Search for missing mass in proton-tagged dilepton events with the ATLAS detector and the AFP spectrometer

Josh Lomas

Thesis submitted for the degree of
Doctor of Philosophy

Particle Physics Group,
School of Physics and Astronomy,
University of Birmingham.

September 10, 2025

Abstract

A search for a missing mass resonance is presented in the channel $pp \rightarrow p'p'\ell\ell X$ using 14.7 fb^{-1} of data from the central ATLAS detector and the ATLAS Forward Proton (AFP) spectrometer, collected in 2017 from pp collisions at $\sqrt{s} = 13$ TeV. The process involves photon-exchange between two protons which remain intact, with central production of a visible boson V decaying into charged leptons and an invisible component X which is undetected in the ATLAS detector. A resonance search is performed in the reconstructed missing mass spectrum, with upper limits set on the signal cross section for three different models, and steps taken to reduce the model dependence of the results. The $V \rightarrow ee$ and $V \rightarrow \mu\mu$ channels and a combination of both channels are considered. Significant improvement in sensitivity is achieved with respect to a previous analysis from CMS, with the addition of a track veto selection which yields a large improvement in background rejection. The uncertainties on the signal cross section upper limits are statistically dominated, with the largest systematic effects originating from estimates of soft-survival probability and track veto signal efficiency.

Quality Assurance (QA) measurements performed on silicon strip sensor components produced for the Inner Tracker (ITk) upgrade for the ATLAS experiment are additionally presented.

Declaration of Author's Contribution

The development and construction of the experiments used to collect the data presented in this thesis was only possible through the collaboration of thousands of people, with hundreds more responsible for successful operation during data-taking, ensuring adequate data quality throughout additional processing steps, and producing the many simulated data samples used in this thesis. Without this widespread effort, none of the work we do in experimental particle physics would be possible, and the resulting measurements are the product of everyone who has contributed.

Chapters 2, 3 and 4 provide an introduction to relevant topics for this thesis, explaining the theoretical particle physics background to give context to the analysis presented, and describing the ATLAS detector and AFP spectrometer used to collect the data. None of the theories, designs or results presented here are the work of the author, with corresponding references given throughout. In particular, Refs. [1, 2] and Ref. [3] were used heavily as sources for the theoretical background and AFP spectrometer description, respectively. However, the author has contributed to data-taking through several AFP on-call shifts, monitoring the detector, responding to issues flagged by the central ATLAS experiment control room and fixing or delegating any problems which arise. More recently, the author has begun co-convening the Forward Proton Combined Performance (CP) working group responsible for coordinating performance studies of AFP data to develop recommendations for analyses.

Chapter 5 describes the theory of semiconductor physics and radiation damage, along with the ongoing ITk upgrade to the ATLAS detector for the next evolution of the LHC, none of which is the work of the author, before covering QA measurements performed on sample strip sensor components. The measurement setup and analysis code were developed by a previous student, with several improvements later made to the code by the author. All of the measurements presented in Section 5.3 were made either directly by or with the help of the author, although the QA process involved the work of many other people, performing irradiation, wirebonding and annealing of sensors, and coordinating the efforts.

Finally, Chapter 6 presents an analysis performed by the author, alongside a team of fellow ATLAS collaboration members. The author worked directly on all components of the analysis, with the exception of the development, generation and validation and the signal models described in Section 6.4. The analysis code was developed by the author from a skeleton provided by another team member, with validation provided through comparison with an existing framework enabled by another collaborator. In addition, significant work was done by other team members on analysis strategy, low-p_{T} studies and reconstruction in data (Section 6.5.3), track veto signal efficiency studies (Section 6.7.1), the determination of theoretical uncertainties (Section 6.7.4), fit optimisation for statistical analysis (Section 6.8). All other studies presented are mainly the work of the author.

Acknowledgements

There's no way I could have finished this PhD without the help and support of an army of people, so there's loads of people I have to thank for getting me to this point.

First to my supervisors, thank you Paul for guiding me these last four years (after having already taught me at undergrad!) and for your patience and knowledge. You've given me endless support throughout my PhD and beyond and if I've (hopefully) picked up a bit of your wisdom then I think I'll be good! Thank you also to Laura for your supervision in my first year, I never thought I'd enjoy hardware so much and you always had an answer for my questions.

Thank you to the rest of the Birmingham group for your support, I've probably asked every single one of you a question at least once. Thanks particularly go to Andrew Kirk, Dave and especially Aleksandra, I wouldn't have made half the progress in my thesis without your tireless help. Working on AFP I've met many brilliant people, thank you to everyone for being so welcoming and knowledgable, particularly Rafal whose technical help on the analysis has been invaluable. To all my colleagues on the Dilepton+X analysis team, if I haven't mentioned you yet, thank you for your contributions and advice, particularly Maura for all of your work and for running the analysis group with me.

Thanks to all the fellow students I've worked alongside over the years, both in the ATLAS office and otherwise. You've been a huge help with and distraction from work, whichever was required at the time. Special thanks are owed to Will for teaching me everything there is to know about ITk QA and to Adrien for doing the same with AFP, both in the analysis and the on-call shifts. I also can't forget all my Bubble chambers teammates, even though we never won the trophy (yet!) the training sessions were often the highlight of my week.

To all my old friends in Leicester and all the new ones I made here in Birmingham, thank you for keeping me entertained and always giving me something to look forward to at weekends.

Last but definitely not least, thank you to my family, especially my parents for always supporting me, taking an interest in my interests and pushing me towards my dreams. You've always made me feel like I can achieve anything, and have given me whatever you could to help me do it, and I will always be so grateful to you for that. And to my new wife (!), thank you Rach for putting up with me, pretending to understand what I do all day, keeping my social calendar filled and most importantly making me happy every day.

We made it!

Everything that happens in life must have a scientific explanation. If you know where to look for it, that is. –The Third Doctor

Contents

1 Introduction **1**

2 Particle Physics Theory and the Physics and Modelling of Hadron Collisions **4**

2.1 The Standard Model 4

2.1.1 Quantum Electrodynamics 6

2.1.2 Quantum Chromodynamics 7

2.1.3 The Weak Interaction and Electroweak Theory 8

2.1.4 Spontaneous Symmetry Breaking and the Higgs Boson 10

2.1.5 The Standard Model Lagrangian 12

2.1.6 Feynman Rules, Renormalisation and Quantum Coupling Effects 13

2.2 Beyond the Standard Model 15

2.3 Collider Physics 20

2.3.1 Proton-Proton Collisions 20

2.3.2 Photon-Photon Collisions and the Equivalent Photon Approximation 23

2.3.3 Intact and Dissociated Final State Protons 25

2.3.4 Soft-Survival 27

2.4 Simulation of Particle Interactions 28

3 The ATLAS Detector at the Large Hadron Collider **33**

3.1 The Large Hadron Collider 33

3.1.1 Luminosity and Pile-up 36

3.2 The ATLAS Experiment 39

3.2.1 Inner Detector 41

3.2.2 Calorimeters 44

3.2.3 Muon Spectrometer 49

3.2.4 Forward Detectors 52

3.2.5 Trigger and DAQ 53

3.3 Physics Object Reconstruction 55

3.3.1 Tracks 56

3.3.1.1 Vertex Reconstruction 61

3.3.2 Electrons and Photons 62

3.3.3 Muons 65

3.3.4 Other Objects 68

4 The ATLAS Forward Proton Spectrometer **70**
4.1 Introduction 70
4.2 Detector 72
4.2.1 Silicon Trackers 75
4.2.2 Time-of-Flight 78
4.2.3 Trigger 80
4.3 Forward Proton Reconstruction 80
4.4 Alignment 84
4.4.1 Local Alignment 84
4.4.2 Global Alignment 87
4.5 Performance 89

5 Quality Assurance for the ATLAS Inner Tracker Upgrade **95**
5.1 Semiconductor Physics 95
5.1.1 Radiation Damage 101
5.2 Inner Tracker Upgrade 104
5.3 Test Structures 106
5.3.1 Test Chips and MD8 107
5.3.2 Inter-strip Resistance Investigation 122
5.3.3 Mini Sensors 127

6 Missing mass search in proton-tagged dilepton events **132**
6.1 Introduction 133
6.2 The Missing Mass Method 135
6.2.1 Missing Mass Resolution 136
6.2.2 Missing Mass Acceptance 137
6.2.3 Event Mixing 138
6.3 Data Overview 139
6.3.1 Blinded Data 140
6.4 Signal Simulation 140
6.4.1 Signal Models 141
6.4.2 Kinematic Comparison 145
6.4.3 Simulated Beamspot Size 146
6.4.4 Updated Missing Mass Resolution 147
6.5 Event Selection 148
6.5.1 Lepton Selection 149
6.5.2 Forward Proton Selection 150
6.5.3 Track Selection 151
6.5.4 Pre-Selection 151
6.5.5 Signal Selection 153
6.5.6 Fiducial Selection 157
6.5.6.1 Track Veto 159
6.5.7 Low-p_{T} Tracks 162
6.5.7.1 Reconstruction in Data 162
6.5.7.2 Selection 164
6.5.7.3 Low+High-p_{T} Track Veto 165

6.6 Background Modelling . . . 166
6.6.1 Data-driven Modelling Method . . . 167
6.6.2 Signal-induced Background . . . 170
6.7 Systematic Uncertainties . . . 171
6.7.1 Central Detector Uncertainties . . . 172
6.7.2 Forward Detector Uncertainties . . . 175
6.7.3 Modelling Uncertainties . . . 179
6.7.4 Theoretical Uncertainties . . . 183
6.7.5 Summary of Systematic Uncertainties . . . 184
6.8 Statistical Analysis . . . 186
6.8.1 Likelihood . . . 186
6.8.2 Fit Procedure . . . 189
6.9 Results . . . 191
6.9.1 Fit to Blinded Data . . . 191
6.9.2 Fit to Data and Cross Section Upper Limits . . . 193
6.9.2.1 High and Low-p_{T} Track Veto Comparison . . . 198
6.9.3 CMS Comparison . . . 200

7 Conclusion 203

A Background Validation with Alternative Model 205
A.1 Simulated Samples . . . 205
A.2 Misidentified Leptons . . . 205
A.3 Exclusive SM Processes . . . 208
A.4 Comparison with Data-Driven Model . . . 209

B Signal-Induced Background Investigation 215

References 237

Chapter 1

Introduction

The field of particle physics aims to develop a complete understanding of our Universe at a fundamental level, which is consistent with all physical observations at every distance scale. From the smallest meaningful distance (the Planck length, 10^{-35} m) up to the size of the observable Universe (10^{27} m), such an understanding would allow accurate predictions of every physical process occurring across this full range. Currently our best attempt comes in the form of the Standard Model (SM) of particle physics, which describes a set of elementary particles which form all the visible matter in the Universe and the interactions between these particles. While this model agrees with the majority of experimental observations, there a several places where it falls short, the most obvious being the lack of explanation for the mysterious Dark Matter, which is implied to exist by many astronomical observations.

This has led the particle physics community to build experiments attempting to observe new physics beyond the SM, in the hope that this provides guidance for what is missing from the current theory. As these experiments become larger, we are able to probe higher energy processes with increasing levels of precision. The largest facility constructed to date is the Large Hadron Collider (LHC), a particle accelerator used to collide together particles at extremely high energies, with detectors such as the ATLAS experiment designed to measure the products of these collisions. Analysis of the collected data allows for a wide range of searches to be performed looking for new physics, to push the boundaries of the SM to their limit. The next push is already planned in the form of the High Luminosity LHC (HL-LHC), which requires upgrades to all detectors, including the Inner Tracker (ITk) upgrade to the ATLAS inner tracking detector.

This thesis presents a search for new physics performed by measuring the missing mass in proton-proton (*pp*) collisions at the LHC with the ATLAS detector. The specific channel being considered is the photon-induced production of a central dilepton system in association with an additional component X which is not detected by the ATLAS detector. Photon-induced interactions occur in peripheral *pp* collisions and allow the associated protons to remain intact, losing energy and being scattered down the LHC beampipe. A specialised detector called the ATLAS Forward Proton (AFP) spectrometer is positioned on either side of the ATLAS detector to measure the energy lost by these protons, allowing the total energy of the central interaction to be determined. By subtracting the measurement of the central dilepton system provided by the ATLAS detector, the mass of the undetected component X can be reconstructed, allowing for a search for any previously unobserved mass resonances in the resulting distribution. This is the first analysis utilising this technique in the ATLAS collaboration.

The theoretical underpinning of the SM is summarised in Chapter 2, in addition to the current shortfalls and potential extensions to account for these. The physics of *pp* collisions at the LHC are presented, with particular emphasis on photon-induced interactions, and the procedure used for simulating particle physics processes at the LHC is explained. The technical details of the LHC and ATLAS experiments are given in Chapter 3, and the reconstruction of physics objects relevant to this analysis from the resulting measurements is explained. Chapter 4 provides an overview of the AFP spectrometer, including the technical details and measured performance of the detector and the principle of proton reconstruction. Work carried out on the Quality Assurance (QA) program for the ATLAS ITk upgrade is detailed in Chapter 5, alongside an explanation of the theoretical background behind silicon particle detectors and radiation damage in silicon. Finally, the analysis introduced above is presented in Chapter 6, including methodology, event selection, modelling of signal and background processes, systematic uncertainties and the resulting upper limits obtained on the cross sections of the considered signal processes.

Chapter 2

Particle Physics Theory and the Physics and Modelling of Hadron Collisions

In this chapter, a theoretical overview of the current Standard Model (SM) of particle physics is presented in Section 2.1, and then several remaining issues with the model and further theories potentially explaining them are discussed in Section 2.2. Section 2.3 presents a summary of proton collider physics such as that occurring at the Large Hadron Collider (LHC), including the resulting photon-photon fusion interactions. Finally, the basic principles used to produce simulations of particle physics processes are summarised in Section 2.4.

2.1 The Standard Model

The Standard Model (SM) of particle physics is the most complete description ever produced of the universe at the smallest scales. It states that all matter and the forces between it are created from 17 types of particle, shown in Figure 2.1. These are fundamental particles, meaning that they are point-like and cannot be further

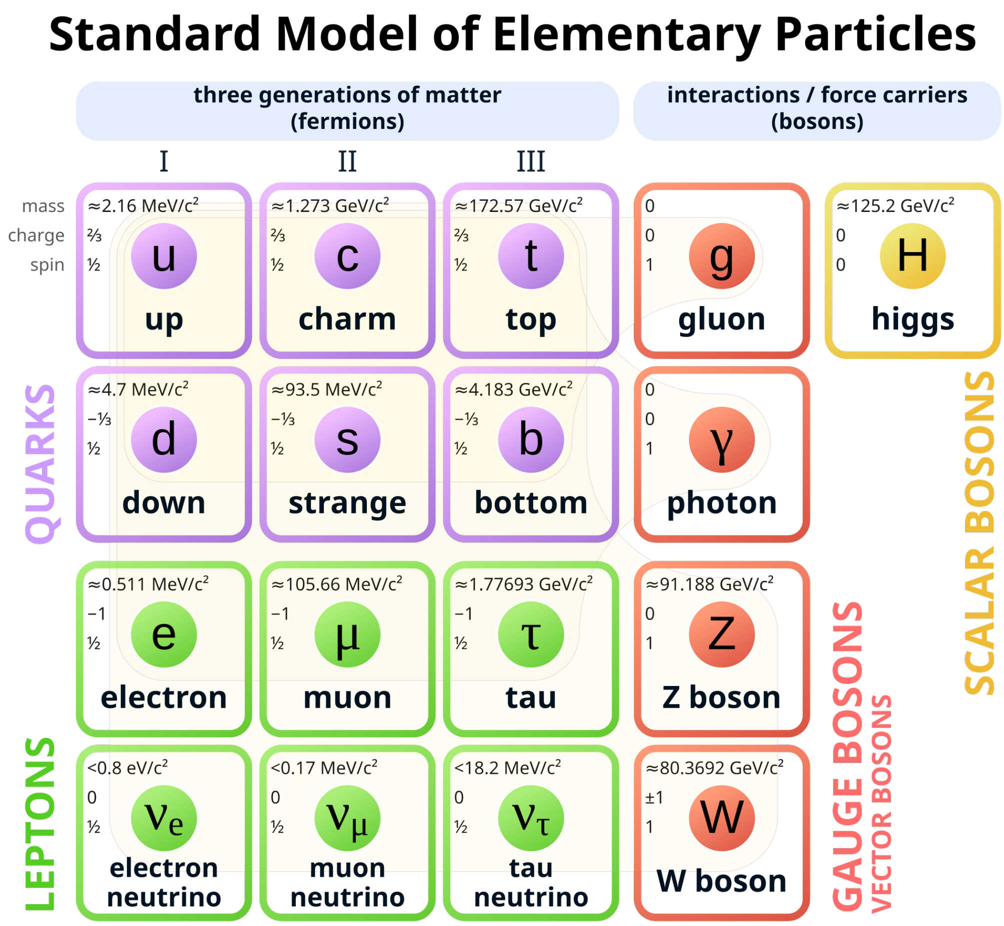

Figure 2.1: Fundamental particles of the Standard Model (SM) of particle physics [4].

broken down. The SM predicts the existence of 12 fermions and 5 bosons. Fermions are spin-$\frac{1}{2}$ particles obeying Fermi-Dirac statistics and made up of 6 quarks and 6 leptons, and their corresponding antiparticles[1], which form all visible matter in the Universe. Bosons are integer-spin particles obeying Bose-Einstein statistics, which mediate the fundamental forces of nature responsible for all interactions taking place between particles. There are four spin-1 bosons, the photon, gluon (coming in 8 varieties) and $W^{\pm}$ and Z^0 bosons and a single spin-0 Higgs boson.

The SM is formulated as a Quantum Field Theory (QFT), which is the simplest theoretical description which can account for both quantum mechanics and special

[1]Every particle has a corresponding antiparticle with the same mass and opposite physical charges such as electrical charge.

relativity, and describes particles as excitations of quantum fields running through the entire Universe. Theories in QFT are defined by their symmetries under "gauge transformations", where the gauge is an abstract frame of reference. Changes in this reference frame are called "gauge transformations", and a theory whose observables are not affected by such a transformation is "gauge invariant", referred to as a "gauge theory". The set of gauge transformations to which a theory is symmetric form a Lie group, of which the simplest example is $U(1)$, the group of all unitary 1×1 matrices. This group has a single member or "generator", consisting of all complex numbers with a magnitude of 1, with the corresponding gauge transformation being a change in the complex phase $\psi(x) \rightarrow e^{i\alpha}\psi(x)$. For constant α this is a global gauge transformation, while for $\alpha(\vec{x}, t)$ which varies in space-time this is a local gauge transformation. For each gauge symmetry present in a theory, a corresponding "gauge field" must be introduced to compensate for the changing parameter and preserve the gauge invariance. In the case of $U(1)$, the group has a single generator and so a single corresponding gauge field exists to maintain the symmetry. If the gauge invariance is local then the gauge field can vary across space-time, creating a corresponding force which results in the observed interactions between particles. The quanta of these fields are called "gauge bosons" [1, 2].

2.1.1 Quantum Electrodynamics

The first QFT was proposed in 1927 by Paul Dirac in the form of Quantum Electrodynamics (QED) [5]. This is an extension of classical electromagnetism (EM), which creates a force between particles based on their intrinsic electric charge, for example binding the negatively charged electrons to the positively charged nucleus in atoms. QED is a gauge theory respecting local $U(1)$ transformations, with the Lagrangian

$$\mathcal{L}_{\text{QED}} = i\bar{\psi}\gamma^{\mu}D_{\mu}\psi - m\bar{\psi}\psi - \frac{1}{4}F^{\mu\nu}F_{\mu\nu}. \tag{2.1}$$

The first two terms describe the kinematics of fermions and interactions between them and the bosonic, spin-1 field A_{μ}, which is introduced via the modified derivative

D_μ in order to maintain the gauge symmetry of the Lagrangian. ψ is a Dirac spinor describing a four-component fermion field, encoding both the spin direction and particle/antiparticle state of the fermion, and γ^μ are the Dirac matrices [6]. The final term is the kinetic energy term for the gauge field, creating a physical gauge boson, the photon.

The Lagrangian does not include a mass term for the gauge boson as this would break the gauge symmetry, so the photon must be massless. Additionally, since the $U(1)$ group is Abelian, meaning that its generators commute ($AB = BA$ for any complex A and B), no self-interaction term arises for the photon. Therefore, the photon itself does not carry electric charge and the Electromagnetic (EM) force has infinite range [1, 2].

2.1.2 Quantum Chromodynamics

Analogously to QED, a QFT describing the strong interaction is obtained with local gauge invariance in the form of Quantum Chromodynamics (QCD). The strong interaction affects only quarks, and is responsible for binding quarks together to form hadrons. These can be mesons such as the π_0, formed of a quark and an antiquark ($q\bar{q}$), or baryons such as the proton and neutron, formed of three quarks (qqq), with four and five-quark-antiquark states having also been observed [7, 8]. The conserved property in QCD is colour charge, which is an intrinsic property of particles similar to electric charge.

The concept of colour charge was introduced by Greenberg in 1964 [9] to explain the existence of hadron states such as uuu, which seem to violate the Pauli exclusion principle, which states that no two fermions may occupy the same quantum mechanical state [10]. This is solved by the existence of three colour charges: red, blue, green; and the corresponding anticolours, with each quark carrying a single one of these colour charges. Other particles such as leptons are “colourless” and therefore do not participate in the strong interaction. All free hadron states observed in nature are also colourless, meaning that they must contain either equal quantities of red, green and blue colour charge, or equal mixtures of colour and anticolour. This must be unaffected by any transformation in the quark colour fields (which can be thought of as rotations in RGB space), giving the local gauge symmetry required in QCD.

Transformations in the quark colour fields are described by the $SU(3)$ group, containing all "special" unitary 3×3 matrices which are traceless (having a determinant of 1). This group has eight generators $T_a = \lambda_a/2$, where λ_a are the 3×3 "Gell-Mann" matrices λ^a for $a \in \{1\ldots8\}$ [2, 11]. The QCD Lagrangian is given by

$$\mathcal{L}_{\text{QCD}} = \sum_f \left(\bar{q}_{fj}\left(i\gamma^\mu\partial_\mu - m\right)q_{fj} - g_s\left(\bar{q}_{fj}\gamma^\mu T_a q_{fj}\right)G^a_\mu\right) - \frac{1}{4}G^a_{\mu\nu}G^{\mu\nu a}, \tag{2.2}$$

where $f \in \{u, b, c, s, t, b\}$ is summed over all quark flavours. As in QED, the first term describes quark kinematics, with the second term describing the interactions between quarks and a set of eight spin-1 bosonic gauge fields G^a_μ introduced to maintain gauge invariance, which couple to quarks with a strength g_s, the strong interaction coupling. Since colour cannot be measured experimentally, the eight individual gluon fields cannot be distinguished, and so all gluonic field excitations are typically referred to collectively as a single "gluon". The final term describes the kinematics of the gluon, with a key difference from QED arising from the fact that the $SU(3)$ group is non-Abelian, meaning that its generators do not commute. This results in self-interaction between gluons via Yang-Mills theory [12], so the strong force has a finite range and gluons themselves must carry colour charge, with each gluon carrying a colour and an anticolour [1, 2].

2.1.3 The Weak Interaction and Electroweak Theory

Measurements of the lifetimes of particles, in addition to the existence of β-decay in atomic nuclei, suggested the existence of a third, weaker fundamental force affecting all fermions, which was named the "weak interaction". This is the only SM force which couples to the neutrino, first postulated by Pauli in 1930 to explain the electron energy spectrum in β-decays and formalised by Fermi in 1934 [13], which is both colourless and electrically neutral.

The weak interaction is defined as a QFT respecting non-Abelian $SU(2)$ local gauge symmetry to the conserved property of weak isospin I. The group consists of three generators $t_a = \sigma_a/2$, where σ_a are the Pauli matrices. Weak interactions are locally gauge invariant under rotations in the three-component weak isospin I which, as for the previous gauge theories, requires the introduction of three spin-1 bosonic gauge fields W^a_μ which couples to fermions with a strength g, the weak interaction coupling. As for QCD, the non-zero commutator of the group's generators introduces self-interaction terms for the three corresponding bosons, and any mass term for the bosons is forbidden under gauge invariance.

It was first observed in the β-decays of polarised (spin-aligned) ^{60}Co nuclei [14], and subsequently proposed by Lee and Yang in 1956 that the weak interaction does not conserve parity [15]. Parity invariance requires the symmetry of interactions under the inversion of all vectors $\mathbf{v} \to -\mathbf{v}$ to the opposite chirality[2]. Therefore, we must consider separately the effect of the weak interaction on fermions with left- and right-handed chirality $\psi_{L,R}$, defined as eigenstates of the projection operators $\hat{P}_{L,R} = \frac{1\mp\gamma_5}{2}$, where $\gamma_5 = i\gamma_0\gamma_1\gamma_2\gamma_3$ is product of the Dirac matrices. They take the form

$$\psi_L = \begin{pmatrix} \nu_{eL} \\ e^-_L \end{pmatrix}, \begin{pmatrix} \nu_{\mu L} \\ \mu^-_L \end{pmatrix}, \begin{pmatrix} \nu_{\tau L} \\ \tau^-_L \end{pmatrix}, \begin{pmatrix} u_L \\ d'_L \end{pmatrix}, \begin{pmatrix} c_L \\ s'_L \end{pmatrix}, \begin{pmatrix} t_L \\ b'_L \end{pmatrix} \tag{2.3}$$

$$\psi_R = e_R,\, \mu_R,\, \tau_R,\, u_R,\, d_R,\, c_R,\, s_R,\, t_R,\, b_R, \tag{2.4}$$

with left-handed spinors forming weak isospin doublets, while right-handed spinors form singlets and do not carry weak isospin. As a consequence of this, right-handed neutrinos do not interact in the SM in any way, having no electric, colour or weak

[2] Chirality is an intrinsic particle property, related to how the particle wave function transforms under rotations

isospin charge. Such neutrinos, if they exist, are referred to as "sterile" neutrinos, and would interact only via gravity [1, 2].

In the 1960s the Glashow-Weinberg-Salam model was developed to unify QED and the weak interaction into a single interaction [16–18]. The extended "electroweak" theory is a gauge theory respecting both symmetries of its composite theories as $SU(2)_L \times U(1)_Y$, conserving both weak isospin in left-handed fermions (L) and hypercharge Y, related to both electric charge and weak isospin as

$$Y = 2\,(Q - I_3)\,. \tag{2.5}$$

However, gauge invariance of this theory forbids mass terms for the weak bosons, which have been measured experimentally to have mass. In addition, asymmetrical mixing between the left- and right-handed fermion states in the electroweak Lagrangian breaks the gauge symmetry for non-zero fermion mass, which has also been experimentally observed. These issues are solved through mechanism of spontaneous symmetry breaking [1, 2].

2.1.4 Spontaneous Symmetry Breaking and the Higgs Boson

Spontaneous Symmetry Breaking (SSB) is achieved through the Brout-Englert-Higgs (BEH) mechanism [19–21], and allows for non-zero fermion and weak gauge boson mass by keeping the Lagrangian gauge invariant but allowing the vacuum (the state which minimises the Hamiltonian) to vary from zero. This has the effect of spontaneously breaking the symmetry of the theory, as the vacuum is not gauge invariant. This is achieved by introducing an $SU(2)$ doublet of complex scalar fields ϕ, with a potential

$$V(\phi) = \mu^2\phi^\dagger\phi + \lambda(\phi^\dagger\phi)^2, \tag{2.6}$$

where μ and $\lambda > 0$ are constant coefficients. For $\mu^2 < 0$ the potential has an infinite number of minima which are invariant under $SU(2)$ symmetry. This symmetry is spontaneously broken by expanding $\phi(x)$ about a particular minimum

$$\phi_1 = \phi_2 = \phi_4 = 0, \quad \phi_3^2 = -\frac{\mu^2}{\lambda} \equiv v^2, \tag{2.7}$$

where v is the "Vacuum Expectation Value" (VEV) [18], giving an $SU(2)$ doublet of the form

$$\phi(x) = \sqrt{\frac{1}{2}} \begin{pmatrix} 0 \\ v + h(x) \end{pmatrix}. \tag{2.8}$$

This gives rise to the scalar Higgs field $h(x)$, whose excitation is the Higgs boson, which allows three of the four electroweak symmetries $SU(2)_L \times U(1)_Y$ to be broken. This implies the existence of three associated massless Goldstone bosons [22], which will allow the three weak bosons to gain mass. However, since ϕ_0 is electrically neutral, the $U(1)_Q$ symmetry originating in QED is maintained and the photon remains massless, as observed [1, 2].

The Higgs field couples to the electroweak gauge fields W_μ^i and B_μ. As a result of this coupling, the Goldstone bosons associated with the three broken symmetries of the theory are absorbed into the gauge fields, becoming the longitudinal components of the weak bosons and giving them mass. The physical electroweak fields are then recovered through mixing of the gauge fields as

$$W_\mu^\pm = \frac{1}{\sqrt{2}} \left(W_\mu^1 \mp i W_\mu^2 \right) \tag{2.9}$$

$$\begin{pmatrix} A_\mu \\ Z_\mu \end{pmatrix} = \begin{pmatrix} \cos\theta_W & \sin\theta_W \\ \sin\theta_W & \cos\theta_W \end{pmatrix} \cdot \begin{pmatrix} B_\mu \\ W_\mu^3 \end{pmatrix}, \tag{2.10}$$

where θ_W is the Weinberg weak mixing angle, related to the electroweak coupling constants as $\theta_W \equiv \tan^{-1}\left(\frac{g_Y}{g_L}\right)$ and the elementary electric charge as $e = g_L \sin(\theta_W) = g_Y \cos(\theta_W)$. Mass terms for the electroweak bosons can then be derived as

$$m_{W^\pm} = \frac{1}{2} v g, \quad m_Z = \frac{1}{2} v \sqrt{g^2 + g'^2}, \quad m_\gamma = 0, \tag{2.11}$$

with the mass of the Higgs boson given at tree level as $m_H = v\sqrt{2\lambda}$. The Higgs field also provides a mechanism for fermions to gain mass [1, 2].

This theory introduces several free parameters whose values are not predicted in the SM, such as μ, λ, v and the corresponding boson masses. Instead, these must be measured experimentally, with the latest $W^{\pm}$ and Z_0 world-averaged mass measurements being $m_{W^{\pm}} = 80.3692 \pm 0.0133$ GeV and $m_Z = 91.1880 \pm 0.0020$ GeV, and the VEV determined to be $v = 246.22$ GeV [23]. With the discovery of the Higgs boson in 2012 by the ATLAS [24] and CMS [25] experiments, this theory was seemingly confirmed, "completing" the SM. The current world-averaged Higgs boson mass is $m_H = 125.20 \pm 0.11$ GeV [23].

2.1.5 The Standard Model Lagrangian

The complete SM is formed of a combination of $SU(3)$ QCD and $SU(2) \times U(1)$ electroweak theory, giving an overall group respecting $SU(3)_C \times SU(2)_L \times U(1)_Y$ local gauge symmetry, with the subscripts indicating the conserved charges of colour, weak isospin and electric charge, respectively. The symmetry is spontaneously broken down to $SU(3)_C \times U(1)_Q$ through the BEH mechanism by the addition of the scalar Higgs field with non-zero vacuum expectation value. The full SM Lagrangian can be constructed from the various components discussed before as

$$\mathcal{L}_{\text{SM}} = \mathcal{L}_{\text{Gauge}} + \mathcal{L}_{\text{Fermion}} + \mathcal{L}_{\text{Higgs}} + \mathcal{L}_{\text{Yukawa}}, \tag{2.12}$$

where the four components are:

- $\mathcal{L}_{\text{Gauge}}$ - describing the gauge bosons via field strength tensors, including kinetic and interaction terms, plus self-interactions for certain bosons.

- $\mathcal{L}_{\text{Fermion}}$ - describing fermion kinematics and their interactions with gauge bosons through both the strong and electroweak forces.

- $\mathcal{L}_{\text{Higgs}}$ - describing the coupling between the weak gauge bosons and the Higgs field, allowing them to gain mass, plus a kinetic term for the Higgs boson.

- $\mathcal{L}_{\text{Yukawa}}$ - describing the coupling of fermions to Higgs field, allowing them to gain mass.

This complete Lagrangian can be used to predict physical quantities such as the cross sections of particle interactions, related to their probability of occurring, as discussed in Section 2.3.

2.1.6 Feynman Rules, Renormalisation and Quantum Coupling Effects

As will be discussed in detail in Section 2.3, the Lagrangian of a theory such as QED can be used to formulate Feynman rules, dictating which interaction vertices are permitted under that theory, and from these how to determine the probability of a given interaction occurring. Such interactions can be illustrated using Feynman diagrams, first used by Feynman in 1949 [26], such as the outlines shown in Figure 2.2. Figure 2.2a shows a Leading Order (LO) or "tree-level" diagram with the

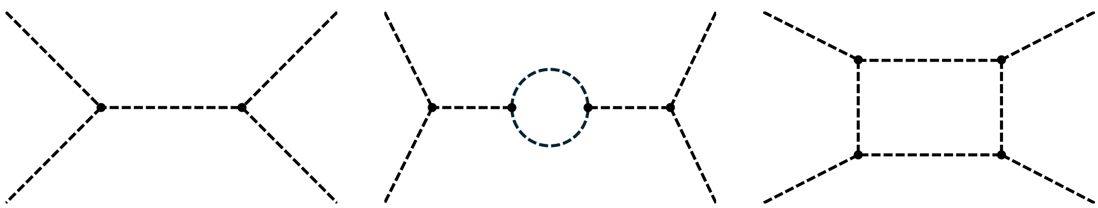

(a) (b) (c)

Figure 2.2: Example (a) Leading Order (LO) and (b) and (c) Next-to-Leading Order (NLO) Feynman diagram outlines.

minimum number of vertices required for the interaction to take place, while Figures 2.2b and 2.2c show Next-to-Leading Order (NLO) contributions including additional vertices. There is no constraint on the number of additional vertices which can be added within such an interaction, although the probability decreases with each addition. Since the momenta of such loops are not constrained, "Ultraviolet (UV)" divergences can be introduced into calculations. These result in infinite computations, and must therefore be resolved through the application of regularisation and renormalisation. Regularisation simply states that the momentum cannot exceed a certain threshold, preventing integration up to infinite momenta, while renormalisation redefines the Lagrangian properties, which do not correspond to physical observables and so can be scaled to absorb these divergences [2].

When renormalisation is applied to QED, the electromagnetic field coupling, and by extension the elementary electric charge q itself, are redefined from constants to being dependent on the energy scale of the interaction. Physically, this is a result of charge screening due to the vacuum polarisation effect, where vacuum fluctuations cause continuous production of fermion-antifermion pairs, which become polarised in the presence of a charged particle. This screens the EM force from the charged particle, reducing its effect for increasing distances. Also referred to as "fine structure", this causes the QED coupling to increases at larger energy scales (or equivalently smaller distances), according to

$$\alpha_{\mathrm{QED}}(Q^2) = \frac{q^2(Q^2)}{4\pi} = \frac{\alpha_{\mathrm{QED}}(\mu_R^2)}{1 - [\alpha_{\mathrm{QED}}(\mu_R^2)/3\pi]\ln(Q^2/\mu_R^2)}, \tag{2.13}$$

for momentum transfer Q^2. Here, μ_R^2 is introduced as the "renormalisation scale", which prevents divergence for large Q^2 [27]. In the limit $Q^2 \to 0$, $\alpha_{\mathrm{QED}} \approx 1/137$.

Similarly, QCD requires renormalisation to prevent divergences, which is achieved by a "running" strong coupling which varies with energy scale. However, a key difference from QED is the presence of gluon self-interactions, which result in an anti-screening effect which competes against the continuous quark-antiquark production occurring due to vacuum fluctuations. This causes the quark-antiquark potential to increase with distance, until it becomes energetically favourable to produce a new quark-antiquark pair from an intermediate gluon. This is the principle of "colour confinement" and manifests in the QCD coupling becoming stronger at larger distance scales (lower energies), resulting in only colourless hadron states being observed in nature. The coupling varies with momentum transfer Q^2 as

$$\alpha_S(Q^2) = \frac{12\pi}{(33 - 2n_f)\ln(Q^2/\Lambda^2_{\mathrm{QCD}})}, \tag{2.14}$$

where n_f is the number of quark flavours (3 in the current understanding of the SM), and $\Lambda_{\mathrm{QCD}} \sim 1$ GeV is the threshold for perturbative calculations to be appropriate for QCD, beyond which α_S becomes too large. Above this energy scale, colour confinement breaks down and free colour charges become possible, achieving "asymptotic freedom" [23, 28].

2.2 Beyond the Standard Model

The SM agrees well with the vast majority of experimental observations, having been extensively tested to high precision by experiments such as those at the LHC, with no disagreements having yet been found. However, many things are still not explained by the SM, for example why all particles have electric charges which are multiples of $q/3$, allowing atoms to be electrically neutral, or why there are three generations of increasingly heavy fermions. Other issues are the neutrino mass, predicted to be zero by the SM but confirmed to be non-zero by the observation of neutrino flavour oscillations, and the CP violation of the Universe, which is implied by our current understanding that at the creation of Universe matter and antimatter were produced in equal quantities, but in the present day the two are hugely imbalanced. In addition, cosmological observations such as the acceleration of the expansion of the Universe and the large scale structures of galaxies and clusters suggest that only around 5% of the energy in the Universe is composed of particles predicted by the SM, with the remainder composed of dark energy (~69%) and dark matter (DM) (~26%), whose origins are currently unknown [29]. In the attempting to explain the origin of dark energy, the fact that the vacuum energy predicted by the SM disagrees by a factor of at least 10^{56} with cosmological measurements is referred to as the "cosmological constant problem". Finally, the SM does not account in any way for gravity, the weakest of the four fundamental forces, which will cause the SM to break down at the highest energy scales where gravity becomes significant relative to the other forces, occurring at the Planck scale of ~10^{19} GeV. The current gravitational theory of general relativity does not necessarily contradict the SM, but unlike the other interactions it cannot be quantised in the form of a QFT [2].

However, perhaps the most convincing argument that the SM is not the final story is how seemingly arbitrary it is, with at least 20 free parameters needed to describe particle masses, mixing angles, couplings and other constants. None of these parameters are predicted by the SM and must instead be set experimentally, leading to "fine-tuning", where although the theory does not predict the values of parameters, it only works when they take very specific values. All this suggests that the SM is in fact a low-energy approximation of some more fundamental "Grand Unified Theory" (GUT). The current SM does not fully unify the electromagnetic, weak and strong interactions, with the $SU(3) \times SU(2) \times U(1)$ group being the product of three disconnected sets of gauge transformations with unrelated coupling constants, whose ratios must be measured experimentally. The interactions can only truly be unified if they are embedded into some larger set of transformations

$$G \supset SU(3) \times SU(2) \times U(1), \tag{2.15}$$

with a single unified coupling g_G related to the individual couplings of each theory, which explains all interactions simultaneously. Indeed, as discussed in the previous section, the SM couplings following renormalisation vary as a function of the energy scale, which allows us to imagine some very large energy scale $Q = \Lambda_G$ at which all three are equal. It can be shown that this "unification scale" is expected to be on the order of 5×10^{14} GeV, far above our current experimental reach [1].

It was shown in 1974 by Georgi and Glashow that the $SU(5)$ group is the smallest potential candidate for a unified group G [30]. This model predicts 24 gauge bosons, consisting of the 12 SM gauge bosons (8 gluons, 3 weak bosons and the photon) and a pair of new super-heavy gauge bosons X and Y which form a weak doublet and are coloured, therefore coupling to both lepton and quark fields. Also called "leptoquarks", these bosons are capable of violating baryon and lepton number, due to the unification of the strong and electromagnetic interactions at this scale. However, their existence would render the proton unstable, which disagrees with current observations. In addition, a problem arises due to the energy scale of this model being so much higher than the electroweak energy scale ($\mathcal{O}(10^2)$ GeV). The tree-level Higgs mass in the SM receives quadratically-divergent corrections from loop diagrams which should push its mass towards the $SU(5)$ scale of $\Lambda_G{\sim}10^{14}$ GeV, the scale on which the leptoquark masses would be produced. However, the observed Higgs mass is around 12 orders of magnitude below this, which is known as the "electroweak hierarchy problem" [2].

One existing Beyond the Standard Model (BSM) theory which could solve this problem is Supersymmetry (SUSY), which suggests that every particle in the SM has a corresponding supersymmetric counterpart or "sparticle", which has equal mass but differs in spin by 1/2. This can result in the loop contributions to the Higgs mass being cancelled out by contributions from the associated sparticle, removing the hierarchy problem. In addition, SUSY provides many new particles which could contribute to DM. However, no experimental evidence has yet been found for SUSY.

Another BSM model theorises the existence of the axion, which arises as a pseudo-Goldstone boson from the spontaneous breaking of an additional global $U(1)$ symmetry imposed in the Peccei-Quinn mechanism [31] in an attempt to explain CP violation in the strong interaction [32–35]. In addition, such a particle could contribute to DM [36, 37]. The axion must be very light and couple to two photons. Some further extensions to the SM also predict the existence of heavier "Axion-like Particles (ALPs)", neutral, spin-0 particles coupling to both fermions and gauge bosons. ALPs couple to fermions through dimension-5 operators proportional to the fermion mass, and to gauge bosons through dimension-5 operators containing derivatives. This results in ALPs only being measurable through their couplings to the gauge or Higgs bosons beyond the energy scale of the top mass [38]. Many previous searches have been performed looking for evidence of ALPs through their coupling to photons, both at the LHC and through other experiments, with a summary of several exclusion limits shown in Figure 2.3. This includes several results

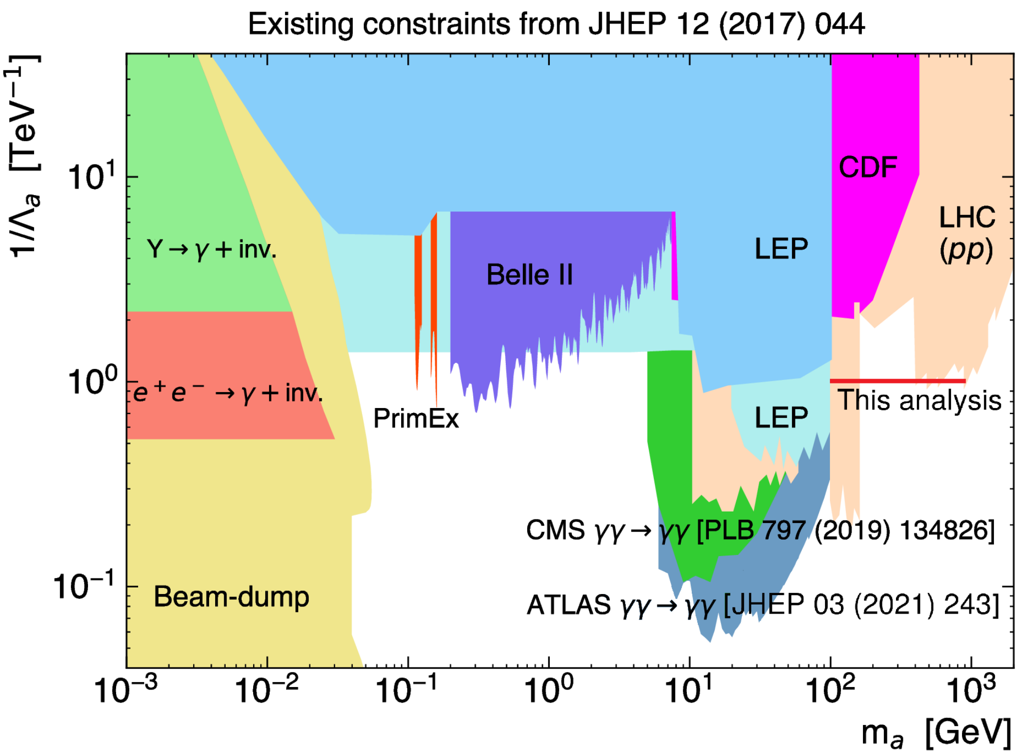

Figure 2.3: Compilation of exclusion limits at 95% CL in the ALP-photon coupling ($1/\Lambda_a$) versus ALP mass (m_a) plane obtained by different experiments, adapted from [39], assuming a 100% ALP decay branching fraction into photons. The phase space probed by the analysis presented in this thesis is shown in red. Recent results from measurements of light-by-light scattering in collisions between heavy nuclei (Pb) are shown from both the ATLAS experiment, the experiment used in this thesis, and the CMS experiment, an equivalent experiment at the same facility. Additionally, results from proton-proton (pp) collisions measured by the same experiments are shown, in a similar phase space to this analysis.

from the ATLAS and CMS experiments at the LHC, in both proton (pp) and heavy nucleon collisions, with higher ALP masses accessible through pp interactions due to their higher collision energy. The LHC is presented in detail in Chapter 3, alongside a description of the ATLAS experiment, which was used to collect the data for the analysis presented in this thesis. Another similar analysis was recently conducted searching for ALPs produced in light-by-light scattering, with the additional use of the AFP spectrometer at the ATLAS experiment. This detector provides measurements of "forward protons" which remain intact following pp interactions, allowing for more detailed measurement of the invisible component of particle interactions, which cannot be directly detected by experiments [40]. The AFP spectrometer is presented in detail in Chapter 4, and is utilised for the analysis presented in this thesis in a similar manner.

The analysis presented in this thesis includes a search for ALPs production from photon-fusion, with the specific model requiring the production of two ALPs, one which is long-lived and is not measured by the ATLAS detector, and another which is short-lived, decaying into a pair of leptons which are measured. Short-lived particles are those with lifetimes below around 10^{-10} s, meaning that they decay before reaching the innermost layer of the ATLAS detector, while long-lived particles are those with lifetimes above around 10^{-7} s, meaning that they are likely to travel beyond the detector before decaying. The phase space probed by this analysis is plotted in Figure 2.3, covering a wide range of currently unexplored phase space. Several analyses have been performed by both the ATLAS [41] and CMS [42, 43] experiments searching for the production of two ALPs, as in the model considered in this thesis, including in the context of short and long-lived ALPs [44]. However, these studies all consider so-called di-ALP production from a Higgs boson, while this analysis considers production from photon-fusion.

In addition to ALP production, this thesis considers two generic models allowing for new BSM physics produced in photon-fusion interactions, as will be discussed further alongside the ALP model in Chapter 6. There have been extensive searches for BSM particles via photon interaction, such as anomalous gauge couplings not allowed by the SM interactions described in Section 2.1 [45–51], specific BSM particles such as DM candidates [52, 53], gravitons [54–56], magnetic monopoles [57] and potential resonances found in previous analyses [58, 59]. Many of these studies utilise forward proton detectors such as the AFP spectrometer, with several generic studies performed on the potential sensitivity to new physics through photon-fusion interactions measured with forward proton detectors [60, 61]. Many of these areas can additionally be probed by the analysis presented in this thesis.

2.3 Collider Physics

2.3.1 Proton-Proton Collisions

Particle accelerators such as the LHC aim to produce a wide range of particle interactions by colliding particles together at high energies; the higher the energy, the wider the range of potential interactions that can be measured. As the name suggests, the LHC collides together hadrons, mainly protons. These are formed of three physical quarks, referred to as "valence quarks", but due to the effects of QCD discussed in Section 2.1, strong interactions between the gluons binding the proton together also produce a much larger number of quark-antiquarks pairs called "sea quarks". The valence quarks, sea quarks and gluons inside the proton are collectively referred to as "partons", any of which can be involved in interactions during pp collisions.

The highest energy pp collisions are achieved in hard scattering interactions between two partons with large fractions of the momenta of their parent protons, where this momentum fraction is represented by the Bjorken variable x. The distributions of parton momenta within the proton as a function of the energy scale are encoded in Parton Distribution Functions (PDFs) [62], which are determined experimentally from Deep Inelastic Scattering (DIS) measurements [63, 64]. For example, the PDF sets used in the generation of simulated signal samples for the analysis presented in this thesis are CT14QED [65], MSHT20 [66, 67] and MMHT2015qed [68], accessed through LHAPDF [69].

The cross section $\sigma_{ij\to X}$ of a particle physics process $ij \to X$ is related to the probability for it to occur within a dataset of a given size, commonly measured in units of "barns" with $1\ \text{b} \equiv 10^{-28}\ \text{m}^2$. The expected number of observed events for a given process is then given by

$$N_{ij\to X} = L\sigma_{ij\to X}, \tag{2.16}$$

where L is the integrated luminosity, corresponding to the size of the dataset as explained in Section 3.1.1. Cross sections are calculated perturbatively, making use of the factorisation theorem, which allows the perturbative and non-perturbative components of the calculations to be considered separately [70, 71]. In the perturbative regime, all incoming partons are summed across and integrated over the allowed momentum space, to give the cross section as

$$\sigma_{pp\to X} = \int dx_i dx_j \sum_{i,j} f_i \left(x_i, \mu_F^2\right) f_j \left(x_j, \mu_F^2\right) \sigma_{ij\to X}, \tag{2.17}$$

where i, j are the incoming partons, $f\left(x, \mu_F^2\right)$ are the corresponding PDFs for each parton and μ_F is the factorisation scale, the threshold above which perturbation is applicable for the cross section calculation, with emissions below this threshold instead absorbed into PDFs. The value taken by μ_F is arbitrary, typically chosen around the energy scale of the hard scatter such that $\mu_F \sim Q$.

The final component of Equation 2.17 is $\sigma_{ij\to X}$, the partonic cross section, giving the probability of producing the desired process from a given pair of partons. This is determined by summing over all possible Feynman diagrams for the process, as determined by Feynman rules [26] defined by the corresponding QFT Lagrangian, and calculating the resulting matrix element $\mathcal{M}$. As discussed in Section 2.1.1, there are an infinite number of Feynman diagrams which exist for any process. Starting from the minimal tree-level diagram, increasing levels of complexity, or "orders", can be considered by adding new intermediate vertices through virtual corrections (loops) or radiation (legs). Each additional vertex contributes an extra power of the relevant coupling constant α to the matrix element, causing more complex diagrams to make diminishing contributions to the overall matrix element, although a larger number of potential diagrams exist for higher orders. To keep the number of calculations finite, the matrix element is only calculated up to a specific order, with tree-level diagrams referred to as LO, diagrams with one extra vertex referred to as NLO, and so on. Additional contributions from higher than the selected order are neglected. The process cross section is then related to the matrix element as $\sigma_{ij \to X} \propto |\sum_i \mathcal{M}_i^2|$, averaging over all potential spin configurations of the initial particles. Figure 2.4 shows a range of cross section measurements performed at the LHC compared to their predictions obtained using the above procedure to NLO or NNLO precision. Some of the cross sections shown are "fiducial" cross

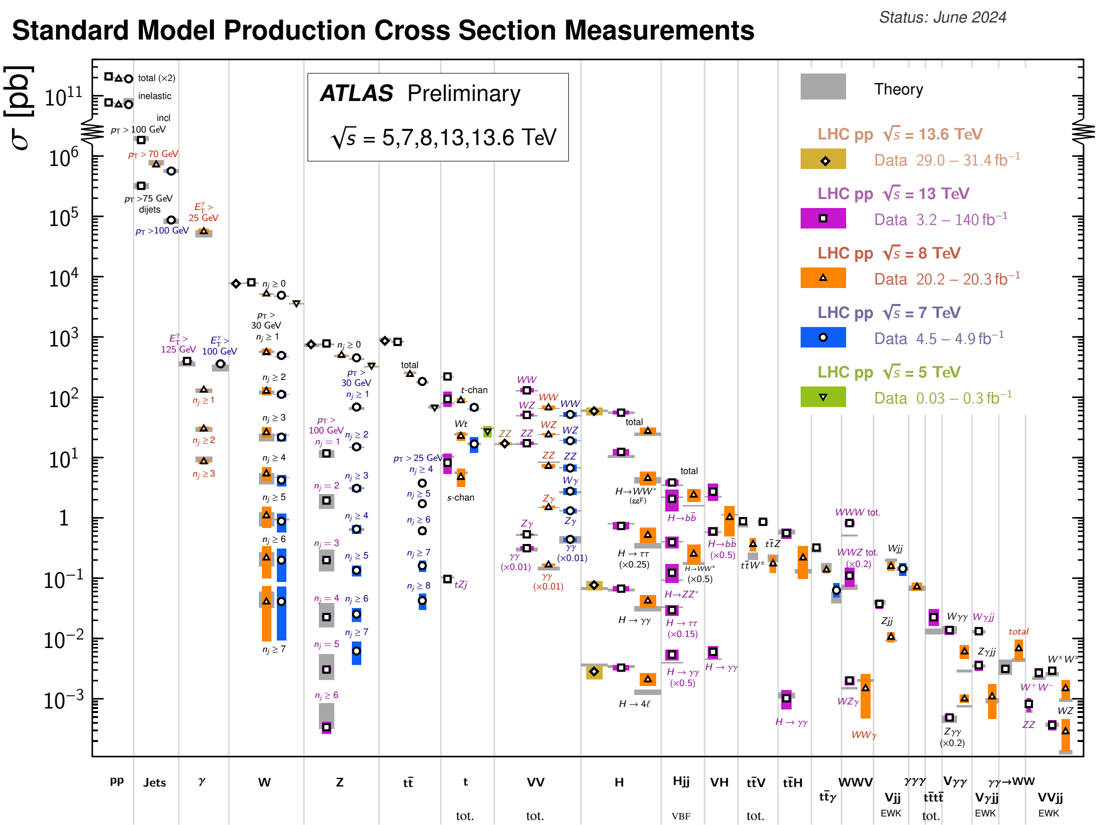

Figure 2.4: Summary of several SM total and fiducial production cross-section measurements [72].

sections, which means that they measure the rate of a process within a limited phase space, typically the region to which the corresponding detector is sensitive. Cross sections without this specification are referred to as "total" cross sections, covering the entire phase space. It is often desirable to measure the "differential" cross section of a process $d\sigma/dx$, which describes the rate of a process as a function of some related kinematic distribution x.

2.3.2 Photon-Photon Collisions and the Equivalent Photon Approximation

Additional contributions are made to hard scatter processes from QED, through photons emitted from the colliding protons, which must be accounted for at the precision now achievable in experiments. Such photons can be considered as extra partons alongside the quarks and gluons in the proton, and included in the corresponding PDFs, generally making much smaller contributions than the QCD partons. These corrections are included in PDF sets provided by several collaborations such as CTEQ [65] and NNPDF [73], which are constrained through measurements of $ep \to e\gamma X$ and Drell-Yan production processes, respectively. These distributions assume high momentum transfer processes with "incoherent" emission of photons from individual quark lines in colliding protons, causing the proton to break apart or "dissociate". However, for low-Q^2 processes, "coherent" emission is also possible whereby the photon is emitted elastically, allowing the proton to remain intact following the interaction. This contribution is most significant at low x, and can be included in the form of additional PDFs [74], or separately through the Equivalent Photon Approximation (EPA).

The EPA replaces the photon PDF with a convolution of equivalent photon fluxes, which arise from the EM field around a proton when it is accelerated to ultra-relativistic speeds. While these fields are emitted radially around a proton at rest, equivalent to a superposition of many low-energy photons, close to the speed of light they become compressed in the direction of travel (by a factor of $\gamma \approx 7000$ at current LHC energies), until they instead resemble a source of coherent high-energy photons emitted from a point-like proton [75, 76]. This effect is illustrated in Figure 2.5. The equivalent photon spectrum generated by a moving particle can be found by integrating the corresponding photon distribution over the transverse momentum up to a threshold $\hat{q}$, below which elastic emission is energetically possible allowing the proton to potentially remain intact. For a typical choice of $\hat{q} = 0.20$

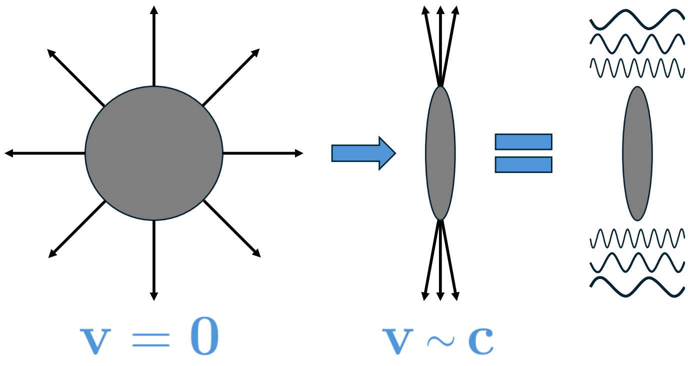

Figure 2.5: Illustration of the principle behind the Equivalent Photon Approximation (EPA). The electric field around a proton at ultra-relativistic speeds becomes compressed such that it resembles a coherent flux of photons. Adapted from [3].

GeV $\approx \Lambda_{\text{QCD}}$ for protons, the equivalent photon spectrum is

$$n(\omega)d\omega = \frac{2Z^2\alpha_{\text{QED}}}{\pi} \ln\left(\frac{\hat{q}\gamma}{\omega}\right)\frac{d\omega}{\omega}, \tag{2.18}$$

for photon energy ω in the limit $\omega \ll \hat{q}\gamma$, where $Z = 1$ is the atomic number of the proton. From here, the cross section for a given central production $pp(\gamma\gamma) \to ppX$ via elastic photon-fusion can be calculated as

$$\sigma(pp(\gamma\gamma) \to ppX) = \int d\omega_1 \int d\omega_2 \sigma(\gamma\gamma \to X) n(\omega_1) n(\omega_2), \tag{2.19}$$

where all possible photon energies $\omega_{1,2}$ are integrated over [77].

An alternative approach to calculating cross sections for photon-initiated processes is possible through the structure function approach [78]. Structure functions describe the internal structure of composite particles, with the proton structure function being well understood through experimental measurements. Applying these functions allows the cross section of photon-initiated production processes to be precisely predicted in a deterministic manner, avoiding the need for PDFs as in the above factorisation approach which are only known to a certain order of precision, introducing uncertainty in the resulting predictions [79]. In addition, the structure function approach allows the cross section to be predicted differentially, as a function of the final-state proton kinematics, and simplifies the inclusion of effects such as proton dissociation and soft-survival, discussed further in Sections 2.3.3 and 2.3.4, respectively. This approach is utilised in the SuperChic generator, discussed further in Section 2.4, which allows uniquely precise determination of the photon-initiated cross section at the percent level [80].

2.3.3 Intact and Dissociated Final State Protons

As discussed above, it is possible for colliding protons to remain intact following an interaction. This is possible if the proton quantum numbers are not changed, and the momentum transfer is sufficiently low. In addition to the exchange of coherent photons discussed above, this can also occur in QCD through the emission of a Pomeron, which can be considered as a non-perturbative collection of low transverse momentum partons forming a colour singlet (e.g. pairs of oppositely coloured gluons). Processes involving the exchange of Pomerons are referred to as "diffractive", and account for a significant portion of the pp interactions occurring in LHC collisions. As discussed above, coherent photons typically have low momentum transfer, resulting in even lower proton p_{T} than is the case for Pomeron exchange. Elastic photon and Pomeron exchange processes leave a common, distinct experimental signature, with the only products being the intact protons in the very forward region, almost parallel to the trajectory of incoming protons, corresponding to very high rapidity $|\eta| \gg 0$[3]. This leaves a "rapidity gap" in the low and mid-rapidity regions in which no particles are produced, although this can be lost if rescattering occurs between the outgoing protons, as discussed in Section 2.3.4.

Central Diffractive (CD) interactions $pp \to pCp$ can also occur, producing a central (low-rapidity) final state C in addition to the high rapidity protons. In QCD this is possible through Double Pomeron Exchange (DPE), where two Pomerons (one from each proton) interact to produce a central state whose colour is discon-

[3]Rapidity is related to the polar angle between an outgoing particle produced in a collision and the accelerator beamline, explained further in Section 3.2, with zero rapidity corresponding to trajectories which are perpendicular to the beamline and higher values corresponding to particles produced more parallel to the beamline.

nected from the protons. This central production can involve interactions between Pomeron constituent partons producing a range of final state particles, or all of the Pomeron energy can go into a single interaction, producing a clean final state with no Pomeron remnants. This latter case is Central Exclusive Production (CEP), and is particularly experimentally interesting due to the energy of the central state being directly related to the initial proton energy loss. The QED equivalent of DPE is double photon exchange, shown in Figure 2.6a, which always leads to CEP as the photon is not composite like a Pomeron, with photon-induced production of some central state C becoming dominant over DPE around $m_C \geqslant 150$ GeV [81]. Measurements of elastic CEP can be used to constrain the total missing mass in the central detector, which can function as a search for BSM physics. This is the principle behind the analysis presented in Chapter 6 of this thesis.

Following an elastic interaction such as double photon exchange, it is possible for one or both of the outgoing protons to become excited. This can cause the proton to break apart into separate partons which immediately hadronise due to QCD confinement, creating a shower of hadronic particles. This process is known as "dissociation", and proceeds via QCD effects in the non-perturbative regime, so it is difficult to predict the kinematic spectra of the outgoing particles. However, since protons involved in photon exchange still maintain significant momentum after radiating a photon, the products of dissociated protons in these interactions are expected to be produced in the same direction, the forward region, maintaining the rapidity gap between the proton dissociation and central systems. Processes involving dissociating protons are referred to as single and double-dissociative, when one or both of the involved protons break apart, respectively. In the case that no dissociation occurs, the process is referred to as elastic. These three scenarios are illustrated in Figure 2.6.

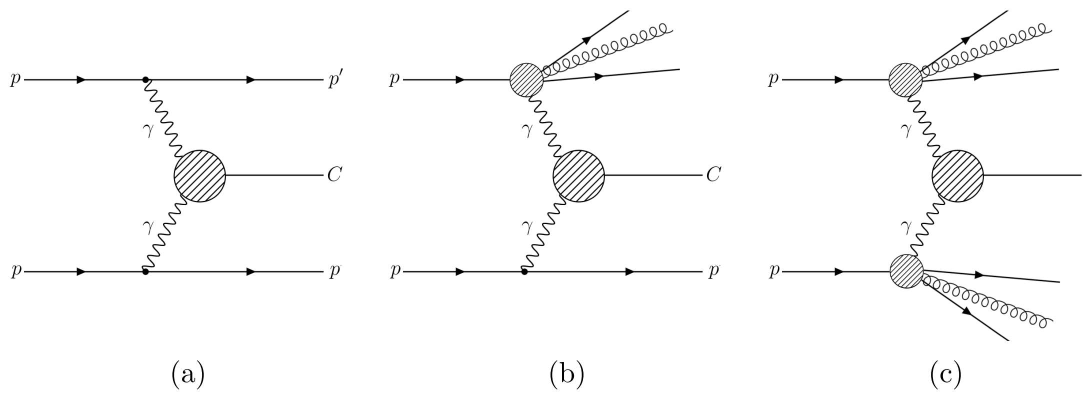

Figure 2.6: Illustration of the three scenarios for photon-induced central exclusive production of some central state C, in the (a) elastic, (b) single-dissociative and (c) double-dissociative channels.

2.3.4 Soft-Survival

The EPA is highly effective at describing photon collisions occurring via proton interactions. However, it does not account for any further interactions between the protons themselves, which can affect the cross-section of photon interactions in addition to the composition of the final state. These proton-proton Multi-Parton Interactions (MPIs) typically occur at low energy scales, and can produce additional particles in the central region which remove the rapidity gap and contaminate the otherwise clean final state produced by photon-induced CEP interactions [82, 83].

Since most MPI proceed via non-perturbative QCD, they cannot be accurately predicted. Instead, this effect is parametrised with a "soft-survival" factor $S \leq 1$, referring to the survival probability of the rapidity gap and corresponding to the probability that no additional particles are produced in the central region, transverse to the initial protons [84]. In most models, soft-survival is dependent on the proton impact parameter (distance of closest approach between the protons) which is inversely proportional to proton p_{T}, with larger impact parameters leading to higher survival probability [80]. This process is very difficult to model accurately, but can be estimated from measurements of the cross sections of relevant processes [85].

2.4 Simulation of Particle Interactions

A key part of measuring processes in particle physics is simulation, which is used to predict the probability and kinematic distributions of a given final state in a particle interaction. These predictions can then be compared to data to determine whether the process is present, and if so whether it is well modelled or understood. This is relevant for both signal (the process being searched for) and background (other known processes with the same final state as the signal) processes, whose predictions must all be combined to make a measurement of the signal. Since particle collisions and subsequent interactions such as those found at the LHC are highly complex, determined by probability distributions acting over a very wide and high-dimensionality phase space, numerical methods must be used to obtain these predictions. Particle physics simulation uses Monte-Carlo (MC) techniques, which use repeated random sampling from known probabilistic distributions of a system to estimate the system properties. So-called MC generators use this technique to generate a sample of events for a given process with random kinematic properties, which average out over large statistics to match the expectation from the model.

There are several steps to producing a simulated MC sample of a physics process, as follows:

1. Generate the matrix element $\mathcal{M}$ of the process of interest by summing over all relevant Feynman diagrams, while accounting for the corresponding PDFs. As discussed in the previous section, this calculation is performed perturbatively only up to a desired order of precision beyond which additional contributions are neglected, in order to maintain a finite number of required calculations. The predicted cross section can be determined from the matrix element using Equation 2.17, as previously discussed.

2. Simulate the "parton shower", which models QCD radiative corrections produced in strong interactions down to the cut-off scale around 1 GeV. At this point, confinement is imposed through hadronisation, where the remaining

partons are grouped into colourless combinations to form hadrons. Both the parton shower and hadronisation steps use specific models which are "tuned" to match experimental observations.

3. Simulate the decay of short-lived particles which will decay before reaching the innermost detector layer, such as top quarks or tau leptons.

4. Overlay MPIs which occur between spectator partons to the primary hard scatter, to form the "Underlying Event" (UE). These processes are non-perturbative, due to large α_S at their typical scale, and must be estimated using phenomenological models which are tuned to the collision conditions. In addition, parton-parton "pile-up" interactions are overlaid, to account for additional scattering occurring between other pp pairs during the same bunch crossing, which is again configured to the collision conditions. Typically, pile-up interactions are soft scatters similar to the UE, with multiple hard scatters per bunch crossing being unlikely.

Figure 2.7 illustrates the entire simulation process for the example process of $t\bar{t}H$ production. Some elements of the process are particularly complex and difficult to simulate accurately. For example the Underlying Event (UE), which needs to account for the dependence on properties such as impact parameter (related to proton p_{T}), in addition to effects such as "colour reconnection" interactions which can occur between partons from different MPIs or the hard scatter itself. Due to the complicated interaction between MPI activity and the products of the hard scatter, the UE must be simulated separately for each event. UE mismodelling is a known issue among many MC generators [87, 88], and is relevant for the analysis presented in this thesis.

There are many different generators which have been developed to produce particle physics MC simulations, with the analysis presented in this thesis using several. Some generators, such as Sherpa [86, 89] and Herwig [90], are multiple-purpose, capable of performing most or all of the above steps in a single program, for a wide range of SM and BSM processes. Another example is Pythia [91, 92], which

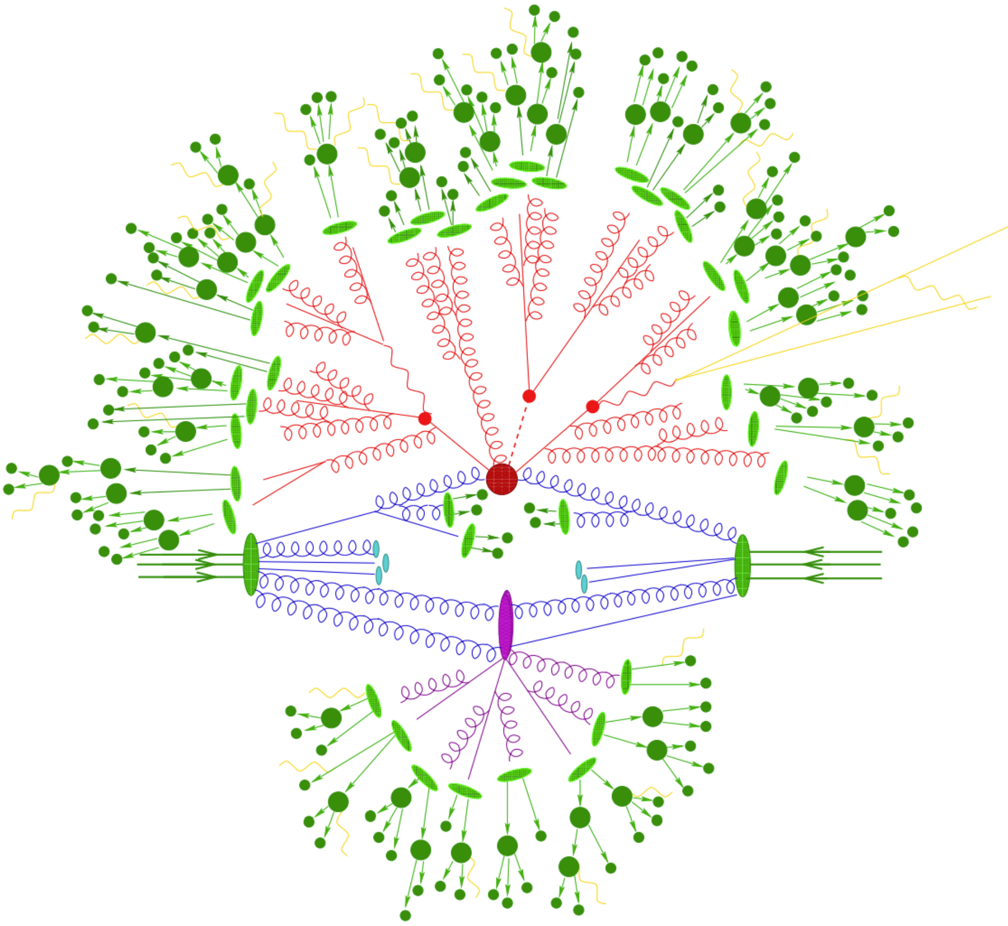

Figure 2.7: Representation of $t\bar{t}H$ production simulated in a Monte-Carlo (MC) event generator, showing the initial hard interaction (central red blob) and resulting quark and Higgs decays (smaller red blobs) and QCD radiative parton showers (red). Overlaid secondary interactions are also shown (purple), with all final state partons below the energy threshold hadronising (light green), before subsequently decaying (dark green). Additional photon radiation occurring at various stages is also shown (yellow) [86].

can simulate both the hard and soft components of hadron and lepton collisions, in both elastic and inelastic cases. Other generators, such as MadGraph [93–95], are matrix element generators which provide only the first step of the above process, with other generators such as Pythia used to provide the parton shower, hadronisation, UE and pile-up simulations.

