaller p_T . Simulated $\mu-\mu\mu$ signals are plotted for three different mass points: $m_N = 5, 7.5, 10$ GeV with $c\tau_N = 10$ mm.	62
5.8	HNL kinematics for simulated $\mu-\mu\mu$ signals for three different mass points: $m_N = 5, 7.5, 10$ GeV with $c\tau_N = 10$ mm.	63

5.9	The DV mass (a) and DV radius (b) distributions for simulated $\mu\text{--}\mu\mu$ signal samples. For the DV mass distribution, three different mass points are shown: $m_N = 5, 7.5, 10$ GeV with $c\tau_N = 10$ mm. For the DV radius plot, the simulated samples have $m_N = 10$ GeV with $c\tau_N = 1, 10$ and 100 mm.	63
5.10	The $\mu\text{--}\mu\mu$ signal efficiency (a) and ratio of the signal efficiency between VSI Leptons and VSI (b) as a function of simulated HNL mass with $3 < m_N < 20$ GeV. Signal samples with $c\tau_N = 1, 10, 100$ mm are shown in red, blue and green respectively. The signal efficiency is defined as the ratio of the selected reconstructed events to the total number of simulated events.	65
5.11	Simulated events for a $\mu\text{--}\mu\mu$, 10 GeV, 10 mm signal in purple and the data-driven same-sign background in grey as a function of the reconstructed HNL mass. The total number of events in the signal and background can be found in the legends in each plot. Events are selected by applying the pre-selection cuts described in section 5.7.	66
5.12	The number of $\mu\text{--}\mu\mu$ 10 GeV, 10 mm signal (a) and two-muon background (b) events for default VSI (blue) and VSI Leptons (orange) vertex configurations. The original number of events is given by the default configuration. Events labels “LRTR3” are the number of events after modifying the track selections as described in table 5.3 without the $\Delta\theta$ and $ z_0 $ correlation selection. The “LRTR3 w Geo cut”, which is the final selection applied in the custom vertex configuration, is the number of events including this final correlation selection. The signal significance (c) is estimated using the number of signal events divided by the square root of the number of background events ($s/\sqrt{b}$).	68
5.13	Signal events for a 10 GeV, 10 mm $\mu\text{--}\mu\mu$ sample (a) and two-muon background events (b) as a function of the two-track forming d_0 cut for VSI (blue) and VSI Leptons (orange) vertex configurations. The original number of events is given by the default configuration. Then the number of events with two-track forming d_0 cut threshold equal to 2, 1.5, 1, 0.5 and 0 mm is shown. The signal significance ($s/\sqrt{b}$) for 10 GeV, $\mu\text{--}\mu\mu$ signal is shown for the 10 mm (c) and 1 mm (d) signal samples as a function of the same d_0 cut.	69
5.14	Signal (a) and background (b) events with exactly two muons in the displaced vertex as a function of the number of tracks (n_{tracks}) in the vertex reconstructed using the VSI configuration with a 1 mm two-track forming d_0 cut. Vertices with $n_{\text{tracks}} = 2$ are selected in the signal region, while events with $n_{\text{tracks}} > 2$ are rejected.	70

5.15	Ratio of the number of selected signal events reconstructed with the custom configuration to the default VSI configuration. Ratios are shown for the 10 GeV signal samples in the $\mu\text{--}\mu\mu$, $\mu\text{--}\mu e$, $e\text{--}e\mu$ and $e\text{--}ee$ channels as a function of the proper decay length.	71
5.16	The three components of the vertex reconstruction efficiency and the total efficiency, measured as a function of radial decay position, L_{xy} , for $\mu\text{--}\mu\mu$, $\mu\text{--}ee$ and $e\text{--}ee$ signals with $m_N=10$ GeV and $c\tau_N=10$ mm.	73
5.17	Two HNL mass distributions for a simulated $\mu\text{--}\mu\mu$ with $m_N = 10$ GeV and $c\tau_N = 10$ mm. The m_{HNL} is computed using m_W fixed to the pole mass of 80.4 GeV (blue) and the median W mass in the allowed kinematic region (purple). Events where the HNL mass has no real solution are set to -1.	77
5.18	The analysis framework pipeline that shows how data and simulated signal events are processed in the displaced HNL analysis.	82
5.19	OS displaced vertices in the VR shown in bins of m_{DV} versus L_{xy} . All vertices are required to pass the VR selections. Leptons in the selected DVs are also required to pass <i>VeryVeryLoose</i> electron and <i>Medium</i> muon identification working points.	90
5.20	OS displaced vertices from data events using the VR selections as a function of L_{xy} . Displaced vertices are separated by flavour into three channels: ee , $e\mu$ and $\mu\mu$. Distributions labeled “All” (solid circles) have no additional selections applied. The distribution with open circles in the ee channel has only the material veto applied and distributions with open squares have the heavy-flavour decay veto applied.	91
5.21	OS and SS displaced vertices in the Validation Region as a function of the DV mass. Leptons in the DVs are required to pass the <i>Medium</i> (<i>VeryVeryLoose</i>) quality working points for muons (electrons). The material map veto is also applied for the ee channel. The ratio of OS to SS vertices is shown in the bottom ratio plot.	93
5.22	Correlation plot of DV mass (m_{DV}) versus radius ($L_{xy} = r_{\text{DV}}$) for the difference between OS and SS vertices in the three DV type channels. The data events studied in these plots are selected from events in the VR. The heavy-flavour decay veto applied in each channel is shown by the solid red line. The z -axis maximum in the plots is saturated at 30 in order to provide a colour scale that can distinguish regions with small negative entries (more SS than OS DVs) from those with small positive entries (more OS than SS DVs).	94
5.23	OS and SS displaced vertices with the heavy-flavour decay veto applied for data events in the VR. The material veto is also applied for displaced ee vertices. . .	95

5.24	The invariant mass between the prompt lepton and same-flavour displaced lepton (m_{ll}) for selected VR data events in the Z mass veto study. Events around the Z mass with $80 \text{ GeV} < m_{ll} < 100 \text{ GeV}$ are identified as Z decays and are removed from the SR.	97
5.25	The HNL mass (m_{HNL}) shapes for the background estimate (shaded purple), DV systematic (solid blue line), prompt lepton systematic (dashed blue line), observed SS events (grey triangles) and observed OS events (black circles) for the six different channels.	100
5.26	SS displaced vertices in the VR with the heavy-flavour decay veto applied. The material veto for ee DVs is also applied. SS DVs are separated by track charge into vertices with two positively charged tracks (SS $++$) and two negatively charged tracks (SS $--$). This study does not include the $\Delta p_T/p_T$ cut for displaced electrons described in section 5.6.	103
5.27	HNL mass distribution for the SS background estimate (orig. plep) and for the shuffled background estimate (random plep), where plep is an abbreviation for “prompt lepton”. The shuffling factor, f , which is defined as the ratio of the number of shuffled events to number of original SS background events, is shown on the plots for each channel.	105
5.28	Tri-lepton mass distribution for the SS background estimate (orig. plep) and for the shuffled background estimate in colours (random plep), where plep is an abbreviation for “prompt lepton”. The shuffling factor, f , which is defined as the ratio of the number of shuffled events to number of original SS background events, is shown on the plots for each channel.	106
5.29	Pre-fit (a-b) and post-fit (c-d) plots for the 2QDH model with inverted-hierarchy (IH) mixing in the “Dirac-limit” (left column) and “Majorana-limit” (right column). The signal region (SR) includes events with $m_{\text{HNL}} < 20 \text{ GeV}$ and the control region (CR) includes events with $20 < m_{\text{HNL}} < 50 \text{ GeV}$	113
5.30	Pre-fit plots for the 1SFH model with muon-only mixing (a-b), electron-only mixing (c-d) and the 2QDH model with normal hierarchy (NH) mixing (e-f) in the “Dirac-limit” (left column) and “Majorana-limit” (right column). The signal region (SR) includes events with $m_{\text{HNL}} < 20 \text{ GeV}$ and the control region (CR) includes events with $20 < m_{\text{HNL}} < 50 \text{ GeV}$	114
5.31	Post-fit plots for the 1SFH model with muon-only mixing (a-b), electron-only mixing (c-d) and the 2QDH model with normal hierarchy (NH) mixing (e-f) in the “Dirac-limit” (left column) and “Majorana-limit” (right column). The signal region (SR) includes events with $m_{\text{HNL}} < 20 \text{ GeV}$ and the control region (CR) includes events with $20 < m_{\text{HNL}} < 50 \text{ GeV}$	115

5.32	The m_{HNL} distribution in the SR and CR for the observed data and the post-fit background for the 2QDH model with inverted-hierarchy mixing in the “Majorana-limit”. Simulated signal samples with three different mass points are also shown normalized to the pre-fit signal strength. After the global fit is performed, the signal strength is zero.	116
5.33	The expected (dashed line) and observed (solid line) 95% confidence level (CL) contours for the 1SFH model with muon-only (a) and electron-only (b) mixing in the “Dirac-limit” and the observed 95% CL contours for the four HNL models listed in table 5.1 in the “Dirac-limit” (c) and “Majorana limit” (d). The limits are shown for the mixing angle ($ U ^2$) as a function of the HNL mass (m_N). The regions enclosed by the contours are excluded. The observed contours for the 1SFH models are also shown for the 1997 DELPHI [68] and the 2019 ATLAS [70] results.	117
5.34	The expected (dashed line) and observed (solid line) 95% CL contours for the mixing angle, $ U ^2$, in the 1SFH model with muon-only (a) and electron-only (b) mixing in the “Majorana-limit” and the contours for the 2QDH model with NH (c-d) and IH (e-f) mixing in the “Dirac-limit” and “Majorana-limit” as a function of the HNL mass, m_N . The green and yellow bands indicate the ranges of expected limits obtained within ± 1 and ± 2 standard deviations of the median limit reflecting uncertainties in the signal and background yields. The regions enclosed by the contours are excluded.	119
6.1	A longitudinal view of the active regions of the ITk detector. The pixel detector is shown in the inner region (smaller radius) and is surrounded by the strip detector. The strip end-cap disks can be seen on either side of the barrel region, coloured with blue and green stripes. Figure taken from [91].	120
6.2	An expanded view of a strip module in the barrel. Petal modules consist of the same basic components, but with different hybrid and sensor geometries. Figure take from [91].	121
6.3	A full sketch of a strip petal and the modules labelled R0-R5. The sensors, in blue, are glued onto the core surface shown in yellow. Light green and yellow chips are housed on the green hybrids and glued onto the sensors. The power boards are shown in orange. Figure modified from [91].	122

6.4	The final validated thermocycling program (a) with a ramp rate of $1.33^{\circ}\text{C}/\text{min}$ and soak times of 15 hours and 30 min at 21°C and $\pm 40^{\circ}\text{C}$ respectively. This plot shows the temperature and RH for the full 46-hour program. The red curves show the measured temperature inside the climate chamber and the blue curves show the RH inside the chamber. The temperature and RH curves (b) for the first thermocycle where the RH spike is the largest. The plot in (b) also indicates the dew point inside the chamber (dashed red line) and the measured pressure differential (green line), which is defined as $\Delta P = P_{\text{in}} - P_{\text{out}}$.	124
6.5	The location of the stable markers glued directly to the carbon surface labeled 1-5.	125
6.6	The petalet hanging in the climate chamber right before thermocycling. . . .	126
6.7	Plots of the surface scans of the sensor on side A and the bus-tape surface of side B before and after thermocycling. The flatness, bow and other sensor measurements are also shown in the text on the right-hand side of the plots. . .	127
6.8	Plot of the displacement vectors ($\Delta \vec{R}$) for the features on side A, including the two hybrids in purple, the sensor in blue and the bus-tape in brown. The vector magnitudes are uniformly scaled by a factor of 2000 to visually examine if there is a common direction for the displacements.	128
6.9	A photo of the first semi-electrical petal. This petal was the prototype used in the thermocycling tests discussed in this section. This photo was taken by the TRIUMF communications team.	129
6.10	The petal coordinate system. The x - and y -axes are defined using the two petal fiducials on the inner and outer petal locators. The origin is defined as the inner fiducial and the x - y plane is defined as the flat surface of the petal locators. . .	130
6.11	Pictures of the different fiducial markers on the sensors, hybrids and bus-tape. These markers are measured by the CMM to determine the position of the various components on the semi-electrical petal.	131
6.12	The support bar in the vertical orientation in the climate chamber. The three attachment points on the bar are where the three petal locators are attached and secured to hang the petal in the chamber.	132
6.13	The petal hanging in the climate chamber before and during thermocycling. . .	133
6.14	Arrow plot and ΔR displacements of the four sensor fiducials (in blue) and the four fiducials on the hybrid (in purple) for one of the R5 sensors on side A. The arrows in the plot are magnified by a factor of 2000 to visually see the direction of motion. The magnitudes and uncertainties of the displacements are shown in the boxes.	134

6.15	Statistical analysis of the roughly 100 different points measured on Side A of the petal. The pull in the y -direction ($\Delta y/\sigma_y$) and the pull x -direction ($\Delta x/\sigma_x$) are shown immediately after thermocycling and the measurements are repeated 48 hours and four weeks later. After 48 hours, the mean of the pull is closer to zero, which indicates that after some time the structures relax and are more consistent with no movement after thermocycling.	134
6.16	The crack on the R4 sensor and the bowing of a hybrid on the same module. Photos are taken immediately after thermocycling the semi-electrical petal 100 times.	136
6.17	Arrow plot and displacements (ΔR and Δz) of the four fiducials (in purple) for one of the hybrids on side A. The arrows in the plot are scaled by a factor of 2000 to visually see the direction of motion. The magnitudes and uncertainties of the displacements are shown in the boxes.	137
6.18	Statistical analysis of the roughly 100 different points measured on Side A of the petal. The pull in the y -direction ($\Delta y/\sigma_y$) and the pull x -direction ($\Delta x/\sigma_x$) are shown immediately after thermocycling and two weeks later. After two weeks, the central distribution is more consistent with zero, except for outliers with large pulls that are from the measurements of the hybrid on the module where the failure occurred.	137
6.19	Current and HL-LHC era DAQ systems for the ATLAS experiment. The FELIX cards will replace the current readout drivers (RODs) and readout systems (ROs) used in the current DAQ system. Figure adapted from [93].	138
6.20	A diagram of the ITk-Strip readout showing how all the custom ASICs (ABC-Star, HCCStar and AMAC) are connected to each other and merged in the end-of-structure (EoS) card before being converted to an optical signal and sent to FELIX.	139
6.21	A sketch of the HCCStar emulator and LCB wrapper integrated into the FELIG data generation structure.	140
A.1	An illustration of the main steps of the shuffling analysis framework that demonstrated the types of files are needed as inputs to each step.	154
A.2	Displaced vertices as a function of DV mass and radius for events produced using the “do-nothing” configuration (in purple) and the original data events (in blue). In the “do-nothing” configuration, the original non-selected tracks are also saved such that, the only difference between each event is the primary vertex. Differences at low L_{xy} can be observed, from the locations of the primary vertices.	156

A.3	Displaced vertices as a function of DV mass and radius (L_{xy}) for events produced using the “do-nothing” configuration with the original non-selected tracks (in purple) and non-selected tracks from host data events (in pink). Differences between the distributions show that using hadrons from a random host-event misses some correlations between the non-selected tracks in the original event and the lepton-only displaced vertices.	157
A.4	DV mass distribution for shuffled and “do-nothing” events for SS and OS displaced vertices. The ratio between shuffled and not shuffled events is the shuffling factor. For this test, with only one collision run from 2016, the expected shuffling factor of 18.8 is the dashed black line in the ratio plot. OS and SS vertices are directly compared in (c) and (d) for the not shuffled and shuffled events respectively.	158
A.5	Data transfer volume on the worldwide LHC computing network as a function of time. The two large spikes in this plot can be attributed to the large amounts of data being transferred for this track shuffling estimate. Figure obtained from the computing monitoring systems.	159

Glossary

1SFH refers to an HNL model used to interpret the results of the analysis described in this thesis. This model has one heavy lepton state that mixes 100% of the time with a single neutrino flavour.

2QDH refers to an HNL model used to interpret the results of the analysis described in this thesis. This HNL model has two quasi-degenerate heavy lepton states with $m_1 \sim m_2$. The lepton states can mix with any flavour of neutrino..

ABCStar refers to one of the custom front-end chips for the ATLAS Inner Tracker. It's role is to convert incoming charge signal into hit information. "ABC" stands for ATLAS Binary Chip and "Star" refers to the independent pathways from each ABCStar to the HCCStar that create a star shape.

AFII refers to the ATLAS Fast Simulation that is used to simulate events in the ATLAS detector. This fast simulation includes a parametrization of the calorimeter response in order to reduce the simulation time.

ALICE is a detector dedicated to studying heavy-ion physics and is one of the four main LHC experiments. It stands for A Large Ion Collider Experiment.

AMAC refers to one of the custom front-end chips for the ATLAS Inner Tracker. It provides monitoring and interrupt functionality. "AMAC" stands for Autonomous Monitor and Control.

ASIC refers to a chip with customized integrated circuits that are designed for particular use cases. It stands for Application-Specific Integrated Circuits.

ATLAS is a general purpose detector and is one of the four main LHC experiments. This detector is used to collect the data analyzed in this thesis. It stands for A Toroidal LHC Apparatus.

BCR refers to a 1-bit command sent to all front-end chips that is synchronous with the LHC bunch crossings. This command resets the bunch crossing counter to zero and is used to ensure front-end synchronization. It stands for Bunch Counter Reset.

Booster is the proton synchrotron booster. It is one of the accelerators in the LHC proton injection chain.

BRAM refers to a module in an field programmable gate array (FPGA) that provides on-chip storage for large amounts of data. It stands for Block Random Access Memory.

Bremsstrahlung is electromagnetic radiation produced by the deceleration or deflection of a charged particle passing through matter near the strong electric fields of nuclei.

BSM refers to a beyond the Standard Model theory that proposes descriptions of physical phenomena that are not currently explained by the Standard Model.

Bus-tape provides the modules in the ATLAS Inner Tracker with low and high-voltage power. Modules are glued directly onto the bus-tape.

Calorimeter-tagged muon refers to a track in the ATLAS ID that matches an energy deposit in the calorimeter consistent with a minimum ionizing particle. This type of muon has the lowest purity of all the muon types.

CC is a charged current interaction mediated by a W^+ or W^- boson.

CERN is the Centre for European Nuclear Research. It is a European research organization that operates the largest particle physics laboratory in the world.

CFF is a common-file format used in the output of the custom track shuffling algorithm. This file consists of a list of tracks and a set of shuffling instructions that maps various combinations of tracks into unique fake events.

CL refers to a Confidence Level. This is the probability that, if the test were repeated many times, the same outcome would be observed.

CMM refers to a coordinate measuring machine. It is used to measure the location in three-dimensional space of an object with respect to a reference point.

CMS is a general purpose detector and is one of the four main LHC experiments. It stands for Compact Muon Solenoid.

Combined muon refers to a track that is reconstructed using a global fit of the hits in the ATLAS ID and MS. During the fit, MS hits may be added or removed in order to improve the fit quality. An *outside-in* approach is used for the reconstruction, starting with the reconstruction of tracks in the MS and then extrapolating inwards to match MS to ID tracks.

- CP** refers to a combined performance group in the ATLAS collaboration that provides recommendations on how to reconstruct and identify standard physics objects.
- CPU** refers to a central processing unit that consists of electronic circuitry that executes instructions from a computer program.
- CR** refers to the region in the HNL search used to constrain the background estimate in the global HNL mass fit. It is defined as the region that passes all the event selections, except that the m_{HNL} is required to be between 20 and 50 GeV.
- CSC** refers to a cathode-strip chamber in the ATLAS MS that is primarily responsible for tracking and measuring the momentum of muons.
- DAQ** refers to a hardware and software system that is used to collect data. It stands for Data Acquisition.
- DELPHI** is a general purpose detector that was one of the four main experiments at the Large Electron-Positron collider built at CERN to study electroweak interactions.
- DP** is the temperature of an environment below which, given the moisture in the chamber, water droplets will start to form. It stands for Dew Point.
- DV** refers to a place of common origin for two or more tracks that is displaced with respect to the primary proton interaction point. It stands for Displaced Vertex.
- ECal** refers to the Electromagnetic Calorimeter that is part of the ATLAS calorimeter system and is primarily responsible for measuring the energy of electrons and photons.
- EoS** refers to the End of Structure cards on each large-scale structure in the ATLAS Inner Tracker strip detector. This card contains, among other things, a custom chip that transmits the clock and control signals and receives data from the front-end chips.
- FCal** refers to the ATLAS Forward Calorimeter that is primarily responsible for measuring the energy of particles in the forward region of ATLAS.
- FE** refers to a front-end readout chip responsible for buffering and reading out data from detectors.
- FELIG** refers to the Front-End Link Generator card used to mimic the ATLAS Inner Tracker Strip readout in the HCCStar emulator that is described in this thesis.
- FELIX** refers to the Front-End Link Exchange card that will be used in LHC Run 3 (see timeline in fig. 3.3) and beyond to send and receive data from front-end chips on the detectors.

Fiducial marker refers to an object used by an imaging machine to determine the location and orientation of an object. For example, during petal mounting, fiducial markers are used to locate the modules and place them onto the petal surface.

FPGA refers to a field programmable gate array that consists of integrated circuits designed to be configurable after manufacturing.

GBT refers to a data protocol developed by CERN for use in data transmission for experiments requiring high data rate links and electronic components capable of sustaining high radiation doses. It stands for GigaBit Transceiver.

Geant4 is a geometry and tracking platform used to simulate the passage of particles through matter. This software toolkit is used to simulate the ATLAS detector system.

GSF refers to a Gaussian-Sum Filter used to re-fit tracks matched to electrons. The GSF takes into account non-linear effects in the trajectory of an electron related to bremsstrahlung.

Hard-scatter refers to the primary interaction of interest in each proton bunch crossing. It is defined as the interaction with the largest momentum transfer.

HCCStar refers to one of the custom front-end chips for the ATLAS Inner Tracker strip detectors. It is the digital interface between the hybrids and the EoS card. “HCC” stands for Hybrid Controller Chip and “Star” refers to the architecture that enables the HCC to build and transmit module wide events.

HEC refers to the Hadronic End-Cap Calorimeter that is part of the ATLAS calorimeter system. It is primarily responsible for measuring the energy of hadrons.

HL-LHC is the High-Luminosity Large Hadron Collider, which is the upgraded LHC machine that will produce proton collisions with ten times the current luminosity.

HLT refers to the high-level trigger and is the second software-based trigger step used in ATLAS to trigger on interesting events.

HNL refers to a heavy neutral lepton and is the subject of the search described in this thesis. HNLs are heavy right-handed states that carry no Standard Model (SM) charges. Introducing HNLs into the SM introduces a mechanism that can explain the light neutrino masses.

Hybrid refers to an electrical circuit board that houses the ATLAS Inner Tracker Strip front-end chips and is glued onto the silicon sensors.

- IBL** refers to the Insertable B-Layer that is part of the ATLAS ID. It is the first layer of the pixel detector and was inserted into the ATLAS ID in 2014 to increase the robustness and precision of the tracking, as well as to mitigate luminosity effects and radiation damage.
- ID** refers to the Inner Detector. It is the innermost detector system of the ATLAS detector and is primarily responsible for reconstructing the trajectory of particles.
- IH** refers to the inverted neutrino mass hierarchy wherein the two heavier mass eigenstates have a small mass difference and the third state has the lightest mass.
- Inside-out** refers to a tracking sequence that initially forms track seeds in the inner silicon detector and then extends the tracks outward in radius.
- ITk** refers to the ATLAS Inner Tracker. It is an all-silicon tracking system that will replace the current ID in time for the start of the HL-LHC. It consists of layers of pixel detectors surrounded by a large-area strip detector system.
- ITk-Strip** refers to the silicon-strip tracking detector and is a sub-system of the ATLAS Inner Tracker. These detectors are the subject of the performance and readout research described in chapter 6.
- Kalman filter** refers to a track reconstruction method to find additional space-points that are compatible with the initial track seed.
- L-1** refers to the level-1 hardware-based trigger that uses information collected from the ATLAS muon and calorimeter detector systems to make decisions about which proton collisions might contain interesting physics events.
- L0a** refers to a beam crossing synchronous pulse that transfers data from the pipeline in the Inner Tracker front-end chips to an event buffer. This data is then sent upstream for further data processing. It stands for Level-0 accept.
- LAr** is liquid argon, which is a cryogenic liquid used to detect particles in various ATLAS detectors, including the tile and forward calorimeters.
- LCB** refers to the Level-0 accept (L0a), Control and Bunch counter reset (BCR) signal. It is a custom data protocol used in the Inner Tracker Strip detectors to transmit L0as, control commands and BCRs.
- LEP** is the Large Electron-Positron collider built at CERN to study electroweak interactions. LEP was closed down in November 2000 for the construction of the LHC in the same tunnels.

LHC is the Large Hadron Collider, which is the world's largest and highest energy particle accelerator and collider. It comprises two superconducting rings and is designed to collide protons at four interaction points.

LHCb is a detector dedicated to studying b-physics and is one of the four main LHC experiments. It stands for Large Hadron Collider Beauty Experiment.

LHLoose refers to an electron identification criterion used to select leptons with a certain efficiency and purity. This identification has a larger efficiency and lower purity than the LHMedium criterion. An identification criterion is also known as a working point.

LHMedium refers to an electron identification criterion used to select leptons with a certain efficiency and purity. An identification criterion is also known as a working point.

LINAC2 is the Linear Accelerator. It is one of the accelerators in the LHC proton injection chain.

LLP refers to a long-lived particle. In this context, long-lived usually means that the particle travels a distance larger than the position resolution of the detector before it decays. This term typically refers to an exotic particle and not an existing SM particle with a long lifetime, such as a b-hadron or muon.

LNC refers to a process or decay that conserves lepton number. It stands for lepton number conserving.

LNV refers to a process or decay that violates lepton number. It stands for lepton number violating.

Loose refers to a muon identification criterion used to select muons with a certain efficiency and purity. This identification has a larger efficiency and lower purity than the medium identification criterion. An identification criterion is also known as a working point.

LRT is the large radius tracking reconstruction algorithm used in ATLAS to reconstruct displaced tracks that do not point back to the interaction point.

Luminosity In particles physics, instantaneous luminosity refers to the number of particles passing per unit area per unit time and thus is a measure of how densely packed the two beams are. The integrated luminosity is the time integral of the instantaneous luminosity.

MC refers to a Monte-Carlo simulation and is computational mathematical technique that is used to model the probability of different outcomes in the particle collision events that are studied in the search described in this thesis.

MDT refers to a monitored drift tube in the ATLAS MS that are primarily responsible for tracking and measuring the momentum of muons.

Medium refers to a muon identification criterion used to select muons with a certain efficiency and purity. An identification criterion is also known as a working point.

MS refers to the ATLAS Muon Spectrometer is the outermost layer of the ATLAS detector. It is primarily responsible for detecting and measuring the momenta of muons.

NC is a neutral current interaction mediated by a Z boson.

NH refers the normal neutrino mass hierarchy wherein the two lightest mass eigenstates have a small mass difference and the third state has a larger mass.

Ntuple refers to a data structure that contains an ordered set of n -elements. Ntuples are used to store physics event information in a format that is simple and quick to process the data events.

OS refers to opposite-sign and is used to describe a vertex that has two tracks with opposite charges.

Outside-in refers to a tracking sequence that starts with track segments in the outer TRT layers and extends the tracks inward toward to interaction point.

Pair production is the formation of a particle and anti-particle pair from a neutral boson. It often refers to the creation of an electron and positron from a photon near a nucleus.

Pileup refers to all the additional collisions that occur at the same time as the hard-scatter primary interaction in a proton-proton collision.

Pixel detector is the silicon pixel detector used to reconstruct tracks in the ATLAS ID.

PS is the Proton Synchrotron. It is one of the accelerators in the LHC proton injection chain.

Pseudorapidity is a spatial coordinate that describes the angle of a particle relative to the z -axis. It is defined as $\eta = -\ln \tan(\theta/2)$. In the massless limit where $m \rightarrow 0$, the difference in pseudorapidity, $\Delta\eta$, is Lorentz invariant under boosts along the z -axis.

QA refers to Quality Assurance tests that study the performance and reliability of detector components during the research, development and production stages.

Rapidity is a kinematic variable defined as $y = \frac{1}{2} \ln(\frac{E+p_z}{E-p_z})$, which is Lorentz invariant for boosts along the z -axis, but is not fixed in physical space.

- RH** refers to the relative humidity. It is defined as the ratio of moisture present in the system relative to moisture that would be present if the air was completely saturated.
- ROD** refers to a custom hardware-based readout driver responsible for processing the data from front-end electronics and passing the data to the central readout system.
- RoI** refers to a region of interest in the ATLAS detector identified by the L-1 trigger for further study in the second high-level trigger step.
- ROS** refers to custom electronics boards that are central to the ATLAS readout system. They receive, buffer and process data events from segments of a given sub-detector and have access to the first-level trigger decisions.
- RPC** refers to a resistive plate chamber in the ATLAS MS that is primarily responsible for triggering on muons.
- SCT** is the Semiconductor Tracker, which is a large-area silicon micro-strip detector used to reconstruct tracks in the ID.
- SM** is the Standard Model, which is a mathematical theory that is currently the best-known description of fundamental particles and their interactions (except for gravity).
- SPS** is the Super Proton Synchrotron. It is one of the accelerators in the LHC proton injection chain.
- SR** refers to the signal region for the search described in this thesis. Events in this region are required to pass the all the selections described in section 5.7.
- SS** refers to same-sign and is used to describe a vertex that has two tracks with the same charge.
- TDAQ** is the trigger and data acquisition system in ATLAS. It is responsible for the triggering and reading out data from various detector systems.
- TGC** refers to a thin-gap chamber in the ATLAS MS that is primarily responsible for triggering on muons.
- Tower** is a region in $\eta \times \phi$ space in the calorimeter. Towers are used to build energy clusters.
- TRIUMF** is Canada's particle accelerator centre located in Vancouver, British Columbia.
- TRT** is the Transition Radiation Tracker, which is the outermost sub-system of the ID. It is primarily responsible for extending track reconstruction to larger radii. It is also used to distinguish electrons and pions based on their differing amounts of transition radiation.

VeryVeryLoose refers to a custom electron identification criterion used to select displaced electrons with a certain efficiency and purity. An identification criterion is also known as a working point.

VR refers to the validation region with negligible signal contamination used to study background events and selected displaced vertices for the background estimate in the HNL search.

VSI refers to a vertex reconstruction algorithm used in ATLAS to reconstruct a place of common origin from tracks that do not point back to the interaction point. It stands for Vertex Secondary Inclusive. This term also refers to the default secondary vertexing configuration used in the ATLAS reconstruction software.

VSI Leptons refers to the default lepton-only secondary vertexing configuration used to reconstructed displaced vertices with tracks that are matched to leptons.

Working point is also known as an identification criterion. In this context it refers to a lepton identification criterion used to select leptons with a certain efficiency and purity.

Acknowledgements

To Alison, thank you for your continuous guidance and support through my PhD. Thank you for constantly pushing me to be a better physicist and encouraging me to “do a little bit” of analysis work, even though I wanted to spend more time in the lab. I will forever be grateful for your mentorship, especially over the past few months while I was figuring out what to do after my PhD. I would also like to thank Matthias for his guidance and support while I was working on the displaced HNL search. I am fortunate to have had many insightful research discussions with you and your mentorship over the past few years has been invaluable. To Colin, thank you for getting me started on the hardware readout project. I always valued our chats and I am grateful for your support in the development of the ITk-Strip readout program at UBC. I am also grateful to Wojtek who got me started with firmware programming.

Thank you to all of the displaced HNL team, especially Robin N. and Christian, who worked tirelessly to complete this physics analysis. Thank you also to my PhD committee members, Gary and Eric, for their feedback during my committee meetings. I am also grateful to all the members of the ATLAS Collaboration that I have had the privilege to work with. I also wish acknowledge the financial support from the Natural Sciences and Engineering Research Council of Canada, which gave me the freedom to pursue a variety of research projects.

My PhD would not have been the same without the members of the UBC ATLAS group. To Robin N., Robin H., Sébastien, Zhengcheng, Félix, Chris, Elham, Vincent, Ryan, Russell and Colm, I will always cherish the memories of our lunch-time chats, camping trips and Glühwein nights in the office. To Robin H., thank you for keeping me sane during the long pandemic months. You are one of my biggest cheerleaders and I am grateful for your endless supply of optimism.

To my parents, you taught me to be curious and to think critically. Your support and advice has given me the confidence to pursue my passions. I am also incredibly lucky to have the love and support of my sister Hope, grandparents, Carol, Ernie, Rosemary and Peter, as well as my second family, Larissa, John, Augustine and Samuel.

Most of all, thank you to Nicholas. You are my hero, not only for editing every page of this thesis with no physics background (while also studying to become a lawyer!), but also for making me believe “this is fine”, even when it felt like there were fires burning everywhere. I can’t wait to see what our next adventure will bring!

Chapter 1

The Standard Model

The Standard Model (SM) is currently the best-known description of fundamental particles and their interactions, with the exception of gravity. While this theory provides precise predictions that are consistent with many experimental measurements, there is also evidence that the SM is incomplete. This chapter provides a basic description of the SM and discusses some its successes and limitations, which provide insight into where to search for new physics.

1.1 Particles in the Standard Model

The matter around us is made up of two groups of fundamental particles, namely leptons (l) and quarks (q). Both groups follow Fermi statistics with half-integer spins and, thus, are also referred to as fermions. The charged leptons, which consist of electrons (e), muons (μ) and taus (τ) are organized into three generations. Each charged lepton also has a corresponding charge-neutral neutrino (ν). The other matter particles, known as quarks, come in six flavours and are also paired in three generations: up (u)/down (d), charm (c)/strange (s) and top (t)/bottom (b). Due to colour confinement, quarks are bound in composite colourless systems known as hadrons. Hadrons can be divided into two sub-groups: mesons, which are made up of quark and anti-quark pairs, and baryons, which consist of three quarks, such as protons and neutrons.

As far as we know, our universe has four fundamental forces: strong, electromagnetic, weak and gravitational. Each force is mediated by a particle (or particles) known as a gauge boson. The four fundamental forces, their mediators and their relative strengths¹ are listed in table 1.1.

There are multiple theories that describe each force. The classical theory of the gravitational force is described by Newton’s law of universal gravitation. The relativistic theory of gravity is described by Einstein’s general theory of relativity. However, gravity’s effect on fundamental particles in the SM is so weak that it is almost always negligible in particle physics experiments. The strong force is responsible for confining quarks and binding protons and neutrons together to form nuclei. The theory of electromagnetic force is known as electrodynamics and governs how light and charged matter particles interact. The weak force is

¹The “strength” of a force depends on its nature and the distance over which it is acting. Table 1.1 provides an approximate order of magnitude for each force to illustrate their relative strengths.

1.1. Particles in the Standard Model

Table 1.1: The four fundamental forces with their relative strengths and mediators. Table from [1].

Force	Strength	Mediator(s)
Strong	10	Gluons
Electromagnetic	10^{-2}	Photon
Weak	10^{-13}	W and Z
Gravitational	10^{-41}	Unknown

responsible for decays of unstable subatomic particles, such as mesons, and plays an important role in nuclear radiation. The electromagnetic and weak forces have been unified into a single theory known as *electroweak* theory. In the SM, all fermions interact via the weak force and charged leptons also interact via the electromagnetic force. Quarks carry both electromagnetic charge and colour², which means quarks interact with the electromagnetic, strong and weak forces. There is one final particle in the SM known as the Higgs boson. This boson was the last SM particle to be discovered and it is a scalar boson that is responsible for giving particles mass. A summary of the elementary particles in the SM is found in fig. 1.1.




Figure 1.1: The elementary particles that make up the Standard Model. Figure adapted from [2].

²Here, colour refers to colour charge, which is a quantum number related to the strong interaction.

1.2 Chiral Fermions

In general, the Lagrangian for a massless fermion is:

$$\mathcal{L} = \bar{\psi} i \gamma^\mu \partial_\mu \psi \quad (1.1)$$

where ψ represents a particle field, γ^μ are the gamma matrices and ∂_μ are partial time and space derivatives. In quantum theory, this Lagrangian can either describe a massless particle ψ_L with left-handed helicity, in which its spin is anti-parallel to its direction of motion, or a massless anti-particle ψ_R with right-handed helicity, in which its spin is parallel to its direction of motion. In general, ψ_L or ψ_R represent ‘chiral’ fermion states, which are not equal to helicity states. However, in the case of a massless fermion (or when $E \gg m$), helicity and chirality are equivalent. Thus, chiral states are often labelled by their handedness.

To describe massive fermions, a Dirac mass term could be written down. This term would mix together left and right-handed fermions and combine them into a single Dirac fermion:

$$\mathcal{L}_{\text{Dirac mass}} = -m_D(\bar{\psi}_R \psi_L + \bar{\psi}_L \psi_R) \quad (1.2)$$

In the SM, all fermions are Dirac fermions, except for the neutrinos whose nature remains unknown. If neutrinos are not Dirac, they could be Majorana fermions with a mass term such as:

$$\mathcal{L}_{\text{Majorana mass}} = -m_M(\bar{\psi}_L \psi_L + \bar{\psi}_R \psi_R) \quad (1.3)$$

The origin of neutrino masses will be discussed in chapter 2, as it motivates the search for the theoretical particle described in this dissertation.