Of particular interest to the analysis presented in this thesis is the SuperChic generator [96], which uses the structure function approach described in 2.3.2 to predict photon-initiated production cross sections. This achieves a higher precision than other generators which use PDFs, and provides a cross section prediction which is fully differential in the final-state proton kinematics [80]. In addition, it is one of the only generators to include soft-survival effects, discussed in 2.3.4, which have been found to agree with experimental measurements within uncertainties. As with MadGraph, SuperChic is interfaced to Pythia for parton shower and hadronisation simulations.

Following the generation step outlined above for a given process, a sample is obtained with exact kinematic information on all particles produced during the corresponding simulated *pp* bunch crossing, referred to as "truth-level" or "generator-level" information. To properly represent what can be measured experimentally from the process, this sample must then be passed through an additional step which simulates the detector response to the truth-level particles, accounting for detector acceptance, efficiency and resolution through specific cuts and smearing.

The analysis presented in this thesis utilises data from the ATLAS detector for measurement of centrally produced particles, and the AFP spectrometer, for measurement of intact forward protons. These detectors are described fully in Chapters 3 and 4, respectively. The response of the ATLAS detector is simulated using a full geometrical model of the detector in Geant4 [97, 98], referred to as "full simulation". In some cases, less precision is required, and instead a simplified "fast simulation" can be used to speed up computation [99]. For the AFP spectrometer, a dedicated fast simulation was developed which simulates the transport of truth-level protons from the MC simulation through the LHC beam optics to the AFP spectrometer, before simulating the subsequent detection with smearing applied. The detection can then be simulated at various levels, either down to silicon hit clusters or tracks, from which the protons are reconstructed. AFP proton reconstruction is discussed in more detail in Section 4.3.

Following reconstruction, several weights are applied to simulated samples, to correct for residual differences between simulation and data, such as the distribution of mean interactions per bunch crossing (pile-up reweighting), trigger efficiencies, and the momentum scale of leptons, as discussed further in Chapter 3.

Chapter 3

The ATLAS Detector at the Large Hadron Collider

The following chapter presents an overview of the LHC (Section 3.1) and A Toroidal LHC ApparatuS (ATLAS) detector (Section 3.2) used to collect the data used in this thesis, including the reconstruction methods used for the various physics objects which are analysed (Section 3.3).

3.1 The Large Hadron Collider

With a total circumference of 26.7 km, the LHC is the largest particle accelerator in the world. Built around 100 m below the surface of the Swiss-French border at Centre Européen pour la Recherche Nucléaire (CERN), in the tunnel originally constructed for the Large Electron-Positron (LEP) collider, it is a circular synchrotron capable of colliding together protons at centre-of-mass energies up to $\sqrt{s} = 14$ TeV.

Protons begin in the form of hydrogen gas, which is ionised using electric fields to leave only the protons. These are then linearly accelerated using electric fields by Linear Accelerator 2 (LINAC2)[1] up to 50 MeV. From here they are injected into the

[1] LINAC2 (50 MeV) served as the first acceleration step of CERN's experiments for 40 years

first circular accelerator stage, the Proton Synchrotron (PS) booster, for acceleration up to 1.4 GeV, and then the PS which accelerates the protons further to 26 GeV. Next, the protons move into the Super Proton Synchrotron (SPS) where they are accelerated up to 450 GeV, and finally injected into the LHC in the form of two opposing beams. Here the beams of protons continue to be accelerated until they reach up to 99.9999991% the speed of light, corresponding to a maximum energy per proton of 7 TeV, thus giving centre-of-mass energies for collisions up to 14 TeV. The LHC is also capable of performing nucleon-nucleon collisions at centre-of-mass energies of up to 5 TeV, which are accelerated via Linear Accelerator 3 (LINAC3) and the Low Energy Ion Ring (LEIR), before being sent to the PS and following the same steps as protons before injection into the LHC. The full CERN accelerator complex as of 2022 is shown in Figure 3.1.

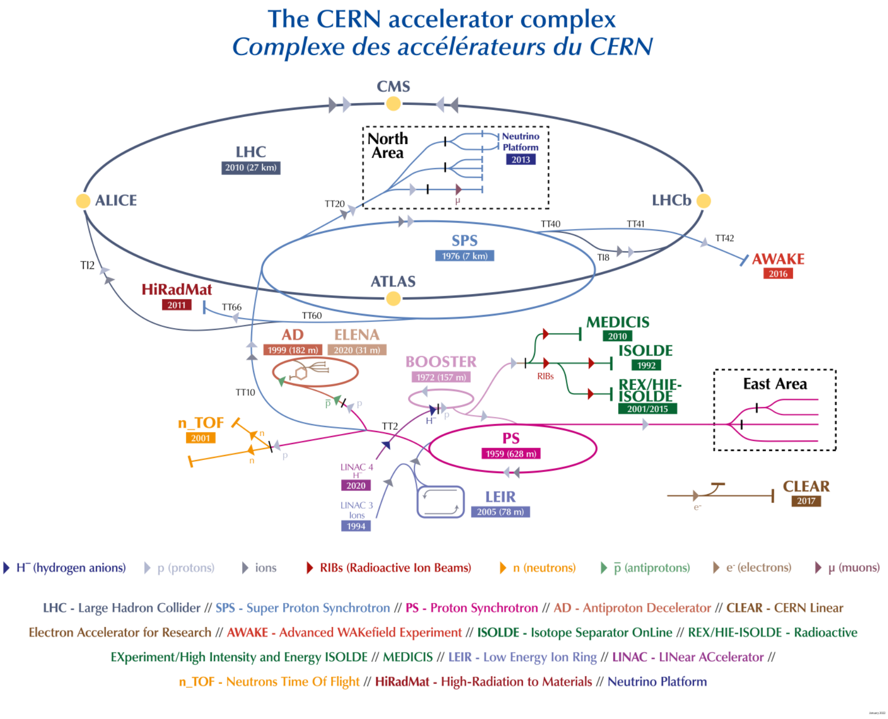

Figure 3.1: The full CERN accelerator complex as of 2022 [100].

The circular trajectory of the proton beams in the LHC is maintained using 1,232 dipole magnets in the form of superconducting electromagnets with niobium-titanium

until it was switched off in 2018; for Run 3 onwards the new Linear Accelerator 4 (LINAC4) (160 MeV) was used.

filaments, which generate a magnetic field of 8.3 T. To induce the superconductivity necessary for the high currents required to produce this field strength (around 11 kA), the magnets are cooled to 1.9 K using liquid helium. In addition, 392 quadrupole magnets are used to focus the beam, squeezing the protons together to create narrow beams, with any sufficiently off-trajectory particles being absorbed by collimators placed around the ring. Like other synchrotron accelerators, the LHC uses Radio Frequency (RF) cavities to accelerate protons. These contain electric fields of 5 MVm^{-1} oscillating (flipping polarity) at a frequency of around 400 MHz. When protons pass through the cavity they are accelerated differently depending on their phase with the oscillating field. Protons which are in-phase (synchronous) with the RF frequency undergo no acceleration, while protons which are out of phase are accelerated (or decelerated) towards the synchronous protons. This leads to the formation of discrete bunches of protons separated by empty space. In the LHC, the proton bunches contain $\sim 10^{11}$ protons and are separated in time by 25 ns (around 10 m physical separation) with a total of 2556 bunches at maximum fill.

At specified locations around the ring, collision points are created by using insertion magnets to cross over the beam paths, causing "bunch crossings" wherever bunches from opposing beams pass through each other. There are 4 such locations, or Interaction Points (IPs), and each one is surrounded by a large particle physics detector. These four experiments are: ATLAS [101] and Compact Muon Solenoid (CMS) [102], which are general-purpose detectors designed to measure as many products of pp collisions as possible, Large Hadron Collider beauty (LHCb) [103], a forward-facing detector designed to study flavour physics and measure CP violation, and A Large Ion Collider Experiment (ALICE) [104], designed to measure the products of heavy ions collisions in the LHC. An important value used to parametrise the extent of squeezing of the beams at the IPs is β^*, the distance at which the beam amplitude is twice that at the IP. This value can be controlled by adjusting the magnet configuration, and in 2017 the LHC typically operated at $\beta^* = 0.3$ m [105].

Data-taking at the LHC is organised into multi-year periods called "Runs". Run

1 took place between 2011-2012 at centre-of-mass energies between 7-8 TeV, and marked the discovery of the Higgs boson [24], Run 2 took place between 2015-2018 at $\sqrt{s} = 13$ TeV, and is the source of all data used for the analysis in this thesis, and Run 3 began in 2022 at $\sqrt{s} = 13.6$ TeV and is planned to continue until mid-2026. Each Run is separated by several-year Long Shutdowns (LSs) to allow for major detector upgrades, and at the end of every year during Runs there are additional breaks called Year-end Technical Stops (YETSs), for minor maintenance and repair work.

3.1.1 Luminosity and Pile-up

In order to study any process with sufficient precision, high statistics, or a large number of events containing that process, are required. The rate $\frac{dN_{\text{pro}}}{dt}$ of any process occurring in the LHC is

$$\frac{dN_{\text{pro}}}{dt} = \sigma_{\text{pro}}\mathcal{L}, \tag{3.1}$$

where σ_{pro} is the cross section of the process and $\mathcal{L}$ is the instantaneous luminosity, which measures the number of colliding particles per unit cross-sectional area per unit time. If we assume identical proton bunches in the LHC beams with Gaussian shape, then the luminosity is

$$\mathcal{L} = \frac{N_b N_p^2 f_{\text{rev}} F}{4\pi\sigma_x\sigma_y}, \tag{3.2}$$

where N_b is the number of bunches per beam, N_p is the number of protons per bunch, f_{rev} is the revolution frequency of the beams, F is a geometrical factor to account for the beam crossing angle (where $F = 1$ indicates a head-on collision) and $\sigma_{x/y}$ are the transverse sizes of the beam at the IP [106]. The LHC has a design luminosity of $1 \times 10^{34}\text{cm}^{-2}\text{s}^{-1}$, however during parts of Run 2 and most of Run 3 it was able to operate at up to 2 times this value [107, 108]. The total amount of data collected over a given time period T is given by the total integrated luminosity L

$$L(T) = \int_0^T \mathcal{L}(t)dt \tag{3.3}$$

and is typically measured in units of fb^{-1}, with the LHC delivering a total of 156 fb^{-1} during Run 2. During operation at such a high luminosity, the probability of having more than one pp interaction per bunch crossing is very high. The mean number of inelastic interactions per bunch crossing is given by

$$\langle \mu \rangle = \frac{\mathcal{L}\sigma_{\mathrm{inel}}}{N_b f_{\mathrm{rev}}}, \tag{3.4}$$

where σ_{inel} is the inelastic cross-section in a pp collision, around 80 mb for 13 TeV collisions. The distribution of μ is shown in Figure 3.2 for 2017 data-taking, when all the data used in this thesis were collected, showing an average of 37.8 with a peak of over 70. In any given bunch crossing at most one hard scatter interaction

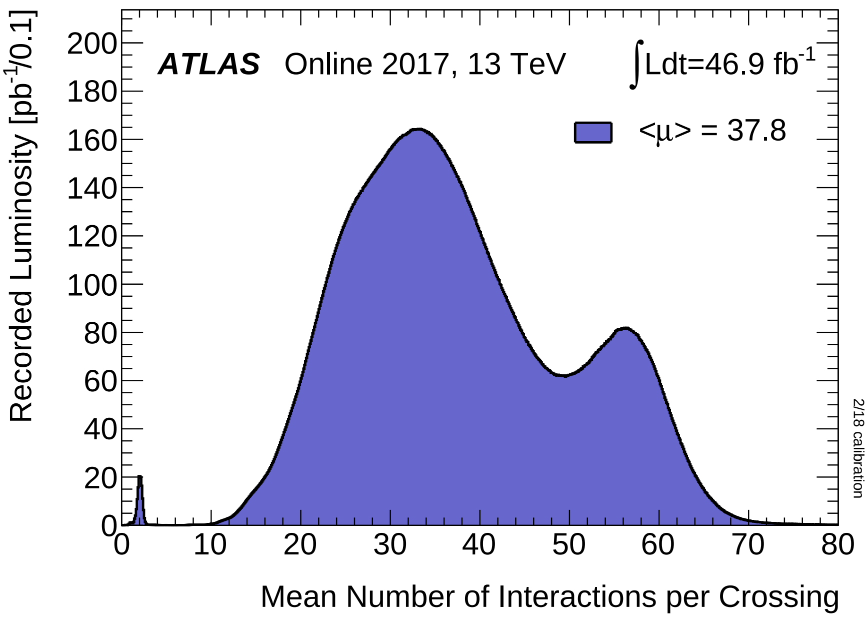

Figure 3.2: Luminosity-weighted distribution of the mean number of interactions per crossing for the 2017 pp collision data at 13 TeV centre-of-mass energy. [109]

is expected with enough interesting features to warrant closer investigation, with the remaining interactions being common soft processes. These are referred to as pile-up interactions.

The long-term plan for the LHC is shown in Figure 3.3 and currently extends until 2040. In particular, the LHC is due to undergo a major upgrade during LS3 between

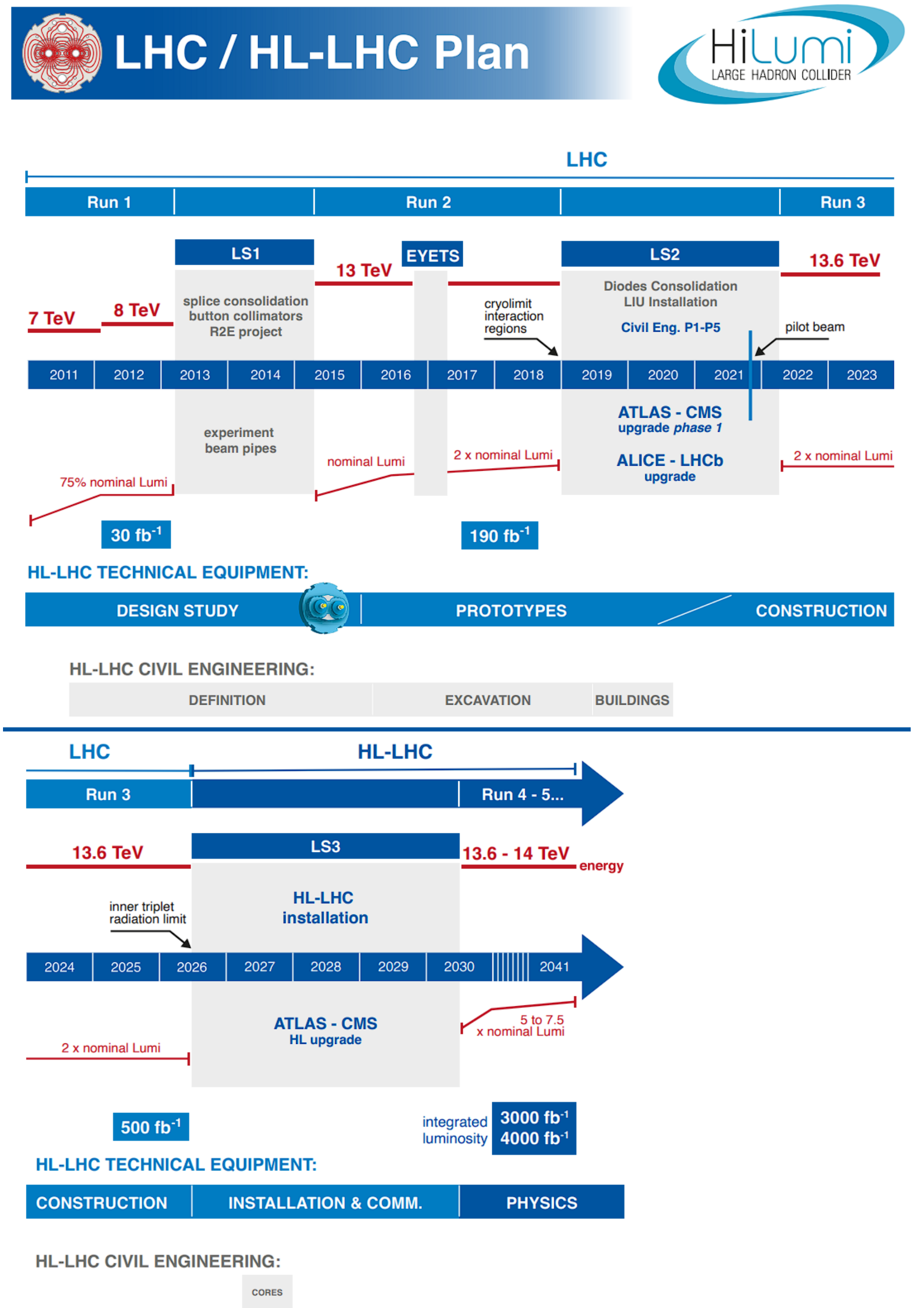

Figure 3.3: Long term schedule for the LHC and future High Luminosity LHC (HL-LHC) as of January 2025 [110].

Runs 3 and 4 (with the latter due to start in mid-2029) to the High Luminosity LHC (HL-LHC). This upgrade plans to increase the levelled instantaneous luminosity at

$\sqrt{s} = 14$ TeV to $5 \times 10^{34}\text{cm}^{-2}\text{s}^{-1}$, with a potential peak of $7 \times 10^{34}\text{cm}^{-2}\text{s}^{-1}$. This will correspond to pile-up of μ~200 and allow the collection of an annual integrated luminosity of around 250 fb^{-1}, with the goal to reach a total dataset of 3000 fb^{-1} by the end of Run 5. This will require extensive upgrades to all of the detectors around the LHC, including the ATLAS detector whose inner tracker upgrade is discussed in detail in Section 5.2.

3.2 The ATLAS Experiment

The ATLAS detector is the largest of the four main detectors placed around the LHC. It is a general purpose particle detector, designed to measure as many different types of particles and corresponding processes as possible. It is cylindrical in shape, forward-backward symmetric and covers almost the entire solid angle around the IP, the position at the centre of the detector where proton-proton collisions take place, allowing it in principle to detect almost every particle produced in these collisions. The detector is composed of a series of concentric layers, each designed to detect a different component of the products of particle interactions. Starting from the IP these layers are:

- the Inner Detector (ID) - designed to measure the trajectories and momenta of charged particles
- the Electromagnetic Calorimeter (ECal) and Hadronic Calorimeter (HCal) - designed to measure the energy of EM and hadronic particles, respectively
- the Muon Spectrometer (MS) - designed to measure the momentum of muons

In addition there is a large solenoid magnet placed between the ID and ECal as well as several toroidal magnets between the HCal and MS, to cause bending in the trajectories of charged particles, allowing their momenta to be determined. These systems are shown together in Figure 3.4 and are all described in more detail in

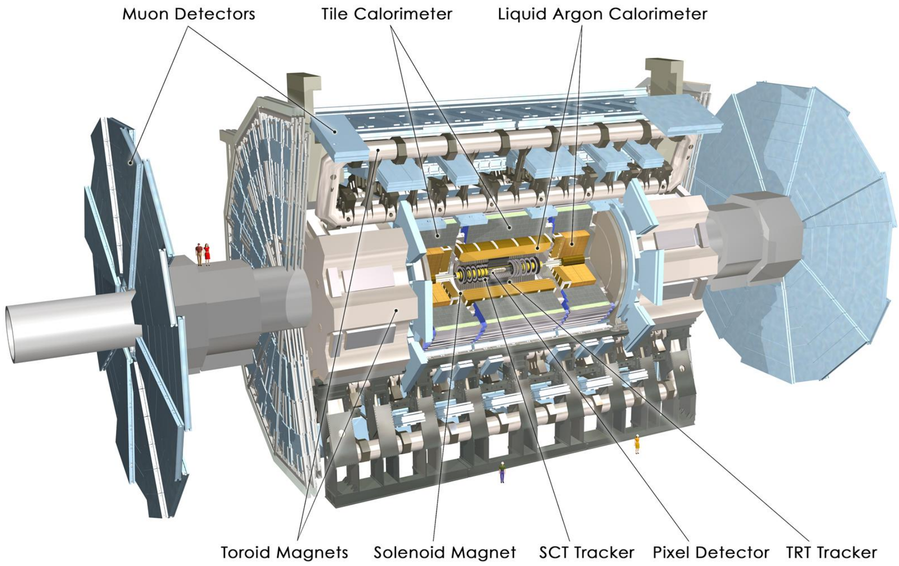

Figure 3.4: Layout of the ATLAS detector [111].

the following sections. The analysis presented in this thesis uses data collected entirely in 2017 during Run 2, so the sections below describe the detector state during this period. During a given year data-taking in the ATLAS experiment is split into Runs lasting on the order of one day (not to be confused with the longer LHC Runs), which are then grouped into Periods lasting on the order of one month. Each run is also split into smaller sections called luminosity blocks, lasting on the order of one minute.

The ATLAS detector has a coordinate system defined with the origin at the IP and the z axis parallel to the beamline. The x axis is horizontal with the positive direction towards the centre of the LHC, and the y axis points vertically upwards. Overall the ATLAS detector uses a right-handed coordinate system, such that the $+z$ direction corresponds to clockwise around the LHC. Due to the cylindrical shape of the detector, it can also be useful to define a cylindrical coordinate system, such that the azimuthal angle ϕ is the angle in the x–y or transverse plane, and the polar angle θ is the angle with respect to the z axis. We can also define the transverse momentum in the x–y plane as $p_{\mathrm{T}} = \sqrt{p_x^2 + p_y^2}$. At hadron colliders, it is more convenient to use the pseudorapidity η in place of θ, defined as

$$\eta = -\ln\left(\tan\frac{\theta}{2}\right). \tag{3.5}$$

This is because the difference $\Delta\eta$ between two particles is invariant under Lorentz transformations, which greatly simplifies translations in the centre-of-mass frame from that of the colliding protons, which are typically balanced in momentum, to that of the colliding partons, which are typically asymmetric. The angular distance ΔR between two particles is then defined as

$$\Delta R = \sqrt{\Delta\eta^2 + \Delta\phi^2}. \tag{3.6}$$

3.2.1 Inner Detector

The innermost section of the ATLAS detector is the ID, designed to measure the trajectories and momenta of traversing charged particles with high precision. The ID is cylindrical in shape, 7 m long with an inner (outer) radius of 33 (1150) mm. There are three concentric sub-detectors forming the ID, in order from the beamline these are the pixel detector, the Semi-conductor Tracker (SCT) and the Transition Radiation Tracker (TRT); with each layer comprised of a barrel section covering the region perpendicular to the beamline and two end-caps, one either side of the IP, to cover the forward region close to the beamline. Together these systems provide good momentum resolution for charged particle tracks with $p_{\mathrm{T}} > 500$ MeV[2] and cover $|\eta| < 2.5$. The layout of the ID is shown in Figure 3.5. The entire system is surrounded by a solenoid magnet, which uses indirectly-cooled aluminium-stabilised superconductor technology to generate a field of 2 T. This field bends the trajectories of charged particles, allowing their momenta to be determined by the angle of the bend and their charge by the direction of the bend. The solenoid is only 44 mm thick to minimise its effect on particle energies [113].

[2] Although lower momentum tracks down to $p_{\mathrm{T}} > 100$ MeV can also be measured and reconstructed, as discussed in Section 6.5.3.

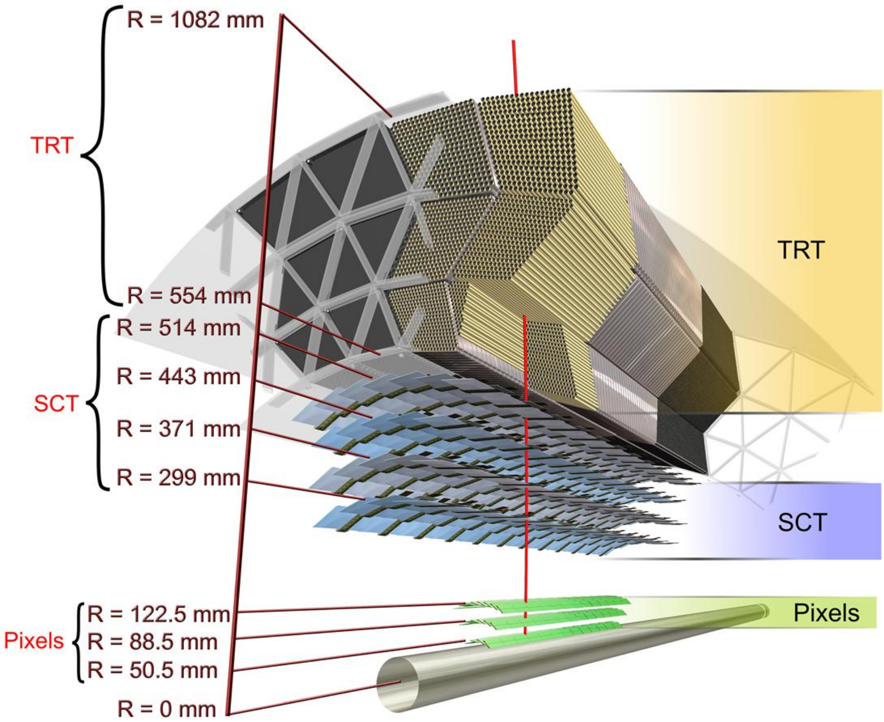

Figure 3.5: Layout of the ATLAS Inner Tracking Detector composed of the Pixel Detector, SCT and TRT [112].

The innermost two layers of the ID utilise silicon tracking technology. The related physics concepts are explored in detail in Section 5.1. The layer closest to the beamline is the pixel detector, with an inner radius of 33 mm, which uses silicon pixel sensors to very precisely measure charged particle trajectories close to the IP, allowing the Primary Vertex (PV) of interactions to be reconstructed. There are four concentric layers of pixel sensors, with the innermost layer, called the Insertable B-Layer (IBL), being added in Run 2 to improve vertex resolution and particle identification for short-lived particles. This is particularly key for B-hadron decays, hence the name [114]. In addition there are three end-cap layers stacked along the z-axis on either side. The three layers around the IBL use planar pixel sensors, with the electrodes implanted on the surface of the silicon. The IBL also incorporates more advanced 3D pixel sensors which use column-like electrodes which penetrate the substrate, giving a faster response time and increased radiation hardness, to better handle the high track density this close to the IP. The same 3D pixel sensors are used in the ATLAS Forward Proton (AFP) spectrometer and so these are explained in more detail in Section 4. In total, there are around 92 million pixel sensors, spread across 1736 barrel modules and 288 end-cap modules. Each pixel is $50 \times 400\,\mu\text{m}^2$ in the external layers and $50 \times 250\,\mu\text{m}^2$ in the IBL, giving a total silicon area of 1.9 m^2. Due to their small size in both dimensions, each pixel can independently make a precise 2D measurement of particle position. Charged particles are expected to leave four hits in the detector, including one in the IBL, yielding an overall precision of ~10 μm [115].

The next layer is the SCT, which uses silicon sensors to measure charged particle position in the same way as the pixel detector. However, since the SCT covers a much wider surface area, in order to minimise cost and material budget less precise strip sensors are used. The SCT is formed of 4 concentric barrel layers and 18 end-cap disks, 9 on either side of the IP, and extends out to 610 mm from the IP. Each strip detector is $6.36 \times 6.40\ \text{cm}^2$ with 768 readout strips and 80 μm between each strip. Therefore, the strip sensors can only independently measure in a single dimension with high precision. To compensate for this, each strip sensor module has two sets of strip sensors glued back-to-back, with a 40 mrad angle between them, such that the combination of the 1D measurements from two adjacent strips provides a precise 2D measurement. There are 4,088 two-sided modules in the detector, with over 6 million total strips, for a total silicon area of ~60 m^2. Charged particles are expected to leave up to eight hits in the detector, for four position measurements, giving an overall precision of 25 μm at best. The silicon sensors in both the pixel detector and SCT are cooled to -10°C to reduce electronic noise and prevent high leakage currents which could damage the sensors [115].

The final layer in the ID is the TRT, which uses a different technique from the other layers to measure particle trajectories, as well as to provide some Particle Identification (PID). This detector is made up of a large number of thin-walled drift tubes, also known as ‘straws”, with a diameter of 4 mm. Each straw is filled with a gas mixture of Xe, CO_2 and O_2, and has at its centre a 0.03 mm diameter gold-plated tungsten wire. When a charged particle enters the gas volume, it will ionise the nuclei, freeing electrons which then drift in an electric field applied across the straw to be collected and read out by the central wire. In addition, polypropylene-polyethylene foils are placed between adjacent straws to induce transition radiation, which is the emission of a photon caused when a charged particle passes between two different dielectric media. These photons then induce an electrical signal via the xenon gas in a similar manner to the tracking hits. The intensity of this radiation is proportional to the Lorentz γ factor of the parent particle, allowing particles of different masses to be distinguished. This is particularly effective in distinguishing between electrons and light hadrons such as pions. The detector is comprised of a single barrel layer containing around 50,000 straws, each 144 cm long and divided in two at its centre to give reduced occupancy and faster readout, and two end-caps with 320,000 straws between them, each 39 cm long with readouts positioned at the outer ends of the straws. Each straw provides a position measurement based on drift time, with a resolution of 170 μm, as well as two independent energy thresholds. This allows lower energy tracking hits to be discriminated from the higher energy transition radiation hits, whose energies also give PID information. The straws are arranged optimally to give an average of 36 hits per charged particle. The weakness of the TRT system is the relatively slow drift time of the straws, creating high dead time such that the detector does not cope well with high track density [115].

3.2.2 Calorimeters

The next section of the ATLAS detector is the calorimeters, which perform "destructive"[3] measurements of particle position and energy, allowing for PID as well as the reconstruction of the Missing Transverse Energy (MET) in an event. The calorimeters are split into two sections, first the ECal, which measures electrons and photons, and second the HCal, which measures hadrons. There is an addi-

[3]A destructive measurement is one which affects the properties of the measured particle, such that no further measurements of the original particle properties can be made.

tional Forward Calorimeter (FCal) as an extension to both calorimeters in the very forward region. The calorimeters are often referred to instead by their detector technology as the Liquid Argon (LAr) calorimeter, made up of the ECal, HCal endcaps and the FCal, and the Tile calorimeter made up of the HCal barrel. Overall, the LAr calorimeter is 13.3 m in length with an outer radius of 2.25 m, while the Tile calorimeter is 12.2 m in length with an outer radius of 4.23 m. Both calorimeters use sampling calorimetry, with alternating passive and active layers. The passive, or absorbing, layer is formed of a high-density material which will interact with incident particles to induce showering, while the active layer collects and measures the energy of showering particles which are not absorbed by the passive layer. The aim of the calorimeter is to completely stop all measured particles within the calorimeter volume, to allow their entire energy to be measured. The calorimeters also form part of the Level 1 (L1) hardware trigger and therefore require very fast readout. The layout of the calorimeters is shown in Figure 3.6.

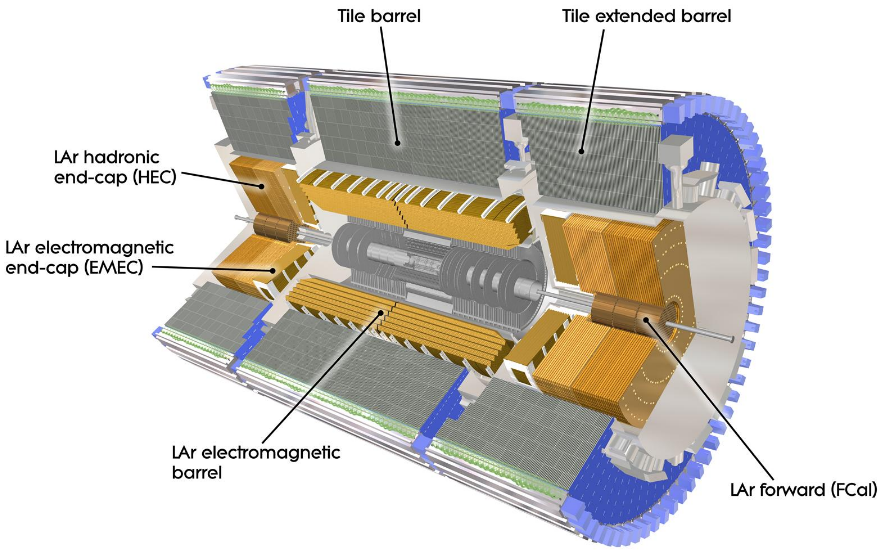

Figure 3.6: Layout of the ATLAS calorimeters composed of the LAr and Tile Calorimeters [116].

The innermost calorimeter, the ECal, is designed to measure the energy of electromagnetic particles, in particular electrons and photons. It is formed of a barrel layer and inner and outer end-caps, which together cover $|\eta| < 3.2$. The ECal uses liquid argon as the active medium for its fast, uniform response and radiation hardness, with lead absorbing layers. The lead induces EM showering from incident particles with $E \gtrsim 1$ GeV via bremsstrahlung, in which photons are emitted from electrons, which then undergo electron-positron pair production ($\gamma \rightarrow e^+e^-$), with the produced leptons emitting further, less energetic photons and so on. The shower continues until ionisation becomes the dominant energy loss mechanism for electrons and the photons are absorbed. The electrons ionise argon atoms in the active layers, producing free electrons which drift in an applied electric field. This induces an electric current, proportional to the energy of the shower particle, within cables contained in vacuum-sealed cylinders running throughout the liquid argon. By combining the energy measurements from all particles within a given shower, the energy of the initial particle can be determined. The layers of the ECal are arranged in an accordion structure to maximise coverage and allow for shorter cables, reducing dead-time, with higher granularity layers positioned closer to the beampipe. A schematic of a single barrel LAr module is shown in Figure 3.7. The barrel section covers the central region ($|\eta| < 1.52$), while the inner and outer end-caps cover the forward region ($1.375 < |\eta| < 2.5$ and $2.5 < |\eta| < 3.2$, respectively). There is an additional thin layer called the presampler on the inside of the innermost barrel and end-cap layers covering $|\eta| < 1.8$, which corrects for the energy lost by particles in material before the calorimeter, mainly in the inner detector. The transition region between the barrel and end-caps ($1.37 < |\eta| < 1.52$, the so-called calorimeter "crack") is used to pass through various service pipes and electronics, and therefore has reduced performance and is sometimes excluded from analyses. Overall, the ECal is between 24-26 X_0 thick, where X_0 is the "radiation length" of the medium, the average distance an electron can travel through it before losing $1/e$ of its energy. This is sufficient to minimise the number of shower particles which escape the ECal without being detected. To maintain the argon in its liquid state, the entire ECal is contained within a cryostat which cools the argon to 89 K. This cryostat is 6.8 m long, with inner and outer radii of 1.15 m and 2.25 m, respectively [117].

The next layer is the HCal, designed to absorb and measure almost all particles

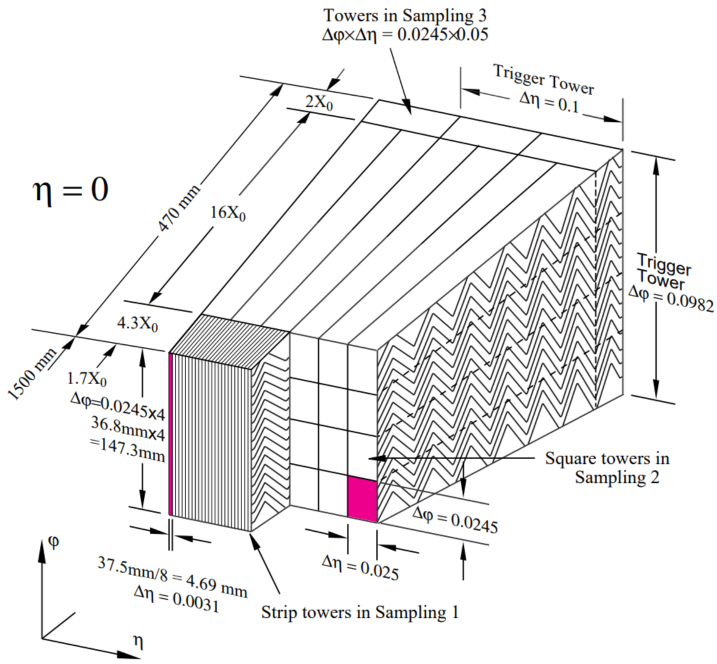

Figure 3.7: Cross section of a barrel module from the LAr calorimeter [117], with three layers of decreasing granularity.

which penetrate the ECal, primarily hadrons. It is composed of a barrel section, covering $|\eta| < 1.7$, and two end-caps covering $1.5 < |\eta| < 3.2$. The HCal barrel is referred to as the Tile calorimeter, as it uses 3 mm thick plastic scintillating tiles as the active layer, along with 14 mm thick layers of steel as the absorber. Hadronic showers are induced when incident particles scatter off nuclei within the steel, creating a shower of additional hadrons which undergo further nuclear interactions or ionisations, including an EM component from the decay of neutral hadrons (e.g. π^0). The shower products then interact with the scintillators to produce UV photons, which are collected by optical fibres. The fibres shift the wavelength of collected photons into the visible spectrum, before delivering them into photomultiplier tubes which amplify them, creating an electric current which can be calibrated to measure the original particle energy. A schematic of a single barrel tile module is shown in Figure 3.8. The tile calorimeter is formed of three concentric layers, a central barrel which is 5.6 m long and contains 64 modules, followed by two extended barrels which are 2.6 m long and contain 64 modules each, giving a total of 420,000 tiles. The tiles in each layer are arranged perpendicular to the beamline and their relative depths are staggered to increase acceptance. To cover the forward region there are Hadronic End Caps (HECs) on either side of the detector, each formed of two wheels using the same LAr technology as the ECal, but with higher-density tungsten absorbing layers to decrease the required volume of the detector in the forward region where there is not much space. Overall, the HCal is ~11λ thick at $\eta = 0$, where λ is the nuclear interaction length, defined as the mean distance a hadronic particle will travel in a given material before undergoing an inelastic nuclear collision. This is sufficient to contain the entire shower produced by an incident particle and minimise the number of particles escaping the HCal without detection, which is known as "punch-through". Similar to the ECal, there is a crack region at $|\eta| = 1$ with reduced performance due to electronics passing between the detector sections [117, 118].

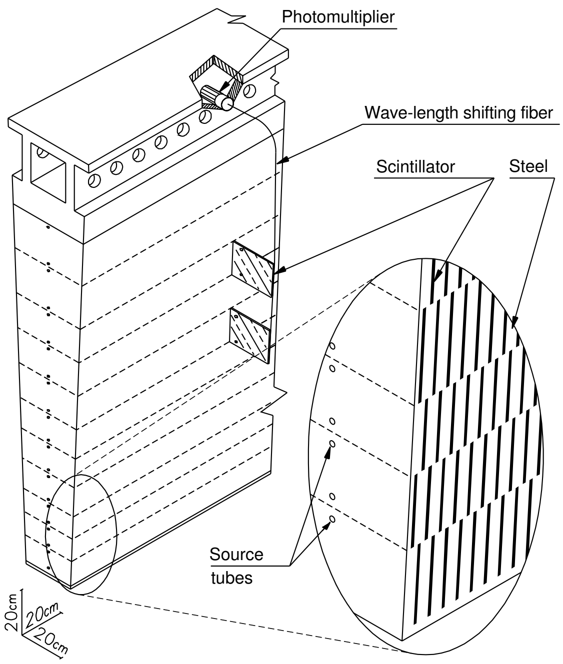

Figure 3.8: Schematic of a barrel module from the Tile calorimeter [119].

To cover the very forward region closest to the beamline between $3.0 < |\eta| < 4.9$, additional FCals are used. Positioned 4.7 m from the IP on either side of the cryostat containing the ECal, each FCal consists of three layers all using the same LAr technology as the ECal. The first layer uses a copper absorber, designed to measure EM showers, and the remaining two layers use a high-density tungsten absorber to measure hadronic showers. Each layer is made up of a metal matrix with regularly spaced longitudinal channels filled with concentric rods and tubes, with liquid argon placed in gaps between the layers, which are as small as 250 μm [117].

3.2.3 Muon Spectrometer

Muons interact more weakly with matter than electrons due to being 200 times more massive, therefore undergoing much less bremsstrahlung, losing less energy and penetrating deeper into material. For this reason, muons typically pass through the calorimeters without being stopped, and so the final layer of the ATLAS detector, the MS, is specifically designed to identify muons and precisely measure their momenta. The MS is made up of Monitored Drift Tubess (MDTs) in the central region and Cathode Strip Chambers (CSCs) in the forward region. In addition, the MS forms part of the L1 hardware trigger, using a combination of Resistive-Plate Chambers (RPCs) in the central region and Thin-Gap Chambers (TGCs) in the forward region. These systems are less precise than the primary MS systems, but have much faster readout as required for triggering, in addition to providing a complementary measurement on muon trajectory. To facilitate the measurement of muon momenta, their trajectories are curved using a system of three large superconducting air-core toroid magnets. In the central region ($|\eta| \leqslant 1.0$) magnetic bending is provided by a large barrel magnet constructed from eight coils which surround the HCal. In the forward region ($1.4 \leqslant |\eta| \leqslant 2.7$), two smaller end-cap magnets are used, which are inserted into both ends of the barrel toroid. In the interval between these regions ($1.0 \leqslant |\eta| \leqslant 1.4$) referred to as the transition region, muon bending is provided by a combination of the two magnetic fields. The toroid system creates a field mostly perpendicular to the muon trajectories, while minimizing the amount of multiple scattering which reduces the resolution of the momentum measurement. The layout of the muon spectrometer and toroid magnets is shown in Figure 3.9. Overall, the muon spectrometer covers up to $|\eta| < 2.7$, measuring points on the muon trajectory to a precision of < 10 μm, corresponding to a momentum resolution of 2-3% across the majority of the muon kinematic range, up to around 10% for $p_T \sim 1$ TeV.

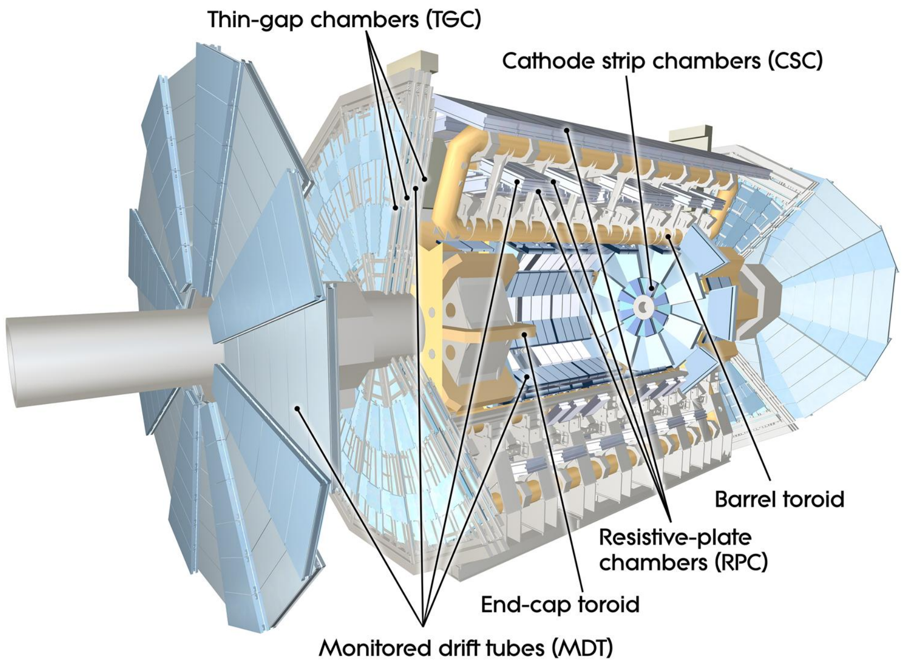

Figure 3.9: Layout of the ATLAS Muon Spectrometer [120].

The MDTs measure muon tracks in the principle bending direction of the toroids, and are formed from 3 cm wide aluminium tubes filled with a gas mixture (CO_2 and Argon), with 50 μm gold-plated tungsten-rhenium wires running down their centre. When muons pass through the gas they ionise atoms, freeing electrons which drift in an applied electric field to the centre of the tubes, inducing an electric current in the wires. They are arranged in 3-8 concentric layers of chambers in the MS barrel in order to maximise the hit rate, with the tubes varying from 0.85-6.5 m in length. In total there are 1,171 chambers with a total of 354,240 tubes, with each tube having a resolution of 80 μm, which corresponds to 35 μm per chamber when the gas is held at a pressure of 3 bars. The CSCs are arranged in three layers of chambers stacked in the z axis in the end-caps at either end of the detector to precisely measure muon tracks in the forward region between $2.0 \leqslant |\eta| \leqslant 2.7$. They use the same detection principle as the MDTs, with the same gas mixture and wires, and use cathode strips as the readout which are arranged orthogonally on adjacent layers. Overall there are 31,000 channels spread across 32 chambers per end-cap, with each chamber giving a resolution of 40 μm and 5 mm in the bending and transverse planes of the toroids, respectively. The difference between the two planes arises from the different readout pitch, and the azimuthal readout running parallel to the anode wires [121]. They also have a fast electron drift time of $\leqslant$ 30 ns giving a timing resolution of 7 ns, and have low neutron sensitivity achieved by minimising the gas volume used, which is needed for the high radiation hardness required close to the beampipe.

The maximum drift time of the MDTs is around 500 ns, much longer than the 25 ns bunch-crossing period for LHC, so additional detectors with faster readout are required to provide triggering. To cover the central region of $|\eta| < 1.05$, the RPCs are formed from pairs of parallel plastic resistive plates with a potential difference held across them, separated by a 2 mm gap filled with a gas volume (mostly tetrafluorethane). Muons passing though the detector ionise atoms in the gas volume, freeing electrons which undergo avalanche multiplication due to the applied electric field between the plates. This generates a signal which is read out via capacitative coupling to metal strips placed at both sides of the detector. There are a total of 380,000 channels spread across 606 chambers, providing overall space and time resolutions of 1 cm and 1 ns, respectively. To perform muon triggering in the forward region between $1.05 \leqslant |\eta| \leqslant 2.4$, TGCs are used. These are similar to the CSCs but with faster readout, using parallel 30 μm thick wires in a gas mixture of CO_2 and n-pentane between cathode plates. Incident muons ionise the gas, and the resulting free electrons drift in an applied electric field to produce a current in the wire. There are a total of 318,000 channels in 3,588 chambers in the TGC. Both types of trigger chamber also provide a “second-coordinate” measurement of track coordinates which is orthogonal to the precision measurement, and roughly parallel to the magnetic field. This is required to allow for offline track reconstruction [122].

3.2.4 Forward Detectors

There are several additional detectors placed much further down the beampipe from the ATLAS detector. These are designed to detect particles with very forward trajectories (high-$|\eta|$) which fall outside the acceptance of all the detectors previously described. Their relative positions along the beamline are shown in Figure 3.10.

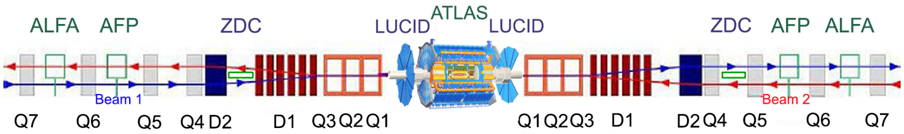

Figure 3.10: The positions of the four forward detectors of the ATLAS experiment, used to measure particles with high-$|\eta|$ and make luminosity measurements [123].

The closest forward detector to the ATLAS detector is Luminosity Cherenkov Integrating Detector (LUCID), positioned 17 m from the IP on either side and covering $5.6 \leqslant |\eta| \leqslant 5.9$. This detector is designed to measure the integrated and per-bunch instantaneous luminosity in the ATLAS detector in real-time, based on the number of inelastic *pp* collisions occurring in the ATLAS detector. This is achieved on each side of the detector by an aluminium vessel filled with C_4F_{10} gas, along with twenty 1.5 m long aluminium tubes angled towards the IP. Incident particles pass through the gas causing it to radiate Cherenkov photons at an angle, which are reflected along the tubes and collected by quartz photomultipliers to create a signal. LUCID has a time resolution of a few ns, which is sufficient to separate individual LHC bunch crossings in order to make instantaneous luminosity measurements [124, 125].

Next along the beamline is the Zero Degree Calorimeter (ZDC), placed 140 m away from the IP on both sides at the point where the two LHC beam pipes split apart, having been merged for interactions in the ATLAS detector. This allows the detector to be placed directly in the pathway of neutral particles with $|\eta| \geqslant 8.3$ (including $\theta = 0$, hence the name). The ZDC is designed to measure neutral particles such as π_0 and neutrons by using a series of tungsten calorimeter modules, with a single module specialised for EM calorimetry and three modules specialised for hadronic calorimetry on either side. The detector is used exclusively for heavy ion runs and can also provide luminosity measurements to complement LUCID [124, 126].

The next forward detector is the AFP spectrometer, which is used to measure the energy loss of intact protons which are scattered due to diffractive or photon-induced interactions in the ATLAS detector [127]. This detector is the basis for the analysis chapter of this thesis, and therefore its description is covered in detail in Chapter 4.

The final forward detector is Absolute Luminosity For ATLAS (ALFA). This is a Roman-pot based detector, like the AFP spectrometer, with two stations inserted vertically into the beampipe either side of the IP at distances of 237 m and 245 m. These stations use scintillating fibres to measure the tracks of scattered protons, in order to provide an absolute luminosity measurement for the ATLAS experiment based on knowledge of the elastic cross section $pp \to pp$. The data can also be used for elastic and soft diffractive cross section measurements [124, 128].

3.2.5 Trigger and DAQ

Proton bunch crossings occur in the ATLAS detector every 25 ns, as discussed in Section 3.1, giving a rate of 40 MHz. Based on the average number of interactions per bunch crossing of $\mu = 37.8$ observed in 2017 [109] this corresponds to an average of 1.5×10^9 pp collisions per second. If every one of these events were recorded it would use around 50 TB of storage space per second. This is impossible to achieve, but also unnecessary since the vast majority of these events are soft processes, which are generally not of interest for physics analyses. A system is therefore required to filter down the millions of events per second to around a thousand, by making selections on event properties which suggest the presence of interesting or rare processes. This is achieved with a two-level Trigger and Data Acquisition (TDAQ) system, made up of an extremely fast hardware level trigger and a slower but more precise software trigger. The components of these systems are shown in Figure 3.11. A trigger "menu" is defined, which lists a series of requirements for events to pass the trigger, with each one aimed towards certain types of interesting processes. Around 1500 event selections were available in the Run 2 trigger menu.

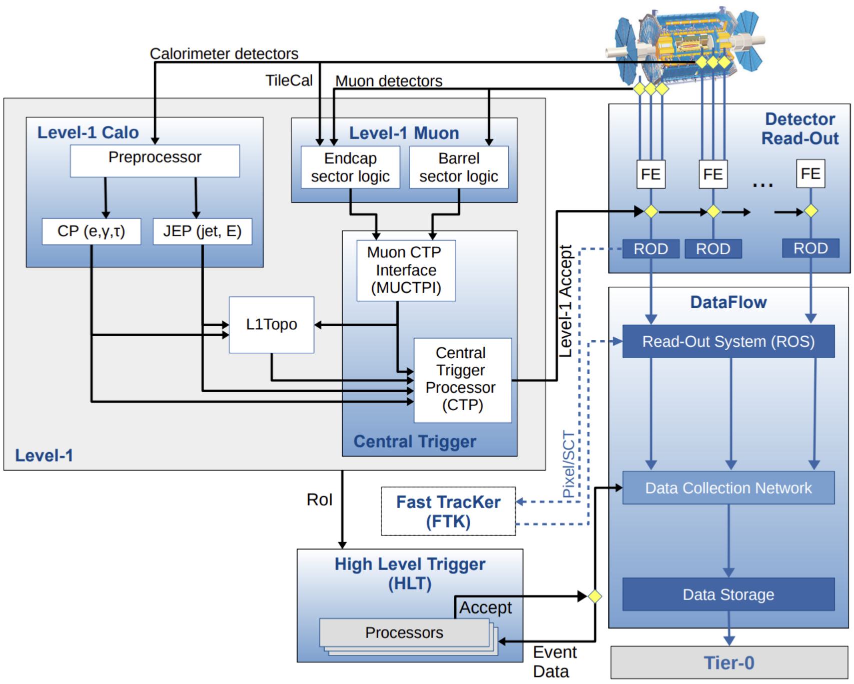

Figure 3.11: The ATLAS TDAQ system in Run 2 showing the components relevant for triggering as well as the detector read-out and data flow [129].

The first level, called L1, is a hardware trigger which reduces the event rate to at most 100 kHz, with a latency of less than 2.5 μs. During this time, the data are pipelined to a system of custom-made electronics which use a combination of information from the calorimeters (L1Calo) and the muon spectrometer (L1Muon). L1Calo identifies the multiplicity and energies of various objects of interest such as electrons, photons and jets [130], while L1Muon estimates the momenta of candidate muons with at least two hits in the MS. These inputs are additionally combined in L1Topo, added for the second half of Run 2, which allows for more complex analysis of objects including the invariant mass and/or angular separation between multiple objects [131]. The geometric locations of potential signals are also determined, called Regions of Interest (RoIs). The information collected by the L1 trigger is passed to the Central Trigger Processor (CTP), which decides whether to accept the event. The CTP also applies pre-scaling, where only a set fraction of certain types of common process are recorded, to keep the trigger rates below the threshold for the second trigger level [132].

The second trigger level, usually referred to as the High-Level Trigger (HLT), is based on software algorithms which are more complex and therefore allow more precise evaluation than the L1 trigger, but at the cost of a significantly increased latency of 200 μs. A farm of 40,000 CPU cores are used to perform optimised versions of the offline object reconstruction algorithms, allowing them to be performed in real-time. These algorithms are grouped into steps to form trigger chains, with each chain seeded by an RoI identified by the L1 trigger. The HLT reduces the event rate to a few kHz, which is sufficiently low to allow all remaining events to be moved to permanent offline storage.

Due to the extreme time limitations for trigger processing, some precision must be sacrificed, which results in small inefficiencies for events containing objects whose kinematics fall close to the threshold of a given trigger. As mentioned in Section 2.4, in simulated events corrections are made on the rates of physics objects as a function of their kinematics, to match this trigger inefficiency observed in data.

3.3 Physics Object Reconstruction

In order to convert the stream of electrical signals from the various subsystems of the ATLAS detector into usable particle measurements, a series of reconstruction procedures are used to create the standard objects used in offline (not during active data-taking) analysis: tracks, photons, electrons, muons, taus, jets and missing transverse energy $E_{\mathrm{T}}^{\mathrm{miss}}$. Figure 3.12 shows the detection principles for each type of common particle produced in pp collisions as they pass through the layers of the ATLAS detector, with each one having a distinct signature used to identify them in reconstruction.

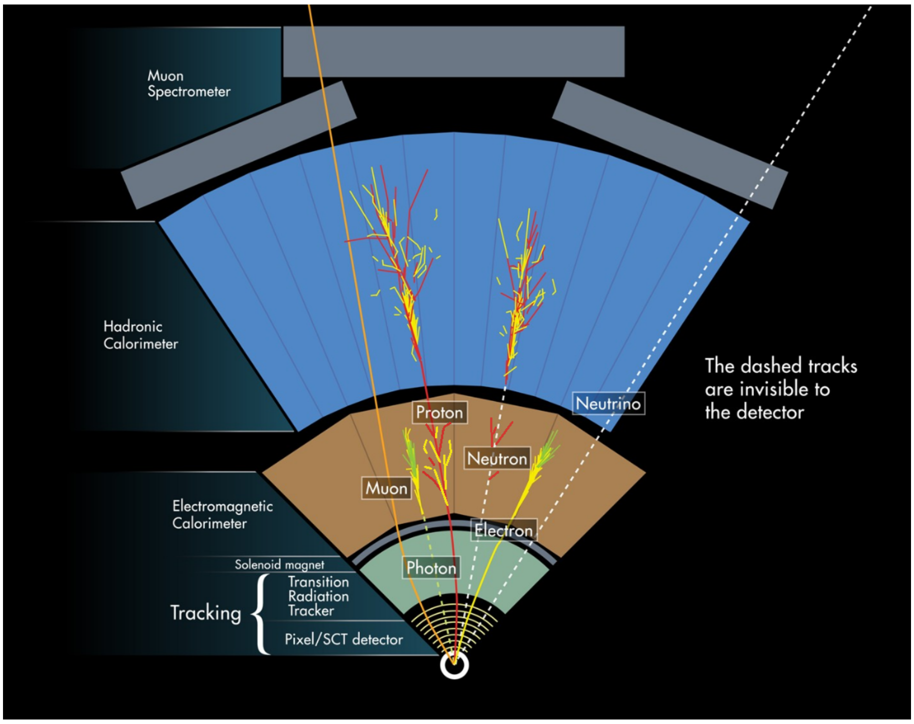

Figure 3.12: Summary of ATLAS experiment reconstruction principles for each particle type [133].

3.3.1 Tracks

Charged particles are detected by the ATLAS ID, leaving hits in each layer which can be combined to reconstruct the trajectory of a particle. A series of reconstructed hits is called a track, and there can be over 1000 tracks per bunch crossing during high pile-up running. Many particles produced in pp collisions decay before reaching the detector, but the ones with sufficient lifetime to leave tracks in the ID are electrons, muons, protons, kaons and pions. The ATLAS experiment uses the perigee representation to parametrise track properties, as shown in Figure 3.13. These are defined relative to a reference point, which is taken from the average position of the pp interactions in an event (the beamspot position). The five parameters are:

- d_0 and z_0 (transverse and longitudinal impact parameters): distances of closest approach of the track in the transverse and longitudinal planes to the reference point
- ϕ and θ: the azimuthal and polar angles of the track momentum at the refer-

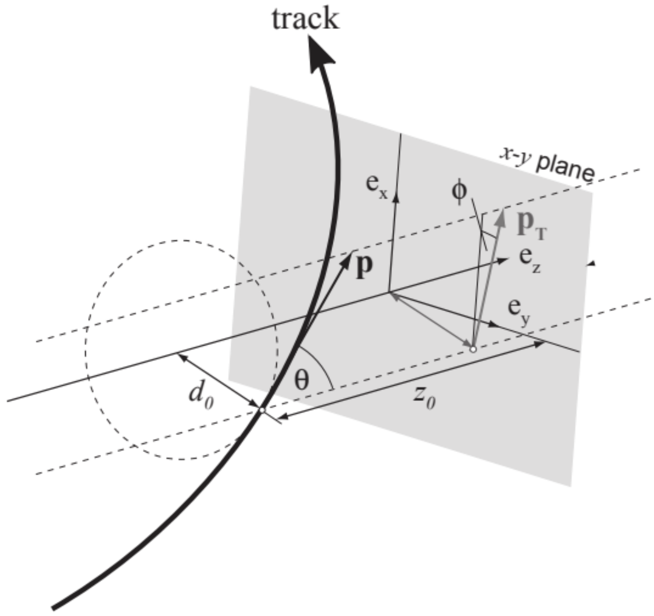

Figure 3.13: Global track parameters with respect to the perigee representation [134].

ence point

- q/p: the ratio of the track charge to the magnitude of its momentum

Tracks are reconstructed using a series of algorithms. First, hits from adjacent channels in the Pixel and SCT detectors are combined into clusters. Pixel clusters already provide a 3D measurement, while pairs of SCT clusters on either side of a sensor module can be combined to form a 3D measurement as discussed in Section 3.2.1; these 3D measurements are called "space-points". Track reconstruction then begins by grouping together triplets of space-points in either the Pixel or SCT which form a consistent track trajectory, these are track seeds (Figure 3.14a). Some loose p_{T} and impact parameter cuts are made on track seeds to remove low-quality tracks early in the process in order to save computation time. Next, sets of detector modules which are expected to contain clusters from track seeds based on their trajectories are built through the rest of the detector to form "search roads" (Figure 3.14b). Track seeds are then extended along these search roads using a combinatorial Kalman filter [135], which searches for adjacent clusters moving in both directions inwards and outwards with respect to the IP, whilst simultaneously smoothing the track tra-

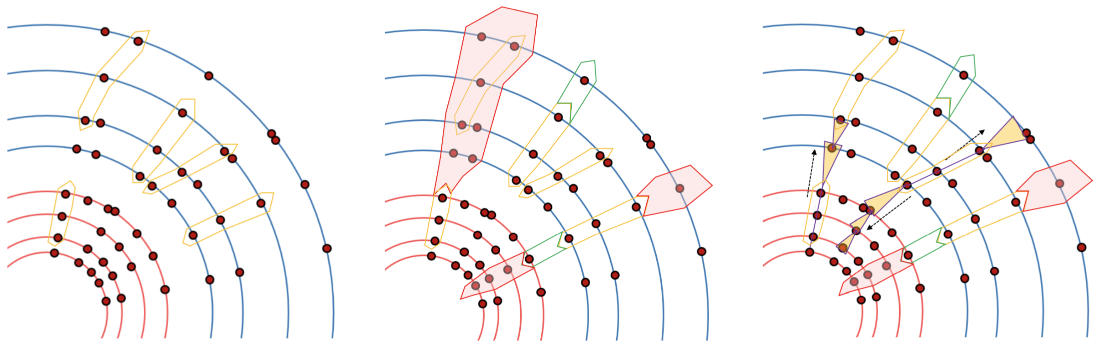

(a) Track seeding (b) Search roads (c) Kalman filter

Figure 3.14: Illustration of the track reconstruction process, with red layers for the pixel detector, blue layers for the SCT and red circles showing Silicon layer hits in the ID. Taken from [134].

jectories (Figure 3.14c). At this stage "bremsstrahlung recovery" is attempted, to account for electron tracks with significant direction changes due to bremsstrahlung emission in the ID. If a track seed fails initial track finding but is within a calorimeter Region of Interest (RoI) then the track finding procedure is repeated, allowing for this additional "kink" in the trajectory.

After the track finding procedure is performed, a set of potential track candidates is produced, however further refinement is still required since the Kalman filter, while computationally fast, is relatively imprecise and doesn't resolve ambiguities such as track overlaps and "fake tracks" (incorrect combinations of unrelated clusters). Track candidates are passed through an algorithm which scores them on a series of quality criteria such as momentum, number of silicon (Pixel or SCT detector) hits, number of shared modules (requires one shared hit for pixel modules and two shared hits for SCT modules) and number of holes (missing hits on active silicon modules in the trajectory of a track). Where several hits are shared between tracks (track overlap) the higher scoring track is kept and others are rejected. Smaller numbers of shared hits are accepted to allow for dense objects such high-energy jets, where track separation may fall below detector granularity. A neural-network based algorithm is then used to update the positions of clusters and corresponding uncertainties, with probabilities assigned for the number of charged particles which have contributed hits towards each cluster. Clusters which are classified as the product of multiple particles are split between track candidates. After the ambiguity resolution is applied, the refined set of track candidates are re-fit using a global χ^2 method.

Finally, an extension into the TRT detector is attempted for each track, following a similar track finding procedure to above, in order to increase momentum resolution and benefit from the PID provided by the TRT as discussed in Section 3.2.1. A search road, seeded by the track candidate, is built into the TRT volume and a Kalman filter is applied. If an extension is successfully performed then the TRT hits are added to the track, which is again refit with a global χ^2 method. If the additional TRT hits lead to a worse fit than before then the original "Si-Only" track is also kept. In addition, if too many of the TRT hits are outliers to the fit or are of too low quality (e.g. tube hits with no leading edge) the track extension is rejected.

The primary reconstruction procedure detailed above is optimised to reconstruct prompt particles produced directly in the primary *pp* interactions. In order to increase acceptance for non-prompt particles produced at a greater distance from the beamline (e.g. electrons produced from photon conversion in the ID) a secondary back-tracking reconstruction step is performed using all detector hits not assigned to tracks in the primary stage. This stage begins with RoIs identified from energy deposits in the ECal with matching hits in the TRT. Pixel and SCT hits close to the trajectory formed by the TRT hits are then grouped to form track seeds, with only two space-points required in this pass compared to three in the primary pass, to allow for shorter secondary tracks. These track seeds are then extended using the same procedure as before (search road, Kalman filter, ambiguity resolution, global χ^2 fit) and then extended back again into the TRT using the original collected hits [134, 136].

Once reconstructed, further cuts can be applied to tracks in order to increase their quality and reduce the rate of fake tracks. This thesis uses the "Loose" track working point which uses the following selections applied by default in the track reconstruction procedure described above:

- $p_{\mathrm{T}} > 500$ MeV
- $|\eta| < 2.5$
- Number of Pixel and SCT clusters on track $\geqslant 7$
- Number of shared modules $\leqslant 1$
- Number of silicon holes $\leqslant 2$
- Number of pixel holes $\leqslant 1$

A tighter selection called "Tight Primary" is also available, which adds the following selections:

- Number of silicon hits $\geqslant 9(11)$ for $|\eta| \leqslant 1.65$ (> 1.65)
- At least one hit on one of the two innermost pixel layers
- No pixel holes

These tightened selections significantly reduce the rate of fake and non-prompt tracks, both improving the accuracy of track reconstruction and suppressing background processes. The efficiency of track reconstruction with each of these working points is shown in Figure 3.15.

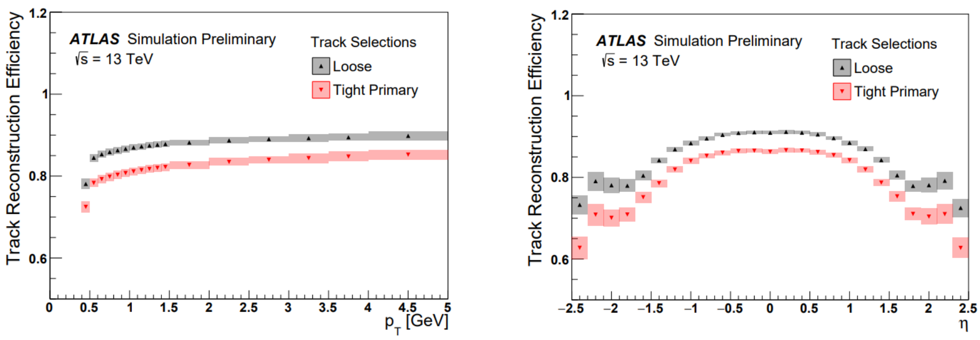

(a) (b)

Figure 3.15: Track reconstruction efficiency for each available working point as a function of (a) p_{T} and (b) η [137].

In simulated MC samples we have access to the truth-level particles producing the tracks, which can then be associated to the reconstructed track clusters wherever in the ID the truth particle deposited energy. The truth-matching probability for a track can then be calculated, which is the likelihood that it resulted from measuring a genuine track found at truth level, as opposed to being a fake track, and is calculated using Equation 3.7:

$$P_{\text{match}} = \frac{10 \cdot N^{\text{common}}_{\text{Pixel}} + 5 \cdot N^{\text{common}}_{\text{SCT}} + 1 \cdot N^{\text{common}}_{\text{TRT}}}{10 \cdot N^{\text{track}}_{\text{Pixel}} + 5 \cdot N^{\text{track}}_{\text{SCT}} + 1 \cdot N^{\text{track}}_{\text{TRT}}}, \tag{3.7}$$

where $N^{\text{common}}_{\text{Pixel, SCT, TRT}}$ are the numbers of hits in each part of the ID which are shared by the track and corresponding truth particle and $N^{\text{track}}_{\text{Pixel, SCT, TRT}}$ are the numbers of hits which form the track. Here each ID hit is weighted according to its importance to track reconstruction. Values of $P_{\text{match}} < 0.5$ suggest a fake track, allowing the rate of fake tracks to be measured for a given selection [137].

3.3.1.1 Vertex Reconstruction

It is very important, especially in high pile-up events, to accurately identify the position of interaction vertices in an event. Generally, there is only one such "Primary Vertex (PV)" per event, where a beam interaction of interest has taken place. Identifying this position helps to isolate the products of the primary interaction from pile-up. Once track reconstruction has taken place, a dedicated vertex reconstruction procedure is performed, which in Runs 1 and 2 used an iterative process whereby track positions are fitted to obtain a seed vertex position and on each iteration less compatible tracks are down-weighted in the fit before the vertex position is recalculated. After this procedure, any vertices with at least two compatible tracks are classified as PV candidates [138]. However, for Run 3 this procedure was replaced with a more complex adaptive multi-vertex fitter to account for the increased pile-up. In the new procedure, initial vertex positions are estimated using track density along the beam axis, with a Gaussian resolution model for the track impact parameter, to estimate the most likely pp interaction vertex position. This position is used as the seed, and all nearby tracks passing a quality selection are assigned to it with a certain weighting. In future iterations tracks can be assigned to multiple different vertices with varying weights. These tracks are then used to fit the vertex position with a weighted adaptive Kalman filter, causing any other linked vertex candidates and corresponding track weights to be refit as well. When a vertex candidate reaches a certain set of thresholds (e.g. isolation from other vertex candidates) it is removed from future consideration, along with any associated tracks. The reduced set of tracks is now refit to generate a new vertex seed and the process is repeated until all seeds have been eliminated from the pool. The final set of vertices are then passed through a selection to reject low quality candidates. This updated method allows tracks to be considered for multiple different vertex candidates, in order to reduce the chance of nearby interactions being incorrectly merged into a single vertex, which becomes more likely in high pile-up conditions [134].

The beamspot is the volume around the IP where the LHC beams cross over and pp interactions are expected to take place. It is reconstructed from an unbinned maximum likelihood fit of the distribution of PVs over a large number of events in around 10 minute intervals of data-taking. Only the PV with the highest $\sum p_{\mathrm{T}}^2$ per event is used. The beamspot is assumed to follow a 3D Gaussian distribution with a typical longitudinal size of 40 mm and a transverse width of $\mathcal{O}(10\,\mu\mathrm{m})$ [138].