1.3 Structure of the Standard Model

The structure of the SM is based on symmetries of local transformations known as gauge transformations. It is described by the gauge group $SU(3)_C \times SU(2)_L \times U(1)_Y$. The strong interactions are described by the $SU(3)_C$ gauge group, where the subscript C denotes the colour. The electroweak interactions are described by the $SU(2)_L \times U(1)_Y$ gauge group, which represents the unification of weak and electromagnetic forces. The subscript L denotes that only left-handed particles participate in weak interactions and the subscript Y denotes the weak hypercharge. This weak hypercharge $Y = Q - T^3$ relates the electric charge Q of the particle and third component of the weak isospin T^3 , all of which are quantities that must be conserved in an electroweak interaction.

In the SM, the strong, weak and electromagnetic gauge fields are represented by G_μ^i , W_μ^j , and B_μ . The quarks are represented by left-handed quark doublets $Q_L = (u_L, d_L)^T$ and right-

1.3. Structure of the Standard Model

handed u_R and d_R singlet states. The leptons are represented by left-handed lepton doublets $L_L = (\nu_L, e_L)^T$ and right-handed e_R singlet states. Leptons do not have corresponding ν_R singlet states because neutrinos are massless in the SM. The values of the electric charge, the third component of the weak isospin and the hypercharge for the fields are summarized in table 1.2.

Table 1.2: The SM fermion fields and their electric charges Q , third components of the weak isospin T^3 and hypercharges Y .

Fermion Type	Field	Q	T^3	Y
Quarks	$Q_L = \begin{pmatrix} u_L \\ d_L \end{pmatrix}$	+2/3	+1/2	1/6
		-1/3	-1/2	
	u_R	+2/3	0	+2/3
	d_R	-1/3	0	-1/3
Leptons	$L_L = \begin{pmatrix} \nu_L \\ e_L \end{pmatrix}$	0	+1/2	-1/2
		-1	-1/2	
	e_R	-1	0	-1

The left and right fermion fields carry different electroweak charges because of the chiral nature of the weak interaction. Therefore, a Dirac mass term of the form in eq. (1.2) cannot be directly added to the Lagrangian. Instead, masses for the quarks and charged leptons are generated via interactions between the fermion and Higgs fields.

The Lagrangian of the SM takes the form:

$$\mathcal{L}_{\text{SM}} = \mathcal{L}_{\text{gauge}} + \mathcal{L}_{\text{Higgs}} + \mathcal{L}_{\text{Yukawa}} \quad (1.4)$$

The gauge Lagrangian $\mathcal{L}_{\text{gauge}}$ consists of kinetic terms for the fermions and gauge bosons:

$$\begin{aligned} \mathcal{L}_{\text{gauge}} = & -\frac{1}{4}(G_{\mu\nu}^i)^2 - \frac{1}{4}(W_{\mu\nu}^j)^2 + -\frac{1}{4}(B_{\mu\nu})^2 \\ & + i\bar{Q}_L \not{D} Q_L + i\bar{u}_L \not{D} u_R + i\bar{d}_L \not{D} d_R \\ & + i\bar{L}_L \not{D} L_L + i\bar{e}_L \not{D} e_R \end{aligned} \quad (1.5)$$

where these terms are products of the gauge field strength tensors and the fermion fields. The terms are written in a compact notation using a covariant derivative of the form:

$$\not{D} = \gamma^\mu (\partial_\mu + ig_s t_C^i G_\mu^i + ig_w t_L^j W_\mu^j + ig' Y B_\mu) \quad (1.6)$$

where g_s is the coupling strength of the strong interaction, g_w is the coupling strength of the weak interaction and g' is the coupling strength for the electromagnetic interaction. The matrices t_L^j form a representation of the $SU(2)_L$ group and can be expressed in terms of the

Pauli matrices, σ^j , such that $t_L^j = \sigma^j/2$. The Gell-Mann matrices t_C^i form a representation of the $SU(3)_C$ group. The Higgs Lagrangian $\mathcal{L}_{\text{Higgs}}$ describes the kinetic and potential terms for the Higgs doublet Φ :

$$\mathcal{L}_{\text{Higgs}} = |D\Phi|^2 - (-\mu^2|\Phi|^2 + \lambda|\Phi|^4) \quad (1.7)$$

where μ and λ are parameters of the Higgs potential. This potential is discussed in more detail in section 1.4. The Yukawa Lagrangian, $\mathcal{L}_{\text{Yukawa}}$, describes the interactions between the scalar bosons and fermions:

$$\mathcal{L}_{\text{Yukawa}} = -F_u \bar{Q}_L \tilde{\Phi} u_R - F_d \bar{Q}_L \Phi d_R - F_e \bar{L}_L \Phi e_R \quad (1.8)$$

where $\tilde{\Phi} = i\sigma_2 \Phi$ and F_u, F_d, F_e are the Yukawa couplings between the Higgs and fermion fields.

1.4 Electroweak Symmetry Breaking

As mentioned in section 1.3, the mass terms for the charged fermions and massive W and Z bosons cannot be added directly to the SM Lagrangian. Instead, these masses are generated via the Higgs mechanism, with the spontaneous breaking of the local $SU(2)_L \times U(1)_Y$ gauge invariance.

In a theory without symmetry breaking, the Higgs Lagrangian, $\mathcal{L}_{\text{Higgs}}$, has a potential of the form $V(\Phi) = \mu^2|\Phi|^2 + \lambda|\Phi|^4$ with $\mu^2 > 0$, where Φ is a doublet in the SM. This is sketched in fig. 1.2 (a). This potential has a unique minimum, or ground state, at $\Phi = 0$ that is symmetric: rotations about this minimum leave the ground state unchanged. For spontaneous symmetry breaking to occur $\mu^2 < 0$. This potential is sketched in fig. 1.2 (b).




Figure 1.2: The potential $V(\Phi) = \mu^2|\Phi|^2 + \lambda|\Phi|^4$ in a theory without (a) and with (b) local symmetry breaking around a minimum of the potential.

It is clear from the figure that there is no longer a unique ground state for the system. Instead, there is a ring of ground states that lie at a radius equal to $\sqrt{\frac{-\mu^2}{2\lambda}} \equiv \frac{v}{\sqrt{2}}$. This value v is known as the “vacuum expectation value” or VEV of the Higgs field. In this new potential, the local field close to any given choice of ground state is not invariant under rotation. This type of symmetry is thus referred to as “spontaneously broken”.

In the Higgs mechanism, it is convenient to represent the Higgs doublet as:

$$\Phi = \frac{1}{\sqrt{2}} \begin{pmatrix} 0 \\ v + h \end{pmatrix} \quad (1.9)$$

This choice minimizes the number of scalar degrees of freedom and expresses Φ in terms of a single scalar Higgs field h . Substituting eq. (1.9) back into the Higgs Lagrangian in eq. (1.7) results in mass terms for the scalar Higgs field h and the massive gauge bosons W and Z with:

$$\begin{aligned} m_h^2 &= 2\lambda v^2, \\ m_W^2 &= \frac{1}{4}g_w^2 v^2 \text{ and} \\ m_Z^2 &= \frac{1}{4}(g_w^2 + g'^2)v^2. \end{aligned} \quad (1.10)$$

1.5 Successes and Limitations of the Standard Model

The SM is an elegant theory that has been experimentally tested at collider experiments to a high degree of precision across a wide range of energy regimes. In fig. 1.3, cross section measurements for various SM processes are shown. These measurements are performed using proton-proton collision data collected by the ATLAS experiment at the Large Hadron Collider. So far the SM predictions have been incredibly successful and show overall excellent agreement between data and theory.

However, there is experimental evidence that suggests this theory is incomplete. For instance, astrophysical observations of gravitational lensing [3], rotational velocities of galaxies [4, 5] and measurements of the cosmic microwave background [6] all suggest that the visible matter described in the SM represents only a small fraction of the total matter present in our universe. The SM does not describe the nature of this “dark” matter. Another example is the ratio of matter to anti-matter present in the universe. Our universe is largely matter-dominated and currently, the SM offers no explanation for this. Most relevant to the search described in this dissertation is the observation of neutrino flavour oscillations [7]. In the SM, neutrinos are strictly massless, due to the absence of right-handed neutrino fields. However, experiments have observed the probability of measuring a neutrino with a definite flavour can vary as the neutrino propagates through space. This flavour oscillation probability is proportional to the mass splitting between the neutrino mass states. Thus, the observation of

nonzero neutrino masses also suggests there exists physics beyond the SM.




Figure 1.3: Summary of several Standard Model cross-section measurements. The measurements are corrected for branching fractions, compared to the corresponding theoretical expectations. Figure taken from [8].

Chapter 2

Heavy Neutral Leptons

Many theories have been proposed for physics beyond the Standard Model (BSM) in order to provide a complete theory of the universe. In particular, theories with several right-handed neutrino states have been proposed to explain (some of the) observed BSM phenomena. Heavy neutral leptons (HNLs) are massive particles that interact similarly to neutrinos with interaction strengths that are described by flavour dependent mixing angles (U_e^2 , U_μ^2 and U_τ^2). HNLs were first proposed in left-right symmetric theories as heavy right-handed states that interact via the weak force [9–12]. Introducing right-handed heavy states also provides a type-I seesaw mechanism that can explain the light SM neutrino masses [13–15]. Not only does this theory provide a natural explanation to the nature of the relatively small neutrino masses (when compared with charged leptons), it also provides an explanation for the asymmetry between matter and anti-matter in the universe through a process known as leptogenesis. In HNL scenarios with masses around the electroweak scale, the seesaw mechanism implies that the neutrino Yukawa couplings are relatively small such that neutrino production in the early universe is very slow. If charge-parity violating oscillations occurred during neutrino production in the early universe, then this could explain the matter-antimatter asymmetry [16]. Moreover, models with three HNLs can also incorporate a dark matter candidate [17, 18]. This chapter provides a brief overview of the seesaw mechanism that can explain why neutrinos have extremely light masses and reviews how HNLs interact with the SM.

2.1 Neutrino Masses From Right-Handed Neutrinos

This section describes how right-handed HNLs can give rise to light SM neutrino masses. This description of the type-I seesaw mechanism is adopted from [19]. The Lagrangian for $\mathcal{N}$ right-handed neutrinos N_{R_I} is described by:

$$\mathcal{L}_{\text{SM+HNL}} = \mathcal{L}_{\text{SM}} + i\bar{N}_{R_I}\gamma^\mu\partial_\mu N_{R_I} - (F_{\alpha I}\bar{L}_\alpha N_{R_I}\tilde{\Phi} + \frac{1}{2}M_I\bar{N}_{R_I}^c N_{R_I}) \quad (2.1)$$

where $\mathcal{L}_{\text{SM}}$ is the usual SM Lagrangian, $F_{\alpha I}$ is the Yukawa couplings matrix, M_I is the mass matrix for the right-handed states N_{R_I} , $L_\alpha = (\nu_\alpha, e_\alpha)^T$ are left-handed lepton doublets with $\alpha = \{e, \mu, \tau\}$, Φ is the Higgs doublet with $\tilde{\Phi} = i\sigma_2\Phi$ and $\bar{N}_{R_I}^c$ is the charge conjugate transpose of the N_I state. The right handed neutrino states carry no standard model charges

(i.e. they are “sterile”) and therefore it is possible to add a Majorana mass term of the form $M_I \bar{N}_{R_I}^c N_{R_I}$ to the Lagrangian, which is usually prohibited by gauge invariance. A summary of the elementary particles in the SM with the addition of $\mathcal{N} = 3$ right-handed sterile neutrinos is found in fig. 2.1.




Figure 2.1: The elementary particles that make up the Standard Model with the addition of $\mathcal{N} = 3$ right-handed sterile neutrinos. Figure adapted from [2].

After electroweak symmetric breaking, the Higgs field Φ obtains a vacuum expectation value v and $\tilde{\Phi} = \frac{1}{\sqrt{2}} \begin{pmatrix} v + h \\ 0 \end{pmatrix}$. Thus, eq. (2.1) can be rewritten as:

$$\mathcal{L}_{\text{HNL}} = \mathcal{L}_{\text{SM}} + i \bar{N}_{R_I} \gamma^\mu \partial_\mu N_{R_I} - ((v + h)/\sqrt{2}) F_{\alpha I} \bar{\nu}_{L_\alpha} N_{R_I} - \frac{1}{2} M_I \bar{N}_{R_I}^c N_{R_I} \quad (2.2)$$

Thus, the effect of the Yukawa interaction term in eq. (2.1) is to generate a Dirac mass term that couples left and right-handed neutrino states and the HNLs “inherit” weak-like interactions with the W and Z bosons that are suppressed by the mixing angles, $\Theta_{\alpha I}$.

The mass terms in eq. (2.2) can always be diagonalized, but in the limit where $|M_I| \gg v|F_{\alpha I}|$, this diagonalization leads to $\mathcal{N}$ heavy and three light states. The eigenvalues of the diagonalized matrix corresponds to the HNL and SM neutrino masses. These light neutrino

masses are suppressed by the HNL-neutrino mixing angles, $\Theta_{\alpha I}$, such that:

$$\Theta_{\alpha I} = \frac{v|F_{\alpha I}|}{M_I} \quad (2.3)$$

This is known as a seesaw mechanism because increasing the eigenvalues increases the masses of the HNLs, while pushing the masses of the SM neutrinos down.

Experimental observations of neutrino flavour oscillations tell us that there are two mass differences for the SM neutrinos [7]. In order for the above description to explain the observed neutrino mass splittings, the number of N_I states must be $\mathcal{N} \geq 2$. The search described in this dissertation uses a model-independent approach where only one HNL is kinematically accessible and, if other HNLs are present in the theory, they are considered sufficiently heavy that they do not affect the analysis. An additional model-dependent interpretation, which considers two HNLs with *nearly degenerate* masses ($M_1 \sim M_2$), is also included. In this case, the HNL mass, $m_N = (M_1 + M_2)/2$. Further motivation for this model-dependent approach is discussed in [20].

2.2 Interactions with the Standard Model

HNLs interact with the SM through mixing with neutrinos. The experimentally relevant observable for this interaction is the mixing of the HNLs with a given neutrino flavour:

$$|U_\alpha|^2 \equiv \sum_I |\Theta_{\alpha I}|^2 \quad (2.4)$$

where $\alpha = e, \mu, \tau$ and I is the number of HNLs. This mixing is also referred to as the “coupling” to a definite flavour neutrino, α , since it measures the strength of the interaction between the two particles. The total mixing between HNLs and neutrinos is:

$$|U|^2 \equiv \sum_{\alpha I} |\Theta_{\alpha I}|^2 \equiv \sum_\alpha |U_\alpha|^2 \quad (2.5)$$

The mixing angles, $\Theta_{\alpha I}$, are related to the neutrino Yukawa couplings, $F_{\alpha I}$. There are measured constraints on the Yukawa couplings that can be expressed as constraints on the values of $|U_\alpha|^2$ that are compatible with neutrino oscillations data. An analysis done in [20] used NuFIT, the most recent global fit to neutrino oscillation data [21], to map out values of the relative strength of the mixing with electron, muon and tau neutrinos, $x_\alpha = |U_\alpha|^2/|U|^2$, that are consistent with neutrino oscillations measurements (see fig. 2.2). Two regions are defined that are consistent with the data assuming inverted and normal neutrino mass hierarchy. Normal (NH) and inverted hierarchy (IH) benchmark models are defined in the two allowed regions according to the ratio between the flavour specific mixing angles (see table 2.1). Ex-

perimental searches typically report the interpretation of their results under the assumption of mixing with a single neutrino flavour. For this reason, electron-only and muon-only mixing scenarios are also defined.³




Figure 2.2: A plot that shows the allowed combinations of $|U_\alpha|^2/|U|^2$ that are consistent with the recent NuFIT [21] fit to neutrino oscillation data. The allowed regions under normal (blue) and inverted (green) hierarchy are shown. The red cross markers represent two benchmark scenarios that are defined for the two different hierarchy assumptions. Figure adapted from [20].

Table 2.1: The relative strength of the mixing, $x_\alpha = |U_\alpha|^2/|U|^2$, with the three neutrino flavours for the inverted (IH) and normal (NH) hierarchy benchmark and single-flavour mixing scenarios.

Model	x_e	x_μ	x_τ
IH	1/3	1/3	1/3
NH	0.06	0.48	0.46
Muon-only	0	1	0
Electron-only	1	0	0

³In the search presented in this dissertation, the acceptance to decays with tau leptons is zero. For this reason, a single tau-only mixing scenario is not considered.

2.3 HNL Decays

The analysis described in this dissertation searches for HNLs produced through a W boson that decays into a lepton and a neutrino, which mixes with the HNL. HNL decays are mediated by the weak interaction via charged current or a neutral current decays. The decay, shown in fig. 2.3, can either be fully leptonic if the bosons decay to lepton final states or semi-leptonic if the bosons decay hadronically. If the charged leptons from the decay of the HNL have the same flavour (i.e. $\beta = \gamma$), then both the charged current and neutral current decays contribute to the total process.




Figure 2.3: Feynman diagrams for HNL production from the decay of a W boson and the HNL decay into leptonic and semi-leptonic final states.

2.4 Lepton Number Violation

The SM Lagrangian has a global $U(1)$ symmetry that corresponds to a conserved quantum number known as the lepton number L , where leptons have $+L$ and anti-leptons have $-L$. If HNLs behave as Majorana fermions, with identical particle and anti-particle states, then lepton number violating (LNV) decays are possible. In some models, LNV decays can be suppressed compared to the lepton number conserving (LNC) processes, while in others, such as a Dirac fermion HNL model, LNV decays are forbidden [20, 22–25]. Diagrams of LNC and LNV decays are shown in fig. 2.4. The distinction between an LNC and LNV decay comes from the relative charge between the first lepton (l_α) and the lepton (l_β or ν_β) that couples to the HNL. For LNC decays, these two leptons must have opposite charges, while for LNV decays, these leptons have the same charge.

Due to the chiral nature of the weak interaction, the HNL is produced with a preferred helicity state (spin polarization) that is correlated with the spin of the lepton l_α . Since the HNL is a very weakly interacting particle (and in particular is electrically neutral), its spin

is preserved during its propagation through space. Thus, when the HNL decays, the spin of the lepton that couples directly to the HNL l_β will be correlated to the spin of the HNL and correspondingly to the spin of l_α . Depending on how the HNL decays (either via an LNC or LNV known decay), the angular distributions of the leptons will be different. Such kinematic differences can have small impacts on the experimental efficiencies of selecting HNL signal events. The impact of spin correlations in LNC and LNV decays on the search described in this dissertation is described in more detail in section 2.4.




Figure 2.4: Lepton number conserving (LNC) and lepton number violating (LNV) diagrams for the charge-current leptonic HNL decay.

2.5 HNL Cross Section

The probability of producing a heavy neutral lepton can be described by the production cross section, σ_N , measured in units of area. To compute the HNL cross section, the mass, m_N , and mixing angles, $|\Theta_\alpha|^2$, need to be specified. The branching ratio for the W boson decaying to an HNL and a charged lepton, l_α , is directly proportional to the coupling strength and can be expressed as [26]:

$$BR(W \rightarrow l_\alpha N) = BR(W \rightarrow l_\alpha \nu) \cdot |\Theta_\alpha|^2 \left(1 - \frac{m_N^2}{m_W^2}\right)^2 \left(1 + \frac{m_N^2}{2m_W^2}\right) \quad (2.6)$$

where $BR(W \rightarrow l_\alpha \nu)$ is the known W branching ratio to a lepton-neutrino pair, Θ_α are the mixing angles between the HNL and SM neutrinos and m_N and m_W are the mass of the HNL and W boson respectively.

This means that the total HNL cross section from the decay of W boson, σ_N , is given by:

$$\begin{aligned}\sigma_N &= \sigma(pp \rightarrow W) \cdot BR(W \rightarrow l_\alpha N) \\ &= \sigma(pp \rightarrow W) \cdot BR(W \rightarrow l_\alpha \nu) \cdot x_\alpha \cdot |\Theta|^2 \cdot \left(1 - \frac{m_N^2}{m_W^2}\right)^2 \left(1 + \frac{m_N^2}{2m_W^2}\right)\end{aligned}\quad (2.7)$$

where $x_\alpha \cdot |\Theta|^2$ has been substituted for $|\Theta_\alpha|^2$.

To compute the rate for the HNL decay to leptons, σ_N is multiplied by the branching ratio for the fully-leptonic HNL decay $BR(N \rightarrow ll\nu)$. This branching ratio takes into account the strength of the mixing angles between the HNL and its decay products. The relevant HNL branching ratios used in the analysis in this dissertation are shown in fig. 2.5 for the fully-leptonic decay channels.




Figure 2.5: Branching ratio of $N \rightarrow ll\nu$ with a sum of all neutrino flavours as a function of m_N for muon-only (a), electron-only (b), NH (c) and IH (d) mixing scenarios. In each mixing scenario, the branching ratio is shown for $N \rightarrow ee\nu$ (blue), $N \rightarrow e\mu\nu$ (orange), $N \rightarrow \mu e\nu$ (pink) and $N \rightarrow \mu\mu\nu$ (purple) decays.

Individual HNL decay branching ratios, $BR(N \rightarrow ll\nu)$ depend on m_N and the relative strengths of the mixing angles: x_e , x_μ and x_τ . The branching ratios are computed assuming one HNL state in the “Majorana-limit” with both LNC and LNV decays contributing equally. These branching ratio values include a sum of all possible flavours of neutrinos. The difference between $BR(N \rightarrow l_\gamma l_\gamma \nu)$ and $BR(N \rightarrow l_\beta l_\gamma \nu)$ (with $\beta, \gamma = e$ or μ) is due to the interference between the charged and neutral current decays in cases where the leptons, l , are identical. In the muon-only (electron-only) mixing scenario, $N \rightarrow ee$ ($N \rightarrow \mu\mu$) decays have a small branching ratio from exclusively neutral current decays and $N \rightarrow e\mu$ ($N \rightarrow \mu e$) decays are not possible. In the multi-flavour mixing scenarios, all decays are possible. In the NH benchmark the mixing with μ/ν_μ is stronger than the mixing with e/ν_e , resulting in larger branching ratios for $N \rightarrow \mu\mu$ and $N \rightarrow \mu e$ decays. The IH benchmark has equal mixing with all three lepton flavours and thus, the difference between the four decays comes from the interference between neutral and charge-current decays.




Figure 2.6: The total mixing $|U|^2$ as a function of HNL mass m_N for the muon-only mixing scenario. Each $|U|^2$ is computed for m_N between 2 and 20 GeV with proper decay lengths $c\tau_N = 1, 10$ and 100 mm.

The proper lifetime of the HNL, τ_N , strongly depends on the mass, m_N , and the mixing angle, $|\Theta|^2$, such that [27]:

$$\tau_N \propto \frac{1}{m_N^5 |\Theta|^2} \quad (2.8)$$

The dependence on the mass in this expression can be derived by approximating the HNL decay as a four-fermion contact interaction where the width, Γ , is proportional to the square of the Fermi constant G_F^2 . The value of G_F^2 has negative four mass dimensions and Γ has one mass dimension. Thus, using dimensional analysis, Γ must be proportional to m^5 . Since τ is

inversely proportional to Γ , the lifetime is proportional to $1/m^5$.

The total mixing angle, $|\Theta|^2$, for one HNL model in the “Majorana-limit” can be analytically computed for a given generated mass and lifetime using a Fermi theory computation described in [28]. These computations include all exclusive decay widths and a sum over all the decay channels that contribute to the process. The mixing angle, $|U|^2 = \sum_{\alpha I} |\Theta_{\alpha I}|^2$, as a function of HNL mass, m_N , for the muon-only mixing scenario is shown in fig. 2.6 for $c\tau_N = 1, 10, 100$ mm. The relationships between $|U|^2$ and m_N in the other mixing scenarios are similar and the corresponding $|U|^2$ have similar orders of magnitude to the muon-only mixing scenario.

Chapter 3

The ATLAS Detector at the Large Hadron Collider

The research described in this dissertation uses proton collision data collected by the ATLAS experiment at the European Centre for Nuclear Research (CERN). The proton collisions are provided to the ATLAS experiment by the Large Hadron Collider (LHC). This chapter gives an overview of the LHC and describes the ATLAS detector.

3.1 The Large Hadron Collider

The LHC is the world's largest particle accelerator and collider. Consisting of two superconducting rings, the LHC is designed to collide protons with a centre-of-mass energy of $\sqrt{s} = 14$ TeV and a nominal luminosity of $L = 10^{34} \text{ cm}^{-2} \text{ s}^{-1}$ [29]. It is installed in a 27 kilometre-long tunnel under the countryside surrounding Geneva, Switzerland. Protons are accelerated in opposite directions in each of the superconducting rings and collide at four interaction points. These points, where the two proton beams cross, are the locations of the four main LHC experiments: ATLAS [30], CMS [31], ALICE [32] and LHCb [33].

Building the LHC at CERN meant that this new machine would be able to reuse the existing tunnels and injection chain that were originally built for the Large Electron Positron (LEP) machine. The injection chain starts with a linear accelerator (LINAC2), where the protons, sourced from hydrogen gas with their electrons stripped off by an electric field, are accelerated to 50 MeV. Then, a series of smaller accelerators gradually increase the energy of the protons. The Proton Synchrotron Booster (Booster), the Proton Synchrotron (PS) and the Super Proton Synchrotron (SPS) accelerate the protons up to 1.4, 25 and 450 GeV, respectively. Thereafter, the protons are injected into the LHC where they can be further accelerated up to a maximum energy of 7 TeV. Figure 3.1 shows an illustration of the path followed by the protons.




Figure 3.1: An illustration of various experimental facilities and beam lines at the CERN accelerator complex relevant to the LHC. The protons are accelerated through the injection chain and into the LHC. The path they follow is: LINAC2 → BOOSTER → Proton Synchrotron (PS) → Super Proton Synchrotron (SPS) → LHC. Figure adapted from [34].

3.1.1 LHC Filling Scheme

The beams in the LHC consist of bunches of protons spaced by 25 ns intervals. This is equal to a bunch crossing rate of 40 MHz. At the LHC’s design luminosity, each bunch of protons contains about 1.15×10^{11} protons. When filling the LHC, not all of these bunch spaces, also known as buckets, are filled with protons. Instead, empty buckets are left in the beam in order to meet various operational constraints, such as allowing time for the injection of protons from the SPS or time to dump the beam if/when this is required. This means that when the LHC is full there are a total of 2808 bunches circulating in each LHC beam. The LHC filling scheme changes with time, but any changes are always done to give ATLAS and CMS the same number of protons collisions. The full filling scheme is described in detail in Ref. [35].

3.1.2 Luminosity

The number of events per second produced at the LHC for a given process is proportional to the cross section of the process of interest (σ_{process}) and the instantaneous luminosity (L)

that is delivered by the LHC over time (t):

$$N_{\text{process}} = \sigma_{\text{process}} \int L dt = \sigma_{\text{process}} \mathcal{L} \quad (3.1)$$

where $\mathcal{L}$ is the time integral of the instantaneous luminosity, known as the integrated luminosity, $\mathcal{L}$. The instantaneous luminosity is given by [36]:

$$L = f \frac{n_1 n_2}{4\pi \sigma_x \sigma_y} F \quad (3.2)$$

where n_1 and n_2 are the number of protons in each bunch, f is the collision frequency and σ_x and σ_y are the root-mean-square of the horizontal and vertical bunch size respectively. Since collisions at the LHC are not exactly head-on, an additional factor F accounts for geometric effects arising from the nonzero crossing angle between the two beams.




Figure 3.2: Total integrated luminosity versus time (a) delivered to ATLAS (green), recorded by ATLAS (yellow), and certified to be good quality data (blue) during Run 2. Distribution of the mean number of interactions per bunch crossing (b) recorded by the ATLAS experiment in the years 2015-2018. Figures taken from [37].

3.1.3 Dataset

Although the LHC was designed to collide protons with a centre-of-mass energy of up to $\sqrt{s} = 14$ TeV, it has not yet been run at this maximum energy. The first LHC run (Run 1) operated at $\sqrt{s} = 7$ TeV (2011) and $\sqrt{s} = 8$ TeV (2012). The second run (Run 2) began in 2015 and finished collecting data in 2018 at a centre-of-mass energy of $\sqrt{s} = 13$ TeV. The search described in this dissertation uses the data collected during Run 2. This dataset contains a total integrated luminosity of 139 fb⁻¹ after accounting for recording efficiency and good detector operation. The evolution of luminosity delivered to the ATLAS experiment is

shown in fig. 3.2 (a). During Run 2, a total of 156 fb^{-1} of stable proton-proton collisions were delivered to ATLAS by the LHC. Out of the total luminosity delivered, 147 fb^{-1} was recorded by ATLAS, due to various data taking inefficiencies. Of this recorded luminosity, 139 fb^{-1} met the requirement that the data was of good quality for all sub-detectors needed for the search presented in this dissertation.

3.1.4 Pileup

When the proton bunches collide at the interaction points around the ring, multiple proton collisions can be recorded. The primary collision of interest, known as the hard-scatter interaction, is the one with the largest momentum transfer. The additional collisions are called pileup and can occur in time or out of time with the hard-scatter event. In-time pileup refers to other collisions that are from the same bunch crossing as the hard-scatter collision. In contrast, out-of-time pileup occurs when readout times for the detector systems are longer than the time between two bunch crossings resulting in the detector sometimes recording energy deposits from earlier or later bunch crossings. Figure 3.2 (b) shows the mean number of interactions per bunch crossing recorded by the ATLAS experiment between 2015 and 2018 at $\sqrt{s} = 13 \text{ TeV}$.

3.1.5 High-Luminosity LHC

The LHC is undergoing a series of upgrades to enable the machine to produce collisions with ten times the current luminosity. This upgraded machine, known as the High-Luminosity LHC (HL-LHC) will provide a challenging new environment for future detector operations. There is a two-and-a-half-year shutdown planned in order to install the necessary upgrades for the current LHC machine and for the various detectors on the LHC ring. One of these detector upgrades for the ATLAS experiment is discussed in detail in chapter 6. The details of the LHC/HL-LHC run schedule are shown in fig. 3.3.




Figure 3.3: The LHC/HL-LHC timeline and run schedule as of Fall 2021. This diagram includes the planned shutdown and upgrade periods. Figure taken from [38].

3.2 The ATLAS Detector

The ATLAS experiment is a general purpose detector designed to search for the Higgs boson over a large mass range, perform precision measurements of the SM and search for evidence of BSM physics. In order to accomplish these objectives, ATLAS is equipped with a variety of very fast and radiation-hard electronics and particle detection elements that provide good object resolution and reconstruction efficiency over a wide range of energies [39]. This section provides an overview of the different detector components of the ATLAS experiment.

3.2.1 Coordinate System

The ATLAS experiment uses a right-handed global coordinate system that defines the locations of the various detector subsystems and describes the kinematics of particles produced in the proton-proton collisions. The origin of the coordinate system is at the nominal proton beam interaction point. The x -axis points from the origin towards the centre of the LHC ring, while the y -axis points up from the interaction point towards the surface. The z -axis is then perpendicular to the $x - y$ plane, pointing along the LHC beam axis. The azimuthal angle ϕ and the polar angle θ are defined in the usual way with ϕ measured around the x -axis and θ measured as the angle between the z -axis and a line segment from the origin to the point of interest.

In hadron colliders, it is rare that the centre-of-mass frame of the interaction is the same as the detector's rest frame, since the individual quarks or gluons that collide are not equally balanced in momentum. For this reason, it is useful to define the pseudorapidity, $\eta = -\ln \tan(\theta/2)$, because, in the massless limit where $m \rightarrow 0$, the difference in pseudorapidity, $\Delta\eta$, is Lorentz invariant under boosts along the z -axis. Alternatively, we can also define the rapidity, $y = \frac{1}{2} \ln(\frac{E+p_z}{E-p_z})$, which is truly Lorentz invariant for boosts along the z -axis, but is not fixed in physical space. This variable is typically used to measure massive objects and is approximately equal to the pseudorapidity η for highly relativistic particles with $m \ll E$. Other commonly used quantities include the transverse momentum, $p_T = \sqrt{p_x^2 + p_y^2}$, transverse energy, $E_T = \sqrt{m^2 + p_T^2}$, and the angular distance between two objects, $\Delta R = \sqrt{\Delta\eta^2 + \Delta\phi^2}$.

3.2.2 Detector Overview

A schematic of the ATLAS detector is shown in fig. 3.4. It is the largest of the four main LHC experiments, measuring 45 m in length, 25 m in height and weighing approximately 7,000 tonnes [39]. ATLAS is a multi-purpose detector that consists of three layers designed to take complementary and redundant measurements in order to fully reconstruct particles emerging from the proton collisions. These layers are symmetrically distributed around the nominal interaction point. In the central region, the sub-detectors are arranged in a cylindrical barrel

structure around the beam pipe, while disk-shaped end-cap structures are located on either side of the barrel.




Figure 3.4: A cut-away view of the the ATLAS detector. Figure taken from [39].

The ATLAS inner detector (ID) is located closest to the beam pipe and is contained in a 2 T solenoid magnetic field [39]. The ID is primarily responsible for the precision tracking of charged particles in the central $|\eta| < 2.5$ region. Moving out in radius, the next system consists of an electromagnetic and hadronic calorimeter designed to measure the energy of charged and neutral particles. The muon spectrometer (MS) is the outermost layer of the ATLAS detector and is designed to detect charged particles (primarily muons) that are not stopped in the calorimeter system. The MS system also independently measures the momentum of these charged particles in the $|\eta| < 2$ region. For this reason, the system also has three toroidal magnets that generate magnetic fields of approximately 0.5 T and 1 T in the barrel and end-cap regions respectively, in order to bend the charged tracks and obtain momentum measurements [39].

3.2.3 Inner Detector

The ATLAS inner detector (ID) system is contained in a 2 T solenoid magnetic field and is designed to provide precision tracking and momentum measurements of charged particles

within the pseudorapidity range of $|\eta| < 2.5$. To accomplish this goal, the ID is divided into three complimentary sub-systems. Closest to the interaction point, where the detector requires the most granularity, is a dense silicon pixel detector. Further away from the interaction point, the semiconductor tracker (SCT) uses silicon-microstrip detectors to reduce material effects and cost [40]. Surrounding the silicon pixel and strip detectors is the transition radiation tracker (TRT), which is made up of straw tube gas detectors. These three detector sub-systems are arranged in concentric circles in the barrel region and in disks in the end-caps. The transverse view of the ID is shown in fig. 3.5.




Figure 3.5: A transverse view in the R - ϕ plane of the barrel region of the inner detector. Tracks from the proton collisions traverse outward in radius, with most of them passing through the four silicon pixel layers, four silicon-microstrip (SCT) layers and approximately 36 straw tubes (TRT). In this figure, the fourth innermost layer of the pixel detector detector (IBL) is not shown. Figure taken from [39].

Pixel Detector

The silicon pixel detector is designed to provide high precision measurements as close to the proton interaction point as possible. When charged particles pass through the semiconducting silicon layers, they create electron-hole pairs that are collected by the sensors. A longitudinal view of the pixel detector is shown in fig. 3.6. The pixel system consists of four layers. The




Figure 3.6: A longitudinal view of the active region of the pixel detector. In this figure, the fourth innermost layer of the pixel detector (IBL) is not shown. Figure taken from [41].

innermost layer is called the insertable B-layer (IBL) and was installed in 2014 between Run 1 and Run 2 to increase the robustness and precision of the tracking and help mitigate luminosity effects [41]. The IBL is located at a radius of $r = 33$ mm and covers the $|\eta| < 3$ region of the detector. It also has the highest granularity of the four layers with a nominal pixel size of $50 \times 250 \mu\text{m}^2$ in $\phi \times z$. The other three layers of the detector have a nominal pixel size of $50 \times 400 \mu\text{m}^2$ in $\phi \times z$ (about 90% of the pixels) [39]. The intrinsic detector resolution of the IBL is 8 μm in ϕ and 40 μm in z , while the remaining pixel layers have a resolution of 10 μm in ϕ and 115 μm in z [40, 41]. The entire pixel detector detector consists of approximately 1744 sensors and 80 million readout channels.

Semiconductor Tracker

The semiconductor tracker (SCT) consists of a barrel detector and two symmetric end-cap detectors with approximately 61 m² of silicon-microstrip sensors [40]. The barrel detector has

four layers of silicon modules placed at a constant radius around the beam-pipe, while the end-cap detectors are located on either side of the barrel perpendicular to the beam. In the central region, each silicon module consists of four detectors with a constant strip pitch of 80 μm . Two detectors are wire bonded together and then two such pairs are mounted together at a 40 mrad stereo angle to provide three-dimensional tracking. This results in an intrinsic detector resolution of 16 μm in ϕ and 580 μm in z . Figure 3.7 illustrates an SCT module with four detectors connected to the hybrid readout electronics and a small stereo angle between the two pairs of detectors.




Figure 3.7: An SCT module showing the 4 silicon detectors connected to a readout hybrid assembly. Figure taken from [40].

Transition Radiation Tracker

The outermost layer of the ID is the transition radiation tracker (TRT) [42]. The TRT uses straw tube gas detectors that have a diameter of 4 mm and can be up to 150 cm in length. The TRT has about 50,000 straws in the barrel region and 320,000 straws in the end-cap region. In total, the TRT has about 420,000 readout channels and has a detector resolution of about 130 μm . Inside each straw tube, a 0.03 mm diameter gold-plated tungsten wire acts as the anode. The straw tubes are filled with a gas mixture that consists of Xe^4 , CO_2 and O_2 [39]. Charged particles that pass through the straw tubes ionize the gas creating electron-ion pairs. These electrons drift toward the anode wire where they are read out by the electronics. Relativistic charged particles crossing the material boundary in the TRT produce transition radiation photons in the X-ray range. These photons are absorbed in the active gas inside the

⁴Since the beginning of Run 2, TRT tubes with high gas leak rates have been replaced with a less expensive Ar-based gas mixture. During Run 1 a number of tubes in inaccessible areas developed leaks. To save costs, Ar was used in these modules, even though it is less efficient at absorbing X-ray photons than Xe [43].

straw tubes and serve as a detecting method for tracking and particle identification. Lighter particles, such as electrons, have larger gamma factors, $\gamma = E/mc^2$, and will produce larger amounts of radiation compared to heavier particles, such as pions. This is used to help with electron identification.

3.2.4 Calorimeter System

Outside of the solenoid magnet is a large volume calorimeter system. Electromagnetic and hadronic calorimeters measure the energy of the incident particles passing through them and provide coverage in the $|\eta| < 4.9$ region. The goal of the ATLAS calorimeter system is to absorb and measure the energy of most particles and limit the number of particles that pass through into the muon system [39]. Figure 3.8 shows a cut-away view of the calorimeter system.




Figure 3.8: A cut-away view of the ATLAS calorimeter system. Figure taken from [39].

Electromagnetic Calorimeter

The electromagnetic calorimeter (ECal) consists of a barrel and two end-cap detectors and is designed to measure the energy of incident electrons and photons [44]. It consists of alternating layers of “active” liquid argon (LAr) and “passive” lead absorbers. This type of calorimeter, with alternating “active” and “passive” layers, is known as a sampling calorimeter. The granularity varies across the different layers of the ECal and is highest in the $|\eta| < 2.5$ region, where tracking information from the ID is also available. The highest granularity of

the system is 0.025×0.025 in $\Delta\eta \times \Delta\phi$.

Measuring the energy of the particles relies on the behavior of an electromagnetic shower. In the high energy regime of the LHC, electrons mainly lose their energy via bremsstrahlung, while photons mainly lose their energy via pair production [45]. As the electrons and photons lose energy in the material, a multiplication of particles occurs. These electromagnetic showers typically occur in the absorbing material and are then detected in the active medium. This cascade of particles can be described in terms of the radiation length of the material X_0 , which is the distance over which an electron loses, on average, $1/e$ of its energy due to radiation. In the case of photons, one radiation length is roughly $7/9$ times the mean distance travelled by a photon before it undergoes pair production. The ATLAS ECal has a thickness of more than $22X_0$ in the barrel and $24X_0$ in the end-cap regions. This provides adequate containment of electromagnetic showers originating from bremsstrahlung and pair production.

Hadronic Calorimeter

The hadronic calorimeter is designed to measure strongly interacting particles such as hadrons. The ATLAS hadronic calorimeter consists of three components: the tile calorimeter, the LAr hadronic end-cap calorimeter (HEC) and the LAr forward calorimeter (FCal). The tile calorimeter is also a sampling calorimeter and uses steel as an absorber and scintillating tiles as the active material. Similar to electromagnetic showers, strongly interacting particles produce hadronic cascades. While an electromagnetic shower is characterized by its radiation length, X_0 , a hadronic shower can be described by its (usually much longer) interaction length, λ , which is the mean distance between hadronic interactions in a given material [45]. The tile calorimeter has three layers and a total thickness of 9.7λ at $\eta = 0$. The highest granularity of the tile calorimeter is 0.1×0.1 in $\Delta\eta \times \Delta\phi$. The HEC has two independent wheels per end-cap and extends out to $|\eta| < 3.2$. It is also a sampling calorimeter with alternating layers of passive copper plates and active LAr. The highest granularity of the hadronic end-cap calorimeter is 0.1×0.1 in $\Delta\eta \times \Delta\phi$. Finally, the forward calorimeter is designed to measure both electromagnetic and hadronic interactions. It covers the very forward region of the detector $3.1 < |\eta| < 4.9$ and its highest granularity is also 0.1×0.1 in $\Delta\eta \times \Delta\phi$. The forward calorimeter is approximately 10λ thick and consists of three layers. The first layer is made of copper and is optimized for electromagnetic measurements, while the two other layers are made of tungsten and liquid argon.

3.2.5 Muon Spectrometer

The outermost layer of the ATLAS detector is the muon spectrometer (MS). Muons require a dedicated system to detect and measure their momenta because, unlike most particles, they are not stopped in the calorimeter system. A particle's energy loss per unit path length largely

depends on its velocity (β) and reaches an energy loss minimum near $\beta\gamma \sim 4$ [45]. Since muons at the LHC are typically produced around this threshold, they are known as minimum ionizing particles. Muons energy loss per unit path length (dE/dx) is a tiny fraction of their energy and therefore they are not contained by the many layers of material in the calorimeter system. A cut-away view of the MS is shown in fig. 3.9.




Figure 3.9: A cut-away view of the ATLAS muon system. Figure taken from [39].

The MS consists of four different detector technologies that are designed to accomplish two main goals. The monitored drift tubes (MDTs) and cathode-strip chambers (CSCs) provide high precision tracking and momentum measurements of the incoming muons, while the thin-gap chambers (TGCs) and resistive plate chambers (RPCs) are primarily used to make initial trigger decisions. The entire system is embedded in a large toroidal magnetic field that is used to bend the charged muons in order to measure their momentum based on their curved trajectories. The magnetic field in the muon system points in the ϕ direction, which is mostly orthogonal to the muon trajectories and deflects charged particles in the η direction.

Precision Muon Tracking Chambers

Two gas detector systems provide precision tracking information in the MS. Monitored drift tube (MDT) chambers are 30 mm drift tubes that are full of a pressurized Ar/CO₂ gas mixture. Charged particles moving in a gas primarily lose energy due to excitation and ionization of nearby gas molecules. This results in electron-ion pairs that are free to move around in the medium. These ionization electrons are then collected at tungsten-rhenium wires located at

the centre of the drift tubes. The time difference between the initial passage of the particle and the arrival time of the ionization electrons at the central wire is related to how far away the incoming particle was from the centre of the tube and is used to provide a position measurement. The maximum electron drift time is about 700 ns. The MDTs provide pseudorapidity coverage up to $|\eta| < 2.7$ and are aligned tangentially to the ϕ direction to provide precision measurements along the z direction. The average spatial resolution of a single tube is about 80 μm .

Cathode-strip chambers (CSCs) are multi-wire proportional chambers that also contain an Ar/CO₂ gas mixture. These chambers are installed in the forward region of ATLAS ($2.0 < |\eta| < 2.7$), where the particle rate is much higher. The anode wires in the chambers are oriented in the radial direction. The cathode plates are segmented into strips that are parallel and perpendicular to the wires in order to provide 2D tracking. The position of the track is interpolated between the charge collected on neighbouring strips. On average the CSC has a position resolution of 60 μm per plate.

Muon Trigger Chambers

The MS also has two detector systems with fast readout to provide information about the muon tracks traversing through the detector that is used to make trigger decisions. In the barrel region ($|\eta| < 1.05$), layers of resistive plate chambers (RPCs) are used for triggering on muons. RPCs are gaseous parallel electrode-plate detectors with a typical position-time resolution on the order of 1 cm $\times$ 1 ns [46]. They are kept at a constant potential difference, which provides an electric field of about 5 kV/mm between the plates. The gas used is a mixture of C₂H₂F₄/Iso-C₄H₁₀/SF₆. Charged particles passing through the plates ionize the gas and the resulting charges are collected on the anode plates. Each RPC station consists of two independent layers, providing measurements in the η and ϕ directions.

The thin-gap chambers (TGCs) in the end-cap region ($1.05 < |\eta| < 2.7$) provide trigger information, as well as the azimuthal position of the tracks to complement the MDT measurements in the radial direction. TGCs are also multi-wire proportional chambers with a high-quenching gas mixture of CO₂ and n-pentane. The TGCs provide a timing resolution of less than 25 ns for 99% of tracks.

3.2.6 Trigger and Data Acquisition

The LHC delivers proton collisions to ATLAS at a rate of 40 MHz. With such a high collision rate, it is impossible (in terms of data readout speed and storage) to record every proton collision event. As a result, the ATLAS trigger and data acquisition system reduces the readout rate down to about 1 kHz of events that are saved for further analysis. The ATLAS experiment currently uses two stages of triggers known as the Level-1 trigger and High-Level trigger. A

schematic of the trigger and data acquisition system is shown in fig. 3.10.




Figure 3.10: The ATLAS trigger and data acquisition system (TDAQ) system in Run 2 showing the various parts of the system relevant to triggering, detector readout and data flow. Figure adapted from [47].

Level-1 Trigger

The Level-1 (L-1) trigger is a hardware-based system that uses information collected from the muon and calorimeter detector systems to make decisions about which proton collisions might contain interesting physics events [47]. The calorimeter information is used to find charged leptons, photons, jets and missing transverse energy, while the muon system information is primarily used to search for high-transverse momentum muons. For every event that is accepted by the L-1 trigger, the front-end (FE) electronics read out the data for that event. The data is initially processed and formatted in the Readout Drivers (RODs) before passing to the Readout System (ROS) where the data is buffered. In addition to selecting interesting events, the L-1 trigger also identifies so-called Regions-of-Interest (RoIs) within the detector that are further studied in the second trigger step. The L-1 trigger reduces the event rate to a maximum detector readout rate of 100 kHz. The maximum latency, which is the time that the detector systems must buffer the data in their front-ends while waiting for an L-1 decision, is $2.5 \mu\text{s}$. There is also a limit to how many events can be read out consecutively to

avoid overlapping readout windows, as well as a limit on the number of events stored within a given number of bunch crossings to prevent the front-ends buffers from overflowing.

High-Level Trigger

The second trigger level, called the high-level trigger (HLT), is software based. Dedicated reconstruction algorithms are used to reconstruct particles and estimate their properties in order to make the final event selection. Typically the HLT trigger decision is made within a few hundred milliseconds, which is long enough to run multiple reconstruction algorithms to extract physics information from the RoIs defined by the L-1 hardware trigger and decide whether the event passes the trigger conditions. If the event passes the L-1 hardware trigger, then the full detector is read out for a final trigger decision. The HLT software is based on the ATLAS offline software called Athena [48]. On average, the HLT trigger rate during an ATLAS data-taking period is 1.2 kHz, which corresponds to about 1.2 GB/s of data being sent to permanent storage.

3.2.7 ATLAS Detector Simulation

The simulation of the ATLAS detector is based on a GEANT4 simulation toolkit [49] and implemented in the ATLAS simulation framework [50]. This framework simulates how stable particles interact with the detector. Energy deposits in the sensitive detector elements are recorded as “hits”. Pileup effects and detector noise are added to the simulation by overlapping additional energy deposits from simulated minimum bias events on top of these “hits”. The trigger algorithms are also simulated and then the events are reconstructed with the same algorithms that are used for actual experimental data (see chapter 4). Simulating the complex detector geometry and detector response is computationally expensive. In fact, due to complex calorimeter shower developments, almost 80% of the full simulation time is used to simulate the calorimeter responses. For this reason, ATLAS fast simulation, or AFII, [51] was developed to parametrize the calorimeter response and reduce the simulation time. For the search described in this dissertation, AFII is used to simulate the signal. AFII works reasonably well because the signal only contains leptons and does not rely on any complex jet substructure variables that are not as well modelled in the ATLAS fast simulation.

Chapter 4

Object Reconstruction

This chapter provides an overview of the reconstruction of physics objects with a focus on those used in the analysis presented in this dissertation. The ATLAS detector is primarily designed to reconstruct prompt particles that are produced at the primary interaction point. As a result, standard reconstruction algorithms have been developed to identify different particles based on the signatures they produce in the various detector components. These standard reconstruction algorithms are developed and maintained by combined performance (CP) groups that provide recommendations to the entire collaboration on how to use standard physics objects in an analysis. These groups also provide calibrations and associated uncertainties for standard objects. In most cases, this means that analyzers directly use reconstructed physics objects, such as muons or electrons, in their analysis and only interact with the reconstruction algorithms if they also work in a CP group. The relevant standard object reconstruction algorithms are described in section 4.1. The reconstruction of basic components used to identify particles, including tracks, vertices and energy clusters is described in section 4.1.2. The reconstruction of electrons and muons is described in section 4.1.3 and section 4.1.4 respectively.

The search described in this dissertation also requires non-prompt object reconstruction. Non-standard reconstruction techniques are required to identify displaced objects and search for displaced decays of exotic particles. Since these objects are not part of the standard reconstruction chain, there are currently no CP groups that provide recommendations for non-standard physics objects. As a result, the analyzers often take on a greater role in the reconstruction and identification of non-standard objects. The reconstruction of non-standard objects relevant to this analysis are described in section 4.2, with a focus on displaced vertex reconstruction in section 4.2.3.

4.1 Standard Reconstruction

4.1.1 Overview

A schematic of different particles interacting with the ATLAS detector is shown in fig. 4.1. Muons pass through the entire detector and mainly interact with the ID and MS. Electrically neutral photons do not interact with the inner detector, but instead form a collimated shower

of particles in the ECal. Electrons⁵ also produce a similar shower in the ECal, but they carry electric charge and therefore also leave tracks in the ID. Neutrons do not interact with the inner detector due to their lack of electric charge and instead deposit their energy in the hadronic calorimeter creating a hadronic shower of particles. Protons cause the same hadronic shower, but also leave tracks in the inner detector.




Figure 4.1: A schematic of different particles interacting with the ATLAS detector in a section of the $x - y$ plane. Figure adapted from [52].

4.1.2 Tracks, Primary Vertices and Energy Clusters

Tracks, primary vertices and energy clusters are the basic inputs used to reconstruct the other physics objects discussed in this chapter. Charged particles passing through the ID and MS leave behind tracks that indicate their trajectories. Primary vertices are the location of the proton-proton interaction points and are reconstructed at the point of origin of a collection

⁵Throughout this dissertation, the term “electron” will usually refer to both electrons and positrons. This is also the case for all other particles and anti-particles (e.g. muons, quarks, protons etc.)

of tracks. Energy clusters are formed from groups of energy deposits left behind by particles that interact with the calorimeters.

Tracks

Standard track reconstruction starts with information from the pixel and SCT detectors. MS track segments are independently reconstructed and are discussed in section 4.1.4. Once track candidates are reconstructed in the silicon detectors, they are extended using the information from the TRT. A detailed description of the tracking reconstruction algorithms can be found in Refs. [53] and [54]. Track reconstruction begins by forming clusters from measurements in the pixel or SCT detectors. Charge deposits that are above a threshold in the pixel and SCT sensors are grouped into clusters. Three dimensional measurements, called space-points, are created from these clusters. Each space-point represents a place where a charge particle passed through the active material in the ID. After clusterization, an iterative combinatorial track-finding algorithm is run using the space-point measurements. Track seeds are formed by defining sets of three space-points. Track candidates are then reconstructed using a Kalman filter [55] to find additional space-points that are compatible with the initial track seed. Each time a new space-point is added, the track candidate is updated. In dense environments, particles pass so close to each other that it is possible to have a number of track candidates that share space-points. To resolve such ambiguities, the track candidates are processed individually and given a score, where a higher score indicates that the track candidate is more likely to represent the trajectory of a charged particle. This ambiguity solver primarily uses measures of track quality, such as the expected cluster multiplicity, the number of holes⁶ and the goodness of fit (χ^2) of the track, to determine the track score. Track candidates that fail to meet basic quality thresholds are rejected. After the track score is calculated, candidates that are assigned to multiple clusters are further processed by a neural network algorithm to split merged clusters. When all the track ambiguities are resolved, the track finding algorithm is run again to perform a final fit. Reconstructed tracks found in the silicon detectors are then used as inputs to find compatible sets of TRT measurements to extend the tracks into the TRT. Up to this point, an *inside-out* sequence is used to find tracks in the inner silicon tracker and extend the tracks outward in radius. While this method is very efficient, ambiguous hits can prevent track seeds from passing the ambiguity solver and tracks coming from secondary decays (e.g. kaon decays) or photon conversions may not have many silicon hits. To recover tracks in these cases, a second reverse *outside-in* sequence is run, starting with track segments in the TRT and extending the tracks inward to find any leftover compatible hits that were not used in the inside-out sequence. The final output of the track reconstruction algorithm is

⁶A track hole is defined as the intersection of a track candidate with a detector element that does not contain a matching cluster, but where one is expected.

a list of ID tracks that can be defined by the following set of five parameters:

$$(d_0, z_0, \phi, \theta, q/|p|) \quad (4.1)$$

where d_0 and z_0 are the transverse and longitudinal impact parameters, ϕ and θ are the usual azimuthal and polar angles and $q/|p|$ is the ratio of the particle's charge to momentum.

Primary Vertices

Primary vertices are defined as points in space where a proton-proton interaction is expected to have occurred. Details of the primary vertex reconstruction can be found in Ref. [56]. Vertex reconstruction can be separated into two steps: vertex finding and vertex fitting. In the former step, reconstructed tracks are associated to candidate vertices. The vertex fitting step then reconstructs the point of origin of these associated tracks to define the vertex position. To reconstruct a vertex, a track selection is applied to construct a list of input tracks and an initial seed position is chosen. The selected tracks and initial seed position are used to determine a best-fit position and then the tracks are weighted up or down depending on how compatible they are with the best-fit position. Once the final vertex position is determined, any track that is not compatible with this position is removed from the vertex. These rejected tracks can then be used to form another vertex. This procedure is repeated with all the remaining tracks in the event until no new vertices are formed.

The primary vertex reconstruction is also responsible for identifying the primary vertex associated to the hard-scatter event. The location of the primary hard-scatter vertex is essential to the reconstruction of the full kinematics of the event. The highest sum of the squared transverse momenta of the tracks ($\sum |p_T|^2$) is used to distinguish the hard-scatter vertex from pileup vertices. This choice is based on the assumption that charged particles produced in a hard-scatter event have, on average, a larger transverse momentum spectra than tracks produced in pileup collisions. For the analysis in this dissertation, the primary hard-scatter vertex is defined as the primary vertex with the highest $\sum |p_T|^2$ of the tracks associated to it.

Energy Clusters

Energy clusters are formed from a set of calorimeter cells that contain energy deposits and share common edges. ATLAS has two different clustering algorithms. The sliding-window reconstruction algorithm was used to build energy clusters for electron and photon reconstruction until the end of 2016. The second topological clustering (topo-cluster) algorithm was used from 2017 onward for electron and photon reconstruction and was also used to build energy clusters for the reconstruction of hadronic showers, also known as jets. For the analysis presented in this dissertation, the primary use of energy clusters is for electron reconstruction

and, for this reason, the remainder of this section will focus on electromagnetic energy cluster algorithms used to reconstruct electrons.

The sliding-window reconstruction method divides the calorimeter in η and ϕ into a grid of fixed elements known as “towers”. Inside each of these towers, the energy of all cells in the longitudinal layers is added to form the “tower energy”. Electromagnetic-energy cluster candidates are seeded using energy deposits that are larger than 2.5 GeV in a window of size 3×5 tower units in $\eta \times \phi$. Then, the centre of the 3×5 window slides around in steps of 0.025, in either the η or ϕ direction, to search for other energy deposits to add to the cluster. This seed-to-cluster process is repeated until this has been done for every element in the calorimeter. The sliding-window algorithm is described in more detail in Ref. [57].

The topological clustering algorithm uses signal-significance patterns to form 3D topological cell clusters, also known as topo-clusters. Cluster seeds are formed from cells that have a high signal significance defined by the ratio of the signal to the expected cell noise. Neighbouring cells above a certain signal significance threshold are iteratively added to the cluster. This process continues until there are no cells left to be added to clusters. For electron reconstruction, topo-clusters are only formed using energy from ECal cells. The topo-cluster algorithm is described in more detail in Ref. [58].

4.1.3 Electrons

Electrons are charged particles that leave tracks in the inner detector and deposit most of their energy in the electromagnetic calorimeter. ATLAS uses tracks and energy clusters to identify and reconstruct electrons. An overview of the electron reconstruction can be found in the Refs. [59] and [60].

Reconstruction

Electrons lose most of their energy due to bremsstrahlung when interacting with material. This results in a collimated electromagnetic shower of particles that is normally reconstructed as part of one energy cluster in the ECal. As a result, it is possible to produce and match multiple tracks that originated from the primary electron to a single energy deposit. Therefore, an electron is defined as an object that consists of an energy cluster with at least one matched track. Since close matching of tracks to energy clusters is required to identify electrons, the reconstruction of electrons occurs in the central precision region of the ATLAS detector with $|\eta| < 2.5$. A schematic illustration of the path of an electron through the detector is shown in fig. 4.2.

Electron reconstruction starts by finding energy clusters (described in section 4.1.2) in the electromagnetic calorimeter with a transverse energy $E_T > 2.5$ GeV. The cluster reconstruction efficiency for electrons ranges from 65% at $E_T = 4.5$ GeV, to 95% at $E_T = 7$ GeV,




Figure 4.2: A schematic of the path of an electron through the ATLAS detector. The solid red trajectory shows the path of the electron traversing through the tracking system (pixel, SCT and TRT layers) and entering the electromagnetic calorimeter. The dashed red trajectory indicates the path of a photon produced by the interaction of the electron with the material in the tracking system. Figure taken from [59].

to more than 99% for $E_T > 15$ GeV, as shown in fig. 4.3 [59]. Next, ID tracks are found using the reconstruction method described in section 4.1.2. Figure 4.3 shows that the track reconstruction also varies as a function of E_T and is 80% efficient at $E_T = 1$ GeV and more than 98% efficient above $E_T = 10$ GeV.

Since charged electrons lose energy when traversing through material, a subsequent fitting procedure, using a Gaussian-sum filter (GSF) [61], takes into account non-linear effects, or “kinks”, in the tracks related to bremsstrahlung. Radiative energy losses decrease the momentum of the electrons, causing them to bend more in the magnetic field. Taking into account this additional energy loss improves the reconstruction of electron track parameters, such as d_0 , that are related to the bending of the track.

A likelihood-based (LH) method is used to identify prompt electrons in order to ensure that only high-quality electrons are used for analysis and to suppress non-prompt backgrounds from photon conversions, jets and heavy flavour decays. The likelihood is calculated using properties of the electron track, information about the electromagnetic shower shape and the compatibility of the energy cluster with the primary electron track. The probability distribution functions for the prompt electron signal and non-prompt backgrounds are derived from simulation and corrected to accurately model data. Electron identification operating points are defined using a likelihood discriminator that increases in value for more “electron-like” objects and requires the tracks satisfy additional tracking requirements.




Figure 4.3: The total reconstruction efficiency for simulated electrons in a single-electron sample as a function of the true transverse energy, E_T , for each step of the electron reconstruction: the energy cluster reconstruction (red triangles), track reconstruction using the Global χ^2 Track Fitter (blue open circles), cluster and track reconstruction using GSF tracking (yellow squares), and the final electron-candidate reconstruction, which includes the track-to-cluster matching (black circles). Since the cluster reconstruction requires that cluster seeds have $E_T > 2.5$ GeV, the total reconstruction efficiency is less than 60% below 4.5 GeV (dashed line). Figure taken from [59].

The prompt electron identification working point that is relevant to the work described in this dissertation is called *LHMedium*. The *LHMedium* working point has the following tracking requirements: at least two hits in the pixel detector with at least one hit in the innermost (functional) layer, at least seven silicon (pixel + SCT) hits and a minimum E_T of 4.5 GeV. The *LHLoose* working point is similar to *LHMedium*, except that a looser cut on the likelihood discriminator is used. *LHVeryLoose* electrons are also similar to those selected by the *LHMedium* and *LHLoose* working points, except there is no longer a required hit on the first working pixel layer and there is only one required pixel hit. The *LHLoose* and *LHVeryLoose* working points are not used in the analysis described in this dissertation, but serve as a comparison for the displaced electron identification discussed in section 4.2.2. The working point *LHTight* is not used in the analysis, but is similar to *LHMedium* except that it has a smaller identification efficiency. Figure 4.4 shows the electron identification efficiency for the *LHLoose*, *LHMedium* and *LHTight* working points as a function of E_T .




Figure 4.4: The electron identification efficiency in $Z \rightarrow ee$ events in data as a function of E_T . The efficiencies are obtained by applying data-to-simulation efficiency ratios measured in $J/\psi \rightarrow ee$ and $Z \rightarrow ee$ events to $Z \rightarrow ee$ simulation. The *LHMedium* operating point efficiency curve is plotted using red squares. Figure taken from [60].

Isolation

In general, high- p_T electrons are expected to be isolated from other particles and, as a result, their detector signatures should be spatially separated from other detector deposits. To measure the isolation, variables are constructed to measure the amount of detector activity near the electron object. Usually this is done by summing the transverse cluster energy (calorimeter-based isolation) or transverse momenta of tracks (track-based isolation) that are in a cone of radius $\Delta R = \sqrt{\Delta\phi^2 + \Delta\eta^2}$ around the electron candidate. In this dissertation, track-based isolation requirements are applied, which require that the ratio of the sum p_T of all tracks in a cone $\Delta R = 0.3$ around the electron to the electron's p_T is less than some threshold. For prompt electrons this threshold is 0.05 ($\sum p_T^{(0.3)}/p_T < 0.05$) and for displaced electrons the threshold is 1.0 ($\sum p_T^{(0.3)}/p_T < 1.0$).

4.1.4 Muons

Muons are reconstructed mainly using information from the ID and MS system. An overview of the muon reconstruction can be found in the Refs. [62] and [63].

Reconstruction

Muon reconstruction begins by independently reconstructing tracks in the ID and MS. Track reconstruction of charged particles in the ID is described in section 4.1.2. In the MS, muon track segments are formed by combining hits from the layers of muon chambers. These track segments are initially fit using a loose pointing constraint based on the location of the primary vertex and the approximate trajectory of a charged particle bending in the magnetic field. Segments are then combined, starting with the ones in the middle layers of the MS, to form muon track candidates. In most regions of the detector, two track segments are required to reconstruct a track candidate and the same segment is initially allowed to be used in several track candidates. A global χ^2 fit is performed using the hits associated to each track candidate to determine if the track passes the χ^2 selection criteria.

Four muon types are defined depending on what detector information is used to reconstruct the muon [63]:

- Combined muon: a track is reconstructed using a global fit of the hits in the ID and MS. During the fit, MS hits may be added or removed in order to improve the fit quality. An *outside-in* approach is used for the reconstruction, starting with the reconstruction of tracks in the MS and then extrapolating inwards to match MS to ID tracks.
- Inside-out muon: an ID track is extrapolated to at least three aligned MS hits. This method does not rely on the MS track reconstruction and therefore recovers some efficiency in regions where the MS coverage is limited. These muons are usually merged together with the combined muons.
- Segment-tagged muon: a track from the ID is classified as a muon if the track can be associated to a segment in the MS. These muons are only used when the muon crosses one layer of the MS.
- Calorimeter-tagged muon: a track from the ID is classified as a muon if it matches an energy deposit in the calorimeter that is consistent with a minimum ionizing particle. This type of muon has the lowest purity of all the muon types.
- Extrapolated muon: a track is based only on information from the MS and has a loose requirement that the track originated from the primary vertex. Extrapolated muons are generally only used to extend the acceptance for muon reconstruction into the $2.5 < |\eta| < 2.7$, which is outside of region covered by the ID.

For the work described in this dissertation, all muon types, except for the lowest purity calorimeter-tagged muons, are initially used to seed the displaced vertex reconstruction (see section 4.2.3). However, in the final event selection only combined muons (including inside-out muons) are used.




Figure 4.5: The muon reconstruction efficiency for the *Medium* identification working point as a function of muon p_T in the region $0.1 < |\eta| < 2.5$. This reconstruction efficiency is obtained with $Z \rightarrow \mu\mu$ and $J/\Psi \rightarrow \mu\mu$ events. The error bars on the efficiencies indicate the statistical uncertainty. The panel at the bottom shows the ratio of the measured to predicted efficiencies, with statistical and systematic uncertainties. Figure taken from [64].

Identification

After reconstruction, muon identification requirements are applied in order to suppress non-prompt muon backgrounds, such as decays of charged hadrons. The relevant working point for the muons used in this dissertation is *Medium* and the reconstruction efficiency for *Medium* muons are shown in fig. 4.5. Selection criteria for high-quality muons include requirements on the number of hits in the ID or MS, the properties of track and the compatibility of the measurements between the two subsystems. Muons originating from the decays of charged hadrons will produce a “kink”, or change in direction, in the reconstructed track. As a result, this can lead to a poor fit quality for the combined track, as well as incompatible momentum measurements between the ID and MS tracks.

Identification criteria, known as working points, are defined mainly using the muon selection efficiency and purity in simulations of $Z \rightarrow \mu\mu$ and $J/\psi \rightarrow \mu\mu$ events. For the identification of combined muons the following variables are used:

- q/p significance: defined as the absolute value of the ratio between the charge, q , and the momentum, p , of the muon divided by the sum in quadrature of the uncertainties on q and p ;
- ρ' : defined as the absolute value of the difference between the transverse momentum measurements in the ID and MS divided by the p_T of the combined track;
- normalized χ^2 of the combined track fit.

The *Medium* working point is used to identify muons in the analysis in this dissertation and is also the default muon working point in ATLAS. This working point is designed to minimize the systematic uncertainties associated with muon reconstruction and calibration. The collection of *Medium* muons only contains combined and inside-out muons that have hits in at least two MS layers and a q/p significance of less than seven to ensure a loose compatibility between ID and MS measurements. The muon reconstruction efficiency with this selection is around 97% for muons with $4 < p_T < 20$ GeV.

Isolation

Prompt muons from the decays of heavy particles, such as W , Z or Higgs bosons, are often spatially isolated from other particles. Muons from semi-leptonic decays of heavy flavoured hadrons are typically produced near a jet and are not well separated from other particles. As a result, the muon isolation, which measures how much detector activity is around the muon, can be a useful variable to select prompt muons. Similar to electrons, a track-based isolation is calculated using the scalar sum of the transverse momenta of all tracks within a cone of radius $\Delta R = \sqrt{\Delta\phi^2 + \Delta\eta^2}$ around the muon candidate. In this dissertation, a track-based isolation requirement is used to select prompt muons where the ratio of the scalar sum of transverse momenta of all tracks in a cone of $\Delta R = 0.3$ to the transverse momentum of the prompt (displaced) muon candidate is $\sum p_T^{(0.3)}/p_T < 0.05(1.0)$.

4.2 Non-Standard Reconstruction

4.2.1 Displaced Tracks

Standard track reconstruction, described in section 4.1.2, is optimized to reconstruct charged particles that originate from the primary interaction point. However, standard track reconstruction efficiency falls off rapidly for charged particles that originate from decays occurring

a significant distance from the interaction point. Therefore, analyses that search for displaced signals reconstruct displaced tracks using a dedicated large radius tracking algorithm. The details of this algorithm can be found in Ref. [65].

Large radius tracking (LRT) is performed after the standard inside-out and outside-in track reconstruction. This dedicated tracking step uses the leftover silicon hits to reconstruct displaced tracks that are not required to point back to the interaction point. LRT follows the same reconstruction strategy as inside-out tracks, but uses relaxed constraints on the track impact parameters (d_0 and z_0) and number of hits. The key differences between standard and large radius track reconstruction are outlined in table 4.1.

Table 4.1: Main differences between standard and large radius tracking (LRT).

Track parameter	Standard	LRT
Maximum d_0	10 mm	300 mm
Maximum z_0	250 mm	1500 mm
Maximum track $ \eta $	2.7	5.0
Minimum track p_T	500 MeV	900 MeV
Minimum silicon hits	7	7
Minimum unshared silicon hits	6	5
Seed extension	Combinatorial	Sequential

The large radius track seeding uses a sequential Kalman filter to prevent the creation of multiple track candidates from different combinations of space-points. After track candidates are reconstructed, the track ambiguity solving step uses the same relaxed constraints, and any tracks that pass the ambiguity solver are extended into the TRT. The resulting collection of displaced tracks is then merged with the collection of standard tracks.

LRT is optimized for high track reconstruction efficiency and this comes at the cost of a large number of fake tracks. This means that the LRT reconstruction algorithm is very CPU intensive and requires a large amount of disk space to store the output track collection. As a result, the LRT reconstruction step is not run by default. Instead, a subset of the recorded data events, based on the output of the standard reconstruction, is selected for this special reprocessing. A description of the selections applied to filter data events for this analysis is found in section 5.7.2.

LRT reconstruction efficiency varies as function of the distance from the primary interaction point. BSM signals with large production radii ($r_{\text{prod}}^2 = x^2 + y^2$) that decay into leptons and hadrons are used to study the reconstruction efficiency:

$$\text{reconstruction efficiency} = \frac{\text{number of signal particles matched to reconstructed tracks}}{\text{number of signal particles}} \quad (4.2)$$

The dependence of the reconstruction efficiency on r_{prod} is shown in fig. 4.6. At low r_{prod} , standard tracking is very efficient and there are not many leftover hits to form tracks in




Figure 4.6: Track reconstruction efficiency for displaced charged particles produced by the decay of long-lived signal particles. The efficiency for standard tracking, LRT and the sum of the two is shown as a function of the radius of production of the displaced particles ($r_{\text{prod}}^2 = x^2 + y^2$). Displaced leptons are shown on the left and displaced hadrons on the right. Figure taken from [65].

the LRT step. However, for $r_{\text{prod}} > 50$ mm, fewer than 20% of hadrons (30% of leptons) from displaced decays would be reconstructed without LRT. Beyond $r_{\text{prod}} > 50$ mm, LRT helps to recover efficiency for displaced tracks. However, the efficiency in fig. 4.6 is still well below 100% for large r_{prod} , even with the addition of large radius tracks. This is due to inefficiencies in the tracking reconstruction algorithm, in particular on the minimum number of silicon hits, material effects from interactions with the detector and the hermetic design of the ATLAS detector that is optimized to reconstruct prompt particles that decay at the primary interaction point.

4.2.2 Displaced Leptons

Standard electron and muon reconstruction and identification algorithms, described in section 4.1.3 and section 4.1.4 respectively, are primarily designed to reconstruct prompt leptons produced at the primary interaction point. Displaced leptons are mainly reconstructed using the same standard algorithms, with the exception that the collection of ID tracks now uses a combination of both standard and large radius tracks. However, in order to reconstruct displaced electrons, the $|d_0|$ and the number of silicon hits requirements, that are applied in addition to the track reconstruction requirements, are removed. This is primarily necessary for electrons, since the number of silicon hits is used to separate electrons and photons. Without the removal of d_0 and hit requirements, many reconstructed displaced electrons end up being classified as photons.

After reconstruction, identification working points are used to select good quality elec-

trons and muons. For muons, the standard working points (described in section 4.1.4) are relatively efficient at identifying muons from the displaced decay of the signal described in this dissertation. As a result, the default muon working points are not modified. However, for displaced electrons, the standard identification working points require the electrons to have hits in the innermost pixel layer or a certain number of hits in the silicon detectors. Thus, the search described in this dissertation uses a custom working point called *VeryVeryLoose* is defined to identify displaced electrons. This working point is similar to the standard *LHVeryLoose* working point, except that any d_0 related requirements are removed, along with the requirements on the number of pixel and silicon hits.

4.2.3 Displaced Vertices

For certain long-lived particle signals, such as those presented in this dissertation, we are also interested in determining a place of common origin for displaced tracks, which is called a displaced vertex (DV). A reconstruction algorithm, known as Vertex Secondary Inclusive (VSI), is used to reconstruct DVs in the ID. ATLAS also has other non-standard reconstruction techniques designed to reconstruct decays in other parts of the detector (e.g. the MS system), but this section will only described DV reconstruction in the ID, as this is the focus of the search described in this dissertation.

This section discusses the VSI algorithm in general, while section 5.5.2 discusses the custom VSI configuration used in the analysis described in this dissertation. The main steps of the VSI algorithm can be found in fig. 4.7. More details about this algorithm can be found in Ref. [66].




Figure 4.7: The main steps in the VSI displaced vertexing algorithm.

Track Selection

The first step in the algorithm is to further select tracks from a full set of reconstructed tracks. This step is very important to maintain a high signal efficiency, while reducing the number of fake vertices and CPU time. By default, the input to the secondary vertexing includes both standard (see section 4.1.2) and large radius (see section 4.2.1) tracks. However, the algorithm is also configurable to alternatively select ID tracks matched to muon or electron

objects (see section 4.1.3 and section 4.1.4 for a description of how tracks are matched to electron and muon objects respectively). Selecting only muon and electron-matched tracks configures VSI to only reconstruct displaced lepton vertices. This lepton-only track selection is also important when reconstructing vertices with electrons, since this method can select the electron tracks that are reconstructed using the GSF instead of the ATLAS Global χ^2 fitter. This is important because GSF electron tracks have improved reconstruction of track bending parameters, such as d_0 , which are crucial to the vertex reconstruction.

Other track selection functions are also included in the algorithm and these selections can be optimized to maximize signal efficiency and reject background vertices. Examples of track cuts that can be applied include track p_T , fit quality (χ^2), d_0 and z_0 . In order for the vertex reconstruction to continue, at least two tracks per event must be selected in the track selection step.

Two-Track Vertex Finding

The reconstruction sequence continues with a two-track vertex finding step, which attempts to form vertices with all compatible two-track pairs. Again, to help reduce the CPU time and the number of fake vertices, this step has a few different selection requirements for each compatible two-track pair to pass to the next vertexing stage. First, a fast fit is done for each track pair to determine an initial vertex position. Then, the position and momentum of the two-track DV are used to apply kinematic cuts using the vertex pointing vector, $\vec{r}_{DV} = \vec{r}(r_{DV}, z_{DV}, \phi_{DV})$, which is defined as the relative position of the fitted vertex with respect to the primary hard-scatter vertex. These cuts are designed to reject vertices that are compatible with the decay of a particle originating from the primary vertex. There is also a two-track vertex forming d_0 cut, which requires that at least one track has a d_0 greater than a user-specified value in order for the two-track pair to be accepted.