3.3.2 Electrons and Photons

Electrons and photons are both detected via EM showers induced and absorbed by the ECal, while electrons additionally leave tracks in the ID due to their electric charge. Reconstructed photons are not used in this thesis so this section will mainly focus on electron reconstruction.

Reconstruction begins by building clusters from energy deposits in the ECal using a topological-cluster algorithm to group neighbouring calorimeter cells if they are significant. At the start of Run 2, a different "sliding-window" approach was used which used fixed-size rectangles to cluster cells, due to limitations in calibration methodology for dynamically-sized clusters [139]. However, since the development of multivariate calibration techniques [140] it has become possible to use this new topo-cluster algorithm, which allows recovery of low-energy bremsstrahlung photons emitted in the ID in electron and photon reconstruction, increasing accuracy [141]. The topo-cluster algorithm selects calorimeter cells based on their significance $\varsigma^{\mathrm{EM}}_{\mathrm{cell}}$, defined as

$$\varsigma^{\mathrm{EM}}_{\mathrm{cell}} = \left| \frac{E^{\mathrm{EM}}_{\mathrm{cell}}}{\sigma^{\mathrm{EM}}_{\mathrm{noise,cell}}} \right| , \tag{3.8}$$

where $E^{\mathrm{EM}}_{\mathrm{cell}}$ is the cell energy and $\sigma^{\mathrm{EM}}_{\mathrm{noise,cell}}$ is the expected background noise. Cells with $\varsigma^{\mathrm{EM}}_{\mathrm{cell}} \geqslant 4$ are grouped to form a seed cluster, and then less energetic neighbouring cells are iteratively added until all surrounding cells have $\varsigma^{\mathrm{EM}}_{\mathrm{cell}} < 2$. Final clusters are required to have $p_{\mathrm{T}} > 400$ MeV to reduce contributions from pile-up and $\pi_0 \to \gamma\gamma$ decays [142]. Next, ID tracks are refit, allowing for bremsstrahlung, and any photon conversion vertices which are present are reconstructed. The tracks and conversion vertices are then matched to the topo-clusters, and a "supercluster" algorithm is run to merge together nearby clusters. The algorithm is different for electrons and photons, and only the electron algorithm is presented here. Electron superclusters are seeded from any cluster with $E_{\mathrm{T}} > 1$ GeV with an associated track with at least 4 silicon hits in the ID. Next, satellite clusters within a window of $\Delta\eta \times \Delta\phi = 0.075 \times 0.125$ around the seed cluster are added in order to recover inefficiency due to bremsstrahlung emission. A second pass is then performed for electrons with an extended window of $\Delta\eta \times \Delta\phi = 0.125 \times 0.3$ (particularly extended in ϕ as the direction of the solenoid field causes bremsstrahlung photons to emit mainly in this direction). For the second pass, satellite clusters are required to share a "best-matched" track with the seed cluster, to discriminate from pile-up.

Once the superclusters are built, the electron and photon objects are reconstructed. An electron is defined as a supercluster with at least one matching ID track, while photons have two categories: converted and unconverted. A converted photon is one which has undergone e^+e^- pair-production in the ID, producing a conversion vertex, and any supercluster matched to a conversion vertex is defined as a converted photon. An unconverted photon is one which does not decay before reaching the ECal, and is therefore defined as any supercluster which matched no tracks or conversion vertices [141, 143]. Conversion occurs for around 20% of photons at low $|\eta|$, up to around 65% for $|\eta| \sim 2.3$ due to increased shower probability from the higher material budget in these regions. It is possible for an object to be reconstructed as both an electron and photon, in which case an ambiguity resolution is applied to remove the overlap, although some overlap is permitted to maintain high efficiency, which can then be refined further at the analysis stage.

Quality working points are applied to reconstructed electron objects in order to reduce the rate of fake electrons. Electron identification is performed using a likelihood discriminant which accounts for the EM shower shape in the ECal and transition radiation in the TRT, while electron isolation is determined using a combination of ID and ECal measurements to sum the energies or momenta of objects within a cone around electrons, which is required to be below a set threshold. The analysis described in this thesis uses at pre-selection the "Loose" identification and "Loose_VarRad" isolation working points, which are described in detail in [143]. In the final selection a tighter identification requirement "LooseBL" is used, which additionally requires at least one hit in the IBL, the innermost layer of the ID described in Section 3.2.1. The reconstruction efficiencies of each available identification and isolation working point are shown in Figures 3.16a and 3.16b as a function of transverse energy.

As mentioned in Section 2.4, several weights and corrections are applied to simulated electrons, to correct for any observed differences in their kinematic distributions between data and simulation. Corrections are applied to the energy scale and resolution, reconstruction efficiency and the efficiency of the identification and isolation working points described above. More detail on the considered corrections and their determination is given in [143].

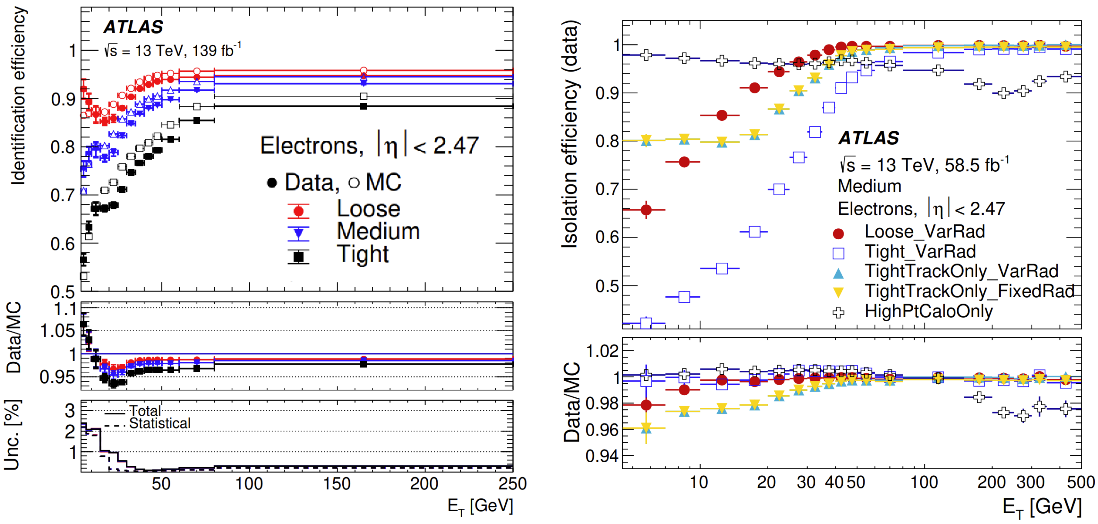

(a) (b)

Figure 3.16: Electron reconstruction efficiencies for each (a) identification and (b) isolation working point, determined from $Z \to e^+e^-$ decays as a function of transverse energy [144].

3.3.3 Muons

Muons are reconstructed using mainly information from the ID, the MS or a combination of the two. Depending on the detector signature, multiple types of reconstructed muon are defined as follows:

- Combined (CB) muons have tracks measured in both the ID and MS, which are combined in a global refit. These are the highest purity reconstructed muons.
- Muon Spectrometer Extrapolated (ME) or Stand-alone (SA) muons have a track measured only in the MS, with no matching ID track, which is loosely extrapolated back to the beamline.
- Segment-Tagged (ST) muons have tracks measured in the ID and in only a single layer of the MS, due to effects such as multiple scattering changing the trajectory of the muon in the detector.
- Calorimeter-Tagged (CT) muons have a track measured in the ID and an

energy deposit in the calorimeters, with no MS hits.

These four types of reconstruction are depicted in Figure 3.17.

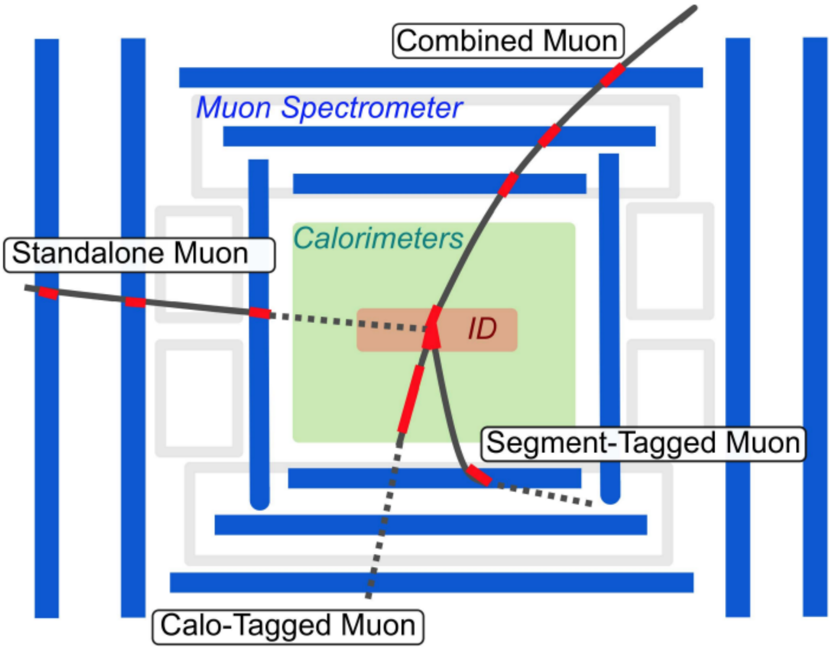

Figure 3.17: Muon reconstruction categories in the ATLAS detector [145].

Track reconstruction in the ID is described above in Section 3.3.1. For reconstruction of tracks in the MS, the procedure begins by combining compatible hits into straight-line segments in each individual layer of the MS using a Hough transform [146]. Segments are then loosely combined into preliminary track candidates using the IP position and information on the magnetic field generated by the toroids. The precision measurements from the MDTs and CSCs (described in Section 3.2.3) in the bending plane of the toroids are combined with complementary measurements from the RPCs and TGCs to reconstruct 3D track candidates. A global χ^2 fit of the muon trajectory is then performed, accounting again for the magnetic field, in addition to potential material interactions and misalignments in the detector. Outlier hits from this fit are removed from the track and any matching hits not yet assigned are added, before an ambiguity resolver similar to the one used for ID tracks is applied to remove any tracks overlapping with higher-quality tracks. The final set of MS tracks is then refit, accounting for positions of the IP and any large energy deposits in the calorimeters, to calculate the momentum of each muon.

As for electrons, quality working points are applied to reconstructed muons to control the rate of fake and non-prompt muons. All available identification and isolation requirements are described in detail in [147]. This thesis uses the "Medium" identification working point for muons, which uses only CB muons. In principle, ME muons are also permitted at $|\eta| > 2.5$, beyond the ID acceptance, although since this analysis uses a lower cut of $|\eta| < 2.4$ for muons these are not included. For the "Medium" working point, muon tracks must be seeded from either the MS and matched to an ID track (outside-in) or from an ID track and matched to at least three MS hits in at least two different MS stations (inside-out). For inside-out muons with $|\eta| < 0.1$, matched hits in only one station are allowed provided that there is at most one precision hole station (an MDT or CSC layer with no hit where one is expected based on the muon trajectory). To ensure an accurate momentum measurement, a loose agreement between the ID and MS measurements is required using q/p compatibility, defined as

$$q/p \text{ compatibility} = \frac{|(q/p)_{\text{ID}} - (q/p)_{\text{MS}}|}{\sqrt{\sigma^2_{(q/p)_{\text{ID}}} + \sigma^2_{(q/p)_{\text{MS}}}}}, \tag{3.9}$$

where $(q/p)_{\text{ID/MS}}$ are the ratios of muon charge q to momentum p in the ID and MS and σ are the corresponding uncertainties. The Medium identification working point requires q/p compatibility < 7. The isolation working point used in this thesis is "PFlowFixedRad_Loose", with the corresponding requirement of

$$\left(p_{\text{T}}^{\text{cone20}} + 0.4 \cdot E_{\text{T}}^{\text{neflow20}}\right) < 0.16 \cdot p_{\text{T}}^{\mu}, \tag{3.10}$$

where $p_{\text{T}}^{\text{cone20}}$ is the summed p_{T} of all tracks with $p_{\text{T}} > 500$ MeV in a cone of fixed size $\Delta R = 0.2$ around the muon and $E_{\text{T}}^{\text{neflow20}}$ is the transverse energy of all neutral particle-flow objects within the same cone [147]. The efficiencies of each identification working point are shown in Figure 3.18a, and the overall efficiency of applying both the identification and isolation working points used in this thesis are shown in Figure 3.18b.

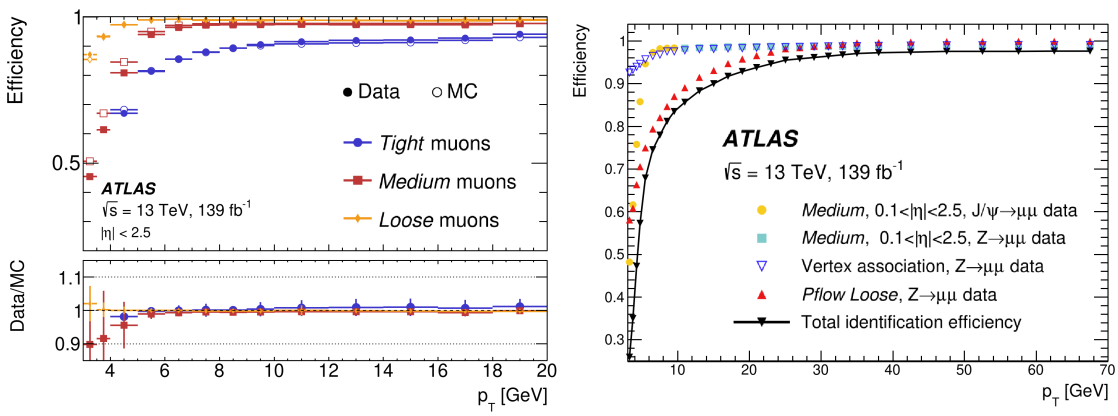

(a) (b)

Figure 3.18: Muon reconstruction efficiencies for (a) each identification working point and (b) the combined identification and isolation working points used in this thesis, determined from $J/\psi \to \mu^+\mu^-$ and $Z \to \mu^+\mu^-$ decays as a function of transverse momentum [147].

As for electrons, a set of weights and corrections are applied to simulated muons, to match their kinematic distributions to those observed in data. Corrections are made on momentum and energy scale and resolution, reconstruction efficiency and the efficiencies of the identification and isolation working points described above. More information on the determination of muon corrections is given in [145].

3.3.4 Other Objects

There are several other standard reconstructed objects in the ATLAS experiment which are not directly relevant to this thesis. Their reconstruction is briefly described below.

Jets: hadronic particles (subject to the strong force) are detected differently depending on their charge, with charged hadrons leaving signals in the ID, ECal and HCal while neutral hadrons (e.g. neutrons) only leave a signal in the HCal. Many particle decays following pp interactions produce quarks, which cannot exist in a free state due to QCD confinement, which forbids coloured free states. Therefore, quarks and gluons produced at the IP hadronise long before reaching the detector, producing an expanding cone of particles moving through the detector, referred to as jets. These objects are reconstructed, typically using an algorithm such as anti-k_{T} [148], and their components combined to obtain the properties of the original parton [149].

Taus: the heaviest type of lepton, taus have very short lifetimes due to their large mass, and so they decay before reaching the detector. Therefore, as for jets, only their decay products can be measured and the tau is reconstructed from the combination of these objects. These decays can be hadronic ($\sim 65\%$) producing hadronic jets similar to those described above, or leptonic ($\sim 35\%$) producing lighter leptons. However, since the tau lifetime is so short, these decays occur too close to the IP to distinguish their products from prompt electrons and muons, so only hadronically decaying taus have dedicated reconstruction in the ATLAS experiment [150]

Missing Transverse Energy (MET) $\mathbf{E}_{\mathbf{T}}^{\mathbf{miss}}$: the remaining SM object, neutrinos, cannot be detected by the ATLAS detector due to their small interaction cross-section. Therefore, the only way to measure these particles is by taking advantage of the negligible transverse momentum of the colliding protons, which means that the vectorial sum of p_{T} over the products of a given pp interaction should cancel out. Taking the inverse vectorial sum of the momentum of all measured objects (electrons, photons, muons, taus, jets and other tracks) gives the remaining transverse energy, which serves as an indirect measurement of this invisible event component [151].

Chapter 4

The ATLAS Forward Proton Spectrometer

The following chapter presents an overview of the ATLAS Forward Proton (AFP) spectrometer used to collect the data on forward protons used in the analysis presented in this thesis. The detector is introduced and its motivations discussed in Section 4.1 and a technical description of the detector is given in Section 4.2, followed by an explanation of the methods for reconstruction of forward protons from AFP spectrometer measurements in Section 4.3. The alignment of the detector is covered in Section 4.4 and finally the performance of the detector during Run 2 is presented in Section 4.5.

4.1 Introduction

The ATLAS Forward Proton (AFP) spectrometer is an ATLAS experiment subdetector designed to measure the momentum of protons which are scattered through tiny angles ($\mathcal{O}(\mu\text{rad})$) after undergoing an interaction in the central ATLAS detector but remaining intact and continuing down the beamline with a deflected trajectory. The detector consists of two arms, each positioned around 200 m down the beamline from the ATLAS detector, and makes use of Roman Pots (RPs), which are vacuum sealed devices attached to motors allowing the detector to be inserted to smaller radii than that of the beampipe in order to approach the LHC beam as closely as possible. This allows the detector to cover the very forward region around $|\eta|\sim 10$ where protons are scattered following diffractive or photon-induced interactions in the ATLAS detector, in which they lose a small fraction of their energy. The main objective of the AFP spectrometer is to measure this energy loss, parametrised as ξ which is defined as

$$\xi = 1 - \frac{E'_p}{E_p}, \tag{4.1}$$

where E'_p is the reduced energy of a scattered proton and E_p is the beam proton energy, which in 2017 was 6.5 TeV. Measuring ξ allows the proton kinematics to be partially reconstructed, and adds an additional component of information about events which can be combined with the central ATLAS detector reconstruction to perform unique physics analyses. An equivalent detector called the Precision Proton Spectrometer (PPS) is operated alongside CMS, originally in collaboration with the TOTal and Elastic Measurement (TOTEM) collaboration, which has since been fully absorbed into CMS [152].

The AFP spectrometer is designed to measure protons which have interacted within the ATLAS detector but remained intact, instead of breaking apart, which is possible if the quantum numbers of the protons are not changed during the interaction, as discussed in Section 2.3.3. This can occur in a range of soft and hard-diffractive processes, as well as in photon-induced interactions, due to the absence of colour exchange. Soft processes are of particular interest since they are relatively poorly understood from a theoretical standpoint, unlike hard interactions which can be studied via the application of QCD factorisation [70]. These soft processes must instead be studied via direct measurements, which can be achieved using the AFP spectrometer. Good understanding of these processes is necessary for accurate simulation of the corresponding interactions, which form backgrounds to many analyses and therefore must be well modelled. In addition, soft processes are responsible for the "underlying-event", discussed in Section 2.4, produced by interactions between spectator partons in the colliding protons which are not involved in the primary hard scattering interaction. The AFP spectrometer can also provide data to differentiate between two possible models for hard diffractive interactions with intact protons: soft colour interaction and resolved Pomeron exchange [127]; as well as studying other hard scattering processes such as jet, γ+jet and Drell-Yan electroweak boson production.

As discussed in Section 2.3, it is possible for scattered protons from elastic interactions to undergo further interaction such as gluon exchange, causing them to break apart. This gives rise to three categories of interaction in the AFP spectrometer: Elastic-elastic (EE), where both protons remain intact, Single Dissociative (SD) and Double Dissociative (DD), where one or both protons dissociate.

In AFP analyses which study CEP processes, discussed in Section 2.3.3, a redundancy which exists between measurements of the central system and the proton energy loss ξ can sometimes be exploited to remove background, e.g. for exclusive dilepton production $pp \rightarrow p\ell^+\ell^-p$, ξ can be additionally calculated from the lepton properties as

$$\xi^{\pm}_{\ell\ell} = \frac{m_{\ell\ell}}{\sqrt{s}} e^{\pm y_{\ell\ell}}. \tag{4.2}$$

CEP via photon-photon exchanges have so far been the subject of two AFP publications, one studying dilepton production [85] and the other ALP production [40]. A third analysis studying dilepton production in association with an invisible event component X is the subject of the analysis presented in this thesis.

4.2 Detector

The AFP spectrometer is comprised of two arms, each placed around 200 m from the IP at the ATLAS detector. These are referred to as Side A (anticlockwise around the LHC from the ATLAS detector) and Side C (clockwise). Each arm is made up of two stations, called Near and Far based on their relative distances from the

ATLAS detector. A general schematic of the detector layout is shown in Figure 4.1. The exact location of each station is optimised for maximal ξ acceptance within

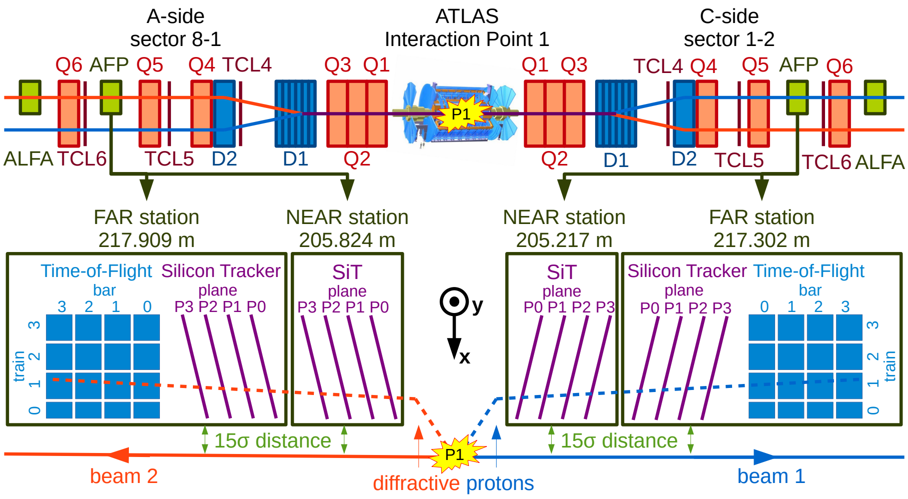

Figure 4.1: General scheme of the AFP spectrometer detectors [153].

the limits dictated by the available space on the LHC beamline between existing components. The station locations for 2017 data-taking with standard optics are summarised in Table 4.1, along with the ξ acceptance of each station, which is dependent on beam parameters, beam apparatus between the IP and the detectors, station locations and the global alignment of the stations, discussed in detail in Section 4.4.2. During operation, the FAR stations are inserted closer to the beam

Station ID	Side	LHC Sector	Proximity	z Position [m]	ξ Acceptance
0	A	8-1	FAR	+217.9	[0.018, 0.12]
1	A	8-1	NEAR	+205.8	[0.028, 0.115]
2	C	1-2	NEAR	-205.2	[0.026, 0.115]
3	C	1-2	FAR	-217.3	[0.019, 0.12]

Table 4.1: Naming conventions, locations relative to the ATLAS IP and acceptances of each AFP station during 2017. The acceptance range corresponds to values of ξ with at least 80% proton reconstruction efficiency.

than the NEAR stations, leading to their increased acceptance.

Each station contains a Silicon Tracker (SiT) module with 4 silicon planes capable of measuring proton x position to within less than 20 μm, described in more detail in Section 4.2.1. Using two stations on each side of the AFP spectrometer allows for more accurate determination of proton ξ and also enables measurement of the proton scattering angle, from which proton p_{T} and the 4-momentum transfer t can be determined, as discussed further in Section 4.3. In addition, each FAR station has a Time-of-Flight (ToF) detector to measure the time taken for protons to travel from the IP to the detector. This improves background rejection by enabling the reconstruction of the proton interaction vertex z position at the IP, as discussed further in Section 4.2.2.

These detectors are mounted in Roman Pot (RP) systems based on the stainless steel cylindrical pot design used by PPS [152], which can be moved horizontally in and out of the beampipe, perpendicular to the beam. This allows precise and flexible control of the distance between the detectors and the beam in order to approach as closely as possible. A schematic of a FAR station RP is shown in Figure 4.2. The

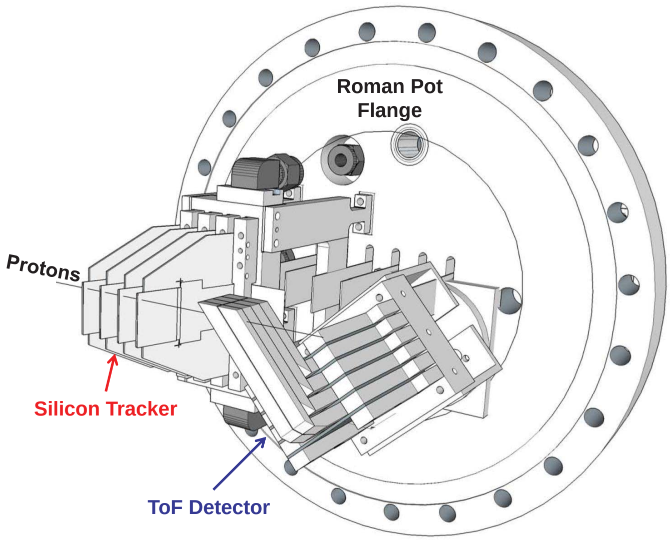

Figure 4.2: Diagram of a FAR station Roman Pot (RP) module with SiT and ToF detectors mounted [153].

pots are connected to the same vacuum chamber used by the LHC via bellows, but with a secondary vacuum for the interior of the RP in order to minimise the risk on the integrity of the LHC vacuum and reduce the deformation of the thin bottom window of the pots. The RPs are limited in how close to the LHC beam they can approach to avoid disturbing the beam. During operation the RPs are inserted such that the SiT edges are within around 12-15σ from the beam centre (~2 mm). Outside of stable beams the stations are kept in a "parked" position around 40 mm from the beam to avoid unnecessary radiation exposure.

As mentioned above, the ξ of each station is partly dependent on beam parameters. These are the crossing angle θ_c (-150 μrad in 2017 [154]) and β^*, which is related to the extent of beam squeezing at the IP (typically 0.3 m in 2017). Another factor is the beam apparatus between the IP and the station, with the magnet layout (beam optics) determining the degree of deflection of scattered protons when they reach the detector, discussed further in [155], and collimators limiting the maximum distance of protons from the beam. There are three collimators between the ATLAS detector and each AFP spectrometer arm: TCL4, 5 and 6 (visible in Figure 4.1) which are set to distances of $\pm 30\sigma$, $\pm 50\sigma$ and $\pm 40\sigma$ from the beam, respectively, and will absorb any scattered protons deflected beyond these distances. These values are optimised to give the AFP spectrometer good acceptance while continuing to protect machine components from radiation damage. Based on the values given in Table 4.1 the overall acceptance for the AFP spectrometer in 2017 was taken to be $0.02 < \xi < 0.12$, but it is recommended to reduce this to $0.035 < \xi < 0.08$ in analyses [156], as the station efficiencies are well understood for double station reconstruction in this region, as discussed further in Section 4.5.

4.2.1 Silicon Trackers

Each AFP station contains a Silicon Tracker (SiT) module with 4 silicon pixel sensors to measure the position of scattered protons, as shown in Figure 4.3. The sensors require both good spatial resolution for precise measurements and very high radiation hardness to withstand extensive, non-uniform irradiation due to their proximity to the beam during operation. Figure 4.4 shows a simulation of the radiation received by a SiT plane during normal operation, with two distinct regions visible. On the

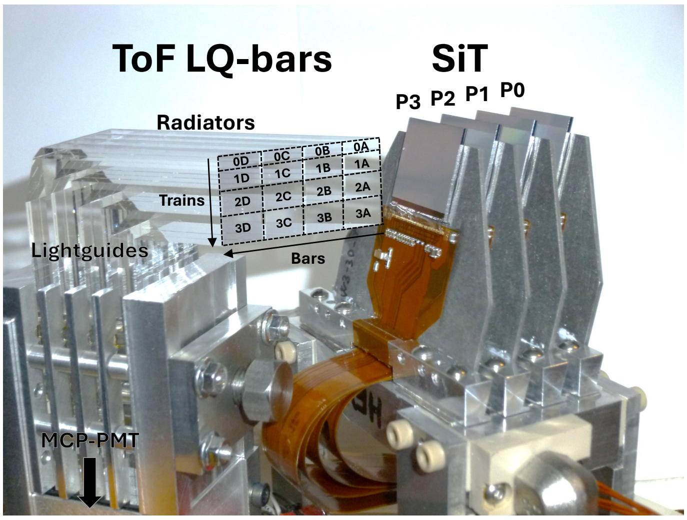

Figure 4.3: Photo of the SiT and ToF detectors for a single FAR station mounted on the RP. Adapted from [153].

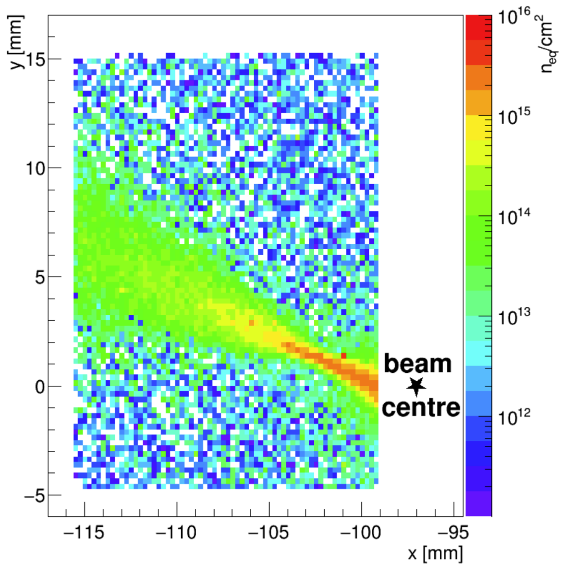

Figure 4.4: Simulation of radiation fluence through the SiT at 212 m [127].

side closest to the beam is a central "line" of diffractively scattered protons with typical energies close to that of the beam protons, with up to 3×10^{15} $\mathrm{n_{eq}/cm^2}$ of fluence expected over a period corresponding to the collection of an integrated luminosity of 100 $\mathrm{fb^{-1}}$. Outside of this region irradiation occurs primarily from ionisation by pair-produced e^+e^- with energies in the MeV-GeV range with an average fluence of 5×10^{12} $\mathrm{n_{eq}/cm^2}$. This variation in irradiation across the sensor over multiple orders of magnitude leads to irregular radiation damage and heating during operation. To cope with these conditions while maintaining sufficient resolution, the pixel sensors used by the ATLAS IBL detector [114] (Section 3.2.1) are used. These sensors differ from standard planar sensors such as those used by the SCT detector) in that they use 3D pixels with column-like n- and p-type electrodes which penetrate the substrate as shown in Figure 4.5, leading to reduced drift path within the sensor. This reduces charge trapping, thus reducing the overall bulk radiation

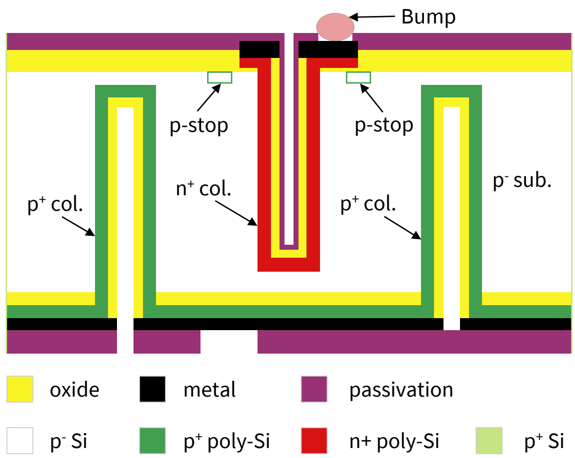

Figure 4.5: Diagram of the 3D silicon pixel design used for the ATLAS IBL and AFP SiT detectors [127].

damage on the sensor [157]. These semiconductor physics and radiation damage concepts are covered in detail in Chapter 5. The silicon planes are selected to minimise the "dead" edge of the sensor closest to the beam to below 100 μm to allow measurement sensitivity as close to the beam as possible to maximise the detector acceptance. Each plane has 336 $\times$ 80 pixels in the x–y directions, with individual pixels measuring 50 $\times$ 250 μm^2, for a total active area per sensor of 16.8 $\times$ 20 mm^2. Sensors are oriented in this way such that the more precise short pixel direction is along the dipole bending axis (x) as the position measurements in this axis are the most important for determining proton energy loss, since scattered protons are deflected in this direction. Vertical (y) proton deflection arises only due to a non-zero beam crossing angle at the IP and is therefore less important. The sensors are 230 μm thick and each of the four sensors in a station is separated from its neighbours by 9 mm. The electrodes are connected via bump-bonding to FE-I4 readout chips, which are wire-bonded to flexible printed circuits in order to read out the signals.

The per-plane resolution of the sensors is 14 (72) μm in the x (y) direction, although this is improved to 6 μm in the x direction by tilting the sensors 14° about the y axis in order to maximise charge sharing between adjacent pixels in the x axis, which allows hits to be measured by multiple pixels, giving a higher resolution.

4.2.2 Time-of-Flight

Each Far station also houses a ToF detector, shown alongside the SiT in Figure 4.3, designed to measure the time taken for protons to travel from their interaction vertex at the IP to the AFP station. This measurement allows the z position of the proton interaction vertex to be determined in double-tag events (where both protons are measured) as

$$z = \frac{c\Delta t}{2} = \frac{c(t_C - t_A)}{2}, \tag{4.3}$$

where t_A and t_C are the proton arrival times measured on Sides A and C of the AFP spectrometer respectively. This principle is demonstrated in Figure 4.6a. Re-

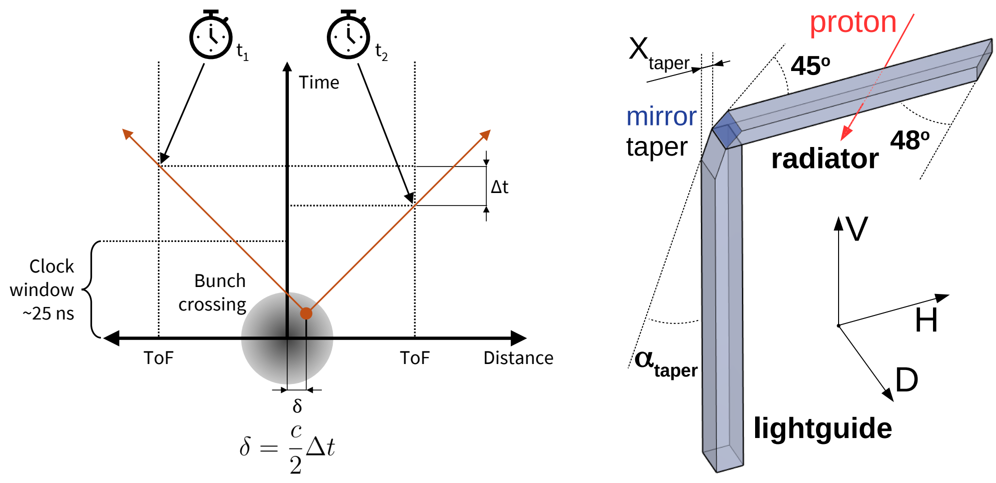

(a) (b)

Figure 4.6: (a) Demonstration of the proton vertex z position reconstruction using ToF measurements on both sides of the AFP spectrometer and (b) diagram of a single LQ-bar used in ToF [153].

construction of the proton vertex allows backgrounds to CD signal processes arising due to pile-up to be rejected with a high efficiency, since protons not originating in a signal interaction will not match with the reconstructed primary vertex of the centrally produced particles. This is particularly important in high-luminosity running, where this combinatorial background due to pile-up becomes significant. For operation during high-luminosity runs during Run 3 and especially following the HL-LHC upgrade, the detector requires extremely precise timing resolution of $\mathcal{O}(10\text{ ps})$ or better, with a spatial acceptance at least covering that of the SiT detectors. Additionally, the detector requires high detection efficiency, segmentation in the x direction for multi-proton timing, high rate capability $\mathcal{O}(5\text{ MHz})$ per segment and high radiation hardness. This is achieved using a Cherenkov detector based on L-shaped light-guiding quartz bars (LQ bars) positioned behind the SiT planes on the outer side of each Far station, which emit Cherenkov photons at a characteristic angle when a scattered proton with an energy close to that of the beam passes through them [158]. Each LQ-bar is formed of two quartz arms glued together at 90°, with mirror tape attached at the boundary to guide photons, as shown in Figure 4.6b. The L-shape is chosen for optimal light collection efficiency within the space constraints determined by the RP size. Four L-shaped Quartz (LQ)-bars are placed in sequence parallel to the beam to form a "train", with each station containing four trains which are oriented to the Cherenkov angle with respect to the proton flight direction, as shown in Figure 4.3. The bar geometry is deigned such that the optical paths of all bars in a given train are equal, to ensure that the signals from a single proton arrive simultaneously at the end of the light-guide. The light-guides are connected to a multi-anode Microchannnel Plate Photomultiplier Tube (MCP-PMT) with 16 channels, which converts the photons to an electrical signal which can be read out. In addition, a reference timing system is used to allow correlation of measurements between the stations on either side of the AFP spectrometer.

4.2.3 Trigger

Several of the physics processes which the AFP spectrometer aims to measure do not leave any significant signature in the central ATLAS detector, and therefore are either not triggered on or heavily prescaled. Therefore, a dedicated AFP L1 trigger is required to record these processes. This can be achieved using both the Far station SiT and ToF detectors, which can be set up to pass their signals to an RF switch, which selects between the two signals and passes one on to CTP. However, in practice, typically only SiT is used, due to performance problems encountered with ToF, as discussed in Section 4.5. The AFP L1 trigger must have extremely low latency due to the delay in signal propagation between the ATLAS detector and AFP spectrometer of $\mathcal{O}(1)\,\mu$s, which takes up a significant portion of the total ATLAS L1 trigger latency of around 2.5 μs.

4.3 Forward Proton Reconstruction

Forward proton tracks are reconstructed using hits measured by the AFP SiT detectors using a similar process to that used by the ID to reconstruct tracks in the ATLAS experiment, described in Section 3.3.1. A hit is registered if a pixel measures a signal above a threshold of around 2000 electrons. Pixel hits within the same plane are then recursively combined with their immediate neighbours in the x direction, since the pixel dimensions and plane rotation makes charge sharing in this direction very likely, as mentioned in Section 4.2.1, and each group of hits forms a cluster. The coordinates of each cluster are taken as the charge-weighted average of the pixel centres corresponding to all hits forming the cluster. At this stage, an inter-plane alignment correction is applied to the positions of each cluster in each plane of a given SiT station, to account for any misalignment between the four planes in that station, which is discussed in detail in Section 4.4.1. Next, tracks are reconstructed from sets of at least two clusters with separations in the x–y plane below 0.5 mm. A linear regression is performed on the cluster positions to determine the track parameters and χ^2 is calculated accounting for the pixel resolution to determine the goodness-of-fit. Finally, a global alignment correction is applied to the reconstructed track coordinates to account for any misalignment between the four AFP stations and the ATLAS detector, as discussed further in Section 4.4.2. There are three available working points for quality requirements on reconstructed proton tracks:

- Loose - 2 clusters per track
- Medium - Loose + hits on 2 different planes per track
- Tight - Medium + no more than 1 cluster hit per plane

The analysis presented in this thesis uses the default Medium working point.

Once tracks are reconstructed in each station, proton objects can be reconstructed. By default this is done using double-station reconstruction, where a track in both the NEAR and FAR stations on a given side are required with a maximum separation of $r < 2$ mm, which is given as

$$r = \sqrt{(x_{\text{FAR}} - x_{\text{NEAR}})^2 + (y_{\text{FAR}} - y_{\text{NEAR}})^2}, \tag{4.4}$$

where $x_{\text{NEAR/FAR}}$ and $y_{\text{NEAR/FAR}}$ are the track x and y coordinates measured in each station. If a pair of tracks fails this selection, single-station reconstruction is also possible using only the FAR station track. The proton properties can be reconstructed by converting the spatial track coordinates (x, y, z) into proton 4-momentum (E, p_x, p_y, pz), which is simplified by assuming ultra-relativistic protons with $E/m \gg 1$ and small-angle scattering such that $p_{\text{T}}/p_z \ll 1$. This leads to proton $E{\sim}p_z$, reducing the number of components to 3: (E, p_x, p_y) or equivalently (E, p_{T}, ϕ). Figure 4.7a shows an example of the dispersion of proton x and y positions in an AFP SiT detector for different values of proton energy loss ξ and p_{T}, with larger ξ values giving increased deflection from the beam and so larger x and y displacements, while for fixed ξ non-zero proton p_{T} causes a larger spread of hits in the x–y plane. Figure 4.7b shows that

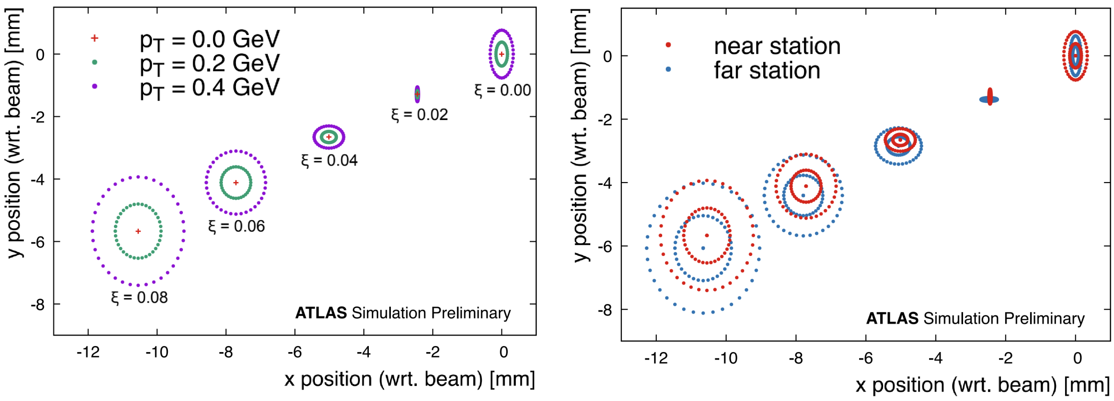

(a) (b)

Figure 4.7: Simulated dispersion of proton x and y positions in the AFP SiT planes for particular values of proton ξ and p_{T} (a) for a particular station and (b) compared between the NEAR and FAR stations on a given side, for equally spaced values of azimuthal scattering angle [156].

there can be a significant difference in measurements between the two stations on a given AFP side for protons with the same parameters. Since the deflection in y is strongly sensitive to the LHC beam crossing angle, which frequently changes, only x is considered for proton reconstruction. The conversion from spatial coordinates to proton properties is done using an inverse conversion, found by simulating the transport of various protons with different properties through the LHC optics between the ATLAS detector and the AFP stations. The simulated x to ξ correspondence can then be fitted to yield a parametrised transport function relating the track position in an AFP station to the energy loss ξ of the proton $x = T(\xi)$ [159, 160]. Different complexities of parametrisation can be used for this function depending on the precision required. For example, the AFP analysis measuring dilepton production via photon exchange uses $x(\xi) = -119\xi - 164\xi^2$ [85], while the analysis presented in this thesis uses a more sophisticated fit of $x(\xi) = -119\xi - 139\xi^2 - 195\xi^3$ for Side A and $x(\xi) = -120\xi - 138\xi^2 - 204\xi^3$ for Side C. For single-station reconstruction, only ξ can be reconstructed in this way, with p_{T} forcibly assumed to be zero. However, if a proton has tracks in both the NEAR and FAR stations then the two sets of spatial measurements, separated by 12 m in z, can be combined to determine the slope in x of the proton trajectory with respect to the beamline, referred to as Δx. A unique mapping then exists between this and the x measurement in either station $(x, \Delta x)$ and the proton properties (ξ, p_x), as illustrated in Figure 4.8. This

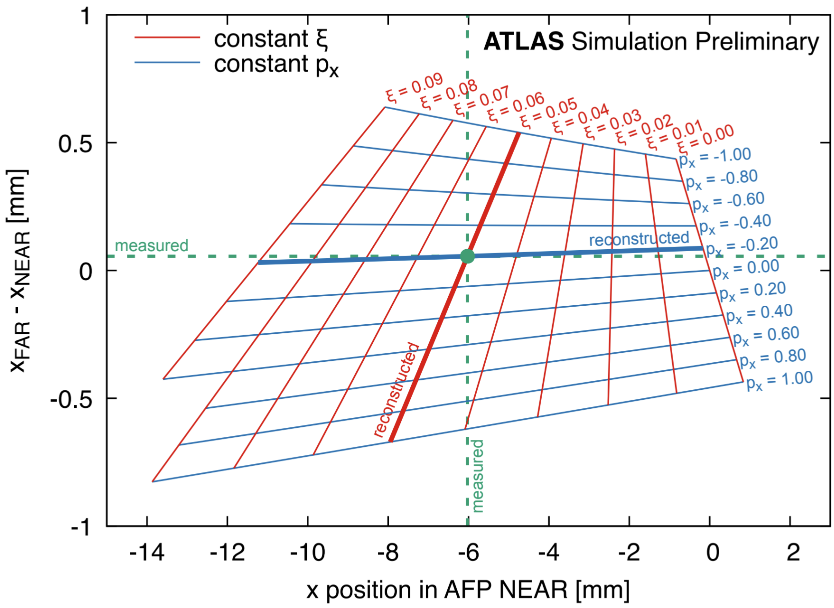

Figure 4.8: Simulated difference Δx between proton x positions measured in the NEAR and FAR stations on AFP Side A and the x coordinate measured in the NEAR station, as a function of proton ξ and p_x, demonstrating the unique mapping existing between these properties [156].

can be parametrised in a similar way to that shown above, in order to convert the pair of x coordinate measurements into ξ and p_{T} measurements for a given proton, where p_y is assumed to be zero due to the lack of measurement precision in this axis [156]. This in turn can be used to determine other proton properties such as the 4-momentum transfer t as

$$t = -\frac{p_{\mathrm{T}}^2}{1-\xi}. \tag{4.5}$$

The overall ξ and p_{T} acceptance of the AFP spectrometer is shown in Figure 4.9. Several uncertainties affect the reconstructed proton ξ values, such as alignment and the variation of the beam crossing angle by up to 50 μrad during a run. The total uncertainty in ξ varies from around 16% at low ξ, dominated by global alignment, to 10% at high ξ, dominated by beam optics uncertainties.

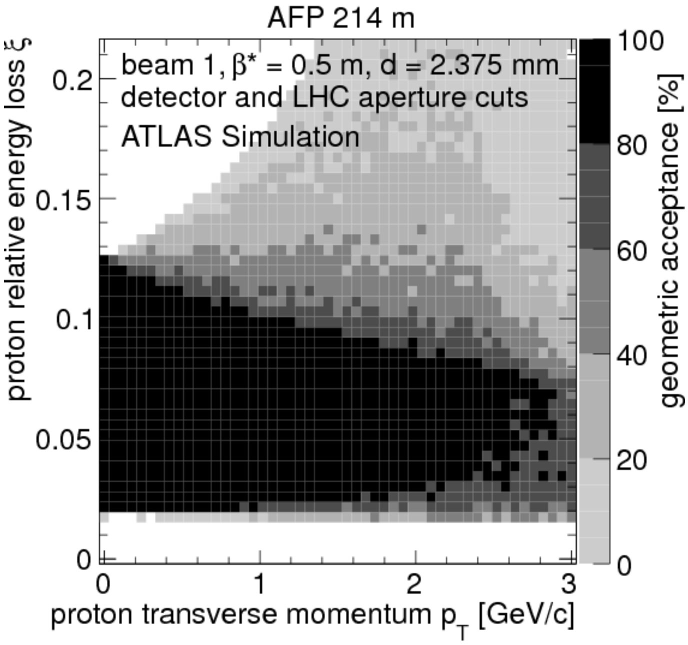

Figure 4.9: ξ and p_{T} acceptance of the AFP spectrometer in standard running conditions [160].

4.4 Alignment

Obtaining an accurate measurement of proton ξ depends heavily on the alignment of the AFP spectrometer, with two categories considered:

- Local Alignment - alignment of the 4 sensor planes within a given SiT station
- Global Alignment - alignment of the 4 AFP stations with the central ATLAS detector

The following sections describe the procedures used to achieve accurate alignment in each of these categories.

4.4.1 Local Alignment

Firstly, the position and orientation of each SiT plane within each station must be precisely measured, so that successive hits from a single proton in each plane can be correctly lined up to accurately reconstruct the track, also referred to as inter-plane alignment. Each plane has 6 degrees of freedom: positions (x, y, z) and rotations (γ, β, α) about the x, y and z axes respectively. The coordinate system is illustrated in Figure 4.10. However, the current alignment procedure only accounts explicitly

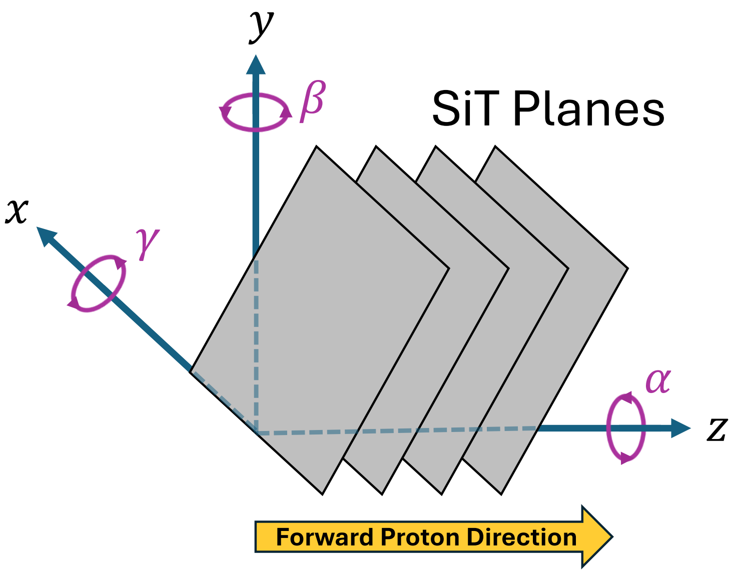

Figure 4.10: Local coordinate system used for the AFP SiT planes, where the origin is defined as the corner of the first SiT plane in a station. Adapted from [3].

for x and y translation and z rotation. The remaining degrees of freedom do not have significant effects, with translation in z only affecting the separation of each plane parallel to the beamline, which has a negligible effect on track reconstruction, and rotations abut the x and y axes having much smaller effects on reconstruction than the z rotation α, which is more likely to cause charge sharing between adjacent pixels and so has a larger impact.

Local alignment is performed using an iterative procedure, which begins with reconstruction of tracks using the method described in Section 4.3 while assuming perfect inter-plane alignment. To simplify the reconstruction it is also assumed that all tracks are parallel to the beamline (having zero slope) and the x and y track coordinates are calculated from the average corresponding coordinates of each cluster forming the track. This simplification has a negligible effect on reconstruction since forward protons originating from the IP have very small polar angles, since the differential cross-section decreases exponentially with increasing squared 4-momentum transfer $|t|$. This also helps to remove shower tracks, which have larger polar angles due to their production occurring closer to the AFP spectrometer, which helps the track reconstruction algorithm to converge. Once preliminary tracks are reconstructed, the residuals are calculated between the cluster centres (defined by the hit positions) and the corresponding best-fit track coordinates in each plane. A fit is then performed to determine the alignment parameters which minimise the mean value of the current set of residuals, the method for which is described in detail in [3], and the alignment is updated to reflect this.

This procedure is then repeated for 10 iterations, after which a cut of $\chi^2/\text{dof} < 2$ is made on each track, calculated between the hit and fitted track positions, to remove outlier tracks originating in background or shower processes and to remove anomalous clusters arising due to noisy pixels. The procedure is then repeated for another 10 iterations, after which all parameters are typically changing by negligible quantities in successive iterations, as demonstrated in Figure 4.11 which shows an example of the full procedure for a single parameter. The alignment parameters for

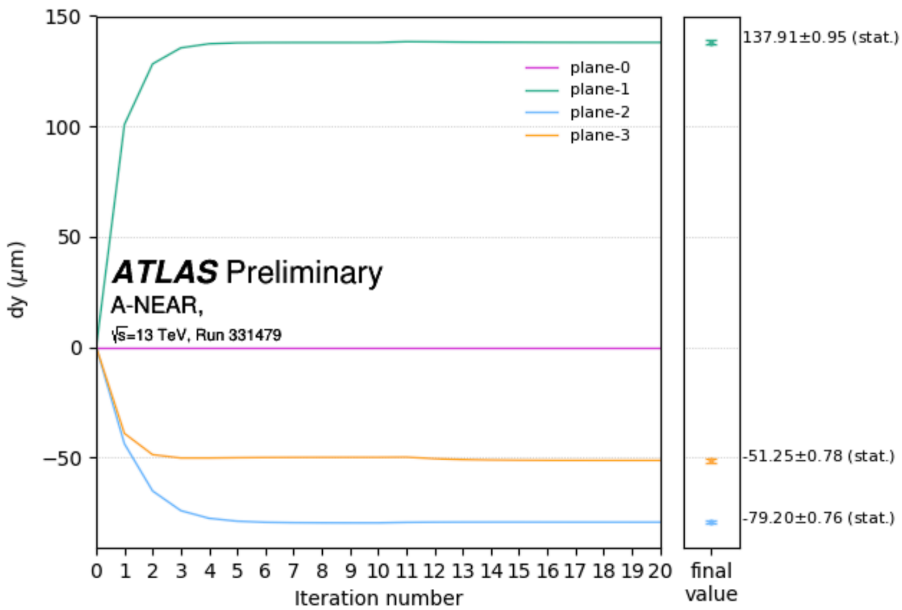

Figure 4.11: Evolution of the alignment parameter corresponding to translation in the y axis over 20 iterations of the inter-plane alignment procedure, with final values for each plane in μm. Each offset is calculated with respect to plane 0 to remove dependence on the global alignment [156].

a given station are determined with respect to plane-0 of that station in order to remove dependence on the global alignment between different stations at this stage.

The inter-plane alignment was repeated at several points during 2017 and the corrections were found to be stable to within $\mathcal{O}(1\,\mu\text{m})$ [3]. The systematic uncertainties on the alignment parameters are determined from the Root Mean Square (RMS) widths of the distributions of the corresponding residuals across a run once the alignment procedure has converged. These are typically of $\mathcal{O}(10\,\mu\text{m})$, which is small compared to the global alignment uncertainty discussed in the next section.

4.4.2 Global Alignment

Next, each AFP station must be aligned overall with the ATLAS detector. In this case, only translation in the x axis is considered, since x is strongly correlated with ξ and this therefore has the largest impact on kinematic reconstruction, with the remaining degrees of freedom left as sources of systematic uncertainty. For each station a single constant is determined defining the offset in x to be applied to all track coordinates in a given station, which depends both on station s and time t, defined as:

$$x\left(s,t\right) = x_{\text{pre-align}} + x_{\text{tracker}} + x_{\text{beam}}\left(s\right) + x_{\text{RP}}\left(s,t\right) + \delta x_{\text{corr}}\left(s\right). \tag{4.6}$$

This correction is formed of several components:

- $x_{\text{pre-align}}$: the raw track position following the inter-plane alignment corrections discussed in the previous section.
- x_{tracker}: the position of the SiT module, estimated from the edge of the active region of the sensor planes to the outer side of the RP floor. This is taken to have a fixed value of -0.5 mm for all stations.
- $x_{\text{beam}}\left(s\right)$: the nominal beam position, measured for each station before long data-taking periods during dedicated Beam-Based Alignment (BBA) runs. In

these runs, a low intensity beam is produced by the LHC and then each RP is gradually moved closer towards the beam until the nearest Beam Loss Monitor (BLM) detects a sharp rate change, at which point the beam is assumed to have been touched [161]. These values are typically of $\mathcal{O}(1\,\text{mm})$ and are continuously monitored throughout the year by Beam Position Monitors (BPMs) [162], remaining stable to within less than 100 μm.

- $x_{\text{RP}}\,(s,t)$: the RP position, corresponding to the distance from the beam to which the RPs are inserted. These are set to fixed values for long periods of time, accounting for both the beam width σ and a safety margin to protect the beam integrity. Table 4.2 shows the values used throughout 2017, which were changed at two points: once due to an agreed 0.5 mm decrease in safety margins following a period of successful running, and a second time due to a change in beam parameters from $\beta^* = 0.4$ m to $\beta^* = 0.3$ m. Any additional run-dependence due to small changes in beam conditions are currently considered to be negligible [156].

Insertion setting	$12\sigma + 0.8$ mm	$12\sigma + 0.3$ mm	$11.5\sigma + 0.3$ mm
Station	Roman Pot position x_{RP} [mm]		
0 (AFar)	-3.16	-2.65	-2.38
1 (ANear)	-4.07	-3.57	-3.60
2 (CNear)	-4.26	-3.76	-3.87
3 (CFar)	-2.96	-2.43	-2.23

Table 4.2: The RP position parameters $x_{\text{RP}}\,(s,t)$ used at different points throughout 2017 data-taking. The RMS beam width σ is around 200 (100) μm at the position of the Near (Far) station [156].

- $\delta x_{\text{corr}}\,(s)$: in-situ corrections to account for any remaining misalignment still present. This is obtained by analysing a highly pure sample of exclusive dimuon events and comparing the x value $x_{\mu\mu}$, which corresponds to the value of $\xi_{\mu\mu}$ calculated from the central muons using Equation 4.2, to the x coordinate of the corresponding proton track measured in the AFP spectrometer. This is demonstrated for a single station in Figure 4.12. The value of the in-situ correction is found to be between -0.22 mm and -0.43 mm for all stations.

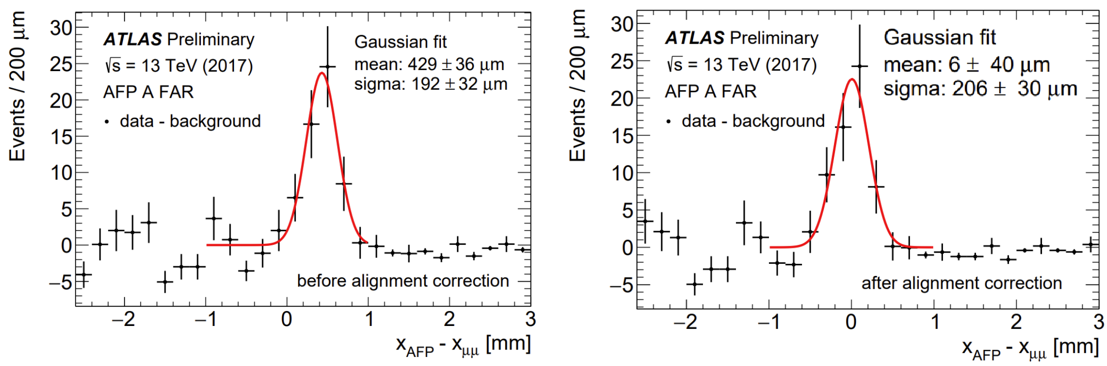

(a) Before in-situ correction (b) After in-situ correction

Figure 4.12: Example of the method for determining the in-situ correction for global alignment, showing the $x_{\mathrm{AFP}} - x_{\mu\mu}$ distribution in the AFar station for exclusive dimuon events (a) before and (b) after applying the correction. The raw signal distribution (left) is fitted to a Gaussian and the fitted mean is taken as the correction [156].

A final "fine-tuning" step is applied following all the above corrections, to explicitly require that the proton p_x distribution is centred on zero in both Far stations. A conservative systematic uncertainty of 300 μm is taken from the RMS widths of the Gaussian fits in in-situ analysis after applying the correction (e.g. Figure 4.12b). This conservative estimate also covers any neglected time dependence in parameters such as the RP position $x_{\mathrm{RP}}(s,t)$. The global alignment procedure was repeated to test the effect of several variations to the in-situ analysis, such as sub-divisions of the sample or changes to cuts, with all variations having effects on the alignment below 100 μm. The global alignment of all stations is found to be stable within 200 μm, which is at the level of the statistical uncertainties in the dataset used for the alignment procedure, when broken down into the sub-divisions.

4.5 Performance

The AFP spectrometer was inserted into the beam in 213 runs between June and November 2017, recording 32.0 fb^{-1} of data out of the 46.9 fb^{-1} recorded by the ATLAS detector, as shown in Figure 4.13. Requiring only data which satisfy the central ATLAS Good Run List (GRL) reduces the dataset to 26 fb^{-1}. The GRL

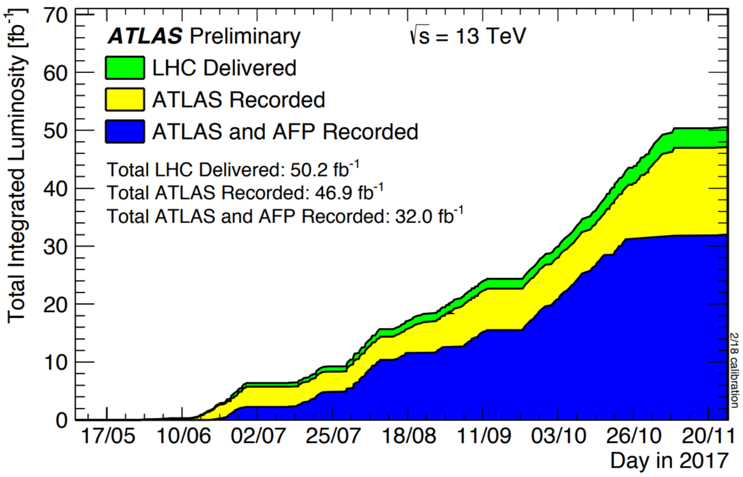

Figure 4.13: Luminosity recorded over time by the ATLAS detector and AFP spectrometer in 2017 [163].

selects the times within each run where the detector is taking high quality data, and is imposed at the level of luminosity blocks, which are ∼1 minute subdivisions of LHC runs with fixed trigger and data acquisition conditions and roughly constant instantaneous luminosity [164]. A dedicated AFP GRL is also applied, with the following requirements:

- All stations are inserted to their nominal data-taking positions
- At least two SiT planes in each station have high voltage on
- Data Acquisition (DAQ) system is fully functional [165]

This is defined as the "Loose" AFP GRL, and reduces the dataset to 19.2 fb^{-1}. Following further analysis, several runs were identified with significant drops in proton track reconstruction efficiency, caused by stations with fewer than 3 active SiT planes. The stricter "Nominal" AFP GRL removes these runs, reducing the available luminosity to the final value of 14.7 fb^{-1}, which is the size of the dataset used for the analysis presented in this thesis.

In order to evaluate the proton reconstruction efficiency, or the probability that a genuine proton scattered from an interaction at the ATLAS detector IP and within the AFP spectrometer acceptance is recorded, two different methods can be used.

The first is "tag-and-probe", where tracks are identified or "tagged" first in a single station, and then the corresponding track is checked or "probed" for in the other station, with the limit of $r < 2$ mm as defined in Section 4.3. The fraction of events where the probe identifies a valid track in the other station is taken as the efficiency. Efficiencies calculated from this method for individual runs throughout 2017 data-taking are shown in Figure 4.14, with on average around 99% efficiency for the NEAR stations and around 96% for the FAR stations. The reduced FAR station uncertainty arises from protons breaking apart after passing through the NEAR station, creating a shower of particles which do not leave a matching track in the FAR station. Another

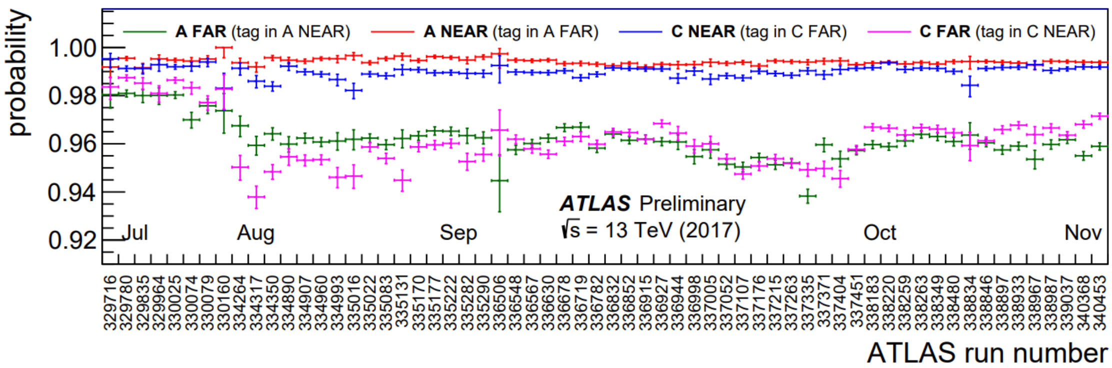

Figure 4.14: Proton reconstruction efficiencies for each AFP station determined from the "tag-and-probe" study, throughout 2017 data-taking as a function of ATLAS run number. The uncertainties shown are statistical [156].

method is to use the same principle but at the level of individual planes within a single station, where for a given plane the tag is formed of compatible hits in the other three planes and the probe is a search for a compatible hit in the plane being measured. This method yields per plane efficiencies consistently above 90% for all planes with high voltage on, giving overall station efficiencies of 99.9% for NEAR stations and 99.7% for FAR stations when a minimum of 2 plane coincidences are required. These efficiencies are higher than those calculated using tag-and-probe between stations, because this method neglects proton showering in between the two stations on each side which occurs due to interaction with the planes or pot windows, leading to no matching track being measured in the FAR station. Combining the station efficiencies gives an overall proton reconstruction efficiency of 0.92 ± 0.02, where the uncertainty arises due to the ξ dependence of the efficiency.

The ξ resolution of the AFP spectrometer has been evaluated in exclusive dilepton analysis, by comparing the measurements from the AFP spectrometer (ξ_{AFP}) with those made using the central ATLAS dilepton measurements with Equation 4.2 ($\xi_{\ell\ell}$) [85]. Figure 4.15 shows the measured difference, with Gaussian fits yielding a resolution of around 10% at the mean measured ξ value of ~0.024 [156].

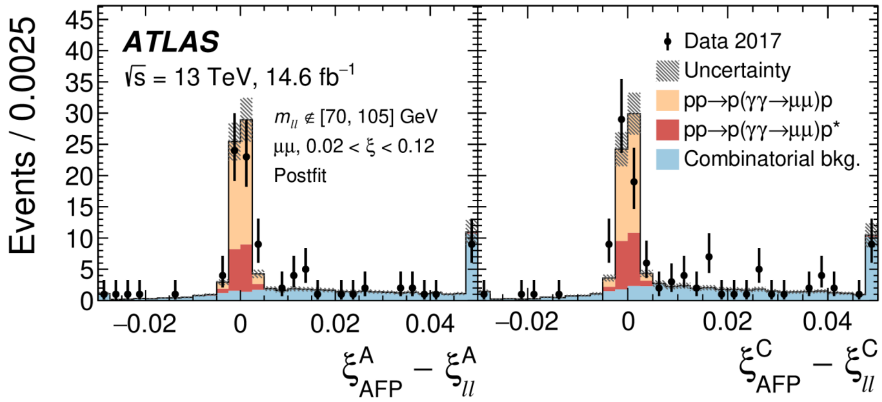

Figure 4.15: Distributions of the difference between ξ measurements from protons in the AFP spectrometer and from the dilepton system in the ATLAS detector in exclusive dilepton production events in the muon channel [156].

The timing resolution of the ToF detectors was measured to be 21 ± 3 ps for Side A and 28 ± 4 ps on Side C, corresponding to a combined reconstructed z vertex resolution of 5.3 ± 0.6 mm. This is in good agreement with the measured resolution obtained when comparing the vertex position reconstructed by ToF z_{ToF} to the primary vertex position measured in the ATLAS detector z_{ATLAS}, as shown in Figure 4.16, giving $\sigma = 6.0 \pm 2.0$ mm. This is at a sufficient level of precision to make a significant contribution to double-tag analyses, such as the one presented in this thesis. However, the measurement efficiency was measured to be extremely low in all ToF trains throughout 2017 data-taking. Figure 4.17 shows the efficiencies determined from AFP calibration stream data in a single run, with Side A not exceeding 20% efficiency and Side C even lower at below 10%. This efficiency falls well below the detector specifications [167], and degraded further throughout the year due to ageing down to sub-percent levels. This poor performance is attributed to exceeding

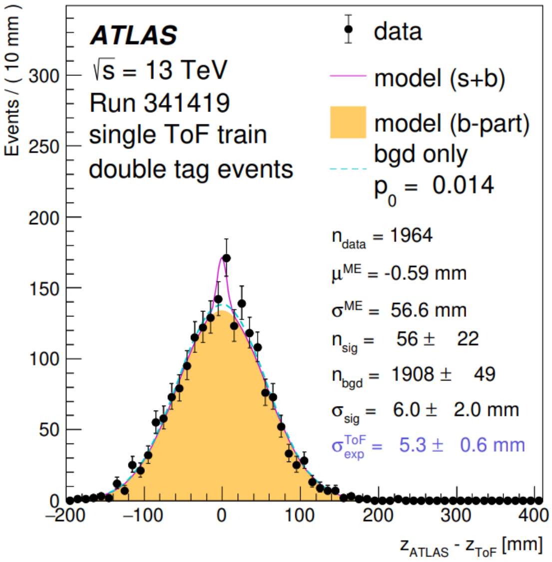

Figure 4.16: The z_{ATLAS}–z_{ToF} distribution in double-tagged events measured in the ATLAS detector in Run 341419. The excess of signal (purple) over the background (yellow) is due to double-tagged events with both protons originating in the same signal process [166].