Next there is a hit pattern consistency check for the two-track vertex. A *classical* hit pattern check has two requirements. The first *requires* ID hits from the tracks just ahead of the vertex position ($r > r_{DV}$) and the second *forbids* hits from the track just behind the vertex position ($r < r_{DV}$), as illustrated in fig. 4.8. The idea here is to check that the two tracks look like they originate from the fitted vertex position. If there are hits behind the vertex, or no hits just ahead of the vertex, then the vertex is not consistent with the decay of a displaced particle.

VSI can also be configured to use an extrapolation method where the *classical* hit pattern consistency checks are complemented by information about disabled ID modules. The ATLAS detector is not perfect and, in any given run, various ID modules can be labelled as “not active” either because they are not functioning properly or may not be collecting data at all. As a result, the efficiency of the hit pattern check that requires hits just in front and forbids hits




Figure 4.8: A schematic illustration of the fake tracks rejection method. In example (a), the vertex is between the pixel B-Layer and Layer-1. The tracks of the reconstructed secondary vertex must not have hits in the layers behind the vertex (i.e. IBL and B-Layer), and must have hits in the closest layer ahead of the vertex (i.e. Layer-1). In example (b), the vertex is close to the Pixel B-Layer. The tracks in this vertex are not allowed to have hits in the IBL, but may have hits in the B-Layer, and must have hits on the Layer-1. Analogous requirements are made on vertices close to the other layers. Figure taken from [66].

just behind the vertex can be affected by disabled modules. To improve the efficiency, run-by-run ID conditions are taken into account in the VSI algorithm using a *disabled* hit pattern check. If a module consistent with a track is “not active” just ahead of the candidate vertex, then the *require* hits criterion is skipped and the vertex will be accepted. On the other hand, if an ID module consistent with a track is “not active” just behind of the candidate vertex, then the *forbidden* hits criterion cannot be determined and the vertex will be rejected.

N-Track Vertex Finding and Vertex Merging

After the two-track vertex finding step, N -track vertices are formed using an incompatibility graph method that determines the sets of tracks that are fully compatible with each other. Here, the term compatible means that the two tracks form a vertex with the required quality. VSI focuses on incompatible track pairs rather than compatible pairs. A list of incompatible track pairs are represented in a graph where the nodes (points) represent tracks and the lines represent incompatible pairs (see fig. 4.9 right). This graph is then used to find the combination of nodes that leave behind only fully-isolated nodes when they are removed from the graph. For the example in fig. 4.9, if nodes (3, 4, 5, 6) are removed from the graph, then only nodes 1 and 2 remain and there are no longer any lines connecting incompatible pairs. This means tracks 1 and 2 form a compatible vertex. Alternatively nodes (1, 2, 6) could be

removed and tracks 3, 4 and 5 would also form a compatible vertex.




Figure 4.9: An example of finding N -track vertices using an incompatibility graph. The layout of tracks and vertices in an event (a) and the incompatibility graph (b) that is formed by connecting two tracks when the two-track vertices are determined to be incompatible. Figure taken from [67].

Once N -track vertices are formed, it is possible that a track may be compatible with multiple vertices as illustrated in fig. 4.10. In this case, there is an ambiguity as to which vertex the track belongs to. At most, each track is allowed to be associated to one vertex and thus this ambiguity needs to be resolved. For each track, a list of vertices that share the track is created. Then, the track that shares the most vertices is identified. The goodness of fit (χ^2) of this track to each of these vertices is calculated. The vertex with the worst χ^2 for the track is found and, if χ^2 is larger than a user-specified threshold, then the track is removed from that vertex and this process is repeated. Otherwise, if the significance of the distance between the two vertices that share the track is smaller than another user-specified threshold, or at least two tracks are shared, then the vertices are merged together. At this point, if the track sharing ambiguity can not be resolved, then the track is simply removed from the vertex with the worst χ^2 . This process is then repeated until no shared tracks remain. Once there are no vertices with shared tracks, a vertex fit is performed and, if the distance between any two vertices is less than another user-specified threshold, then these vertices are also merged.

Track Attachment

At this point, all the tracks that were *not* selected in the track selection step are not yet associated to any DV. The final step in this reconstruction algorithm attempts to recover lost tracks, by running the track attachment step. While this has the potential to recover signal tracks with poor quality that did not pass the track selection, this step can potentially attach




Figure 4.10: An example of track sharing by multiple vertices. Figure taken from [67].

irrelevant or fake tracks. This step allows non-selected tracks to be attached or associated to the vertex provided that the tracks meet some basic quality checks (p_T , track χ^2 etc.) and that they are close enough to the existing candidate vertex (both in d_0 and z_0 significance). This final track attachment step can be turned on or off.

Displaced Vertex Variables

After the DV reconstruction is completed, the final vertices are recorded and a few different variables are calculated for the displaced vertices and the tracks associated to them. For example, every track stores if it was initially selected in the track selection step or not. This is useful when calculating the VSI track selection efficiency. Another variable records if the track was attached to the vertex in the final track attachment step of VSI or not. This can be useful in some analyses to identify fake or irrelevant tracks that are attached to the vertices and help regain some signal efficiency. Finally, new track variables, calculated with respect to the reconstructed DV position, are computed. These variables include track kinematics (d_{0DV} , z_{0DV} , p_{TDV} , η_{DV} , a and ϕ_{DV}), the square of the uncertainty for d_0 , z_0 and q/p (σ_{d_0} , σ_{z_0} and σ_p) and the goodness of fit of the track to the vertex (χ^2). These variables are useful since the original track quantities are calculated with respect to the primary vertex, but for displaced tracks this prompt-particle assumption is no longer valid.

In addition to the displaced track quantities, DV quantities are also computed. The DV kinematics are found using the sum of the momentum of the associated tracks using the track parameters calculated with respect to the DV (p_{TDV} , η_{DV} , ϕ_{DV}). This same momentum sum along with a pion mass assumption for each track is used to compute the invariant mass of the DV (m_{DV}). The radial position (L_{xy}) of the vertex is also used as a measure of vertex displacement. This variable is computed using the x and y position of the vertex, $L_{xy} = \sqrt{x^2 + y^2}$.

Chapter 5

Search for Heavy Neutral Leptons

This chapter summarizes a search for long-lived heavy neutral leptons (HNLs) performed with data from the ATLAS detector at the LHC. Many BSM theories predict new particles that have short lifetimes. These signatures decay so quickly that are considered *prompt*. This means that the distance the particle travels before it decays is less than or equal to the position resolution of the detector trying to measure it. In contrast, searches for *metastable* or long-lived particles (LLPs) probe regions of phase space within resolvable distances for the detectors and look for exotic decays that would be rejected by most prompt searches. As discussed in chapter 2, the HNL lifetime depends on its mass m_N and its mixing angle with SM neutrinos $|\Theta|^2$ (see eq. (2.8)). HNLs with GeV scale mass and mixing angles between 10^{-3} and 10^{-6} would produce sufficiently displaced decays that could be detected at high-energy collider experiments. These signals are particularly interesting as they could be found in relatively under-explored regions of phase space.

This analysis searches for a long-lived HNL produced when a W boson decays into a lepton (muon or electron) and a neutrino that mixes with the HNL. This analysis targets final states that are consistent with a fully-leptonic HNL decay. The experimental signature consists of an isolated prompt lepton (e/μ) from the W boson decay accompanied by a displaced vertex from the HNL decay with either two muons ($\mu\mu$), two electrons (ee) or one muon and one electron ($e\mu$) between 4 and 300 mm from the primary vertex.

This chapter begins with a brief overview of experimental searches for HNLs and then describes the analysis strategy and results of the ATLAS search for displaced HNLs performed using proton-proton collision data collected between 2015 and 2018.

5.1 Previous Searches for HNLs

Previous searches for HNL signals on the electroweak scale include analyses with the DELPHI experiment at the Large Electron-Positron collider and most recently with collider experiments at the LHC. The DELPHI search studied 10^6 events from Z decays and placed constraints in the HNL mass range between 3 and 50 GeV [68]. More recently, with the LHC's very high rate of W boson production, it is possible to search for HNLs using W decays and identify HNLs using their fully or semi-leptonic decay modes.

CMS has published an analysis that searches for events with three prompt charged leptons and searches for HNLs in the mass range of 1 GeV - 1.2 TeV [69]. At masses below 20 GeV, the objects in the decay are increasingly displaced and prompt searches become inefficient. Thus, ATLAS has both prompt and displaced HNL searches that provide constraints on the HNL mass up to 50 GeV [70]. Both CMS and ATLAS have also published papers searching for general displaced signatures, but most of these searches do not have sensitivity to HNLs due to high- p_T requirements or a requirement that two or more displaced vertices are constructed in the same event.

A 2019 ATLAS result previously searched for displaced HNLs using pp data collected in 2016 and set limits on the mixing angle to muon neutrinos, $|U_\mu|^2$, and excluded models with $|U_\mu|^2$ between 10^{-3} and 10^{-6} for HNL masses between 4.5 and 10 GeV [70]. The range of models excluded by the 2019 displaced HNL search are shown in the 95% confidence level (CL) red exclusion curves in fig. 5.1. The 2019 prompt HNL analysis limits are also shown in fig. 5.1 and extend out to 50 GeV.





Figure 5.1: Observed 95% confidence level (CL) exclusion limits for HNL to muon neutrino mixing angle $|U_\mu|^2$ versus HNL mass m_N for the prompt signature (the region above the black line is excluded) and the displaced signature (the region enclosed by the red line is excluded). The solid lines show limits assuming lepton-number violation (LNV) for half of the decays and lepton-number conservation (LNC) for the other half. The long-dashed line shows the limits assuming LNC for all decays. The green and yellow bands indicate the ranges of expected limits obtained within ± 1 and ± 2 standard deviations of the median limit reflecting uncertainties in the signal and background yields. Figure taken from [70].


Beam-dump experiments at CERN and Fermilab also search for HNLs with smaller mass (e.g. $m_N < 2$ GeV). Summary plots of the sensitivity to HNLs are shown in fig. 5.2. In each plot, dominant mixing to either electron, muon or tau neutrinos is assumed. These plots do not include the 2019 ATLAS prompt and displaced search results (or the results from the

5.1. Previous Searches for HNLs


search described in this dissertation) because these results were published after the summary plots were prepared.



(a) Electron-only mixing



(b) Muon-only mixing



(c) Tau-only mixing

Figure 5.2: Sensitivity to HNLs with dominant mixing to first, second and third lepton generations. Current bounds on the mixing (filled regions) and prospects for new results in the next 5-15 years are shown as a function of the HNL mass m_N . Figures taken from [71].

5.2 Search Overview

The signal model in this search is a long-lived HNL produced from a W boson decay. The analysis signature consists of a prompt lepton that is produced with the HNL and a displaced vertex from the fully-leptonic decay of the HNL (see fig. 2.3). The analysis considers six channels that are defined as $\alpha - \beta\gamma$ according to the flavour of the prompt lepton (α) and the flavours of the two charged leptons produced in the HNL decay ($\beta\gamma$): $\mu - \mu\mu$, $\mu - \mu e$, $e - e\mu$, $e - ee$, $e - \mu\mu$ and $\mu - ee$. Both LNC and LNV decays are considered in this analysis.

Simulated signal samples are generated to cover a mass range from 3 to 20 GeV and proper decay lengths ($c\tau_N$) of 1, 10, 100 mm. HNL theory predicts a relationship between the mixing angle, mass and lifetime (see eq. (2.8)). Beyond ~ 20 GeV, HNL models that experiments at the LHC are currently sensitive to are increasingly prompt. Therefore, this displaced search targets models with $m_N < 20$ GeV. At the lower HNL mass limit, signals with $m_N < 3$ GeV are overwhelmed by backgrounds.

This analysis exploits the prompt lepton from the W boson decay to make use of the ATLAS single-lepton triggers. Selected events are reconstructed using displaced tracking algorithms (see section 4.2.1). Signal selections on electron and muon identification working points, the mass and radius of the displaced vertex, the HNL mass and invariant mass of the tri-lepton system are optimized for the six decay channels.

The primary background is from leptons that randomly cross in the ID that are reconstructed as a two-lepton DV. For ee vertices, interactions with the detector material and support structures also form a source of background. At lower masses, there are also SM resonances, which contribute to the background.

This analysis includes many improvements compared to the 2019 ATLAS result. It now considers additional channels, including channels with prompt electrons and channels with different flavours between the prompt and displaced leptons (e.g. $\mu - ee$ and $e - \mu\mu$). The custom vertexing configuration used to reconstruct lepton-only DVs is also novel, as is the discriminatory variable that reconstructs the HNL mass m_{HNL} using the available information from the three charged leptons in the signal combined with the W boson mass and the HNL flight direction. Details about the custom vertexing configuration and HNL mass reconstructions can be found in section 5.5.2 and section 5.5.3 respectively. A new background estimation method was also developed to provide a more sophisticated way to estimate the dominant random track crossing background that reduces some of the statistical limitations faced in the 2019 ATLAS analysis. This background estimation method is discussed in section 5.8.6. Finally, this search also sets the first experimental limits on two benchmark models that the LHC-long-lived-particle community developed in recent years to target HNLs under the inverted and normal neutrino mass hierarchy models. These benchmark results, as well as results for the muon-only and electron-only mixing scenarios, are presented in section 5.11.

5.3 HNL Models

Eight HNL models are used to interpret the results in this search. They include models with single-flavour and multi-flavour mixing in various lepton number conserving/violating (LNC/LNV) scenarios. In the “Dirac-limit”, all of the HNL decays are LNC, while in the “Majorana-limit,” half of the decays are LNC and half of the decays are LNV. As discussed in section 2.4, the primary difference between LNC and LNV decays comes from spin correlations that can alter the kinematics of the leptons in the HNL decay. The results consider four different mixing scenarios (see section 2.2) in two HNL models: one HNL model with single-flavour mixing (1SFH) and two quasi-degenerate HNL model (2QDH). In the 2QDH model, the mass splitting between the pair of HNLs is smaller than the resolution of the detectors and, thus, the analysis is unable to distinguish between the two particles. See section 2.1 for more details about the motivation for the 2QDH model. Table 5.1 summarizes the eight models.

Table 5.1: A summary of the eight models used in the interpretation of the analysis results. Models are defined according to the number of HNLs (1SFH or 2QDH), the fraction of LNC and LNV decays and the relative strength of the mixing between the lepton generations. The inverted (IH) and normal (NH) hierarchy benchmarks are used for 2QDH models only, while the 1SFH models have single-flavour mixing.

Model	Fraction of LNV decays	Fraction of LNC decays	(x_e, x_μ, x_τ) $(x_\alpha = U_\alpha ^2/ U ^2)$
1SFH muon-only in the “Dirac-limit”	0	1	(0 , 1 , 0)
1SFH muon-only in the “Majorana-limit”	$\frac{1}{2}$	$\frac{1}{2}$	
1SFH electron-only in the “Dirac-limit”	0	1	(1 , 0 , 0)
1SFH electron-only in the “Majorana-limit”	$\frac{1}{2}$	$\frac{1}{2}$	
2QDH IH mixing in the “Dirac-limit”	0	1	$(\frac{1}{3}, \frac{1}{3}, \frac{1}{3})$
2QDH IH mixing in the “Majorana-limit”	$\frac{1}{2}$	$\frac{1}{2}$	
2QDH NH mixing in the “Dirac-limit”	0	1	(0.06, 0.48, 0.46)
2QDH NH mixing in the “Majorana-limit”	$\frac{1}{2}$	$\frac{1}{2}$	

The cross section calculation in section 2.5 assumes one HNL in the “Majorana-limit”. The cross sections for the other models in the “Dirac-limit” and those with more than one HNL (i.e. 2QDH model), can be computed as multiples of this Majorana cross section σ_M . Table 5.2 describes the relationships between the models. For an HNL in the “Majorana-limit”, both LNC and LNV decays are allowed and, thus, for in the “Dirac-limit”, with only LNC decays, there are half as many decay channels open. This increases the lifetime by a factor of two and thereby enhances the cross section by the same factor. For the 2QDH models, twice as many channels are opened because there are now two HNL particles that contribute to the process. Thus, the cross section in models with a pair of quasi-Dirac HNLs is twice as large as the one HNL model.

Table 5.2: The cross section for LNC and LNV decays for one Majorana HNL, one Dirac HNL, a quasi-Dirac HNL pair with “Majorana-like” decays and a quasi-Dirac HNL pair with “Dirac-like” decays. The total cross section is the sum of the LNC and LNV contributions.

	1SFH “Majorana-limit”	1SFH “Dirac-limit”	2QDH “Majorana-limit”	2QDH “Dirac-limit”
LNC decays	σ_M	$2\sigma_M$	$2\sigma_M$	$4\sigma_M$
LNV decays	σ_M	0	$2\sigma_M$	0
Total	$2\sigma_M$	$2\sigma_M$	$4\sigma_M$	$4\sigma_M$

5.4 Data and Signal Samples

5.4.1 Data Samples

This search was performed using 139 fb^{-1} of LHC proton-proton collision data collected by the ATLAS experiment between 2015 and 2018. A subset of these data events are selected and processed using the specialized displaced track reconstruction described in section 4.2.1. All data events are required to pass basic event cleaning, which includes only selecting events that are considered to be “good” for physics analysis (see discussion about data quality in section 3.1.3). Events are also required to have a reconstructed primary vertex with at least two associated tracks with $p_T > 500 \text{ MeV}$.

5.4.2 Signal Samples

Monte Carlo (MC) signal samples were generated using a PYTHIA8⁷ $pp \rightarrow W$ boson production model. W bosons (of both charges) are forced to decay into a lepton (μ or e) and a custom heavy neutral lepton ($W \rightarrow lN$). The HNL is given a long lifetime and set to decay leptonically ($N \rightarrow ll\nu$). The HNL decay is modeled as a weak decay mediated by either a virtual W or Z boson (see fig. 2.3). If the charged leptons in the HNL decay have the same flavour (e.g. $\alpha - \gamma\gamma$), then both diagrams contribute to the decay. The interference between the neutral and charged current decays is taken into account in the computation of branching ratio, $BR(N \rightarrow ll\nu)$, that is used to compute the HNL cross section in section 2.5.

Figure 5.3 shows the Feynman diagrams that contribute to models with single-flavour mixing. For each channel, events are generated with the following HNL masses: $m_N = 3, 4, 4.5, 5, 7.5, 10, 12.5, 15, 17.5, 20 \text{ GeV}$. Each mass point is generated with three different proper decay lengths: $c\tau_N = 1, 10, 100 \text{ mm}$.

⁷Signal samples were obtained using PYTHIA 8.212 [72] with the ATLAS A14 set of tuned parameters [73] and the NNPDF2.3LO PDF set [74].




Figure 5.3: Feynman diagrams included in the combinations for the muon-only mixing ($\mu-\mu\mu$, $\mu-\mu e$ and $\mu-ee$) and electron-only mixing ($e-e\mu$, $e-ee$ and $e-\mu\mu$) scenarios. Neutral current decays are also possible for the $\mu-\mu\mu$ and $e-ee$ channels.

The six channels that are generated for each mass and proper decay length are:

- $\mu-\mu\mu$: W boson decays to a muon and HNL ($W \rightarrow \mu N$) and the HNL decays to a muon pair and a neutrino ($N \rightarrow \mu\mu\nu$),
- $\mu-\mu e$: W boson decays to a muon and HNL ($W \rightarrow \mu N$) and the HNL decays to a muon-electron pair and a neutrino ($N \rightarrow \mu e\nu$),

- $\mu-ee$: W boson decays to a muon and HNL ($W \rightarrow \mu N$) and the HNL decays to an electron pair and a neutrino ($N \rightarrow ee\nu$),
- $e-ee$: W boson decays to an electron and HNL ($W \rightarrow eN$) and the HNL decays to an electron pair and a neutrino ($N \rightarrow ee\nu$),
- $e-e\mu$: W boson decays to an electron and HNL ($W \rightarrow eN$) and the HNL decays to an electron-muon pair and a neutrino ($N \rightarrow e\mu\nu$) and
- $e-\mu\mu$: W boson decays to an electron and HNL ($W \rightarrow eN$) and the HNL decays to a muon pair and a neutrino ($N \rightarrow \mu\mu\nu$).

Depending on the relative size of the mixing angles ($|U_e|^2$, $|U_\mu|^2$ and $|U_\tau|^2$) in the HNL model that is considered, the decay channels will have various contributions to the overall signal strength. For models where the HNL mixes with both muon and electron flavoured neutrinos (e.g. IH and NH benchmarks), two additional channels contribute. These channels are:

- $\mu-e\mu$: W boson decays to a muon and HNL ($W \rightarrow \mu N$) and the HNL decays to an electron-muon pair and a neutrino ($N \rightarrow e\mu\nu$), and
- $e-\mu e$: W boson decays to an electron and HNL ($W \rightarrow eN$) and the HNL decays to a muon-electron pair and a neutrino ($N \rightarrow \mu e\nu$).

Feynman diagrams for these additional two channels, as well as the charged-current decays for the $\mu-ee$ and $e-\mu\mu$ channels that are also possible in a mixed-flavour mixing scenario, are shown in fig. 5.4. Experimentally, the leptons in the displaced vertex are indistinguishable and, therefore, in models with both nonzero muon and electron mixing angles, the $\mu-\mu e$ ($e-e\mu$) and $\mu-e\mu$ ($e-\mu e$) signals are combined into one channel. These processes only differ due to the kinematics differences from spin correlations and the relative strength of the mixing with the SM neutrinos $x_\alpha = |U_\alpha|^2/|U|^2$. Thus, the event weighting procedure, described in section 5.4.3, is used to compute separate weights for the $\mu-\mu e$ and $\mu-e\mu$ processes. A combined weight that is equal to the sum of the individual weights is used for models where both the $\mu-\mu e$ and $\mu-e\mu$ processes are allowed. The same method is used for the $e-e\mu$ samples to get a combined weight for $e-e\mu$ and $e-\mu e$ processes.

The analysis is not sensitive to HNL mixing with tau neutrinos, $|U_\tau|^2$. Tau reconstruction is typically more difficult than electron or muon reconstruction due to the short proper decay length of the tau lepton ($c\tau_\tau \sim 0.1$ mm) and the complicated nature of the decay products. The sensitivity to $|U_\tau|^2$ through the leptonic tau decays to muons and electrons was studied using $N \rightarrow \tau\tau\nu$ samples. These samples were run through the analysis using the selections from all six channels. In all cases, less than one signal event was observed for tau decays after the

final analysis selections were applied. Thus, in models where $|U_\tau|^2$ is nonzero, all tau channels are assumed to have zero acceptance.




Figure 5.4: Additional Feynman diagrams for signal processes that contribute to the signal strength if both $|U_\mu|^2$ and $|U_e|^2$ are nonzero. Diagrams are shown for the charged-current interaction only.

5.4.3 Modelling Spin Correlations

Spin correlations (described in section 2.4) between the prompt lepton and the HNL are not included in the PYTHIA8 signal generation. However, these correlations are taken into account in MADGRAPH using the HEAVYN model [75, 76]. The impact of spin correlations in the modelling of HNL decays was only brought to light after all of the PYTHIA8 samples were produced. Due to the significant amount of time and computing resources required to validate and re-simulate all of the signal samples with MADGRAPH, an event weighting method is used to account for the small kinematic differences between LNC and LNV decays.

MC simulations model the probability of different outcomes by computing matrix elements for the processes. To take spin correlations into account, the MC generator should sample from

a distribution that is proportional to:

$$f(x) = |M_{\text{production} + \text{decay}}|^2 \quad (5.1)$$

where x denotes the set of all relevant kinematic variables that the matrix element, M , depends on. Instead, PYTHIA8 factorizes the production and decay matrix elements and samples from a distribution that is proportional to:

$$g(x) = |M_{\text{production}}|^2 \times |M_{\text{decay}}|^2 \quad (5.2)$$

To correct for this sampling effect, a weight that is equal to $w(x) = f(x)/g(x)$ is applied to each simulated event. This event weight can be computed analytically. The matrix elements in $f(x)$ and $g(x)$ depend on the mass of the HNL, the 4-momentum of the W boson and the 4-momentum of the prompt and displaced leptons. To simplify the computation of the matrix elements, light lepton (e/μ) masses are neglected and the matrix elements are reported up to a constant.

Six matrix elements are computed that include spin correlations. Three matrix elements are computed for LNC decays that correspond to the charge current (CC), neutral current (NC) or charge plus neutral current (CCNC) contributions. Similar matrix elements are also computed for LNV decays. The size of the contribution from each of the CC, NC and CCNC decays depends on the process and which mixing angles are nonzero. Thus, the computation of the weight, $w(x)$, depends on the lepton number conserving behaviour of the decay (i.e. LNC or LNV), the flavours of the leptons in the HNL decay and the relative strengths of the mixing angles, $x_\alpha = |U_\alpha|^2/|U|^2$.

For $e-\gamma\gamma$ and $\mu-\gamma\gamma$ channels, with $\gamma\gamma = ee$ or $\mu\mu$, the following event weight is computed:

$$w_{CC+NC} = \frac{x_\gamma |M_{CCNC}|^2 + (1 - x_\gamma) |M_{NC}|^2}{|M_{\text{PYTHIA}}|^2} \quad (5.3)$$

where $|M_{CCNC}|^2$ is the matrix element for CC plus NC decays proportional to $|U_\gamma|^2$, $|M_{NC}|^2$ is the contribution from NC decays proportional to $|U_{\alpha \neq \gamma}|^2$, $|M_{\text{PYTHIA}}|^2$ is the matrix element generated by PYTHIA8 and x_γ is the relative strength of the mixing angle of the lepton flavour in the decay.⁸ This event weight is used for $\mu-\mu\mu$, $\mu-ee$, $e-ee$ and $e-\mu\mu$ channels.

For processes with different flavours of leptons in the HNL decay (i.e. $\beta \neq \gamma$), only the

⁸For example, in a $\mu-ee$ decay in a model where the mixing to all three lepton generations is nonzero, the $|M_{CC+NC}|^2$ term would take into account all the diagrams (CC and NC) that are proportional to $|U_e|^2$ and the $|M_{NC}|^2$ term would take into account NC-only decays where a muon or tau neutrino couples to the HNL (i.e. $N \rightarrow \nu_{l \neq e} Z^*$; $Z^* \rightarrow e^+ e^-$).

charge current decay contributes and the weight is:

$$w_{CC} = \frac{|M_{CC}|^2}{|M_{PYTHIA}|^2} \quad (5.4)$$

where $|M_{CC}|^2$ is the matrix element for the charged current decays. This event weight is used for $\mu-\mu e$ and $e-e\mu$ channels.

The validation of this method is performed by comparing samples generated with PYTHIA8 and MADGRAPH⁹ Comparisons of the lepton p_T distributions between MADGRAPH samples, the weighted and unweighted PYTHIA8 samples are shown in fig. 5.5 for the $\mu-\mu\mu$ channel with CC and NC decays.




Figure 5.5: Comparison of the lepton p_T distributions for LNC decays between MADGRAPH samples (purple) and the weighted (pink) and unweighted (gray) PYTHIA8 samples for a $\mu-\mu\mu$ signal using the NC+CC event weights. Samples are generated with $m_N = 10$ GeV and $c\tau_N = 10$ mm signal sample.

⁹MADGRAPH signal samples were obtained with MADGRAPH5_AMC@NLO v2.9.3 [77] using the HEAVYN model [75, 76]. The PYTHIA8 version description can be found in section 5.4.2.

Figure 5.6 shows the p_T comparisons for the $\mu-\mu e$ channel with CC decays only. Good agreement between MADGRAPH and the weighted PYTHIA8 samples indicates that the reweighting method can be used to successfully include spin correlations in the model and, therefore, no additional signal uncertainty is necessary to account for this effect.




Figure 5.6: Comparison of the lepton p_T distributions for LNC decays between MADGRAPH samples (blue) and the weighted (orange) and unweighted (gray) PYTHIA8 samples for a $\mu-\mu e$ signal using the CC event weights. Samples are generated with $m_N = 10$ GeV and $c\tau_N = 10$ mm signal sample.

5.4.4 Signal Kinematics

Kinematic distributions of simulated particles involved in the HNL production and decay are shown for a $\mu-\mu\mu$ signal in figs. 5.7 to 5.9. In simulation, the kinematic differences between the different channels are negligible because leptons with momenta on the order of a GeV can be treated as equivalent (massless) objects. Three different mass variables are used in this analysis. The DV mass, m_{DV} , is the invariant mass of the two displaced charged leptons. The

DV mass does not include the neutrino momentum from the HNL decay and thus, m_{DV} will always underestimate the simulated HNL mass, m_N . The tri-lepton mass, m_{lll} , is computed using the kinematics of the three charged leptons in the signal. It is also missing the neutrino momentum and is always less than the W boson mass. The HNL mass in fig. 5.8 (a) is calculated using the HNL flight direction and the momenta of the other charge leptons in the HNL decay to constrain two out of three degrees of freedom in the neutrino momentum. The final degree of freedom can be solved for using energy-momentum conservation of the W boson decay. Details about this HNL mass calculation can be found in section 5.5.3.




Figure 5.7: Simulated lepton p_T distributions for the three charged leptons in the HNL event and the invariant mass of the tri-lepton system (m_{lll}). The leading displaced lepton is defined as the lepton in the DV with the largest p_T , while the subleading displaced lepton has the smaller p_T . Simulated $\mu-\mu\mu$ signals are plotted for three different mass points: $m_N = 5, 7.5, 10$ GeV with $c\tau_N = 10$ mm.




Figure 5.8: HNL kinematics for simulated $\mu\text{--}\mu\mu$ signals for three different mass points: $m_N = 5, 7.5, 10$ GeV with $c\tau_N = 10$ mm.




Figure 5.9: The DV mass (a) and DV radius (b) distributions for simulated $\mu\text{--}\mu\mu$ signal samples. For the DV mass distribution, three different mass points are shown: $m_N = 5, 7.5, 10$ GeV with $c\tau_N = 10$ mm. For the DV radius plot, the simulated samples have $m_N = 10$ GeV with $c\tau_N = 1, 10$ and 100 mm.

5.5 Event Reconstruction

The experimental displaced HNL signature consists of two main building blocks: a prompt lepton and a two-lepton DV. The prompt lepton is reconstructed using the standard electron and muon reconstruction algorithms described in section 4.1.3 and section 4.1.4 respectively and the two-lepton DV is reconstructed using the VSI algorithm described in section 4.2.3. However, since the default VSI configuration was not optimized to reconstruct two-lepton DVs, a custom VSI configuration was developed for this analysis. The details of the custom VSI configuration can be found in section 5.5.2.

In addition to standard kinematic variables (e.g. p_T , η , ϕ), this analysis also calculates the HNL mass, m_{HNL} , using the kinematics of the three charged leptons in the signal, as well as information about the flight direction of the HNL and the W boson mass. The details of this mass variable can be found in section 5.5.3.

5.5.1 Displaced Vertex Reconstruction for the HNL Signal

The displaced vertex reconstruction, described in section 4.2.3, produces two collections of vertices by default. These collections are created by modifying the VSI algorithm’s configuration settings to apply different cuts and control which steps in the reconstruction sequence are run. The collection used by most ATLAS displaced vertex analyses, simply referred to as “VSI”, initially selects tracks from the collection of ID tracks (both standard and large-radius). In order to reduce the number track pairs, a 2 mm two-track forming d_0 cut is used. This cut requires that at least one track has a $d_0 > 2$ mm in order to be defined as a compatible two-track pair. This configuration also runs the final track attachment step with non-selected tracks.

The second collection of vertices, called “VSI Leptons”, only selects lepton-matched tracks. The details of how tracks are matched to electron and muon objects can be found in section 4.1.3 and section 4.1.4 respectively. Since the number of track pair combinations is now reduced by selecting lepton-matched tracks, this configuration has no requirements on the d_0 of the tracks. Since the purpose of this configuration is to reconstruct vertices with displaced leptons, the final track attachment step with non-selected tracks is not run, since these tracks consist of mostly non-lepton tracks.

The signal efficiency using VSI and VSI Leptons vertices is shown in fig. 5.10 (a) for simulated HNL signals with varying lifetimes and masses. This efficiency is defined as the ratio of the selected reconstructed events to the total number of simulated events. Selecting events using VSI Leptons vertices has a larger signal acceptance for all masses and lifetimes when compared to VSI. Figure 5.10 (b) shows the signal efficiency using VSI Leptons vertices is up to 230% of the VSI signal efficiency for $c\tau_N = 1$ mm signals.




Figure 5.10: The $\mu\text{--}\mu\mu$ signal efficiency (a) and ratio of the signal efficiency between VSI Leptons and VSI (b) as a function of simulated HNL mass with $3 < m_N < 20$ GeV. Signal samples with $c\tau_N = 1, 10, 100$ mm are shown in red, blue and green respectively. The signal efficiency is defined as the ratio of the selected reconstructed events to the total number of simulated events.

There are two main differences between the VSI and VSI Leptons configurations that result in this increase in signal efficiency. First, the VSI Leptons configuration has no requirement of the d_0 of the tracks when forming compatible two-track pairs. This is especially important for signals that are not very displaced, since the tracks from the HNL decay will have a small d_0 . As a result, the observed increase in efficiency is largest for the 1 mm lifetime signals (see the red points in fig. 5.10 (b)). The second effect comes from the non-selected track attachment step that is only run in the VSI configuration. In this configuration, reconstructed signal vertices may have additional tracks attached to the vertex either from fake tracks or pileup tracks near the HNL decay. Since the signal selection requires that the displaced vertex has exactly two tracks, any signal vertex with additional attached tracks is rejected by this two-track requirement.

While the signal acceptance is larger for events reconstructed with the VSI Leptons configuration, the sensitivity of the analysis depends on both the signal and background acceptance. In fig. 5.11, the signal and data-driven same-sign background estimate is shown as a function of the reconstructed HNL mass.¹⁰ The lepton-only vertexing configuration improves the signal efficiency by 20% for this $\mu\text{--}\mu\mu$ signal model and up to 300% for other benchmark signals. However, this signal efficiency increase comes at the cost of over a factor of 20 increase in the two-muon DV data-driven background.

¹⁰The details of the reconstructed HNL mass variable can be found in section 5.5.3 and a description of the data-driven same-sign background estimate is described in section 5.8.




Figure 5.11: Simulated events for a $\mu\mu$, 10 GeV, 10 mm signal in purple and the data-driven same-sign background in grey as a function of the reconstructed HNL mass. The total number of events in the signal and background can be found in the legends in each plot. Events are selected by applying the pre-selection cuts described in section 5.7.

5.5.2 Custom Displaced Vertex Configuration

With the default lepton-only vertexing configuration, the analysis could not benefit from the increase in signal acceptance due to the large increase in the background. Therefore, a custom vertex configuration was developed to maintain most of the signal efficiency increase and greatly reduce the background. The three main areas in the VSI algorithm that were optimized for the two-lepton displaced HNL signal are the track selection, two-track vertex forming d_0 cut and track attachment. These three optimizations are discussed below.

The custom vertex configuration that is used to reconstruct vertices for this analysis is based on the VSI Leptons configuration with additional track selections applied to both standard and LRT tracks and a two-track forming d_0 cut of 1 mm. This custom configuration also runs the final track attachment step.

Track Selection Optimization

It was identified that HNL signal events can be lost due to fake tracks being attached or merged with the signal vertices. Identifying and removing these fake tracks prior to vertexing improves the signal acceptance. As will be discussed in section 5.8, the dominant background comes from random track crossings that form a two-lepton vertex. Removing non-signal-like tracks before vertexing reduces the probability of a random crossing and therefore also reduces the background.

At the same time as the vertex optimization studies were being performed, large-radius track reconstruction was also undergoing optimization in preparation for the start of the LHC

Run 3.¹¹ As a result, new LRT quality selections had been developed specifically to reduce the number of fake tracks. The additional selections on the track parameters that were added to the custom VSI configuration were inspired by these Run 3 LRT improvements and are listed in table 5.3.

Table 5.3: The main differences in the track selection between the default VSI configuration and the lepton-only selection used in the custom VSI configuration in this analysis.

Track Type	Track Parameter	Default Value	Custom Value
Standard & LRT	Lepton matching	Not required	Required
	Min. # Silicon hits	7	8
LRT	Max. $ z_0 $	1000 mm	500 mm
	Min. # SCT hits	2	7
	Max. # SCT holes	2	1
	Max. χ^2/dof	50	9
	$\Delta\theta$ and $ z_0 $ correlation	n.a.	$\Delta\theta \leq 1$ or $ z_0 < 200$ mm

The impact of the additional track selections on the signal and background can be seen in fig. 5.12. Applying additional selections removes fake or poor quality tracks and reduces the number of two-lepton VSI background events by a factor of 2 and the number of VSI Leptons background events by 20%. This reduction in the background is achieved while maintaining a similar number of signal events. To measure the performance of the additional track selections, the signal significance was estimated as the number of signal events divided by the square root of the number of background events ($s/\sqrt{b}$). Adding the additional track selections increases the signal significance, indicating that these selections improve the analysis sensitivity.

Two-Track Vertex Forming d_0 Cut Optimization

As described in section 5.5.1, one of the main differences between the VSI and VSI Leptons configurations is the removal of the two-track forming d_0 cut that is used to determine the compatibility of the two tracks. Studying simulated HNL signal events shows that about 80% (20%) of events with 1 mm (10 mm) lifetimes have two tracks with $d_0 < 2$ mm. This means that 80% (20%) of the time, signals with 1 mm (10 mm) lifetimes will fail to pass the two-track compatibility threshold in the VSI configuration.

Signal and background events were studied for different two-track forming d_0 cuts ranging from 0 mm to 2 mm. It was observed that changing the d_0 cut from 2 mm to 0 mm can increase the signal efficiency by a factor of 1.2 to 3 (see fig. 5.13 (a)). However, the number of background events also increases as the d_0 cut decreases (see fig. 5.13 (b)). To find one configuration for all signal models, a balance must be found when optimizing the vertexing configuration for signals with different lifetimes. Using the signal significance ($s/\sqrt{b}$) as a

¹¹Run 3 refers to the period of data collection that is planned to start in 2022. See fig. 3.3 for a timeline of the LHC machine.

metric (fig. 5.13 (c) and (d)), it was determined that a 1 mm d_0 cut gave the best balance between increasing the sensitivity to low lifetime signals, while minimally impacting longer lifetime samples.




Figure 5.12: The number of $\mu\text{--}\mu\mu$ 10 GeV, 10 mm signal (a) and two-muon background (b) events for default VSI (blue) and VSI Leptons (orange) vertex configurations. The original number of events is given by the default configuration. Events labels “LRTR3” are the number of events after modifying the track selections as described in table 5.3 without the $\Delta\theta$ and $|z_0|$ correlation selection. The “LRTR3 w Geo cut”, which is the final selection applied in the custom vertex configuration, is the number of events including this final correlation selection. The signal significance (c) is estimated using the number of signal events divided by the square root of the number of background events ($s/\sqrt{b}$).




Figure 5.13: Signal events for a 10 GeV, 10 mm $\mu\text{-}\mu\mu$ sample (a) and two-muon background events (b) as a function of the two-track forming d_0 cut for VSI (blue) and VSI Leptons (orange) vertex configurations. The original number of events is given by the default configuration. Then the number of events with two-track forming d_0 cut threshold equal to 2, 1.5, 1, 0.5 and 0 mm is shown. The signal significance ($s/\sqrt{b}$) for 10 GeV, $\mu\text{-}\mu\mu$ signal is shown for the 10 mm (c) and 1 mm (d) signal samples as a function of the same d_0 cut.

Track Attachment

The plot of the number of background events as a function of the d_0 cut (fig. 5.13 (b)) shows that even when the d_0 cut is the same, the VSI Leptons background is still about four times larger compared to VSI. However, as described in section 5.5.1, the final difference between the configurations is that VSI runs the final track attachment step and attaches non-selected tracks to the vertex. This means that non-lepton tracks that are compatible with a two-lepton pair will be added to the vertex and increase the number of tracks (n_{tracks}) in the vertex such that $n_{\text{tracks}} > 2$.

To study the impact of the track attachment step, DVs with exactly two lepton-matched tracks are selected. These events are then separated according to the total number of tracks in the vertex. Vertices that consists of two leptons (i.e. $n_{\text{tracks}} = 2$) are selected as signal (or signal-like) events. Any vertex with $n_{\text{tracks}} > 2$ is rejected from the signal region. In fig. 5.14 (a), about four signal events are removed from the signal region because additional tracks are attached to the two-track signal vertex. However, in the background distribution in fig. 5.14 (b), there are many vertices with $n_{\text{tracks}} > 2$ that have exactly two-leptons. Since the VSI Leptons configuration does not run this additional track attachment step, many of these vertices with $n_{\text{tracks}} > 2$ will be reconstructed with exactly two lepton-matched tracks and be selected in the signal region.

Without the additional track attachment step, information about the non-lepton tracks near by the background vertices is missing. This missing information is what causes the factor of four increase in background events that is observed in fig. 5.13 (b) when the two configurations are run with the same d_0 cut. Running the track attachment in the VSI algorithm improves the signal significance, since the reduction in background is much larger than the small amount of signal events that are lost.




Figure 5.14: Signal (a) and background (b) events with exactly two muons in the displaced vertex as a function of the number of tracks (n_{tracks}) in the vertex reconstructed using the VSI configuration with a 1 mm two-track forming d_0 cut. Vertices with $n_{\text{tracks}} = 2$ are selected in the signal region, while events with $n_{\text{tracks}} > 2$ are rejected.

Performance of the Custom Displaced Vertex Configuration

A summary plot of the signal efficiency increase observed with the custom vertex configuration is shown in fig. 5.15. This plot shows the ratio of the the number of signal events for the custom to the default VSI configuration is shown for the four single-flavour mixing channels ($\mu-\mu\mu$, $\mu-\mu e$, $e-e\mu$ and $e-ee$) as a function of lifetime. For most signal channels and lifetimes, the custom configuration has larger signal yields, especially for 1 mm samples, where the $e-ee$ channel has a factor of 3.8 increase in acceptance. Longer lifetime samples are minimally impacted by the additional selections applied to LRT tracks.




Figure 5.15: Ratio of the number of selected signal events reconstructed with the custom configuration to the default VSI configuration. Ratios are shown for the 10 GeV signal samples in the $\mu-\mu\mu$, $\mu-\mu e$, $e-e\mu$ and $e-ee$ channels as a function of the proper decay length.

To evaluate the vertex reconstruction efficiency for the signal, a matching procedure is used to determine if the simulated HNL decays are reconstructed as displaced vertices. This matching procedure follows the same definitions adopted in Ref. [78]. HNL decays are considered to be reconstructed if the tracks associated to a displaced vertex are matched to particles originating from the HNL. The matching of tracks to particles is based on a weighted scoring of shared detector hits between the track and particle trajectories. A track is considered matched to a particle if this score is greater than 0.5. For each pair of a reconstructed vertex

position, v , and the simulated decay vertex position, l , a matching score s is computed, which uses the magnitude of the track p_T as a weight. This score is given by:

$$s(v, l) \equiv \frac{\sum_{i \in \text{tracks in } l} (p_T^{(i)} | \text{descendent of signal decay } l)}{\sum_{i \in \text{tracks in } v} p_T^{(i)}} \quad (5.5)$$

The total vertex efficiency can be broken up into several components including the acceptance, the track selection efficiency and the core efficiency of the VSI algorithm. The acceptance, seed and core efficiencies are defined as follows,


- **Acceptance ($\mathcal{A}$):** fraction of signal decays with at least two reconstructed tracks with $p_T > 1$ GeV, a z -position of $|z| < 2720$ mm and a transverse distance from the origin of $L_{xy} < 563$ mm.
- **Seed efficiency (ϵ^{seed}):** fraction of signal decays with at least two tracks passing the VSI track selection requirements.
- **Core efficiency (ϵ^{core}):** fraction of signal decays with a reconstructed vertex that has a matching score $s(v, l) > 0.5$.

The total efficiency (ϵ^{tot}) is then given by:


$$\epsilon^{\text{tot}} = \mathcal{A} \cdot \epsilon^{\text{seed}} \cdot \epsilon^{\text{core}} \quad (5.6)$$

These four efficiencies ($\mathcal{A}$, ϵ^{seed} , ϵ^{core} and ϵ^{tot}) are shown in fig. 5.16 as a function of radial decay position, L_{xy} , for 10 GeV, 10 mm signal for the $\mu-\mu\mu$, $\mu-\mu e$ and $e-ee$ channels. The total efficiency peaks between 15-25% depending on the signal sample. The efficiency at very low values of L_{xy} is suppressed due to the minimum $|d_0|$ of 1 mm and at high values of L_{xy} due to the requirements on the number of hits per track.


Overall the total vertexing efficiency of the HNL signal is low compared to other long-lived DV signals because the HNL signature only has two tracks in the DV. If either of the tracks is not reconstructed or the two reconstructed tracks do not form a good quality vertex, then the signal is rejected. In contrast other long-lived displaced vertex signals can have many tracks and even if one track is missing, the signal vertex can still be reconstructed without this track.




(a) Acceptance, $\mathcal{A}$



(b) Seed efficiency, ϵ^{seed}



(c) Core efficiency, ϵ^{core}



(d) Total efficiency, ϵ^{tot}

Figure 5.16: The three components of the vertex reconstruction efficiency and the total efficiency, measured as a function of radial decay position, L_{xy} , for $\mu\text{--}\mu\mu$, $\mu\text{--}ee$ and $e\text{--}ee$ signals with $m_N=10$ GeV and $c\tau_N=10$ mm.

5.5.3 Reconstructed HNL Mass

The HNL mass (m_{HNL}) can be obtained using energy-momentum conservation in the HNL production ($W \rightarrow N l_1$) and decay ($N \rightarrow l_2 l_3 \nu$), where l_1 is the prompt lepton and l_2 and l_3 are the charged leptons in the DV. The problem can be summarized with the following equations. Four-momentum conservation in the N decay gives:

$$p_N^\mu = p_2^\mu + p_3^\mu + p_\nu^\mu \equiv p_{23}^\mu + p_\nu^\mu \quad (5.7)$$

Four-momentum conservation in the W decay gives:

$$p_W^\mu = p_1^\mu + p_N^\mu = p_1^\mu + p_{23}^\mu + p_\nu^\mu \quad (5.8)$$

It will also be useful to define:

$$p_{23}^2 = E_{23}^2 - |\vec{p}_{23}|^2 = m_{23}^2 \quad (5.9)$$

$$p_{23}^\parallel \equiv \vec{p}_{23} \cdot \hat{v} \quad (5.10)$$

$$p_{23}^\perp \equiv |\vec{p}_{23} - p_{23}^\parallel \hat{v}| \quad (5.11)$$

where m , E , and $|\vec{p}|$ are the mass, energy, and momentum-vector magnitude of the particles indicated by their subscript and $\hat{v}$ is the flight direction of the HNL given by the vector connecting the PV and DV.

The components of the neutrino momentum are not measured. Equivalently, due to eqs. (5.7) and (5.8), we could treat the components of p_N or p_W (respectively) as the unknowns. Either way, the HNL flight direction provides a constraint on the neutrino momentum components orthogonal to $\hat{v}$, leaving two remaining degrees of freedom. The simplest solution is to solve the problem using the neutrino degrees of freedom (E_ν and $p_\nu^\parallel$) and computing the various quantities in a reference frame that is rotated with respect to the lab frame.

The solution to the HNL mass is presented in the coordinate system, $k = (\hat{x}', \hat{y}', \hat{z}')$, which is rotated with respect to the ATLAS coordinate system (see section 3.2.1), such that the origin of the k -frame is at the PV and the z' -axis points along the flight direction of the HNL. The definition of this coordinate system is:

$$\hat{z}' = \hat{v}, \quad \hat{x}' = \frac{\vec{p}_{23} \times \hat{z}'}{|\vec{p}_{23} \times \hat{z}'|}, \quad \hat{y}' = \hat{z}' \times \hat{x}' \quad (5.12)$$

Energy-momentum conservation in the W and N decays can be expressed in terms of $p_\nu \cdot \hat{v} = \alpha$

using:

$$\vec{p}'_{23} \equiv \vec{q} \quad (5.13)$$

$$E'_{23} = \sqrt{q^2 + m_{23}^2} \quad (5.14)$$

$$\vec{q} = (0, |\vec{p}_{23} \times \hat{z}'| \equiv q_{\perp}, \vec{p}_{23} \cdot \hat{z}' \equiv q_z) \quad (5.15)$$

$$\vec{p}_{\nu} = (0, -q_{\perp}, \alpha) \quad (5.16)$$

$$E_{\nu} = \sqrt{q_{\perp}^2 + \alpha^2} \quad (5.17)$$

Equation (5.8) gives:

$$m_W^2 = m_1^2 + m_{\nu}^2 + m_{23}^2 + 2p'_1 \cdot (p'_{23} + p'_{\nu}) + 2p'_{23} \cdot p'_{\nu} \quad (5.18)$$

where:

$$p'_1 \cdot (p'_{23} + p'_{\nu}) = E'_1(E'_{23} + E'_{\nu}) - p'_{1,z}(q_z + \alpha) \quad (5.19)$$

$$p'_{23} \cdot p'_{\nu} = E'_{23}E'_{\nu} + q_{\perp}^2 - q_z\alpha \quad (5.20)$$

In the energy regime of interest, with particles that have GeV momenta, both the mass of the prompt lepton (m_1) and the neutrino (m_{ν}) can be set to zero. Rearranging eq. (5.18) to solve for E_{ν} gives:

$$E'_{\nu} = A + B\alpha \quad (5.21)$$

where

$$A = \frac{\frac{m_W^2 - m_{23}^2}{2} - E'_1 E'_{23} + p'_{1,z} q_z - q_{\perp}^2}{E'_1 + E'_{23}}, \quad B = \frac{p'_{1,z} + q_z}{E'_1 + E'_{23}} \quad (5.22)$$

Subtracting eq. (5.21) from eq. (5.17) gives the following quadratic expression α :

$$(B^2 - 1)\alpha^2 + 2AB\alpha + A^2 - q_{\perp}^2 = 0 \quad (5.23)$$

Using the quadratic formula, the solution for α is:

$$\alpha = \frac{-AB \pm \sqrt{(B^2 - 1)q_{\perp}^2 + A^2}}{(B^2 - 1)} \quad (5.24)$$

This gives two solutions for the z component of the neutrino momentum, α . Both solutions were studied using simulated HNL events and it was observed that the solution with the smaller $|p_N^{\vec{z}}|$ typically led to a value for m_{HNL} that was closer to the simulated HNL mass

m_N . This solution often corresponds to the forward emission of the neutrino with respect to the HNL decay in the lab frame. Therefore, in the definition of the HNL mass used in the analysis, the positive radical is chosen as the solution for α .

From eq. (5.7) and the definitions in eqs. (5.13) to (5.17), the expression for the HNL mass in terms of α is:

$$m_{\text{HNL}}^2 = m_{23}^2 + 2p'_\nu \cdot p'_{23} \quad (5.25)$$

$$= m_{23}^2 + 2E'_{23} \sqrt{q_\perp^2 + \alpha^2} + 2q_\perp^2 - 2q_z \alpha \quad (5.26)$$

Substituting the expression for α in eq. (5.24) into eq. (5.26) gives the solution for the HNL mass.

Picking the W mass

Equation (5.26) gives a solution for the HNL mass in terms of the known energies and momenta of the two charged particles in the HNL decay (E'_{23} , $\vec{p}'_{23} \equiv \vec{q}$) and the component of the neutrino momentum in the z' direction, α . Equation (5.24) provides an analytical solution for α , given the momenta of all three charged leptons and the mass of the W boson, m_W .

ATLAS has measured the W boson pole mass to be $M_W = 80.370 \pm 0.019$ GeV [79] and this measurement is combined in Ref. [80] with results from LEP to provide a measurement for the W boson width, $\Gamma_W = 2.195 \pm 0.083$ GeV. Since the W mass has a width, if m_W is fixed to the pole mass in eq. (5.18), then it is possible m_{HNL} has no real solution. Instead of rejecting these events, m_W is set equal to the median m_W in the kinematically allowed region defined by the charged lepton kinematics. This ensures that α (and correspondingly m_{HNL}) will always have a real solution. The difference in the HNL mass distribution between fixing the m_W to the pole mass and using the median m_W in the kinematically allowed region can be seen in fig. 5.17. Fixing the W mass results in events where the HNL mass has no real solution. These are represented in fig. 5.17 by the entries at $m_{\text{HNL}} = -1$.

To define the kinematically allowed region, the minimum W mass that is consistent with the charged lepton decay products ($m_{W,\text{min}}$) is computed. From eq. (5.18), the mass of the W boson is given by:

$$m_W^2 = m_{23}^2 + 2(E'_1 E'_{23} + E'_\nu(E'_1 + E'_{23}) - p'_{1,z} q_z + q_\perp^2 - \alpha(p'_{1,z} + q_z)) \quad (5.27)$$

and $m_{W,\text{min}}$ occurs where:

$$\frac{d m_W^2}{d\alpha} = (E'_1 + E'_{23}) \frac{dE'_\nu}{d\alpha} - (p_{1,z} + q_z) = 0 \quad (5.28)$$




Figure 5.17: Two HNL mass distributions for a simulated $\mu\text{--}\mu\mu$ with $m_N = 10$ GeV and $c\tau_N = 10$ mm. The m_{HNL} is computed using m_W fixed to the pole mass of 80.4 GeV (blue) and the median W mass in the allowed kinematic region (purple). Events where the HNL mass has no real solution are set to -1.

Using $\frac{dE_\nu}{d\alpha} = \frac{d\sqrt{q_\perp^2 + \alpha^2}}{d\alpha} = \frac{\alpha}{E_\nu}$ in eq. (5.28), the value of α that gives the minimum m_W is:

$$\frac{d\frac{m_W^2}{2}}{d\alpha} = (E'_1 + E'_{23}) \frac{\alpha}{E_\nu} - (p'_{1,z} + q_z) = 0 \quad (5.29)$$

$$\frac{\alpha}{E_\nu} = \frac{p'_{1,z} + q_z}{E'_1 + E'_{23}} = B \quad (\text{same } B \text{ as in eq. (5.23)}) \quad (5.30)$$

$$\frac{\alpha^2}{E_\nu'^2} = B^2 \quad (5.31)$$

$$\alpha^2 = B^2 E_\nu'^2 = B^2 (q_\perp^2 + \alpha^2) \quad (5.32)$$

$$\alpha = \pm \frac{q_\perp B}{\sqrt{1 - B^2}} \quad (5.33)$$

Substituting $\alpha = \frac{q_\perp B}{\sqrt{1 - B^2}}$ back into eq. (5.27), the minimum W boson mass is:

$$m_{W,\min}^2 = m_{23}^2 + 2 \left(E'_1 E'_{23} + (E'_1 + E'_{23}) \sqrt{q_\perp^2 + \frac{q_\perp^2 B^2}{1 - B^2}} - p'_{1,z} q_z + q_\perp^2 - (p'_{1,z} + q_z) \frac{q_\perp B}{\sqrt{1 - B^2}} \right) \quad (5.34)$$

The cumulative probability for the W boson to have a mass greater than $m_{W,\min}$ is used to find the median of the remaining distribution. The probability density function (f) for m_W is proportional to:

$$f(m_W^2) \propto \frac{1}{(m_W^2 - M_W^2)^2 + M_W^2 \Gamma_W^2} \quad (5.35)$$

Therefore, the cumulative distribution function (F) is

$$F(m_W^2) = \frac{1}{\pi} \arctan \left(\frac{m_W^2 - M_W^2}{M_W \Gamma_W} \right) + \frac{1}{2} \quad (5.36)$$

The midpoint of the allowed kinematic region has a value of

$$F_{\text{med}} = \frac{1 + F(m_{W,\text{min}}^2)}{2} \quad (5.37)$$

Rearranging eq. (5.36) for m_W^2 gives

$$m_W^2 = M_W^2 + \Gamma_W M_W \tan \left(\pi \left[F - \frac{1}{2} \right] \right) \quad (5.38)$$

Substituting $F = F_{\text{med}}$, in eq. (5.38) gives an expression for the median W mass in the kinematically allowed region

$$m_{W,\text{med}}^2 = M_W^2 + \Gamma_W M_W \tan \left(\pi \left[\frac{1 + F(m_{W,\text{min}}^2)}{2} - \frac{1}{2} \right] \right) \quad (5.39)$$

This value of $m_{W,\text{med}}$ is used in eq. (5.24) to solve for α .

5.6 Object Definitions

Events in this search are selected using three types of physics objects: displaced vertices, muons and electrons. The definitions for each object are as follows:

Displaced vertices

Displaced vertices are reconstructed from standard and large-radius tracks following the procedure described in section 4.2.3. A custom lepton-only displaced vertex configuration, described in section 5.5.2, is used to perform the reconstruction. At the analysis level, only displaced vertices with exactly two tracks matched to displaced leptons (e/μ) are selected.

Prompt muons

Prompt muons are reconstructed as described in section 4.1.4. The prompt muon in the event is defined as the highest p_T muon with $|d_0| < 3$ mm and $|z_0 \sin(\theta)| < 0.5$ mm that does not overlap with any tracks in the displaced vertex. This muon is required to be a combined muon and pass the *Medium* quality working point. The muon is also required to have a $p_T > 3$ GeV. If the selected prompt muon overlaps with a prompt electron within $\Delta R = \sqrt{\Delta\eta^2 + \Delta\phi^2} < 0.05$, then the event is rejected.

Displaced muons

Displaced muons are tracks in a displaced vertex that are matched to muon objects. The matching of tracks to muon objects is described in section 4.1.4. A displaced muon must be matched to a combined muon that passes the *Medium* quality working point. The muon is also required to have a $p_T > 3$ GeV. If a displaced muon overlaps with an electron, and the matched muon passes at least the *Loose* quality working point, then the track is defined as a muon. If the track is not matched to a muon or an electron, then the vertex is rejected.

Prompt electrons

Prompt electrons are reconstructed as described in section 4.1.3. The prompt electron in the event is defined as the highest p_T electron with $|d_0| < 3$ mm and $|z_0 \sin(\theta)| < 0.5$ mm that does not overlap with any tracks in the displaced vertex. The prompt electron is required to pass the *LHMedium* quality working point. The electron is also required to have a $p_T > 4.5$ GeV. If the selected prompt electron overlaps with a prompt muon within $\Delta R < 0.05$, then the event is rejected.

Displaced electrons

Displaced electrons are tracks in a displaced vertex that are matched to electron objects. The matching of tracks to electron objects is described in section 4.1.3. Displaced electrons are required to pass the custom *VeryVeryLoose* quality working point. The electron is also required to have a $p_T > 4.5$ GeV and have a $\Delta p_T / p_{T, \text{track-only}} < 0.5$, where $\Delta p_T = p_{T, \text{track-only}} - p_{T, \text{calibrated electron}}$. If a displaced electron overlaps with a muon and the matched muon does not pass the *Loose* quality working point, then the track is defined as an electron if the matched electron passes the *VeryVeryLoose* quality working point. If the track is not matched to a muon or an electron, then the vertex is rejected.

5.7 Event Selection

In this search for displaced HNLs, a series of requirements are used to select events of interest. These selections are separated into a set of pre-selection requirements that are, for the most part, defined by properties of the displaced HNL signal or by the detector acceptance. Then, signal region (SR) selections are applied to maintain sensitivity to the signal, while also reducing the backgrounds (non-signal physics processes) as possible. Events that meet all the pre-selection and SR requirements are analyzed in the results shown in section 5.11.

This section begins with a description of the main analysis framework that is used to select data and simulated signal events (see section 5.7.1) and then describe the pre-selection requirements and SR selections in section 5.7.2 and section 5.7.3 respectively. Additional

regions, described in section 5.7.4, are also defined to primarily study backgrounds in the data.

5.7.1 Analysis Framework

Simulated signal and data events are processed through the analysis framework to select interesting events. The first steps in the framework, including the trigger, filter, reconstruction and derivation, all use analysis software that is written by collaboration members who are experts in the areas of triggers and object reconstruction. These steps are designed to reconstruct the physics objects of interest and reduce the number of events that need to be processed in the remaining steps to save time and computing resources. Generally speaking, these steps are common to most physics analyses and therefore, to streamline this process, they are usually run by a central data processing team. The subsequent steps are usually shared by a few analysis teams and are customized to apply events selections for each analysis. Since these steps are custom to the analysis, analysis team members usually play a larger role in the processing of data from the derivation step to produce the final results.

This search for displaced HNLs uses a special stream of data from the ATLAS experiment that initially applies an event selection, known as a “filter”, to select events that are potentially interesting for the analysis. This filter step is required in order to run the displaced tracking reconstruction described in section 4.2.1, since LRT cannot be run over all the full dataset because it is computationally expensive and requires a large amount of disk space to store all the tracking information.

The diagram in fig. 5.18 shows how simulated signal and data events are processed in the analysis framework. The eight main steps in the framework are as follows,

1. **Trigger:** It is impossible, in terms of data readout speed and storage, for ATLAS to process every proton-proton collision event. As a result, triggers are designed to select events of interest for further processing (see section 3.2.6 for more details about triggers). In this analysis, standard single muon and electron triggers are used to select data events. In the signal, it is usually the prompt lepton produced in the W boson decay that passes the trigger requirements.
2. **Filter:** A displaced HNL filter is used to select interesting data events that are then reprocessed to reconstruct large-radius tracks. Only about 10% of the full dataset in the LHC Run 2 is selected by all the analysis filters included in this step. For more details see section 5.7.2.
3. **Reconstruction:** Events that are selected by the filter are then reconstructed using the algorithms described in chapter 4. In this step, DVs are reconstructed using the default VSI configurations discussed in section 5.5.1.

4. **Derivation:** A derivation step is used to skim away information that is not required for the displaced HNL analysis. This reduces the size of the output files by 98% and also reduces the number of events that need to be reprocessed saving time and computing resources. This derivation step also re-runs the displaced vertex reconstruction using the custom VSI configuration described in section 5.5.2. A basic event selection is applied in this step that requires the event to have at least one DV reconstructed using the custom VSI configuration.
5. **Filter Emulator:** A filter emulator is run to estimate the impact of the data filter on simulated signal events. See section 5.7.2 for more details.
6. **Ntuples:** Events from the derivation step are processed and stored in ordered sets that contain n-elements. These data structures, known as “ntuples”, store the necessary information about each event, including the prompt lepton and displaced vertex information. This step stores the information in such a way that it is relatively fast to further process the events to study the impact of the event selection in data and simulation.
7. **Histograms:** After the ntuples are produced, a histogram-making step applies remaining events selections. This step also saves mini-ntuples that contain a smaller subset of the information only for selected events. These mini-ntuples are used as an input in the statistical analysis framework.
8. **Statistical analysis:** The final step in the analysis pipeline is the statistical analysis of selected events in the signal and background. The details of this framework are described in section 5.10.

5.7.2 Pre-Selections

All data events are required to pass basic event cleaning and data quality requirements discussed in section 5.4.1. In addition to this basic event cleaning, a series of pre-selection requirements are defined to select events of interest in the HNL search. A full list of the pre-selection requirements can be found in table 5.8. These selections include:

Trigger

Single muon and electron triggers are used to select events in this analysis. The list of triggers is shown in table 5.4. These triggers have a minimum p_T requirement of 20-26 GeV (depending on the year). The trigger efficiency for the HNL signal is $\sim 30 - 35\%$. Since the prompt lepton spectrum does not have a strong dependence on the HNL mass (see fig. 5.7), this trigger efficiency is fairly constant for all HNL mass and lifetime signal models.




Figure 5.18: The analysis framework pipeline that shows how data and simulated signal events are processed in the displaced HNL analysis.

Table 5.4: Single muon and electron triggers from 2015-2018 that are used to select events of interest.

Lepton	Trigger name	2015–2016	2017	2018
muon	HLT_mu20_loose_L1MU15	X		
	HLT_mu26_ivarmedium	X	X	X
electron	HLT_e24_lhmedium_L1EM20VH	X		
	HLT_e26_lhtight_nod0_ivarlose	X	X	X

Filter

Four HNL filters are used to select events for the special data reprocessing that includes large-radius track and displaced vertex reconstruction. These filters require at least one prompt and one displaced lepton. The filters are defined for the various combinations of prompt and displaced muons and electrons. Their requirements are listed in table 5.6.

5.7. Event Selection

The filter rate is defined as the number of data events that pass the filter selections divided by the total number of events. The filter rate for the four HNL filters, as well as the overall rate that includes all analysis filters in the data reprocessing, is listed in table 5.5.

Table 5.5: The HNL filter and overall filter rate for a set of test events in each of the four data taking years (2015-2018).

Year	HNL 4-filter rate	Overall filter rate
2015	0.61%	3.04%
2016	2.09%	7.00%
2017	3.77%	10.95%
2018	2.38%	10.49%

Table 5.6: Requirements of the four HNL analysis filters. Definitions of the lepton selections (e.g. ID or isolation requirements) used in the filter can be found in the descriptions of the lepton reconstruction algorithms in section 4.1.3 and section 4.1.4.

		Prompt μ displaced μ filter	Prompt e displaced μ filter	Prompt e displaced e filter	Prompt μ displaced e filter
HLT trigger		Single muon triggers	Single electron triggers		Single muon triggers
Prompt lepton	p_{T}	$> 28 \text{ GeV}$			
	$ \eta $	< 2.5			
	Muon type	Combined	-		Combined
	Electron ID	-	$LHLoose$		-
	Isolation	$\sum p_{\text{T}}^{(0.3)}/p_{\text{T}} < 0.05$			
Displaced lepton	p_{T}	$> 5 \text{ GeV}$			
	$ \eta $	< 2.5			
	Muon type	Combined, standalone, segment-tagged		-	
	Electron ID	-		None	
	Isolation	$\sum p_{\text{T}}^{(0.3)}/p_{\text{T}} < 1.0$			
	$ d_0 $	$> 0.1 \text{ mm}$ for combined muon with $\chi^2/\text{dof} < 5$		$> 1 \text{ mm}$	

One unique feature about processing of the ATLAS data is that there is not an equivalent filter step in the processing of simulated signal events. This can be seen visually in fig. 5.18 where the simulated signal events only enter the analysis pipeline at the reconstruction step. As a result, a filter emulator selection is also applied using the same selections listed in table 5.6. However, while the selections used in the filter and filter emulator are the same, the objects, such as muons and electrons, have been re-fit using large radius tracks. As a result, differences on the order of 3% are observed between data and simulated signal events.

Therefore, a filter uncertainty is assigned to account for this difference in the event selection. More details about this uncertainty can be found in section 5.9.

Displaced Vertex

Events are required to include a displaced vertex that is reconstructed using the custom lepton-only vertex configuration described in section 5.5.2. The DV is required to have exactly two tracks matched to displaced leptons. The DV is also required to be within a fiducial volume of $4 \text{ mm} < L_{xy} < 300 \text{ mm}$. The definitions of displaced leptons (e/μ) can be found in section 5.6.

Prompt Lepton

Selected events are required to have one prompt lepton (either a muon or electron depending on the channel). The definition of a prompt lepton can be found in section 5.6.

Trigger Matching

The trigger matching selection requires that at least one lepton in the event passes the single lepton trigger requirements. This lepton is referred to as the “trigger-matched” lepton. About 90% of the time, the prompt lepton is the trigger-matched lepton. However, there is also a small acceptance to signal events with where the displaced leptons fires the trigger. This recovers some efficiency to events where the prompt lepton has p_T below the trigger’s minimum p_T thresholds. Thus, the trigger matching selection requires at least one prompt or displaced *Medium* (*LHMedium*) muon (electron) is trigger-matched and has $p_T > 27 \text{ GeV}$.

5.7.3 Signal Region Selection

After applying the pre-selection requirements in section 5.7.2, additional selections are applied that further reduce the number of events in the SR. These selections can be divided up into selections that are required by the HNL signature and selections that are applied to remove signal-like events from background processes. Selections defined by the HNL signature include:

DV Charge

HNLs are neutral particles and will decay to particles whose charges sum to zero. Thus, in the leptonic decay mode, the two charged tracks will have opposite charge. The DV charge selection requires that the selected DV has exactly two opposite-sign (OS) tracks.

DV Type

Leptons in the displaced vertex can have three different combinations of flavours depending on the HNL decay channel. These combinations include two muons ($\mu\mu$), two electrons (ee)

or one muon and one electron ($e\mu$). The DV type selection uses the displaced lepton object definitions in section 5.6 to select DVs with tracks that are matched to the corresponding lepton flavours in each channel.

Tri-Lepton Mass

The tri-lepton mass (m_{lll}) is the invariant mass of both the prompt lepton from the W boson decay and the two displaced leptons in the vertex from the HNL decay. The tri-lepton mass is reconstructed without the information about the kinematics of the neutrino in the HNL decay. As a result, this variable underestimates the W boson mass for simulated signal events (see fig. 5.7). Due to the the minimum p_T requirement in the trigger selection, the tri-lepton mass distribution starts to fall off around 40 GeV. Thus, the tri-lepton mass selection requires $40 < m_{lll} < 90$ GeV.

HNL Mass

The reconstruction of the HNL mass (m_{HNL}) is described in detail in section 5.5.3. The HNL mass selection requires that each event has $m_{\text{HNL}} < 20$ GeV. HNL models that the LHC is currently sensitive to with $m_N > 20$ GeV are prompt. Thus, this displaced analysis is not sensitive to models with HNL masses larger than 20 GeV

Prompt and Same-Flavour Displaced Lepton Charge

For HNL decays that conserve lepton number, the prompt lepton and the displaced lepton that couples to the HNL must have opposite charges (see section 2.4). This analysis can only separate LNC and LNV decays using an explicit charge requirement for the 1SFH model in the $\mu-\mu e$ and $e-e\mu$ channels where the displaced leptons are experimentally distinguishable. In these channels for 1SFH models with only have LNC decays (e.g. in the “Dirac-limit”), an OS charge requirement between these leptons keeps 100% of the signal and removes approximately half of the expected number of background events.

In addition to the selections defined by the HNL signature, a series of selections are applied to remove contributions from background processes. These selections include the cosmic muon, heavy-flavour decay and Z mass vetoes discussed in section 5.8. These selections rely on a few additional variables including the DV mass (m_{DV}), the DV radius (L_{xy}) and the invariant mass of the prompt lepton plus one of the displaced leptons (m_{ll}). A summary of the selections applied in each channel can be found in table 5.7 and a list of all selections can be found in table 5.8.

Table 5.7: Signal region selections that are applied in each channel.

Selection	$e-ee$	$\mu-ee$	$e-e\mu$	$\mu-\mu e$	$e-\mu\mu$	$\mu-\mu\mu$
Pre-selection: prompt μ prompt e	x	x	x	x	x	x
DV charge	x	x	x	x	x	x
Prompt + disp. l charge (single-flavour mixing with LNC decays)			x	x		
DV type: ee $e\mu$ $\mu\mu$	x	x	x	x	x	x
Tri-lepton mass	x	x	x	x	x	x
HNL mass	x	x	x	x	x	x
Cosmic muon veto	x	x	x	x	x	x
Material veto	x	x				
Heavy-flavour decay veto: $m_{\text{DV}}^{\text{cut}}$ $m_{\text{DV}}-L_{xy}$ cut	x	x	x	x	x	x
Z mass veto	x		x	x		x

Table 5.8: Description of the selection criteria used to defined the SR. All criteria are further explained within this chapter.

Level	Selection	Value
Pre-selection	Event cleaning Primary vertex Trigger Trigger-matched lepton Filter Prompt lepton Trigger matched lepton Displaced lepton-only vertex Number of tracks in DV Fiducial volume	Standard ATLAS event cleaning At least one with $p_T > 500$ MeV At least one single muon or electron trigger At least one <i>Medium</i> (μ) or <i>LHMedium</i> (e) lepton with $p_T > 27$ GeV At least one HNL filter At least one At least one At least one 2 $4 \text{ mm} < L_{xy} < 300 \text{ mm}$
SR selection	DV charge Prompt + disp. l charge DV type Tri-lepton mass HNL mass Cosmic muon veto Material veto Heavy-flavour decay veto Z mass veto	Opposite-sign tracks Opposite-sign leptons (For 1SFH models with only LNC decays) ee , $e\mu$ or $\mu\mu$ vertex $40 \text{ GeV} < m_{lll} < 90 \text{ GeV}$ $m_{\text{HNL}} < 20 \text{ GeV}$ $\sqrt{(\Sigma\eta)^2 + (\pi - \Delta\phi)^2} > 0.05$ Applied for ee DVs only $m_{\text{DV}} > 5.5 \text{ GeV}$ ($\mu\mu$ DVs) or $m_{\text{DV}}-L_{xy}$ cut (ee or $e\mu$ DVs) $m_{ll} < 80 \text{ GeV}$ or $m_{ll} > 100 \text{ GeV}$, for same-flavour opposite-sign leptons

5.7.4 Additional Regions

ATLAS searches are designed to be blind analyses in order to avoid biases when studying the data. To perform a blind analysis, the data is analyzed in such a way that the final answer is not known until the complete analysis strategy is designed and approved by members of the collaboration. However, this does not mean that the data is never studied in the design of the analysis. Instead, additional regions are defined that have negligible contamination from the signal events to study the data. The regions in this search for HNLs include events with same-sign vertices, a validation region and a control region.

Same-Sign Vertices

The search for displaced HNLs requires a two-lepton DV to be reconstructed in the ID. Since this two-lepton vertex originates from a neutral particle, the signal requires that the charges of the two leptons in the vertex have opposite signs. Therefore, selecting DVs with two same-sign (SS) leptons will have negligible signal contamination and this selection can be used to study contributions from background processes.

Control Region

The control region (CR) is defined using the same analysis selections described in table 5.8, except that the HNL mass selection is inverted to require that $20 < m_{\text{HNL}} < 50$ GeV. This region is used to study OS and SS DVs in a region close to the SR. The number of observed OS events in the CR is also used to constrain background estimate in the HNL mass fit (see section 5.10 for more details about the fit). In the region with $m_{\text{HNL}} > 20$ GeV, the signal is not excluded, but its acceptance multiplied by its cross section with the current dataset yields significantly fewer than one signal event. Therefore, there is no signal contamination in this inverted-HNL-mass CR.

Validation Region

The validation region (VR), also known as the inverted-prompt-lepton VR, is defined by vetoing events that contain a prompt electron or muon. HNL signal events produced from a W boson decay should always have a prompt lepton and thus, this region has a negligible signal contribution. The VR has higher statistics than the SR and is used to study background physics processes that produce OS displaced vertices. This region is also used to select a collection of DVs with two random crossing lepton-matched tracks that are used to estimate the background (see section 5.8.6 for more details). The list of selections that define the validation region can be found in table 5.9.

Table 5.9: Definition of the inverted-prompt-lepton VR. A prompt lepton is defined as any muon or electron object with *Loose* or better quality that has $|d_0| < 3$ mm and $|z_0 \sin(\theta)| < 0.5$ mm.