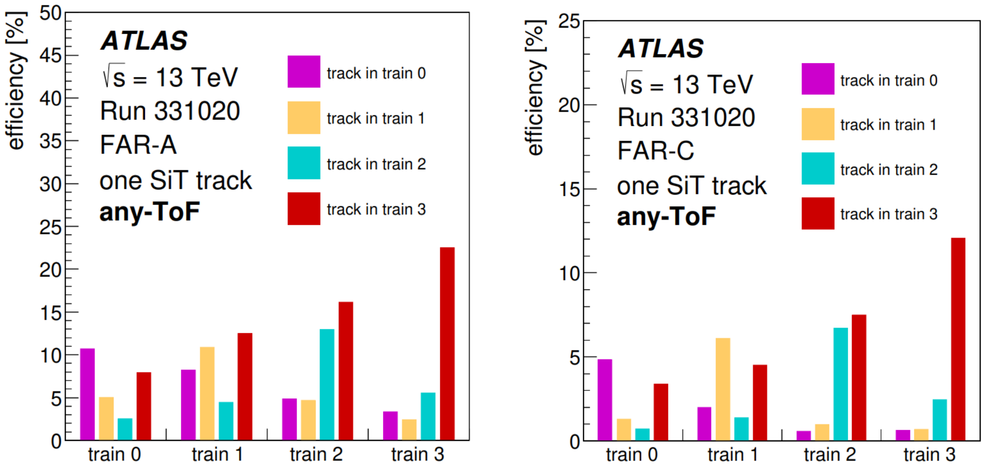

(a) (b)

Figure 4.17: ToF train efficiencies determined using AFP calibration stream data in ATLAS Run 331020 in the (a) AFar and (b) CFar stations, in events with exactly one reconstructed SiT track [166].

the lifetimes of the MCP-PMTs used in ToF [166], and renders ToF data unsuitable for use in high pile-up runs. Therefore, ToF information is not considered in the analysis presented in this thesis. A new design is now in use in Run 3 to prevent these issues [168] and a significant amount of usable ToF data have now been taken, although still at a lower efficiency than for the AFP SiTs.

The AFP spectrometer was also operational during 2018 data-taking. However, a timing issue occurred causing a decorrelation between the data from the AFP spectrometer and the central ATLAS detector, unfortunately preventing the use of these data in any physics analyses. As a result, the analysis presented in this thesis uses only 2017 data.

Chapter 5

Quality Assurance for the ATLAS Inner Tracker Upgrade

This chapter presents a summary of the work performed during the first year of the author's PhD on the Quality Assurance (QA) program for the upcoming Inner Tracker (ITk) upgrade to the ATLAS Inner Detector (ID). The theory of semiconductor physics behind the silicon sensors used by the current and future ID is covered in Section 5.1 and an overview of the ITk upgrade is given in Section 5.2. Finally, a summary of the measurements performed by the author towards the ITk QA program is presented in Section 5.3.

5.1 Semiconductor Physics

In its pure form, silicon has the properties of a semiconductor, which means that its conductivity falls between that of a conductor (such as a metal) and an insulator (such as ceramic). Another property of semiconductors is the "band-gap" within the structure of the energy bands which can be occupied by valence electrons. Electrons are forbidden from occupying this region, instead being confined to the low-energy valence band below the gap, and the high-energy conduction band above it. At low temperatures, all electrons are confined to the valence band, so there are no free charge carriers and the semiconductor has low conductivity. However, at higher temperatures electrons can become excited and rise to the conduction band, which leaves behind a positively-charged "hole" in the valence band. Both the energised electron and the corresponding hole are then free to diffuse around the lattice, contributing to the concentration n_i of free charge carriers (where i corresponds to electrons or holes) and thus the conductivity of the semiconductor. Through this process, the conductivity of semiconductors such as silicon becomes strongly temperature dependent. However, this can be improved via a process called "doping".

Silicon atoms are tetravalent, having four valence electrons in their outer shells, so they each form four covalent bonds. Therefore, solid silicon has an extremely stable lattice structure, similar to diamond, with all valence electrons confined to these bonds. Doping is the introduction of other atoms with different numbers of valence electrons into the silicon lattice, referred to as "impurities". Trivalent atoms, such as boron, only form three covalent bonds, and therefore leave a hole in the lattice where a bond is missing, creating "p-type" silicon with holes as the "majority charge carrier". Pentavalent atoms, such as phosphorous, form five covalent bonds and therefore contribute an extra free electron to the lattice after forming four covalent bonds, creating "n-type" silicon with electrons as the majority charge carrier. These impurities create extra levels within the band gap, decreasing the energy which is required for charge carriers to move to an excited state.

When p-type and n-type silicon are joined together, they form a p–n junction, in which the mobile electrons and holes can diffuse across the junction and recombine. This creates two regions of fixed charge either side of the junction, positive on the n-type side and negative on the p-type side, in turn creating a Space Charge Region (SCR) at the junction with an effective space charge density of $N_{\text{eff}} = N_D - N_A$, where N_D and N_A are the concentration of acceptors (holes on the p-type side) and donors (electrons on the n-type side), respectively. Due to the transition across the junction between the differing potentials V_n and V_p created on the n and p-type sides, an electric field is generated across the SCR, which acts in the opposite direction to the diffusion current of charges across the junction. Once enough recombination has occurred and the SCR is large enough, the current due to the resulting electric field perfectly balances the diffusion of charge carriers across the junction, creating an equilibrium state where the net rate of diffusion and recombination falls to zero and on either side of the junction there is a stable region of fixed charge carriers referred to as the "depletion region" with width W_d. This process is outlined in Figure 5.1. The depletion region has a "built-in" voltage V_0 which maintains the equilibrium

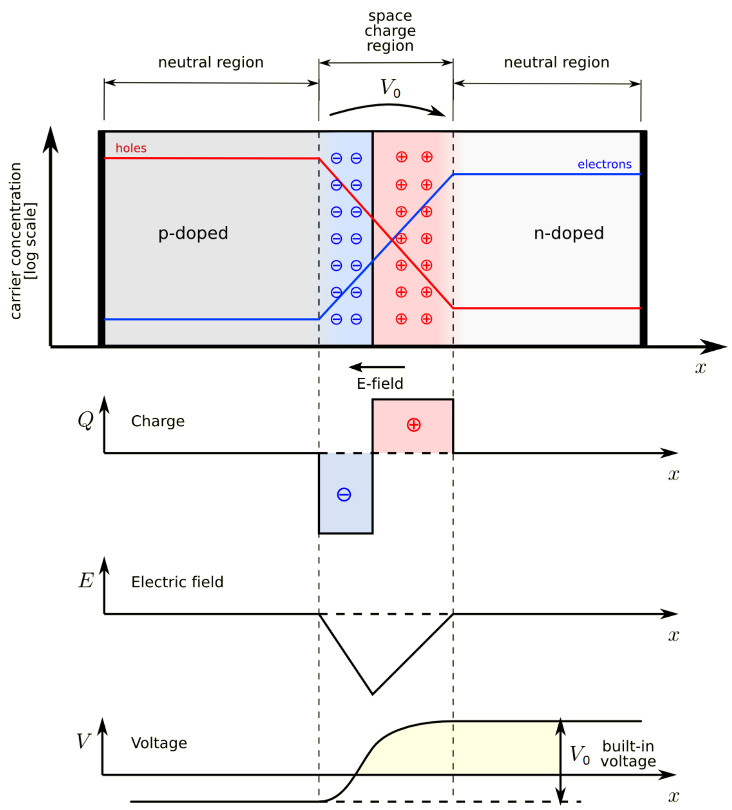

Figure 5.1: Diagram showing the charge carrier concentration across a p–n junction, and the resulting electric field and built-in voltage [169].

across the junction, corresponding to the difference between the constant potentials on the n and p-type sides. This is also related to the doping concentrations N_A and N_D as

$$V_0 = V_T \ln\left(\frac{N_A N_D}{n_i^2}\right), \tag{5.1}$$

with the thermal voltage $V_T = \frac{k_B T}{q}$, where k_B is the Boltzmann constant, T is the absolute temperature of the p–n junction and q is the charge of an electron [170].

The *p*–*n* junction can be treated as a diode. When an external bias voltage V_{bias} is applied across the junction in the forward direction (high voltage on the p-side) this opposes the built-in electric field across the junction, which reduces the width of the depletion region that can be held at equilibrium. This also increases the flow of majority charge carriers across the junction, creating a current I defined by the Shockley diode equation as

$$I = I_0 \left(e^{\frac{V_{\text{bias}}}{V_T}} - 1 \right), \tag{5.2}$$

where I_0 is the reverse leakage current of the *p*–*n* junction [171]. This forward current will increase exponentially with $V_{\text{bias}} \gg V_T$. Conversely, if a reverse bias in the opposite direction is applied, it instead strengthens the built-in electric field, allowing for further recombination and increasing the width of the depletion region. This produces a large potential barrier, preventing the flow of current across the junction. Equation 5.2 shows that for negative $V_{\text{bias}} \gg V_T$ the exponential term will tend to zero, leaving the reverse leakage current $I = -I_0$. This current arises from the random excitation of electrons by thermal energy, which results in a small concentration of free charge carriers able to diffuse across the junction. Therefore, the leakage current is strongly temperature dependent, characterised as

$$I_0(T) \propto T^2 e^{-\frac{E_{\text{eff}}}{2k_B T}}, \tag{5.3}$$

where $E_{\text{eff}} = 1.21$ eV [172]. At low temperatures the leakage current is typically very small, however very high reverse bias voltage can create a strong enough electric field across the junction to accelerate thermally exited electrons to a sufficiently high energy to cause further ionisation within the lattice. If the energy is high enough, electrons freed by this ionisation will cause further ionisation, creating an avalanching process known as "breakdown" which leads to extremely high current across the junction. These diode *I*–*V* properties are illustrated in Figure 5.2.

When an energetic charged particle passes through the depletion region of a reverse biased *p*–*n* junction it will cause ionisation, creating electron/hole pairs which

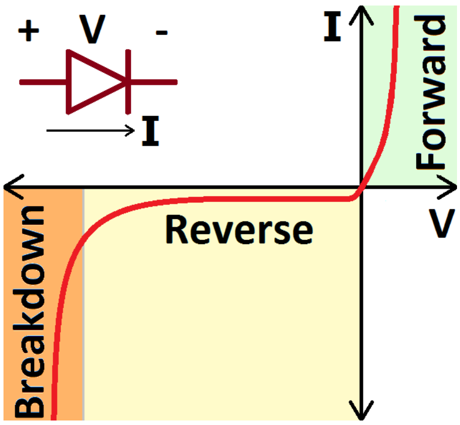

Figure 5.2: The I–V behaviour of a diode for a large range of forward and reverse bias voltages. Adapted from [173].

then drift along the electric field created by the applied bias voltage. If conductive electrodes (e.g. aluminium) are placed nearby, the movement of these charges will induce a current in the electrodes. This is the principle behind the silicon charged particle sensors used in the ATLAS detector, AFP spectrometer and other particle physics experiments. Example diagrams of particle sensors using this principle are shown in Figure 5.3. Particles will only cause the ionisation required for detection

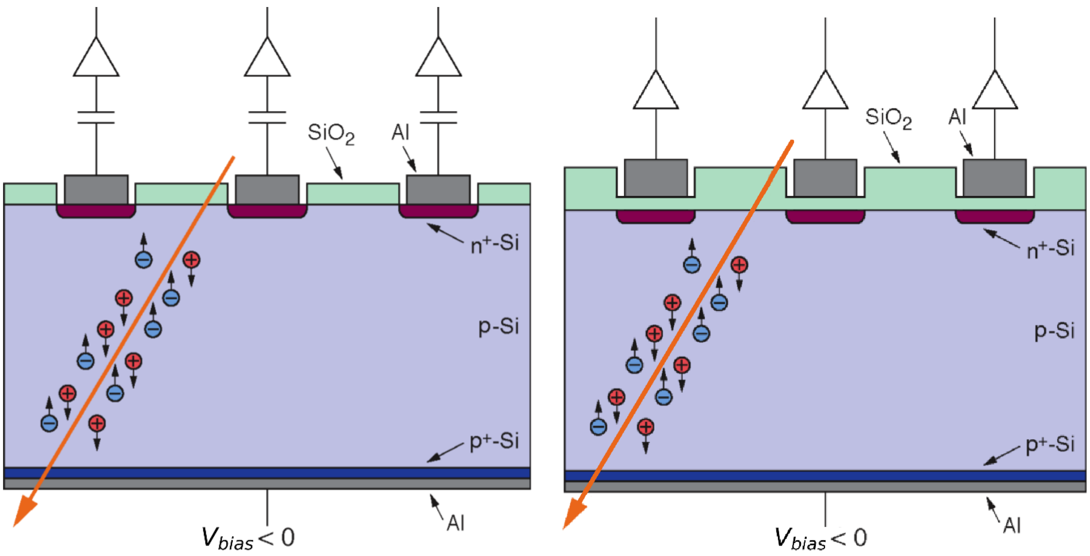

(a) DC Coupled (b) AC Coupled

Figure 5.3: A basic silicon strip sensor using a reverse biased p–n junction with (a) DC and (b) AC coupling. Based on [174].

if they pass through the depletion region of the silicon, and the resulting charge is proportional to the width of the depletion region. Therefore, the sensor efficiency is directly related to the depletion width W_d, which is itself directly proportional to the bias voltage according to

$$W_d = \sqrt{\frac{2\varepsilon_{\mathrm{Si}}\left(N_A + N_D\right)}{qN_AN_D}\left(V_0 + V_{\mathrm{bias}}\right)}, \tag{5.4}$$

where $\varepsilon_{\mathrm{Si}}$ is the absolute permittivity of silicon [175]. The depletion width is maximised at Full Depletion (FD), when the depletion region covers the entire volume of silicon, giving the maximum possible charge per hit, which occurs at the depletion voltage $V_{\mathrm{bias}} = V_{\mathrm{FD}}$. Particle detectors in experiments such as the ATLAS experiment need high granularity to measure the exact position of each particle hit, which requires segmentation of the sensors, as shown in Figure 5.3. This is achieved by making the majority of the volume of a sensor from a single type of silicon, p-type in the diagram, referred to as the “bulk”, and then forming the junction by adding a smaller quantity of the other type of silicon on top, n-type here, referred to as the “implant”. These small implants can then be precisely segmented into pixels (of roughly square dimension, giving precision 2D position measurements) or strips (much longer in one axis, giving precise measurements in only one-dimension). In order to compensate for the reduced thickness of the implant layer of silicon and maintain the same depletion width for the same bias voltage, the implant layer’s doping concentration must be increased relative to the bulk in order to satisfy the relation

$$N_Dx_n = N_Ax_p, \tag{5.5}$$

arising from the requirement for zero overall charge in the depletion region, where x_n and x_p are the widths of the depletion region on the n and p-type sides of the p–n junction respectively. This more highly doped silicon is referred to as e.g. n^+-type. Additional layers of highly doped silicon with the same type as the bulk (p^+-type in Figure 5.3) can be deposited between adjacent segments, in order to stop electrons flowing between them and avoid charge sharing, which can diminish signal strength. This is referred to as a p-stop or p-spray. The implants are then connected to the electrodes via a capacitor which allows for a current to be induced by incident particles.

Figure 5.3 shows two different methods for connecting the implants and electrodes. Figure 5.3a shows a Direct Current (DC) coupled sensor, where the electrodes are directly connected to the implants and a dedicated coupling capacitor is placed further up the circuit. Figure 5.3b shows an Alternating Current (AC) coupled sensor, which has an insulating layer of oxide (e.g. SiO_2) between the implants and electrodes, forming a built-in capacitor with the oxide acting as the dielectric. An additional layer of metal is attached on the opposite side of the bulk to allow for biasing, known as the backplane, which is typically connected via a highly doped layer of the same type of silicon as the bulk to ensure good ohmic contact.

5.1.1 Radiation Damage

When exposed to high levels of radiation such as those present in the LHC experiments and particularly in the future HL-LHC, silicon particle sensors become damaged in two main ways. The first is bulk damage, caused by the displacement of atoms from the lattice by energetic particles, and the second is surface damage, caused by ionisation in the oxide layers of a sensor, forming defects at the interface between the oxide and the implants. These forms of damage both decrease the effectiveness of sensors in some way, and experiment design therefore seeks to mitigate these effects however possible. This damage can be reversed to an extent by a process of heating the sensor referred to as annealing, which produces several temperature dependent effects which can remedy the effects of radiation damage.

Bulk Damage

Energetic particles which are incident on the sensor can collide with atoms in the silicon lattice, causing Non-Ionising Energy Loss (NIEL) and eventually leading to the displacement of atoms from the lattice. This leaves a range of defects in the lattice, as shown in Figure 5.4a, such as vacancies, where atoms are missing from lattice sites, and interstitials, where atoms are free to diffuse through the lattice.

These defects also create new energy levels in the silicon band gap, changing its

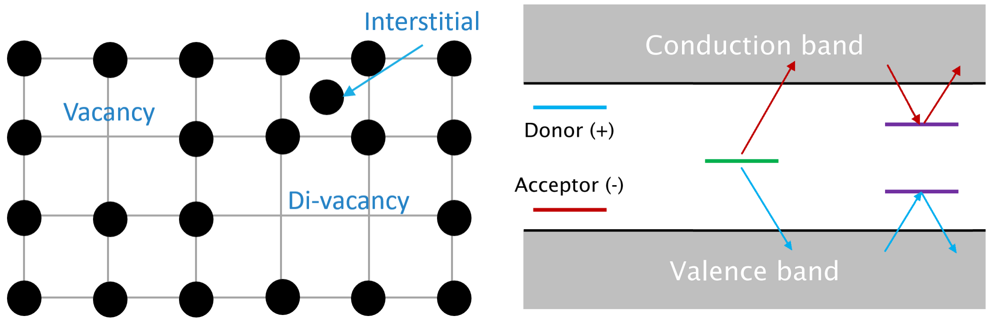

(a) (b)

Figure 5.4: (a) Defects which form in the silicon lattice due to bulk damage. [Based on Figure 2.1 from [176]] and (b) the effects due to the resulting new energy levels in the silicon band gap [176, 177].

properties via three main macroscopic effects, illustrated in Figure 5.4b.

- Creation of extra holes and the removal of free electrons: creates negative space charge, making silicon more p-type. This increases V_{FD}, decreasing sensor efficiency at a given voltage. This can also result in type inversion in high-resistivity n-type silicon if the threshold of $N_D > N_A$ is crossed, such that the silicon inverts from n-type to become p-type.

- Formation of defect energy levels near the middle of the band gap, referred to as "generation centres", which emit extra electron/hole pairs. This creates additional leakage current, increasing the noise and power consumption of sensors.

- Creation of energy levels near the valence and conduction bands which "trap" charge carriers over significant timescales, preventing them from diffusing and contributing to the signal for a particle "hit". This weakens the signal produced from a hit by a particle of a given energy and is the dominant damaging effect in highly irradiated sensors [176, 177].

The amount of NIEL caused by a given radiation source is dependent on both the type and energy of the irradiating particles. However, comparisons can be made between the damage caused by different radiation sources via the NIEL hypothesis, which states that the effects from radiation damage scale linearly with NIEL independently of their spatial and energetic distribution, depending only on the initial number of primary defects in the lattice. Therefore, the damage caused by different radiation sources can be scaled via their NIEL (number of lattice displacements) to a common factor. This is chosen as the damage caused by the equivalent fluence of 1 MeV neutrons, such that radiation quantities are given in terms of their "1 MeV neutron-equivalent fluence" in units of $n_{\mathrm{eq}}\mathrm{cm}^{-2}$ [172].

Surface Damage

Radiation can cause ionisation of atoms in the layer of oxide deposited on top of the bulk of sensors, creating additional electron/hole pairs which are quickly separated by the high electric field that exists in the oxide. Shallow energy levels in the oxide result in a large hole capture cross section, 10^6 times higher than for electrons, such that while the produced electrons quickly drift out of the oxide, the holes instead drift slowly towards the interface between the oxide and the bulk, where they become trapped by defects in the oxide. This creates fixed positive charge at the bulk-oxide interface, which can attract electrons from doped silicon causing a current to flow between adjacent segmented implants along the interface. This decreases the resistance between adjacent segments, increasing noise and cross-talk which reduce the performance of the detector. In addition, the resulting charge sharing decreases the sensitivity of sensors by splitting small signals between multiple segments, causing them to fall below the detection threshold. This effect is particularly relevant for AC coupled sensors where the bulk-oxide interface separates the implants from the electrode, so signals are more likely to be lost due to trapping at the interface [176, 177].

Annealing

Annealing is the process of heating a radiation-damaged silicon sensor in order to repair some of the damage through several mechanisms. It can mitigate the effects of bulk damage by causing an increase in donor-like defects which can recombine with free acceptor states, reducing the negative charge and full depletion voltage. It is also shown to continuously decrease leakage current and decreases the trapping rate for electrons. However, it has the opposite effect on holes, increasing their trapping rate. Furthermore, long-term annealing can activate additional acceptor-like states which counteract the beneficial effects of annealing and eventually result in an overall worsening in performance, referred to as reverse annealing. Annealing can also energise electrons to drift towards the bulk-oxide interface and recombine with trapped holes, reducing surface damage, although this is a very slow process requiring extensive annealing to see positive effects.

5.2 Inner Tracker Upgrade

As discussed in Section 3.1.1, the LHC is being upgraded for Run 4 onwards to the HL-LHC, with a planned increase in instantaneous luminosity by around a factor of 3 to allow for faster data collection. This corresponds to an increase in interactions per bunch-crossing μ from around 60-70 in Run 3 to ~200. The current ATLAS Inner Detector (ID) system is composed of the silicon-based Pixel and SCT detectors, using pixel and strip sensors respectively, and the gas-straw TRT detector discussed in Section 3.2.1. This system does not have the required precision or granularity to resolve tracks at the very high density resulting from this level of pile-up. In addition, the increased particle flux will result in higher radiation exposure for the detector, resulting in damage which the current system cannot withstand while maintaining sufficient performance across the planned 10 years of HL-LHC operation. Therefore, an upgrade to the entire ID is planned, called the Inner Tracker (ITk), which replaces all current systems with a new all-silicon design. The ITk detector, shown in Figure 5.5, is separated into two main components, the inner silicon pixel detector and outer silicon strip detector. The basic existing structure of the ID is maintained, with a layered barrel parallel to the beamline covering the central region and layers of disks forming perpendicular

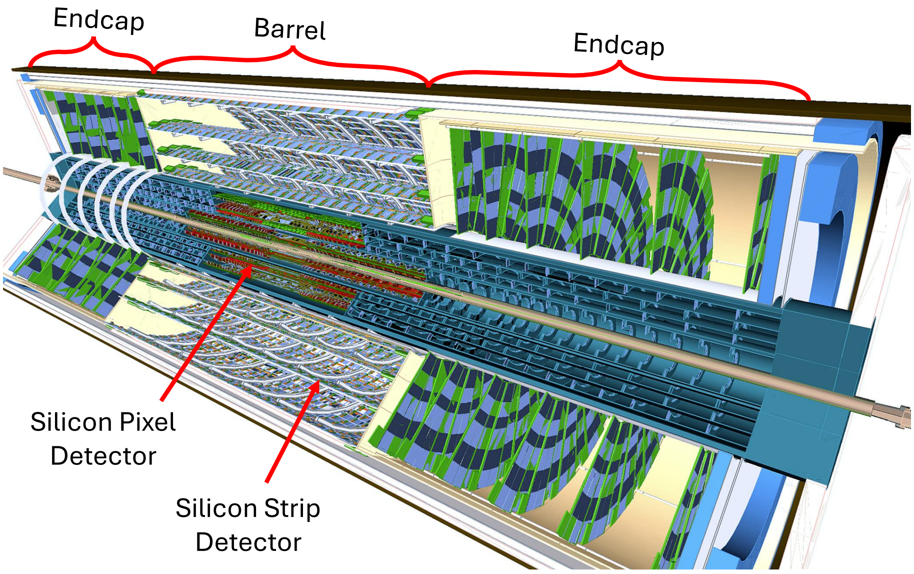

Figure 5.5: Simulation of the ITk layout, comprised of silicon pixel and strip layers, with a barrel and endcap structure [178].

endcaps in both forward regions. Closest to the beamline, the ITk pixel detector barrel is composed of five layers of silicon pixel sensors positioned between 39-271 mm from the IP, with pixel sizes of $50 \times 50\,\mu\text{m}^2$ or $25 \times 100\,\mu\text{m}^2$. This is surrounded by the ITk strip detector, made up of four layers of silicon strip sensors extending up to 1000 mm from the IP, in order to completely replace the existing TRT. The inner two layers use short strips (SS) 24.1 mm long and 75.5 μm wide, while the outer two layers use long-strips (LS) 48.2 mm in length, sacrificing some precision for a less complex and expensive system. The end-caps are built from petal-design disks, with four rings of pixel sensors covering up to $|\eta| < 4$ and six disks of strip sensors covering up to $|\eta| < 2.7$. All strip sensors use double-sided modules, with a small angle between sensors on either side to provide a 2D measurement, the same principle currently used by the SCT. Overall, the ITk will have a total silicon area of 14 m^2 in pixels, over ten times more than the existing pixel detector, and 165 m^2 in strips, around three times more than the existing SCT [178, 179]. The layout and rotation of each layer, shown in Figure 5.6, has been optimised to provide at least 13 hits for all tracks with $|\eta| < 2.7$, to ensure accurate reconstruction of inner detector tracks and their corresponding vertices.

The remainder of this chapter focuses on the ITk strip detector. The design of

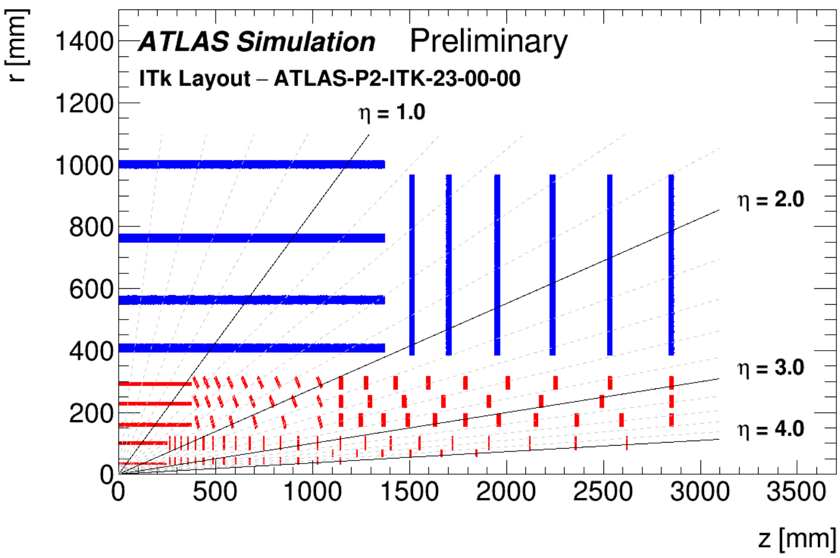

Figure 5.6: Schematic of the ITk layout, with pixel layers shown in red and strip layers shown in blue [180].

the new strip sensors is optimised for operation in a high radiation environment. They use n-type implants in p-type bulk (n-in-p, as in Figure 5.3), which has been shown to give larger, clearer signals following irradiation than p-in-n sensors, as used by the current SCT. This is also motivated by the increased trapping time for holes caused by annealing, which is less problematic in n-in-p sensors where the movement of electrons produces signals. They are also designed with AC coupling to take advantage of the reduction in leakage current, since the built-in capacitors only respond to the changing current produced by an incident charged particle, and not the constant leakage current produced in the bulk.

5.3 Test Structures

As part of the production of the strip sensors for the ITk, a Quality Assurance (QA) program is carried out to perform measurements of silicon test structures and verify that they meet the required specifications. The first year of the author's PhD was spent performing these measurements and improving the corresponding measurement setup for the contribution from Birmingham to the QA program. Tests are performed on dedicated test structures produced on the edge of the silicon wafers used for the main ITk sensors, both before and after irradiation, to ensure that the specifications will also be met following several years of intense irradiation during operation. The maximum expected fluence of radiation from the HL-LHC is estimated at 4.8×10^{14} n_{eq}/cm^2 in the barrel of ITk, with a Total Ionising Dose (TID) of 217 kGy, and 8.1×10^{14} n_{eq}/cm^2 in endcaps, with a TID of 333 kGy. Accounting for a safety factor of 2 in case of an underestimate, the test structures are therefore irradiated to 1.6×10^{15} n_{eq}/cm^2 and 660 kGy. Several radiation sources are used, with either a combination of neutron irradiation to cause bulk damage and gamma irradiation to cause ionising damage, or proton irradiation to achieve both effects simultaneously. Following irradiation, sensors are annealed at 60°C for 80 minutes to simulate the periodic warm-ups which are planned for the ITk during operation, to take advantage of the beneficial effects of annealing discussed in Section 5.1.1. The rest of the time, irradiated sensors are stored at −15°C to freeze out any further annealing effects. Three test structures are considered:

- Monitor Diodes (MD8s): simple diodes formed of a p–n junction
- Test chips: collections of several components from the main sensors on a single piece of silicon
- Mini sensors: miniaturised versions of the main strip sensors

5.3.1 Test Chips and MD8

Test chips are collections of various components which are used in the main strip sensors, and are produced alongside MD8s. The layout of the test chip and MD8 is shown in Figure 5.7, with each measured component labelled. In addition to the MD8, these components are:

- Bias Resistors: a collection of polysilicon bias resistors

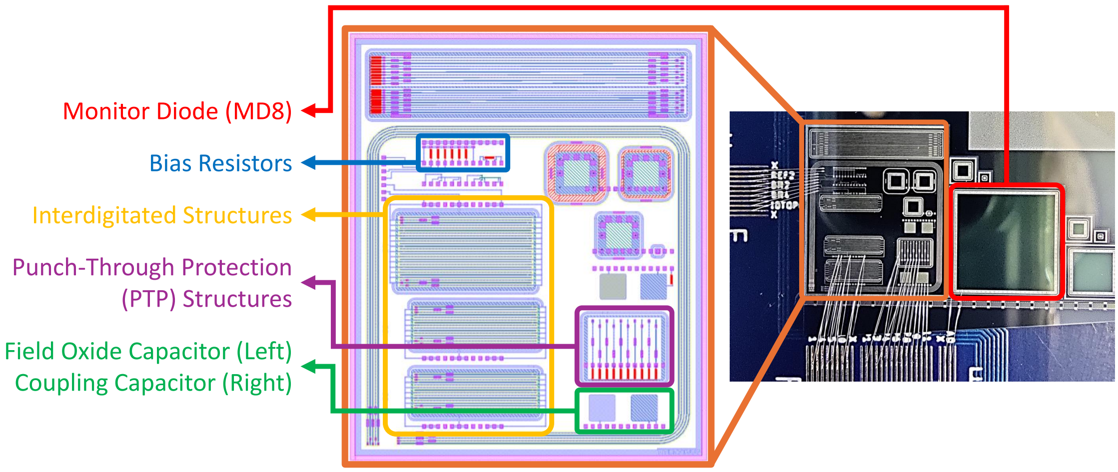

Figure 5.7: Test chip and MD8 layout.

- Interdigitated Structures: sets of interspaced strips where each strip is isolated from its nearest neighbours

- Punch-Through Protection (PTP) Structures: a set of channels connected at one end to a bias ring via bias resistors, with the other end isolated from the bias ring by a layer of oxide

- Coupling Capacitor: a square coupling capacitor

The test chip also contains several other unused structures. Test chips are fixed to custom-built Printed Circuit Boards (PCBs) via conductive tape and each component is wire-bonded to specified inputs to allow measurements to be made. The PCB is then connected to the measurement setup and placed inside a small box in a climate chamber, which cools the components to $-20°$C throughout the measurements. The box is flushed with nitrogen to ensure low humidity, as any condensation could result in damage to the components. The measurement setup uses a switching matrix, controlled via Python code, which allows any necessary connections to be made between the components on one axis and the instruments on the other, without manually changing connections. This allows various measurements of different components to be taken in sequence automatically. The measurement setup is shown in Figure 5.8.

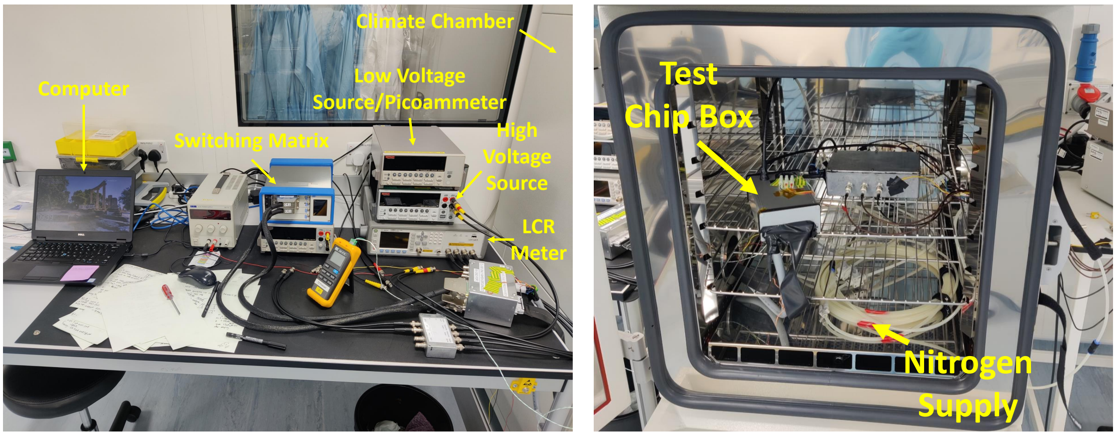

(a) (b)

Figure 5.8: (a) Test chip bench with the different measuring instruments labelled and (b) inside of the climate chamber where the box holding a wire-bonded test chip is placed for measurements.

MD8

MD8s are simple diodes formed of n^+-type implants on a p-type bulk, essentially large versions of the n-in-p junctions used in the main sensors. They have an active area of 0.5095 cm^2, and are surrounded by a guard ring and edge structure to isolate them from neighbouring components. Two tests are performed on MD8s, first a reverse bias voltage is applied across the diode and varied from $0 \leqslant V_{\text{bias}} \leqslant 700$ V (where here and throughout this section, voltage is implicitly negative, applied in the reverse direction). The leakage current through the diode is measured, and required to be below 100 μAcm^{-2} (again, implicitly negative) for irradiated MD8s for $V_{\text{bias}} \leqslant 500$ V, to ensure that the sensors will have sufficiently low noise and power consumption. In addition, the MD8s are required to display no breakdown for $V_{\text{bias}} \leqslant 700$ V, to ensure the sensors can be fully depleted to achieve full efficiency. Initially, the ITk sensors will be biased only to 500 V, although this may be increased to as high as 700 V to compensate for the effects of radiation damage. The presence of breakdown is assessed using the k-factor

$$k = \frac{dI/dV}{I/V}, \tag{5.6}$$

where dI/dV is the gradient of the I–V curve obtained under reverse bias. A significant spike in this distribution indicates a breakdown, as demonstrated in Figure 5.9 for an MD8 failing specification. In contrast, Figure 5.10 shows a well-behaved

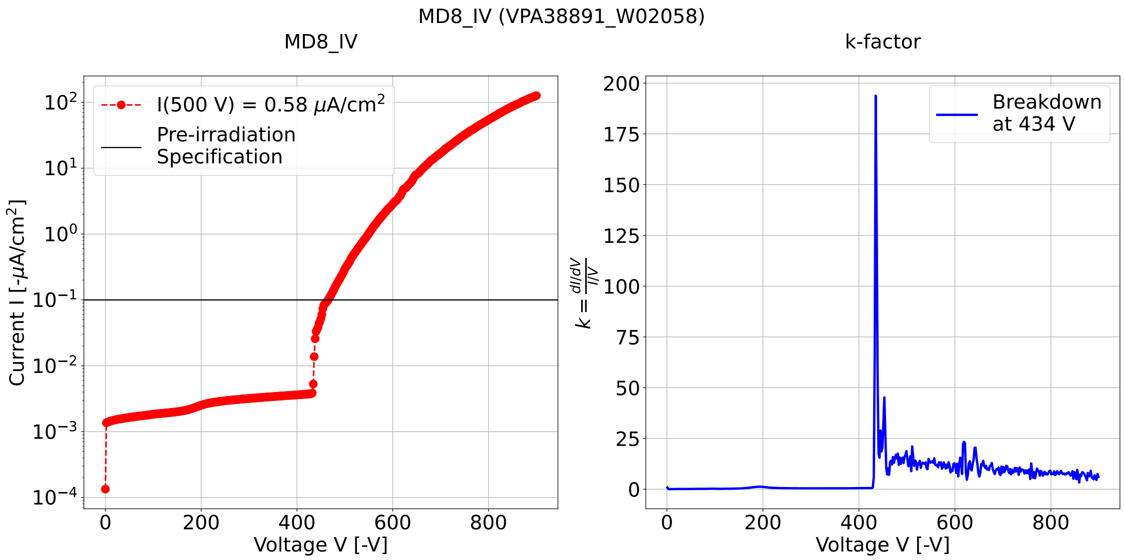

Figure 5.9: Example MD8 I–V curve and corresponding k-factor of an unirradiated test chip, showing early breakdown at $V_{\text{bias}} = 434$ V and relatively high leakage current of 0.58 μA/cm^2 at 500 V, failing both specifications for unirradiated test chips.

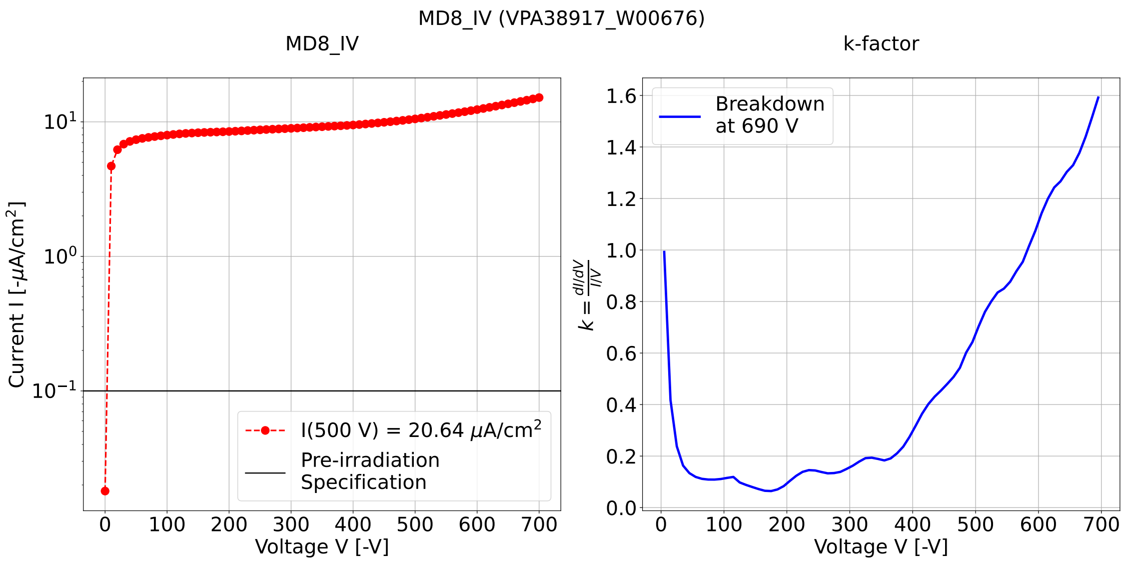

Figure 5.10: Example MD8 I–V curve and corresponding k-factor of an irradiated test chip, showing no breakdown below $V_{\text{bias}} = 700V$ and relatively low leakage current of 20.64 μA/cm^2 at 500 V, passing both specifications for irradiated test chips.

MD8 I–V with low leakage current and no breakdown. Figure 5.11 shows the post-irradiation leakage current measured for all MD8s investigated by the author during their contribution to the QA program, with all falling within the specifications.

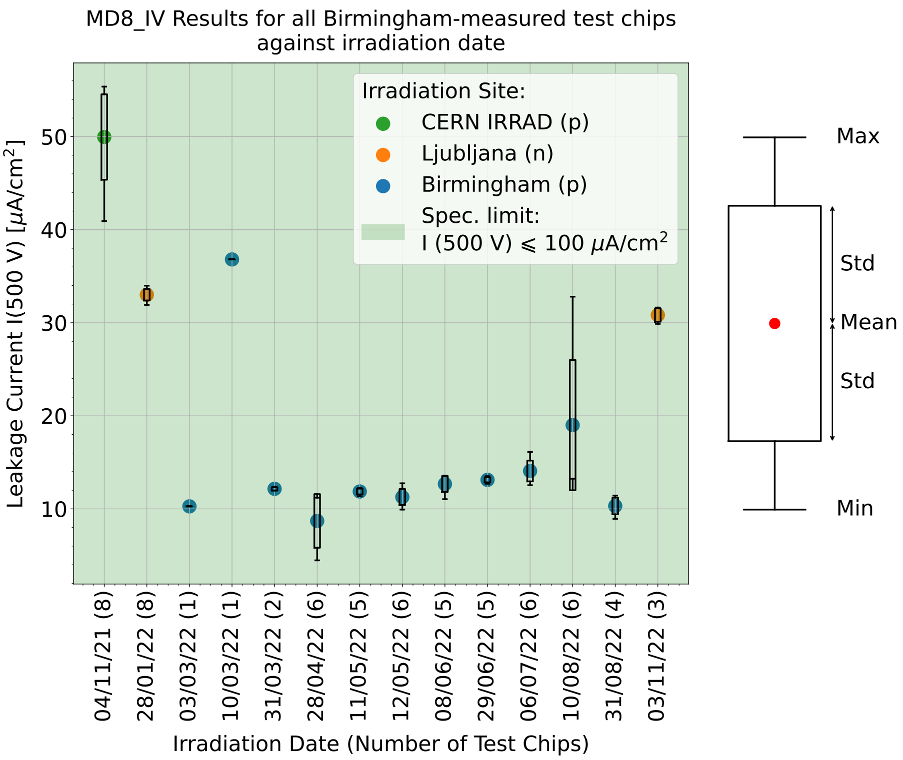

Figure 5.11: Leakage current measured at $V_{\text{bias}} = 500$ V for all irradiated MD8s investigated by the author, plotted against irradiation date. Combined results are shown for batches of test chips which were irradiated together, with the number of test chips in each batch indicated alongside the date of each irradiation. The locations of each irradiation are also indicated, with the CERN Proton Irradiation Facility (IRRAD) and Birmingham MC40 irradiating using protons and Ljubljana using a combination of neutrons and photons intended to be equivalent. For each batch, the result is presented as a box plot, with the mean, minimum and maximum measurements and the standard deviation all shown, as demonstrated to the right of the plot.

The second measurement of MD8s is of the capacitance across the diode for the same range of applied reverse bias voltages, in order to determine the full depletion voltage V_{FD}. Equation 5.4 gives the relation between the bias voltage and the depletion width W_d. Since the ITk sensors use n^+-type implants which are significantly more highly doped than the p-type bulk, $N_D \gg N_A$ and so this simplifies to

$$W_d = \sqrt{\frac{2\varepsilon_{\text{Si}}}{qN_A}(V_0 + V_{\text{bias}})}. \tag{5.7}$$

The depletion region is saturated with non-mobile charge carriers, and so it acts as a dielectric insulator between the n-type and p-type silicon, causing the p–n junction to act as a parallel plate capacitor with width W_d and corresponding capacitance

$$C = \frac{\varepsilon A}{W_d} = \sqrt{\frac{\varepsilon q N_A}{2\,(V_0 + V_{\text{bias}})}}A, \tag{5.8}$$

where A is the active area of the sensor. Therefore, for high $V_{\text{bias}} \gg V_0$, capacitance is proportional to $V_{\text{bias}}^{-1/2}$, so the quantity $1/C^2$ is directly proportional to V_{bias}. Once full depletion is reached for $V_{\text{bias}} \geqslant V_{\text{FD}}$, W_d stops increasing and so C becomes constant for increasing V_{bias}. Therefore, V_{FD} can be extracted from the distribution of $1/C^2$ against V_{bias} as the point of transition between non-zero and zero gradient linear fits. Figure 5.12 shows an example distribution, with the intersection between the two linear fits taken as the measurement of V_{FD}[1]. The ITk specifications require

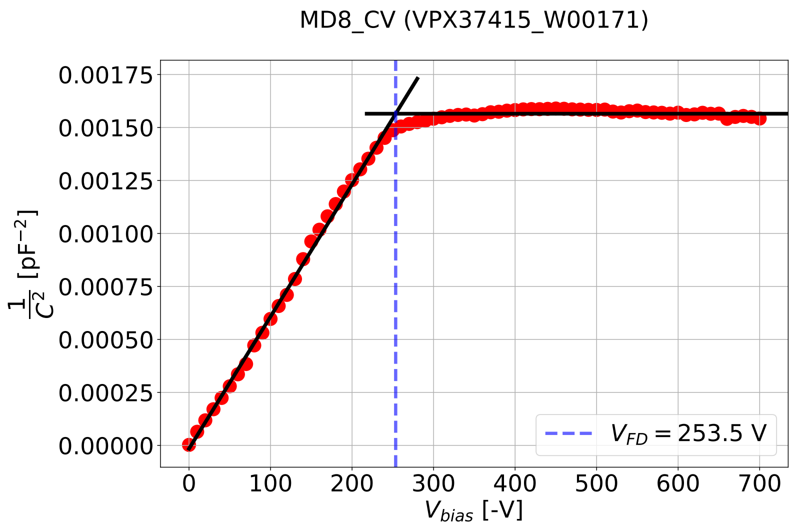

Figure 5.12: MD8 C–V showing two linear fits of the regions above and below V_{FD}, with their intersection giving a measurement of V_{FD}.

$V_{\text{FD}} < 350$ V for unirradiated MD8s to ensure full depletion is reached before sensor

[1] In this fit, the transition point between the constant and linear regions was taken as the point at which the value of $1/C^2$ reached within 1/16 of its maximum measured value.

breakdown.

Bias Resistors

The first test chip components are a set of three polysilicon resistors identical to those that separate the strip implants from the bias ring on the main sensors, referred to as the bias resistors. Voltage is applied across the resistors and varied in the range $-5 \leqslant V \leqslant 5$ V, the resulting current is measured and the resistance is extracted as $R = \frac{dV}{dI}$. The measurements are temperature corrected from the measurement temperature of $-20°$C to the specification temperature of $20°$C via

$$R(T_1) = R(T_0)e^{\alpha\left(\frac{1}{T_1} - \frac{1}{T_0}\right)}, \tag{5.9}$$

where $\alpha = 312.2$ K was been obtained in previous studies [172]. The specification requires the average R_{bias} across the three resistors to fall within the range 1.5 ± 0.5 MΩ, with a maximum of 0.5 MΩ between the minimum and maximum values. Figure 5.13a shows an example I–V curve with extracted R_{bias} measurements. Figure 5.13b shows the results from every test chip investigated by the author, with 98% of test chips passing the specifications.

Interdigitated Structures

The next test chip components are the interdigitated structures. Each structure is composed of two sets of n^+-type strips implanted on a p-type bulk, arranged such that each strip is neighboured on either side by strips from the other set. Adjacent strips are separated by a layer of oxide and the whole structure is surrounded by a bias ring and a guard ring. An example interdigitated structure is shown in Figure 5.14. Each test chip contains three structures, referred to as UP, MID and LOW based on their relative positions on the test chip, and the strips have identical length and pitch to the main sensors. The structures are fully depleted with a reverse bias voltage of 500 V and then voltage is applied between the two sets of strips and varied from $-5 \leqslant V_{\text{int}} \leqslant 5$ V. The resulting inter-strip current I_{int} and capacitance C_{int} of each structure are measured.

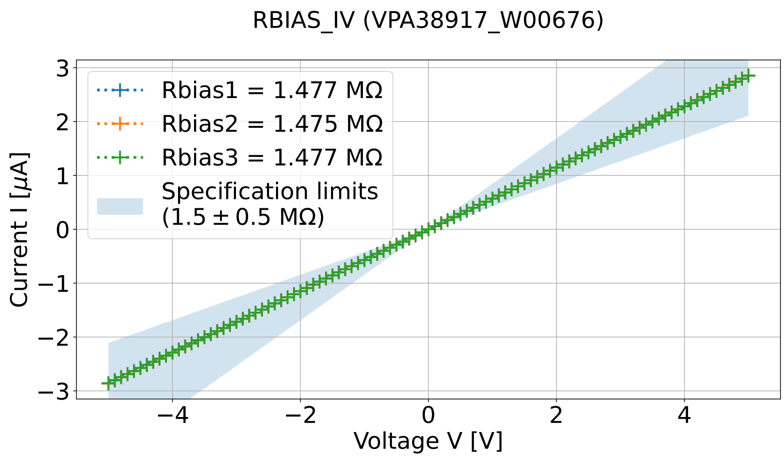

(a)

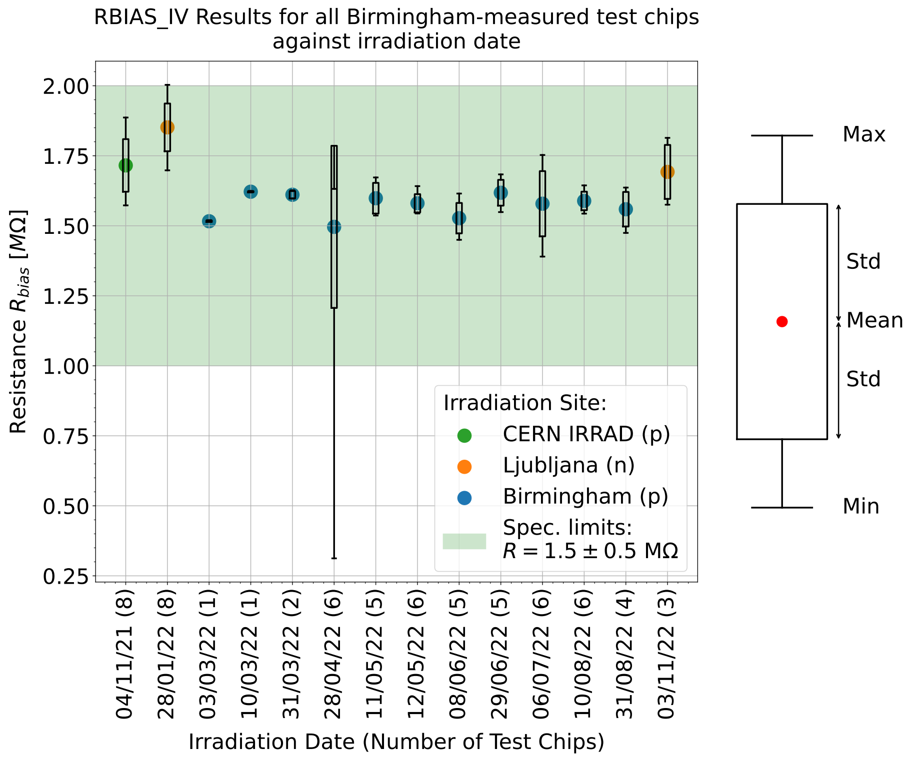

(b)

Figure 5.13: (a) Example I–V plots for each bias resistor on an irradiated test chip, with extremely similar results obtained for each structure, such that the resulting distributions overlap and not all are visible. (b) R_{bias} measurements for every test chip investigated by the author, as a function of irradiation date. Combined results giving the mean, minimum and maximum measurements and the standard deviation are shown for batches of test chips which were irradiated together, with the number of test chips in each batch indicated alongside the date of each irradiation. The location and particle type of each irradiation are also indicated.

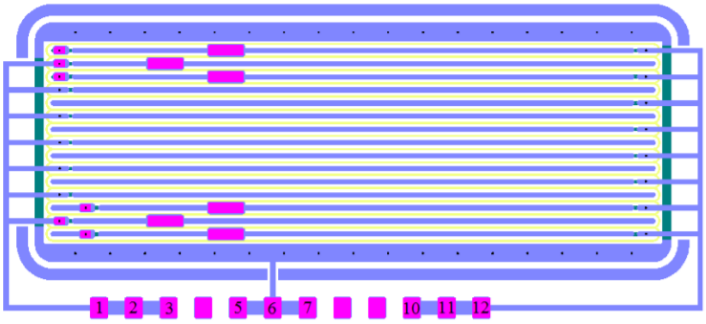

Figure 5.14: Diagram of an interdigitated structure, with two sets of isolated n-type strip implants arranged on top of a p-type bulk such that each strip is neighboured on either side by strips from the other set. The structure is surrounded by a bias ring to allow a voltage to be applied to the bulk, and then a guard ring to isolate the structure from other test chip components. Connecting pads (purple, numbered) are used to make electrical connection to each set of strips through Pads 1-3 and 10-12, respectively, and to the bias ring via Pads 5-7.

The inter-strip resistance R_{int} is extracted from the I–V curve, and the specification requires that $R_{\text{int}} \geqslant 10 \cdot R_{\text{bias}}$, i.e. $R_{\text{int}} \geqslant 15$ MΩ, for every structure. This is to ensure sufficient isolation between neighbouring strips in the main detector, to avoid cross-talk which would cause reduced detector precision and lower charge collection per strip. Figure 5.15a shows the I–V curves measured for each structure on an example sensor, and Figure 5.15b shows the results of all measurements performed by the author. Generally, very low values of R_{int} were obtained, with only around 20.0% of each type of structure passing the specification and a wide spread between measurements of the three structures on the same test chip. Several additional investigations were performed into this poor performance, detailed in Section 5.3.2.

Inter-strip capacitance it also measured across the same range of voltages. A correction is applied to the raw measurements to subtract the intrinsic capacitance of the measurement setup, obtained from measurements taken with an empty PCB. Corrected values are divided by the length of the corresponding structure to give measurements of capacitance per unit length, with the specification requiring $C_{\text{int}}(0\,V) \leqslant 2$ pF/cm, low enough to avoid excess noise. Figure 5.16a shows an example measurement across all three structures, showing consistent values as a function of voltage and generally within specification. The MID structure results

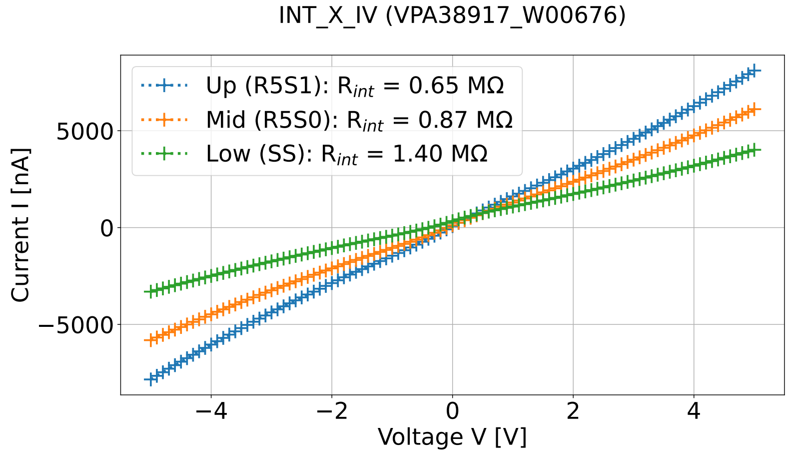

(a)

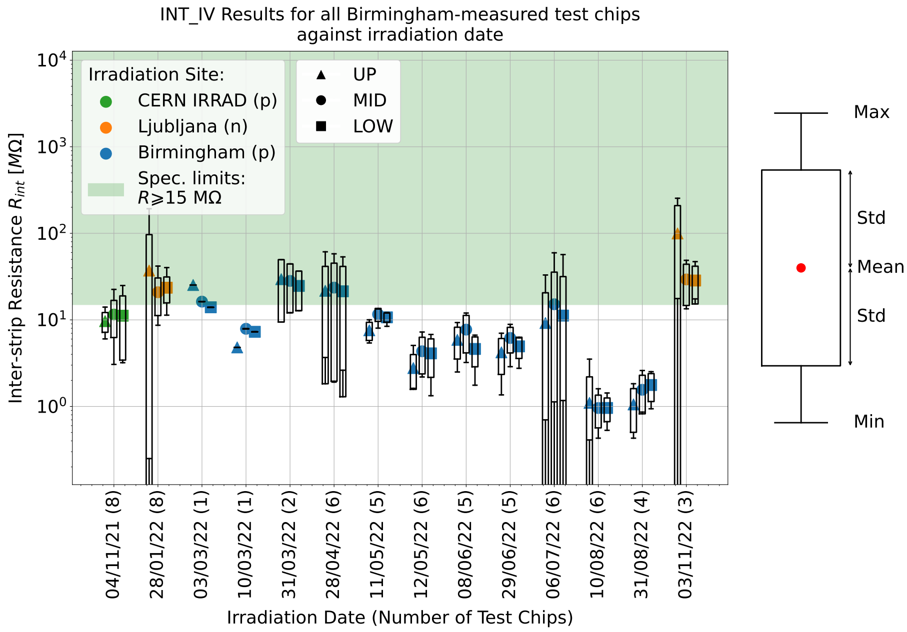

(b)

Figure 5.15: (a) Example inter-strip I-V curves for each interdigitated structure on an irradiated test chip and (b) R_{int} measurements for every test chip investigated by the author, split by structure type, as a function of irradiation date. Combined results giving the mean, minimum and maximum measurements and the standard deviation are shown for batches of test chips which were irradiated together, with the number of test chips in each batch indicated alongside the date of each irradiation. The location and particle type of each irradiation are also indicated.

are slightly too high, which was a common trend seen across a significant proportion of test chips. Figure 5.16b shows the results for all test chips investigated by the author, with 100% of UP and LOW structures and 80.3% of MID structures passing the specification.

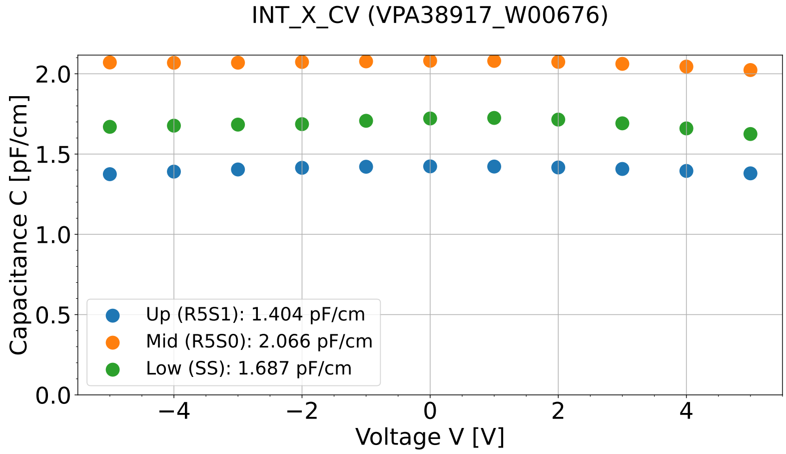

(a)

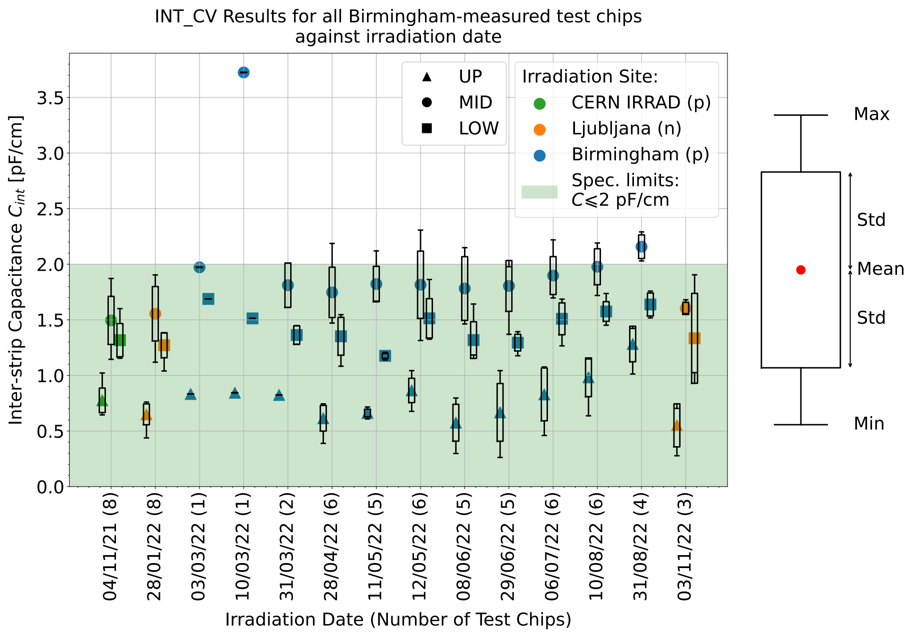

(b)

Figure 5.16: (a) Example inter-strip C–V plots for each interdigitated structure on an irradiated test chip and (b) C_{int} measurements for every test chip investigated by the author, split by structure type, as a function of irradiation date. Combined results giving the mean, minimum and maximum measurements and the standard deviation are shown for batches of test chips which were irradiated together, with the number of test chips in each batch indicated alongside the date of each irradiation. The location and particle type of each irradiation are also indicated.

Coupling Capacitor

The coupling capacitors are square n^{+}-type implants on top of the p-type bulk, with an oxide layer deposited on top, followed by a metal contact. This forms a parallel plate capacitor matching those found on the main sensors connecting the implants to the electrodes, with the oxide as the dielectric. The leakage current I_{CPL} and capacitance C_{CPL} across the structure are measured.

Leakage current is measured as a function of voltage up to 100 V, with an initial high granularity check up to 10 V to ensure there is no dielectric breakdown at low voltages, referred to as a pinhole. The specification requires $I_{\mathrm{CPL}}(100\ \mathrm{V}) \leqslant 10$ nA to ensure that no dielectric breakdown occurs which could allow a large current to flow to the read-out electronics, causing damage. Figure 5.17a shows an example I–V curve, with very low leakage current at all voltages, well within the specification. The observed fluctuations are attributed to measurement uncertainty, which is significant at such low values of measured current. This was consistently found across the majority of test chips, as shown in Figure 5.17b, with 93.4% passing the specification, and four structures failing due to pinholes.

The coupling capacitance C_{CPL} is measured across the structure with no voltage applied. To account for the intrinsic capacitance of the measurement setup, a correction C_{ref} is measured between reference pads which are not connected to the structure (Pads 7-9 and 12 in Figure 5.18). The measurement C_{meas} is then taken between Pads 10 and 11, and the corrected coupling capacitance per unit length is calculated as

$$C_{\mathrm{CPL}} = \frac{C_{\mathrm{meas}} - C_{\mathrm{ref}}}{\ell}, \tag{5.10}$$

where $\ell = 3.4$ cm is the equivalent length of the capacitor on the main sensor. The specifications require $C_{\mathrm{CPL}} \geqslant 20$ pF/cm, to ensure that it is much greater than the inter-strip capacitance, to allow high charge collection efficiency at the electrodes, giving clear signals. Figure 5.19 shows the results of all measurements performed by the author, with all passing the specification.

PTP Structure

The Punch-Through Protection (PTP) structure, shown in Figure 5.20 is composed of a series of n^{+}-type strip implants deposited on the p-type bulk, all connected

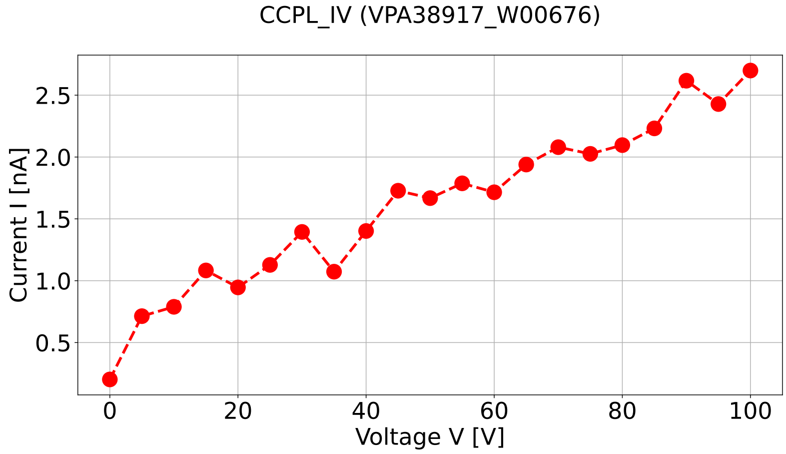

(a)

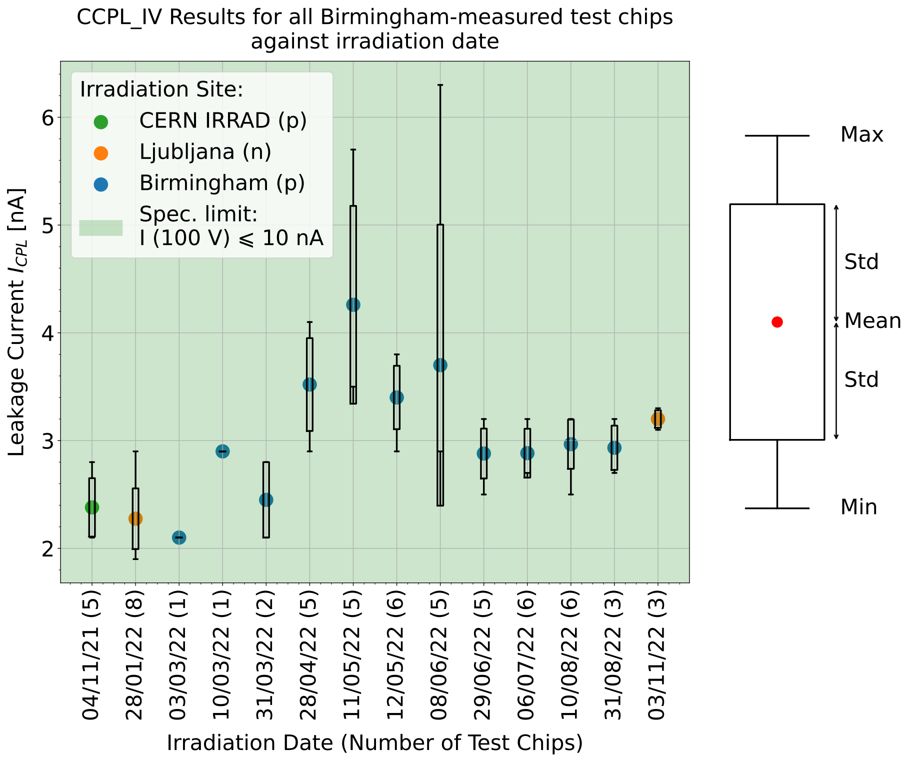

(b)

Figure 5.17: (a) Example leakage current I-V curve for the coupling capacitor on an irradiated test chip and (b) I_{CPL} measurements for every test chip investigated by the author, as a function of irradiation date. Combined results giving the mean, minimum and maximum measurements and the standard deviation are shown for batches of test chips which were irradiated together, with the number of test chips in each batch indicated alongside the date of each irradiation. The location and particle type of each irradiation are also indicated.

at one end via bias resistors to a bias ring surrounding the structure, and isolated at the other end by a layer of oxide. The structure is reverse biased to -500 V

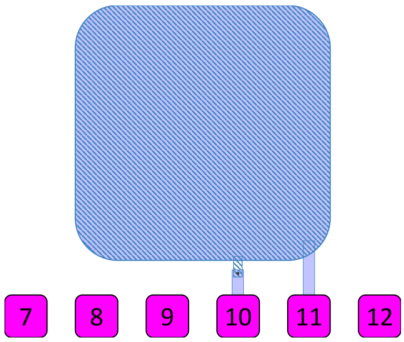

Figure 5.18: Diagram of the coupling capacitor, showing the n-type implant deposited on top of the p-type bulk. Connecting pads (purple, numbered) are used to make electrical connections, with the p and n-type sides of the capacitor connected to Pads 10 and 11, respectively, and additional unconnected reference pads provided to allow a reference measurement of the capacitance of the measurement setup to be obtained.

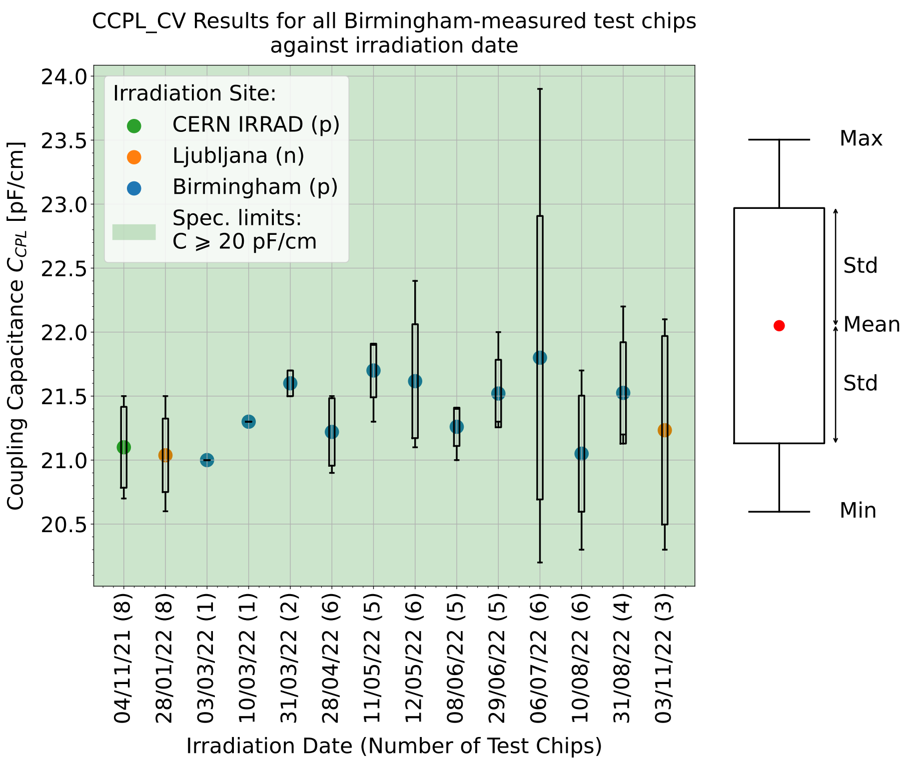

Figure 5.19: C_{CPL} measurements for every test chip investigated by the author, as a function of irradiation date. Combined results giving the mean, minimum and maximum measurements and the standard deviation are shown for batches of test chips which were irradiated together, with the number of test chips in each batch indicated alongside the date of each irradiation. The location and particle type of each irradiation are also indicated.

to fully deplete the bulk and voltage is applied between the strips and the bias rail. The current is measured and the voltage increased until the effective resistance $R_{\text{eff}} = \frac{V}{I}$ falls below half the resistance of the bias resistors, i.e. $R_{\text{eff}} < R_{\text{bias}}/2$, i.e. $R_{\text{eff}} <= 0.75\,\text{M}\Omega$. The voltage at which this threshold is crossed is taken as the punch-through voltage V_{PTP}, corresponding to the voltage at which the resistance of the oxide layer insulating the strips from the bias rail falls below that of the bias resistors. This measurement is performed on each of the central five strips (Pads 4-8 in Figure 5.20) and the specifications require the average value to fall within $5 \leqslant V_{\text{PTP}} \leqslant 50$ V. The sensors are designed such that this punch-through will

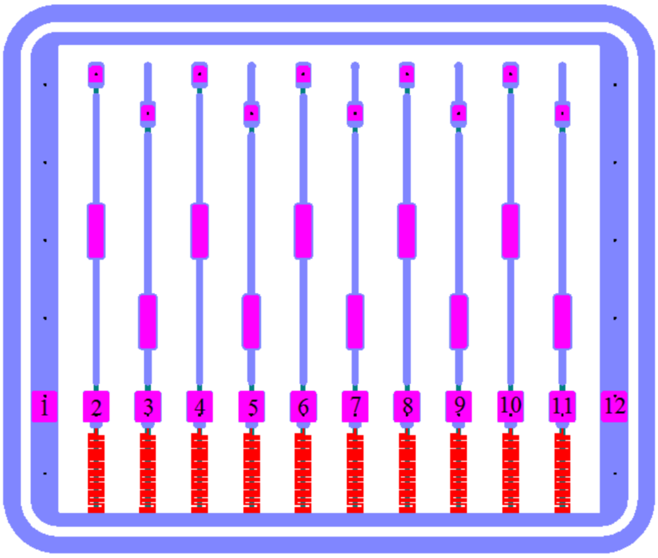

Figure 5.20: Diagram of the PTP structure, composed of n^{+}-type strip implants on top of a p-type bulk, surrounded by a bias ring which allows a voltage to be applied to the bulk. Each strip is connected to the bias ring via a bias resistor (red), with connecting pads (purple, numbered) to allow electrical connections to the bias ring and each individual strip. The whole structure is surrounded by a guard ring to isolate it from other test chip components.

occur at voltages significantly below the minimum dielectric breakdown voltage of the coupling capacitors of 100 V, to divert high currents which could be produced by scattered high energy beams safely into the bias rail, away from the sensitive read-out electronics, to prevent damage. Figure 5.21a shows an example I–V curve, with the corresponding R_{eff} calculated in Figure 5.21b. This shows the expected behaviour, with initially constant resistance which decreases sharply as the oxide isolation breaks down, with the punch-through threshold reached at around 28 V.