Level	Selection	Value
Pre-selection	Event cleaning	Standard ATLAS event cleaning
	Primary vertex	At least one with $p_T > 500$ MeV
	Filter	Pass <i>any</i> analysis filter
	Number of prompt leptons	0
DV selection	Displaced lepton-only vertex	At least one
	Cosmic veto	$\sqrt{(\Sigma\eta)^2 + (\pi - \Delta\phi)^2} > 0.05$
	Number of track in vertex	2
	Fiducial volume	$4 < L_{xy} < 300$ mm
	DV type	ee , $e\mu$ or $\mu\mu$ vertex

5.8 Backgrounds

Backgrounds are defined as physics processes that contribute to the selected events in the signal region that are not from the production and decay of an HNL. Five sources of backgrounds are identified that produce two-lepton OS displaced vertices and satisfy the pre-selections defined in section 5.7.2. These sources include:

1. Vertices from random track crossings (see section 5.8.1)
2. Vertices from material interactions (see section 5.8.2)
3. Decays from metastable particles, such as b-hadrons (see section 5.8.3)
4. $Z \rightarrow ll$ decays (see section 5.8.4)
5. Muons from cosmic-ray interactions (see section 5.8.5)

The background estimation technique in this analysis uses a data-driven method to determine the contribution in the SR from random track crossings. This method is discussed in detail in section 5.8.6. This assumes that all other sources of backgrounds are negligible compared to the random track crossing background or have been removed by dedicated cuts in the event selection. In particular, backgrounds from metastable decays are rejected from our signal region by DV mass and radius selections. Material interactions are either negligible for channels with displaced muons (e.g. $e\mu$ or $\mu\mu$ vertices) or rejected by a material veto (e.g. ee vertices). Events from $Z \rightarrow ll$ decays are rejected by a Z mass veto cut and muons from cosmic-ray interactions are rejected using a cosmic muon veto. Section 5.8.6 outlines the studies that have been performed to validate these assumptions.

5.8.1 Random Track Crossings

Random track crossings occur when two uncorrelated tracks accidentally cross in the inner detector sufficiently close to each other that the vertexing algorithm associates them to a common displaced vertex. While the probability of a two-lepton random track crossing is relatively low, this source of background is considered to be irreducible, since the DV mass is unconstrained.

One way to estimate this background is to use the number of random crossings with two SS tracks as an estimate for the number of random crossings with two OS tracks. This method assumes that the probability for two tracks to form a random crossing vertex is independent of the charge of the two tracks. Hence it is equally likely that random crossings will form a SS or OS vertex and thus, the OS random crossing background can be estimated using SS displaced vertices. Validation of this charge assumption is described in section 5.8.6. Similar techniques are used to estimate accidental crossings of tracks in other ATLAS searches, such as a super-symmetry search [81] and the 2019 ATLAS displaced HNL search [70]. However, this SS background estimation method is statistically limited and, as a result, it is not used as the final background estimate in this analysis. Nonetheless, this method is used to study the backgrounds and validate assumptions and event selections in the analysis.

Instead of using the SS estimate, backgrounds from random crossings are estimated using a novel data-driven approach that mixes prompt leptons and displaced vertices from *different* events. This method assumes that the prompt lepton is unrelated to the random crossing DV. Thus, events created from all possible pairs of prompt leptons and DVs can be used to increase the statistics of the background estimate. This assumption has been validated using the prompt lepton shuffling technique and SS background estimate. For more details, see the validation studies in section 5.8.6.

A full description of the data-driven estimate for random track crossings using the prompt lepton shuffling method and the corresponding background uncertainties is found in section 5.8.6. The predicted number of background events in the signal and control regions from random track crossings are listed in table 5.13.

5.8.2 Material Interactions

Interactions with the inner detector material (both active layers and the support structures) can produce two-lepton vertices at some distance from the interaction point. This creates a source of displaced lepton vertices that is unrelated to the decay of an HNL. Material interactions are removed from the SR using DV mass cuts and a dedicated material map veto for channels with displaced ee vertices.

Many of the material interactions that produce ee displaced vertices have a reconstructed displaced vertex mass of $m_{\text{DV}} < 2$ GeV. This can be seen in fig. 5.19 (c), which shows DV mass

as a function of the DV radius for OS ee displaced vertices in the statistically enhanced VR (see section 5.7.4 for a description of the VR). Large clusters of events near the material layers (dashed grey lines) can be removed by requiring that the m_{DV} is at least 2 GeV. Figure 5.19 (b) and (c) shows that material effects for channels with displaced muons are not dominant.




Figure 5.19: OS displaced vertices in the VR shown in bins of m_{DV} versus L_{xy} . All vertices are required to pass the VR selections. Leptons in the selected DVs are also required to pass *VeryVeryLoose* electron and *Medium* muon identification working points.

However, it is also possible that material interactions can produce DVs with $m_{\text{DV}} > 2$ GeV, which makes it difficult to distinguish between these material interactions and HNL signal decays, even with the other vertex selections applied. Therefore, a *material veto* uses two three-dimensional maps of the ID material to reject vertices that decay within the material layers. The first map, which extends in radius to $L_{xy} < 150$ mm, is constructed from low-mass vertices in the data (with selections to remove low-mass SM resonances) [82]. The second map, which extends in radius from $150 \text{ mm} < L_{xy} < 300$ mm, is reconstructed using vertices in simulation to create a more complete map than statistics will allow in data [82]. Both maps extend out to $|z| = 300$ mm and so the material veto also requires that the DV has

$|z| < 300$ mm. The impact of the material veto for ee OS vertices in the VR is shown in fig. 5.20 (a).




Figure 5.20: OS displaced vertices from data events using the VR selections as a function of L_{xy} . Displaced vertices are separated by flavour into three channels: ee , $e\mu$ and $\mu\mu$. Distributions labeled “All” (solid circles) have no additional selections applied. The distribution with open circles in the ee channel has only the material veto applied and distributions with open squares have the heavy-flavour decay veto applied.

Peaks in the background for ee vertices can be seen, which correspond to regions that are more material-dense. After applying the material veto, many of the material interactions that peak at the detector layers are gone. Due to imperfect detector modelling, the material map does not remove all the material interactions. However, the SR selections also include m_{DV} cuts designed to remove vertices from metastable particle decays (see the heavy-flavour decay veto in section 5.8.3). After applying the full analysis selections (including the heavy-flavour decay veto) the background from material interactions is considered to be negligible. Any uncertainty in this assumption is taken into account by a background DV systematic uncertainty that is derived for the nominal background estimate accounting for differences

between SS and OS vertices (see section 5.8.6 for more details).

Applying the material veto also removes any signal decays near material layers. This results in roughly a 30 – 40% loss in signal efficiency. In fig. 5.20 (b) and (c) large peaks are not observed near the material layers in the VR for $\mu\mu$ and $e\mu$ displaced vertices and, therefore, the material veto is not applied for channels with at least one displaced muon.

5.8.3 Metastable Particle Decays

SM particles produced in proton-proton collisions at the LHC with lifetimes near the displaced HNL signal region, and potentially result in two-lepton vertices with a DV radius between 4 and 300 mm, are charm, bottom and strange hadrons, as well as the tau lepton.

- Hadrons with b quarks, also called b-hadrons, can have decay lengths of a few mm when they are highly boosted and/or in the tail in the exponential decay distribution. These particles have masses of the order of 5 to 6 GeV, thus encroaching on the signal region. Double-semi-leptonic decays $b \rightarrow c + W(\rightarrow l + \nu)$ followed by $c \rightarrow d + W(\rightarrow l + \nu)$ can also produce contributions to all three DV types ($e\mu$, $\mu\mu$ and ee).
- J/Ψ mesons have masses of 3.1 GeV and can decay to pairs of same-flavour OS leptons ($\mu^+\mu^-$ or e^+e^-). These decays can be displaced if the J/Ψ s are produced by other heavy-flavour particles. These would mostly come from b-hadron decays in order to be displaced enough to contaminate the signal region. They are separated here from other b-hadron decays because their DV mass would be narrowly peaked at the J/Ψ mass.
- Strange baryons and the tau lepton have lifetimes near the signal region, however they have smaller masses, between 1 and 2 GeV. Moreover, strange baryons almost always decay into hadrons, thus almost never produce vertices containing leptons. As such, they do not contribute to the background in this search.
- The Υ meson is another relatively massive (9.5 GeV) particle, which can decay into a lepton pair. Its production cross section is small and its lifetime is very short. Since the Υ meson is not produced in b-hadron decays, one would not expect it to give rise to displaced vertices.

Contributions from metastable decays are studied in the VR and separated according to the DV type ($\mu\mu$, $e\mu$ or ee). In particular, when looking at the DV mass distribution of the selected events in fig. 5.21, a heavy-flavour continuum spectrum is observed that falls off around 5.5 GeV and a narrow contribution from J/Ψ events is observed in the same flavour channels (ee and $\mu\mu$). The background estimation method described in section 5.8.6 assumes that only random crossings contribute to backgrounds in the SR. Therefore, analysis selections are designed to remove additional OS contributions from metastable decays from the SR.




Figure 5.21: OS and SS displaced vertices in the Validation Region as a function of the DV mass. Leptons in the DVs are required to pass the *Medium* (*VeryVeryLoose*) quality working points for muons (electrons). The material map veto is also applied for the ee channel. The ratio of OS to SS vertices is shown in the bottom ratio plot.

The correlation between m_{DV} and L_{xy} in fig. 5.19 shows that OS DVs from metastable decays typically occur at smaller radii. Based on the distributions in fig. 5.19, an $m_{\text{DV}} > 5.5$ GeV selection would reject the metastable decays. However, this completely removes sensitivity to HNL decays with $m_{\text{DV}} < 5.5$ GeV and significantly impacts the sensitivity to models with $5.5 \text{ GeV} < m_N < 10 \text{ GeV}$.¹² Instead, a selection that removes events with DVs that have small m_{DV} and L_{xy} would remove metastable decays while maintaining some sensitivity to long-lived HNL signals with $m_N < 5.5$ GeV.

To exploit the correlation between m_{DV} and L_{xy} in the background, a diagonal cut in the $m_{\text{DV}}-L_{xy}$ plane is optimized using the difference between OS and SS vertices as a measure of

¹²The displaced vertex mass is reconstructed from the two charged leptons from the HNL decay (i.e. the neutrino is not reconstructed) and, thus, the displaced vertex mass in signal always underestimates the simulated HNL mass (see fig. 5.9).

how different the two distributions are. Based on the properties of random crossings, if there is good agreement between SS and OS vertices, then this implies that only random crossings remain. Therefore, regions of phase space where there is a large difference between OS and SS vertices have contributions from non-random processes, while regions with small differences are dominated by random track crossings. The result of the optimization study is shown in fig. 5.22.




Figure 5.22: Correlation plot of DV mass (m_{DV}) versus radius ($L_{xy} = r_{\text{DV}}$) for the difference between OS and SS vertices in the three DV type channels. The data events studied in these plots are selected from events in the VR. The heavy-flavour decay veto applied in each channel is shown by the solid red line. The z -axis maximum in the plots is saturated at 30 in order to provide a colour scale that can distinguish regions with small negative entries (more SS than OS DVs) from those with small positive entries (more OS than SS DVs).

In the $e\mu$ channel, a diagonal cut with $m_{\text{DV}} > 7 \text{ GeV} (1 - \frac{L_{xy}}{150 \text{ mm}})$ for vertices with m_{DV} between 2 and 5.5 GeV (red line in fig. 5.22 (b)) is sufficient to remove the OS events in the small m_{DV}/L_{xy} region. Below 2 GeV, there is still an excess of OS events at large radii and above 5.5 GeV, b-hadron decays are not expected to contribute. This same diagonal cut is also

sufficient to remove contributions from metastable decays in the ee channel. However, in the $\mu\mu$ channel (fig. 5.22 (a)), the diagonal cut fails to remove contributions from resonant J/Ψ decays. Changing the slope of the diagonal cut further reduces the additional phase space that is gained by including this cut. The signal efficiency increase for channels with $\mu\mu$ DVs using a modified diagonal cut was not sufficient to significantly increase the analysis sensitivity to $\mu\mu$ decays. Therefore, the channels with $\mu\mu$ DVs have a flat $m_{\text{DV}} > 5.5$ GeV selection applied.

After applying the heavy-flavour decay veto, DVs from data events in the validation region are compared to check the agreement between the SS and OS distributions. Good agreement between OS and SS DVs (a ratio of 1) in fig. 5.23, indicates that random track crossings dominate. With the heavy-flavour decay veto applied, the background from metastable decays is considered to be negligible. Any uncertainty in this assumption is taken into account by the DV background systematic described in section 5.8.6. A summary of the heavy-flavour decay veto applied in the SR can be found in table 5.10.





 (a) $\mu\mu$ DVs

 (b) $e\mu$ DVs

 (c) ee DVs

Figure 5.23: OS and SS displaced vertices with the heavy-flavour decay veto applied for data events in the VR. The material veto is also applied for displaced ee vertices.

Table 5.10: Description of the heavy-flavour decay veto.

DV type	Selection
$\mu\mu$	$m_{\text{DV}} > 5.5 \text{ GeV}$
ee or $e\mu$	$m_{\text{DV}} > \begin{cases} 5.5 \text{ GeV, if } L_{xy} \leq 225/7 \text{ mm} \\ 7 \text{ GeV}(1 - \frac{L_{xy}}{150 \text{ mm}}), \text{ if } 225/7 \text{ mm} < L_{xy} < 750/7 \text{ mm} \\ 2 \text{ GeV, if } L_{xy} \geq 750/7 \text{ mm} \end{cases}$

5.8.4 Leptonic Z Boson Decays

Leptonic decays of Z bosons contribute to the background if one of the leptons from the decay is identified as the prompt lepton and the other lepton forms a vertex with a third lepton. This third lepton could be from a random crossing or an asymmetric photon conversion where only one of the leptons is reconstructed. The random crossing background estimate, described in section 5.8.6, assumes that the prompt lepton is unrelated to the leptons in the displaced vertex. This assumption is not true for Z decays where the prompt lepton and one of the leptons in the vertex are produced by a Z boson decay and, therefore, these events are removed from the signal region with a dedicated Z mass veto. The Z decay background is only present in channels where the prompt lepton and at least one of the displaced leptons have the same flavour (i.e. $\mu-\mu\mu$, $e-ee$, $\mu-\mu e$ and $e-e\mu$). The Z mass veto is not applied to channels where the prompt and displaced leptons have different flavours (i.e $\mu-ee$ and $e-\mu\mu$).

To study this background, $Z \rightarrow ll$ events are selected in data using the pre-selection requirements and the heavy-flavour decay veto (described in section 5.8.3). Selected events must have a two-lepton SS vertex. The prompt and same-flavour lepton in the displaced vertex must also have opposite charges. Events that do not meet the last charge requirement cannot be produced in a neutral Z boson decay. The mass between this pair of leptons (m_{ll}) in the selected data events is shown in fig. 5.24. Z decays are identified as events with m_{ll} near the Z boson mass of 91 GeV. For channels with indistinguishable displaced leptons (e.g. $\mu-\mu\mu$ and $e-ee$), the displaced lepton with the highest p_T is paired with the prompt lepton to reconstruct the di-lepton mass, m_{ll} . It is assumed that the displaced lepton with the highest p_T is more likely to come from the Z decay. This assumption was validated using data events to show that the highest p_T displaced lepton usually comes from the Z decay.

The Z mass veto selection is summarized in table 5.11. Events from Z decays are removed using a m_{ll} cut in a window around the Z mass. This selection has a negligible impact on our signal acceptance because, in signal, the same-flavour lepton pair originated from a W boson decay and typically has a di-lepton less than 80 GeV.




Figure 5.24: The invariant mass between the prompt lepton and same-flavour displaced lepton (m_{ll}) for selected VR data events in the Z mass veto study. Events around the Z mass with $80 \text{ GeV} < m_{ll} < 100 \text{ GeV}$ are identified as Z decays and are removed from the SR.

Table 5.11: Description of the Z mass veto.

Channel	Selection
$\mu-\mu\mu$ $\mu-\mu e$ $e-e\mu$ $e-ee$	$m_{ll} < 80 \text{ GeV}$ or $m_{ll} > 100 \text{ GeV}$ if prompt and same-flavour leptons have opposite charges.
$e-\mu\mu$ $\mu-ee$	No veto applied.

5.8.5 Muons from Cosmic-Ray Interactions

Muons from cosmic-ray interactions in the atmosphere can penetrate through the earth and interact with the ATLAS detectors. When cosmic muons enter the magnetic field in the ID, they bend in the ϕ direction. If a vertex is reconstructed from two segments of a muon that traverses the detector from the top down, then it can be reconstructed as an opposite-sign

vertex. These cosmic ray muon interactions can occur at any distance from the interaction point and therefore produce a background source of OS displaced vertices. However, since these two tracks come from the same energetic muon, they should be almost perfectly back-to-back in $\eta - \phi$ space. This means that $\Delta\phi - \pi$ and $\sum \eta$ of the two tracks should be close to zero.

A minimum separation between the tracks in the displaced vertex is required to remove back-to-back muons from cosmic ray interactions. The cosmic veto stipulates that the two leptons within the displaced vertex have a separation, $\sqrt{(\sum \eta)^2 + (\pi - \Delta\phi)^2} > 0.05$. This selection is based on a study of cosmic-ray muon interactions that was performed in the context of another ATLAS displaced vertex search [83]. A similar selection was also used in the 2019 ATLAS displaced HNL search that analyzed pp data from 2016 [70]. This selection is 100% efficient in signal and is therefore applied to all channels.

5.8.6 Data-Driven Estimate of Random Track Crossings Using a Prompt Lepton Shuffling Method

The main drawback of the data-driven SS background estimate for random crossings (described in section 5.8.1) is that it is statistically limited. The probability of two leptons accidentally crossing and forming a displaced vertex is very small and, in addition to the vertex being reconstructed, the random crossing DV also needs to be paired with a “good quality” prompt lepton that passes all the pre-selection requirements. Given that the signal distribution for the HNL mass (see section 5.5.3) peaks at the simulated mass, it would be ideal to have a background estimate with a smooth distribution as a function of this variable to provide a strong discriminant between signal and background. To do this requires an increase in the statistics of the random track crossing background estimation method.

The data-driven SS background estimate provides one prompt lepton for every two-lepton SS vertex. However, if we assume that the prompt lepton is unrelated to the DV, then we can pair each two-lepton vertex with a prompt lepton from a different event and re-calculate the HNL and tri-lepton masses. The validation of the assumption that the prompt lepton is unrelated to the random crossing DV is described in section 5.8.6. Once it is established that the pairing of the prompt lepton and displaced vertex is *random*, pairing leptons with vertices from *different events* can be used to form a larger collection of events to increase the statistics of the background estimate.

The increase in statistics achieved with this method is limited by the number of selected prompt leptons and the number of selected random crossing DVs. To address the first limitation, a “loose” collection of prompt leptons is selected from data events using the pre-selection requirements and the presence of a SS DV with the same DV type (i.e. $\mu\mu$, $e\mu$, ee) that has no lepton identification or DV mass requirements. The prompt leptons selected using this

method were found to have similar kinematics (e.g. p_T , η , ϕ) to prompt leptons in the SR and thus, this selection does not bias the shuffling estimate. The second limitation is addressed by selecting random crossing DVs from data events in the VR. Events in this region do not contain prompt leptons and therefore have a statistically enhanced number of random crossing DVs from real data events.

The nominal background estimate shuffles OS vertices selected in the VR with the “loose” collection of prompt leptons. Combining each OS DV with each lepton, the shuffled lepton (SL) sample has at least 2000 times the number of events in the original event sample for each channel, resulting in a high-statistics background model. The SL sample is initially normalized to the number of events from the SS background estimate, such that:

$$N_{SR}^B = N_{SR}^{SL} \times \frac{N_{SR}^{SS} + N_{CR}^{SS}}{N_{SR}^{SL} + N_{CR}^{SL}} \quad (5.40)$$

$$N_{CR}^B = N_{CR}^{SL} \times \frac{N_{SR}^{SS} + N_{CR}^{SS}}{N_{SR}^{SL} + N_{CR}^{SL}} \quad (5.41)$$

where N_{SR}^B (N_{CR}^B) is the background prediction in the signal (control) region, N_{SR}^{SS} (N_{CR}^{SS}) is the number of same-sign DV events in the signal (control) region and N_{SR}^{SL} (N_{CR}^{SL}) is the number of shuffled lepton events in the signal (control) region. In the statistical fit described in section 5.10, this normalization is unconstrained, so this choice simply provides an initial starting point for background estimate. Distributions of the shuffled background estimate are shown in fig. 5.25 as a function of the HNL mass. The SR selection applied in these plots is listed in section 5.7.3, except that the $m_{\text{HNL}} < 50$ GeV. In this way, the signal and control regions can be studied together.

Uncertainties in the Background Estimate

Two sources of systematic uncertainties are identified in the prompt lepton shuffling background estimation method. The dominant uncertainty comes from the assumption that the selected DVs used in the shuffling method are from random track crossings. If the selected DVs are not from random crossings, then this provides a source of bias in the estimate. Random track crossings should have uncorrelated tracks and the distributions from random crossing events should be independent of track charge. Therefore, the difference between SS and OS distributions can be used as a measure of the uncertainty on the background prediction. SS vertices from the VR are also selected and shuffled with the “loose” collection of prompt leptons to produce the DV systematic distributions shown in fig. 5.25.

The difference between the nominal and DV systematic distributions is used as a shape uncertainty on the background estimate. The better the shape agreement is between the two

distributions, the smaller the systematic uncertainty is. The shape agreement is the worst in channels with $\mu\mu$ DVs and is limited by the number of DVs in the VR. This results in the largest DV systematic uncertainty in the $\mu-\mu\mu$ and $e-\mu\mu$ channels.




Figure 5.25: The HNL mass (m_{HNL}) shapes for the background estimate (shaded purple), DV systematic (solid blue line), prompt lepton systematic (dashed blue line), observed SS events (grey triangles) and observed OS events (black circles) for the six different channels.

The second source of systematic uncertainty is due to the finite number of prompt leptons used in the shuffling method. A prompt lepton uncertainty is derived using a different collection of prompt leptons and comparing the shapes between the nominal and systematic distributions. In the nominal method, a “loose” collection of prompt leptons is selected using events with a SS DV of the *same* DV type (e.g. in the $\mu\text{--}\mu\mu$ background estimate, events with SS $\mu\mu$ vertices are used to select prompt muons). To account for differences in the HNL mass shape due to the finite number of prompt leptons, the OS vertices are also shuffled with prompt leptons from SS events with a *different* DV type (e.g. in the $\mu\text{--}\mu\mu$ prompt lepton systematic, events with SS $e\mu$ vertices are used to select prompt muons). The prompt lepton systematic distributions are also shown fig. 5.25 and are also used as a shape uncertainty on the background estimate.

The impact of the background systematics in the SR are summarized in table 5.12. For each channel, the difference between the nominal and systematic uncertainty predictions divided by the nominal prediction is shown for the DV and prompt lepton uncertainties.

Table 5.12: The relative systematic uncertainties on the background estimate, estimated by creating shuffled events from SS DVs or from prompt leptons taken from events with different DV channels. The total uncertainty is the quadrature sum of the two contributions. The uncertainties are treated as uncorrelated among the different channels.

Channel	Relative Uncertainty (%)		
	DV	Prompt l	Total
$e\text{--}ee$	35	<1	35
$\mu\text{--}ee$	27	<1	27
$e\text{--}e\mu$	2	1	2
$\mu\text{--}\mu e$	5	5	7
$e\text{--}\mu\mu$	70	2	70
$\mu\text{--}\mu\mu$	79	3	79

Background Predictions

The nominal background and systematic uncertainty distributions are shown for m_{HNL} in fig. 5.25. For reference, the observed SS and OS data events are also shown in these plots. The shuffled background estimate is fairly smooth as a function of the HNL and tri-lepton mass. Assuming that the random crossing events do not have a particular shape in the m_{HNL} distribution, other than those introduced by the kinematic cuts, the upper tails of the reconstructed mass distribution are used to validate that the background modeling matches what is observed in data. The CR defined in section 5.7.4, with m_{HNL} between 20 and 50 GeV, is used to provide additional constraints on the background prediction. This region was also used to validate the background estimation method by comparing OS and SS events prior to

looking at the OS events in the SR.

The background predictions prior to the fit, as well as the observed number of SS and OS events, are shown in table 5.13. The shuffled background predictions are listed with the combined statistical and systematic uncertainty, while the observed number of background events (SS and OS) are shown with statistical uncertainties only.

Table 5.13: Estimated background predictions with statistical and systematic uncertainties in the signal region and the control region. The observed yields are shown for OS DVs and SS DVs.

Channel	Signal region			Control region		
	Background Prediction	Observed yield		Background Prediction	Observed yield	
		SS	OS		SS	OS
$e-ee$	0.6 ± 0.2	1	2	5.4 ± 0.2	5	2
$\mu-ee$	0.3 ± 0.1	0	1	3.7 ± 0.1	4	1
$e-e\mu$	2.3 ± 0.1	0	0	10.7 ± 0.1	13	5
$\mu-\mu e$	0.9 ± 0.1	0	2	4.1 ± 0.1	5	13
$e-\mu\mu$	1.5 ± 1.1	1	1	2.5 ± 1.1	3	3
$\mu-\mu\mu$	2.4 ± 1.4	0	2	3.6 ± 1.8	6	9

Background Estimate Validation

Both the SS and shuffled lepton background estimate assumes that the background events in the SR are dominated by random track crossings combined with the association of a prompt lepton. The shuffled lepton estimate also assumes that the association of a prompt lepton to a shuffled vertex is random. In the SS estimate, the displaced vertex is already associated with a prompt lepton in the same event and no assumptions need to be made. This section will discuss the studies performed to validate these assumptions.

Assumption 1: there is no dependence on track charge in the creation of random crossing vertices.

- Under the assumption that only random crossings survive after applying the heavy-flavour decay veto (and material veto for ee DVs), SS events in the VR should not have any charge asymmetry between $++$ and $--$ DVs. In fig. 5.26, no significant asymmetry is observed between $++$ and $--$ DVs in the $e\mu$ and ee channels. Although the available statistics in the $\mu\mu$ channel is limited, there is no physics motivated reason why $\mu\mu$ vertices should behave differently compared to the other two channels.
- A custom track algorithm, described in appendix A, is used to randomly shuffle tracks

from different events and produce displaced vertices. This algorithm randomly pairs leptons from different events to create a set of fake events. Thus, vertices reconstructed using this method form a pure sample of random track crossings. The track shuffling studies in appendix A (in particular fig. A.4) demonstrate that SS and OS distributions match at better than 20%, including key kinematic variables.

- **Conclusion:** random track crossings are independent of track charge. If the background in our signal region consists only of random crossings, then SS distribution is a good estimate of the OS background.




Figure 5.26: SS displaced vertices in the VR with the heavy-flavour decay veto applied. The material veto for ee DVs is also applied. SS DVs are separated by track charge into vertices with two positively charged tracks (SS $++$) and two negatively charged tracks (SS $--$). This study does not include the $\Delta p_T/p_T$ cut for displaced electrons described in section 5.6.

Assumption 2: backgrounds are dominated by random crossings, except in the low m_{DV} region (below 5.5 GeV) where heavy-flavour decays dominate.

- Using the fact that random crossings are independent of track charge, good agreement between SS and OS distributions indicates that only random crossings remain. In

fig. 5.21, no significant differences between SS and OS distributions are observed above $m_{\text{DV}} \sim 5.5$ GeV. Below 5.5 GeV, heavy-flavour decays dominate, leading to an excess of OS displaced vertices. After applying the heavy-flavour decay veto described in fig. 5.23, good agreement is obtained (within the available statistics).

- The number of SS and OS events are also compared in the CR (see table 5.13) and no significant difference between the two types of events is observed.
- An inverted tri-lepton mass validation region is used to study data event in the tails of the m_{lll} distribution (i.e. $m_{lll} < 40$ or $m_{lll} > 90$ GeV). This region has a signal to background ratio of less than 5%, which makes it possible to study the data events without significant signal contamination. Although this region has low statistics due to the limited number of events that are selected by the filter discussed in section 5.7, no significant differences are observed between SS and OS DVs.
- Conclusion: OS and SS distributions show good agreement in the VR and CR, which implies that the background is dominated by random crossings. The DV systematic, described in section 5.8.6, takes into account any uncertainties from this assumption.

Assumption 3: the pairing of a prompt lepton with a displaced vertex from accidental track crossings is random.

- The shuffling background method relies on artificially creating fake shuffled events and randomly pairing a shuffled displaced vertex with a prompt lepton. Using this method, SS vertices are shuffled with prompt leptons and the shape of the resulting distribution is compared to the shape of the SS background estimate. This comparison is performed in a looser signal region that does not have the heavy-flavour decay veto applied in order to increase the statistics in the shape comparison. Events are selected using a minimum 2 GeV DV mass cut and the pre-selection requirements. The results are shown for the m_{HNL} and m_{lll} variables in fig. 5.27 and fig. 5.28 respectively. The shuffling factor, f , which is defined as the ratio of the number of shuffled events to number of SS background events, is also included in the plots.
- A Kolmogorov-Smirnov (KS) goodness-of-fit test [84] is performed to quantify the shape agreement between the two distributions. Table 5.14 shows the resulting KS probabilities. Two distributions are considered incompatible if the KS probability is much less than one. The results of the KS test indicate that the shuffled and SS background estimate have compatible m_{HNL} and m_{lll} shapes.
- Conclusion: good shape agreement between the shuffled and SS background estimates indicates that the pairing of a prompt lepton with an accidental crossing two lepton displaced vertex is random.




Figure 5.27: HNL mass distribution for the SS background estimate (orig. plep) and for the shuffled background estimate (random plep), where plep is an abbreviation for “prompt lepton”. The shuffling factor, f , which is defined as the ratio of the number of shuffled events to number of original SS background events, is shown on the plots for each channel.




Figure 5.28: Tri-lepton mass distribution for the SS background estimate (orig. plep) and for the shuffled background estimate in colours (random plep), where plep is an abbreviation for “prompt lepton”. The shuffling factor, f , which is defined as the ratio of the number of shuffled events to number of original SS background events, is shown on the plots for each channel.

Table 5.14: The KS test probabilities obtained when comparing the shuffled and SS background estimate HNL mass (m_{HNL}) and tri-lepton mass (m_{lll}) distributions. The KS test is done separately for each channel. Two distributions are considered incompatible if the KS probability is much less than one.

	KS probability					
	$\mu-\mu\mu$	$\mu-\mu e$	$\mu-ee$	$e-e\mu$	$e-ee$	$e-\mu\mu$
m_{HNL}	0.43	0.99	0.94	0.80	0.99	0.99
m_{lll}	0.20	0.23	0.73	0.62	0.89	0.45

5.9 Signal Systematic Uncertainties

This section focuses on systematic uncertainties in the signal prediction from both experimental and theoretical effects. Both up and down variations ($\pm 1\sigma$) are modeled for most of the uncertainties. The sources of systematic uncertainties in the background estimate are discussed in section 5.8.6.

Many of the experimental systematics are well understood and can be evaluated with central tools developed by CP groups. Examples of such systematics include lepton trigger, reconstruction and identification uncertainties. However, due to the non-standard displaced nature of the signal, other uncertainties must be evaluated separately. A summary of the systematic uncertainties and their contributions to the overall uncertainty of the signal selection efficiency are listed in table 5.15. Each uncertainty in the table is briefly discussed below.

Table 5.15: Table summarizing the systematic uncertainties that impact the signal prediction. Systematic uncertainties in this table are the impacts on the signal yields in the SR.

Type	Source	Maximum Selection Efficiency Uncertainty [%]					
		$\mu-\mu\mu$	$\mu-\mu e$	$\mu-ee$	$e-ee$	$e-e\mu$	$e-\mu\mu$
Experimental	Integrated luminosity				2		
	Pileup				3		
	Filter				3		
	Tracking				3		
	Displaced vertexing	11	21	19	20	28	9
	Displaced lepton identification	5	7	7	7	6	4
	Trigger	< 1	1	< 1	< 1	< 1	< 1
	Lepton reconstruction and identification	4	9	12	17	15	2
	W cross section and modelling				3		
Theoretical	HNL branching ratio				5		
	Total				8 – 34		

Displaced Vertexing

The displaced vertex reconstruction efficiency can vary between data and simulation. The impact of this uncertainty is evaluated using displaced kaon decays. Kaons are produced in relative abundance at the LHC and have relatively long mean lifetimes of $c\tau = 26.84$ mm. This provides a way to measure the difference in the vertex reconstruction efficiency for the HNL signal samples with $c\tau_N = 1, 10, 100$ mm).

The ratio of the number of selected K_s displaced vertices in data to simulated events is calculated as a function of the displaced vertex p_T and radius. These distributions are normalized such that the data to simulation ratio is equal to one in the first bin with the smallest radius. This assumes that the tracking and displaced vertex reconstruction is well-understood close to the interaction point. Any uncertainty in this low-radius bin is taken into account by the track reconstruction uncertainty. The difference between the data to simulation ratio and one is symmetrized and reported as the $\pm 1\sigma$ DV reconstruction uncertainty. The maximum size of this uncertainty varies as a function of the DV p_T and radius for each mass-lifetime signal between 9 and 28%.

Track Reconstruction

Evaluation of the track reconstruction efficiency is performed using a “track-dropping” method. The probability of reconstructing a track is determined as a function of p_T and η . Next, HNL decays are selected using the SR selections and this nominal selection efficiency is compared to the efficiency obtained with a certain fraction of tracks removed (dropped) from the event. When a displaced track is removed from the vertex, the event no longer passes the two-track vertex selection and, thus, the event is not selected. The overall difference between the nominal and “track-dropping” selection efficiencies is symmetrized and reported as a 3% uncertainty.

Displaced Lepton Identification

Prompt lepton reconstruction and identification uncertainties are well-understood. However, the displaced nature of the leptons in the HNL decay introduces an additional uncertainty that is not captured in the standard uncertainties. In particular, the lepton identification efficiencies become smaller as the $|d_0|$ of the lepton-matched track becomes larger. This effect is studied in simulation because there is no good source of displaced leptons in the data to model this behaviour. To quantify the impact, distributions of signal events are studied with and without the lepton identification criteria applied (*Medium* for muons and *VeryVeryLoose* for electrons) as a function of the lepton-matched track $|d_0|$. The two distributions are normalized so that the ratio is equal to one for events with $0 < |d_0| < 3$ mm. For leptons with $|d_0| < 3$ mm, the standard identification uncertainties should be sufficient. This choice of normalization quantifies the uncertainty of the identification for leptons with larger $|d_0|$. The majority of

HNL signal events have small $|d_0|$ and therefore this $|d_0|$ extrapolation uncertainty is usually small. However, for signals with lower masses and longer lifetimes, the uncertainty can be as large as 7%.

Filter

A filter is used to select data events prior to the displaced track and vertex reconstruction (see section 5.7.2). To evaluate the impact of the filter on simulated events, a filter emulator is used. However, there is a mismatch in the signal efficiency due to the different physics objects used by the data and emulated filters. The impact of this uncertainty is evaluated using simulated events that are reconstructed with and without the large-radius tracks. A 3% uncertainty is applied to account for this discrepancy.

Pileup

While pileup is modeled in simulation, the simulated collision conditions are not identical to those experienced during the data collection. A centrally managed tool is used to re-weight signal events and take into account the impact of pileup during the various data periods. This uncertainty is reported to be 3%.

Integrated Luminosity

The uncertainty in the combined integrated luminosity for the data collected between 2015 and 2018 is 2% [85]. This value is obtained using a dedicated luminosity detector [86].

Prompt Lepton Trigger, Reconstruction and Identification

The systematic uncertainties associated with the lepton calibration, identification, reconstruction and trigger efficiencies are implemented using the central analysis tools. These uncertainties are estimated using $Z \rightarrow ll$ and $J/\Psi \rightarrow ll$ events for high- p_T (20-1000 GeV) and low- p_T (3/4.5-20 GeV¹³) leptons respectively. The analysis selection requires one trigger-matched lepton (either prompt or displaced) and the trigger uncertainty is calculated for the trigger-matched lepton in the event. Standard reconstruction, calibration and identification uncertainties are also applied to each of the three leptons in the event, based on the definition of each object in chapter 4.

HNL Cross Section

Two uncertainties are assigned to the HNL cross section calculation in eq. (2.7). The first is a 3% experimental uncertainty on the ATLAS measurement of the $W \rightarrow l\nu$ cross section. The

¹³The minimum low- p_T threshold is 3 GeV for muons and 4.5 GeV for electrons.

second is a 5% theoretical uncertainty on the HNL branching ratio. Previous HNL searches used a parameterization of the mass, mixing angle and lifetime relationship described in Ref. [27] to compute the $N \rightarrow ll\nu$ branching ratio. However, this parameterization is estimated to differ by up to 30% from the most precise calculations to date. In this analysis, branching ratios are analytically computed for each mass and lifetime sample, taking into account perturbative higher-order corrections. These calculations, before including the corrections, are within 1% of the values generated by MADGRAPH. After including higher-order corrections, the differences between the analytic computations and MADGRAPH vary between 3 and 8% for HNLs with masses between 2.5 and 20 GeV. It is known that MADGRAPH neglects these higher-order corrections and is consequently wrong. Nevertheless, a conservative 5% systematic uncertainty is assigned to the HNL branching fraction, $BR(N \rightarrow ll\nu)$.