Figure 5.22 shows the results from all test chips investigated by the author, with all

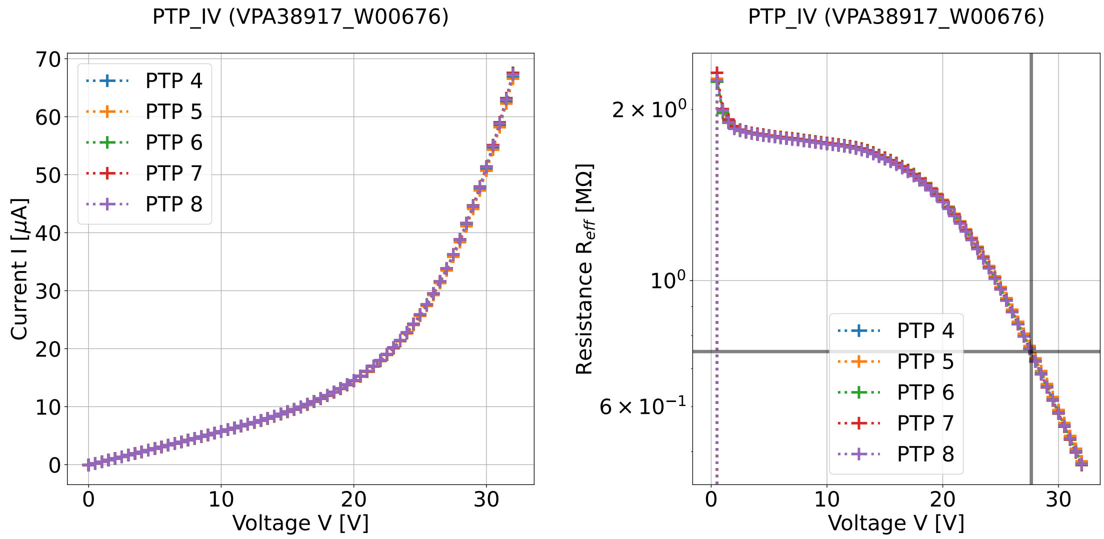

(a) (b)

Figure 5.21: (a) Example I–V plots and (b) calculated R_{eff} for each PTP structure measured on an irradiated test chip. Extremely similar results are obtained for each structure, such that the plotted distributions overlap and not all are visible.

showing consistent results and passing the specification.

5.3.2 Inter-strip Resistance Investigation

As shown in Figure 5.15b, the inter-strip resistance measurements obtained in Birmingham were consistently very low, with many falling below the specification. However, similarly low measurements were not seen by other institutes performing the same QA measurements. To try to understand this difference, and the cause of the extremely low inter-strip resistance measurements obtained in Birmingham, a series of investigations were undertaken, primarily by the author. As a first step, two test chips from a single irradiation were sent to Toronto to be measured in their equivalent QA setup, with the remaining four measured in Birmingham. Table 5.1 shows the results, with three of the four test chips measured in Birmingham recording very low values of R_{int}, with the only exception being `VPX37415-W172` from a much earlier production in a batch which has been extensively studied and shown to perform well. On the other hand, both test chips measured in Toronto showed

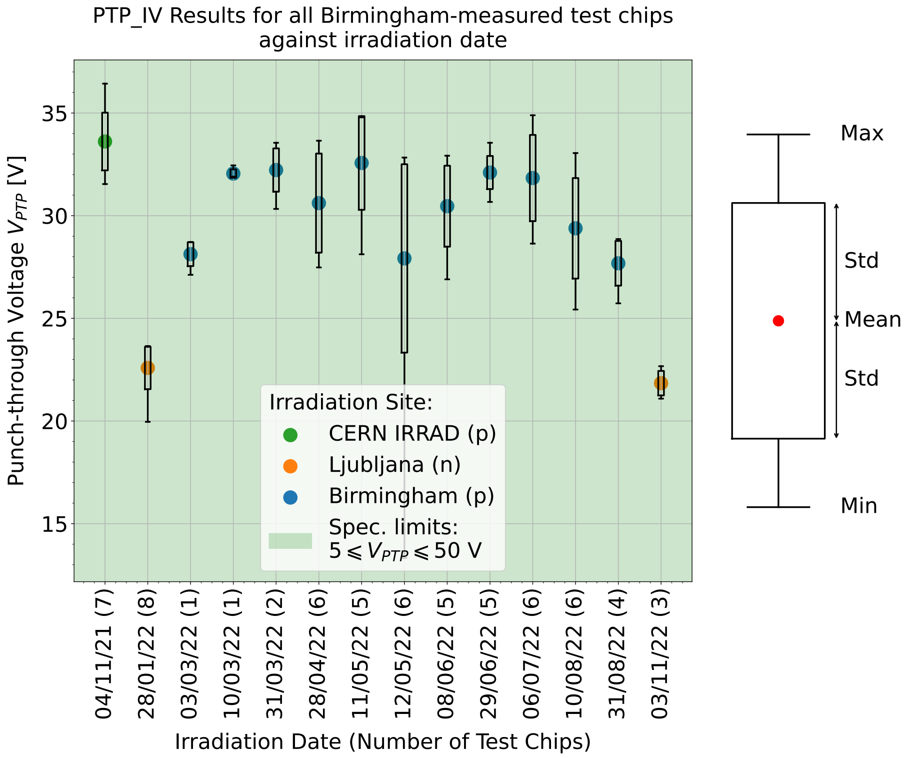

Figure 5.22: V_{PTP} measurements for every test chip investigated by the author, as a function of irradiation date. Combined results giving the mean, minimum and maximum measurements and the standard deviation are shown for batches of test chips which were irradiated together, with the number of test chips in each batch indicated alongside the date of each irradiation. The location and particle type of each irradiation are also indicated.

Test Chip	VPX37415-W172	VPA38906-W277		VPA38208-W439	
Measurement Site	Birmingham				Toronto
Temperature [°C]	-20		-27	-20	-27
R_{int} UP [MΩ]	61.15	3.68	7.60	6.53	17.60
R_{int} MID [MΩ]	57.97	1.97	4.06	1.91	4.11
R_{int} LOW [MΩ]	53.44	2.62	5.43	3.39	9.02

Test Chip	VPA38699-W1901	VPA38690-W1696		VPA38697-W1791	
Measurement Site	Birmingham	Toronto	Birmingham	Toronto	Birmingham
Temperature [°C]	-20	-27			
R_{int} UP [MΩ]	7.30	67.74	23.13	82.84	28.27
R_{int} MID [MΩ]	10.48	164.80	46.29	65.21	22.89
R_{int} LOW [MΩ]	10.04	156.64	44.61	36.42	14.23

Table 5.1: Results of R_{int} measurements made in Birmingham and Toronto for test chips from the same irradiation, at different temperatures.

very high values of R_{int}, well above specification. This caused concerns about the accuracy of the measurement setup in Birmingham and so, as a cross-check, the two test chips were returned from Toronto to Birmingham, while VPA38208-W439 was sent to Toronto, and all were remeasured. The test chip remeasured in Toronto showed higher R_{int} values by around a factor of three, although it was noted at this point that Toronto perform QA measurements at the lower temperature of -27°C, which is expected to cause the resistance to increase by around a factor of two. This was confirmed upon remeasurement of VPA38906-W277 in Birmingham at -27°C, resulting in twice the measured R_{int}. Therefore, the two test chips originally measured in Toronto were remeasured in Birmingham at -27°C. However, they still yielded reduced R_{int} compared to the Toronto measurements, although significantly higher than the other test chips measured in Birmingham, mostly passing the specification. This suggested that there was some physical difference between these two test chips and the rest of the batch, perhaps due to additional annealing or humidity exposure during transport between the two institutes.

To assess whether humidity exposure post-annealing, such as that received in the fairly high humidity cold storage at Birmingham, could be having an effect, another test chip from the control batch VPX37415-W174, which has been previously measured, was stored in a dry cabinet overnight to thoroughly dry it, and then remeasured. The results are shown in Table 5.2, with no significant difference in R_{int} observed. Therefore, to go even further in ensuring the sensor was dry, it was

Test Chip	VPX37415-W174		
Conditions	Normal	Overnight Drying	Overnight Annealing (1100 minutes)
R_{int} UP [MΩ]	5.07	5.09	13.65
R_{int} MID [MΩ]	4.98	5.00	12.35
R_{int} LOW [MΩ]	3.90	3.91	8.92

Table 5.2: Results of R_{int} measurements of the test chip VPX37415-W174 in normal conditions and after extensive drying and annealing.

annealed overnight for a total of 1100 minutes at 60°C. After remeasuring again, significantly higher R_{int} values were obtained, suggesting that a change had occurred within the sensor. This could be due to the extra annealing received, although previous measurements consistently show lower measured R_{int} following annealing, and the effects of annealing have been shown to saturate after the standard 80 minutes.

Since pre-annealing measurements of R_{int} are consistently above specification, and only drop to the very low values observed following annealing, it was concluded that any physical changes causing the particularly low R_{int} measured in Birmingham must occur during or following annealing.

To test this, five test chips from a new irradiation were measured using a special procedure:

1. Remove the unannealed, irradiated test chip from the normal cold storage
2. Place the test chip in the drying cabinet for at least 30 minutes to warm up in a dry environment
3. Fix and wire-bond the test chip to a PCB
4. Immediately anneal the test chip on the PCB
5. Upon removal from the oven, place the test chip straight into the climate chamber and begin drying immediately, to ensure that it does not cool down in a humid environment
6. Cool to -20°C in a nitrogen-flushed, low humidity environment and measure the test chip

This ensures that the test chips spend no time in the cold storage following annealing, avoiding the possibility of humidity related effects on the inter-strip resistance. However, the results still showed very low values of R_{int}, with none passing the specification. Therefore, it was considered next that the humidity exposure could occur during the annealing process, which is performed in a lab grade oven without humidity control. To test this, three test chips from the same irradiation were instead annealed at room temperature (21°C) in a nitrogen-flushed dry cabinet at $< 5\%$ humidity. The annealing time was increased to compensate for the decrease in temperature. The literature conversion factor from the required annealing time at 60°C to that at room temperature is 325, based on previous studies of p-in-n

sensors [181, 182]. However, more recent studies on n-in-p sensors suggest a lower factor of 100 [183]. Therefore, the test chips were measured after both 100 and 325 times the normal annealing time (5.6 and 18 days), and for two of the test chips an additional intermediate measurement was made after 2.6 days to more closely study the time dependence of the change in R_{int} with annealing. The results are shown in Figure 5.23, with no significant deviation from the results obtained with the normal annealing procedure, and the most significant drop in R_{int} occurring after the initial intermediate annealing period. This confirms that the normal annealing

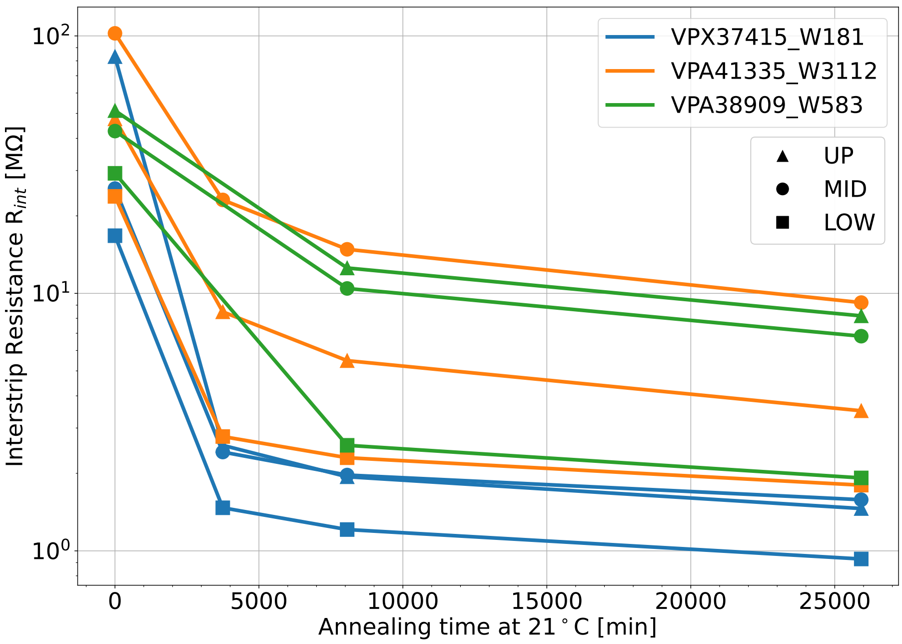

Figure 5.23: Results of R_{int} measurements after a series of total annealing times at room temperature (21°C) and low humidity (5%).

conditions are not the cause of the very low inter-strip resistance measurements. This concluded the author's investigations into this effect. Based on subsequent investigation, it has been concluded that the low inter-strip resistance occurred due to anomalous surface damage during proton irradiation using the MC40 cyclotron in Birmingham. Therefore, this facility is no longer used for ITk test chip irradiations, and sensors irradiated using a combination of neutron and gamma sources and measured in Birmingham now show much higher R_{int} values, well within the specifications.

5.3.3 Mini Sensors

The final component measured for the ITk QA program is the “mini sensor”, miniaturised versions of the main ITk strip sensors, with a p-type bulk and n^{+}-type strip implants. In Birmingham, these are irradiated with protons using the MC40 cyclotron to a fluence of $1.6 \times 10^{15} \mathrm{n_{eq}cm^{-2}}$. Following this they are annealed for 80 minutes at 60°C, the same as the other components, and their Charge Collection Efficiency (CCE) is measured.

Mini sensor measurements are performed using a specialised system called “AliBaVa” [184], consisting of a motherboard and connected daughter board, to which the mini sensor is connected, with each strip separately wire-bonded to a corresponding read-out. The daughter board and sensor are placed beneath a ^{90}Sr β^{-} radiation source which is incident on the sensor surface, and the total flux of electrons detected by the mini sensor is measured and compared to measurements from a scintillator positioned below the daughter board. During measurements, the bulk of the mini sensor is reverse biased to -500 V to achieve full depletion and maximal charge collection and efficiency, and the daughter board and sensor are cooled in a freezer to $\leqslant -25$°C to minimise the resulting leakage current and prevent any further annealing. The freezer is additionally flushed with nitrogen to keep the humidity below 10% to prevent condensation. The full measurement setup is shown in Figure 5.24.

Measurements

The primary measurement made of mini sensors is the Charge Collection Efficiency (CCE), although this actually measures the magnitude of the charge collected by the sensor, instead of measuring any efficiency. Measurements of CCE and leakage current are performed at reverse bias voltages in the range $100 \leqslant V_{\mathrm{bias}} \leqslant 1000$ V. To measure the collected charge at each voltage, first a “pedestal” run is performed to determine the baseline electronic noise and background radiation detected

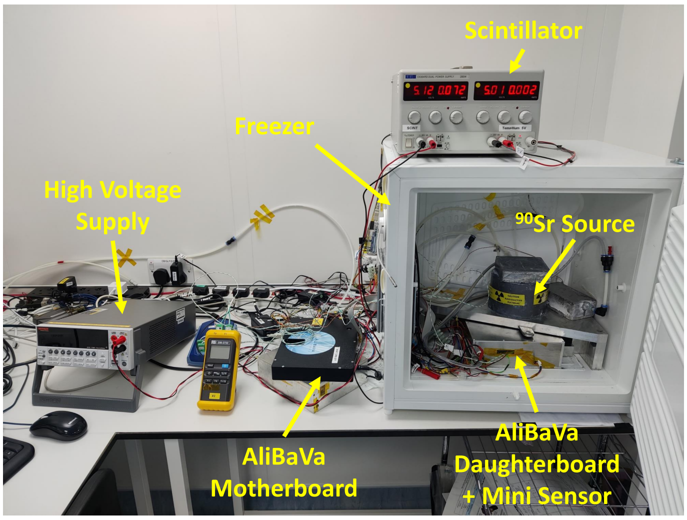

Figure 5.24: Mini sensor measurement setup, with the AliBaVa daughter board placed in the freezer, along with a ^{90}Sr β^- source, with the AliBaVa motherboard also connected to a scintillator and high voltage source-meter.

by the system, which is subtracted from a following "radiation" source run which records the collected charge from the first 100,000 charged particles incident on the sensor. Figure 5.25 shows two example spectra produced from these measurements at different bias voltages, which are fitted to a convolution of Gaussian and Landau distributions, to account for the sensor resolution and the stochastic nature of electron energy loss, respectively. The peak value $\langle q \rangle$ is extracted, corresponding to the most probable value of collected charge. The spectrum in Figure 5.25 obtained with $V_{\text{bias}} = 1000$ V has a much higher peak, as expected due to the wider depletion region caused by the higher reverse bias voltage.

The AliBaVa read-out chip produces an analogue signal which undergoes Analogue-to-Digital Conversion (ADC) before being outputted by the motherboard, giving the units shown in Figure 5.25. This is then converted to units of electrons (e), proportional to the number of electron-hole pairs produced in the depletion region of a sensor by an incident charged particle, using the expected value of 23,050 electron-hole pairs produced by a Minimally Ionising Particle (MIP) passing through a fully depleted, unirradiated sensor with the dimensions of the mini sensor. Here, a MIP is a particle with minimal energy loss per unit distance travelled in a medium (described by the Bethe-Bloch equation [175]), similar to high energy particles pro-

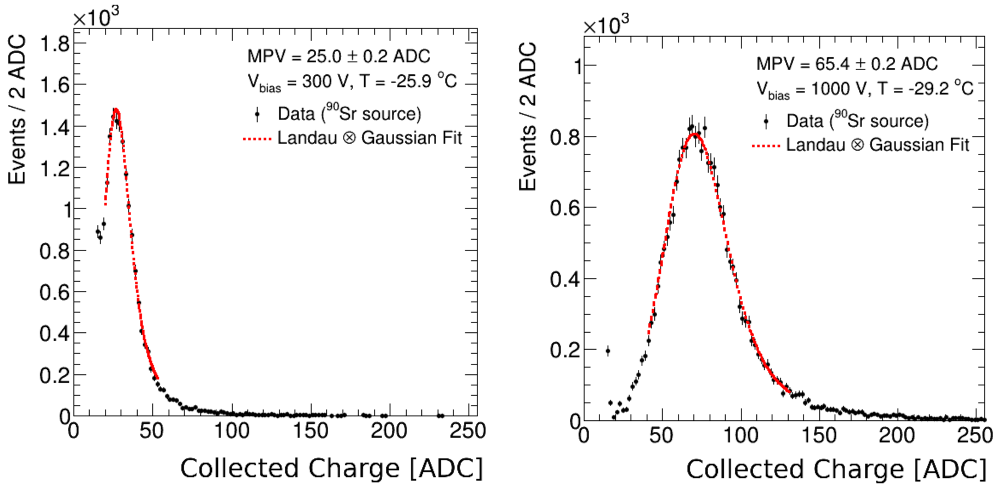

(a) $V_{\text{bias}} = 300$ V (b) $V_{\text{bias}} = 1000$ V

Figure 5.25: Collected charge spectra obtained for 100,000 electron hits on a mini sensor with the backplane reverse biased to (a) 300 V and (b) 1000 V, with fitted Landau ⊗ Gaussian curves shown in red, and the corresponding best-fit parameters given.

duced in LHC collisions. The peak collected charge $\langle q_{\text{ADC}} \rangle_{\text{FD}}$ is measured in ADC counts for a fully depleted sensor, and then used to convert an arbitrary ADC measurement $\langle q_{\text{ADC}} \rangle$ to the equivalent value in electron units $\langle q_{\text{electrons}} \rangle$ as

$$\langle q_{\text{electrons}} \rangle = \frac{23{,}050}{\langle q_{\text{ADC}} \rangle_{\text{FD}}} \langle q_{\text{ADC}} \rangle . \tag{5.11}$$

An additional temperature correction must be applied to measurements, as the mini sensor temperature consistently decreases throughout data-taking as the freezer continues to cool down once crossing the -25°C threshold, which affects the gain of the read-out chips. To obtain the correction, the CCE of a fully depleted, unirradiated sensor is measured at a range of temperatures from room temperature down to the measurement temperature. The temperature dependence of the read-out gain is then fitted to a linear distribution as shown in Figure 5.26, and the resulting correction is applied to the raw CCE measurements before use in Equation 5.11.

Figure 5.27 shows an example of the temperature-corrected CCE measurements in units of ke (1000 e) for a series of irradiated mini sensors as a function of bias voltage, both before and after annealing. In all cases the collected charge increases with

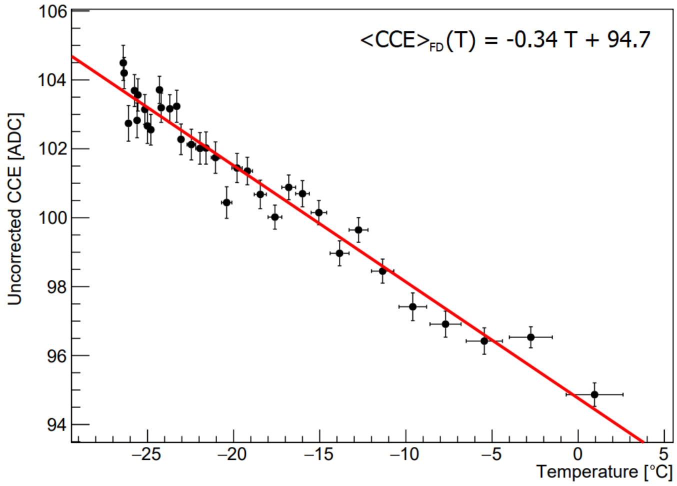

Figure 5.26: Linear fit of uncorrected peak collected charge in ADC counts as outputted by the read-out chip as a function of the measurement temperature.

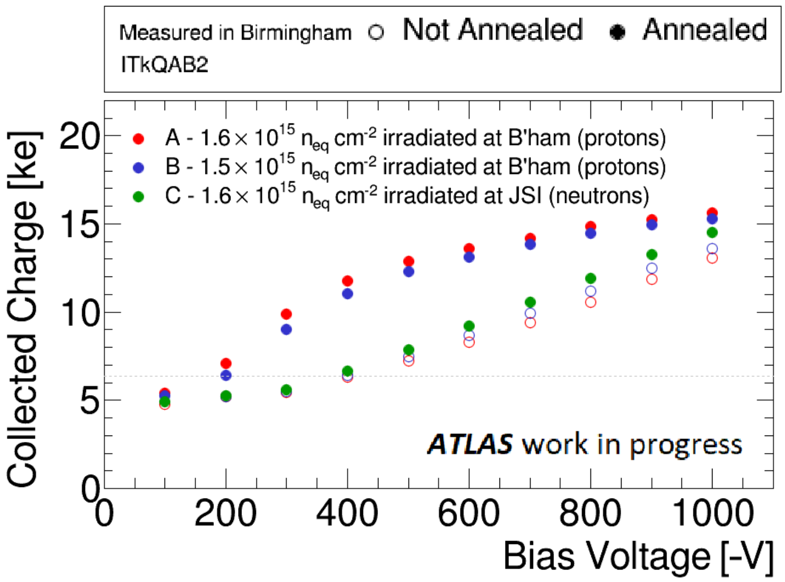

Figure 5.27: Temperature-corrected most probable collected charge in electron units plotted as a function of voltage for three irradiated mini sensors A, B and C, before and after annealing. A and B were irradiated with protons at Birmingham and C was irradiated using neutrons at the Jožef Stefan Institute (JSI) in Ljubljana.

bias voltage due to the increased depletion width, levelling out when full depletion is reached. The collected charge is much higher for irradiated sensors following annealing, due to the resulting partial repair of the bulk damage received by the sensors. It can also be seen that the neutron irradiated sensor `VPX37420-W281` has much lower collected charge than the proton irradiated sensors, suggesting that neutron irradiation causes greater bulk damage. All three sensors fall well below the expected charge collection of 23.05 ke discussed above, due to the negative effects of radiation damage, although all pass the ITk specifications which require $\langle q_{\text{electrons}} \rangle(500\text{ V}) \geqslant 6.35$ ke. Figure 5.28 shows the results from all irradiated mini sensors measured during the author's involvement with the QA program, with 100% passing the specifications.

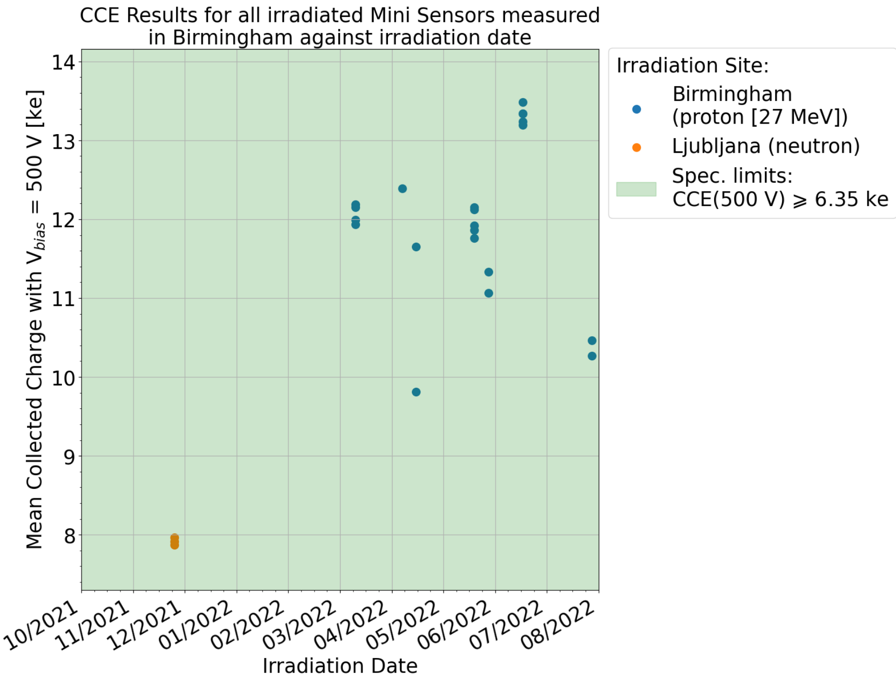

Figure 5.28: Most probable collected charge for every mini sensor investigated by the author in Birmingham, with a reverse bias voltage of 500 V applied. The irradiation site and particle type of each sensor is shown, with most irradiated using 27 MeV protons in Birmingham, and the rest irradiated using neutrons at JSI in Ljubljana. The specification requirements are highlighted, with all measured sensors passing the requirements.

Chapter 6

Missing mass search in proton-tagged dilepton events

This chapter presents the first ATLAS collaboration search for BSM physics using the missing mass method enabled by forward proton data measured by the AFP spectrometer. The search is performed in the channel $pp \to p\,(\gamma\gamma \to \ell\ell + X)\,p$, in which both protons remain intact, and formed the majority of the author's work during their PhD studies. The data used for the search comprise a total integrated luminosity of 14.7 fb^{-1} collected during pp collisions in 2017 at a beam energy of $\sqrt{s} = 13$ TeV. The missing mass method is explained in Section 6.2, then the real and simulated datasets used in the analysis are described in Sections 6.3 and 6.4, respectively, including three different signal models, with the signal region defined in Section 6.5. The process of modelling the background for this process is presented in Section 6.6, and all considered sources of systematic uncertainty are detailed in Section 6.7. Finally, Section 6.8 describes the statistical procedure employed and Section 6.9 presents the resulting upper limits set on the cross sections of the considered signal models, with large improvements observed over a similar CMS analysis [185] across most of the search space.

6.1 Introduction

The SM does a good job of explaining the majority of particle physics observations, and how these affect the Universe at a fundamental level. However, there are several areas where this model falls short, as discussed in Section 2.2. This leads particle physicists to search for signals which contradict the current SM by giving either evidence of new particles or conflicting measurements of parameters. One largely unexplored area in which to search for new physics utilises forward detectors, such as the AFP spectrometer, and the missing mass method. As described in Chapter 4, the AFP spectrometer is able to measure protons which undergo interactions in the ATLAS detector causing them to lose energy remain intact, such that they scatter out of the LHC beam into a region where they can be detected. In events where both interacting protons are measured by the AFP spectrometer, the total energy available to the central interaction and any associated particle production can be determined. The missing mass method, described in Section 6.2, combines this with measurements of all visible central particles detected by the ATLAS detector, which can be subtracted from the total interaction energy, to determine the total energy and corresponding invariant mass produced in the central interaction which is not detected by the ATLAS detector. This is more sophisticated than the missing transverse energy already measured by experiments such as ATLAS [186–188] and CMS [189–191] (explained in Section 3.3.4), as it includes the longitudinal component. This missing mass could potentially originate from BSM particles produced in the interaction which are invisible to detectors, such as DM candidates.

This chapter details an analysis searching for new physics using the missing mass method with data measured using the ATLAS detector in combination with the AFP spectrometer, in the channel $pp \to p\,(\gamma\gamma \to \ell\ell + X)\,p$. This process, illustrated in Figure 6.1, involves double photon emission between two protons, which remain intact, resulting in the central production of a visible boson V decaying leptonically and some undetected event component X. The missing mass method is used to reconstruct the kinematic properties of X, which is free to be any undetected particle,

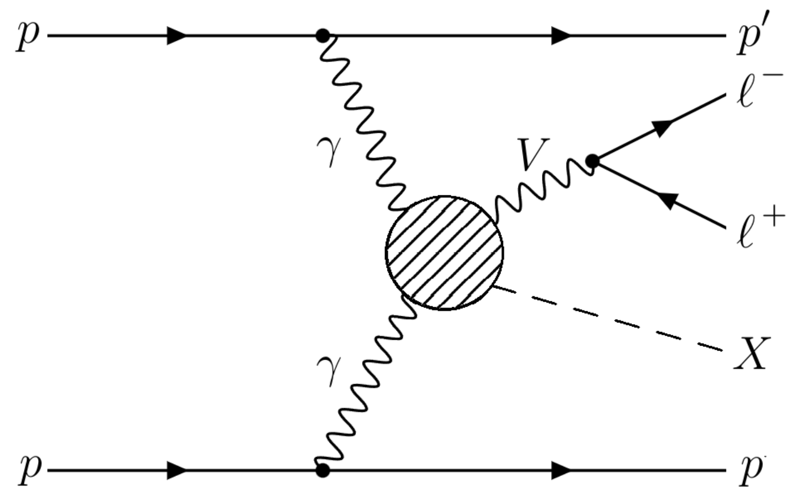

Figure 6.1: Feynman diagram of the signal process $pp \to p\,(\gamma\gamma \to \ell\ell + X)\,p$ considered in this analysis.

or multi-particle system. This could include heavy SUSY particles or Dark Matter (DM) candidates. The missing mass m_X is determined for all events passing the selection and a bump-hunt is performed in the resulting distribution, searching for any significant excess of data over the expected background, which could indicate new physics. Since exact knowledge of the properties of X is not required, this can function as a model-independent search.

Several previous analyses have utilised proton tagging to make measurements, from both the ATLAS [85] and CMS [192–194] collaborations, including several similar BSM searches [40, 195]. A similar analysis to the one presented in this thesis has been performed by the CMS-TOTEM collaboration [185], which considered three channels with V as a Z boson decaying to either electrons or muons or V as a photon. This analysis observed no significant excess over the background, instead setting upper limits on the hypothesised signal cross section. The analysis presented in this thesis is the first ATLAS collaboration analysis of this kind, utilising AFP spectrometer data with the missing mass method, and in the absence of any observation aimed to set improved limits relative to CMS. Only the electron and muon decay channels of the visible boson are considered, due to their clean signatures compared to the tau lepton decay channel.

6.2 The Missing Mass Method

The missing mass method combines proton energy loss measurements from the AFP spectrometer with central particle measurements from the ATLAS detector to reconstruct the four-momentum which is missing in the visible central production. This analysis applies this method to the signal process $pp \to p\,(\gamma\gamma \to \ell\ell + X)\,p$ shown in Figure 6.1. The AFP spectrometer is used to measure the energy loss of the intact signal protons ΔE_{p_A} and ΔE_{p_C}, where p_A and p_C refer to the protons detected on the A and C sides of the AFP spectrometer, respectively. In events where both signal protons are detected, the four-momentum of the interacting photon pair $\mathbf{p}_{\gamma\gamma}$ can be reconstructed, assuming that the p_{T} of the protons is negligible. The four-momentum of the dilepton system $\mathbf{p}_{\ell\ell}$ produced from the decay of V and measured in the ATLAS detector can then be subtracted to give the missing four-momentum in the event $\mathbf{p}_X$, from which the missing mass m_X can be obtained as

$$m_X^2 = (E_{\gamma\gamma} - E_{\ell\ell})^2 - (\overrightarrow{p_{\gamma\gamma}} - \overrightarrow{p_{\ell\ell}})^2 \tag{6.1}$$

$$= \left[\begin{pmatrix} \Delta E_{p_A} + \Delta E_{p_C} \\ 0 \\ 0 \\ \Delta E_{p_A} - \Delta E_{p_C} \end{pmatrix} - \begin{pmatrix} E_{\ell\ell} \\ p_x^{\ell\ell} \\ p_y^{\ell\ell} \\ p_z^{\ell\ell} \end{pmatrix}\right]^2 . \tag{6.2}$$

The missing mass m_X is then used as the observable for this analysis. The proton energy loss ΔE_p for a given proton is related to the fractional energy loss ξ as

$$\Delta E_p = \xi E_p, \tag{6.3}$$

where E_p is the beam proton energy, equal to 6.5 TeV in the data considered in this analysis.

6.2.1 Missing Mass Resolution

The resolution of the missing mass method using the AFP spectrometer was investigated in order to determine the optimal binning which could be used for the final distribution[1]. Several initial studies were carried out before the fully simulated signal MC samples described in Section 6.4 were available.

An analytical estimate was performed by propagating the estimated proton ξ resolution of the AFP spectrometer (10%), expected to be the dominant source of uncertainty in the missing mass calculation, through Equation 6.1, resulting in an estimated fractional resolution of 7%. A further study to include the contribution from the uncertainties in lepton measurement was performed using simulated $ZZ \to \ell\ell\ell\ell$ samples, by randomly selecting one of the Z bosons to be the missing mass, and reconstructing it from the other event components using the missing mass method. Generator-level proton information was used, as no AFP simulation was available at this time, so the proton reconstruction uncertainty was not included. However, this yielded extremely poor reconstruction, with a fractional resolution approaching 100%. This implied that the missing mass method is not effective for small masses on the order of the Z mass. To test higher missing masses, the two highest p_{T} leptons were instead selected, regardless of which Z decay they originated from, and these were collectively treated as the missing mass and reconstructed with the missing mass method from the other two leptons in conjunction with the corresponding protons. This yielded more promising results, suggesting a resolution between 5-10% for measurements of missing mass above around 400 GeV, with the fractional resolution increasing significantly below this threshold.

A more accurate measurement of the missing mass resolution obtained from fully simulated signal samples is presented in Section 6.4.

[1] In the end a single-binned approach was adopted for the final fits due to low statistics, as discussed in Section 6.9

6.2.2 Missing Mass Acceptance

The acceptance range in the missing mass reconstructed in this analysis depends mainly on the acceptance of the AFP spectrometer and the corresponding proton ξ cut applied in the selection. The maximum measurable missing mass for a given upper limit on ξ can be estimated from Equations 6.1-6.3 by considering the “best case scenario” in which $\xi_A = \xi_C$, allowing the lowest possible maximum ξ value for a given missing mass. Taking an approximated linear relation between the dilepton T and the proton energy loss, estimated from data, the relation shown in Figure 6.2 is obtained between the upper ξ limit and the maximum measurable m_X. Both the

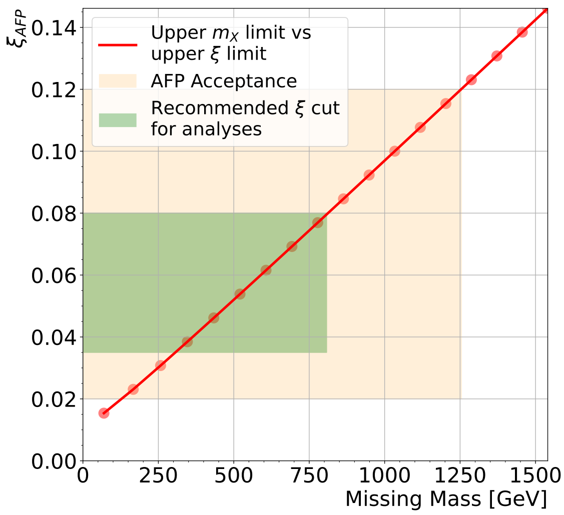

Figure 6.2: Missing mass acceptance region of the AFP spectrometer, where the red line is a rough estimate of the maximum missing mass value which can be obtained as a function of the highest proton ξ measured in an event, allowing the upper limit on m_X acceptance of the AFP spectrometer to be extrapolated from its ξ acceptance. The corresponding limit for the tightened ξ range considered in this analysis is additionally shown.

broadest ξ acceptance of the AFP spectrometer and the tighter ξ cut applied in this analysis (explained in Section 6.5.5) are overlaid. This suggests that events with $m_X > 1250$ GeV cannot be measured by the AFP spectrometer at all, as at least one proton will fall outside of the detector acceptance. This limit drops further for the tightened ξ cut to around 800 GeV. Up to this limit, the reconstruction efficiency is expected to drop for increasing m_X, as the probability of at least one proton being outside of the selection increases.

6.2.3 Event Mixing

The primary background in analyses combining data from the ATLAS detector with data from the AFP spectrometer is a combinatorial background. This is produced when central products produced in SM processes which match those in the targeted final state, a pair of leptons in this analysis, are combined in reconstruction with unassociated AFP protons originating in independent pile-up interactions. This background is described in more detail in Section 6.6, and it can be modelled using a data-driven method called event mixing. This is an established procedure, used to model this background in several other analyses involving AFP data [40, 85]. The process begins by selecting an event-shift, denoted i, between 1 and $N-1$, where N is the number of events in the data sample. Next, the central and forward proton components of each event in data are separated, and for a given value of i the central component of each event is shifted and combined with the proton data from the event i positions along in the dataset. This procedure is illustrated in Figure 6.3 for an event-shift $i = 2$, where the central information from the first event (1) is combined with the proton information from the third event ($1 + i = 3$), and so on. The process is fully circular, looping back to the beginning of the

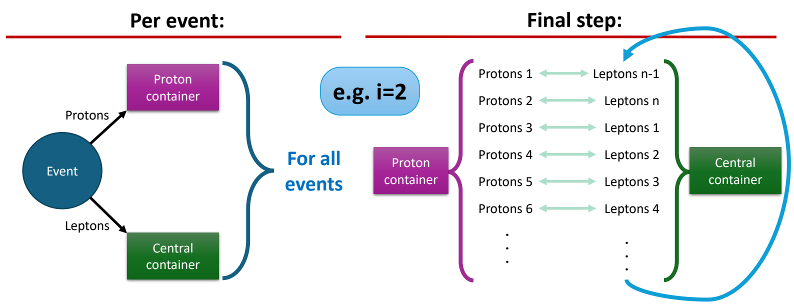

Figure 6.3: Event mixing procedure used to produce the data-driven model of the combinatorial background for this analysis, with an example event shift of $i = 2$.

dataset such that the proton information from the ith event is combined with the central information from the Nth event. For any event-shift $i \geqslant 1$, the central and proton data which are combined to reconstruct the final state are uncorrelated, originating from different events and thus different processes. This exactly simulates the combinatorial background in which the two components originate from different processes, irrespective of the dataset composition in terms of signal and background. A sample with $i = 0$ corresponds to the unaltered data sample for the analysis, in which the central and proton data are correlated.

6.3 Data Overview

The analysis uses pp collision data collected at a centre-of-mass energy of $\sqrt{s}$ = 13 TeV during Run 2 of the LHC, with the ATLAS detector measuring centrally produced particles and the AFP spectrometer measuring forward protons. The analysis is limited to only data collected during 2017, as in other years the AFP spectrometer was either inoperational or had a critical data-taking issue, as described in Section 4.5. Both the ATLAS and nominal AFP GRLs, described in Section 4.5, are applied to the dataset, giving an available integrated luminosity of 14.7 fb^{-1}.

The analysis uses data which have been skimmed according to the STDM7 derivation, which aims to select dilepton events and includes AFP information. The skimming requires that events fire any of the unprescaled signal and double lepton triggers in the ATLAS experiment menu, and contain at least two electrons or muons which pass several loose p_{T}, η and quality cuts which are looser than baseline selections in the analysis. The triggers applied in the analysis are the unprescaled single and double lepton triggers with the loosest available p_{T} thresholds. These require either:

- A single, loosely isolated muon with $p_{\mathrm{T}} \geqslant 26$ GeV
- A pair of muons with $p_{\mathrm{T}} \geqslant 14$ GeV
- A single, loosely isolated electron with $p_{\mathrm{T}} \geqslant 26$ GeV
- A pair of electrons with $p_{\mathrm{T}} \geqslant 17$ GeV

6.3.1 Blinded Data

During the analysis the data were blinded, which involves modifying the data in some way to hide the true observable, preventing any features present in the data early in the analysis process from biasing analysers during optimisation of event selections. This analysis used a novel blinding approach applying the event-mixing procedure discussed in Section 6.2.3 to produce a single event-mixed sample with an event-shift of $i = 1$ which was used as the blinded dataset. As a reminder, this corresponds to a dataset in which the central and proton components are taken from two different events, separated by one position in the unaltered dataset. This ensures that the central and proton components of the data originate from different unrelated events, so once they are combined they do not contain information on the true missing mass distribution in the dataset. This was beneficial compared to more standard blinding approaches as it allowed the signal region in other key variables to be observed before unblinding, and eliminated the need for any control regions in which one or more signal selections are flipped, simplifying the analysis.

6.4 Signal Simulation

The analysis considers three different signal models, all with two intact protons, a visible boson V decaying to two oppositely charged leptons and an undetected particle X in the final state. The first two models, $Z + X$ and $Z + H'$, are generic models in which V is set to be a Z boson. The third targets a specific BSM scenario with ALPs, in which V is set to be a short-lived ALP decaying into leptons. Since two intact protons are required in the final state, only the elastic production mode is considered, with the single and double-dissociative production modes treated as part of the combinatorial background.

6.4.1 Signal Models

The first simple model is generated using SuperChic [96], and directly simulates the elastic photon-induced production of an invisible particle X along with a Z boson, which decays into electrons or muons. The final state is produced via a four-point interaction, as illustrated in Figure 6.4. This model was used to reproduce the

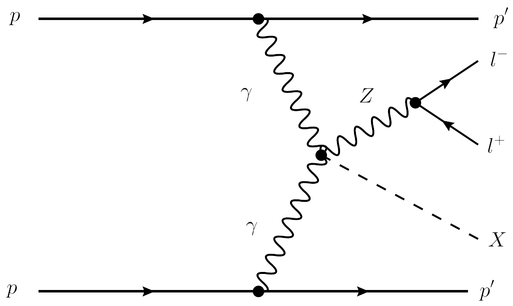

Figure 6.4: Feynman diagram for $Z + X$ production via a four-point photon interaction.

model used by CMS in their equivalent analysis [185], and therefore uses a matching parametrisation, with the ZX invariant mass generated with a probability proportional to $e^{-\tau \cdot m_{ZX}}$, where τ is an arbitrary model parameter. A value of $\tau = 0.04$ GeV^{-1} is chosen, again to match with CMS, although the value of this parameter was found to have a negligible impact on the final kinematic distributions. Several signal samples were produced varying the decay channel of the Z between electrons and muons, and varying the missing mass m_X between 300 and 900 GeV, in 100 GeV intervals. The generation is performed in SuperChic 5.1, with parton showering and hadronisation simulated in Pythia 8.310 [91]. As discussed in Section 2.4, SuperChic is one of very few MC generators which can simulate soft-survival effects in photon-induced processes, which are described in Section 2.3.4 and encode the probability $1 - S$ of further interactions occurring between the protons, causing additional particle production and removing the rapidity gap. However, since the $Z + X$ model is implemented in SuperChic as a simplified model, this effect is not included, instead having a constant $S = 1$ across all events. Therefore, the effect of soft-survival is instead estimated using a similar simulated sample of exclusive dilepton production $\gamma\gamma \to \ell\ell$ generated in SuperChic, which does include this effect.

Figure 6.5 shows the effect on event yield in this process of turning the soft-survival effect on and off in SuperChic generation, as a function of the mass of the central system $m_{\ell\ell}$. This is equivalent to m_{ZX} in the current analysis, as both are equal to the diphoton mass $m_{\gamma\gamma}$. The corresponding ratio was parametrised as a first-order

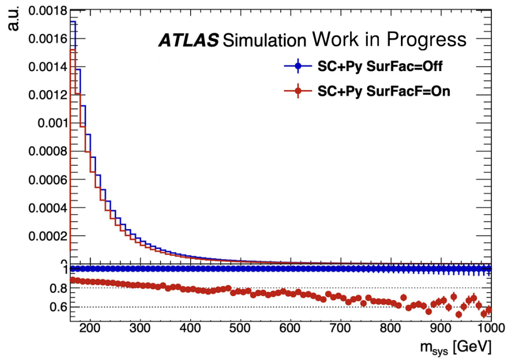

Figure 6.5: Effect of turning soft-survival effects on and off in a SuperChic simulated sample of exclusive dilepton production $\gamma\gamma \to \ell\ell$, as a function of the mass of the central system $m_{\ell\ell}$.

polynomial to limit dependence on low-statistics regions, and the resulting fitted ratio $(0.9387 - 0.000365 m_{ZX})$ was then applied as an event-by-event weight to the raw $Z + X$ signal sample as a function of the generator-level mass of the central system m_{ZX}. This effectively simulated the effect of non-unity soft-survival in this signal sample. The soft-survival factor in SuperChic additionally has a weak dependence on the rapidity of the central system y_{ZX}, as shown in Figure 6.6a. However, the

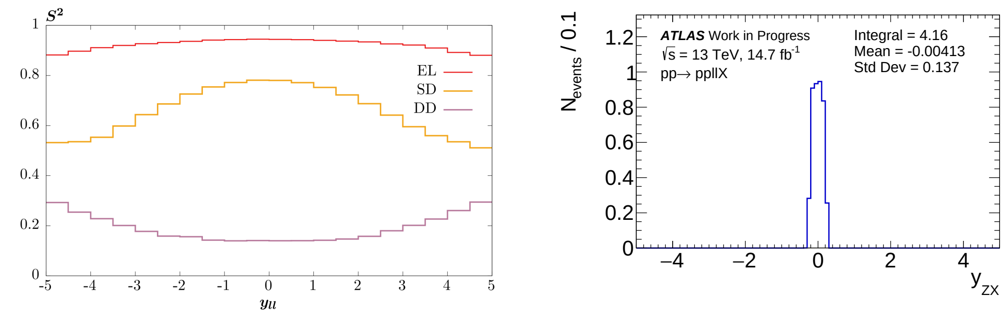

(a) (b)

Figure 6.6: (a) The simulated dependence on the soft-survival probability for exclusive dilepton production as a function of the central rapidity $y_{\ell\ell}$ [80], equivalent to y_{ZX} in the current analysis, for EE events as used here as well as SD and DD events. (b) The rapidity distribution for the SuperChic $Z + X$ signal model with $m_X = 700$ GeV in the muon channel.

rapidity observed in the simulated $Z+X$ samples, shown in Figure 6.6b for a representative signal mass, is consistently close to zero, due to the tight proton ξ cut applied in the signal selection which forces events to be highly symmetric in their energy distribution. The corresponding region in Figure 6.6a shows a flat distribution for low values of $|y_{\ell\ell}|$ indicating a negligible dependence on rapidity. Therefore, the effect of the rapidity of the central system on the soft-survival factor is not considered further.

The other generic model considered is generated using MadGraph [93], and uses the SM process $\gamma\gamma \to ZH$ as a basis, with the Higgs boson replaced by a generic scalar referred to as H', which is assigned a tunable mass and forced not to decay. The Z boson decays leptonically, and all couplings are kept at their default SM values for simplicity. This process can occur in the SM (if H' is replaced by an SM Higgs boson) via loop-induced diagrams, two of which are shown in Figure 6.7. Several

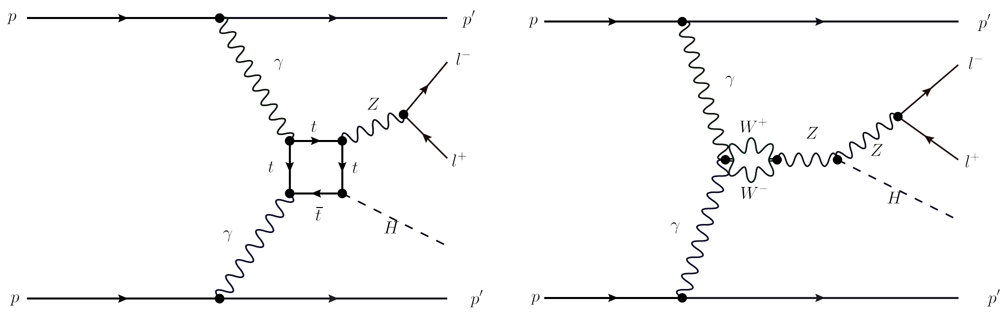

(a) (b)

Figure 6.7: Representative loop-induced Feynman diagrams contributing to $Z+H'$ production.

samples were produced, varying the missing mass between $100 \leqslant m_X \leqslant 1000$ GeV in 100 GeV steps. The generation is performed in MadGraph 2.9.5 and parton showering and hadronisation is performed in Pythia 8.306.

The final model considered in this analysis probes an existing BSM theory discussed in Section 2.2 in the form of ALPs, hypothetical pseudoscalar particles considered to be potential candidates for dark matter, as discussed in Section 2.2. The so-called "di-ALP" model simulates the photon-induced production of two distinct, electrically neutral ALPs, a short-lived particle S_1, decaying leptonically into either e^+e^- or $\mu^+\mu^-$ with equal probability, and a long-lived particle S_2, which is invisible to the ATLAS detector. The process is shown in Figure 6.8. There are several

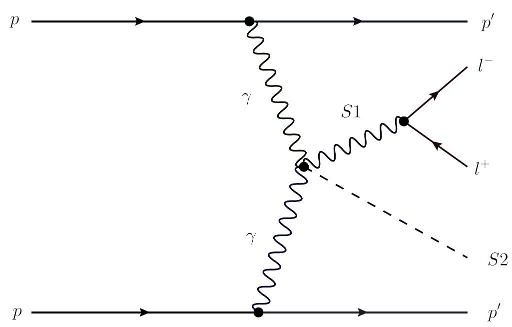

Figure 6.8: Feynman diagram for di-ALP production via photon fusion.

free parameters in this model: the mass of the both ALPs and their couplings to the photons. The mass of the short-lived ALP S_1 was set to the Z mass to make a clearer comparison to the other models, and both couplings were set to 1 for simplicity, as these are not expected to affect the kinematics of the final state, only the overall signal strength. Samples were produced varying the missing mass m_X (i.e. that of S_2) between 200 GeV and 900 GeV in 100 GeV intervals. This model is implemented using the FeynRules package [196, 197], which is then passed to MadGraph 2.9.5 for event generation, with parton showering and hadronisation simulated with Pythia8 8.306.

The MadGraph generator does not implement soft-survival effects, and therefore these are not considered for either of the MadGraph models, instead this effect is considered as a systematic for these samples, as discussed in Section 6.7.

Table 6.1 gives an overview of all simulated signal samples produced for this analysis. Also shown is the filter which was applied at generator level to the forward protons in each event. This was used to increase the generation efficiency for events passing the event selection considered by the analysis, to reduce the total number of generated events required for a given sample size within the signal region. The filter places requirements on proton fractional energy loss $\xi < 0.15$ and proton $p_{\mathrm{T}} \leqslant 1.5$ GeV, with particularly the ξ cut removing a large proportion of potential signal events

Process	Generator	UEPS	Filter	m_X Range [GeV]
$\gamma\gamma \to Z\,(ee) + X$	SuperChic v5.1	Pythia 8.310	$\xi < 0.15$, $p_{\mathrm{T}}^{\mathrm{proton}} < 1.5$ GeV	$300 \leqslant m_X \leqslant 1000$
$\gamma\gamma \to Z\,(\mu\mu) + X$	SuperChic v5.1	Pythia 8.310	$\xi < 0.15$, $p_{\mathrm{T}}^{\mathrm{proton}} < 1.5$ GeV	$300 \leqslant m_X \leqslant 1000$
$\gamma\gamma \to S1\,(\ell\ell) + S2$	MadGraph v2.9.5	Pythia 8.306	$\xi < 0.15$, $p_{\mathrm{T}}^{\mathrm{proton}} < 1.5$ GeV	$200 \leqslant m_X \leqslant 900$
$\gamma\gamma \to Z\,(\ell\ell) + H$	MadGraph v2.9.5	Pythia 8.306	$\xi < 0.15$, $p_{\mathrm{T}}^{\mathrm{proton}} < 1.5$ GeV	$100 \leqslant m_X \leqslant 1000$

Table 6.1: Overview of the simulated signal samples considered in the analysis, the corresponding programs used to perform the generation and Underlying Event and Parton Shower (UEPS) simulation steps, applied generator-level filters and the generated ranges of hypothesised signal masses.

which would be impossible to detect with the AFP spectrometer. A corresponding weight is calculated for each value of m_X corresponding to the efficiency to retain events following the application of this generator-level filter, to be applied to any resulting signal cross section measurements.

6.4.2 Kinematic Comparison

Figure 6.9 shows a comparison between the three signal models across several kinematic distributions, for a signal mass of $m_X = 500$ GeV. This demonstrates the difference in final state kinematics between the models, with the MadGraph models having a significantly higher p_{T} dilepton system, and correspondingly higher proton energy loss measured in the AFP spectrometer. However, it can be seen that these two effects balance out in their influence on m_X to give very similar missing mass distributions, with comparable widths.

The wide range of potential kinematics across the considered signal models enhances the sensitivity of the analysis to many potential sources of new physics, whose underlying kinematics could cover a similarly wide spectrum. This allows the analysis to be more model-independent, allowing a general search to be performed.

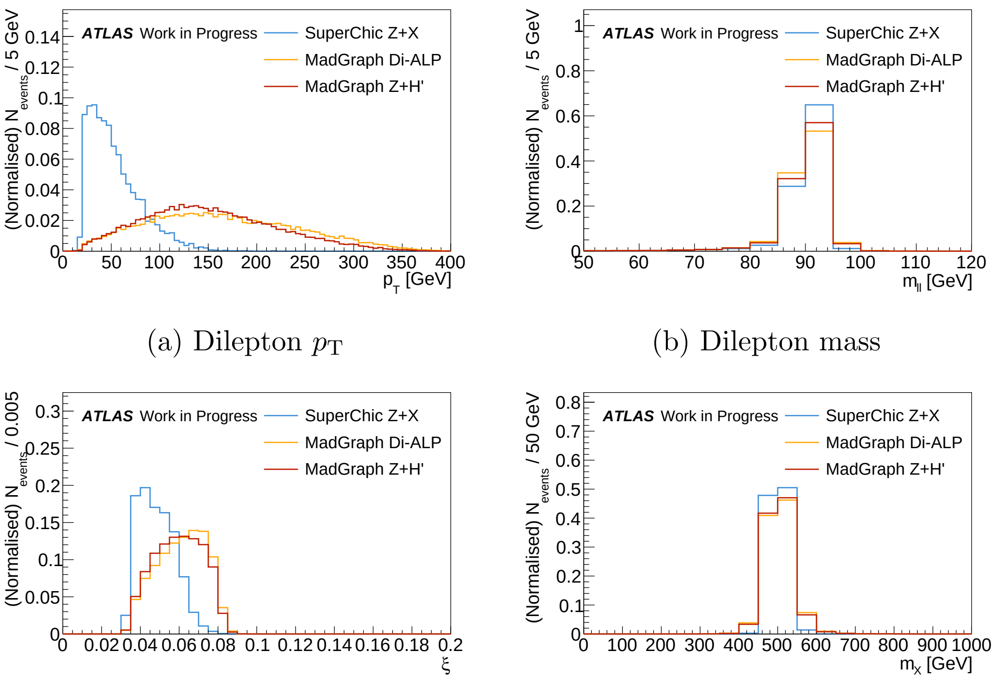

(a) Dilepton p_{T}

(b) Dilepton mass

(c) Proton ξ

(d) Missing mass m_X

Figure 6.9: Unit normalised comparison of generator-level kinematic distributions generated for each signal model with a hypothesised signal mass of $m_X = 500$ GeV. For the MadGraph di-ALP model, m_{S_1} is set to the Z boson mass.

6.4.3 Simulated Beamspot Size

The beamspot, as described in Section 3.3.1.1, is the volume around the ATLAS detector IP where the two LHC beams cross over, allowing pp interactions to occur. The simulated signal samples use a constant longitudinal beamspot size of 35 mm, which is very close to the mean beamspot size observed during 2017 data-taking of 36.94 mm. However, as shown in Figure 6.10, the beamspot size in data varied significantly between around 30-45 mm across the year. This affects the simulated track density around the IP, with an increased beamspot size giving reduced track density, which in turn affects the signal efficiency of the track veto selection described in Section 6.5. This can introduce a small deviation in the simulated track veto signal efficiency compared to estimates made from data, although this is expected to be negligible compared to other sources of uncertainty such as the modelling of pile-up interactions on which the simulated track veto signal efficiency is based. These

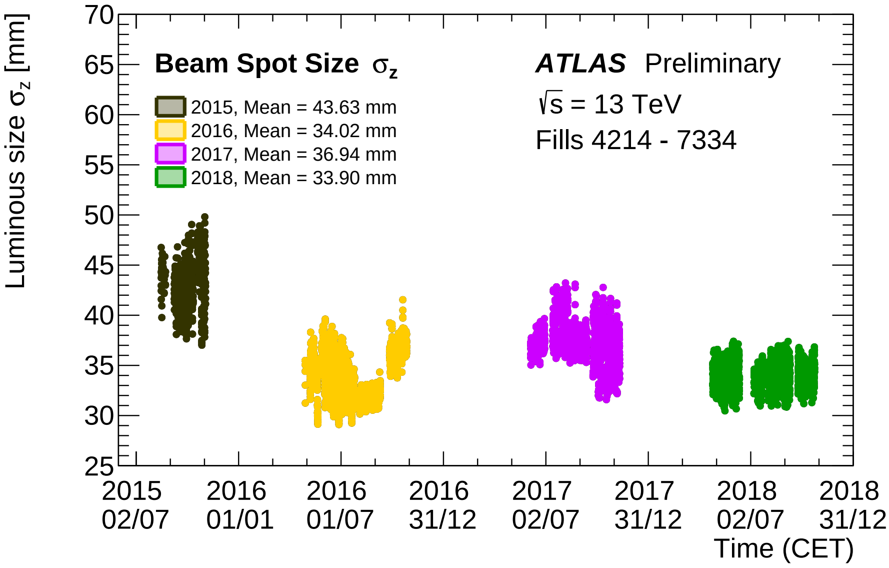

Figure 6.10: The size of the luminous region in the ATLAS detector during $\sqrt{s} = 13$ TeV pp collisions in Run 2. The data points are the hourly average of results of a maximum likelihood fit to the spatial distribution of primary vertices collected over a two minute period. The luminosity weighted average size is provided for each year.

effects are all covered by a single systematic uncertainty, which is discussed in detail in Section 6.7.1.

6.4.4 Updated Missing Mass Resolution

Using the fully simulated samples, an updated measurement of the resolution of the reconstructed missing mass was determined for each signal model using Equation 6.1. The resulting distribution was fitted to a Gaussian function using very fine binning to extract the width σ_{m_X} of the distribution. Figure 6.11 shows the width as a function of the hypothesised signal mass for each generated mass in each signal model. These plots show a consistent width of at most 50 GeV for almost all models, in both considered lepton channels, and so this was chosen as the optimal bin width to use for the missing mass distribution in the final fits. As will be shown in Section 6.5.6 (see Figure 6.17) the exceptional point $m_X = 1000$ GeV for the SuperChic $Z + X$ model occurs due to a very low fiducial selection efficiency. This mass point is not considered in the final result and so this does not affect

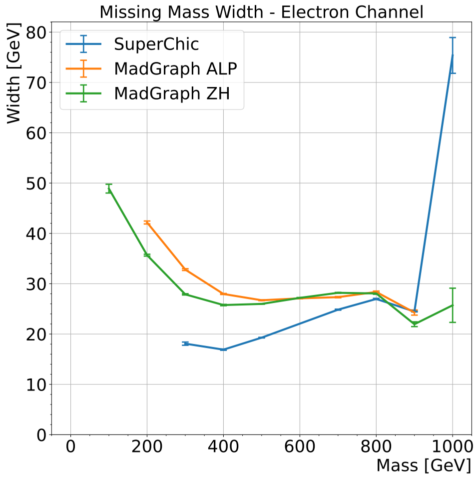

(a) Electron channel

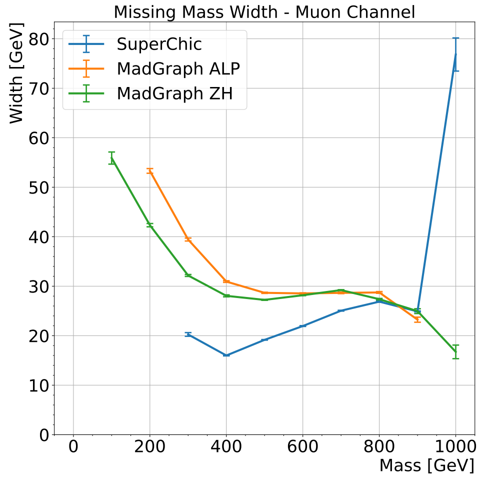

(b) Muon channel

Figure 6.11: Width of the reconstructed missing mass distribution in simulated signal samples as a function of hypothesised signal mass, for each signal model and in each lepton channel.

the chosen binning. This result is in reasonable agreement with the preliminary estimates presented in Section 6.2.1, although the resolution determined here is actually better than previously suggested by those studies.

6.5 Event Selection

Reconstructed events are required to pass several levels of selection in order to be considered for the potential presence of signal. First, basic event-cleaning is performed to remove invalid or corrupted events, before a loose pre-selection is applied to select only dilepton events matching with the expected signal final state. This dataset uses looser cuts than the final signal selection to provide a large dataset on which to perform studies during the analysis, and to construct the data-driven background model discussed in Section 6.6. Finally, a tight signal selection is applied to remove background and increase the sensitivity of the analysis.

The initial dataset is obtained by applying the ATLAS and AFP GRLs to the 2017 dataset, as described in Section 6.3, leaving a total integrated luminosity of 14.7 fb^{-1}. Several event-level vetoes are made to reject bad or corrupted events due to issues in the LAr, Tile and SCT detectors, in addition to incomplete events missing certain elements of reconstruction. Checks are performed to remove duplicate events, which can occur due to glitches during dataset creation, by ensuring that all events have a unique pairing of run number and event number within that run. Finally, events are required to have at least one reconstructed Primary Vertex (PV) with at least two associated nominal ID tracks with $p_{\mathrm{T}} > 500$ MeV, as described in Section 3.3.1.1.

6.5.1 Lepton Selection

Electrons are reconstructed using the procedure described in Section 3.3.2. Candidate electrons are required to pass kinematic cuts of $p_{\mathrm{T}} > 18$ GeV and $|\eta| < 2.47$, with the p_{T} cut designed to be slightly tighter than the trigger requirement described in the previous section, to remove the effects of any trigger inefficiency. They must additionally satisfy the Loose identification and isolation working points, described in more detail in Section 3.3.2. Electrons undergo an energy calibration to optimise the response of the detectors, and electrons reconstructed using "bad" clusters, affected by the presence of dead readout channels or masked cells in the LAr calorimeter, are removed. To ensure that the electron tracks are close to the primary vertex, track-to-vertex association requirements are placed on the impact parameters measured by the IBL of $\left|d_0^{\mathrm{BL}}\right| / \sigma\left(d_0\right) < 5$ and $\left|\Delta z_0^{\mathrm{BL}} \sin\theta\right| < 0.5$ mm. The track parameters d_0, z_0 and θ are defined in Section 3.3.1.

Muons are reconstructed following the procedure given in Section 3.3.3. Pre-selected muons must pass kinematic cuts of $p_{\mathrm{T}} > 15$ GeV and $|\eta| < 2.4$, where again the p_{T} cut is chosen to be slightly tighter than the trigger requirement following the recommendations. Muon candidates must satisfy the Loose selection and isolation working points described in Section 3.3.3, and pass track-to-vertex selections of $\left|d_0\right| / \sigma\left(d_0\right) < 3$ and $\left|\Delta z_0 \sin\theta\right| < 0.5$ mm. Momentum calibration corrections are also applied to muons.

Overlap removal of pre-selected leptons is performed to select between multiple leptons sharing the same ID track. Electrons sharing a track with a muon are rejected, unless the muon is reconstructed as a Calorimeter-Tagged (CT) muon (explained in Section 3.3.3), in which case the muon is rejected. In cases where multiple electrons share the same track, only the highest p_{T} electron is kept. Overlap removal of leptons against jets is also applied, where the considered jets pass the following basic selections:

$$|\eta| < 2.4, \begin{cases} 30 < p_{\mathrm{T}} < 60 \text{ GeV, JVT} > 0.5 \\ p_{\mathrm{T}} > 60 \text{ GeV} \end{cases}$$

$$2.4 < |\eta| < 4.9,\ p_{\mathrm{T}} > 30 \text{ GeV}$$

where the Jet-Vertex Tagger (JVT) is a multivariate discriminator for jets defined in [198]. Leptons are rejected against these jets if they are within $\Delta R < 0.4$, where $\Delta R = \sqrt{\Delta\eta^2 + \Delta\phi^2}$. However, it was found that this overlap removal is negligible when the exclusivity requirement is imposed via the track veto, which removes most events containing jets.

6.5.2 Forward Proton Selection

Forward proton reconstruction is performed using AFP data as described in Section 4.3. This analysis uses the Medium quality working point described in the same section, requiring reconstructed AFP tracks to have hits in at least two different SiT planes. Pre-selected protons are required to have fractional energy loss in the range of AFP acceptance at $0.02 < \xi < 0.12$, referred to hereafter as “loose” protons, with “tight” protons referring to the tightened requirement of $0.035 < \xi < 0.08$ used in the final signal selection. Protons reconstructed using either single or double-station reconstruction are accepted for the pre-selection, to maximise statistics for this sample.

Feature	Electron/Muon Event Criterion
p_{T}	$> 18/15$ GeV
$\lvert\eta\rvert$	$< 2.47/2.4$
Identification	Loose
Isolation	Loose
$\left\lvert d_0^{\mathrm{BL}}\right\rvert / \sigma\left(d_0\right)$	$< 5/3$
$\left\lvert \Delta z_0^{\mathrm{BL}} \sin\theta \right\rvert$	< 0.5
Proton quality	Medium
Proton ξ	$0.02 < \xi < 0.12$
Proton reconstruction	Single or double-station

Table 6.2: Summary of object pre-selection for candidate events in the electron and muon channels.

6.5.3 Track Selection

Inner detector tracks in addition to the leptons are selected in this analysis for use in the track veto, described in Section 6.5.6.1. Tracks reconstructed using the standard ATLAS experiment procedure described in Section 3.3.1 are required to have $p_{\mathrm{T}} > 500$ MeV and $\lvert\eta\rvert < 2.5$, and must satisfy the Loose quality working point described in the same section. Overlap removal with the tracks comprising the candidate leptons is performed by using direct ID matching for muon tracks and a $\Delta R < 0.01$ cut for electron tracks. An additional cut of $\lvert d_0\rvert < 0.5$ mm is imposed to reduce the rate of fake tracks, increasing the performance of the track veto.