5.10 Statistical Analysis

A global fit for the signal strength μ is performed using the CL_s prescription [87] and implemented using a statistics package called TRExFitter [88]. A likelihood function, $\mathcal{L}(\alpha) = P(x|\alpha)$, is used to define a probability function for a hypothesis α given an observation x . In this analysis, the probability of observing of n events is tested using a signal plus background hypothesis, $\mu S + \mu_b B$, where μ is the strength of the signal and μ_b is the unconstrained background normalization. The likelihood function is:

$$\mathcal{L}(n|\mu, \mu_b, \vec{\theta}) = \prod_{i \in \text{bins}} P(n_i | \mu S(\vec{\theta}) + \mu_b B(\vec{\theta})) \cdot \underbrace{\prod_{j \in \text{NP}} G(\theta_j)}_{\text{systematics}} \quad (5.42)$$

In this expression, the likelihood of observing n events given an expected number of μS signal plus $\mu_b B$ background events is given by a Poisson distribution P . Here, a signal strength of $\mu = 0$ corresponds to a background-only hypothesis, while $\mu = 1$ corresponds to the nominal signal hypothesis. The expected number of signal events is computed using the simulated samples discussed in section 5.4.2 and the predicted background shape is given by the event shuffling method described in section 5.8.6. The systematic uncertainties described in sections 5.8.6 and 5.9 are accounted for by introducing nuisance parameters $\vec{\theta}$ into the likelihood. This method uses Gaussian probability density functions to constrain the expected number of events and take into account fluctuations that occur due to systematic errors. In addition, the background model includes a normalization factor μ_b that controls the overall strength of the background B . The background prediction for each channel has an independent μ_b whose value is determined in the global fit.

The distribution of the main observable, m_{HNL} , is split into two bins: a CR bin with $20 < m_{\text{HNL}} < 50$ GeV and an SR bin with $m_{\text{HNL}} < 20$ GeV. Including the CR bin in the

likelihood directly constrains the background prediction in the SR, which is allowed to float freely in the fit. The likelihood is maximized to determine the best-fit value for the global signal strength μ , which is shared across all channels that contribute to the signal model.

In the statistical analysis, the profile likelihood ratio is used as the test statistic q_μ . Out of convenience, the negative log of the profile likelihood ratio is used to allow for a minimization instead of a maximization prompt and multiplied by two by convention:

$$q_\mu = -2 \ln \lambda(\mu) \quad (5.43)$$

The profile likelihood ratio is defined as:

$$\lambda(\mu) = \frac{\mathcal{L}(\mu, \hat{\hat{\theta}}_\mu)}{\mathcal{L}(\hat{\mu}, \hat{\theta}_\mu)} \quad (5.44)$$

where $\mathcal{L}$ is the likelihood, μ is the signal strength, $\hat{\mu}$ and $\hat{\theta}$ are estimators for the signal strength and nuisance parameters that maximize the likelihood and $\hat{\hat{\theta}}_\mu$ is a conditional maximum-likelihood estimator for the nuisance parameters given a specific μ . This profile likelihood ratio can take on values between zero and one, where larger values indicate more disagreement between the data and the signal-plus-background hypothesis. The probability of the background fluctuating such that it is equal or larger than the observed number of data events is given by the p-value p_0 , such that:

$$p_0 = \int_{q_{\text{obs}}}^{\infty} f(q_0|0, \hat{\theta}_0) dq_0 \quad (5.45)$$

where $f(q_0|0, \hat{\theta}_0)$ is the probability distribution function in the background-only hypothesis q_0 with $\mu = 0$ and q_{obs} is the observed test statistic. A similar p-value for the signal-plus-background hypothesis q_μ is defined as:

$$p_\mu = \int_{q_{\text{obs}}}^{\infty} f(q_\mu|\mu, \hat{\theta}_\mu) dq_\mu \quad (5.46)$$

Instead of using the p-value p_μ to exclude the signal hypothesis, the CL_s method [87] is used. This method is a conservative frequentist approach that avoids excluding signals where the experiment lacks sensitivity. For example, if the expected number of signal events is very small compared to the background prediction, and the observed number of events has a downward fluctuation, then the signal-plus-background hypothesis can be excluded. This effect is mitigated in the CL_s method by deriving a test statistic from the ratio of the signal-plus-background to the background-only hypothesis. This test statistic equals:

$$\text{CL}_s = \frac{p_\mu}{p_0} \quad (5.47)$$

If the signal-plus-background and background-only hypotheses are well separated (i.e. the experiment has good sensitivity), then $p_0 \gg p_\mu$ and the exclusion will be very similar to the exclusion computed with p_μ only. However, if the two distributions are similar, then the CL_s is approximately one, which reduces the probability of excluding the signal. The 95% CL on the signal strength can be found by varying μ_{Limit} until $\text{CL}_s(\mu_{\text{Limit}}) = 0.05$. Signal models are excluded if $\mu_{\text{Limit}} < 1$.

To test each signal hypothesis, the TRExFitter software package is used to generate ten thousand pseudo-experiments for each value of μ with $\hat{\mu}$ constrained between 0 and 100 to create signal-plus-background ($f(q_\mu|\mu, \hat{\theta}_\mu)$) and background-only ($f(q_0|0, \hat{\theta}_0)$) distributions. From these toy experiments, the p-values p_0 and p_μ are calculated and μ_{Limit} is computed for each of the generated mass and lifetime samples discussed in section 5.4.2. Based on the signal samples that are excluded, limits on $|U|^2$ as a function of m_N can be extracted for the various scenarios listed in table 5.1. A cubic spline interpolation method is used to create smooth contours with $\pm 1\sigma$ and $\pm 2\sigma$ bands in the $|U|^2$ versus m_N plane.

5.11 Results

This section discusses the results of the statistical analysis described in section 5.10. Figure 5.29 shows the signal and background yields before (pre-fit) and after (post-fit) the global fit is performed for the 2QDH model with inverted-hierarchy mixing. Good agreement is observed between the background prediction and the observed data events. This true for all signal models. The pre-fit and post-fit plots for the rest of the models are shown in figs. 5.30 and 5.31 respectively. No significant excess is found in the data and the best-fit for the signal strength μ is consistent with zero in all scenarios. The fit also has no significant deviations from zero for the fitted values of the nuisance parameters. The systematic uncertainties on the shuffled background estimate are the only nuisance parameters that are pulled from their nominal value, but they are all within one standard deviation of the expected value.

In fig. 5.32, the m_{HNL} distribution for the post-fit background and the expected signal normalized to the pre-fit signal strength and the observed events are shown. The pre-fit signal strength normalization is used because, after the global fit is performed, the signal strength is zero and the signal distributions would not be visible.

The post-fit background predictions and the background scale factors μ_b are listed in table 5.16 for the 2QDH model with IH mixing in the “Majorana-limit”. Due to the negligible signal yield found by the fit, the background predictions are essentially identical for the different signal hypotheses. Table 5.16 also shows the background prediction and observed event-yield in the single-muon (single-electron) mixing scenarios for the $\mu\text{--}\mu e$ ($e\text{--}e\mu$) channel. These two channels are shown in the “Dirac-limit” with the additional LNC configuration: $l_\alpha^\pm - l_\alpha^\mp l_\gamma^\pm$. For more details about this LNC configuration see section 5.7.3.

Table 5.16: Post-fit background predictions, observed event-yields and scale factors μ_b with statistical and systematic uncertainties in the control and signal regions. The results are shown for the 2QDH model with inverted-hierarchy mixing in the “Majorana-limit”. The results are also shown for all final states and for the 1SFH Dirac-limit, LNC configuration $l_\alpha^\pm - l_\gamma^\mp l_\gamma^\pm$.

Model	Channel	Signal region		Control region		μ_b
		Background	Observed	Background	Observed	
2QDH (Majorana-limit)	$e-ee$	0.4 ± 0.3	2	3.6 ± 1.8	2	$0.7 \pm_{-0.3}^{0.4}$
	$\mu-ee$	0.2 ± 0.14	1	1.8 ± 1.3	1	$0.5 \pm_{-0.3}^{0.4}$
	$e-e\mu$	0.9 ± 0.4	0	4.1 ± 1.9	5	$0.4 \pm_{-0.2}^{0.2}$
	$\mu-\mu e$	2.8 ± 0.8	2	12.3 ± 3.2	13	$3.0 \pm_{-0.7}^{0.9}$
	$e-\mu\mu$	1.2 ± 0.9	1	2.8 ± 1.6	3	$1.0 \pm_{-0.4}^{0.6}$
	$\mu-\mu\mu$	2.2 ± 1.5	2	8.7 ± 2.9	9	$1.9 \pm_{-0.5}^{0.6}$
1SFH (Dirac-limit)	$e^\pm - e^\mp \mu^\pm$	0.6 ± 0.3	0	2.4 ± 1.4	3	$0.5 \pm_{-0.2}^{0.3}$
	$\mu^\pm - \mu^\mp e^\pm$	1.9 ± 0.6	0	8.1 ± 2.6	10	$2.9 \pm_{-0.8}^{1.0}$




Figure 5.29: Pre-fit (a-b) and post-fit (c-d) plots for the 2QDH model with inverted-hierarchy (IH) mixing in the “Dirac-limit” (left column) and “Majorana-limit” (right column). The signal region (SR) includes events with $m_{\text{HNL}} < 20$ GeV and the control region (CR) includes events with $20 < m_{\text{HNL}} < 50$ GeV.




Figure 5.30: Pre-fit plots for the 1SFH model with muon-only mixing (a-b), electron-only mixing (c-d) and the 2QDH model with normal hierarchy (NH) mixing (e-f) in the “Dirac-limit” (left column) and “Majorana-limit” (right column). The signal region (SR) includes events with $m_{\text{HNL}} < 20$ GeV and the control region (CR) includes events with $20 < m_{\text{HNL}} < 50$ GeV.




Figure 5.31: Post-fit plots for the 1SFH model with muon-only mixing (a-b), electron-only mixing (c-d) and the 2QDH model with normal hierarchy (NH) mixing (e-f) in the “Dirac-limit” (left column) and “Majorana-limit” (right column). The signal region (SR) includes events with $m_{\text{HNL}} < 20$ GeV and the control region (CR) includes events with $20 < m_{\text{HNL}} < 50$ GeV.




Figure 5.32: The m_{HNL} distribution in the SR and CR for the observed data and the post-fit background for the 2QDH model with inverted-hierarchy mixing in the “Majorana-limit”. Simulated signal samples with three different mass points are also shown normalized to the pre-fit signal strength. After the global fit is performed, the signal strength is zero.

The signal strength is found to be zero after the global fit is performed. Therefore, limits are set on the mixing angles squared $|U|^2$ in various HNL scenarios for masses in the approximate range of $3 < m_N < 15$ GeV. The observed limits are within the $\pm 1\sigma$ band of the expected limits for all models.

The observed and expected limits with $\pm 1\sigma$ and $\pm 2\sigma$ uncertainty bands for the 1SFH model with muon-only and electron-only mixing in the "Dirac-limit" are shown in fig. 5.33. Summary plots of the observed 95% CL contours for the four models in the "Dirac-limit" and "Majorana-limit" are also shown in fig. 5.33. The observed and expected limits for the rest models are shown in fig. 5.34. The strongest limits are observed for the 2QDH model with IH multi-flavour mixing. This scenario has equal mixing strength between all neutrino flavours and thus, the combination of all six channels leads to the strongest exclusion limits.




Figure 5.33: The expected (dashed line) and observed (solid line) 95% confidence level (CL) contours for the 1SFH model with muon-only (a) and electron-only (b) mixing in the "Dirac-limit" and the observed 95% CL contours for the four HNL models listed in table 5.1 in the "Dirac-limit" (c) and "Majorana limit" (d). The limits are shown for the mixing angle ($|U|^2$) as a function of the HNL mass (m_N). The regions enclosed by the contours are excluded. The observed contours for the 1SFH models are also shown for the 1997 DELPHI [68] and the 2019 ATLAS [70] results.

The unusual shape of the limit contour is related to the relationship between the mixing angle, lifetime and mass described in eq. (2.8). The analysis is sensitive to models with $c\tau_N$ between 1 and 100 mm. For models with $3 < m_N < 5$ GeV, this lifetime range correspond to mixing angles between 10^{-3} and 10^{-5} . For longer lifetimes models (with smaller mixing angles), there are not enough decays in the detector volume for the analysis to have sensitivity. As the mass increases, models that the search is sensitive to are too short-lived and fail the minimum L_{xy} requirements, resulting in a limit in the sensitivity beyond ~ 15 GeV. The dip in the contour around 5 GeV is a result of the L_{xy} -dependent m_{DV} selection in the heavy-flavour decay veto (see section 5.8.3), which removes events with small m_{DV} and L_{xy} . This selection is required to remove decays from metastable particles, but also limits the available phase space for models with short lifetimes and $2 < m_N < 5$ GeV. This selection was not used in the statistically limited 2019 ATLAS analysis, which explains the shape difference between the two ATLAS results. Regardless, this region of phase space is well-covered by the 1997 DELPHI results, as shown in fig. 5.33 (a).

Observed limits on the square of the HNL to ν_μ mixing angle $|U_\mu|^2$ for the 1SFH model exceed the limits previously set by ATLAS in both the low-mass (~ 1 GeV improvement) and high-mass (~ 4 GeV improvement) region. The results also extend the limits to mixing angles that are approximately three times smaller than the 2019 ATLAS limits. The limits on the square of the HNL to ν_e mixing angle $|U_e|^2$ are set for the first time by ATLAS for a long-lived HNL. The CMS collaboration submitted their first displaced HNL results for single-flavour mixing scenarios [89]. This analysis uses 138 fb^{-1} of data collected by the CMS experiment between 2015 and 2018. The CMS and ATLAS limits are comparable in the muon-only single flavour mixing scenarios, except at large values of $|U_\mu|^2$ where CMS excludes a larger region of phase space. One major difference between the two searches is that CMS has a different track reconstruction workflow, which does not require a minimum L_{xy} selection. This means that CMS can search for prompt and displaced HNL signatures in the same analysis. This tracking workflow also makes it easier to study the heavy-flavour decay background in simulation. The CMS analysis uses simulated MC samples to estimate the heavy-flavour decay background contribution between 0 and 4 cm [90]. The ATLAS analysis instead has a minimum L_{xy} selection of 4 mm and uses an L_{xy} -dependent m_{DV} selection to remove the heavy-flavour decay background, which impacts the signal sensitivity at larger $|U_\mu|^2$ (shorter $c\tau_N$). However, this region of phase space is well-covered by the DELPHI experiment exclusion limits (see fig. 5.2). In the electron-only mixing scenarios, CMS also excludes a similar region at larger values of $|U_e|^2$ that also overlaps with the DELPHI limits. However, at smaller values of $|U_e|^2$, the ATLAS $|U_e|^2$ limits exceed the CMS limits by about a factor of two for $5 < m_N < 14$ GeV. The 139 fb^{-1} ATLAS results also include a new interpretation for realistic 2QDH models with multi-flavour mixing scenarios that assume inverted and normal neutrino-mass hierarchies. To date, this is the first experiment that has directly explored 2QDH models.




Figure 5.34: The expected (dashed line) and observed (solid line) 95% CL contours for the mixing angle, $|U|^2$, in the 1SFH model with muon-only (a) and electron-only (b) mixing in the “Majorana-limit” and the contours for the 2QDH model with NH (c-d) and IH (e-f) mixing in the “Dirac-limit” and “Majorana-limit” as a function of the HNL mass, m_N . The green and yellow bands indicate the ranges of expected limits obtained within ± 1 and ± 2 standard deviations of the median limit reflecting uncertainties in the signal and background yields. The regions enclosed by the contours are excluded.

Chapter 6

ATLAS Inner Tracker Upgrade

6.1 ATLAS Inner Tracker

The LHC is undergoing a series of upgrades in order to produce ten times its current luminosity. For more information about the High-Luminosity LHC (HL-LHC), see section 3.1.5. In order to meet the requirements of this challenging new environment, an all-silicon tracking system, called the Inner Tracker (ITk), will be installed inside the ATLAS solenoid magnet and replace the current inner detector (pixel + SCT + TRT). Figure 6.1 shows a sketch of the ITk. More details about the ITk can be found in Ref. [91].




Figure 6.1: A longitudinal view of the active regions of the ITk detector. The pixel detector is shown in the inner region (smaller radius) and is surrounded by the strip detector. The strip end-cap disks can be seen on either side of the barrel region, coloured with blue and green stripes. Figure taken from [91].

The ITk will consist of a pixel detector near the beam line that is roughly twice the radius and four times the length of the current pixel detector. The ITk-Pixel detector has five barrel layers and five layers of inclined or vertical rings that will provide coverage in the forward region. This will be surrounded by a segmented strip tracking detector covering over three times the silicon area of the current SCT detector. The ITk-Strip detector consists of four

layers in the barrel region and six disks in the end-cap. The work presented in this dissertation focuses on detector quality assurance and the readout for the ITk-Strip detector.

6.1.1 Inner Tracker Strip Detector

The ITk-Strip detector is segmented into four barrel layers with one end-cap region on each side of the barrel containing six end-cap disks. The barrel extends from -1400 to 1400 mm along the z -axis and consists of silicon sensors with 24.1 mm short strips with a pitch of 75.5 μm . In the end-caps, the strip pitch can vary from 70-80 μm . Module pairs are combined to implement stereo angles of 52 mrad in the barrel and 40 mrad in the end-caps [91]. The basic unit of the ITk-Strip detector is a module shown in fig. 6.2. A module consists of a sensor and one or two electrical circuit boards, called hybrids, that house the front-end applied specific integrated circuits (ASICs) called ABCStar and HCCStar. These front-end chips are discussed in more detail in section 6.3. Modules are constructed by gluing the hybrid on the sensors and connecting wire bonds that are 25 μm in diameter between the read-out chips and the silicon strips on the sensors. A power board also sits on the hybrid to supply power to the front-end chips.




Figure 6.2: An expanded view of a strip module in the barrel. Petal modules consist of the same basic components, but with different hybrid and sensor geometries. Figure taken from [91].

The work in this dissertation focuses on strip detectors in the end-cap region. An end-cap disk is made up of 32 petal-shaped wedges. Each petal consists of nine sensors, with six

different geometries labelled R0-R5 that are mounted on each side of a petal core structure. The petal core structure consists of a carbon fibre structural base containing cooling channels, kapton bus-tape that provides electrical connections and silicon modules. A full sketch of the petal is shown in fig. 6.3.




Figure 6.3: A full sketch of a strip petal and the modules labelled R0-R5. The sensors, in blue, are glued onto the core surface shown in yellow. Light green and yellow chips are housed on the green hybrids and glued onto the sensors. The power boards are shown in orange. Figure modified from [91].

6.2 Quality Assurance for ITk-Strip Petals

The ITk-Strip detector needs to operate reliably for ten to fifteen years. Therefore, reliability testing of all the various detector components is an essential part of the detector design and construction program. Quality Assurance (QA) tests study the performance and reliability of components during the research and development phase and during production. In particular, during pre-production, large-scale, multi-module prototypes of ITk-Strip components help demonstrate that the construction of all components meet certain quality thresholds. QA tests are generally “destructive” in nature, which means they involve more extreme stress testing of the components than what would typically be experienced during detector operation. Examples of stress tests include operating at elevated or lowered temperatures, rapid thermocycling and operating at elevated voltages [91]. For the purposes of this dissertation, only the thermocycling QA tests for the petal will be discussed.

Thermocycling is a type of stress test that repeatedly elevates and lowers the temperature in a controlled environment. The goal of thermocycling the petal is to repeatably stress the various components and study whether they fail. Examples of failures for the petal would be object movement (e.g. movement of the sensors or hybrids) or failure of components (e.g.

broken wire bonds or glue failures).

A thermocycling QA program was designed to thermocycle a loaded petal in a climate chamber and measure features, such as sensor bow, sensor positions etc., before and after the thermocycling, to determine their stability. Thermocycling tests are performed in a controlled environment known as a climate chamber. For this test, the climate chamber needed to accommodate the entire petal structure, which is approximately 800 mm in length. Part of the design of this QA program was to procure and test a new climate chamber suited to the needs of the ITk petal production at TRIUMF, Canada's particle accelerator centre. The details of the climate chamber testing can be found in section 6.2.1.

The thermocycling program was initially designed with a smaller prototype version of the petal called a petalet. The petalet is made out of the same materials as a full-sized petal, but only has room for one sensor, called the R0 sensor, on either side of the structure. The design and assembly of the petalet, as well as the thermocycling tests are discussed in section 6.2.2. Once a full-size prototype petal structure was built, the petal was also subjected to the same thermocycling tests. These tests are discussed in section 6.2.3.

6.2.1 Thermocycling Program

During thermocycling, it is important to maintain a dry environment so that there is no risk of water condensing on the detector. For this reason, it was extremely important to find a machine that included a dry air system. Careful validation of the thermocycling program and monitoring of the moisture levels in the chamber was done to ensure the safety of the petal when undergoing temperature cycling tests. This validation included a monitoring system set up to measure the temperature, dew point and relative humidity in the chamber. A sensor measured the internal air temperature of the chamber and the *relative humidity* (RH) of the air, which is defined as the ratio of the amount of moisture present in the system relative to the amount of moisture that would be present if the air was completely saturated. Since the latter amount is dependent on the temperature of the system, it is more useful to consider the dew point of the system. The *dew point* (DP) is the temperature below which, given the moisture in the chamber, water droplets will start to form. Given the temperature (T) and relative humidity (RH), the dew point temperature in the chamber can be calculated using the Magnus formula shown in eq. (6.1). For the temperature range from 45°C to 60°C, the Magnus parameters are given by $\beta = 17.62$ and $\lambda = 243.12^\circ\text{C}$ [92].

$$DP(T, RH) = \frac{\lambda(\ln(\frac{RH}{100}) + \frac{\beta T}{\lambda + T})}{\beta - (\ln(\frac{RH}{100}) + \frac{\beta T}{\lambda + T})} \quad (6.1)$$

In addition to the sensor inside the chamber, a second sensor was used to measure the pressure differential between the inside and the outside of the chamber ($\Delta P = P_{\text{in}} - P_{\text{out}}$),

as well as the temperature and RH outside of the chamber. This external monitoring system was used to record any changes in the external temperature/humidity that might influence the internal chamber environment. The pressure monitoring was also used to see if there were points in the thermocycling program where the pressure differential changed from positive to negative. This change in pressure could result in moist air from the outside being sucked into the system, so we wanted to see if there were places in the thermocycling program that were susceptible to these types of pressure changes.




Figure 6.4: The final validated thermocycling program (a) with a ramp rate of $1.33^{\circ}\text{C}/\text{min}$ and soak times of 15 hours and 30 min at 21°C and $\pm 40^{\circ}\text{C}$ respectively. This plot shows the temperature and RH for the full 46-hour program. The red curves show the measured temperature inside the climate chamber and the blue curves show the RH inside the chamber. The temperature and RH curves (b) for the first thermocycle where the RH spike is the largest. The plot in (b) also indicates the dew point inside the chamber (dashed red line) and the measured pressure differential (green line), which is defined as $\Delta P = P_{\text{in}} - P_{\text{out}}$.

When testing the thermocycling program, the *soak time* (how long the system sits at the same temperature), and *temperature ramp time* (how quickly you change the temperature of the system) had the most impact on the dryness of the system. In order to prevent condensation from forming on and damaging the detector, it was decided to find a program that would maintain a dew point at least 10°C below the air temperature inside the chamber. Before finalizing the thermocycling program, it was important to optimize the soak and ramp times to achieve a program that met the dew point criterion and took a reasonable amount of time to complete 10 cycles between $\pm 40^\circ\text{C}$. The 46-hour program, shown in fig. 6.4), was established to meet the thermal cycling test and detector safety requirements.

6.2.2 Thermocycling the Petalet

The petalet is a prototype version of the petal structure that was designed for initial testing. The overall size of the petalet is 15×18 cm. The petalet has the same components as the petal and consists of a carbon fibre base, a kapton bus-tape, and space for two R0 modules. However, due to the limited availability of components at the time, the prototype version we tested, only had one module mounted on side A. Stable reference markers were also glued onto the structure to define a global coordinate system on the petalet. The location of the reference markers are shown in fig. 6.5.




Figure 6.5: The location of the stable markers glued directly to the carbon surface labeled 1-5.




Figure 6.6: The petalet hanging in the climate chamber right before thermocycling.

To investigate the stability of the structures over time, various measurements of the locations of the components and properties of the sensor were recorded. Features on the sensor and hybrids are sometimes referenced to as fiducial markers since they are used by imaging machines to locate the objects when loading them on the petal. Measurement scans were done using a coordinate measuring machine (CMM) to locate features on the sensor, hybrid and bus-tape. Next, the petalet was placed in the climate chamber and the thermocycling program (described in section 6.2.1) was run. For the test with the petalet, the structure was hung from the wire rack in the climate chamber using copper beryllium alloy wire with a thickness of 180 μm (see fig. 6.6). After the program was finished, the measurements were repeated and the results were examined to study any changes that occurred during thermocycling.

Measurements of the sensor flatness and bow are shown in table 6.1. The flatness is measured by scanning across the surface of the sensor and taking x , y , z measurements at each point. The flatness is then reported by the CMM as the width of a zone bounded by two parallel planes that contains all the measured data points of the surface. The sensor bow, given by eq. (6.2), is also a measure of how flat a sensor is.

$$\text{bow} = \frac{1}{4}(\sum \text{corners}_z) - \text{centre}_z \quad (6.2)$$

Scans of the surface on side A and B before and after thermocycling are shown in fig. 6.7. These plots show how the surface shape of the sensor changed before and after thermocycling. The flatness and bow measurements (for the R0 sensor on side A) are also recorded in the text on the right-hand side of these plots, along with additional sensor measurements, such as the sensor length, radius etc. We observed that the flatness on both sides and sensor bow on Side A are the same before and after thermocycling.

Table 6.1: Measurement results for flatness of Side A and B, as well as the bow for the sensor on side A.

Measurement Type	Before (μm)	After (μm)
Flatness (Side A)	131 ± 3.78	131 ± 3.78
Sensor Bow (Side A)	37.2 ± 0.54	37.6 ± 0.54
Flatness (Side B)	390 ± 3.78	389 ± 3.78




Figure 6.7: Plots of the surface scans of the sensor on side A and the bus-tape surface of side B before and after thermocycling. The flatness, bow and other sensor measurements are also shown in the text on the right-hand side of the plots.

The locations of the features on the petalet are measured using the coordinate system defined by the stable reference markers. Defining a vector from the origin of the coordinate system to each feature as $\vec{R}$, we measure $\Delta\vec{R} = (\vec{x}' - \vec{x}, \vec{y}' - \vec{y})$, where unprimed coordinates are measured before thermocycling and the primed coordinates are measured after. To visualize the displacement, a plot the displacement vectors ($\Delta\vec{R}$) is shown in fig. 6.8. In this visualization we are looking to see if the displacements of the features have a common direction that might arise, for example, from gravity pulling uniformly on the module during thermocycling. The magnitudes of the displacements were on the order of a few μm with no common direction and consistent with no movement within the measurement uncertainty of around 3 μm . Based on our analysis of the results all components on the petalet passed the thermocycling stress test.




Figure 6.8: Plot of the displacement vectors ($\Delta\vec{R}$) for the features on side A, including the two hybrids in purple, the sensor in blue and the bus-tape in brown. The vector magnitudes are uniformly scaled by a factor of 2000 to visually examine if there is a common direction for the displacements.

6.2.3 Thermocycling the Semi-Electrical Petal

While the petalet was a useful structure to design and test the thermocycling program because it was a small-scale object that was more readily available, the goal of the thermocycling program was to stress test large-scale objects that are as similar as possible to the final production design. The structure used in the full-size petal tests is known as the first semi-electrical petal (see picture in fig. 6.9). The semi-electrical petal is fully loaded on both sides and consists of the following modules:

- R5: Dummy sensors that include lithography, reference markers and bond pads
- R4: Dummy sensors with no lithography or reference markers; module was built with cracked dummy sensors
- R1-R3: Dummy sensors that include lithography, reference markers and bond pads
- R0: Fully electrical sensors




Figure 6.9: A photo of the first semi-electrical petal. This petal was the prototype used in the thermocycling tests discussed in this section. This photo was taken by the TRIUMF communications team.

Petal measurements before and after thermocycling were performed using the CMM in a dedicated measurement frame. This frame is constructed out of aluminum and consists of a flat surface where three petal locators rest. These locators are found on the outer left and right edges the petal and on the inner edge of the petal. During assembly, these locators are what will hold the petals to the assembly structure to construct the end-cap disks. The three petal locators are then secured in the measurement frame using 3D printed locking pins. A

coordinate system is set up using the petal locators and the two petal fiducial markers on the inner and one of the outer petal locators. The petal fiducials are tiny circles in the carbon fibre structure that are used during loading to construct a coordinate system that can be used to precisely place the modules on the carbon fibre structure. The x - and y -axes are defined using the petal fiducials and the x - y plane is defined as the surface of the petal locators. The origin is set to be the inner petal fiducial. A sketch of the coordinate system used for the petal measurements is found in fig. 6.10.






Figure 6.10: The petal coordinate system. The x - and y -axes are defined using the two petal fiducials on the inner and outer petal locators. The origin is defined as the inner fiducial and the x - y plane is defined as the flat surface of the petal locators.

The CMM manual tool is used to measure different markers on the sensors, hybrid and bus-tape on both sides of the petal. There are nine sensors per side and a total of five measurements are taken per sensor, with one in each corner, as well as a point in the centre of the sensor to be used for bow measurements. On each side the petal, there are a total of 13 hybrids with four measurements taken per hybrid. Finally, eight points on the kapton bus-tape are also


measured. In total, 210 measurements are recorded. Pictures of the different features that were measured on the petal are shown in fig. 6.11.




(a) electrical sensor F




(b) dummy sensor F



(c) hybrid F



(d) hybrid circle F



(e) bus-tape circle F

Figure 6.11: Pictures of the different fiducial markers on the sensors, hybrids and bus-tape. These markers are measured by the CMM to determine the position of the various components on the semi-electrical petal.

It should also be noted that while measurements were taken on each of the modules, R2 in particular had some fiducials obscured and others that were difficult for the CMM to measure. The R4 modules also have dummy mechanical sensors that did not have any fiducial markers and therefore these sensors could not be measured. As a result, these modules are ignored

and so not all 210 measurements are included in the analysis.


The sources of measurement uncertainty include the CMM measurement uncertainty for manual tool and circle fitting and the petal inner fiducial position (x_0, y_0) that is used to establish the coordinate system. The measurement uncertainty for each feature is estimated as the standard deviation (σ) from ten repeated measurements of the x - y position. The CMM measurement uncertainties and the petal fiducial uncertainty can be found in table 6.2. The fiducial markers on the dummy sensors were not as well defined as the other sensors so the sensor fiducial markers are recorded as two different uncertainties.

Table 6.2: Measurement uncertainties for the features on the semi-electrical petal.

Feature Type	σ_x (μm)	σ_y (μm)
F on hybrid	1.53	1.37
Circle on hybrid	2.73	1.89
F on sensor	1.53	1.11
F on dummy sensor	3.48	3.06
Circles on bus-tape	3.97	3.81
Petal fiducial	0.4	0.8



(a) Support bar



(b) 3D locking pins

Figure 6.12: The support bar in the vertical orientation in the climate chamber. The three attachment points on the bar are where the three petal locators are attached and secured to hang the petal in the chamber.

Since the petal is heavier than the petalet structure, the beryllium alloy wire used to hang the petalet in the climate chamber did not support the weight of the full-sized petal. Instead, a metal bar was built to hang the petal in the climate chamber. This bar can be hung horizontally or vertically to allow the petal to hang in different orientations. 3D printed locking pins (similar to the pins used in the measurement frame) were used to fix the petal to the support bar (see fig. 6.12). Pictures of the semi-electrical petal hanging in the climate change before and during the thermocycling test can be found in fig. 6.13.




Figure 6.13: The petal hanging in the climate chamber before and during thermocycling.

Immediately after thermocycling, the petal was brought back to the clean room and placed in the measurement frame to measure the location of the features. Initially, the analysis of the measurements taken before and after showed worrying $3\sigma \Delta R$ displacements that were consistent with movement due to gravity. This can be seen in the arrow plot of one of the R5 sensors on side A in fig. 6.14. The arrows indicate the direction of motion of the fiducial marker marker, while the magnitudes of the displacements are shown in the coloured boxes.

A statistical analysis of the roughly 100 different points on each side of the petal was also done. The “pull” of two measurements is defined as $(x_1 - x_2)/\sigma_x$, where σ_x is the uncertainty on the two measurements. For a set of measurements that are consistent with no movement, the pull plots should have a Gaussian distribution centred at zero with width one. These pull plots for the x and y measurements of the features on the petal are shown in fig. 6.15.




Figure 6.14: Arrow plot and ΔR displacements of the four sensor fiducials (in blue) and the four fiducials on the hybrid (in purple) for one of the R5 sensors on side A. The arrows in the plot are magnified by a factor of 2000 to visually see the direction of motion. The magnitudes and uncertainties of the displacements are shown in the boxes.




Figure 6.15: Statistical analysis of the roughly 100 different points measured on Side A of the petal. The pull in the y -direction ($\Delta y/\sigma_y$) and the pull x -direction ($\Delta x/\sigma_x$) are shown immediately after thermocycling and the measurements are repeated 48 hours and four weeks later. After 48 hours, the mean of the pull is closer to zero, which indicates that after some time the structures relax and are more consistent with no movement after thermocycling.

Immediately after thermocycling, the mean of the y -pull was negative, which indicates that, on average, the measurements were trending down in the y -direction. This motion was especially worrying since it was consistent with the direction of gravity when the petal was hanging in the climate chamber during thermocycling. However, the measurements were repeated 48 hours later and the mean of the pull in the y -direction was closer to zero indicating that the structures seemed to be relaxing back into the position that they occupied prior to thermocycling. The measurements were repeated four weeks later and the distributions were found to be more consistent with zero movement.

We concluded that moisture effects were the cause of the displacements observed immediately after thermocycling. In the climate chamber, the petal is kept in very dry storage. Dry air is constantly pumped into the chamber to maintain a low dew point and prevent moisture buildup on the detector (see discussion of dew point in section 6.2.1). This removal of moisture from the environment can cause structures on the petal to shrink. However, upon return to the clean room environment, where the measurements are done with the CMM, the structures re-expand and return to their original positions.

This moisture effect theory was further confirmed with tests performed with the petal in a horizontal orientation in the climate chamber. Similar order of magnitude displacements along the long axis of the petal were observed immediately after thermocycling. Since the petal was hanging horizontally, these displacements were no longer consistent with movement due to gravity. The conclusion from these tests is that nothing untoward happened to the petal after 10 thermal cycles and while stressing the petal, we observed a well-known phenomenon that drying out the petal can result in small displacements of the various structures.

As mentioned at the start of this section, QA tests are designed to stress the structures beyond what would typically be experienced during detector operation. Since the structures relaxed back to their original position, and no permanent damage occurred, even after being subjected to a much drier environment than we expect to operate the detector in, this test was considered a success.

100 Thermocycles Test

The goal of these tests is to thermally stress the petal to determine the reliability of its design and construction. Since the petal passed the initial 10 thermal cycles stress test, a subsequent test was run where the climate chamber program was set to run 100 thermal cycles instead of 10. After 100 thermal cycles, visual inspection of the petal showed that one of the hybrids on the R4 modules has bowed significantly. The glue under this hybrid had failed, causing the structure to bow and a crack to develop on the right-hand side of the R4 sensor (see fig. 6.16).

The analysis of the measurements taken before and after the 100 cycles test was consistent with this crack and bowing occurring during the test. Figure 6.16 shows the arrow plots for this sensor and the large ΔR and Δz displacements that occurred during thermocycling. The pull plots in fig. 6.18 show that the rest of the measured features, not permanently damaged by the bowing and cracking, exhibited the previously observed contraction after removing the petal from the dry climate chamber. After two weeks the structures relaxed back to their original positions and distributions were consistent with no movement, indicating that all structures (except for R4) passed the long-term stress test. The outliers in the pull plots were verified to be measurements from the R4 hybrid where the failures (cracking and bowing) occurred. A similar crack also occurred on side B on another R4 module.

Although the 100 thermal cycles test observed two failures of the R4 modules, these modules were always “compromised”. They were built using dummy sensors that were cracked even before the semi-electrical petal underwent the thermocycling stress tests. This result is not all that surprising, since these R4 modules were not built using the final production sensors.

We concluded from this test that all modules that did not use dummy sensors (i.e. all modules except R4) passed the long-term stress test. However, this test also demonstrated that it would be useful to perform a long-term stress test (similar to this 100 thermocycle test) on a fully electrical petal with no dummy/cracked sensors.




Figure 6.16: The crack on the R4 sensor and the bowing of a hybrid on the same module. Photos are taken immediately after thermocycling the semi-electrical petal 100 times.







Figure 6.17: Arrow plot and displacements (ΔR and Δz) of the four fiducials (in purple) for one of the hybrids on side A. The arrows in the plot are scaled by a factor of 2000 to visually see the direction of motion. The magnitudes and uncertainties of the displacements are shown in the boxes.




(a) y -pull immediately after thermocycling



(b) y -pull two weeks later



(c) x -pull immediately after thermocycling



(d) x -pull two weeks later

Figure 6.18: Statistical analysis of the roughly 100 different points measured on Side A of the petal. The pull in the y -direction ($\Delta y/\sigma_y$) and the pull x -direction ($\Delta x/\sigma_x$) are shown immediately after thermocycling and two weeks later. After two weeks, the central distribution is more consistent with zero, except for outliers with large pulls that are from the measurements of the hybrid on the module where the failure occurred.