6.5.4 Pre-Selection

The event pre-selection for candidate leptons and protons is summarised in Table 6.2. Pre-selected events are required to have at least two candidate leptons of the same flavour and opposite charge (e^+e^- or $\mu^+\mu^-$), and either exactly one loose proton per side of the AFP spectrometer or exactly one tight proton per side (see next paragraph). Figure 6.12 illustrates the different scenarios for the number of protons detected on each side of the AFP spectrometer, and the resulting necessity for requiring exactly one proton per side. In an event where at least one side of

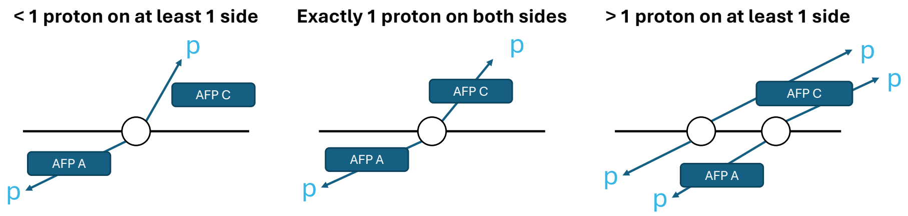

Figure 6.12: Examples of the different possible scenarios for a given event of how many protons are detected on each side of the AFP spectrometer.

the AFP spectrometer has no proton reconstructed (left in Figure 6.12) the missing mass method cannot be applied, since the total energy available to the central interaction cannot be determined without measuring both signal protons. However, frequently at least one side of the AFP spectrometer will have multiple protons reconstructed (right in Figure 6.12) due to multiple diffractive interactions producing pile-up protons, which creates the dominant combinatorial background for this analysis. In such events, it is not possible to distinguish which of the reconstructed protons, if any, originated in the signal interaction, as kinematic matching (as is possible for exclusive dilepton production with Equation 4.2) cannot be performed with this central state due to the presence of missing mass. Therefore, the only state which can be accepted is one where exactly one proton is detected on both sides of the AFP spectrometer, allowing the missing mass to be reconstructed, with the assumption that both measured protons originate in the signal interaction. In fact, this is not always the case, as explored in more detail in Section 6.6.2.

This requirement of exactly one proton per side complicates the proton component of the event pre-selection, as the signal region requires a tighter proton ξ selection of $0.035 < \xi < 0.08$, as discussed in the next section. If the pre-selection simply required exactly one loose proton per side, then the signal region requiring exactly one tight proton per side would not be fully contained in the corresponding sample. An example event which passes the signal selection, but would not pass this version of the pre-selection due to the presence of an additional loose pile-up proton, is illustrated in Figure 6.13. Therefore, to avoid missing such events in the pre-selection, the requirement is extended to allow for either exactly one loose proton per side of the AFP spectrometer or exactly one tight proton per side.

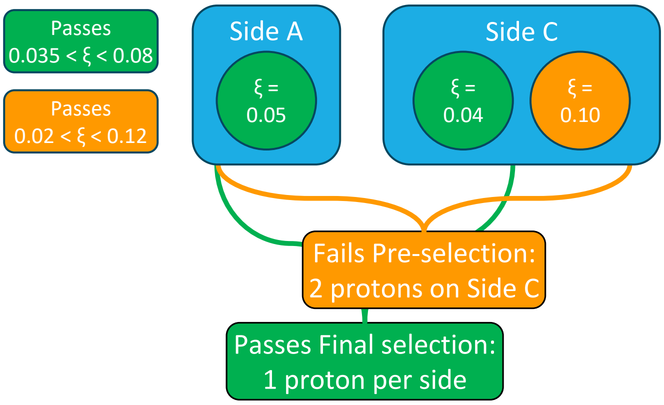

Figure 6.13: Example event with more than one loose proton per side, but exactly one tight proton per side, which passes the signal selection but fails a pre-selection requiring exactly one loose proton per side.

6.5.5 Signal Selection

The final signal selection, which is applied to pre-selected events to separate potential signal events from the majority of the background, is summarised in Table 6.3. The identification working points for electrons and muons are tightened to LooseAndBLayer and Medium, respectively, to increase the quality of the leptons considered for signal events. Optimisation studies were performed into the different available lepton selection and isolations working points, and no significant effect was observed on the final results of the analysis, so these were chosen as a balance between quality and higher statistics. The dilepton system is then selected from the highest p_{T} lepton in the event and the corresponding particle of the same flavour and opposite charge with the next highest p_{T}, determining the lepton channel to which the event contributes.

To remove quarkonium resonances, arising from the leptonic decays of mesons such as J/ψ $(c\bar{c})$ and Υ $(b\bar{b})$ and observed particularly in the muon channel, a cut on the total invariant mass of the dilepton system is imposed at 50 GeV. Despite all the considered signals having the mass of the visible boson at the value of the Z boson mass, and the corresponding distribution in data being dominated by the Z boson resonance, a specific cut targeting this was not employed in order to remain independent of this particular model component (allowing, for example, variation of

Feature	Signal Region Criterion
Electrons	
Kinematic	$p_{\mathrm{T}} > 18$ GeV
Identification	LooseAndBLayer
Muons	
Kinematic	$p_{\mathrm{T}} > 15$ GeV
Identification	Medium
Dilepton	
Charge	$q_{\ell_1} + q_{\ell_2} = 0$
Kinematic	$p_{\mathrm{T}} > 20$ GeV
Mass	$m_{\ell\ell} > 50$ GeV
Track number	$N^{0.5\ \mathrm{mm}}_{p_{\mathrm{T}}>500\ \mathrm{MeV}} = 0$
Protons	
Kinematic	$0.035 < \xi < 0.08$ NEAR station track $x < -3.5$ mm
Number of stations	Double only

Table 6.3: Summary of the selection for signal events in each signal region, in addition to the pre-selection detailed in Table 6.2.

the S_2 mass in the di-ALP model).

In addition, to take advantage of the significant difference between the transverse momentum of the dilepton system observed for signal events in Figure 6.9 and for data where the visible boson is typically produced with minimal transverse momentum, a cut of $p_{\mathrm{T}}^{\ell\ell} > 20$ GeV is imposed. An optimisation study was performed into the value of this cut by determining the expected cross section upper limits for each signal model with dilepton p_{T} cuts of 0, 10, 20 and 30 GeV applied. The ratio of the obtained limits was then found between progressively increasing thresholds to determine the highest value where improved limits are consistently observed. The limit ratios are shown in Figure 6.14. Improvement in sensitivity is consistently

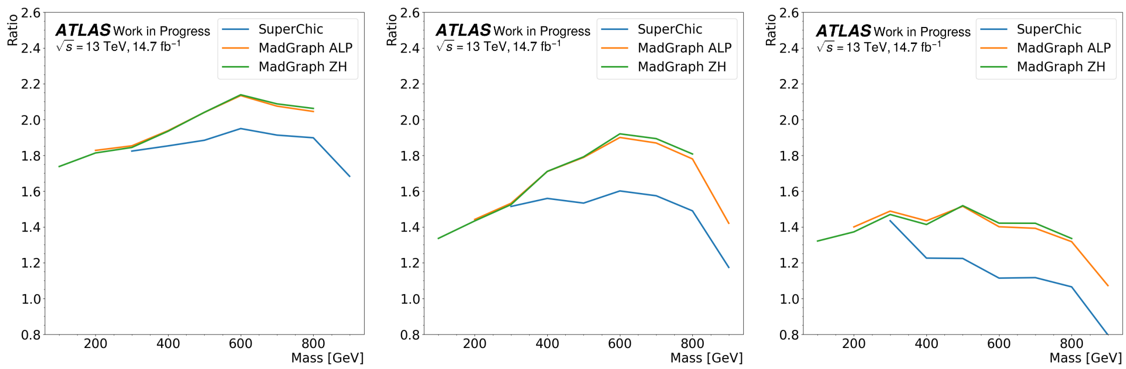

(a) Limits Ratio 0 GeV:10 GeV (b) Limits Ratio 10 GeV:20 GeV (c) Limits Ratio 20 GeV:30 GeV

Figure 6.14: Comparison of the ratio of the expected cross section upper limits obtained for each signal model between progressively increasing dilepton p_{T} cuts.

observed for progressively higher thresholds, until the threshold is increased from 20 to 30 GeV, at which point some models start to lose sensitivity due to extremely low statistics. Therefore, the value of 20 GeV was chosen as a balance between all considered signal models, resulting in an overall 2-4 times improvement in sensitivity compared to having no dilepton p_{T} cut. The final central signal criterion is the track veto, described in detail in Section 6.5.6.1.

For the proton component of the events, the signal selection accepts only protons reconstructed from tracks in both AFP spectrometer stations, and additionally requires that the track segment x coordinate in the NEAR detector satisfies $x < -3.5$

mm, to ensure high quality reconstruction. A tighter cut of $0.035 < \xi < 0.08$ is imposed on signal protons, to restrict the selection to the region in which the proton reconstruction efficiency is well understood, as explained in Section 4.2. In addition, this was observed to reduce the significance of the time dependence of the shape of the missing mass spectrum. Figure 6.15 shows the m_X distributions in data separately for each data-taking period throughout 2017, with each period corresponding to a subsection of events recorded within a relatively small time window, such that any differences in detector conditions should be minimised. Small differences are

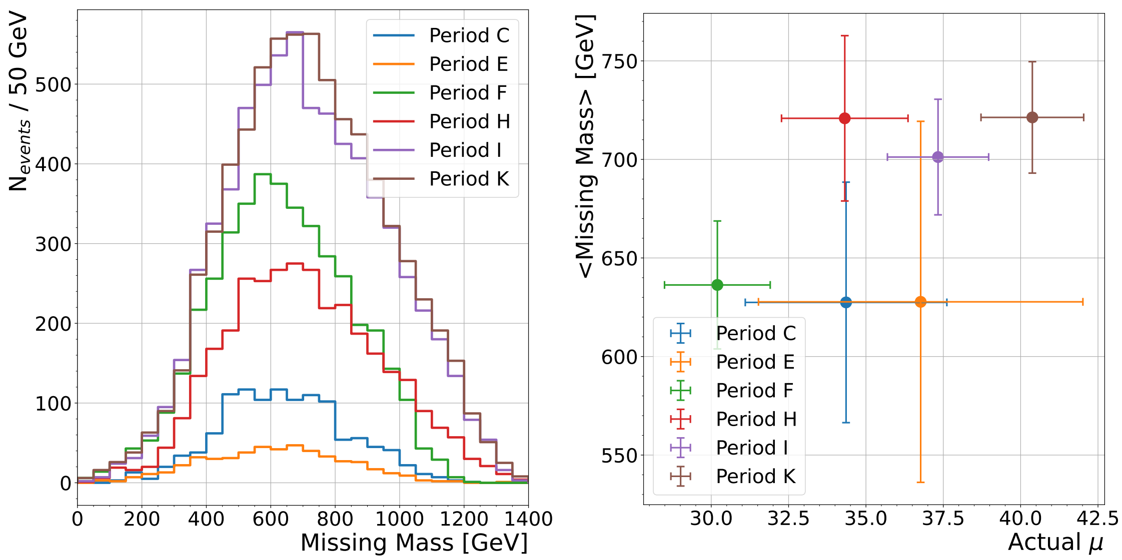

(a) Missing mass distributions (b) Mean m_X as a function of μ

Figure 6.15: (a) Missing mass distributions and (b) average missing mass (m_X) versus average interactions per bunch crossing (μ) separated by data-taking period, with the loose $0.02 < \xi < 0.12$ proton selection applied. The different normalisations in (a) are expected due to the total integrated luminosity differing between each data-taking period. The error bars in (b) correspond to errors on the mean values.

observed between the m_X distributions in each period, with slightly different widths and peak values. This is emphasised in Figure 6.15b, which plots the average missing mass in each period against the average interactions per bunch crossing μ within that period. There is some evidence here for two distinct sets of data: Periods C-F have a lower average $m_X \sim 630$ GeV, while Periods H-K have a higher average $m_X \sim 710$ GeV; with no obvious dependence on the level of pile-up. This transition suggests a change in conditions between periods F and H, and is likely to be explained by the change in the AFP spectrometer insertion settings which occurred twice during 2017 data-taking as detailed in Table 4.2. The closer approach of the detector to the beam changes the acceptance of the detector, allowing higher values of m_X to be reconstructed with higher efficiency. To remove this time dependence and keep the observable distribution consistent across the full dataset, the proton acceptance is restricted to the tightened range of $0.035 < \xi < 0.08$. As shown in Figure 6.16, this results in the differences in mean reconstructed m_X values between the different periods dropping to around 20 GeV, well below the resolution measured for the missing mass method in Section 6.2.1. This tightened cut reduces the upper limit of

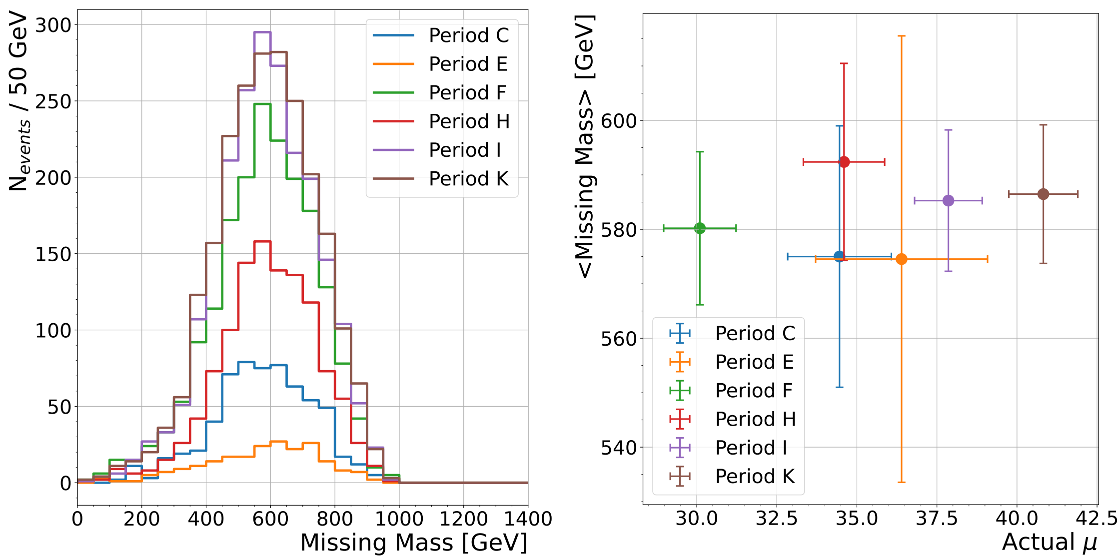

(a) Missing mass distributions (b) Mean m_X as a function of μ

Figure 6.16: (a) Missing mass distributions and (b) average missing mass (m_X) versus average interactions per bunch crossing (μ) separated by data-taking period, with the tightened $0.035 < \xi < 0.08$ proton selection applied.

reconstructed missing mass observed in data from around 1200 GeV to 1000 GeV, as expected from the study in Section 6.2.2.

6.5.6 Fiducial Selection

The analysis is performed with respect to a limited fiducial volume, which is a defined region of phase space corresponding to the sensitive regions of the detectors with well-understood efficiency, which yield high signal efficiency and low background. This volume is defined following the signal region event selections described in Section 6.5.5. This is required due to the wide kinematic spectrum of the ZX system in each of the considered signal models, causing many generated signal events to fall outside of detector acceptance, particularly the AFP spectrometer acceptance of $0.02 < \xi < 0.12$, as discussed in 6.4. Table 6.4 summarises the defined fiducial volume, with the corresponding selection applied to the generator-level kinematics of the final state signal particles. For the final fits shown in Section 6.9, simulated

Feature	Criterion
Electrons	Same flavour, opposite charge $p_{\mathrm{T}} > 18$ GeV $\lvert\eta\rvert < 2.47$
Muons	Same flavour, opposite charge $p_{\mathrm{T}} > 15$ GeV $\lvert\eta\rvert < 2.4$
Dilepton system	$m_{\ell\ell} > 50$ GeV $p_{\mathrm{T}}^{\ell\ell} > 20$ GeV
Protons	$0.035 < \xi < 0.08$

Table 6.4: Summary of fiducial volume selection criteria for signal events.

signal samples are normalised with respect to this fiducial volume, such that the resulting limits set on the cross section correspond to fiducial cross sections, only measured within this volume of phase space. Signal events falling outside this fiducial volume are considered as an additional background component, as discussed in detail in Section 6.6.

The generator-level efficiency of the fiducial selection for each signal model is shown in Figure 6.17. These efficiencies account for the efficiency of the generator-level forward filter described above. To verify that these do not introduce any selection bias, additional samples of the SuperChic $Z + X$ model were produced without the generator-level filter applied at hypothesised signal masses of $m_X = 300$, 600 and 1000 GeV and their fiducial selection efficiency was determined. As shown in Figure 6.17, these efficiencies match the filtered sample efficiencies within statistical

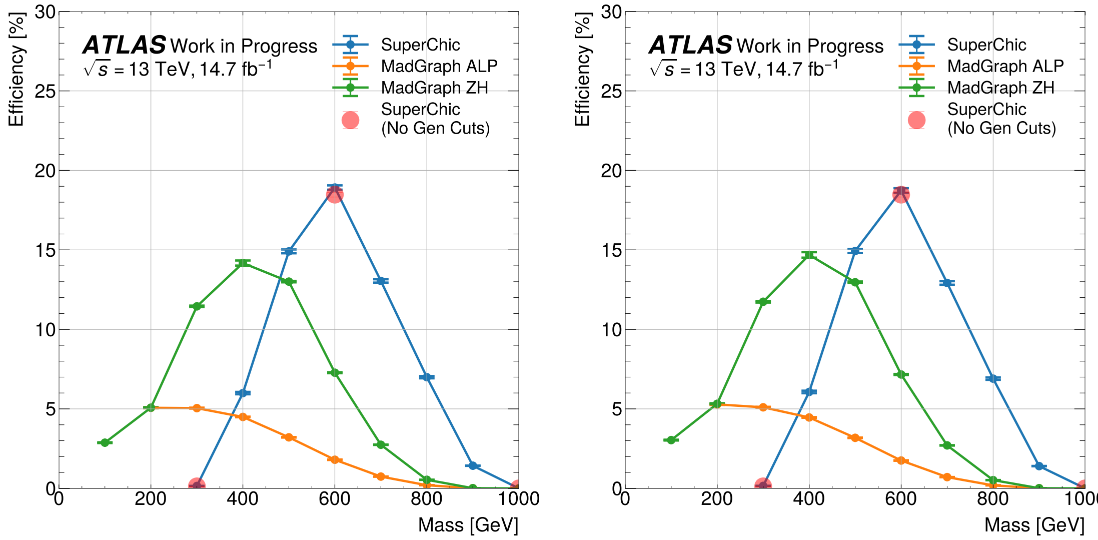

(a) Muon channel (b) Electron channel

Figure 6.17: Fiducial selection efficiency of each signal model in the (a) muon and (b) electron channel. Efficiencies calculated from additional SuperChic samples without generator level filters applied are overlaid to verify that the filter efficiency is correctly accounted for.

fluctuations for the limited points at which they are evaluated, verifying that no bias is present. These plots also show extremely low efficiencies for the SuperChic model with $m_X > 900$ GeV and both MadGraph models for $m_X > 800$ GeV, demonstrating that sensitivity is lost for signals above these thresholds due to the limited acceptance of the AFP spectrometer. Therefore, signal masses beyond these limits are not considered for the final result.

6.5.6.1 Track Veto

The most important background suppression requirement utilised in the analysis is the track veto, which requires that there are no ID tracks present in addition to the tracks comprising the signal leptons within a certain threshold of $|z_0^{\text{track}} - z_0^{\ell\ell}|$ around the dilepton vertex. Here z_0^{track} is the track z position relative to the primary vertex of the event, and $z_0^{\ell\ell} = (z_0^{\ell_1} + z_0^{\ell_2})/2$ is the dilepton vertex z position, taken as the average of the z positions of the two signal leptons $\ell_{1,2}$. The track veto, denoted $N_{\text{tracks}}^{\text{window}} = 0$, removes non-exclusive processes (with inner detector activity

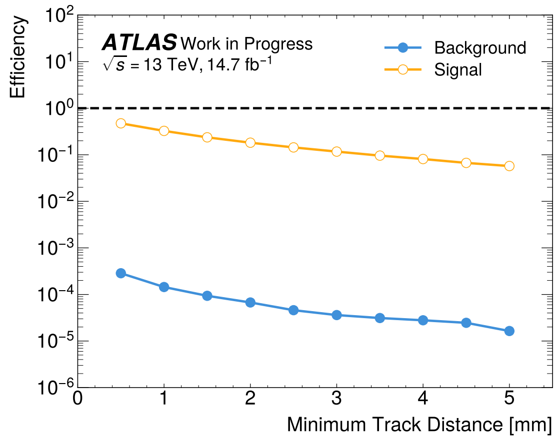

Figure 6.18: The estimated efficiency for signal and background processes in data with respect to the track veto selection with different considered window sizes (corresponding to the minimum allowed track distance from the dilepton vertex).

in addition to the two signal leptons) with very high efficiency. This applies to the majority of the processes which contribute to the background for this analysis, with additional particles such as jets produced alongside the dilepton system.

The estimated efficiency for signal and background processes with respect to this selection is shown in Figure 6.18, for different considered window sizes, with larger window sizes yielding lower selection efficiency due to a higher probability of finding an additional track within larger windows. This demonstrates the effectiveness of this selection, with over 99.9% of background events removed even for the loosest considered window size of 0.5 mm.

However, a small proportion of signal events are also removed by this selection, even though the signal is exclusive and so does not produce any additional tracks in the central detector. This arises from random independent pile-up vertices producing tracks within the set window around the dilepton vertex. The probability of this occurrence for a given window size is estimated from data by selecting a random z position in each event and testing whether any tracks not originating from the signal process fall within the given window of this position. To ensure the distribution of these randomly selected points is representative of the expected signal vertex distribution, the primary vertex position of the previous event in the dataset is used. This method determines the expected "exclusive efficiency" for a given window size, which is shown alongside the expected background efficiency in Figure 6.18. An alternative pile-up based method to estimate the exclusive efficiency was used in several previous analyses using track veto selections [85, 87]. This method was tested and gave compatible results with the method described above. An optimisation study was performed for the size of the track veto window between the potential values plotted above, and the window size was found to have a negligible effect on the resulting sensitivity, with the increased background removal for larger windows balanced by the reduced signal efficiency. Therefore, the minimum considered window size of 0.5 mm was chosen to maximise the available statistics.

The track veto signal efficiency can also be estimated directly from the simulated signal samples as the fraction of events passing the selection, as shown in Figure 6.19 as a function of the signal mass. Consistent efficiencies are observed across all signal

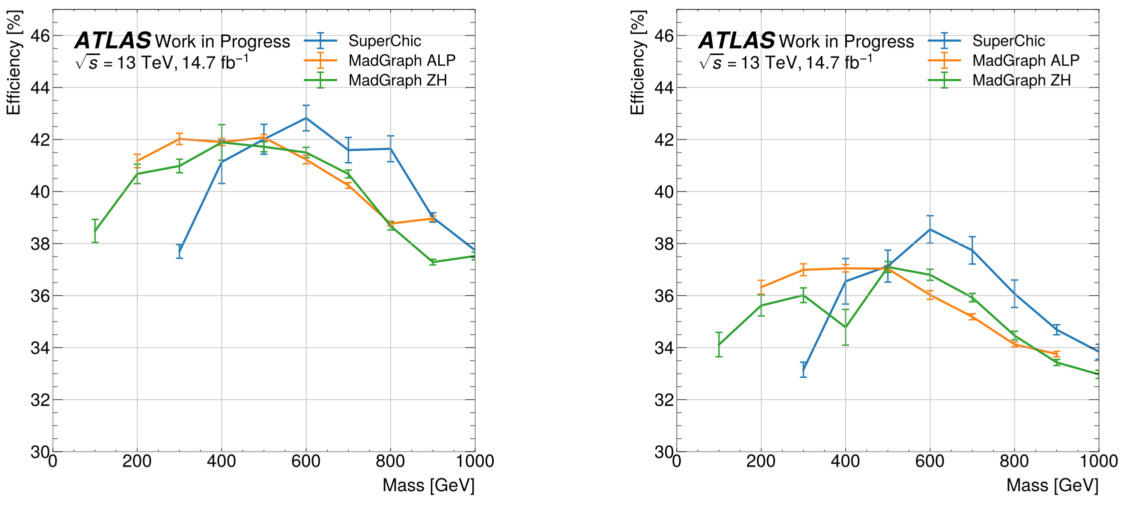

(a) Muon channel (b) Electron channel

Figure 6.19: Track veto signal efficiency observed directly in simulated signal samples as a function of signal mass, for each model and lepton channel.

masses, which is expected since the presence of additional tracks arises from pile-up interactions independent of the signal properties. Efficiencies of 34(38) $\pm$ 3% are observed in the electron (muon) channel, with the lower efficiency in the electron channel likely due to extra tracks produced within the window from the signal electrons via bremsstrahlung photon emission, followed by electron pair production.

Slight differences are observed between the values of track veto signal efficiency estimated from data and those from simulated signal, arising partly due to the differing beamspot size between data and simulation described in Section 6.4, and partly due to small mismodelling effects in the simulation. Therefore, a systematic is assigned to the signal normalisation to account for this difference, as described in detail in Section 6.7.1.

6.5.7 Low-p_T Tracks

6.5.7.1 Reconstruction in Data

During the analysis the use of "low-p_T tracks" was investigated as a potential means to increase sensitivity. These are ID tracks with $100 < p_T < 500$ MeV, which are not provided by the default ATLAS experiment track reconstruction described in Section 3.3.1, which only reconstructs tracks with $p_T > 500$ MeV, hereafter referred to as "high-p_T tracks". Low-p_T tracks are reconstructed using a recently developed technique requiring an additional processing step for data [199]. By adding low-p_T tracks into consideration for the track veto selection described in Section 6.5.6.1, much softer tracks from the additional event activity expected in background processes can be included. This is currently being investigated as an extension of an existing ATLAS collaboration analysis measuring photon-induced W^+W^- production [87], with up to a factor of 5 increase in background suppression observed.

The dedicated low-p_T track reconstruction process is highly computationally expensive and slow, and therefore it is used in conjunction with an "event picking" service which selects only specified events for reconstruction, creating a reduced sample size to reduce computation time. A filtered data sample was produced by cutting on the maximum number of standard high-p_T ATLAS tracks with $p_T > 500$ MeV within the 0.5 mm window of the dilepton vertex in each event. Since any event passing the eventual "low+high-p_T" track veto including low-p_T tracks would necessarily have zero high-p_T tracks within the veto window, this filtered event sample includes every potential event which could enter the signal region. Events with higher numbers of high-p_T tracks within the veto window than the cut value will not enter the signal region, and therefore it is not worth performing the additional reconstruction step for these events. Figure 6.20 shows the number of events passing the pre-selection, as described in Section 6.5.4, with different limits of additional high-p_T tracks within the 0.5 mm veto window. A threshold of $N^{0.5\text{ mm}}_{p_T>500\text{MeV}} \leqslant 15$ was chosen

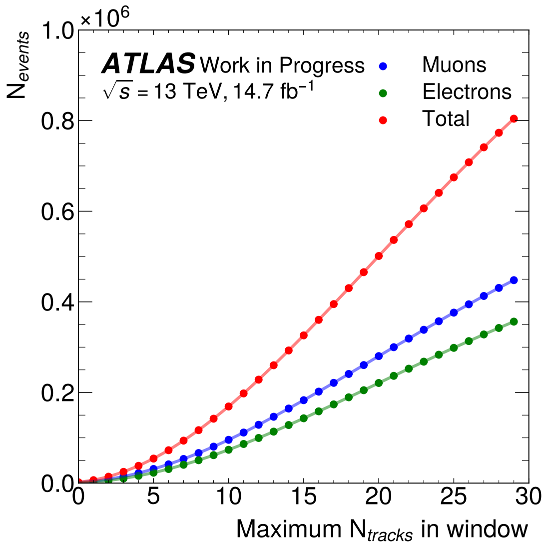

Figure 6.20: Number of data events as a function of the maximum number of additional high-p_T tracks with 0.5 mm of the dilepton vertex.

yielding around 3×10^5 events across both channels, which is low enough to perform low-p_T track reconstruction efficiently while maximising statistics for more detailed background studies within the analysis.

The filtering process does not introduce any bias into the data, since the final selection uses a tighter cut of $N^{0.5\text{ mm}}_{p_T>500\text{ MeV}} = 0$. However, it does affect the data-driven background model produced with the event-mixing procedure described in Section 6.2.3, as it defines the sample of pre-selected data events available for mixing, altering which events are combined in the background model. However, for a background model averaged over a sufficiently large number of orthogonal event-mixed samples, any difference due to this effect will become negligible, which was verified by comparing several filtered and unfiltered background models produced with different ranges of event-shift i, as shown in Figure 6.21. Initially, for a single

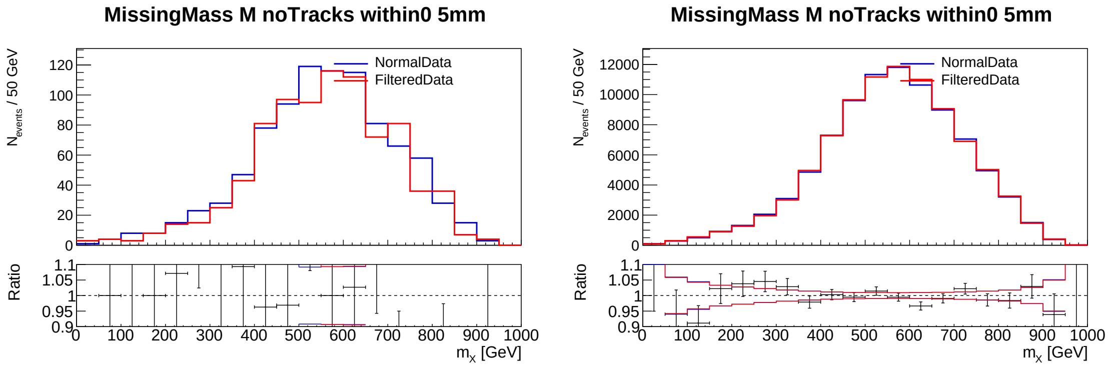

(a) $i = 1$ (b) $2 \leqslant i \leqslant 101$

Figure 6.21: Comparison of background models produced using normal data versus filtered data with a cut of $N^{0.5\text{ mm}}_{p_{\mathrm{T}}>500\text{ MeV}} \leqslant 15$ imposed on the number of additional high-p_{T} tracks within 0.5 mm of the dilepton vertex, using different numbers of summed event-mixed samples. Results are shown for the muon channel. The red and blue lines on the ratio plots show the statistical uncertainty ($\sqrt{N}$) of the corresponding distributions in the top panels. The same level of agreement is observed in the electron channel.

event-mixed sample, there are visible differences between the background models due to statistical fluctuations. However, when a high number of orthogonal event-mixed samples are combined across the range of event-shifts $2 \leqslant i \leqslant 101$, the differences become increasingly small, verifying that no bias is introduced by the filtering process. This is the range of event-shifts which is used for the data-driven background model, as described in detail in Section 6.6.1, and so no bias is expected.

6.5.7.2 Selection

Several quality control criteria are applied to low-p_{T} tracks analogous to the Loose quality working point applied to high-p_{T} tracks described in Section 6.5.3. These selections are as follows:

- $p_{\mathrm{T}} > 100$ MeV
- Number of SCT holes $\leqslant 2$
- Number of pixel holes $= 0$

- Number of pixel hits + dead modules $\geqslant 4$
- $d_0/\sigma_{d_0} < 3.0$ for tracks with $p_T > 250$ MeV
- There must be a hit in either the IBL or the B-Layer if no hit in the IBL is expected

The following additional criteria are also imposed to significantly reduce the rate of fake tracks, which are prevalent when reconstructing tracks with such low energy:

- $|\eta| < 2.5$
- $|d_0| < 1$
- $|z_0| < 1$
- $p_T > \max(120, 81.5/\sin(\theta))$ MeV

Overlap removal for low-p_T tracks with the tracks associated with the lepton candidates is performed using a $\Delta R < 0.01$ cut.

6.5.7.3 Low+High-p_T Track Veto

As introduced above, the selected low-p_T tracks were included into consideration for the track veto to give the "low+high-p_T" track veto selection. The signal and background efficiencies of this tightened track veto were estimated in simulated Z+jets samples for the same range of window sizes considered above for the standard track veto, now referred to as the "high-p_T" track veto. These are plotted alongside the efficiencies estimated for the high-p_T track veto in Figure 6.22. This consistently shows an improvement in background rejection by around a factor of five over the high-p_T track veto across all window sizes as expected, with only a 10-20% decrease in signal efficiency, where the ratio is generally largest for smaller window sizes. The optimisation study performed above was repeated with the inclusion of low-p_T tracks, and the same results were obtained, with no significant effect of the

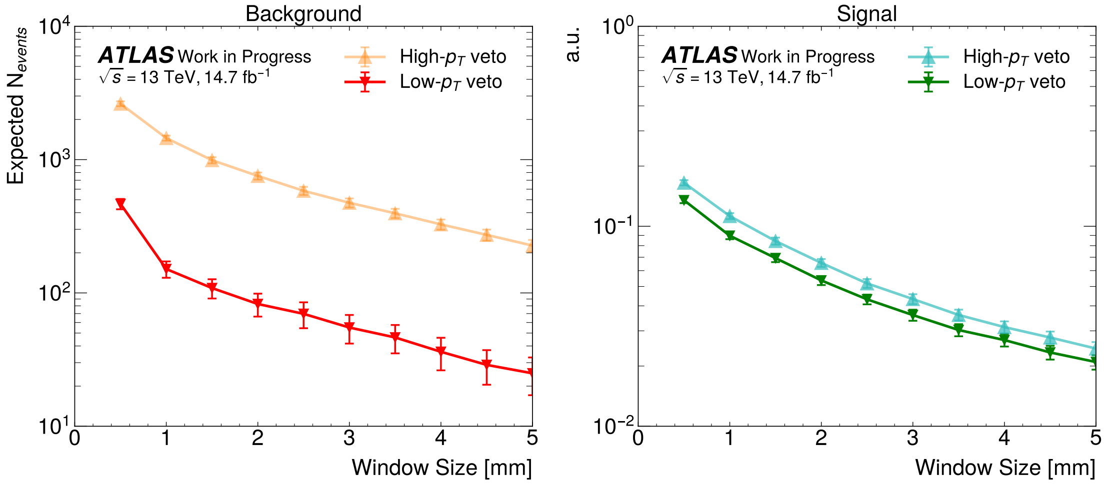

Figure 6.22: Comparison of estimated background and signal event yields in simulated Z+jets events obtained when applying high-p_{T} versus low+high-p_{T} track vetoes for different considered window sizes. The signal event yield is scaled to an arbitrary cross section, as only the ratio between the two veto yields is considered.

chosen window size on the resulting sensitivity. Therefore, the 0.5 mm window was maintained to maximise the available statistics.

The results of the analysis obtained applying both the standard high-p_{T} track veto and the extended low+high-p_{T} track veto are considered separately. This results in two corresponding signal regions, with the high-p_{T} signal region corresponding to the selection given in Table 6.3. The tighter low-p_{T} signal region corresponds to the same selection, with the only difference being the lower p_{T} threshold for tracks considered in the track veto, changing the corresponding requirement in Table 6.3 to $N^{0.5\text{ mm}}_{p_{\mathrm{T}}>100\text{ MeV}} = 0$.

6.6 Background Modelling

The dominant source of background in the analysis is a combinatorial background produced when central dilepton systems arising from non-signal SM processes are wrongly combined in reconstruction with unassociated protons originating in independent pile-up interactions. This is illustrated in comparison with the signal process in Figure 6.23, with each pile-up proton typically coming from different single-diffractive processes. Such processes have cross sections on the order of tens

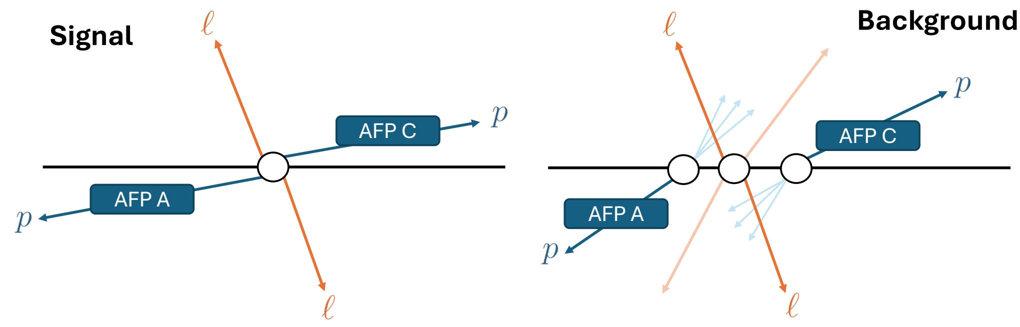

Figure 6.23: Comparison of signal events to the dominant combinatorial background process for the analysis.

of millibarns, so when combined with relatively high cross section central dilepton production processes this produces a large background. Several SM processes are considered which contribute to the central component of the background:

- Z+jets production
- Top quark production: Wt and $t\bar{t}$
- Diboson production: WW, WZ and ZZ
- Photon-induced dilepton production: $\ell\ell$, WW
- Misidentified leptons

Since all the above processes give similar kinematic distributions in the final state under the signal selection, they are treated as a single combined background, hence the “combinatorial” nature of the background.

6.6.1 Data-driven Modelling Method

The combinatorial background is modelled in the analysis using a data-driven method called event mixing, described in detail in Section 6.2.3. The use of this data-driven approach is hugely advantageous, as all events are drawn directly from data and so the resulting kinematic distributions and event yields match exactly to the expectation from data. This removes the need for any simulation to be used in the background modelling, which simplifies the analysis as the corresponding corrections required for matching simulated distributions to data, including the overall event yield, do not need to be considered. Furthermore, the associated uncertainties and mismodelling effects discussed in Section 2.4 do not affect the background modelling. This results in very small systematic uncertainties in the final result, with the majority arising from the signal simulation, which have much smaller effects on the results than uncertainties in the background modelling.

Additionally, the event-mixing procedure allows many different uncorrelated samples to be produced, since each value of event-shift i yields a unique combination of the lepton and proton data. These background samples can then be combined and averaged to produce a much larger statistics model of the background, reducing the statistical uncertainty significantly. The event-shift can take any value in the range $0 \leqslant i \leqslant N-1$, with $i = 0$ corresponding to the unaltered dataset which could contain signal, and $i > 0$ corresponding to a different model of the background for each i. As described in Section 6.3.1, a sample with event-shift $i = 1$ was used as the blinded dataset during the analysis, so only orthogonal samples with $i \geqslant 2$ were considered for the background model. A study was performed on the estimated statistical uncertainty of the background model as a function of the total number of orthogonal background samples (each corresponding to a different value of i) which are averaged over to produce the full high-statistics model. This study is described in detail in Section 6.7.3. As a result of this study, it was chosen to use a total of 100 background samples, for a balance between minimising uncertainty and computation time. Therefore, event-mixed samples with event-shifts range $2 \leqslant i \leqslant 101$ are averaged over to model the combinatorial background.

In order to ensure that the data-driven background model correctly accounts for all expected contributions to the combinatorial background, an alternative background model was produced using simulation. All of the contributions to the central dilepton background listed above were included using MC samples, except for the misidentified lepton contribution which was modelled using a same-sign lepton selection in data. This validation is described in detail in Appendix A, and the resulting comparison between the missing mass distributions from the data-driven and simulated models are shown in Figure 6.24 for each lepton channel. Good agreement

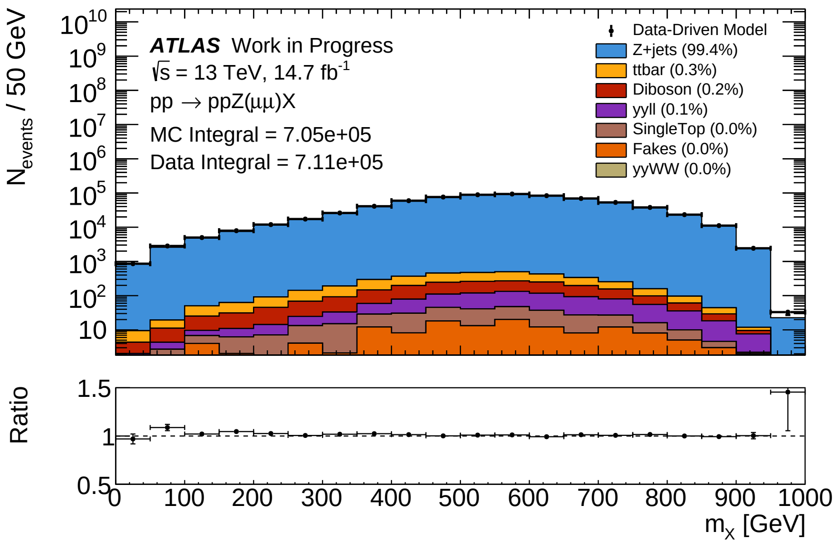

(a) Missing mass m_X

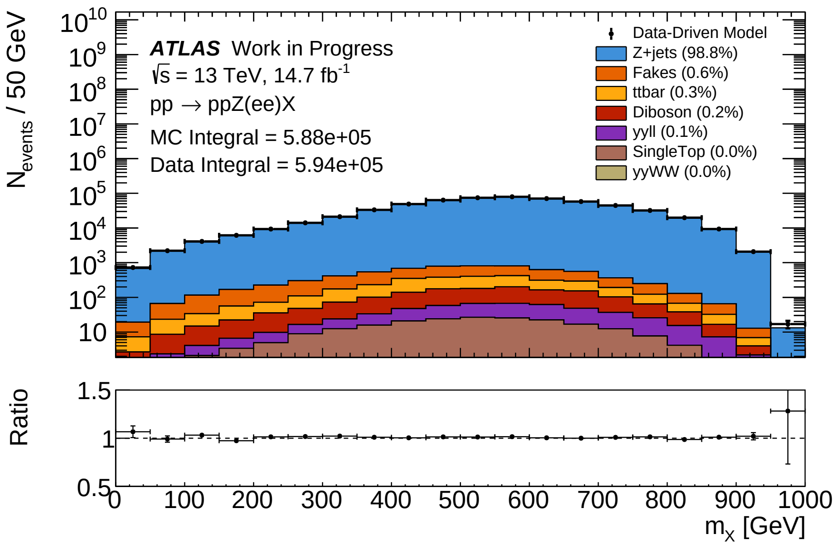

(b) Missing mass m_X

Figure 6.24: Reconstructed missing mass m_X distributions from the combined simulated background model produced with all considered background contributions, after all signal selections are applied except for the track veto and dilepton p_{T} cut, in the (a) muon and (b) electron channels.

is observed within the expected degree of statistical fluctuation in both channels, validating the accuracy of the data-driven model.

6.6.2 Signal-induced Background

If a signal exists, then an additional, signal-induced background arises from signal events where at least one of the signal protons does not fall within the signal selection of $0.035 < \xi < 0.08$. If a single pile-up proton from an independent interaction which does fall within the signal selection is instead measured on the corresponding side of the AFP spectrometer, then the reconstruction will be carried out with this incorrect proton, leading to an inaccurate value of the missing mass. This "mismatched proton" effect is illustrated in Figure 6.25. Typically, this occurs in events with very

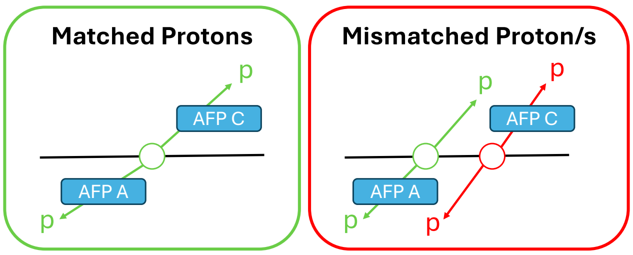

Figure 6.25: Example of an event with both signal protons correctly measured by the AFP spectrometer (left) and an event where one of the signal protons is missed by the AFP spectrometer and a single pile-up proton from a background interaction is measured in its place (right).

high signal proton ξ, and since high proton energy loss is balanced in signal events by high dilepton four-momentum, when a mismatched, lower ξ proton from a pile-up interaction is instead used in Equation 6.1 this can sometimes lead to negative values of reconstructed m_X^2. Clearly this is non-physical, resulting in imaginary values of m_X. To illustrate this effect, lepton kinematic distributions are compared between simulated signal events from a single sample with positive and negative values of reconstructed m_X^2 in Figure 6.26, showing the expected relation of higher lepton momenta in events with negative reconstructed m_X^2. These events are easily discarded in the selection, by requiring $m_X^2 > 0$ at reconstructed level. However,

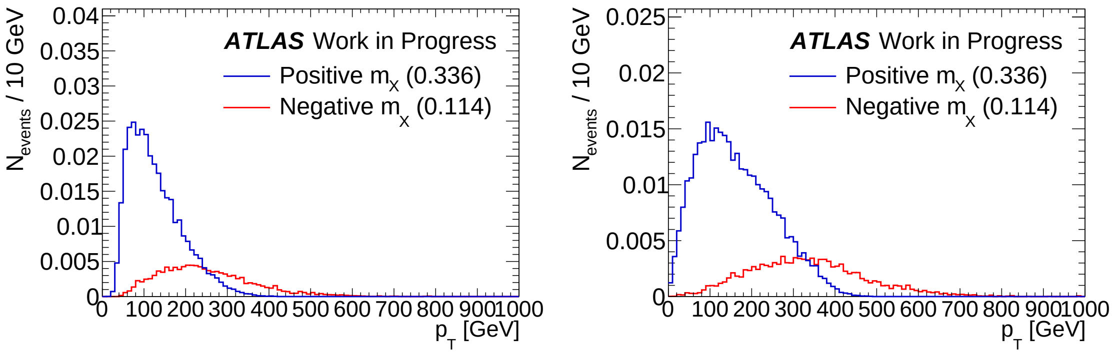

(a) $p_{\mathrm{T}}\ (\ell_1)$ (b) $p_{\mathrm{T}}\ (\ell\ell)$

Figure 6.26: Comparison of lepton kinematic distributions in events with positive versus negative reconstructed missing mass.

often reconstruction using an incorrect pile-up proton will lead to a positive value of m_X^2, giving a valid missing mass, which is more difficult to handle.

A detailed investigation was carried out into the nature of this background and how to mitigate it, as presented in Appendix B. Ultimately, it was found that the fiducial selection applied to signal samples as described in Section 6.5.6 removes most of this background for the majority of hypothesised signal masses. For models with $m_X = 900$ GeV, 10% of the signal yield is still found to originate from this mismatching process. However, a narrow mass window is applied in the final fits, as described in Section 6.8.2, which further removes the signal-induced background. Any remaining contribution is therefore negligible in the final fits, and so it is not included.

6.7 Systematic Uncertainties

Two main categories of uncertainty must be considered in the analysis: statistical and systematic. Statistical uncertainties relate to the finite size of the considered dataset, which leads to random fluctuations in event yields due to the stochastic nature of particle interactions. Any other uncertainty which arises due to the experimental methodology must be accounted for as a systematic uncertainty. This covers uncertainties related to detector resolution and calibration, reconstruction methodology and modelling of signal and background processes, which can affect both the shape and overall yield of the observable distribution. The following sections describe the sources of systematics uncertainty considered for this analysis and the methodology for determining their effects on event yields.

6.7.1 Central Detector Uncertainties

Several experimental systematic uncertainties relating to the central detector are considered, which are mostly investigated by parametrising their effect on the simulated signal samples. Uncertainties arise in lepton reconstruction from the calibrations applied to energy and momentum scales and the track resolutions of the ID and MS. In addition, as discussed in Section 3.3, a range of scale factors are applied to simulated samples to match lepton reconstruction efficiencies in simulated samples to those observed in data, determined via tag-and-probe analyses. The corresponding uncertainties are included as systematics, with reconstruction, identification and isolation efficiencies considered for both lepton flavours and track-to-vertex association additionally considered for muons. Each scale factor is varied by its uncertainty in simulation and the analysis chain is rerun, to determine the effect on the final observable distribution. The same procedure is used to account for uncertainties in the modelling of trigger inefficiencies for both lepton flavours. The scale factor applied to match the distribution of pile-up interactions between simulation and data carries an uncertainty of 4%, which is additionally propagated as a systematic uncertainty, and an uncertainty of 1.13% on the total integrated luminosity of the selected dataset is applied [200].

The final central systematic considered is the uncertainty on the signal efficiency of the track veto selection. This arises due to uncertainties in the modelling of pile-up interactions, and less significantly due to the difference in the distributions of beamspot size between the samples, as discussed in Section 6.5.6.1. This uncertainty is taken from the difference observed between efficiency estimates made from data and simulated signal samples. This systematic also covers the uncertainty in the modelling of the tracking efficiency and fake-rate.

The track veto signal efficiency uncertainty is determined by comparing the efficiency measured directly from the simulated samples, the "lepton vertex efficiency", to the pile-up based exclusive efficiency estimate described in Section 6.5.6.1, which selects a random position on the z axis in each event and assesses whether a signal vertex in that position would pass the track veto. Both methods are applied to the simulated signal samples to allow for a direct comparison, and the results are averaged across all signal models and masses to give high statistics, since it was shown in Section 6.5.6.1 that the signal efficiency is consistent with being independent of both signal model and mass. Figure 6.27 shows the resulting efficiency estimates from each method, measured as a function of the mean number of interactions per bunch crossing μ. The track veto signal efficiency is strongly correlated to the value of μ, since it arises

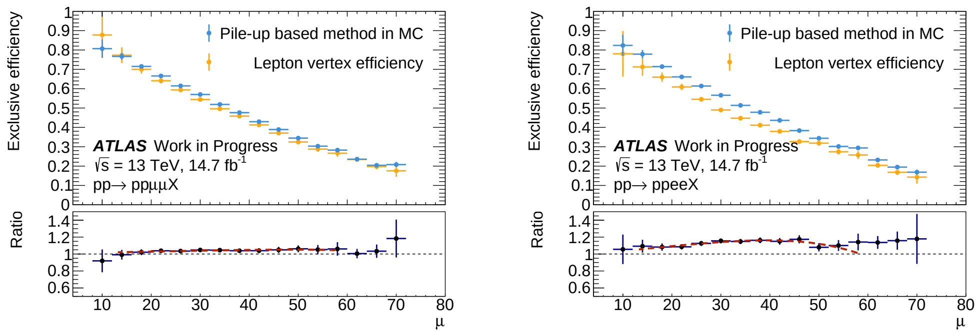

(a) Muon channel (b) Electron channel

Figure 6.27: Comparison of track veto signal efficiencies estimated as a function of the mean number of interactions per bunch crossing μ, using a pile-up based approach and calculated directly from the lepton vertex, in simulated data in the (a) muon and (b) electron channels. The ratio is fitted to a quadratic polynomial to reduce statistical fluctuations, and the fit range is limited to within $\pm 2\sigma$ of the mean value of μ to remove outliers.

directly from the random coincidence of pile-up tracks with the signal vertex within the closed 0.5 mm window, which becomes increasingly likely with a larger number of pile-up vertices. The efficiency distributions are fitted to second-order polynomials to reduce the effect of statistical fluctuations. The fit is restricted to the region of μ within $\pm 2\sigma$ of the average value measured in data (corresponding to $\mu = 35 \pm 22$), to remove outliers. The maximum deviation of the ratio from unity in the resulting fits, shown overlaid in Figure 6.27, is taken as the track veto signal efficiency uncertainty. The uncertainty is determined separately in each lepton channel, due to the different reconstruction methods for each lepton flavour, resulting in uncertainties of 5.2% in the muon channel and 12.9% in the electron channel.

The electron channel uncertainty is significantly higher. This arises due to a reduced direct efficiency estimate from simulated signal vertices compared to muons, since the pile-up based methods agree between the two lepton channels. This is following an adjustment to the track selection for this study, which increased the ΔR threshold given in Section 6.5.3 used to remove tracks corresponding to the signal leptons from consideration for the track veto, from 0.01 to 0.1 for the electron channel. This is done to remove a significant proportion of tracks observed in the low ΔR region around the dielectron vertex which was not observed in the muon channel. This peak is attributed to pair production from bremsstrahlung photons emitted from the signal electrons, and caused an even higher uncertainty to initially be observed in the electron channel. The remaining difference is attributed to the track-to-vertex association cuts applied for electrons described in Section 6.5.1, which have been observed in previous analyses to cause this discrepancy between the lepton channels [85, 87]. An additional contribution is expected from less energetic bremsstrahlung photons producing electrons at higher ΔR that cannot be separated from other tracks.

The above study was repeated with the inclusion of low-p_{T} tracks yielding slightly increased uncertainties of 6.5% in the muon channel and 16.0% in the electron channel. In the final results a combination of both lepton channels is additionally considered, for which the electron uncertainty is used because it is larger.

6.7.2 Forward Detector Uncertainties

Several dedicated systematic uncertainties are considered related to both the AFP spectrometer and forward proton reconstruction, falling into two categories. Those in the first category are similar to the central systematics listed above, accounting for uncertainties on measured quantities. The uncertainty of 300 μm on the global alignment procedure detailed in Section 4.4.2 is included as a systematic, with the SiT station positions shifted in either direction by this uncertainty before rerunning the proton reconstruction, to determine the effect on the final observable. In addition, an uncertainty of 4 mrad on the local rotational alignment of the planes within each station is included in the same manner. The uncertainty on the proton transport simulation used to determine proton energy loss ξ from the x position measurement provided by the AFP spectrometer, as discussed in Section 4.3, is accounted for by recalculating the proton properties using alternative transport simulations in which the beam slope is altered by $\pm 50\,\mu$rad. Finally, a systematic is included accounting for the uncertainty on the reconstruction of simulated protons in the AFP spectrometer. As discussed in Section 2.4, smearing is applied to the properties of truth-level protons to simulate the resolution of AFP reconstruction. The degree of applied smearing was optimised to match the reconstructed resolution observed in data as closely as possible, and ± 0.05 mm is taken as a conservative estimate of the systematic for this property. Truth-level signal protons were reconstructed using the corresponding altered smearing values and the effect on the signal event yield was observed.

The second category includes several systematics which alter the proton reconstruction method relative to the nominal procedure detailed in Section 4.3. The considered systematics and their effects on proton reconstruction are the following:

- **CLUST_NEIGHBOUR** modifies the search space for neighbouring AFP SiT plane hits when reconstructing clusters from looking only in rows (x) to looking in both rows and columns ($x-y$).

- **TRK_FIND_DIST** modifies the maximum allowed distance between clusters in the track reconstruction from 0.5 mm to 0.4 mm.
- **TRK_FIND_CLUST** modifies the minimum required number of clusters in the track reconstruction from 2 to 3.
- **TRK_SEL_MATCH** modifies the cut on the distance between track positions in the $x - y$ plane in the proton reconstruction from 2 mm to 1 mm.

The effect of each of these systematics can be evaluated for both simulated signal and data, to verify that the expected variations are observed.

Very small effects are observed for the CLUST_NEIGHBOUR and TRK_FIND_DIST systematics, which have minor effects on reconstructed cluster position and track reconstruction, respectively. A larger effect is observed for TRK_FIND_CLUST, which is designed to probe any inefficiency in the AFP spectrometer SiT planes which is not accounted for by the dedicated GRL by requiring an extra cluster per track over the default two, which typically will require three out of four planes in a station to have hits. Based on the estimated efficiency of individual SiT planes of at least 90% discussed in Section 4.5, the probability of observing a hit in at least two planes in a given station from an incident proton is $> 99\%$, which reduces to around 95% for the tighter requirement of hits in at least three planes. This can be propagated to suggest a 5-10% effect on event yields from the tightened selection imposed by TRK_FIND_CLUST.

The effect of the TRK_SEL_MATCH systematic changes depending on the sample tested, as this variation tightens the requirement in double-station reconstruction on the transverse distance between the tracks in each station. This should have a minimal effect on intact protons, such as those expected from the signal process, and these are scattered at very small angles with respect to the beamline and therefore should not move a significant distance in the transverse plane in the gap between two AFP spectrometer stations. However, in background processes, protons can interact in the sensitive layers or the RP windows, causing a shower of particles to form between the two stations. This leads to shower particles which are scattered at high angles, which fail the tightened track matching criteria. Therefore, a large effect is observed for this systematic in data and background samples, but not in signal samples.

An issue was encountered with the TRK_FIND_CLUST systematic in data, which showed an unexpectedly large effect of around 95% due to only two out of four planes in the AFar station being active for large parts of 2017 data-taking. This can be seen in Figure 6.28, which shows the number of reconstructed clusters in each AFP SiT station for double-station reconstructed protons in nominal data. The mean number

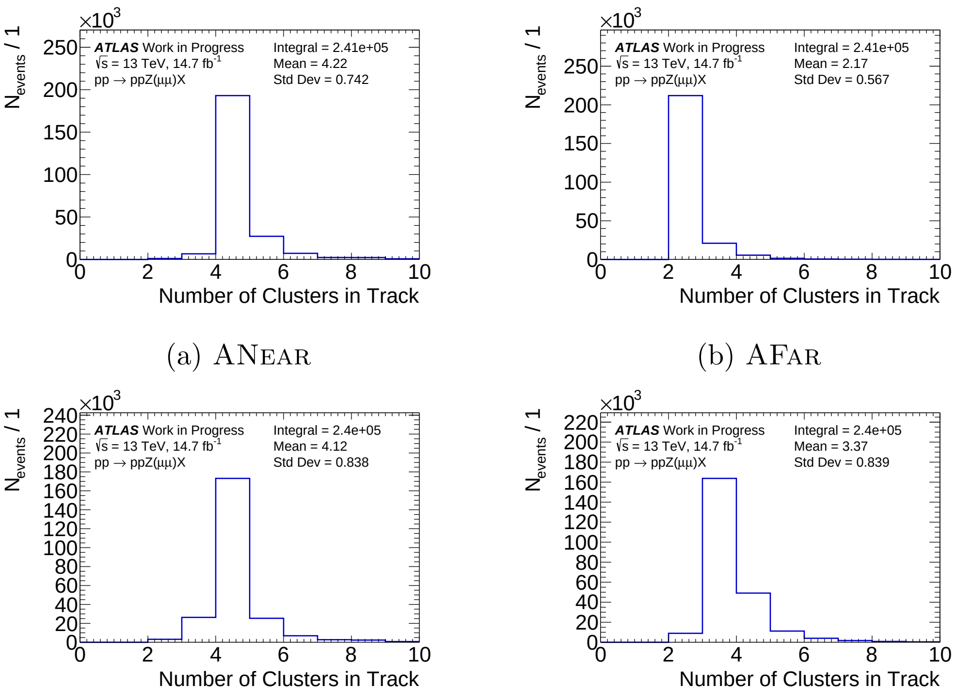

(a) ANear (b) AFar

(c) CNear (d) CFar

Figure 6.28: Number of reconstructed clusters per proton track in each AFP SiT station for double-station reconstructed protons across the nominal dataset.

of clusters reconstructed in AFar is very low at 2.17, compared to 3.37 for the CFar station. This means that when the proton track requirement is increased from at least 2 clusters to at least 3 clusters per track, only 12% of previously reconstructed protons on Side A are accepted, compared to 95% efficiency on Side C. To decouple the effect which the TRK_FIND_CLUST systematic is meant to study from this issue with AFar in 2017, the measured efficiency of Side A must be ignored. Instead, the effect of the tightened cluster requirement is estimated by assuming that without this issue, the efficiency of Side A with respect to this tightened requirement would be the same as observed on Side C (95%). Applying this estimation yields corrected overall efficiencies of the tightened cluster requirement of 90.5% in the muon channel and 90.3% in the electron channel. This correction was also determined bin-by-bin across the missing mass distribution to account for any shape dependence, to obtain the distributions shown below.

Figure 6.29 shows the effects observed in data from each systematic variation discussed above. This exercise was performed both on blinded and unblinded data,

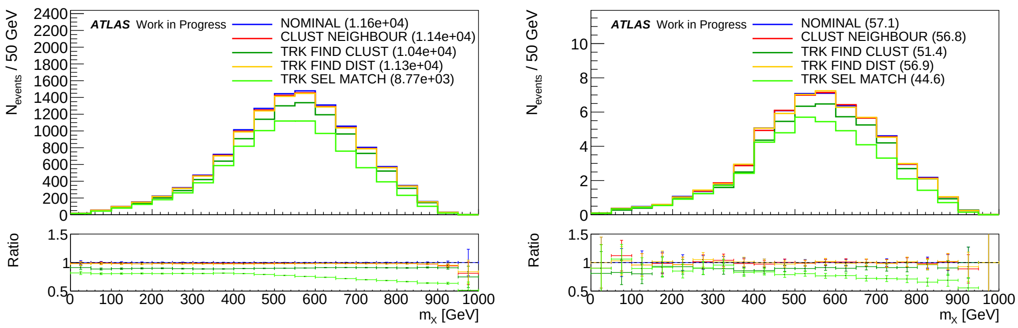

(a) Before track veto

(b) After high-p_{T} track veto

Figure 6.29: Missing mass distributions in blinded data for each AFP systematic which affects proton reconstruction methodology (a) before applying the track veto selection and (b) after applying the high-p_{T} track veto.

and the data-driven background model, with compatible variations observed in every case.

In the case of the simulated signal samples, smaller variations on the order of a few percent are observed for TRK_FIND_CLUST, with the other systematics having negligible effects, as expected. The effect of these systematics on the simulated signal are included alongside the other AFP systematics listed above in the final fits.

6.7.3 Modelling Uncertainties

As described in Section 6.6.1, a data-driven background model is produced using the event-mixing procedure presented in Section 6.2.3. A high-statistics model is produced by averaging over multiple samples with different values of event-shift i, which by construction are all orthogonal to each other. Modelling systematics affecting the data-driven background model were considered from several potential sources.

The statistical uncertainty of the background model was evaluated using bootstrapping, whereby many replicas of the sample are created, with each event in each sample assigned a random weight drawn from a Poisson distribution with a mean of 1. The statistical uncertainty on the observable distribution of the background model can then be determined from the standard deviation of the height of each bin among the replicas. Figure 6.30 shows the results of bootstrapping for background models produced by averaging over $N = 1, 10$ and 100 orthogonal event-mixed samples. As expected, for larger numbers of samples, the statistical uncertainty

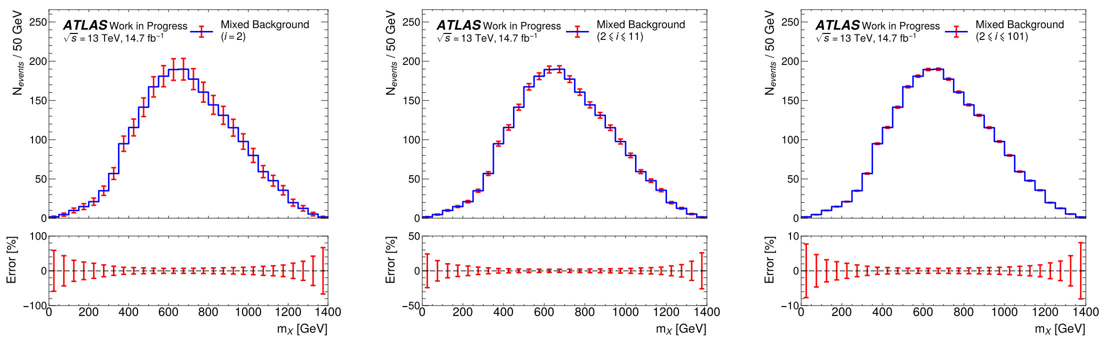

(a) $N = 1$ (b) $N = 10$ (c) $N = 100$

Figure 6.30: Bin-by-bin statistical uncertainty determination for the event-mixed background model estimated using bootstrapping, for different numbers of averaged orthogonal event-mixed samples N. The uncertainties are shown for each bin in the top panel, with the fractional uncertainty (divided by the absolute value) shown in the bottom panel.

decreases significantly, from around 10% for a single sample to 1% for 100 samples, with a slow variation across the missing mass distribution depending on individual bin height. Further improvement can be obtained using even more samples, with around 0.3% uncertainty estimated for 1000 samples. However, generating such a large number of samples is computationally expensive, and therefore 100 samples were chosen for a balance between low uncertainty and low computation time.

In addition to the statistical uncertainty, several potential sources of systematic uncertainty arise from the event-mixing procedure. Firstly, any time dependence of the data-taking conditions for the ATLAS detector or the AFP spectrometer would introduce potential differences between the events which are mixed together to form the background model, which could affect the final observable. However, this was already investigated in Section 6.5.5 and it was found that the tightened ξ selection in the signal region relative to the detector acceptance reduced any time variance of the mean reconstructed missing mass to below the resolution of the missing mass method. Therefore, this uncertainty is considered negligible.

A second potential effect is the differing pile-up distributions between events which are mixed together, which could affect the background model in which most protons are assumed to originate in pile-up interactions. The difference in mean interactions per bunch crossing $|\Delta\mu|$ was evaluated between the events which are mixed together in the background model, for a range of total sample numbers N used in the model, as shown in Figure 6.31. The average value of $|\Delta\mu|$ increases for larger N, which

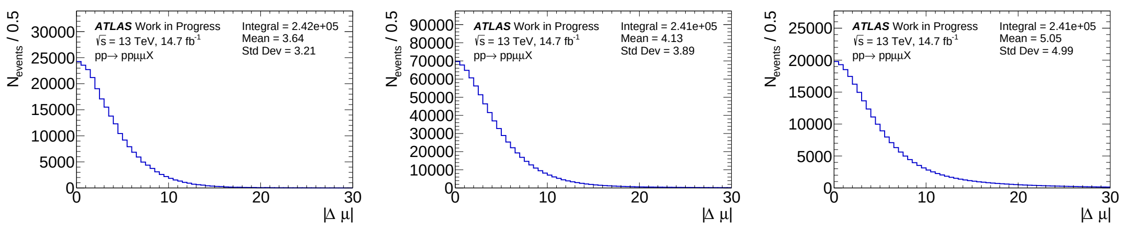

(a) $i = 1$ (1 sample) (b) $2 \leqslant i \leqslant 11$ (10 samples) (c) $2 \leqslant i \leqslant 101$ (100 samples)

Figure 6.31: Difference $|\Delta\mu|$ between the average number of interactions per bunch crossing for the central event and the proton event used in the event-mixing procedure, for different total numbers of combined samples.

follows from these samples including event-mixing with larger values of event-shift i, corresponding to combining events which were recorded further apart in time. Therefore, there is more opportunity for variation in the pile-up distribution between mixed events at large i within $i < N$. However, the increase in mean $|\Delta\mu|$ of around 1.5 between $N = 1$ and $N = 100$ is small compared to the full range of μ values recorded during 2017 data-taking (Figure 3.2). To determine whether high values of $|\Delta\mu|$ can affect the final observable, events with low values of $|\Delta\mu| < 5$ were compared to events with high values of $|\Delta\mu| > 10$ in their m_X distributions. Figure 6.32 shows the results, with the distributions normalised to unity as the respective selection efficiencies are different between the two distributions. The shapes

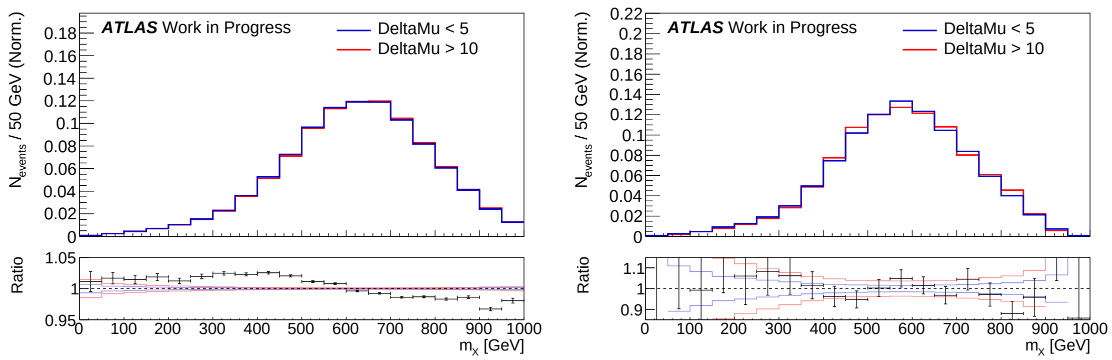

(a) Before track veto (b) After track veto

Figure 6.32: Comparison between unit-normalised missing mass distributions (normalised) (a) before and (b) after track veto using events with low $|\Delta\mu| < 5$ and high $|\Delta\mu| > 10$. The expected statistical fluctuations in each case are shown in red/blue on the ratio plot.

of the distributions match closely, although a small but significant shape effect is observed for the signal selection before the track veto. However, the shape effect is not visible following the application of the track veto, with the remaining difference between the samples consistently falling within the expected level of statistical fluctuation, as shown in the bottom panel. Therefore, it was concluded that higher values of $|\Delta\mu|$ have a negligible effect on the final observable.

Another potential uncertainty arises from the data-driven background model being produced using only double-sided mixing, where both protons from a given event are exchanged with protons from a different event. This covers most potential background processes, but does not account for any background processes where the protons involved in the central production do not dissociate and are measured by the AFP spectrometer, such as in the EE and SD channels of photon-induced dilepton production $pp \rightarrow p\gamma\gamma p^{(*)} \rightarrow p\ell\ell p^{(*)}$. To determine whether any shape difference is expected in the background from such events, a data-driven model was produced instead using single-sided mixing, where only protons from a single side of the AFP spectrometer, randomly chosen for each event[2], are replaced with independent protons from a different event, allowing any protons surviving the central background interaction to be reconstructed with the correlated central products. The resulting missing mass distributions are compared to the standard background model employing double-sided event-mixing in Figure 6.33.

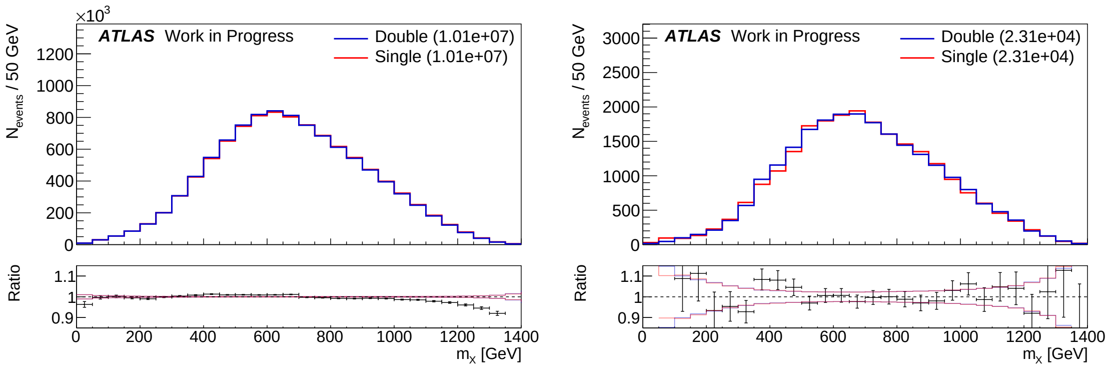

(a) Before track veto

(b) After track veto

Figure 6.33: Comparison between missing mass distributions (a) before and (b) after track veto when mixing both protons per event (double-mixing) and only on one side (single-mixing). The expected statistical fluctuations in each case are shown in red/blue on the ratio plot.

Before the track veto selection is applied, a visible shape effect is observed in the ratios near the high and low m_X limits, likely indicating the presence of an underlying background process with a single intact proton detected in the AFP spectrometer, such as diffractive Z production. However, this is not visible when the track veto is applied, with the ratios between the two distributions now falling consistently within the expected statistical fluctuations. Therefore, it was concluded that this effect is negligible in the final observable, likely due to the dissociating proton in such SD background processes producing additional central tracks which fall within the track veto window, removing the event.

[2]The effect of which side of the AFP spectrometer is chosen was investigated, with no significant difference observed, consistent with the expectation that intact protons from SD processes are evenly distributed between the two sides.

6.7.4 Theoretical Uncertainties

Systematic uncertainties are also considered relating to the choice of the baseline setup for the simulation of signal events, including the soft-survival model and parton shower settings. The effect of soft-survival is included in signal samples through a parametrisation obtained from photon-induced dilepton production simulated in SuperChic, which includes soft-survival effects not implemented in the $Z + X$ model. This parametrisation, described in Section 6.4, is then applied as an event-by-event weight to signal events as a function of the central system mass m_{ZX}. The soft-survival factor is only applied directly in this way to SuperChic $Z+X$ signal samples, since the weights are estimated using this generator. For both MadGraph models the effect is instead included as a downwards-only systematic uncertainty, taken from the parametrisation as a function of the hypothesised signal mass for each model. For the SuperChic samples, a dedicated systematic was required to account for the effect of different potential models for soft-survival which are implemented in SuperChic [80]. The fitted linear polynomial giving the parametrisation (which was used to reduce dependence on statistical fluctuations) has negligible uncertainties which do not contribute to this systematic: $p_0 = 0.9387 \pm 0.0016$ and $p_1 = -0.0003653 \pm 0.0000060$, where p_0 and p_1 are the gradient and y-intercept, repetitively. For the modelling uncertainty, measurements performed in a previous analysis of the exclusive $\gamma\gamma \to \mu^+\mu^-$ cross-section were used to determine an appropriate systematic uncertainty, based on their comparison to SuperChic predictions. Figure 6.34 shows the results of this comparison, with a maximum deviation of around 20% observed for higher dimuon masses [201]. Although the system masses considered here are well below most ZX system masses considered in this analysis, this uncertainty was nevertheless considered to be sufficiently conservative to cover any potential mismodelling. This is supported by the AFP dilepton production analysis, which observed a 10-20% difference between the SuperChic prediction and cross section measurement in each lepton channel [85]. Therefore, an uncertainty of 20% was chosen, to be applied to the overall normalisation of the simulated signal samples. For signal models with estimated soft-survival factors S above 0.8, the

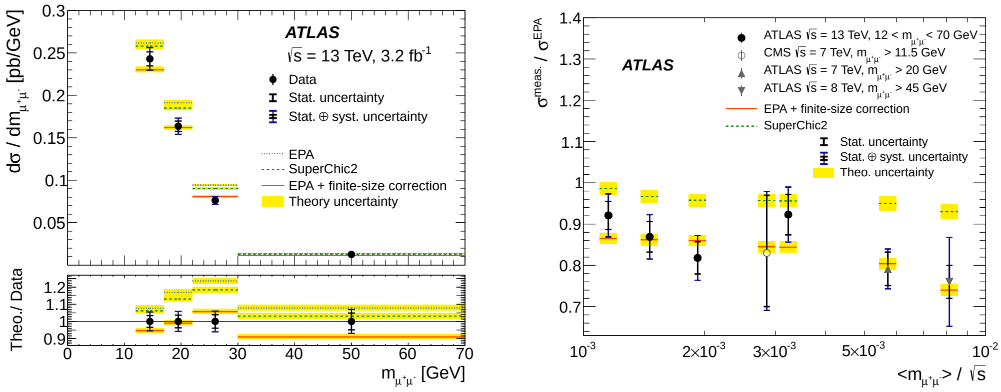

(a) (b)

Figure 6.34: (a) Exclusive $\gamma\gamma \to \mu^+\mu^-$ differential fiducial cross-section measurements as a function of dimuon invariant mass $m_{\mu\mu}$. (b) Comparison of the ratios of measured and predicted cross-sections to the bare EPA calculations as a function of the average dimuon invariant mass scaled to the pp centre-of-mass energy used [201].

upwards variations were limited to $1 - S$ in order to prevent consideration of any unphysical soft-survival factors above one.

The effect of variations in the parton shower settings for Pythia across the simulated signal samples was investigated for pre-selected events, with no significant difference observed between the missing mass distributions. Therefore, this effect was concluded to be negligible.

6.7.5 Summary of Systematic Uncertainties

The full list of systematic uncertainties described in this section are listed in Table 6.5, with systematics affecting the signal and background models shown separately. For each signal systematic, the effect of each uncertainty is shown separately for each signal model as the full range of event yield changes observed across the generated mass range.

The background uncertainties are all either negligible or very small, with systematics accounting for lepton reconstruction and simulation also having generally small effects. The most significant systematics are the uncertainties on the estimated soft

Systematic Uncertainty	**Effect on Event Yield**		
	Signal systematics		
Signal model:	SuperChic	MadGraph di-ALP	MadGraph $Z + H'$
Soft survival factor	20.0%	0.0-39.0%	0.0-39.0%
Track veto signal eff.	12.9%	12.9%	12.9%
Pile-up reweighting	3.1-4.6%	3.6-4.4%	3.7-4.4%
Luminosity	1.7%	1.7%	1.7%
AFP proton transport	8.8-70.3%	15.7-43.0%	13.3-49.8%
AFP global alignment	5.7-63.3%	12.2-43.3%	11.2-53.8%
AFP smearing	4.8-37.6%	8.7-22.0%	7.6-24.7%
AFP track matching	1.2-2.5%	1.3-1.8%	1.3-1.6%
AFP track finding	0.6-1.2%	$< 0.1\%$	0.5%
Electron energy scale	3.9-5.3%	3.9-5.8%	3.6-4.5%
Electron resolution	3.9-5.3%	3.9-5.8%	3.5-4.5%
Electron trigger eff.	0.6%	0.6%	0.6%
Electron isolation eff.	0.6%	0.5-0.6%	0.5-0.6%
Electron trigger scale factor	0.6%	0.5-0.6%	0.5-0.6%
Electron reconstruction eff.	0.0-0.6%	0.1-0.6%	0.1-0.6%
Electron ID eff.	0.0-0.6%	0.0-0.6%	0.0-0.6%
Muon energy scale	3.9-5.2%	3.8-5.7%	3.6-5.5%
Muon resolution	3.8-5.2%	3.8-5.7%	3.5-5.7%
Muon momentum scale	3.8-5.2%	3.8-5.7%	3.5-5.5%
Muon trigger eff.	0.4-0.7%	0.3-0.7%	0.2-0.7%
Muon isolation eff.	0.4-0.6%	0.3-0.6%	0.2-0.5%
Muon reconstruction eff.	0.2-0.5%	$< 0.1\%$	$< 0.1\%$
Signal statistics	1.1-3.4%	0.7-2.1%	0.5-1.8%
	Background systematics		
Background statistics	1.0%		
Event mixing time dependence	$< 0.1\%$		
Event mixing pile-up distribution	$< 0.1\%$		
Event mixing single-sided	$< 0.1\%$		

Table 6.5: Summary of all considered systematic uncertainties for the signal and background models, and their effects on the event yields of the corresponding samples. For the signal samples each generator is shown separately, and the range of absolute observed changes in event yield across all generated signal masses is given.

survival factor and track veto signal efficiency. Additionally, several of the AFP related systematics are very significant, although their effect is very mass dependent, with very wide ranges of effects shown. This occurs because these systematic variations cause the peak of the missing mass distribution to shift away from the generated signal mass, causing a significant change in the number of events falling within the signal region, particularly for very low or high signal masses.

6.8 Statistical Analysis

With the final signal selection applied, the validated data-driven background model can be combined with each signal model to determine the prediction of the event yield for the corresponding signal normalisation. This can then be fitted to the data to determine the most likely value of the signal normalisation, in addition to the highest value which would be compatible with the observed data to a given confidence level, corresponding to upper limits on the signal cross section. The following section describes the statistical methods applied to perform fits and extract cross section limits in the analysis, with the results presented in Section 6.9.

6.8.1 Likelihood

Particle interactions are stochastic processes, meaning that an individual interaction cannot be predicted, but over a large enough sample size the total number of interactions can be predicted to some precision. This prediction follows the discrete Poisson distribution, which describes the probability P of observing k events given an expectation of λ as

$$P\left(k|\lambda\right) = \frac{\lambda^k e^{-\lambda}}{k!}. \tag{6.4}$$

In the context of this analysis, the background model and a given signal model give predictions for the number of background and signal events N_b and N_s respectively, given some initial assumption on the rates of the corresponding processes. So the sum of these values gives the total predicted yield, $\lambda = N_b + N_s$, with the test value k given by the observed number of events in data N_d. Substituting these values into Equation 6.4 gives the probability or "likelihood" of observing N_d events given the prediction of $N_b + N_s$:

$$L\left(N_d|N_b + N_s(\mu_s)\right) = \frac{(N_b + N_s)_d^N e^{-(N_b+N_s)}}{N_d!}. \tag{6.5}$$

We can then find the values of N_b and N_s which maximise this probability given N_d. The expected yield of a given process i within a dataset covering a total integrated luminosity $\mathcal{L}$ can be written as

$$N_i = \mu_i \times A_i \times \mathcal{L} \times \sigma_i, \tag{6.6}$$

where σ_i is the predicted process cross section, A_i is the acceptance factor of the applied event selection to the process and μ_i is the "strength" of the process. When measuring the rate of a particle physics process, it is this strength parameter which is optimised, and it can be thought of as the ratio between the predicted and observed values of the process cross section $\mu_i = \sigma_{i,\mathrm{obs}}/\sigma_i$. Here, σ_i is either the cross section predicted in the SM from the Feynman rules as discussed in Section 2.3, or some arbitrary value close to the expected value of $\sigma_{i,\mathrm{obs}}$ for BSM processes without any realistic prediction. It is the latter which is considered in this analysis. So, when the likelihood function given in Equation 6.5 is maximised, it is the strength parameters of the contributing processes μ_b and μ_s which are allowed to float freely in the fit. The predicted background yield N_b is in fact the sum of all contributing background processes B such that $N_b = \sum_B \mu_B N_B$, with each processes having its own corresponding strength parameter μ_B. In this analysis, there are only two considered contributions, the combinatorial background and the signal-induced background. However, as will be discussed in the next section, a single bin fit is used for the final result, giving only a single observed value of N_d. Therefore, only a single degree of freedom can be included in the fit, which must be the signal strength μ_s as this is the objective measurement or "parameter of interest" in the analysis. Therefore, the strength parameters of both background contributions are fixed to 1. This should not introduce any inaccuracy into the fit results since the combinatorial background is modelled using data and so should already have an accurately predicted normalisation, while the signal-induced background is shown in Section 6.6.2 to be negligible within a mass window such as the one applied in the final fits.

The fit should also take the systematic uncertainties into account, which is done via Nuisance Parameters (NPs) α_p, which parametrise the effect of each systematic on the signal and background yields as $N \to N + \sum_p \alpha_p N_p$, where N_p is the predicted change in the yield caused by each systematic p, and the effects of all systematic variations are summed over. The NPs are taken to be normally distributed as $N(0, 1)$ and included in the likelihood function by multiplying by the normal probability distribution for each parameter, giving the overall likelihood function as

$$L\left(N_d | N_b + N_s(\mu_s)\right) = \frac{(N_b + N_s)_d^N e^{-(N_b+N_s)}}{N_d!} \times \prod_p \frac{1}{\sqrt{2\pi}} e^{-\frac{\alpha_p^2}{2}}, \tag{6.7}$$

where the predicted signal yield is now a function of the parameter of interest μ_s. Typically, the logarithm of the likelihood is instead considered, such that the product of the probability distributions for each NP is converted to a summation, and the resulting maximum becomes independent of any constant scale factors. Additionally, the negative value of log likelihood $-\ln L$ is used so that the function can instead be minimised, which can lead to easier convergence in fits [3, 202, 203].

When making a measurement, the best-fit value of the signal strength is determined as the value corresponding to the minimum of the negative log likelihood. From this best-fit parameter, the resulting cross section measurement can be calculated. However, in many analyses searching for new physics, no signal is found, resulting in a zero or negative best-fit value of signal strength. This suggests that either the considered process does not occur at all, or that it has a sufficiently low cross section such that observing zero events in a dataset of the considered size is consistent with the predicted rate. Assuming the latter scenario, an upper limit can be derived on the cross section, below which the predicted rate is consistent with zero to a certain degree of probability. This is done by varying the signal strength parameter upwards from zero, and determining the corresponding probability from Equation 6.7, referred to as the p-value. The upper limit on μ_s is then taken as the value beyond which the p-value drops below a certain threshold $1 - C$, where C is the Confidence Level (CL) of the limit. The higher the CL, the lower the chances that the true process cross section could be above the determined upper limit. Typically CLs of 95% or 99% are chosen, corresponding approximately to $\pm 2\sigma$ or $\pm 3\sigma$ deviations from the mean in a normal distribution, respectively.

6.8.2 Fit Procedure

The likelihood fit is performed using the TRExFitter package, in which signal and background models and associated systematic variations are included as template histograms through the HistFactory program [204], fits of the likelihood function are performed with RooFit [205] and cross section upper limits are calculated with RooStats [206]. Fits are performed on the missing mass distribution, and initially a binned approach was used applying the 50 GeV bin width across the full mass range $0 \leqslant m_X \leqslant 1000$ GeV as motivated by the studies in Sections 6.2.1 and 6.2.2. However, due to very low statistics in both data and the signal and background models, primarily due to the use of the dilepton p_{T} cut and track veto selection discussed in Section 6.5, a single-binned approach instead had to be adopted. A mass window is defined for each signal model, centred on the hypothesised signal mass, with the size of each window varied depending on the available statistics in each case. In the high-p_{T} track veto signal region, the default mass window is set to 100 GeV either side of the hypothesised signal mass for most mass points (e.g. $400 \leq m_X \leq 600$ GeV for a 500 GeV signal model), but asymmetric windows are used for the 100 GeV ($0 \leq m_X \leq 300$ GeV) and 900 GeV ($700 \leq m_X \leq 1000$ GeV) models, due to low statistics in these mass regions. These binnings are summarised in Table 6.6, alongside the binnings applied in the low+high-p_{T} signal region, which use a wider default bin width of 150 GeV either side of the signal mass due to

Mass [GeV]	High-p_{T} track veto	Low+high-p_{T} track veto
100	0,300	0,500
200	100,300	0,500
300	200,400	100,500
400	300,500	250,550
500	400,600	350,650
600	500,700	450,750
700	600,800	550,850
800	700,900	650,950
900	700,1000	700,1000

Table 6.6: Binning used for signal models with each considered hypothesised signal mass, in each signal region. Wider bins are used in the low+high-p_{T} track veto signal region to compensate for reduce statistics.

lower statistics, and were similarly optimised in the low and high mass regions. The nominal signal cross-section predictions calculated by the respective generators were not applied, since they were rough estimates lacking physical motivation, typically falling well below the sensitivity achievable by this analysis. Instead, the simulated signal samples were normalised to arbitrary cross-sections between 10 fb and 1 pb before the fit. These values were chosen for each model such that the post-fit value of μ_s is close to 1, in order to help the fits to converge.

Generally, systematic variations are expected to have the opposite effect depending on which direction the variation is applied in. For example, varying the lepton trigger efficiency downwards by its uncertainty reduces the observed signal yield, while varying it upwards increases the yield. However, some systematic variations such as the optics parametrisation in AFP reconstruction and the momentum scale of muons, were observed to have the same effect on the yield regardless of which direction the variation was applied in. This is expected in some situations, for example if the nominal value of a parameter is optimised by finding the maximum of some related distribution, then a similar deviation is expected when varying that parameter in either direction of the optimal value. However, this causes issues for the fit, which assumes that the nominal prediction falls in between the two variations. Therefore, in these situations symmetrisation was applied, which used the mean of the absolute values of the upwards and downwards variations and applied the corresponding shift directly to the nominal distribution in each direction. Systematic uncertainties which had below a 0.5% effect on event yield were removed from the fit. This is called pruning, and is necessary to avoid overfitting. This typically included most systematics parametrising uncertainties in scale factors applied to correct lepton reconstruction and trigger efficiencies in the simulation.

As presented in the next section, no significant deviation from the background-only hypothesis was observed, and so upper limits were instead set on the signal cross-section for each model. The limits were determined at a 95% CL using the CL_s method, described in [207], by applying the asymptotic formulae used to simplify limit computation [203]. Limits were set in each lepton channel separately, and for a combination of both channels.

6.9 Results

The section presents the results of applying the statistical procedure outlined in Section 6.8 to the analysis. The final distributions and upper limits on the signal cross section for each considered model are shown, with a comparison to the results from the equivalent analysis performed by CMS [185].

6.9.1 Fit to Blinded Data

First, the statistical procedure was tested on the blinded data sample described in Section 6.3.1, to ensure that the cross section limits obtained were consistent with expectations and no anomalies were observed. Figure 6.35 shows a comparison in the signal region between the missing mass distributions in the blinded dataset and the data-driven background model. The distributions in both lepton channels agree to within the expected degree of statistical fluctuation, as is expected since the blinded dataset is created using the same event-mixing procedure as the background, just with a different value of event-shift i. The resulting blinded upper limits on the

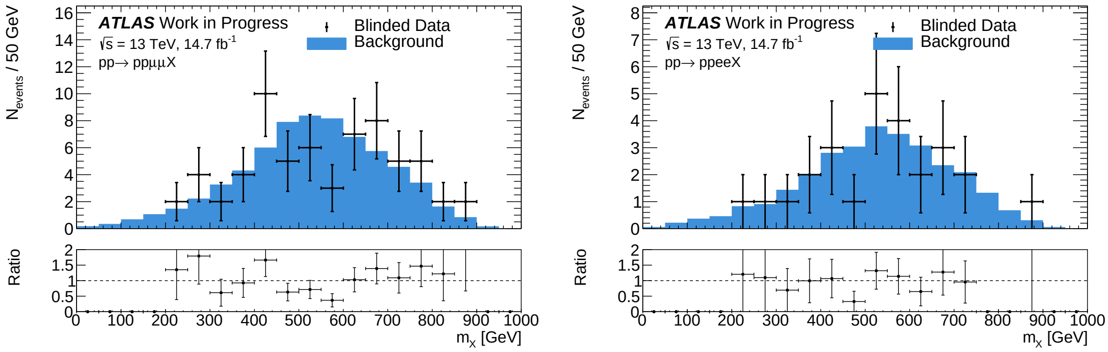

(a) Muon channel (b) Electron channel

Figure 6.35: Comparison of the missing mass distribution between blinded data and the data-driven background model in the (a) muon and (b) electron channels with all signal selections applied.

signal cross section for the SUPERCHIC $Z + X$ model are shown in Figure 6.36. The results follow the expected smooth distribution, with no significant deviation

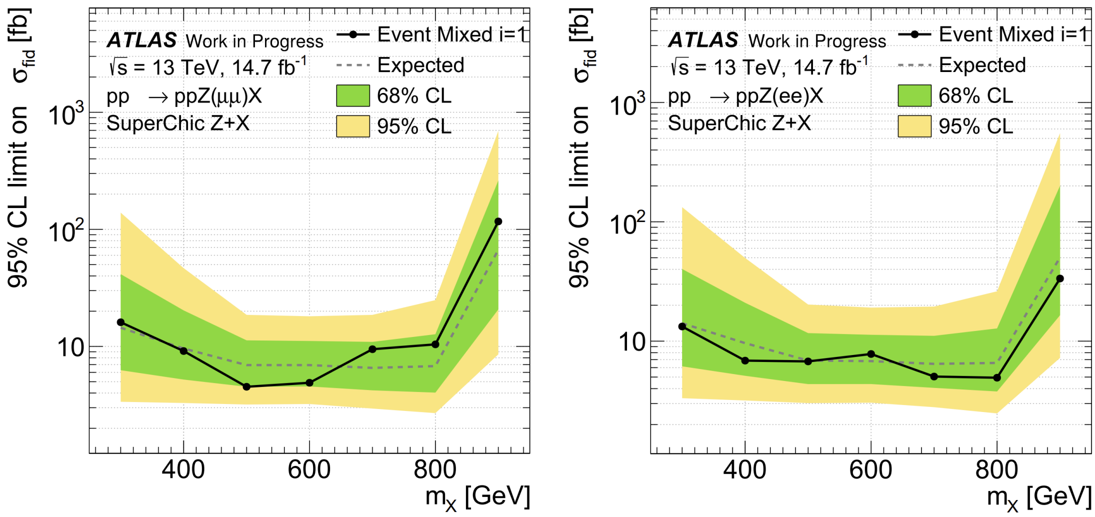

(a) Muon channel (b) Electron channel

Figure 6.36: Upper limits on signal cross section set using blinded data with an event-shift of $i = 1$ for the SUPERCHIC $Z+X$ signal model with the standard track veto applied in the (a) muon and (b) electron channels.

observed from the expected limits.

On the basis of the above results, the analysis procedure was validated and the unblinded dataset (corresponding to an event-shift of $i = 0$, i.e. no shifting applied) could be studied directly with the same procedure.

6.9.2 Fit to Data and Cross Section Upper Limits

Before performing any fits, the data and background model were compared at various stages of the signal selection process, to check for any disagreement which could indicate the presence of signal. Figure 6.37 shows the comparison for key kinematic distributions with only the pre-selection applied, representing a basic search space where no deviation is expected. Each lepton channel is considered separately, with

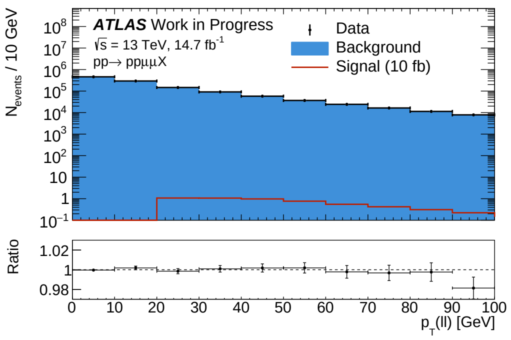

(a) Dilepton p_{T}

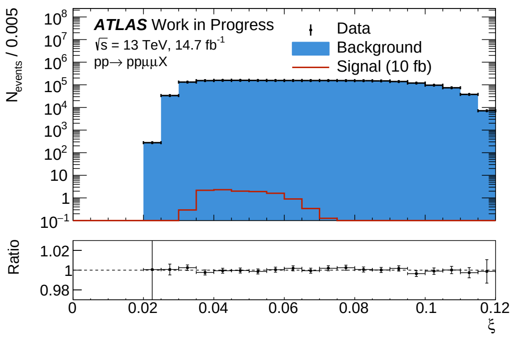

(b) Proton ξ

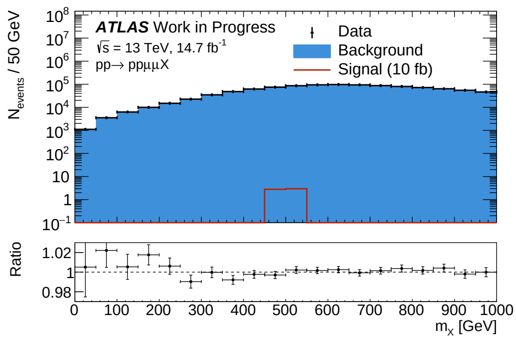

(c) Missing mass m_X

Figure 6.37: Comparison between unblinded data and the data-driven background model in the muon channel with only pre-selection cuts applied for distributions of (a) dilepton pair p_{T}, (b) proton ξ and (c) missing mass m_X. The expectations for a signal with a hypothesised mass of $m_X = 500$ GeV using the SuperChic $Z + X$ model are overlaid and normalised to a cross section of 10 fb.

the muon channel shown as an example in Figure 6.37. Good agreement is observed between the distributions as expected, with a similar level of agreement present in the electron channel. The expectations for a signal with a hypothesised signal mass of 500 GeV using the SuperChic $Z + X$ model is overlaid, normalised to a fiducial cross section of 10 fb. Similar levels of agreement are observed when each signal selection is applied on top of the pre-selection, with the dilepton mass and p_{T} cuts, the tightened proton ξ cut and the track veto selection each tested separately.

Figure 6.38 shows the same comparison with all signal selections applied for the high-p_{T} signal region. These distributions are taken from a combination of both

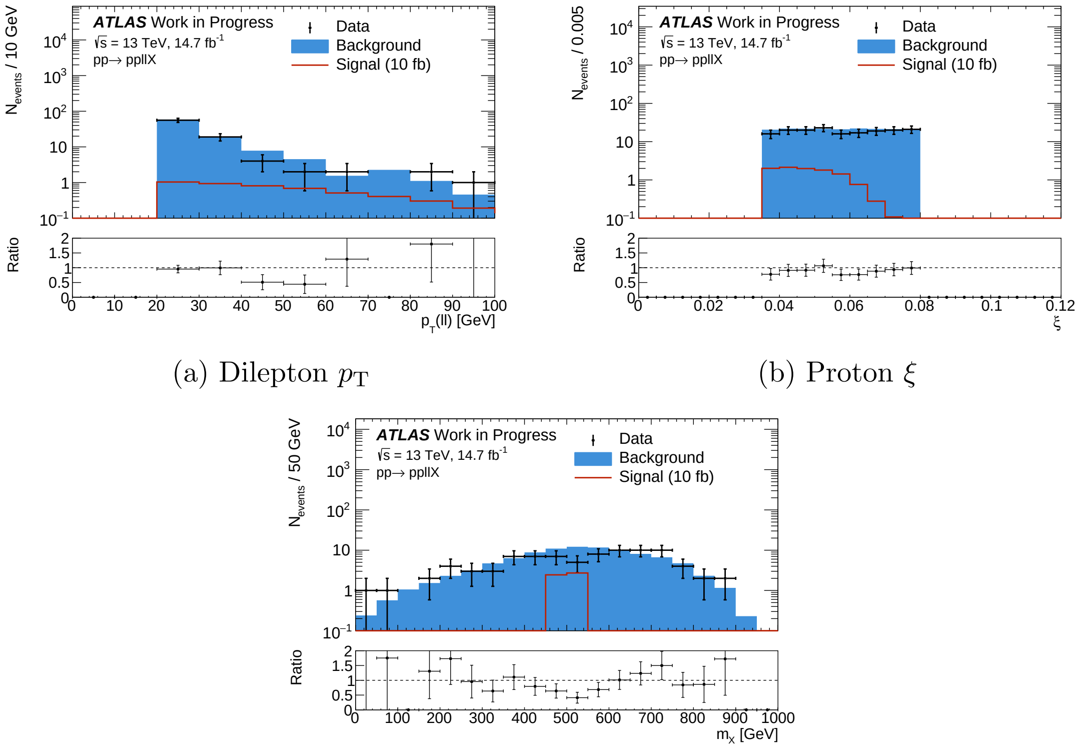

(a) Dilepton p_{T}

(b) Proton ξ

(c) Missing mass m_X

Figure 6.38: Comparison between unblinded data and the data-driven background model in the combined lepton channel with the pre-selection and all signal region cuts applied, for distributions of (a) dilepton pair p_{T}, (b) proton ξ and (c) missing mass m_X. The expectations for a signal with a hypothesised mass of $m_X = 500$ GeV using the SuperChic $Z + X$ model are overlaid and normalised to a cross section of 10 fb.

lepton channels, to give more statistics for the comparison. Again, generally good agreement is observed within the expected level of statistical fluctuation, suggesting that no signal is present.

The single-bin fit was then performed within a mass window for each signal model, as detailed in Section 6.8.2. The pre and post-fit signal event yields for the high-p_{T} track veto signal region are plotted alongside the background model and observed data for each signal model and mass in Figure 6.39. It can be seen that around

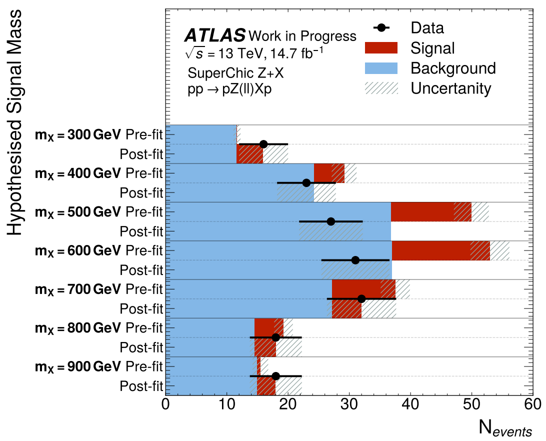

(a) SuperChic $Z + X$

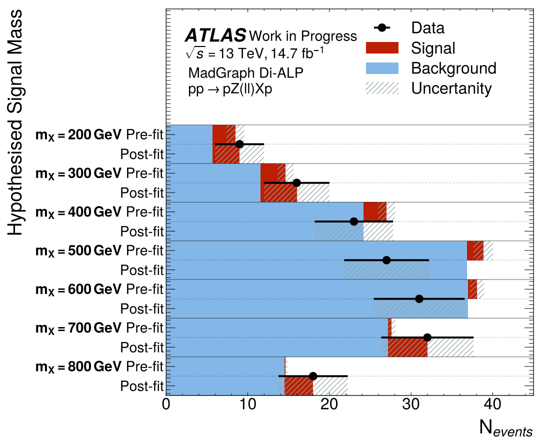

(b) MadGraph di-ALP

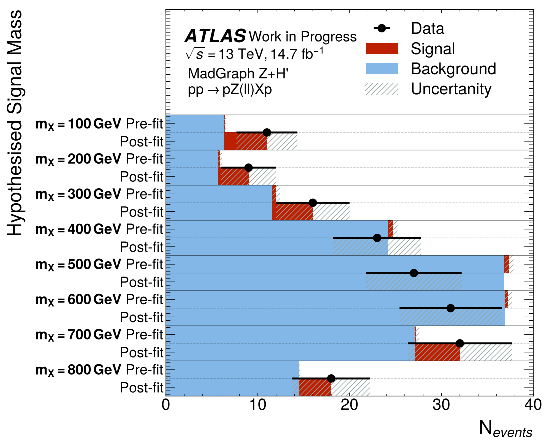

(c) MadGraph $Z + H'$

Figure 6.39: Summary of all pre and post-fit distributions for the signal+background model for each tested signal mass, in the combined lepton channel, for the (a) SuperChic $Z + X$, (b) MadGraph di-ALP and (c)MadGraph $Z + H'$ models. Fits use a mass window of 100 GeV either side of the hypothesised signal mass, with the exception of 100 GeV and 900 GeV models, which use larger windows of $0 \leq m_X \leq 300$ GeV and $700 \leq m_X \leq 1000$ GeV, respectively. All pre-fit signals are shown normalised to 25 fb.

half the time there is room between the background prediction and observed data to allow for a small signal following the fit. However, in as many cases the number of observed events in data falls below the background prediction, resulting in a negative best-fit signal strength. Therefore, any potential signal is concluded to be just the result of statistical fluctuation in the observed dataset, with no observed deviations above $\pm 2\sigma$ from the predicted event yield.

The effect of systematic uncertainties on the final result is evaluated in several ways. First, the best-fit value and associated uncertainty of each corresponding nuisance parameter is determined, as shown in Figure 6.40 for all systematics which are not pruned, for the 500 GeV mass point in each signal model, for the combined lepton channel. All best-fit nuisance parameters follow the expected distribution very

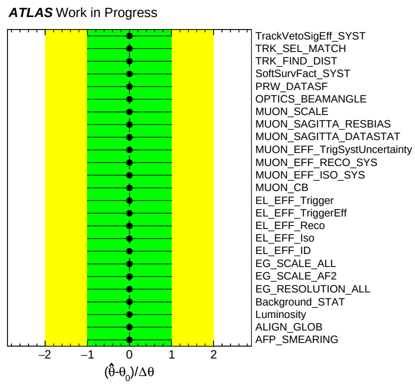

(a) SuperChic $Z+X$ combined channel

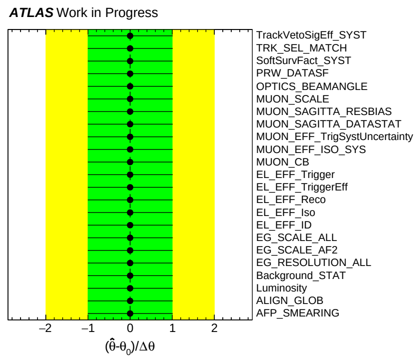

(b) MadGraph di-ALP combined channel

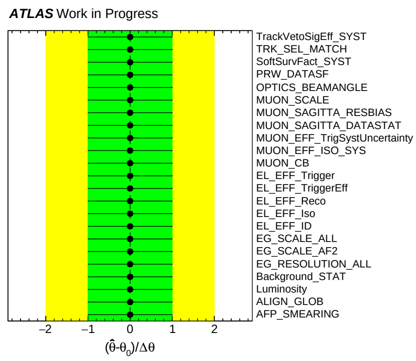

(c) MadGraph $Z+H'$ combined channel

Figure 6.40: Systematic pulls for the 500 GeV signal for each signal model in each channel in the low+high-p_{T} track veto signal region.

closely, with mean values of 0 and variance of 1, suggesting that the corresponding effects have a very small impact on the final fit. This is expected since most systematics considered in the fit only affect the signal prediction, which has a small contribution to the overall yield. The impact of each systematic on the uncertainty of the post-fit signal strength is evaluated using "ranking plots", as shown in Figure 6.41 for the same signal models considered above. These show the difference between the nominal post-fit value of signal strength and that obtained with a given nuisance parameter varied by its corresponding uncertainty in each direction, both

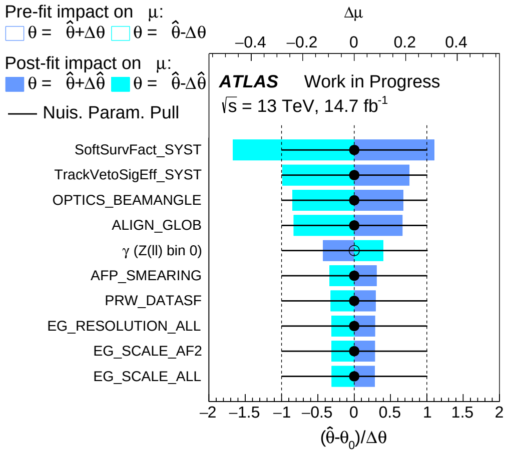

(a) SuperChic $Z + X$ combined channel

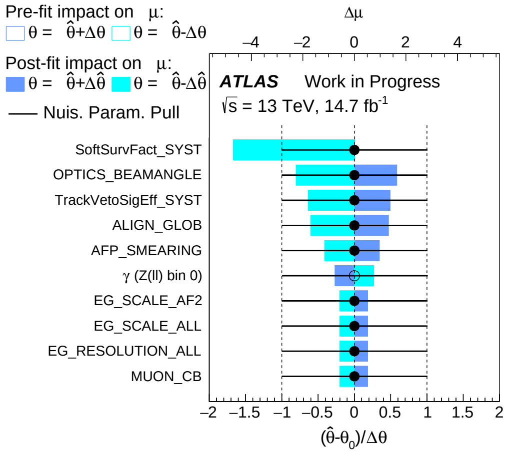

(b) MadGraph di-ALP combined channel

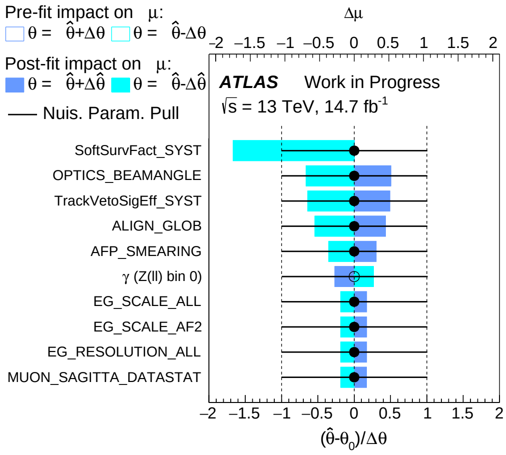

(c) MadGraph $Z + H'$ combined channel

Figure 6.41: Systematic ranking plots for each signal model with $m_X = 500$ GeV, in the combined channel. The highest ranked systematics are those having the largest impact on the final value of signal normalisation, with the soft-survival uncertainty, track veto signal efficiency and AFP spectrometer optics and alignment consistently being the highest ranked. In the case of MadGraph models, the soft-survival uncertainty is applied as a downwards-only variation, resulting in an asymmetric pull.

before and after the fit is performed. The resulting impacts are then ordered, with the top ten highest impact systematics shown on the plots. The results match with expectations based on the sizes of the corresponding uncertainties, with the highest ranked soft-survival uncertainty set at 20% for SuperChic models and around 27% for the MadGraph models shown here, as discussed in Section 6.7.4. The track veto signal efficiency uncertainty is consistently the second-highest ranked, set at 12.9% for the combined lepton channel as discussed in Section 6.7.1. The AFP optics parametrisation and global alignment uncertainties are highly ranked, as expected due to the sensitivity of the proton reconstruction to these variables, and the conservative global alignment uncertainty of 300μm, as discussed in Section 4.4.2. Finally, the statistical uncertainty of the signal model $\gamma(Z(\ell\ell)$ bin 0) is consistently ranked in the top five, due to the low selection efficiency resulting from the tight signal selection which was adopted. This could be mitigated by generating larger signal samples, although the computational and time costs would be extensive.

The upper limits set on the signal cross section for all tested mass points in each signal model are shown in Figure 6.42 for the high-p_{T} track veto signal region. Each lepton channel is shown individually and in combination for each signal model. Again, no significant excess is observed over the predicted background yield, with limits set on the fiducial signal cross-section on the order of 10 fb for the SuperChic $Z+X$ model and MadGraph di-ALP models, and on the order of 100 fb for the MadGraph $Z+H'$ model. Flat limits are obtained despite the large variation in signal selection efficiency as a function of mass, due to the application of the fiducial volume to which the signals are normalised.

6.9.2.1 High and Low-p_{T} Track Veto Comparison

During the analysis, the intention was to set limits in both the high-p_{T} and low+high-p_{T} track veto signal regions, with the low+high-p_{T} track veto expected to yield improved sensitivity due to the increase in signal to background ratio, as discussed in Section 6.5.7.3. However, due to extremely small statistics in both the data and signal and background models following the inclusion of low-p_{T} tracks in the veto, the fits were found to be very unstable, requiring the sizes of the mass windows used for fits to be significantly increased, as recorded in Table 6.6. As a result of this, the corresponding limits lost sensitivity, with the ratio between the limits obtained in each signal region shown in Figure 6.43. Ratios above one here indicate an improvement in sensitivity when low-p_{T} tracks are included. While small gains of around 10% can be seen for some models and mass assumptions, several models

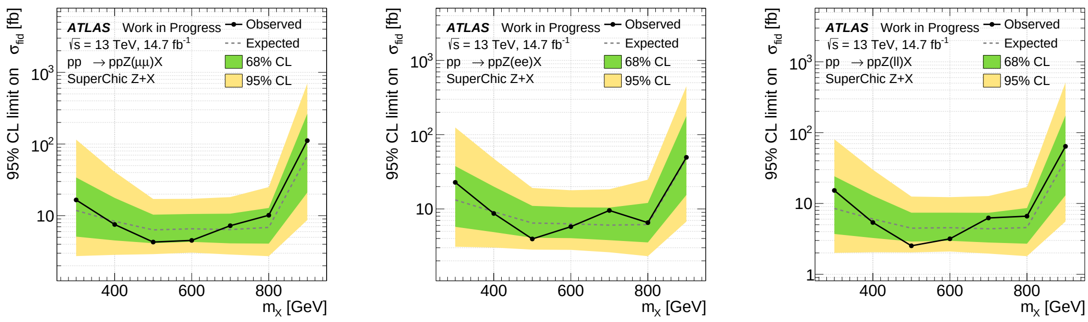

(a) SuperChic $Z + X$ muon channel (b) SuperChic $Z + X$ electron channel (c) SuperChic $Z + X$ combined channels

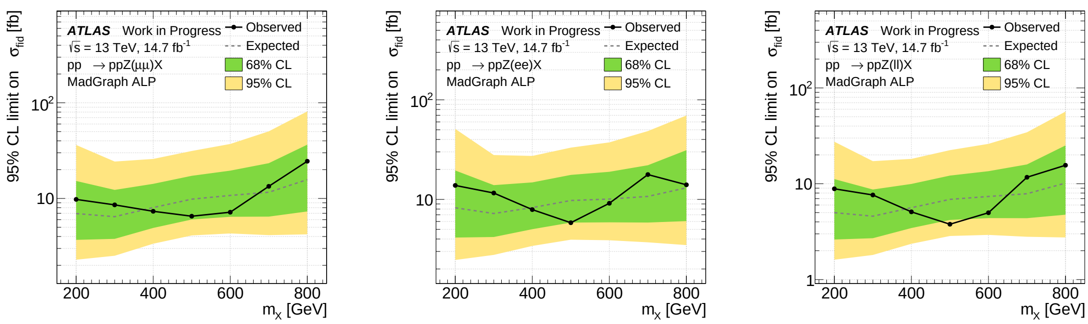

(d) MadGraph di-ALP muon channel (e) MadGraph di-ALP electron channel (f) MadGraph di-ALP combined channels

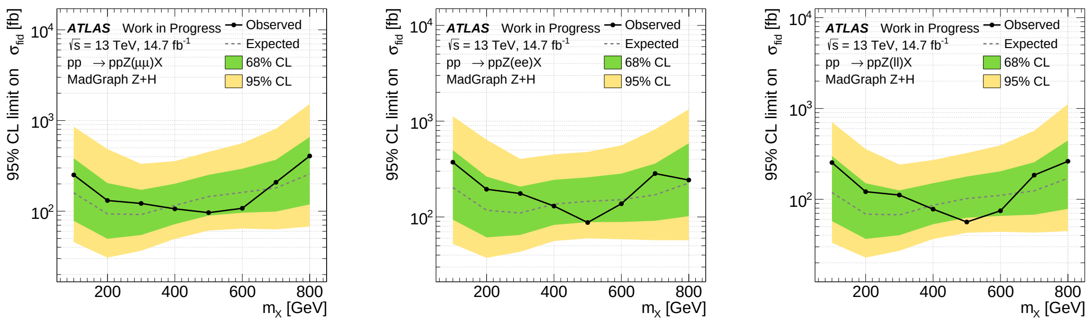

(g) MadGraph $Z+H'$ muon channel (h) MadGraph $Z + H'$ electron channel (i) MadGraph $Z + H'$ combined channels

Figure 6.42: Observed and expected upper limits on the signal cross section set for each signal model with the standard track veto applied, with lepton channels shown separated and combined, and with all systematics and scale factors included.

lose significant sensitivity. The 100 GeV MadGraph $Z + H'$ model is particularly heavily effected in this way, as it required the largest increase in mass window size and so loses almost half the previously observed sensitivity. For this reason, it was decided to discard the low+high-p_{T} signal region from the final results, and instead use only the standard high-p_{T} track veto.

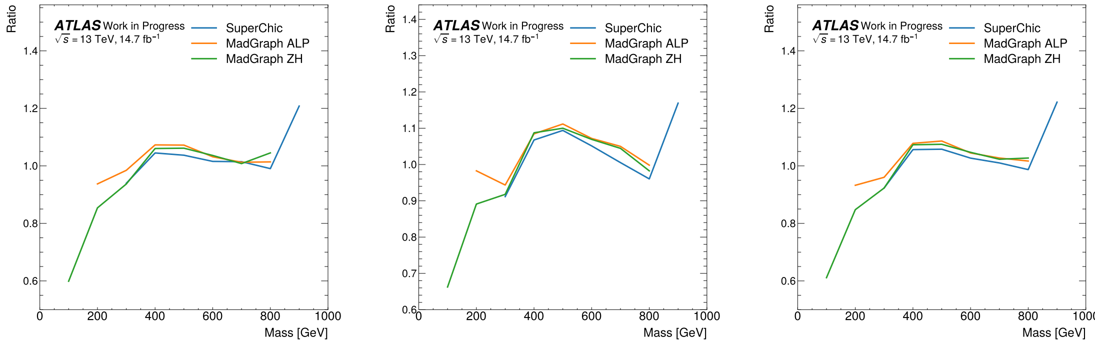

(a) Muon channel (b) Electron channel (c) Combined channel

Figure 6.43: Ratio of the expected limits obtained using only high-p_{T} tracks ($p_{\mathrm{T}} >$ 500 MeV) to those obtained with both high- and low-p_{T} tracks ($p_{\mathrm{T}} >$ 100 MeV) included in the track veto. Ratios greater than one indicate improved sensitivity when low-p_{T} tracks are included. Results are shown separately for the (a) muon, (b) electron, and (c) combined channels.

To investigate what level of improvement could be expected from the low+high-p_{T} track veto in a future analysis with more statistics, the high-p_{T} track veto results were recalculated using the wider mass windows used for the low+high-p_{T} track veto signal region in Table 6.6, to remove the effect of the different binning on the resulting limit ratio. The updated ratios are shown in Figure 6.44, and now a consistent 20-30% improvement in sensitivity in observed across all signal models, demonstrating the effectiveness of the inclusion of low-p_{T} track for a future Run 3 version of this analysis, which would have around 10 times higher statistics based on current estimates.

6.9.3 CMS Comparison

The results shown above were compared to those obtained in the equivalent CMS analysis published in 2022 [185], to determine the improvement achieved by this analysis, primarily due to the addition of the track veto selection. The comparison is made only for the SuperChic $Z + X$ signal model, which was designed to match the model employed by CMS, although this model was only simulated to generator-level, with acceptance and efficiency scale factors applied for a simplified measure of detector response. The SuperChic $Z + X$ and CMS models were compared

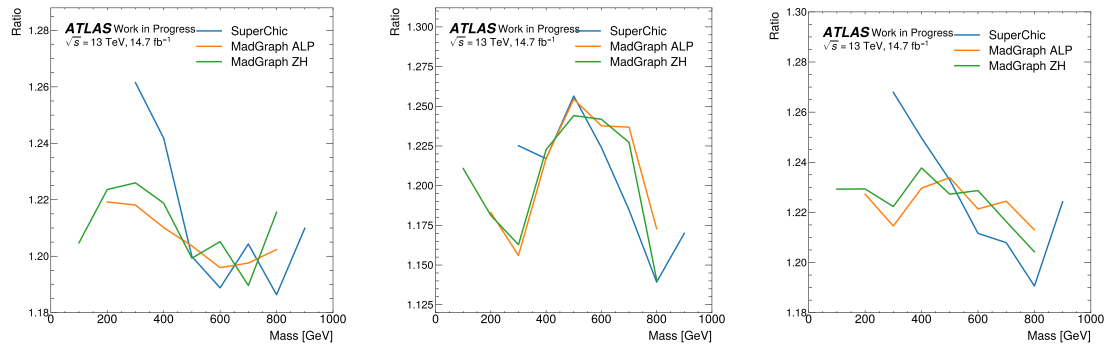

(a) Muon channel (b) Electron channel (c) Combined channel

Figure 6.44: Ratio between the expected limits obtained obtained with the high and low+high-p_{T} track vetoes applied, with both results using the low+high-p_{T} track veto binning. Ratios above 1 indicate an improvement with the low+high-p_{T} track veto applied.

Feature	Criterion
Leptons	$\geqslant 2$ same-flavour leptons (e or μ), opposite electric charge $p_{\mathrm{T}}(\ell_1) > 30$ GeV, $\lvert\eta\rvert(\ell_1) < 2.4$ $p_{\mathrm{T}}(\ell_2) > 20$ GeV, $\lvert\eta\rvert(\ell_2) < 2.4$
Dilepton	$\lvert m_{\ell\ell} - m_Z\rvert < 10$ GeV $p_{\mathrm{T}}^{\ell\ell} > 40GeV$
Protons	$0.02 < \xi_A < 0.16$ and $0.03 < \xi_C < 0.18$

Table 6.7: Summary of fiducial volume selection criteria for signal events in the CMS missing mass analysis [185].

at generator-level and found to match very closely, validating a direct comparison. Furthermore, CMS used a different fiducial volume to the one defined for this analysis in Table 6.4, due to their differing selection strategy and detector acceptance. The CMS fiducial selection is given in Table 6.7. Soft-survival factors are omitted from the SuperChic $Z + X$ model for this comparison, since these were not considered by CMS.

Figure 6.45 shows the comparison of the obtained fiducial cross section upper limits between the two analyses. This analysis is observed to have improved sensitivity by up to a factor of ten for most of the mass points which are common between the two results, despite the CMS analysis using a dataset with a total integrated luminosity of 37.2 fb^{-1}, around 2.5 times larger than the dataset used in this analysis.

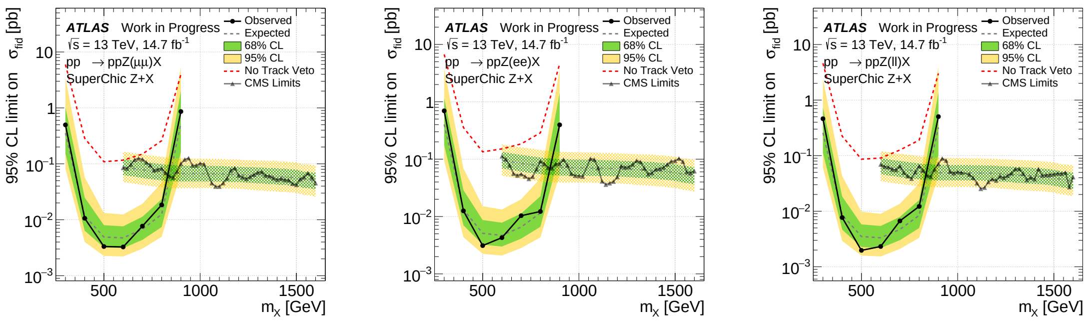

(a) Muon channel (b) Electron channel (c) Combined channel

Figure 6.45: Comparison of the observed limits for the SuperChic $Z + X$ model between this analysis and the CMS results from [185]. The limits from this analysis are scaled to match the fiducial region used in the CMS study, and soft-survival factors are not included. The expected limits obtained in this analysis with no track veto selection are overlaid. The comparison is shown for the (a) muon, (b) electron, and (c) combined channels.

The shape of the limits is very different to that observed in Figure 6.42 due to the different fiducial volume, which is less complementary to our signal selection, leading to relative sensitivity loss at tail masses. The CMS analysis was able to probe hypothesised signal masses up to 1600 GeV due to the PPS, their equivalent forward detector, having much high acceptance of up to $\xi < 0.20$. However, this analysis considered signal masses significantly below the minimum of 600 GeV considered by CMS, giving good complementarity between the two results. The estimated results from this analysis obtained without using the track veto selection are overlaid on Figure 6.45, emphasising the gain in sensitivity achieved with this selection, by around a factor of ten across all signal masses, allowing the relative improvement compared to the CMS analysis.

Chapter 7

Conclusion

This thesis has presented work carried out towards the author's PhD, on both an analysis of ATLAS experiment data and the QA program for the ATLAS ITk upgrade.

A search was performed for new physics in photon-induced dilepton production events, using 14.7 fb^{-1} of data collected at $\sqrt{s} = 13$ TeV by the ATLAS detector and the AFP spectrometer. The inclusion of the AFP spectrometer allowed a first-of-its-kind measurement in the ATLAS collaboration of the total missing mass in each event by comparing the energy loss of scattered protons and the resulting momenta of the visible products. Resonance searches performed in the missing mass distribution allowed upper limits to be set on the cross sections of three potential signal models, including a model probing the theorised ALP, a possible candidate for dark matter. Signal masses between 300 and 900 GeV were considered, limited by the acceptance of the AFP spectrometer.

One of the signal models considered follows the same model employed by an equivalent analysis performed by CMS [185], allowing a direct comparison between the results to be made. This analysis showed large improvement in sensitivity across the majority of common mass points between the analyses, primarily due to the addition in this analysis of a track veto selection, removing events with additional inner detector tracks reconstructed close to the dilepton vertex. This selection yielded around 99.9% background reduction with only a 60% reduction in signal, resulting in a factor of ten increase in sensitivity. The uncertainties on the expected cross section upper limits are dominated by statistical effects, with the largest systematic effects originating from estimates of soft-survival probability and track veto signal efficiency, along with AFP spectrometer alignment and proton reconstruction uncertainties.

An extended track veto was investigated including low-p_{T} tracks with $100 < p_{\mathrm{T}} < 500$ MeV not included in the standard ATLAS ID track reconstruction. A recently developed dedicated reconstruction step [199] was used to obtain these objects in data, which were estimated to offer up to a factor of five times improvement in background rejection, with only 20% additional signal loss. However, due to extremely low statistics in the final signal region with this selection, the expected level of improvement was not observed, with the standard track veto instead presented as the main result. The potential gains for a future higher-statistics analysis from low-p_{T} tracks were estimated at between 20-30%, and should therefore be considered for a Run 3 extension of this analysis which will have approximately 10 times more data.

Appendix A

Background Validation with Alternative Model

In order to ensure that the data-driven background model correctly accounts for all expected contributions to the combinatorial background, an alternative background model was produced using simulation. All of the contributions to the central dilepton background listed in Section 6.6 were included using MC samples, except for the misidentified lepton contribution which was modelled using a same-sign lepton selection in data.

A.1 Simulated Samples

The simulated samples used to model the majority of the central background processes, and their respective generators, are summarised in Table A.1. These simulated samples were overlaid with pile-up protons randomly sampled from a database produced from 2017 data, to simulate the proton component of the combinatorial background. Several alternative samples simulating Z+jet production were investigated, as shown in Table A.1, to study the mismodelling of the underlying event and its effect on the track veto efficiency, as discussed in Section A.4.

A.2 Misidentified Leptons

Misidentified leptons arise from errors in reconstruction, typically when an object such as a jet is misidentified as a lepton. There is also a contribution from genuine leptons which are measured to have the wrong charge, referred to as charge flipping, although this is only expected to occur for high-p_{T} electrons which leave straight

Process	Generator	UEPS	Slice/Filter
$Z (\to ee)$ +jets $Z (\to \mu\mu)$ +jets	Sherpa v2.2.14	Sherpa v2.2.14	$p_{\mathrm{T}} > 3.5$ GeV, $\|\eta\| < 2.7$, $N_\ell >= 2$, $m_{\ell\ell} > 40$
$Z (\to ee)$ +jets (alternative) $Z (\to \mu\mu)$ +jets (alternative)	Powheg-Box v4	Pythia v8.244	$p_{\mathrm{T}} > 3.5$ GeV, $\|\eta\| < 2.7$, $N_\ell >= 2$
$Z (\to ee)$ +jets (alternative) $Z (\to \mu\mu)$ +jets (alternative)	Powheg-Box v5	Herwig v7.2.1	$p_{\mathrm{T}} > 3.5$ GeV, $\|\eta\| < 2.7$, $N_\ell >= 2$
$t\bar{t}$	Powheg-Box v4	Pythia v8.244	$p_{\mathrm{T}} > 3.5$ GeV, $\|\eta\| < 2.7$, $N_\ell >= 2$
Wt	Powheg-Box v4	Pythia v8.244	$p_{\mathrm{T}} > 3.5$ GeV, $\|\eta\| < 2.7$, $N_\ell >= 2$
$W\bar{t}$	Powheg-Box v4	Pythia v8.244	$p_{\mathrm{T}} > 3.5$ GeV, $\|\eta\| < 2.7$, $N_\ell >= 2$
$VV \to \ell\ell\nu\nu$	Sherpa v2.2.14	Sherpa v2.2.14	$p_{\mathrm{T}} > 3.5$ GeV, $\|\eta\| < 2.7$, $N_\ell >= 2$
$VV \to \ell\ell\ell\nu$	Sherpa v2.2.14	Sherpa v2.2.14	$p_{\mathrm{T}} > 3.5$ GeV, $\|\eta\| < 2.7$, $N_\ell >= 2$
$VV \to \ell\ell\ell\ell$	Sherpa v2.2.14	Sherpa v2.2.14	$p_{\mathrm{T}} > 3.5$ GeV, $\|\eta\| < 2.7$, $N_\ell >= 2$
$\gamma\gamma \to \ell\ell$ (EE)	MadGraph v2.9.5	Pythia v8.245	$p_{\mathrm{T}} > 3.5$ GeV, $\|\eta\| < 2.5$, $N_\ell >= 2$
$\gamma\gamma \to \ell\ell$ (DS)	MadGraph v2.9.5	Pythia v8.245	$p_{\mathrm{T}} > 3.5$ GeV, $\|\eta\| < 2.5$, $N_\ell >= 2$
$\gamma\gamma \to \ell\ell$ (SD)	MadGraph v2.9.5	Pythia v8.245	$p_{\mathrm{T}} > 3.5$ GeV, $\|\eta\| < 2.5$, $N_\ell >= 2$
$\gamma\gamma \to WW$ (EE)	MadGraph v2.9.5	Pythia v8.245	$p_{\mathrm{T}} > 3.5$ GeV, $\|\eta\| < 2.5$, $N_\ell >= 2$
$\gamma\gamma \to WW$ (DS)	MadGraph v2.9.5	Pythia v8.245	$p_{\mathrm{T}} > 3.5$ GeV, $\|\eta\| < 2.5$, $N_\ell >= 2$
$\gamma\gamma \to WW$ (SD)	MadGraph v2.9.5	Pythia v8.245	$p_{\mathrm{T}} > 3.5$ GeV, $\|\eta\| < 2.5$, $N_\ell >= 2$

Table A.1: Overview of the simulated samples used to model the central dilepton component of the combinatorial background for the analysis, the corresponding programs used to perform the generation and UEPS simulation steps, and applied generator-level filters.

tracks in the ID making charge determination difficult, whereas muons leave additional tracks in the MS which remove any ambiguity in charge determination. An adjusted event selection searching for same-sign lepton pairs is effective for estimating the misidentified lepton contribution, since such a final state is not expected for any SM process, and therefore must arise from misidentified leptons. This method for determining the misidentified lepton contribution to background contributions is also used in the ATLAS collaboration analysis of photon-induced dilepton production using the AFP spectrometer [85]. Following the pre-selection, around 25000 same-sign events are present in the dataset (0.7% as many as the regular selection), with around 95% in the electron channel. With all signal selections except for the track veto applied this drops to around 4000 events. The kinematic distributions in each lepton channel are compared in Figure A.1. There is a clear Z boson resonance

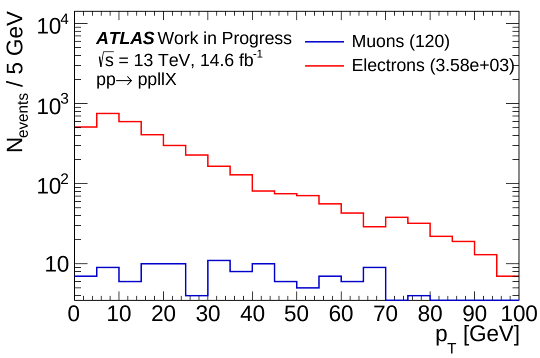

(a) Dilepton p_T

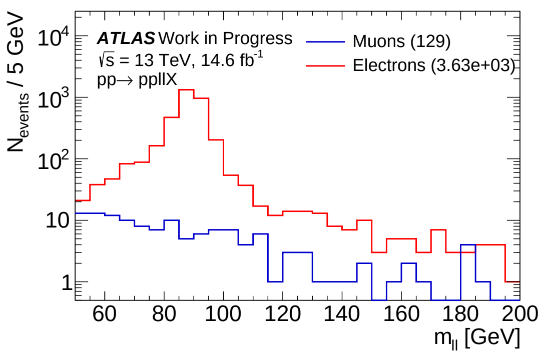

(b) Dilepton mass

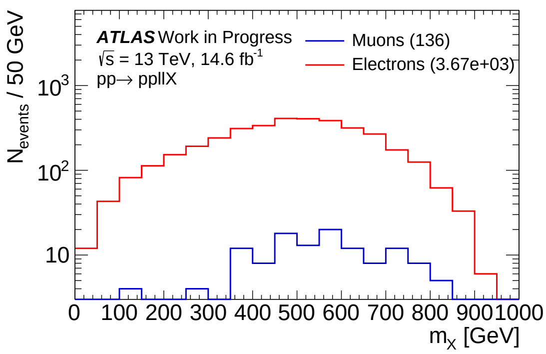

(c) Missing mass

Figure A.1: Kinematic distributions after signal selection, except for track veto, from a misidentified lepton background model created using a same-sign lepton selection in data.

in the electron channel occurring due to charge flipping, which is not present in the muon channel. Following the application of the track veto, zero same-sign events remain, since the exclusivity requirement removes events with any additional tracks close to the leptons, which are expected from misidentified jets.

A.3 Exclusive SM Processes

Two of the lowest cross section processes considered for the combinatorial background are the photon-induced production of leptons $\gamma\gamma \to \ell\ell$ and of leptonically decaying W bosons $\gamma\gamma \to WW$. Following the pre-selection, there are $\mathcal{O}(100)$ $\gamma\gamma \to \ell\ell$ events remaining, whereas only $\mathcal{O}(1)$ $\gamma\gamma \to WW$ events remain, due to the lower process cross section, in addition to the low branching ratio for both W bosons to decay to the same lepton flavour (around 1%). Therefore, the WW component is negligible relative to the dilepton component, so only the latter is considered further here.

Despite their relative rarity compared to the other processes, particularly Z+jets production, photon-induced dilepton events require special consideration. Since the central dilepton system is produced via CEP through double photon exchange, no additional central tracks are produced and so this process is only affected by the track veto criterion to the same extent as the signal process. Therefore this background becomes more significant than most of the higher rate central contributions following the application of this requirement. In addition, since the quantum numbers of the protons are not changed in this process, there is the possibility of the interacting protons remaining intact. As discussed in Section 2.3, such a process can have three channels depending on the final state of the interacting protons: EE, SD and DD. These three channels are illustrated for photon-induced dilepton production in Figure A.2, with the sum of the EE and SD cases recently measured using AFP data [85]. In the case of double dissociation, neither proton remains intact and

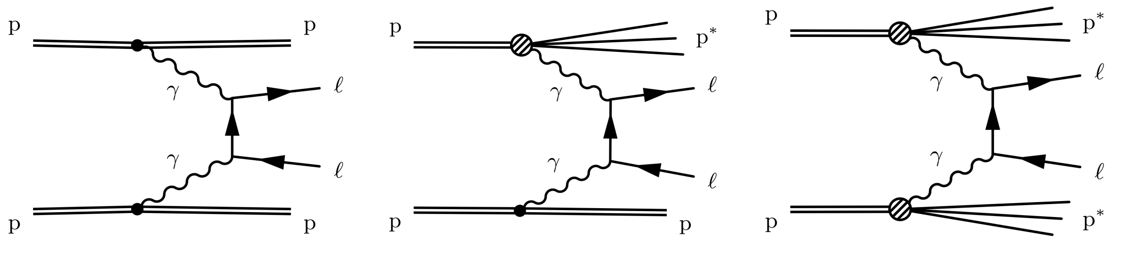

(a) Exclusive (elastic-elastic) (b) Single-dissociative (c) Double-dissociative

Figure A.2: Feynman diagrams showing dilepton production via photon fusion in (a) exclusive $pp \to p(\gamma\gamma \to \ell\ell)p$ (b) single dissociative semi-exclusive $pp \to p^*(\gamma\gamma \to \ell\ell)p$ and (c) double dissociative $pp \to p^*(\gamma\gamma \to \ell\ell)p^*$ topologies [85].

so any protons detected by the AFP spectrometer for reconstruction under the signal selection would originate from pile-up interactions. Therefore, this process is covered by the data-driven background model. However, if either proton remains intact, the case for the EE and SD channels, then it is possible for these protons to be detected by the AFP spectrometer and used for reconstruction. In this case, since the protons are correlated with the central system, this would not be covered by the data-driven background model, which assumes no correlation between the two event components. In the case of single dissociation, this could be modelled using a single-sided mixing approach, where only one of the two protons from each event is shifted to another event, allowing the possibility that the remaining proton was involved in the central production process. The effect of single-sided event mixing is investigated in Section 6.7.3, with no significant difference in shape observed in the m_X distribution. Therefore, this is not expected to cause any deviation from the data-driven background with double-sided event mixing.

The contributions from this process with intact protons will not be accounted for by the data-driven background only if the protons reconstructed in the AFP spectrometer correspond to the protons originating in this process. Figure A.3 shows the generator-level distributions of the proton fractional energy loss ξ for the EE and SD components, in each lepton channel, for pre-selected events without the fiducial selection applied. In both channels it can be seen that $< 10\%$ of the protons

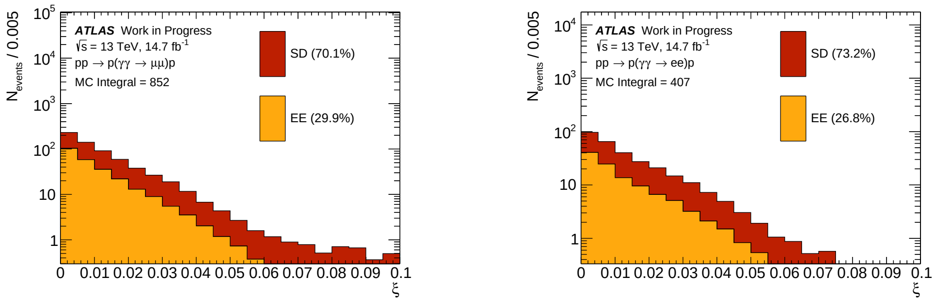

(a) Muon channel

(b) Electron channel

Figure A.3: Generator-level ξ distributions of protons which remain intact following Elastic-elastic (EE) and Single Dissociative (SD) photon-induced dilepton production in the (a) muon and (b) electron channels. Distributions are shown following pre-selection and without the fiducial selection applied.

fall within the signal selection ($0.035 \leqslant \xi \leqslant 0.08$), and therefore we expect most reconstructed events from these processes to contain only pile-up protons. This means that they are not affected by the correlation between the proton and central systems, and are therefore covered by the data-driven background model. Indeed, when matching was performed between reconstructed and generator level protons to check event-by-event whether the protons producing the central $\gamma\gamma \to \ell\ell$ system are detected in the AFP spectrometer, it was found that none of the events contained such a proton, and instead they are always replaced in reconstruction by a pile-up proton. This confirms that the data-driven background model fully covers these processes, even in the exclusive case.

A.4 Comparison with Data-Driven Model

Figures A.4 and A.5 show the combined kinematic distributions across all considered simulated background processes after all signal selections except for the track veto and dilepton p_{T} cut are applied, in each lepton channel. The distributions from

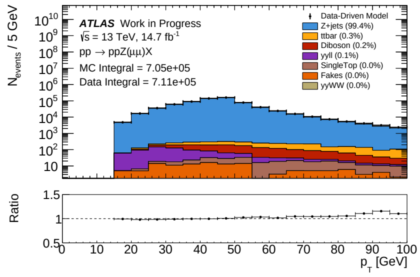

(a) Leading lepton p_{T}

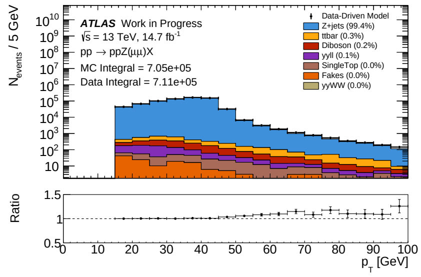

(b) Sub-leading lepton p_{T}

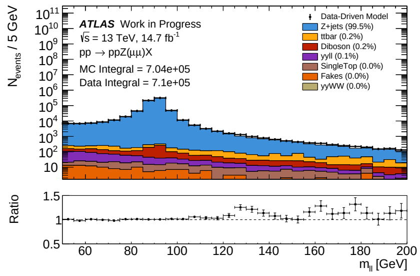

(c) Dilepton mass

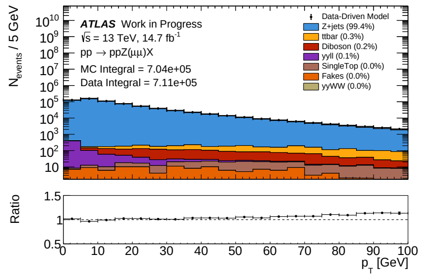

(d) Dilepton p_{T}

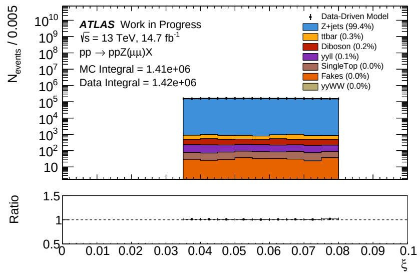

(e) Proton ξ

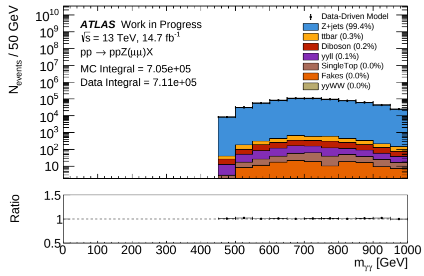

(f) $\gamma\gamma$ mass

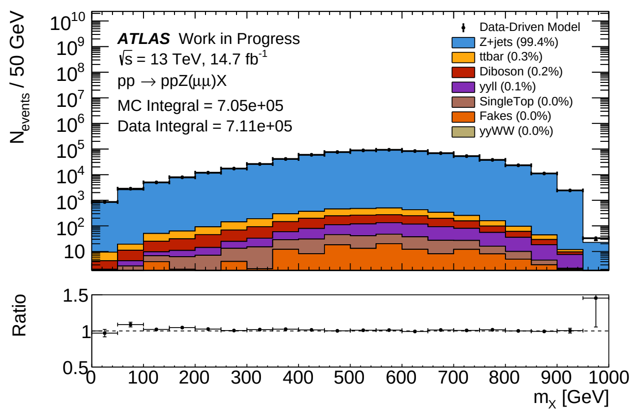

(g) Missing mass m_X

Figure A.4: Kinematic distributions from the combined simulated background model produced with all significant background contributions included, after all signal selections are applied except for the track veto and dilepton p_{T} cut, in the muon channel.