6.3 ITk-Strip Readout

The ITk-Strip detector will consist of approximately 60 million channels that need to be read out at a frequency of 1 MHz [91]. This provides a challenging readout requirement on the order of 10^{13} channels per second. The current inner tracker has readout rate of roughly 10^{10} channels per second. This added complexity required a re-design of the ATLAS data acquisition system (DAQ). The current ATLAS DAQ (fig. 6.19 (a)) consists of front-end electronics and downstream readout systems that are custom to each detector in ATLAS. The new DAQ system (fig. 6.19 (b)) will remove some of the dependency on custom electronics and replace them with a single board common to each detector that uses commercially available field-programmable gate arrays (FPGAs). These so-called Front-End Link Exchange (FELIX) cards are advantageous because they are detector independent and provide a scalable architecture suited to a large complicated DAQ system.




Figure 6.19: Current and HL-LHC era DAQ systems for the ATLAS experiment. The FELIX cards will replace the current readout drivers (RODs) and readout systems (ROSs) used in the current DAQ system. Figure adapted from [93].

In addition to the FELIX cards, the ITk-Strip readout also includes front-end ASICs that are designed specifically for the detector. There are three main ASICs housed on the silicon modules. They include the ABCStar, the HCCStar and the AMAC. Only two of the three (the HCCStar and AMAC) chips are required to communicate with the FELIX cards. Their functionality and production chip names are summarized in table 6.3. Depending on the layout of the detector, between nine and eleven ABCStar chips are connected to a single HCCStar chip. The HCCStar packages the data and sends the information downstream. The HCCStar is the only communication link between the ABCStars and the software system and therefore the HCCStar is also responsible for configuring registers and reporting back the status of the ABCStar chips to the central DAQ system.

Information from many HCCStar chips on a single structure, such as a petal, is merged

together on an end-of-structure (EoS) card and converted to an optical signal. The EoS card also contains other chips (IpGBT and VL+) that convert the data into a custom data protocol and into an optical signal. The diagram in fig. 6.20 focuses on the upstream path where the data from the silicon sensor is buffered and read out by FELIX. However, in reality, there are also streams of data which flow in the other direction and pass information from the DAQ system through FELIX to the front-end chips. The data lines connecting the AMAC and FELIX are also not represented in this diagram.

Table 6.3: Overview of the custom made ASICs for the ITk Strip Detector [91].

Acronym	Full Name	Basic Functionality	Production Chip
ABC	ATLAS Binary Chip	Converts incoming charge signal into hit information.	ABCStar
HCC	Hybrid Controller Chip	Interface between ABC and bus-tape.	HCCStar
AMAC	Autonomous Monitor and Control Chip	Provides monitoring and interrupt functionality.	AMAC-II




Figure 6.20: A diagram of the ITk-Strip readout showing how all the custom ASICs (ABCStar, HCCStar and AMAC) are connected to each other and merged in the end-of-structure (EoS) card before being converted to an optical signal and sent to FELIX.

6.3.1 HCCStar Emulator

Although the final ITk-Strip detector will consist of 60 million channels, this number of working readout channels will not be ready until close to the detector installation date. To understand how all the information will be integrated together, various demonstrators have been built to test the front-end electronics. Using these systems, we can gain a better understanding of the ITk-Strip readout and prepare for the eventual operation of the detector. To gain experience with a large-scale ITk-Strip readout system, a firmware emulator was designed to mimic both the detector and the front-end chips.

The logic for the HCCStar emulator was designed and tested in a simulation and then

integrated into an FPGA. The emulator generates ITk-Strip-like data packets that are sent through optical fibres to a FELIX card. With the FPGA resources on the current card used to build the emulator, up to 120 HCCStar chips can be emulated, which is equivalent to approximately 300,000 channels. With more cards, this firmware can be used to emulate the number of channels on the order of the whole ITk-Strip detector.

The HCCStar emulator was designed to mimic the physics data-taking mode of the real HCCStar chips. In physics mode, the primary function of the HCCStar is to transmit the data collected by the silicon sensors when a trigger command is received. During detector operation, this trigger command is sent to the front-ends from the central DAQ system when the decision has been made that a specific collision event should be saved for further analysis (see section 3.2.6 for more details about the ATLAS trigger system). This means the emulator needs to respond and send back data, in the same format as a real HCCStar chip, when it receives a trigger command.




Figure 6.21: A sketch of the HCCStar emulator and LCB wrapper integrated into the FELIG data generation structure.

A schematic of the main firmware blocks used in the HCCStar emulator can be found in fig. 6.21. The blocks specific to ITk-Strip data structures are integrated into an existing data generation infrastructure that converts the emulated data into optical signals. This firmware is called the front-end link generator (FELIG) and was designed to test the functionality of the FELIX cards. The main unit of the firmware is called the FELIG lane wrapper. The logic contained inside this wrapper is duplicated up-to 24 times on separate optical channels. The

data in and out of the lane wrapper are sent through a GBT¹⁴ block. This block converts a custom CERN data protocol into bits of information that are passed to the rest of the data generation logic. This block is also responsible for encoding the data into the GBT protocol to send back to FELIX.

There are two main blocks of firmware that are specific to the data format used in the ITk-Strip DAQ system. The first is the LCB (L0a, Control and BCR) wrapper. This firmware functionality is also contained in the real HCCStar ASIC and is responsible for receiving triggers and commands sent via FELIX and decoding them into a format that the HCCStar understands. In physics mode, the three main pieces of required information are the L0a trigger, tag and BCR. The L0a trigger is analogous to the L-1 trigger used in the current DAQ system and is the name of the hardware-based trigger that will be used in ATLAS during the operation of the HL-LHC. The tag is a data event identifier that is used to uniquely locate event data in a buffer. The tag is passed to the front-ends by the central DAQ system to tell the front-ends which data event to send back. This tag is then packed and sent along with the data to ensure that the correct collision event is received. The BCR is the bunch counter reset. The LHC is filled with bunches of protons that collide every 25 ns (see section 3.1.1 for more details about the LHC filling scheme). Every detector in ATLAS is sent a global 40 MHz clock that is synchronized with these bunch crossings. The front-ends use this clock to count the bunch crossings. This associates each bunch crossing with a number and enables the front-ends to uniquely identify events. However, these counters are not infinite, so a BCR signal is sent to reset all these counters to align them with the LHC orbit.

The second ITk-Strip specific firmware is the HCCStar emulator block. This firmware was written to mimic the data transmission of a real HCCStar ASIC. This block receives the L0a, tag and BCR from the LCB wrapper and builds data events. A block memory (BRAM) is used to store emulated detector hit information that is loaded into memory by the software. Once the data packets have been built, they are encoded and sent back to FELIX.

The HCCStar emulator has been successfully operated using ITk software. The software is able to send triggers and receive data back from the emulator. Custom data packets can be loaded into the BRAM to test various aspects of the DAQ system. Since the data generator is scalable, the emulator can be used to mimic multiple HCCStars in one demonstrator system. This setup can be used to perform large-scale tests of the ITk-Strip DAQ system and prepare the system for the eventual operation.

¹⁴GigaBit Transceiver (GBT) architecture is a data protocol developed at CERN to be used in data transmission for experiments that require high data rate links and electronic components capable of sustaining high radiation doses. [94]

Chapter 7

Conclusion

The SM is a highly successful theory that describes fundamental particles and their interactions. With the discovery of the Higgs boson at the LHC in 2012, the final piece of the SM was found [95]. However, experimental evidence suggests that the SM is not a complete description of our universe. For example, it does not predict the existence of dark matter, nor does it explain the observations of neutrino flavour oscillations. Since the detection of the Higgs boson, no new fundamental particles have been discovered. Searches for BSM particles are therefore essential to the LHC physics program.

This dissertation presented a search for a new particle that provides a mechanism which could explain the very light neutrino masses. This search for displaced HNLs was performed using 139 fb^{-1} of proton-proton collision data collected with the ATLAS detector. No significant excess of events were observed. Constraints on the mixing angles between HNLs and SM neutrinos were placed in the challenging long-lived region of phase space. Models with mixing angles between $10^{-3} - 10^{-7}$ were excluded in different HNL scenarios for masses in the approximate range of $3 < m_N < 15 \text{ GeV}$. These results extended the reach of previously published ATLAS results on HNL mixing with muon neutrinos and placed the first ATLAS constraints on the mixing with electron neutrinos for a long-lived HNL. New interpretations of the results were also included for realistic multi-flavour mixing scenarios that are consistent with neutrino flavour oscillations.

In the next LHC data taking period (Run 3), ATLAS LLP searches will benefit from the recent optimization of the LRT software. Until now, the reconstruction of displaced tracks required significant computing resources and an extensive amount of disk space to record all the track information. The LRT software has now been optimized to ensure that displaced tracks can be reconstructed using the available computing resources. For the design pileup in Run 3, the computer processing time has been reduced by a factor of two [96]. This means that all ATLAS data will now be available for LLP searches to analyze. This software optimization will enable searches to significantly improve their signal to background ratios and to design better validation and control regions. This new workflow will also enable analysis designs that can bridge the gap between prompt and displaced searches.

For the displaced HNL search, this will simplify the study of the backgrounds. This opens the possibility of estimating heavy-flavour decay backgrounds using MC simulations and improving the sensitivity to low-mass HNLs with shorter lifetimes. Future searches can also be

extended by considering semi-leptonic HNL decay channels. Thus, the combination of Run 3 data with the existing dataset has the potential to extend ATLAS's sensitivity to displaced HNLs.

However, if HNLs weakly interact with SM neutrinos, their observation will require significantly more data. Thus, the quest for new physics will heavily rely on the successful operation and upgrades of the accelerator and detectors. This dissertation also presented detector performance and readout research for the ATLAS ITk. Installing new detectors with enhanced tracking performance will enable ATLAS to take full advantage of the LHC luminosity upgrades. A test was developed to study the dimensional stability of the ITk silicon strip detectors. A large-scale object (petal) underwent temperature cycling to stress the detector components. These tests showed that the petal did not have any points of failure, such as movement of the silicon sensors with respect to the support structures. The primary function of the ITk is to reconstruct tracks in dense environments during the operation of the HL-LHC. The dimensional stability of the detector components is therefore crucial to its design.

The ITk-Strip detector will have approximately 60 million channels that need to be read out at a rate on the order of 10^{13} channels per second. To study the demands of this challenging readout requirement, a firmware emulator was developed to mimic both the detector and the readout system. This firmware and software system sends ITk-Strip-like data packets and runs test scans to check the physics performance of the detector. Large-scale setups with the real detector components will not be ready until the detector is built. The emulator demonstrated the feasibility of the readout system without requiring real detector components.

Bibliography

- [1] D. Griffiths. *Introduction to Elementary Particles*. Wiley, Weinheim, 2 edition, 2008.
- [2] S. N. Gninenko, D. S. Gorbunov, and M. E. Shaposhnikov. Search for GeV-scale sterile neutrinos responsible for active neutrino oscillations and baryon asymmetry of the Universe. *Adv. High Energy Phys.*, 2012:718259, 2012.
- [3] A. Eddington F. Dyson and C. Davidson. IX. A determination of the deflection of light by the sun’s gravitational field, from observations made at the total eclipse of May 29, 1919. *Philosophical Transactions of the Royal Society*, 220:291–333, 1920.
- [4] F. Zwicky. Republication of: The redshift of extragalactic nebulae. *General Relativity and Gravitation*, 41(1):207–224, 2009.
- [5] S. Smith. The Mass of the Virgo Cluster. *Astrophysics Journal*, 83:23, 1936.
- [6] N. Aghanim, Y. Akrami, M. Ashdown, J. Aumont, C. Baccigalupi, M. Ballardini, A. J. Banday, R. B. Barreiro, N. Bartolo, and et al. Planck 2018 results. *Astronomy & Astrophysics*, 641:A6, 2020.
- [7] G. Bellini, L. Ludhova, G. Ranucci, and F. L. Villante. Neutrino oscillations. *Adv. High Energy Phys.*, 2014:191960, 2014.
- [8] ATLAS Collaboration. Standard Model Summary Plots June 2021. Technical report, 2021.
- [9] J. C. Pati and A. Salam. Lepton number as the fourth ”color”. *Phys. Rev. D*, 10:275–289, 1974.
- [10] R. N. Mohapatra and J. C. Pati. ”Natural” left-right symmetry. *Phys. Rev. D*, 11:2558–2561, 1975.
- [11] R. N. Mohapatra and J. C. Pati. Left-right gauge symmetry and an ”isoconjugate” model of CP violation. *Phys. Rev. D*, 11:566–571, 1975.
- [12] G. Senjanovic and R. N. Mohapatra. Exact left-right symmetry and spontaneous violation of parity. *Phys. Rev. D*, 12:1502–1505, 1975.

- [13] J. Schechter and J. W. F. Valle. Neutrino Decay and Spontaneous Violation of Lepton Number. *Phys. Rev.*, D25:774, 1982.
- [14] J. Schechter and J. W. F. Valle. Neutrino Masses in $SU(2) \times U(1)$ Theories. *Phys. Rev.*, D22:2227, 1980.
- [15] R. N. Mohapatra and G. Senjanović. Neutrino mass and spontaneous parity nonconservation. *Phys. Rev. Lett.*, 44:912–915, 1980.
- [16] M. Drewes, D. Garbrecht, D. Gueter, and J. Klaric. Testing the low scale seesaw and leptogenesis. *JHEP*, 08:018, 2017.
- [17] T. Asaka, S. Blanchet, and M. Shaposhnikov. The ν MSM, dark matter and neutrino masses. *Phys. Lett.*, B631:151–156, 2005.
- [18] T. Asaka and M. Shaposhnikov. The ν MSM, dark matter and baryon asymmetry of the universe. *Phys. Lett.*, B620:17–26, 2005.
- [19] S et al. Alekhin. A Facility to Search for Hidden Particles at the CERN SPS: the SHiP physics case. *Rept. Prog. Phys.*, 79(12):124201, 2016.
- [20] J. Tastet, O. Ruchayskiy, and I. Timiryasov. Reinterpreting the ATLAS bounds on heavy neutral leptons in a realistic neutrino oscillation model, 2021.
- [21] I. Esteban, M.C. Gonzalez-Garcia, M. Maltoni, T. Schwetz, and A. Zhou. The fate of hints: Updated global analysis of three-flavor neutrino oscillations. *JHEP*, 2020(9), 2020.
- [22] J. Kersten and A. Smirnov. Right-Handed Neutrinos at CERN LHC and the Mechanism of Neutrino Mass Generation. *Phys. Rev. D*, 76:073005, 2007.
- [23] J. Tastet and I. Timiryasov. Dirac vs. Majorana HNLs (and their oscillations) at SHiP. *JHEP*, 04:005, 2020.
- [24] M. Drewes, J. Klarić, and P. Klose. On lepton number violation in heavy neutrino decays at colliders. *JHEP*, 11:032, 2019.
- [25] G. Anamiati, M. Hirsch, and E. Nardi. Quasi-Dirac neutrinos at the LHC. *JHEP*, 10:010, 2016.
- [26] M.C. Gonzalez-Garcia, A. Santamaria, and J.W.F. Valle. Isosinglet-neutral heavy-lepton production in z-decays and neutrino mass. *Nuclear Physics B*, 342(1):108–126, 1990.

- [27] M. Gronau, C. N. Leung, and J. L. Rosner. Extending limits on neutral heavy leptons. *Phys. Rev. D*, 29:2539–2558, 1984.
- [28] K. Bondarenko, A. Boyarsky, D. Gorbunov, and O. Ruchayskiy. Phenomenology of GeV-scale heavy neutral leptons. *JHEP*, 2018(11), 2018.
- [29] L. R. Evans and P. Bryant. LHC Machine. *Journal of Instrumentation*, 3:S08001. 164 p, 2008.
- [30] ATLAS Collaboration. The ATLAS Experiment at the CERN Large Hadron Collider. *Journal of Instrumentation*, 3:S08003. 437 p, 2008.
- [31] CMS Collaboration. The CMS experiment at the CERN LHC. The Compact Muon Solenoid experiment. *Journal of Instrumentation*, 3:S08004. 361 p, 2008.
- [32] ALICE Collaboration. The ALICE experiment at the CERN LHC. A Large Ion Collider Experiment. *Journal of Instrumentation*, 3:S08002. 259 p, 2008.
- [33] LHCb Collaboration. The LHCb Detector at the LHC. *Journal of Instrumentation*, 3(LHCb-DP-2008-001. CERN-LHCb-DP-2008-001):S08005, 2008.
- [34] E. Mobs. The CERN accelerator complex. <https://cds.cern.ch/record/2197559>, 2016.
- [35] R. Bailey and P. Collier. Standard Filling Schemes for Various LHC Operation Modes. Technical Report LHC-PROJECT-NOTE-323, 2003.
- [36] M. Thomson. *Modern particle physics*. Cambridge University Press, Cambridge, 2013.
- [37] ATLAS Collaboration. Luminosity pulic results run-2. <http://aiweb.techfak.uni-bielefeld.de/content/bworld-robot-control-software/>, 2019.
- [38] HL-LHC Collaboration. HL-LHC Project Schedule. <https://project-hl-lhc-industry.web.cern.ch/content/project-schedule/>, 2021. Visited on Oct. 11, 2021.
- [39] HL-LHC Project. The ATLAS Experiment at the CERN Large Hadron Collider. *Journal of Instrumentation*, 3:S08003. 437 p, 2008. Also published by CERN Geneva in 2010.
- [40] ATLAS Collaboration. ATLAS inner detector: Technical Design Report. Technical report, 1997.

- [41] ATLAS Collaboration. ATLAS Insertable B-Layer Technical Design Report. Technical Report CERN-LHCC-2010-013. ATLAS-TDR-19, 2010.
- [42] ATLAS TRT Collaboration.
The ATLAS Transition Radiation Tracker Proportional Drift Tube: Design and Performance. *Journal of Instrumentation*, 3(02):P02013–P02013, 2008.
- [43] ATLAS Collaboration. Performance of the ATLAS Transition Radiation Tracker in Run 1 of the LHC: tracker properties. *JINST*, 12:P05002, 2017.
- [44] ATLAS Collaboration. ATLAS liquid-argon calorimeter: Technical Design Report. Technical report, 1996.
- [45] C. Grupen and B. Shwartz. *Particle Detectors; 2nd ed.* Cambridge Monographs on Particle Physics, Nuclear Physics, and Cosmology. Cambridge University Press, Cambridge, 2008.
- [46] ATLAS Collaboration. ATLAS muon spectrometer: Technical design report. 1997.
- [47] ATLAS Collaboration. Operation of the ATLAS trigger system in Run 2. *Journal of Instrumentation*, 15(10):P10004–P10004, 2020.
- [48] ATLAS Collaboration. Athena. <https://doi.org/10.5281/zenodo.2641997>, 2019.
- [49] S. Agostinelli et al. GEANT4: A simulation toolkit. *Nucl. Instrum. Meth.*, A506:250–303, 2003.
- [50] ATLAS Collaboration. The ATLAS Simulation Infrastructure. *The European Physical Journal C*, 70(3):823–874, 2010.
- [51] ATLAS Collaboration. The simulation principle and performance of the ATLAS fast calorimeter simulation FastCaloSim. Technical Report ATL-PHYS-PUB-2010-013, 2010.
- [52] J. Pequeno and P. Schaffner. How ATLAS detects particles: Diagram of particle paths in the detector. <http://cds.cern.ch/record/1505342>, 2013.
- [53] ATLAS Collaboration. Performance of the atlas track reconstruction algorithms in dense environments in lh run 2. *The European Physical Journal C*, 77(10), 2017.
- [54] T. Cornelissen, M. Elsing, I. Gavrilenko, W. Liebig, E. Moyse, and A. Salzburger. The new ATLAS track reconstruction. *Journal of Physics: Conference Series*, 119(3):032014, 2008.

- [55] R. Fruhwirth. Application of Kalman filtering to track and vertex fitting. *Nucl. Instrum. Meth. A*, 262:444–450, 1987.
- [56] ATLAS Collaboration. Reconstruction of primary vertices at the ATLAS experiment in Run 1 proton–proton collisions at the LHC. *The European Physical Journal C*, 77(5), 2017.
- [57] W. Lampl, S. Laplace, D. Lelas, P. Loch, H. Ma, and et al. Calorimeter Clustering Algorithms: Description and Performance. Technical Report ATL-LARG-PUB-2008-002. ATL-COM-LARG-2008-003, 2008.
- [58] ATLAS Collaboration. Topological cell clustering in the ATLAS calorimeters and its performance in LHC Run 1. *The European Physical Journal C*, 77(7), 2017.
- [59] 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. *The European Physical Journal C*, 79(8), 2019.
- [60] ATLAS Collaboration. Electron and photon performance measurements with the ATLAS detector using the 2015–2017 LHC proton-proton collision data. *Journal of Instrumentation*, 14(12):P12006–P12006, 2019.
- [61] ATLAS Collaboration. Improved electron reconstruction in ATLAS using the Gaussian Sum Filter-based model for bremsstrahlung. Technical Report ATLAS-CONF-2012-047, 2012.
- [62] ATLAS Collaboration. Muon reconstruction performance of the ATLAS detector in proton–proton collision data at $\sqrt{s} = 13$ TeV. *The European Physical Journal C*, 76(5), 2016.
- [63] ATLAS Collaboration. Muon reconstruction and identification efficiency in atlas using the full run 2 pp collision data set at $\sqrt{s} = 13$ TeV, 2020.
- [64] ATLAS Collaboration. Muon reconstruction performance of the ATLAS detector in proton–proton collision data at $\sqrt{s} = 13$ TeV. *The European Physical Journal C*, 76(5), 2016.
- [65] ATLAS Collaboration. Performance of the reconstruction of large impact parameter tracks in the inner detector of ATLAS. ATL-PHYS-PUB-2017-014, 2017.
- [66] 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.

- [67] H. Oide. VrtSecInclusive: An inclusive displaced vertex reconstruction package in release 21. Technical Report ATL-COM-INDET-2018-006, 2018.
- [68] DELPHI Collaboration. Search for neutral heavy leptons produced in Z decays. *Z. Phys.*, C74:57–71, 1997. [Erratum: *Z. Phys.*C75,580(1997)].
- [69] CMS Collaboration. Search for heavy neutral leptons in events with three charged leptons in proton-proton collisions at $\sqrt{s} = 13$ TeV. *PRL*, 120(22):221801, 2018.
- [70] ATLAS Collaboration. Search for heavy neutral leptons in decays of W bosons produced in 13 TeV pp collisions using prompt and displaced signatures with the ATLAS detector. *JHEP*, 10:265, 2019.
- [71] J. Beacham, C. Burrage, D. Curtin, and A. De Roeck et al. Physics beyond colliders at CERN: Beyond the standard model working group report. *Journal of Physics G: Nuclear and Particle Physics*, 47(1):010501, 2019.
- [72] Torbjörn Sjöstrand, Stefan Ask, Jesper R. Christiansen, Richard Corke, Nishita Desai, Philip Ilten, Stephen Mrenna, Stefan Prestel, Christine O. Rasmussen, and Peter Z. Skands. An introduction to PYTHIA 8.2. *Comput. Phys. Commun.*, 191:159, 2015.
- [73] ATLAS Collaboration. ATLAS Pythia 8 tunes to 7 TeV data. ATL-PHYS-PUB-2014-021, 2014.
- [74] Richard D. Ball et al. Parton distributions with LHC data. *Nucl. Phys. B*, 867:244, 2013.
- [75] Celine Degrande, Olivier Mattelaer, Richard Ruiz, and Jessica Turner. Fully-Automated Precision Predictions for Heavy Neutrino Production Mechanisms at Hadron Colliders. *Phys. Rev. D*, 94(5):053002, 2016.
- [76] Silvia Pascoli, Richard Ruiz, and Cedric Weiland. Heavy neutrinos with dynamic jet vetoes: multilepton searches at $\sqrt{s} = 14$, 27, and 100 TeV. *JHEP*, 06:049, 2019.
- [77] J. Alwall, R. Frederix, S. Frixione, V. Hirschi, F. Maltoni, O. Mattelaer, H. S. Shao, T. Stelzer, P. Torrielli, and M. Zaro. The automated computation of tree-level and next-to-leading order differential cross sections, and their matching to parton shower simulations. *JHEP*, 07:079, 2014.
- [78] ATLAS Collaboration. Performance of vertex reconstruction algorithms for detection of new long-lived particle decays within the ATLAS inner detector. Technical Report ATL-PHYS-PUB-2019-013, 2019.

- [79] ATLAS Collaboration. Measurement of the W -boson mass in pp collisions at $\sqrt{s} = 7$ TeV with the ATLAS detector. *Eur. Phys. J. C*, 78:110, 2018.
- [80] Particle Data Group. Review of particle physics. *Phys. Rev. D*, 98:030001, 2018.
- [81] 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*, 92:072004, 2015.
- [82] 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*, 97:052012, 2018.
- [83] ATLAS Collaboration. Search for displaced vertices of oppositely charged leptons from decays of long-lived particles in pp collisions at $\sqrt{s} = 13$ TeV with the ATLAS detector. 2019.
- [84] D. E. Knuth. *The art of computer programming: Volume 2, Seminumerical algorithms*. Addison-Wesley, Upper Saddle River, New Jersey, third edition, 1997.
- [85] ATLAS Collaboration. Luminosity determination in pp collisions at $\sqrt{s} = 13$ TeV using the ATLAS detector at the LHC. ATLAS-CONF-2019-021, 2019.
- [86] G. Avoni et al. The new LUCID-2 detector for luminosity measurement and monitoring in ATLAS. *Journal of Instrumentation*, 13(07):P07017, 2018.
- [87] A. L. Read. Presentation of search results: the CL_s technique. *J. Phys. G*, 28:2693, 2002.
- [88] ATLAS Collaboration. Trexfitter documentation.
<https://trexfitter-docs.web.cern.ch/trexfitter-docs/>.
- [89] CMS Collaboration. Search for long-lived heavy neutral leptons with displaced vertices in proton-proton collisions at $\sqrt{s} = 13$ tev, 2022.
- [90] CMS Collaboration. Search for long-lived heavy neutral leptons with displaced vertices in pp collisions at $\sqrt{s} = 13$ TeV with the CMS detector. *CMS-PAS-EXO-20-009*, 2021.
- [91] ATLAS Collaboration. Technical Design Report for the ATLAS Inner Tracker Strip Detector. Technical report, 2017.
- [92] The Sensiron Company. Application note dew-point calculation. http://irtfweb.ifa.hawaii.edu/~tcs3/tcs3/Misc/Dewpoint_Calculation_Humidity_Sensor_E.pdf.

- [93] F. P. Schreuder.
FELIX: The New Detector Readout System for the ATLAS Experiment.
URL: https://indico.cern.ch/event/608587/contributions/2614211/attachments/1519973/2374031/TWEPP2017_FELIX.pdf, 2017.
- [94] P. Moreira, R. Ballabriga, S. Baron, S. Bonacini, and et al. The GBT Project.
Technical report, 2009.
- [95] ATLAS Collaboration. A Particle Consistent with the Higgs Boson Observed with the ATLAS Detector at the Large Hadron Collider. *Science*, 338:1576, 2012.
- [96] ATLAS Collaboration. Software Performance of the ATLAS Track Reconstruction for LHC Run 3. Technical report, May 2021.

Appendix A

Data-driven Estimate of Random Track Crossings Using a Track Shuffling Method

This method uses the number of random crossings from tracks in different events as an estimate for the number of random crossings from tracks in the same events in the signal region. This method assumes that the probability for two tracks to form a random crossing vertex is independent of whether the tracks come from the same event or from different events. Thus, the number of background events from random track crossings can be estimated by combining tracks from different data events and producing fake vertices from them.

Displaced vertices reconstructed from accidental track crossings is an inherently random process and, due to the relatively small number of lepton-matched tracks (on average two per event) that pass the VSI track selection in each data event, the accidental crossing of two lepton tracks is very rare. Thus, the track shuffling method can artificially enhance the number of random crossing DVs and can be used to provide a background estimate. Additionally, this method also serves as a useful tool to study the properties of random track crossings because this method randomly pairs tracks from different events creating a pure sample of random crossing events.

To estimate the number of two-lepton random crossing DVs in the signal region, four regions are defined that are labelled A through D. In this method, a SR is defined (region A) and the background is estimated by counting the number background events in three other regions (regions: BCD) where there is no expected signal contamination. These regions are listed in table A.1.

This estimation method, known as the “ABCD method”, assumes that the distribution in the two axes (AB vs. CD and AC vs. BD) is uncorrelated and that all the sources of background listed in section 5.8, except for random track crossings, are negligible. Then, the number of events in the signal region A is computed as,

$$\frac{N_A}{N_B} = \frac{N_C}{N_D} \Rightarrow N_A = \frac{N_C}{N_D} \cdot N_B \quad (\text{A.1})$$

where N_α is the number of events in region α .

Table A.1: The different regions used in the random track crossing background estimation.

	Signal Region (Prompt lepton)		Validation Region (Prompt lepton veto)	
	OS DV	SS DV	OS DV	SS DV
Tracks from same event	A	B	A'	B'
Tracks from different events	C	D	C'	D'

A shuffling factor is also defined as the ratio between shuffled and not shuffled events, as $f = N_D/N_B$ and provides a metric of how many more random crossing DVs can be produced via this track shuffling method. Larger shuffling factors would result in smaller statistical uncertainties in the background estimate.

A.1 Shuffling Analysis Framework

The shuffling analysis framework uses a custom shuffling algorithm to create a list of tracks that are shuffled together from different data events. It also makes use of the HNL analysis framework (described in section 5.7.1) to create and analyze displaced vertices from a list of shuffled tracks. The shuffling analysis framework can be divided into four steps (see fig. A.1):

1. **Make ntuples:** This step uses the filtered data event with displaced reconstruction as an input and applies track selections to save a list of lepton-matched track in an ntuple. In this step, information about the lepton objects matched to each track, including lepton identification criteria and lepton type (electron or muon), is also saved in the ntuple. Information about the run number and event number for each track is also included.
2. **Make CFF files:** The shuffling algorithm creates a common-file format (CFF) file that contains a list of lepton-matched tracks (selected in the previous step) and a list of shuffling instructions that maps various combinations of tracks into unique fake events. The main requirement in this shuffling step is to maximize the number of fake events produced, while also ensuring that any pair of leptons is used in *at most* one fake event (For more details, see appendix A.2).
3. **Derivation:** This step is similar to the derivation step in the main analysis framework. A collection of tracks is constructed using the list of tracks saved in the CFF file. This collection of tracks is used as an input to the VSI vertexing algorithm, which reconstructs displaced vertices using the shuffled tracks. The rest of the information about the event (e.g. the primary vertex, additional non-lepton tracks etc.) is retrieved from a so-called

“host-event” that is supplied by an auxiliary data file. At this point, a fake event has been created with displaced vertices that consist of lepton-matched tracks from different original events.

4. **Make ntuples and histograms:** Once fake events have been reconstructed, they are analyzed and saved in ntuples and mini-ntuples using the same steps as the main-analysis framework.




Figure A.1: An illustration of the main steps of the shuffling analysis framework that demonstrated the types of files are needed as inputs to each step.

A.2 Track Shuffling Problem

To produce fake events with accidental crossings of two lepton tracks, a map of the various combinations of tracks into unique fake events must be produced. This problem can be summarized as follows,

The final requirement is the key to the shuffling problem. This ensures that the same accidental crossing of two tracks is not created in multiple fake events. In practice, computational constraints prevent all of the lepton-matched tracks from being shuffled simultaneously. Instead, the lepton-matched tracks are divided into chunks of tracks with $P < N$ leptons per

Given a collection of N leptons
and given k , a desired number of tracks per fake event,
maximize the number of fake events from the N leptons
requiring that any pair of leptons is used in at most one fake event.

chunk. These chunks of P leptons are shuffled independently. Uniquely shuffling a set of N objects is a conceptually straightforward combinatorics problem, but due to the constraint that “any pair of leptons is used at most one fake event”, it is very challenging to solve mathematically.

The solution employed in the shuffling algorithm is to solve this “shuffling problem” in advance for a set of P leptons and create a so-called *shuffling basis* that pre-determines which leptons should be used in each fake event. One key aspect of this shuffling basis is that it uses information about the number of electrons and muons in each event to create a set of instructions to produce a list of shuffled lepton-matched tracks with a distribution that matches the input lepton distribution. The shuffling algorithm then chunks the selected lepton-matched tracks into groups of size P and uses the shuffling basis to create a set of instructions that determines which leptons will be used in each fake event.

A.3 Validation of the Track Shuffling Method

To validate the shuffling analysis framework, a “do-nothing” configuration is implemented to bypasses the random shuffling part of the framework and create instructions that produce events with the same lepton-matched tracks as the original events. Thus, when the collection of tracks is created in the derivation step, it consists of the *exact same* tracks that were originally selected by the vertexing algorithm in the original data events. The vertexing algorithm is deterministic and, if the same event information and input tracks are used, then the same vertices should be reconstructed. However, there are a few key differences in this “do-nothing” configuration that can slightly alter the collection of reconstructed displaced vertices. While the selected tracks are the same in the real and “fake” events, the rest of the information about the event, including the primary vertex and the non-selected tracks used the VSI track attachment step, is retrieved from a *random* auxiliary “host-event”. Careful study of the dependence of the vertex collection on the primary vertex and non-lepton tracks is therefore required.

The VSI algorithm has two checks that depend on the primary vertices in the event. First, in the two-track vertex forming step, there is primary vertex compatibility condition for each track. This will reject the vertex if either of the tracks is determined to be compatible with the primary vertex. Second, in the track attachment step, if the reconstructed lepton-only vertex is < 0.5 mm from the hard-scatter or any pileup vertex, then the track attachment procedure

is skipped. This second check has a negligible impact on vertices in the fiducial volume of $4 \text{ mm} < L_{xy} < 300 \text{ mm}$. The original primary vertices are passed along in the “do-nothing” configuration and so this dependence on the primary vertices can cause small differences in the vertex reconstruction.

To study the impact of selecting a different primary vertices, events created using the “do-nothing” configuration are compared to the original data events in fig. A.2. Small differences on the order of 10-20% can be observed mainly at low DV radius where there is a stronger dependence on the location of the primary vertex.




Figure A.2: Displaced vertices as a function of DV mass and radius for events produced using the “do-nothing” configuration (in purple) and the original data events (in blue). In the “do-nothing” configuration, the original non-selected tracks are also saved such that, the only difference between each event is the primary vertex. Differences at low L_{xy} can be observed, from the locations of the primary vertices.

In the custom VSI configuration, the track attachment of non-selected tracks (mainly hadrons) is used in the vertexing algorithm to reduce the number of two-lepton displaced vertices. The dependence on the non-selected tracks is checked using the “do-nothing” configuration where tracks are taken from a random host-event or from the original data event. Comparing the distributions in fig. A.3 shows that there are some missing correlations between the non-selected tracks and the selected lepton-matched tracks that results in a 10-20% difference for vertices with $m_{DV} > 5 \text{ GeV}$.




Figure A.3: Displaced vertices as a function of DV mass and radius (L_{xy}) for events produced using the “do-nothing” configuration with the original non-selected tracks (in purple) and non-selected tracks from host data events (in pink). Differences between the distributions show that using hadrons from a random host-event misses some correlations between the non-selected tracks in the original event and the lepton-only displaced vertices.

A.4 Small-Scale Test of the Track Shuffling Method

The data-driven estimate of random track crossings using a track shuffling method was tested using one data run from 2016. In this small-scale test, events were selected in the VR and processed with the “do-nothing” and shuffled configurations to produce both same-event and different-event collections respectively. SS and OS vertices for the shuffled and non-shuffled events are separated in the A’B’C’D’ regions as defined in table A.1. In this test run, a total of 3.703×10^5 original events were processed and 6.967×10^6 shuffled events were created, which gives a shuffling factor of $f = 18.8$. The results from the small-scale shuffling test are shown in fig. A.4. The ratio between the shuffled and not shuffled events has a good agreement with the expected shuffling factor (see ratio plots in fig. A.4). Shuffled vertices form a pure random crossing sample and, in fig. A.4 (d), excellent agreement between SS and OS vertices is observed. This also demonstrates that random crossings are independent of track charge.

While this small scale test was considered a success, numerous challenges arose when trying to scale up the system that mainly stemmed from having to process large amounts of data (on the order of hundreds of terrabytes) in a non-standard way. ATLAS’s main data processing framework is designed to efficiently run over large data files to analyze collision events. However, in the framework described in fig. A.1, the large data files are used as an auxiliary input and the main event information is stored in the custom CFF files.

A.4. Small-Scale Test of the Track Shuffling Method




Figure A.4: DV mass distribution for shuffled and “do-nothing” events for SS and OS displaced vertices. The ratio between shuffled and not shuffled events is the shuffling factor. For this test, with only one collision run from 2016, the expected shuffling factor of 18.8 is the dashed black line in the ratio plot. OS and SS vertices are directly compared in (c) and (d) for the not shuffled and shuffled events respectively.

During the small-scale test of the track shuffling method, this non-standard way of processing the data caused large surges in data transfer volume on the worldwide LHC computing grid. The two large spikes in fig. A.5 are attributed to data being transferred for the track shuffling estimate. This unique way of processing the data single-handedly helped computing experts commission protocols to prevent these types of surges, which can cause major interruptions in the central computing network. Due to the difficulties faced when scaling up the system, this background estimate was not pursued for the final background estimate in the analysis. Nonetheless this track shuffling method shows promising results for the estimation of random two-track crossing backgrounds.




Figure A.5: Data transfer volume on the worldwide LHC computing network as a function of time. The two large spikes in this plot can be attributed to the large amounts of data being transferred for this track shuffling estimate. Figure obtained from the computing monitoring systems.