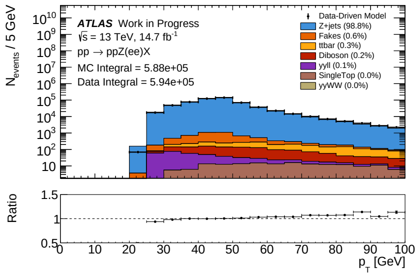

(a) Leading lepton p_{T}

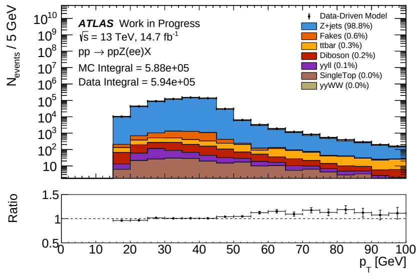

(b) Sub-leading lepton p_{T}

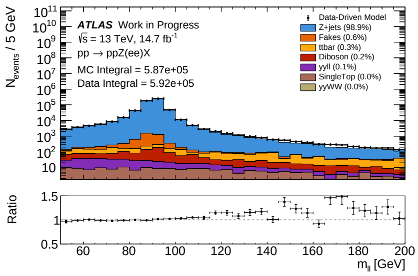

(c) Dilepton mass

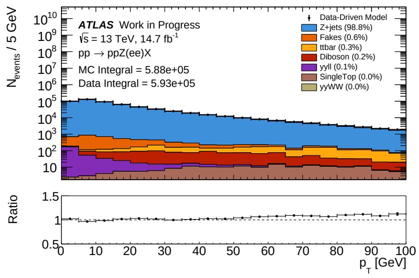

(d) Dilepton p_{T}

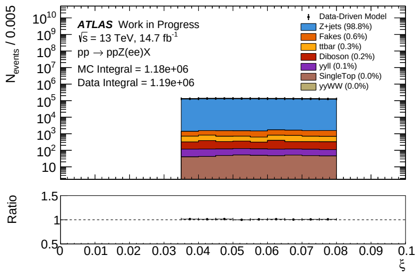

(e) Proton ξ

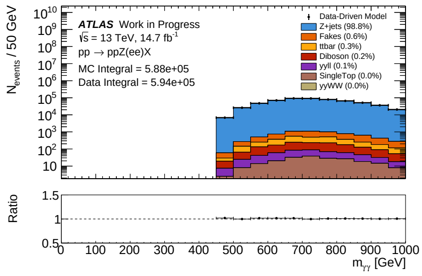

(f) $\gamma\gamma$ mass

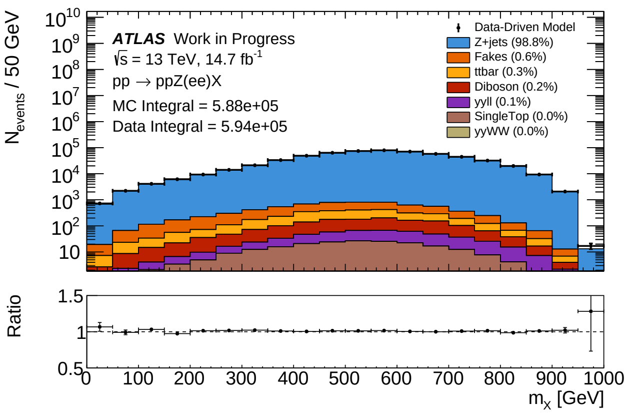

(g) Missing mass m_X

Figure A.5: Kinematic distributions from the combined simulated background model produced with all significant background contributions included, after all signal selections are applied except for the track veto and dilepton p_{T} cut, in the electron channel.

the data-driven background model are overlaid, showing good agreement within the expected statistical fluctuations across all distributions. The central component is dominated by Z+jets production, with all other processes contributing less than 1% to the total yield.

The track veto is omitted in Figures A.4 and A.5 due to the known mismodelling of the Underlying Event (UE) in simulation. The UE produces the majority of additional tracks in background events which cause them to be removed by the track veto. This mismodelling leads to an overestimate of the number of events passing the track veto, resulting in an inaccurate normalisation of the simulated background model. This mismodelling was investigated in the dominant Z+jets process, and the Sherpa samples used for the distributions shown above were found to overestimate the number of events passing the track veto by around a factor of 5 compared to data. Two alternative simulated samples were additionally tested, as listed in Table A.1, both with Powheg as the matrix element generator interfaced with another program for UEPS simulation. The first alternative sample used Pythia for the UEPS simulation, and had an event yield following the application of the track veto criterion which exceeded that observed in data by a factor of 2. The second alternative sample used Herwig for the UEPS simulation, and an excess of 25% was observed in the event yield over data following the track veto requirement. These values are consistent with those found in previous analyses [87] and emphasise the benefit of using a fully data-driven background model in this analysis, for which such mismodelling issues are not a concern. The dilepton p_{T} cut is also omitted in Figures A.4 and A.5 as this distribution is known to be mismodelled in Sherpa, as detailed in [208].

The dilepton p_{T} and track veto selections are reintroduced in Figures A.6 and A.7, and the corresponding distributions are normalised to remove any mismodelling effects on the overall normalisation of the simulated background model. It can be seen that following the track veto the previously dominant Z+jets background is heavily suppressed, and the $\gamma\gamma \rightarrow \ell\ell$ process now becomes significant. This is due to the exclusive nature of this background, with no additional tracks produced alongside the dilepton pair, giving the same visible final state as the signal process and making this process almost unaffected by the track veto. Although the overall event yields cannot be compared due to the normalisation, the shapes of the distributions remain well matched between the simulated and data-driven models, validating that the data-driven model properly covers all of the considered contributing processes.

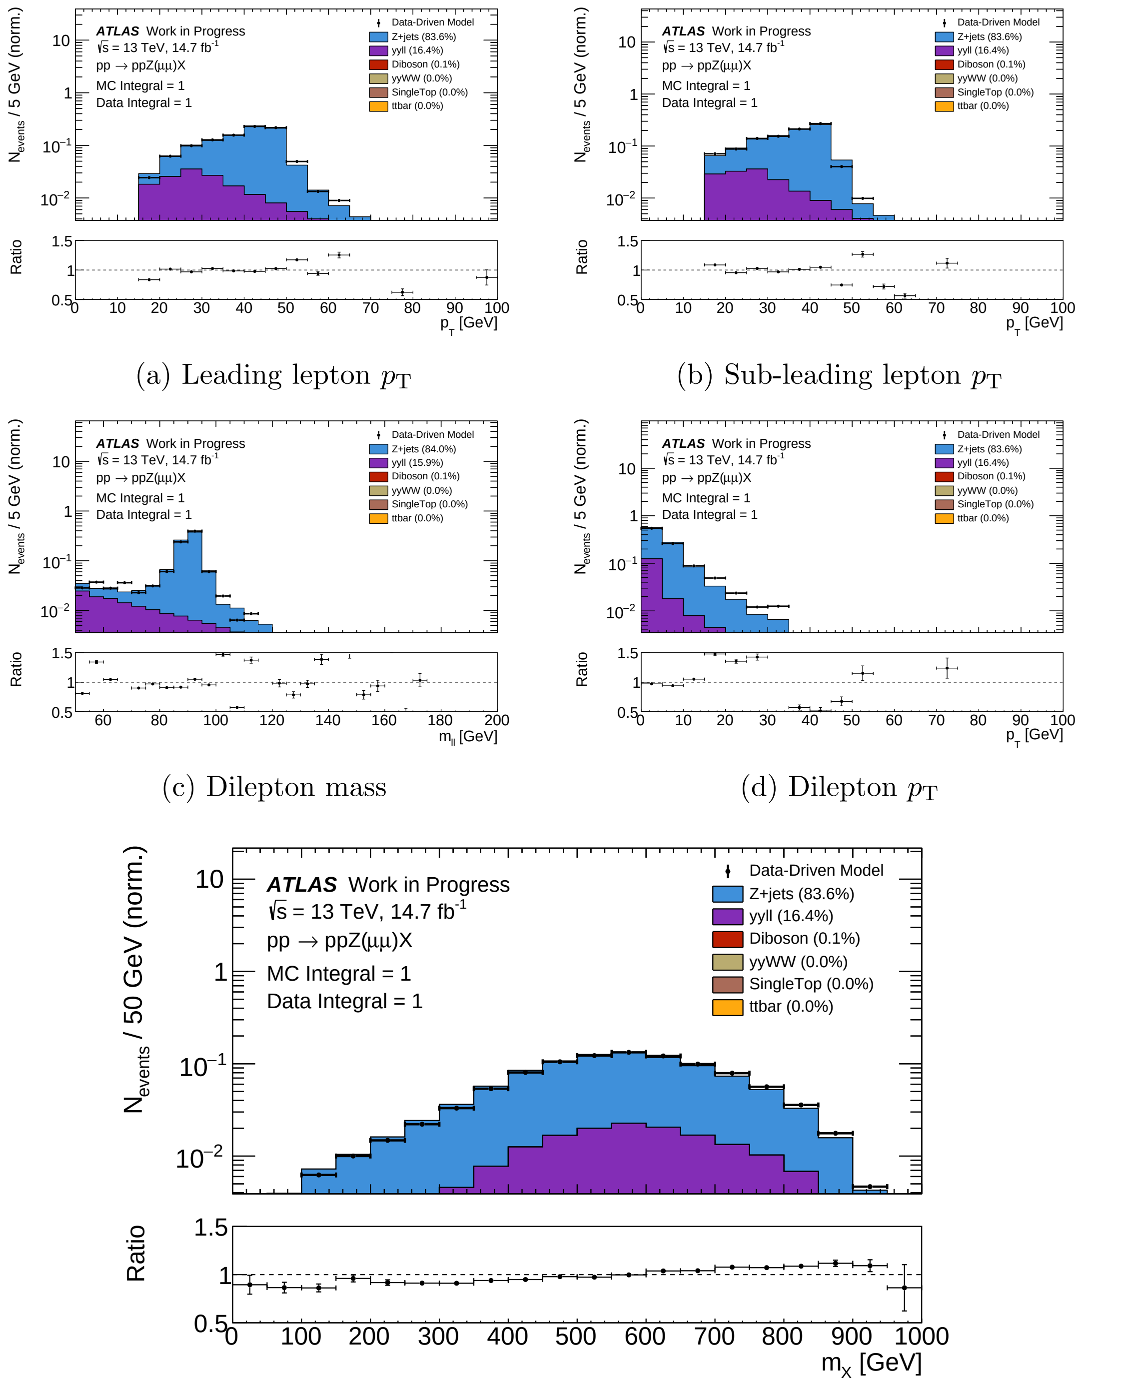

(a) Leading lepton p_{T}

(b) Sub-leading lepton p_{T}

(c) Dilepton mass

(d) Dilepton p_{T}

(e) Missing mass m_X

Figure A.6: Kinematic distributions after the high-p_{T} track veto is applied, from the total simulated background model produced with all significant background contributions included and compared to data-driven model, in the muon channel.

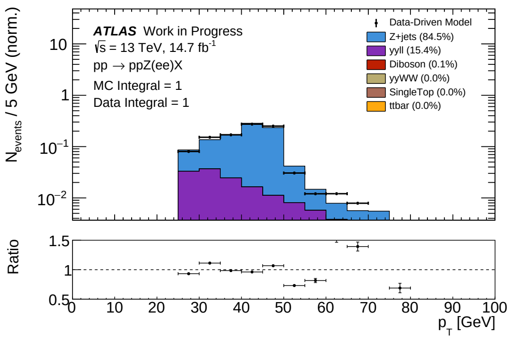

(a) Leading lepton p_{T}

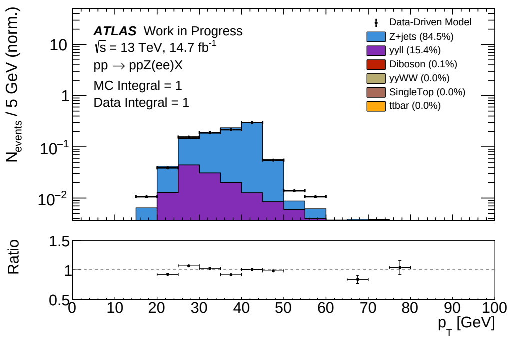

(b) Sub-leading lepton p_{T}

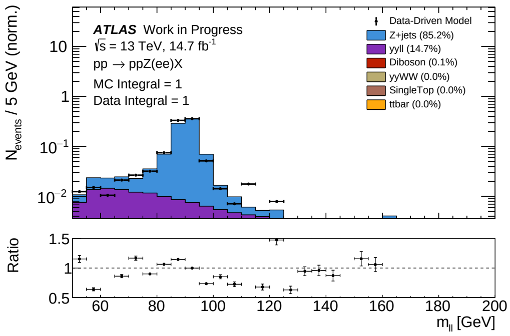

(c) Dilepton mass

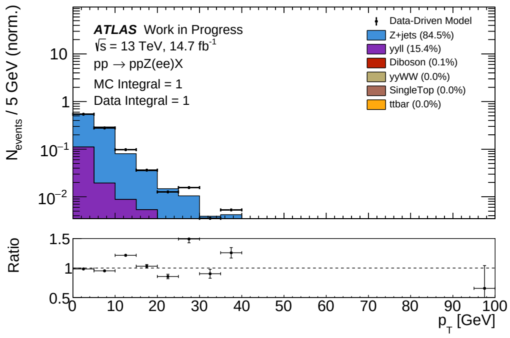

(d) Dilepton p_{T}

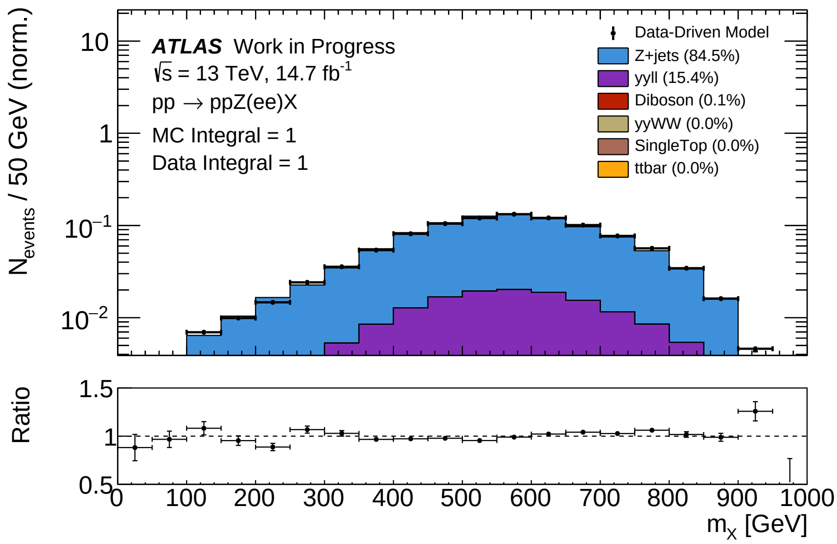

(e) Missing mass m_X

Figure A.7: Distributions after high-p_{T} track veto is applied, from total simulated background model produced with all considered background contributions and compared to data-driven model, in the electron channel.

Appendix B

Signal-Induced Background Investigation

Several investigations were undertaken into the nature of the signal-induced background introduced in Section 6.6.2, which occurs when signal protons are replaced in the missing mass reconstruction by uncorrelated protons from independent pile-up interactions.

By matching the generator-level signal protons to the protons reconstructed by the AFP spectrometer, the proportion of signal events with positive reconstructed m_X^2 but mismatched protons due to the above effect was determined for each considered signal model and mass. Figure B.1 shows that the mismatching effect is observed for every signal model tested, becoming significant at low and high hypothesised signal mass. This follows from the effect arising due to signal protons falling outside the signal ξ selection, which is most common at these signal mass points. For very high signal masses $m_X > \sim 900$ GeV, the proportion of mismatched signal events approaches 100%. This can be understood from the acceptance estimates presented in Section 6.2.2 which suggest that at least one signal proton is always expected to be outside of the selection at around this mass threshold. The mismatching effect is minimised at missing masses between 400-600 GeV depending on the signal model, as this region falls in the "sweet-spot" with high probability that both signal protons are within the selection range.

Figure B.2 shows a comparison between the reconstructed missing mass distributions for simulated signal events where all signal protons are correctly reconstructed (matched), where a single signal proton is replaced by one from pile-up (single mismatched) and where both signal protons are replaced by pile-up (double mismatched). Only distributions for the SuperChic $Z+X$ model are shown, although these are similar to of the equivalent distributions in the other models. It can be seen in Figure B.2a, which shows a combination of all events, that the mismatched proton effect forms a wide resonance resembling the combinatorial background un-

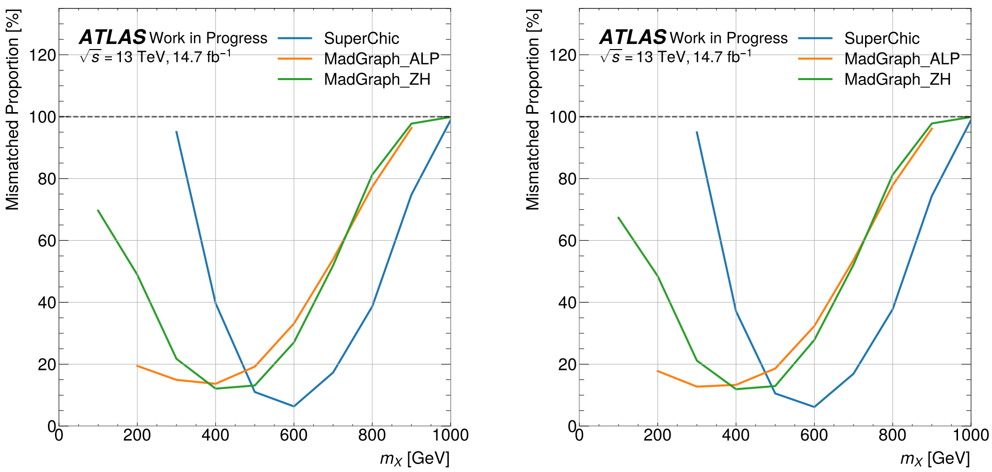

(a) Muon channel

(b) Electron channel

Figure B.1: Proportion of events with $m_X > 0$ with (a) both reconstructed protons matched to the truth-level signal protons and (b) at least one reconstructed proton mismatched to a pile-up proton, for each signal model as a function of mass.

derlying the true signal peaks. The isolated signal resonances, shown in Figure B.2b, have relatively narrow widths below 100 GeV, which follows from the resolution of 50 GeV determined in Section 6.2.1, with the varying integrals for each mass corresponding to the respective selection efficiencies of each model. Figures B.2c and B.2d show the wide background resonance expected from this mismatching effect, which closely resembles the combinatorial background process. In fact, the case of double mismatching is exactly equivalent to the combinatorial background, with both signal protons replaced by pile-up protons, and therefore this component will be accounted for by the data-driven background model, if any signal is present in the dataset. However, the single mismatching effect is not fully accounted for, due to the presence of a single signal proton in the reconstruction, and this appears to have some effect on the shape of the missing mass distribution for these events compared to the double mismatching, with the peak of the distribution shifting slightly for each signal mass. This component therefore requires special consideration. As shown by the relative normalisations of the distributions in Figures B.2c and B.2d, single mismatching is the dominant contribution, so additional steps should be taken to mitigate this background.

Due to the majority of this background leading to reconstructed missing mass values away from the narrow signal peak, the effect of defining a mass window to select only events with reconstructed missing mass within 50 GeV of the hypothesised signal mass was considered. Figure B.3 shows the resulting effect on the proportion of signal events with mismatched protons for each signal model, with a large reduction in the signal-induced background observed across all models. However, the tight mass window also causes significant loss of statistics in genuine signal events, so another method is preferable which affects only the background process.

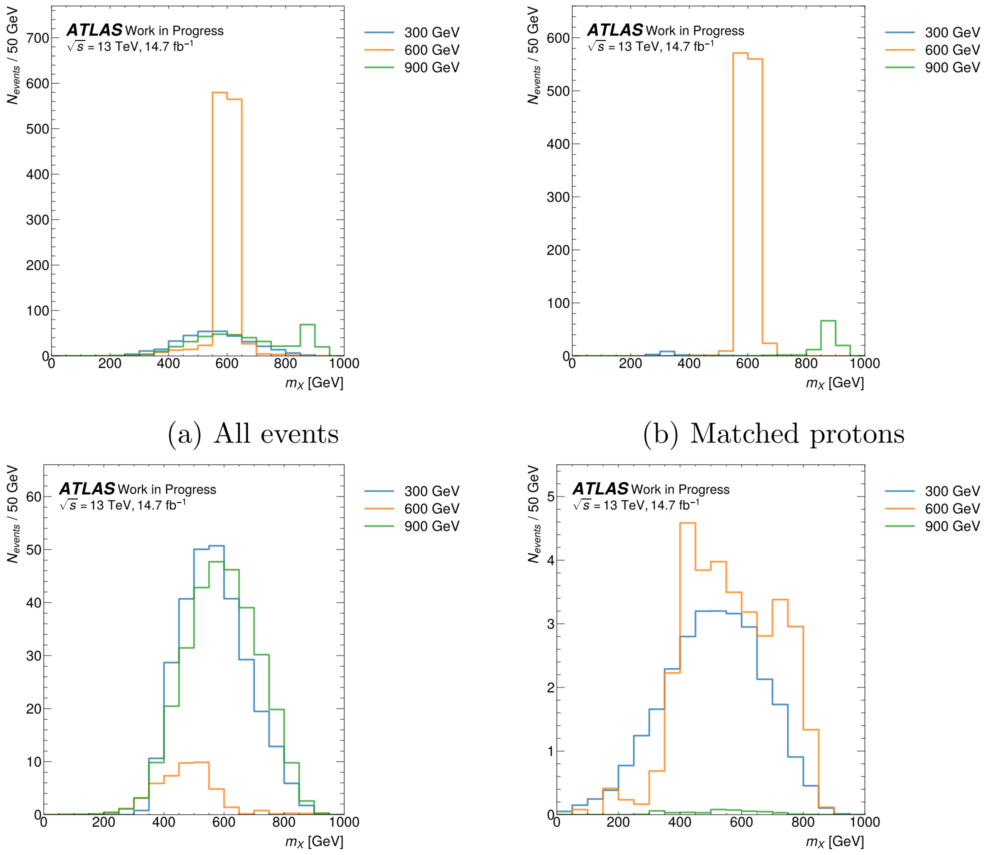

(a) All events (b) Matched protons

(c) Single Mismatched Protons (d) Double Mismatched Protons

Figure B.2: Reconstructed missing mass distributions for SuperChic signal model at selected mass points for (a) all events, (b) only events with both reconstructed protons matched to the truth level signal protons, (c) only events with exactly one reconstructed proton which is not matched to a truth level signal proton and (d) only events with both reconstructed protons not matched to a truth level signal proton.

This was achieved through the inclusion the signal proton ξ selection to the fiducial selection given in Table 6.4. Since this selection is applied at generator-level, it removes any signal events where the reconstructed signal protons are not likely to fall within the signal selection and so could contribute to the signal-induced background, while leaving the signal events falling within selection unaffected. Figure B.4 shows the resulting effect on the proportion of signal events with mismatched protons, and indeed a similar reduction to that observed with the mass window is achieved, without the reduction in statistics for matched signal events. It can additionally be observed here that all sensitivity is lost for models with $m_X > 900$ GeV as expected, since very close to 100% of these events originate in the signal-induced background. Events which fall outside the fiducial selection, including the signal-induced background, are treated as an additional background component.

These distributions show that a significant proportion of $\mathcal{O}(10\%)$ of signal events in

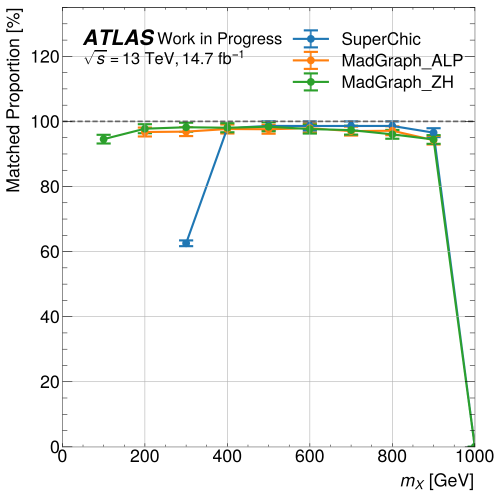

(a) Matched protons, muon channel

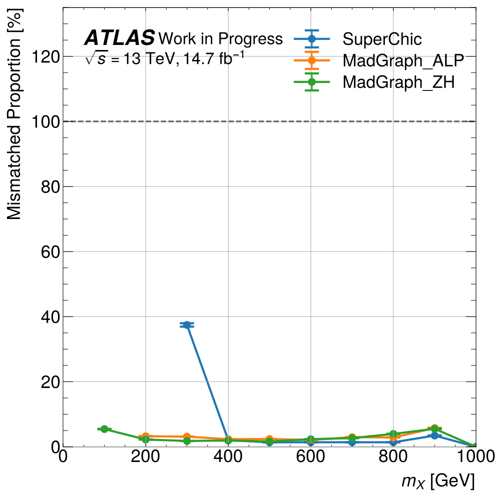

(b) Mismatched protons, muon channel

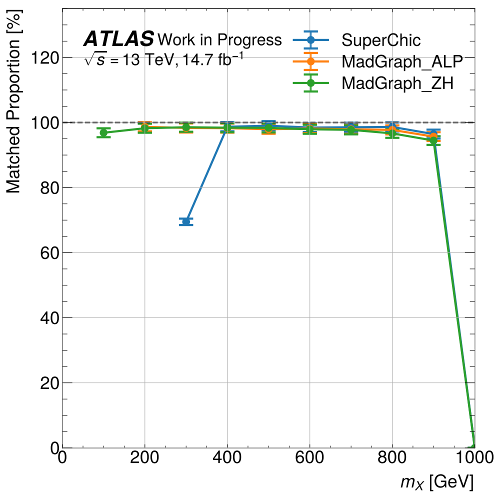

(c) Matched protons, electron channel

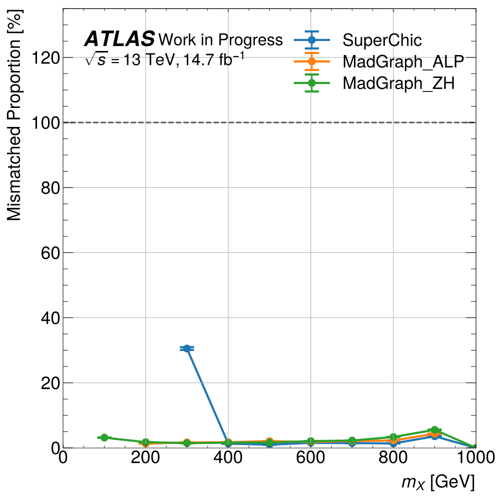

(d) Mismatched protons, electrons

Figure B.3: Proportion of events with $m_X > 0$ with (left) both reconstructed protons matched to the generator-level signal protons and (right) at least one reconstructed proton mismatched to a pile-up proton, when a mass window is applied 50 GeV either side of the hypothesised signal mass for a given model, for each signal model as a function of mass in each lepton channel.

all models with a hypothesised signal mass of 900 GeV are still reconstructed using mismatched protons following the application of the fiducial selection. However, this high proportion is attributed mainly to the low overall statistics in these samples following the application of the fiducial selection. This ultimately results in these models being dropped from the final result for all models except the SuperChic $Z+X$ model, due to the resulting fits being unstable. Additionally, in the final fits a mass window is used in conjunction with the fiducial selection, which reduces this contribution to negligible levels even for models with $m_X = 900$ GeV.

Initially, the normalisation of this background component was left to float freely

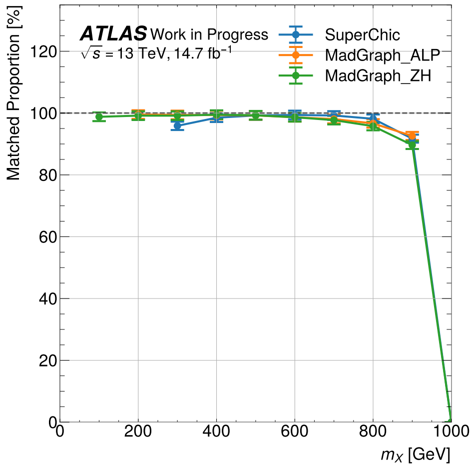

(a) Matched protons, muon channel

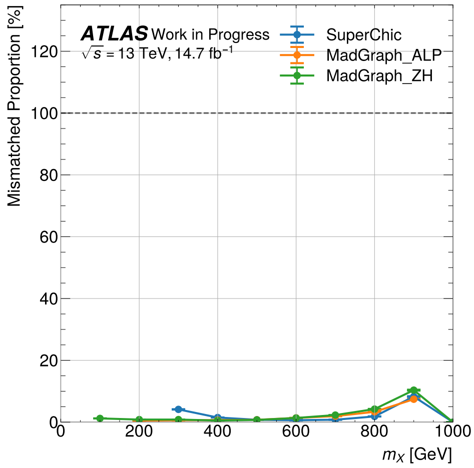

(b) Mismatched protons, muon channel

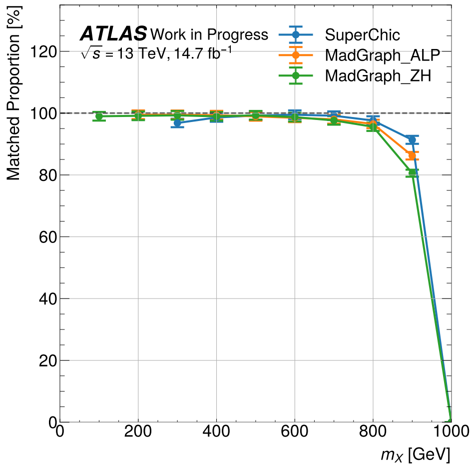

(c) Matched protons, electron channel

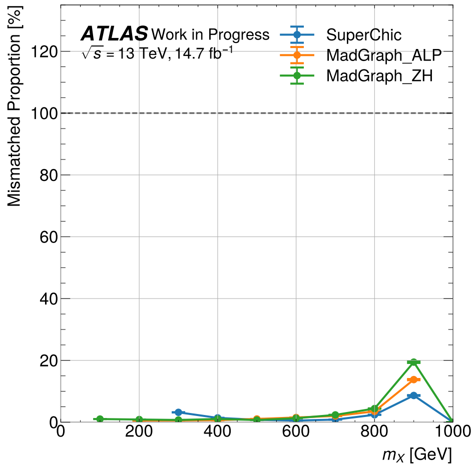

(d) Mismatched protons, electrons

Figure B.4: Proportion of events with $m_X > 0$ with (left) both reconstructed protons matched to the truth-level signal protons and (right) at least one reconstructed proton mismatched to a pile-up proton, when the fiducial selection is applied, for each signal model as a function of mass in each lepton channel.

in the fit of the overall signal and background prediction to the observed dataset presented in Section 6.9. This was done to compensate for any potential mismodelling arising from the random nature of this background. However, as discussed in Section 6.9, a single-bin approach was eventually adopted which required only a single free parameter, the signal normalisation, to be present in the fit. Therefore, the normalisation of this background is fixed in the final fits, however the single-bin approach additionally uses a mass window which is slightly wider than the one tested above, which heavily suppresses the remaining signal-induced background such that this contribution becomes negligible in the final results.

References

[1] F. Halzen & A. D. Martin; *Quarks and Leptons: An Introductory Course in Modern Particle Physics.* (1984). 8 citations on pages ii, 6, 7, 8, 10, 11, 12, and 16

[2] P. Langacker; *The Standard Model and Beyond.* (2010). 10 citations on pages ii, 6, 7, 8, 10, 11, 12, 14, 16, and 17

[3] S. Clawson; “The light at the end of the tunnel gets weaker: Observation and measurement of photon-induced W^+W^- production at the ATLAS Experiment”; (2022). `https://cds.cern.ch/record/2851404`; presented 2022. 8 citations on pages ii, 24, 85, 86, 87, 188, 239, and 241

[4] MissMJ and Cush; “Standard Model of Elementary Particles”; (2019). `https://commons.wikimedia.org/wiki/File:Standard_Model_of_Elementary_Particles.svg`; under the CC-BY-SA-3.0. 2 citations on pages 5 and 239

[5] P. A. M. Dirac & N. H. D. Bohr; “The quantum theory of the emission and absorption of radiation”; Proc. R. Soc. Lond. Ser. A **114**, pp. 243–265 (1927). Cited on page 6

[6] P. A. M. Dirac & R. H. Fowler; “A theory of electrons and protons”; Proc. R. Soc. Lond. Ser. A **126**, pp. 360–365 (1930). Cited on page 7

[7] LHCb Collaboration; “Observation of the Resonant Character of the $Z(4430)^-$ State”; Phys. Rev. Lett. **112** (2014). Cited on page 7

[8] LHCb Collaboration; “Observation of $J/\psi p$ Resonances Consistent with Pentaquark States in $\Lambda_b^0 \to J/\psi K^- p$ Decays”; Phys. Rev. Lett. **115** (2015). Cited on page 7

[9] O. W. Greenberg; “Spin and Unitary-Spin Independence in a Paraquark Model of Baryons and Mesons”; Phys. Rev. Lett. **13**, pp. 598–602 (1964). Cited on page 7

[10] W. Pauli; "On the connection between the closure of electron groups in the atom and the complex structure of the spectra"; Z. Physik **31**, pp. 765–783 (1925). Cited on page 7

[11] M. Gell-Mann; "The Eightfold Way: A Theory of strong interaction symmetry"; (1961). Cited on page 8

[12] C. N. Yang & R. L. Mills; "Conservation of Isotopic Spin and Isotopic Gauge Invariance"; Phys. Rev. **96**, pp. 191–195 (1954). Cited on page 8

[13] E. Fermi; "Attempt at a theory of β-rays"; I. Z. Physik **88**, pp. 161–177 (1934). Cited on page 9

[14] C. S. Wu *et al.*; "Experimental Test of Parity Conservation in Beta Decay"; Phys. Rev. **105**, pp. 1413–1415 (1957). Cited on page 9

[15] T. D. Lee & C. N. Yang; "Question of Parity Conservation in Weak Interactions"; Phys. Rev. **104**, pp. 254–258 (1956). Cited on page 9

[16] W. J. Salam, A.; "Weak and electromagnetic interactions"; Nuovo Cim **11**, pp. 568–577 (1959). Cited on page 10

[17] S. L. Glashow; "Partial Symmetries of Weak Interactions"; Nucl. Phys. **22**, pp. 579–588 (1961). Cited on page 10

[18] S. Weinberg; "A Model of Leptons"; Phys. Rev. Lett. **19**, pp. 1264–1266 (1967). 2 citations on pages 10 and 11

[19] P. W. Higgs; "Broken Symmetries and the Masses of Gauge Bosons"; Phys. Rev. Lett. **13**, pp. 508–509 (1964). Cited on page 10

[20] F. Englert & R. Brout; "Broken Symmetry and the Mass of Gauge Vector Mesons"; Phys. Rev. Lett. **13**, pp. 321–323 (1964). Cited on page 10

[21] G. S. Guralnik *et al.*; "Global Conservation Laws and Massless Particles"; Phys. Rev. Lett. **13**, pp. 585–587 (1964). Cited on page 10

[22] J. Goldstone *et al.*; "Broken Symmetries"; Phys. Rev. **127**, pp. 965–970 (1962). Cited on page 11

[23] S. Navas *et al.*; "Review of Particle Physics"; Phys. Rev. D **110** (2024). 2 citations on pages 12 and 15

[24] ATLAS Collaboration; "Observation of a new particle in the search for the Standard Model Higgs boson with the ATLAS detector at the LHC"; Phys. Lett. B **716**, pp. 1–29 (2012). 2 citations on pages 12 and 36

[25] CMS Collaboration; "Observation of a new boson at a mass of 125 GeV with the CMS experiment at the LHC"; Phys. Lett. B **716**, pp. 30–61 (2012). Cited on page 12

[26] R. P. Feynman; "Space-Time Approach to Quantum Electrodynamics"; Phys. Rev. **76**, pp. 769–789 (1949). 2 citations on pages 13 and 21

[27] R. P. Feynman; "A Relativistic Cut-Off for Classical Electrodynamics"; Phys. Rev. **74**, pp. 939–946 (1948). Cited on page 14

[28] A. Deur *et al.*; "The QCD running coupling"; Prog. Part. Nucl. Phys. **90**, pp. 1–74 (2016). Cited on page 15

[29] C. Balazs *et al.*; "A Primer on Dark Matter"; (2024). `https://arxiv.org/abs/2411.05062`; 2411.05062. Cited on page 15

[30] H. Georgi & S. L. Glashow; "Unity of All Elementary-Particle Forces"; Phys. Rev. Lett. **32**, pp. 438–441 (1974). Cited on page 16

[31] R. D. Peccei & H. R. Quinn; "CP Conservation in the Presence of Pseudoparticles"; Phys. Rev. Lett. **38**, pp. 1440–1443 (1977). Cited on page 17

[32] S. Weinberg; "A New Light Boson?" Phys. Rev. Lett. **40**, pp. 223–226 (1978). Cited on page 17

[33] F. Wilczek; "Problem of Strong P and T Invariance in the Presence of Instantons"; Phys. Rev. Lett. **40**, pp. 279–282 (1978). Cited on page 17

[34] R. D. Peccei & H. R. Quinn; "CP Conservation in the Presence of Pseudoparticles"; Phys. Rev. Lett. **38**, pp. 1440–1443 (1977). Cited on page 17

[35] J. E. Kim & G. Carosi; "Axions and the strong CP problem"; Rev. Mod. Phys. **82**, pp. 557–601 (2010). Cited on page 17

[36] G. Bertone *et al.*; "Particle dark matter: evidence, candidates and constraints"; Phys. Rep. **405**, pp. 279–390 (2005). Cited on page 17

[37] N. Du *et al.*; "Search for Invisible Axion Dark Matter with the Axion Dark Matter Experiment"; Phys. Rev. Lett. **120** (2018). Cited on page 17

[38] C. Baldenegro *et al.*; "Searching for axion-like particles with proton tagging at the LHC"; JHEP **2018** (2018). Cited on page 17

[39] ATLAS Collaboration; "Measurement of light-by-light scattering and search for axion-like particles with 2.2 nb^{-1} of Pb+Pb data with the ATLAS detector"; JHEP **2021** (2021). 2 citations on pages 18 and 239

[40] ATLAS Collaboration; "Search for an axion-like particle with forward proton scattering in association with photon pairs at ATLAS"; JHEP **2023**, p. 234 (2023). 4 citations on pages 19, 72, 134, and 138

[41] ATLAS Collaboration; "Search for new phenomena in events with at least three photons collected in pp collisions at $\sqrt{s} = 8$ TeV with the ATLAS detector"; Eur. Phys. J. C **76** (2016). Cited on page 19

[42] CMS Collaboration; "Search for Exotic Higgs Boson Decays $H \to aa \to 4\gamma$ with Events Containing Two Merged Diphotons in Proton-Proton Collisions at $\sqrt{s} = 13$ TeV"; Phys. Rev. Lett. **131** (2023). Cited on page 19

[43] CMS Collaboration; "Search for the exotic decay of the Higgs boson into two light pseudoscalars with four photons in the final state in proton-proton collisions at $\sqrt{s} = 13$ TeV"; JHEP **2023** (2023). Cited on page 19

[44] ATLAS Collaboration; "Search for short- and long-lived axion-like particles in $H \to aa \to 4\gamma$ decays with the ATLAS experiment at the LHC"; Eur. Phys. J. C **84** (2024). Cited on page 19

[45] E. Chapon *et al.*; "Anomalous quartic $WW\gamma\gamma$, $ZZ\gamma\gamma$ and trilinear $WW\gamma\gamma$ couplings in two-photon processes at high luminosity at the LHC"; Phys. Rev. D **81** (2010). Cited on page 19

[46] O. Kepka & C. Royon; "Anomalous $WW\gamma$ coupling in photon-induced processes using forward detectors at the CERN LHC"; Phys. Rev. D **78** (2008). Cited on page 19

[47] S. Fichet & G. von Gersdorff; "Anomalous gauge couplings from composite Higgs and warped extra dimensions"; JHEP **2014** (2014). Cited on page 19

[48] S. Fichet *et al.*; "Probing new physics in diphoton production with proton tagging at the Large Hadron Collider"; Phys. Rev. D **89** (2014). Cited on page 19

[49] C. Baldenegro *et al.*; "Probing the anomalous $\gamma\gamma\gamma Z$ coupling at the LHC with proton tagging"; JHEP **2017** (2017). Cited on page 19

[50] C. Baldenegro & S. Fichet; "Anomalous gauge interactions in photon collisions at the LHC and the FCC"; (2017). `https://arxiv.org/abs/1708.07531`; 1708.07531. Cited on page 19

[51] S. İnan; "Dimension-six anomalous $tq\gamma$ couplings in $\gamma\gamma$ collision at the LHC"; Nucl. Phys. B **897**, pp. 289–301 (2015). Cited on page 19

[52] I. Şahin & S. İnan; "Probe of unparticles at the LHC in exclusive two lepton and two photon production via photon-photon fusion"; JHEP **2009**, p. 69 (2009). Cited on page 19

[53] S. Fichet; "Shining light on polarizable dark particles"; JHEP **2017** (2017). Cited on page 19

[54] S. Atağ *et al.*; "Extra dimensions in $\gamma\gamma \to \gamma\gamma$ process at the CERN-LHC"; JHEP **2010** (2010). Cited on page 19

[55] R. S. Gupta; "Probing quartic neutral gauge boson couplings using diffractive photon fusion at the LHC"; Phys. Rev. D **85** (2012). Cited on page 19

[56] H. Sun; "Large extra dimension effects through light-by-light scattering at the CERN LHC"; Eur. Phys. J. C **74** (2014). Cited on page 19

[57] L. N. Epele *et al.*; "Looking for magnetic monopoles at LHC with diphoton events"; Eur. Phys. J. Plus **127** (2012). Cited on page 19

[58] S. Fichet *et al.*; "Scattering light by light at 750 GeV at the LHC"; Phys. Rev. D **93** (2016). Cited on page 19

[59] S. Fichet *et al.*; "Measuring the Diphoton Coupling of a 750 GeV Resonance"; Phys. Rev. Lett. **116** (2016). Cited on page 19

[60] S. Fichet *et al.*; "Light-by-light scattering with intact protons at the LHC: from standard model to new physics"; JHEP **2015** (2015). Cited on page 19

[61] S. Fichet; "Prospects for new physics searches at the LHC in the forward proton mode"; (2015). `https://arxiv.org/abs/1510.01004`; 1510.01004. Cited on page 19

[62] A. M. Cooper-Sarkar; "Parton Distribution Functions for discovery physics at the LHC"; (2023). `https://arxiv.org/abs/2302.11788`; 2302.11788. Cited on page 20

[63] H1 and ZEUS Collaborations; "Combination of Measurements of Inclusive Deep Inelastic $e^{\pm}p$ Scattering Cross Sections and QCD Analysis of HERA Data"; (2015). `https://arxiv.org/abs/1506.06042`; 1506.06042. Cited on page 20

[64] Z. P. Zhang; "Determination of proton parton distribution functions using ATLAS data"; Technical report; CERN; Geneva (2022). `https://cds.cern.ch/record/2835612`. Cited on page 20

[65] C. Schmidt *et al.*; "CT14QED parton distribution functions from isolated photon production in deep inelastic scattering"; Phys. Rev. D **93** (2016). 2 citations on pages 20 and 23

[66] S. Bailey *et al.*; "Parton distributions from LHC, HERA, Tevatron and fixed target data: MSHT20 PDFs"; Eur. Phys. J. C **81** (2021). Cited on page 20

[67] T. Cridge *et al.*; "QED parton distribution functions in the MSHT20 fit"; Eur. Phys. J. C **82** (2022). Cited on page 20

[68] L. A. Harland-Lang *et al.*; "Ad Lucem: QED parton distribution functions in the MMHT framework"; Eur. Phys. J. C **79** (2019). Cited on page 20

[69] A. Buckley *et al.*; "LHAPDF6: parton density access in the LHC precision era"; Eur. Phys. J. C **75** (2015). Cited on page 20

[70] J. C. Collins *et al.*; "Factorization of Hard Processes in QCD"; (2004). `https://arxiv.org/abs/hep-ph/0409313`; hep-ph/0409313. 2 citations on pages 21 and 71

[71] J. C. Collins & D. E. Soper; “The Theorems of Perturbative QCD”; Annu. Rev. Nucl. Part. Sci. **37**, pp. 383–409 (1987). Cited on page 21

[72] ATLAS Collaboration; “Standard Model Summary Plots June 2024”; Technical report; CERN; Geneva (2024). `https://cds.cern.ch/record/2903866`. 2 citations on pages 22 and 239

[73] NNPDF collaboration; “Parton distributions for the LHC run II”; JHEP **2015** (2015). Cited on page 23

[74] R. M. Martin, A.D.; “The photon PDF of the proton”; Eur. Phys. J. C **74** (2014). Cited on page 23

[75] C. A. Bertulani; “Electromagnetic interaction of ultrarelativistic heavy ions”; Phys. Rev. A **63** (2001). Cited on page 23

[76] V. Budnev *et al.*; “The two-photon particle production mechanism. Physical problems. Applications. Equivalent photon approximation”; Phys. Rep. **15**, pp. 181–282 (1975). Cited on page 23

[77] M. I. Vysotsky & E. V. Zhemchugov; “Equivalent photons in proton-proton and ion-ion collisions at the Large Hadron Collider”; Physics-Uspekhi **62**, pp. 910–919 (2019). Cited on page 24

[78] T. Han *et al.*; “Structure-function approach to vector-boson scattering in pp collisions”; Phys. Rev. Lett. **69**, pp. 3274–3277 (1992). Cited on page 24

[79] L. Harland-Lang; “The proton in high definition: revisiting photon-initiated production in high energy collisions”; JHEP **2020** (2020). Cited on page 24

[80] L. A. Harland-Lang, *et al.*; “A new approach to modelling elastic and inelastic photon-initiated production at the LHC: SuperChic 4”; Eur. Phys. J. C **80** (2020). 6 citations on pages 25, 27, 31, 142, 183, and 245

[81] S. Fichet *et al.*; “Light-by-light scattering with intact protons at the LHC: from standard model to new physics”; JHEP **2015** (2015). Cited on page 26

[82] L. A. Harland-Lang *et al.*; “The photon PDF in events with rapidity gaps”; Eur. Phys. J. C **76** (2016). Cited on page 27

[83] V. A. Khoze *et al.*; “Multiple interactions and rapidity gap survival”; J. Phys. G: Nucl. Part. Phys. **45**, p. 053002 (2018). Cited on page 27

[84] M. Trzebinski; “Diffractive Physics at the LHC”; Ukr. J. Phys. **64**, pp. 772–775 (2019). Cited on page 27

[85] ATLAS Collaboration; “Observation and measurement of forward proton scattering in association with lepton pairs produced via the photon fusion mechanism at ATLAS”; Phys. Rev. Lett. **125** (2020). 12 citations on pages 27, 72, 82, 92, 134, 138, 161, 174, 183, 207, 208, and 249

[86] T. Gleisberg *et al.*; "Event generation with SHERPA 1.1"; JHEP **2009**, p. 7 (2009). 3 citations on pages 29, 30, and 239

[87] ATLAS Collaboration; "Observation of photon-induced W^+W^- production in pp collisions at $\sqrt{s} = 13$ TeV using the ATLAS detector"; Phys. Lett. B **816** (2021). 5 citations on pages 29, 161, 162, 174, and 212

[88] ATLAS Collaboration; "Measurement of distributions sensitive to the underlying event in inclusive Z boson production in pp collisions at $\sqrt{s} = 13$ TeV with the ATLAS detector"; Eur. Phys. J. C **79** (2019). Cited on page 29

[89] E. Bothmann *et al.*; "Event generation with Sherpa 2.2"; SciPost Phys. **7** (2019). Cited on page 29

[90] M. Bähr *et al.*; "Herwig++ physics and manual." Eur. Phy. J. C **58**, pp. 639–707 (2008). Cited on page 29

[91] T. Sjöstrand *et al.*; "An introduction to PYTHIA 8.2"; Comput. Phys. Commun. **191**, pp. 159–177 (2015). 2 citations on pages 29 and 141

[92] C. Bierlich *et al.*; "A comprehensive guide to the physics and usage of PYTHIA 8.3"; (2022). `https://arxiv.org/abs/2203.11601`; 2203.11601. Cited on page 29

[93] J. Alwall, R. Frederix, S. Frixione *et al.*; "The automated computation of tree-level and next-to-leading order differential cross sections, and their matching to parton shower simulations"; JHEP **2014** (2014). 2 citations on pages 30 and 143

[94] J. Alwall *et al.*; "Computing decay rates for new physics theories with FeynRules and MadGraph 5_aMC@NLO"; Comput. Phys. Commun. **197**, pp. 312–323 (2015). Cited on page 30

[95] R. Frederix *et al.*; "The automation of next-to-leading order electroweak calculations"; JHEP **2018** (2018). Cited on page 30

[96] L. A. Harland-Lang *et al.*; "Photon-Photon Collisions with SuperChic"; CERN Conf. Proc. **1**, p. 59 (2018). 2 citations on pages 30 and 141

[97] S. Agostinelli *et al.*; "Geant4 - a simulation toolkit"; Nucl. Instrum. Methods Phys. Res. A **506**, pp. 250–303 (2003). Cited on page 31

[98] G. Aad *et al.*; "The ATLAS Simulation Infrastructure"; Eur. Phys. J. C **70**, pp. 823–874 (2010). Cited on page 31

[99] A. Hasib; "ATLAS Fast Simulation - from classical to deep learning"; (2022). Cited on page 31

[100] E. Lopienska; "The CERN accelerator complex, layout in 2022." (2022). 2 citations on pages 34 and 239

[101] ATLAS Collaboration; "ATLAS: Technical proposal for a general-purpose *pp* experiment at the Large Hadron Collider at CERN"; (1994). Cited on page 35

[102] CMS Collaboration; "CMS, the Compact Muon Solenoid: Technical proposal"; (1994). Cited on page 35

[103] LHCb Collaboration; "LHCb technical proposal: A Large Hadron Collider Beauty Experiment for Precision Measurements of CP Violation and Rare Decays"; (1998). Cited on page 35

[104] ALICE Collaboration; "ALICE: Technical proposal for a large ion collider experiment at the CERN LHC"; (1995). Cited on page 35

[105] ATLAS Collaboration; "Luminosity determination in *pp* collisions at $\sqrt{s}$ = 13 TeV using the ATLAS detector at the LHC"; Technical report; CERN; Geneva (2019). `https://cds.cern.ch/record/2677054`. Cited on page 35

[106] P. Grafström & W. Kozanecki; "Luminosity determination at proton colliders"; Prog. Part. Nucl. Phys. **81**, pp. 97–148 (2015). Cited on page 36

[107] R. Steerenberg *et al.*; "Operation and Performance of the Cern Large Hadron Collider During Proton Run 2"; in "Proc. 10th International Particle Accelerator Conference (IPAC'19), Melbourne, Australia, 19-24 May 2019," Number 10 in International Particle Accelerator Conferences ; pp. 504–507 (JACoW Publishing, Geneva, Switzerland) (2019); ISBN 978-3-95450-208-0. Cited on page 36

[108] ATLAS Collaboration; "The ATLAS experiment at the CERN Large Hadron Collider: a description of the detector configuration for Run 3"; JINST **19** (2024). Cited on page 36

[109] ATLAS Collaboration; "Public ATLAS Luminosity Results for Run-2 of the LHC - 2017 pp Collisions"; (2017). `https://twiki.cern.ch/twiki/bin/view/AtlasPublic/LuminosityPublicResultsRun2`. 3 citations on pages 37, 53, and 239

[110] CERN; "The HL-LHC project"; (2025). `https://hilumilhc.web.cern.ch/content/hl-lhc-project`. 2 citations on pages 38 and 240

[111] J. Pequenao; "Computer generated image of the whole ATLAS detector"; (2008). `https://cds.cern.ch/record/1095924`. 2 citations on pages 40 and 240

[112] J. Pequenao; "Computer generated image of the ATLAS inner detector"; (2008). `https://cds.cern.ch/record/1095926`. 2 citations on pages 42 and 240

[113] ATLAS Collaboration; "ATLAS central solenoid: Technical design report"; (1997). Cited on page 41

[114] M. Capeans *et al.*; "ATLAS Insertable B-Layer Technical Design Report"; Technical report (2010). https://cds.cern.ch/record/1291633. 2 citations on pages 42 and 77

[115] ATLAS Collaboration; "ATLAS inner detector: Technical design report. Vol. 1"; (1997). 2 citations on pages 43 and 44

[116] J. Pequenao; "Computer Generated image of the ATLAS calorimeter"; (2008). https://cds.cern.ch/record/1095927. 2 citations on pages 45 and 240

[117] ATLAS Collaboration; "ATLAS liquid argon calorimeter: Technical design report"; (1996). 5 citations on pages 46, 47, 48, 49, and 240

[118] ATLAS Collaboration; "ATLAS tile calorimeter: Technical design report"; (1996). Cited on page 48

[119] J. Proudfoot & T. Davidek; "Tile module schematic"; (2008). https://twiki.cern.ch/twiki/bin/view/AtlasPublic/PublishedTilecalFigures. 2 citations on pages 48 and 240

[120] J. Pequenao; "Computer generated image of the ATLAS Muons subsystem"; (2008). https://cds.cern.ch/record/1095929. 2 citations on pages 50 and 240

[121] ATLAS Collaboration; "The ATLAS Experiment at the CERN Large Hadron Collider"; JINST **3** (2008). Cited on page 51

[122] ATLAS Collaboration; *ATLAS muon spectrometer: Technical Design Report*; Technical design report. ATLAS (CERN, Geneva). (1997). Cited on page 51

[123] ATLAS Collaboration; "ATLAS Forward Detectors"; (2018). https://cds.cern.ch/record/2627582; general Photo. 2 citations on pages 52 and 240

[124] L. Fabbri; "Forward Detectors in ATLAS: LUCID, ZDC and ALFA"; p. 166 (2009). 2 citations on pages 52 and 53

[125] Viazlo, O and the ATLAS LUCID Collaboration; "ATLAS LUCID detector upgrade for LHC Run 2"; Technical report; CERN; Geneva (2015). https://cds.cern.ch/record/2062038. Cited on page 52

[126] P. Jenni *et al.*; "Zero Degree Calorimeters for ATLAS"; Technical report; CERN; Geneva (2007). https://cds.cern.ch/record/1009649. Cited on page 53

[127] L. Adamczyk *et al.*; "Technical Design Report for the ATLAS Forward Proton Detector"; Technical report (2015). https://cds.cern.ch/record/2017378. 5 citations on pages 53, 72, 76, 77, and 240

[128] P. Jenni *et al.*; "ATLAS Forward Detectors for Measurement of Elastic Scattering and Luminosity"; (2008). Cited on page 53

[129] ATLAS Collaboration; "Operation of the ATLAS trigger system in Run 2"; JINST **15** (2020). 2 citations on pages 54 and 240

[130] R. Achenbach *et al.*; "The ATLAS Level-1 Calorimeter Trigger"; Technical report; CERN; Geneva (2008). https://cds.cern.ch/record/1080560. Cited on page 54

[131] ATLAS Collaboration; "Performance of the ATLAS Level-1 topological trigger in Run 2"; Eur. Phys. J. C **82**, p. 7 (2022). Cited on page 54

[132] W. Buttinger; "The ATLAS Level-1 Trigger System"; Technical report; CERN; Geneva (2012). https://cds.cern.ch/record/1456546. Cited on page 55

[133] S. Mehlhase; "ATLAS detector slice (and particle visualisations)"; (2021). 2 citations on pages 56 and 240

[134] ATLAS Collaboration; "ATLAS Tracking Software Tutorial"; (2024). 5 citations on pages 57, 58, 59, 62, and 240

[135] R. Frühwirth; "Application of Kalman filtering to track and vertex fitting"; Nucl. Instrum. Methods Phys. Res. A **262**, pp. 444–450 (1987). Cited on page 57

[136] ATLAS Collaboration; "Performance of the reconstruction of large impact parameter tracks in the ATLAS inner detector"; Technical report; CERN; Geneva (2017). https://cds.cern.ch/record/2275635. Cited on page 59

[137] ATLAS Collaboration; "Early Inner Detector Tracking Performance in the 2015 data at $\sqrt{s} = 13$ TeV"; Technical report; CERN; Geneva (2015). https://cds.cern.ch/record/2110140. 3 citations on pages 60, 61, and 240

[138] ATLAS Collaboration; "Reconstruction of primary vertices at the ATLAS experiment in Run 1 proton-proton collisions at the LHC"; Eur. Phys. J. C **77** (2017). 2 citations on pages 61 and 62

[139] W. Lampl *et al.*; "Calorimeter Clustering Algorithms: Description and Performance"; Technical report; CERN; Geneva (2008). https://cds.cern.ch/record/1099735. Cited on page 63

[140] ATLAS Collaboration; "Electron and photon energy calibration with the ATLAS detector using LHC Run 1 data"; Eur. Phys. J. C **74** (2014). Cited on page 63

[141] ATLAS Collaboration; "Electron and photon reconstruction and performance in ATLAS using a dynamical, topological cell clustering-based approach"; Technical report; CERN; Geneva (2017). `https://cds.cern.ch/record/2298955`. 2 citations on pages 63 and 64

[142] ATLAS Collaboration; "Performance of the ATLAS electromagnetic calorimeter for $\pi^0 \to \gamma\gamma$ and $\eta \to \gamma\gamma$ events"; Technical report; CERN; Geneva (2010). `https://cds.cern.ch/record/1273999`. Cited on page 63

[143] ATLAS Collaboration; "Electron and photon performance measurements with the ATLAS detector using the 2015-2017 LHC proton-proton collision data"; JINST **14** (2019). Cited on page 64

[144] ATLAS Collaboration; "Electron and photon efficiencies in LHC Run 2 with the ATLAS experiment"; JHEP **2024** (2024). 2 citations on pages 65 and 240

[145] S. Rettie; "Muon identification and performance in the ATLAS experiment"; Technical report; CERN; Geneva (2018). `https://cds.cern.ch/record/2626330`. 3 citations on pages 66, 68, and 240

[146] J. Illingworth & J. Kittler; "A survey of the hough transform"; CVGIP **44**, pp. 87–116 (1988). Cited on page 66

[147] ATLAS Collaboration; "Muon reconstruction and identification efficiency in ATLAS using the full Run 2 *pp* collision data set at $\sqrt{s} = 13$ TeV"; Eur. Phys. J. C **81** (2021). 3 citations on pages 67, 68, and 240

[148] M. Cacciari *et al.*; "The anti-ktjet clustering algorithm"; JHEP **2008**, p. 63 (2008). Cited on page 69

[149] S. Schramm; "ATLAS Jet Reconstruction, Calibration, and Tagging of Lorentz-boosted Objects"; Technical report; CERN; Geneva (2017). `https://cds.cern.ch/record/2291608`. Cited on page 69

[150] ATLAS Collaboration; "Reconstruction, Identification, and Calibration of hadronically decaying tau leptons with the ATLAS detector for the LHC Run 3 and reprocessed Run 2 data"; Technical report; CERN; Geneva (2022). `https://cds.cern.ch/record/2827111`. Cited on page 69

[151] ATLAS Collaboration; "The performance of missing transverse momentum reconstruction and its significance with the ATLAS detector using 140 fb^{-1} of $\sqrt{s} = 13$ TeV *pp* collisions"; (2024). `https://arxiv.org/abs/2402.05858`; 2402.05858. Cited on page 69

[152] CMS-TOTEM Collaboration; "CMS-TOTEM Precision Proton Spectrometer"; Technical report (2014). `https://cds.cern.ch/record/1753795`. 2 citations on pages 71 and 74

[153] ATLAS Collaboration; "AFP Figures"; https://twiki.cern.ch/twiki/bin/view/Atlas/AFP_Figures. 5 citations on pages 73, 74, 76, 78, and 240

[154] LHC Optics Working Group; "LHC Optics Web: LHC Run II pp physics optics"; https://abpdata.web.cern.ch/abpdata/lhc_optics_web/www/opt2017/. Cited on page 75

[155] M. Trzebiński; "Machine optics studies for the LHC measurements"; in "Photonics Applications in Astronomy, Communications, Industry, and High-Energy Physics Experiments 2014," , volume 9290, edited by R. S. Romaniuk ; International Society for Optics and Photonics (SPIE) (2014). Cited on page 75

[156] ATLAS Collaboration; "Performance of the ATLAS Forward Proton Spectrometer during High Luminosity 2017 Data Taking"; Technical report; CERN; Geneva (2024). https://cds.cern.ch/record/2890974. 10 citations on pages 75, 82, 83, 86, 88, 89, 91, 92, 237, and 241

[157] J. Lange *et al.*; "3D silicon pixel detectors for the ATLAS Forward Physics experiment"; JINST **10** (2015). Cited on page 77

[158] L. Nozka *et al.*; "Design of Cherenkov bars for the optical part of the time-of-flight detector in 44"; Opt. Express **22**, pp. 28984–28996 (2014). Cited on page 79

[159] R. Staszewski & J. Chwastowski; "Transport simulation and diffractive event reconstruction at the LHC"; Nucl. Instrum. Methods Phys. Res. A **609**, pp. 136–141 (2009). Cited on page 82

[160] R. Staszewski *et al.*; "Alignment-related Effects in Forward Proton Experiments at the LHC. Alignment-related Effects in Forward Proton Experiments at the LHC"; Nucl. Instrum. Methods Phys. Res., A **801**, pp. 34–43 (2015). 3 citations on pages 82, 84, and 241

[161] G. Valentino *et al.*; "Semiautomatic beam-based LHC collimator alignment"; Phys. Rev. ST Accel. Beams **15** (2012). Cited on page 88

[162] G. Valentino *et al.*; "Final implementation, commissioning, and performance of embedded collimator beam position monitors in the Large Hadron Collider"; Phys. Rev. Accel. Beams **20** (2017). Cited on page 88

[163] ATLAS Collaboration; "Public Forward Detector Plots for Collision Data"; https://twiki.cern.ch/twiki/bin/view/AtlasPublic/ForwardDetPublicResults. 2 citations on pages 90 and 241

[164] ATLAS Collaboration; "ATLAS data quality operations and performance for 2015–2018 data-taking"; JINST **15** (2020). Cited on page 90

[165] M. Kocian; “Readout and trigger for the AFP detector at ATLAS experiment”; JINST **12** (2017). Cited on page 90

[166] ATLAS Collaboration; “Performance of the ATLAS forward proton Time-of-Flight detector in Run 2”; JINST **19** (2024). 2 citations on pages 93 and 241

[167] ATLAS Collaboration; “Performance of the ATLAS Forward Proton Time-of-Flight Detector in 2017”; Technical report; CERN; Geneva (2021). https://cds.cern.ch/record/2749821. Cited on page 92

[168] T. Sykora; “ATLAS Forward Proton Time-of-Flight Detector: LHC Run2 performance and experiences”; JINST **15** (2020). Cited on page 93

[169] Wikimedia; “P-n junction Equilibrium Graphs”; (2007). https://commons.wikimedia.org/wiki/File:Pn-junction-equilibrium-graphs.png; under the CC-BY-SA-3.0. 2 citations on pages 97 and 241

[170] C. Poole & I. Darwazeh; “Chapter 11 - Microwave semiconductor materials and diodes”; in “Microwave Active Circuit Analysis and Design,” , edited by C. Poole & I. Darwazeh ; pp. 355–393 (Academic Press, Oxford) (2016); ISBN 978-0-12-407823-9. Cited on page 97

[171] B. G. Streetman & S. K. Banerjee; *Solid State Electronic Devices, 7th Edition* (Pearson). (2014). Cited on page 98

[172] M. Moll; “Displacement Damage in Silicon Detectors for High Energy Physics”; IEEE Trans. Nucl. Sci. **65**, pp. 1561–1582 (2018). 3 citations on pages 98, 103, and 113

[173] Jimblom; “Diodes”; (2022). https://learn.sparkfun.com/tutorials/diodes/real-diode-characteristics. 2 citations on pages 99 and 241

[174] V. Manzari; “Silicon Detectors Lecture 2”; (2015). https://indico.cern.ch/event/453690/sessions/99350/attachments/1184199/1726998/2015-11_SiliconDetectors_manzari_Lecture2.pdf. 2 citations on pages 99 and 241

[175] G. Lutz; *Semiconductor Radiation Detectors: Device Physics* (Springer Berlin) (2007). 2 citations on pages 100 and 128

[176] F. Hartmann; *Evolution of Silicon Sensor Technology in Particle Physics*; *Springer Tracts in Modern Physics*, volume 275 (Springer) (2017). 3 citations on pages 102, 103, and 242

[177] L. Gonella; “Radiation Hardness of Silicon Detectors”; (2022). https://indico.cern.ch/event/1129266/contributions/4821464/attachments/2445851/4191110/Si-radiation-hardness_1-2.pdf. 3 citations on pages 102, 103, and 242

[178] ATLAS Collaboration; "Technical Design Report for the ATLAS Inner Tracker Strip Detector"; Technical report; CERN; Geneva (2017). `https://cds.cern.ch/record/2257755`. 2 citations on pages 105 and 242

[179] ATLAS Collaboration; "Technical Design Report for the ATLAS Inner Tracker Pixel Detector"; Technical report; CERN; Geneva (2017). `https://cds.cern.ch/record/2285585`. Cited on page 105

[180] ATLAS Collaboration; "ITk Pixel Layout Updates"; (2020). `https://atlas.web.cern.ch/Atlas/GROUPS/PHYSICS/PLOTS/ITK-2020-002/`. 2 citations on pages 106 and 242

[181] I. Mandić *et al.*; "Annealing effects in n^+–p strip detectors irradiated with high neutron fluences"; Nucl. Instrum. Methods Phys. Res. A **629**, pp. 101–105 (2011). Cited on page 126

[182] M. Moll; "Radiation damage in silicon particle detectors: Microscopic defects and macroscopic properties"; (1999). Cited on page 126

[183] L. Wiik-Fuchs *et al.*; "Annealing studies of irradiated p-type sensors designed for the upgrade of ATLAS phase-II strip tracker"; Nucl. Instrum. Methods Phys. Res. A **924**, pp. 128–132 (2019). Cited on page 126

[184] R. Marco-Hernández *et al.*; "ALIBAVA: A portable readout system for silicon microstrip sensors"; (2007). Cited on page 127

[185] CMS Collaboration; "A search for new physics in central exclusive production using the missing mass technique with the CMS detector and the CMS-TOTEM precision proton spectrometer"; Eur. Phys. J. C **83**, p. 827 (2023). 10 citations on pages 132, 134, 141, 191, 200, 201, 202, 203, 238, and 249

[186] ATLAS Collaboration; "Search for new phenomena in final states with large jet multiplicities and missing transverse momentum using $\sqrt{s} = 13$ TeV proton-proton collisions recorded by ATLAS in Run 2 of the LHC"; JHEP **2020** (2020). Cited on page 133

[187] ATLAS Collaboration; "Search for electroweak production of charginos and sleptons decaying into final states with two leptons and missing transverse momentum in $\sqrt{s} = 13$ TeV pp collisions using the ATLAS detector"; Eur. Phys. J. C **80** (2020). Cited on page 133

[188] ATLAS Collaboration; "Search for direct production of electroweakinos in final states with one lepton, jets and missing transverse momentum in pp collisions at $\sqrt{s} = 13$ TeV with the ATLAS detector"; JHEP **2023** (2023). Cited on page 133

[189] CMS Collaboration; "Search for physics beyond the standard model in final states with a lepton and missing transverse energy in proton-proton collisions at $\sqrt{s} = 8$ TeV"; Phys. Rev. D **91** (2015). Cited on page 133

[190] CMS Collaboration; "Search for new physics in the lepton plus missing transverse momentum final state in proton-proton collisions at 13 TeV center-of-mass energy"; Technical report; CERN; Geneva (2021). https://cds.cern.ch/record/2758647. Cited on page 133

[191] CMS Collaboration; "Search for new physics in the τ lepton plus missing transverse momentum final state in proton-proton collisions at $\sqrt{s} = 13$ TeV"; JHEP **2023** (2023). Cited on page 133

[192] CMS Collaboration; "Observation of proton-tagged, central (semi)exclusive production of high-mass lepton pairs in pp collisions at 13 TeV with the CMS-TOTEM precision proton spectrometer"; JHEP **2018** (2018). Cited on page 134

[193] CMS-TOTEM Collaboration; "Search for central exclusive production of top quark pairs in proton-proton collisions at $\sqrt{s} = 13$ TeV with tagged protons"; JHEP **2024** (2024). Cited on page 134

[194] CMS-TOTEM Collaboration; "Search for high-mass exclusive diphoton production with tagged protons in proton-proton collisions at $\sqrt{s} = 13$ TeV"; Phys. Rev. D **110** (2024). Cited on page 134

[195] CMS-TOTEM Collaboration; "Search for high-mass exclusive $\gamma\gamma \to WW$ and $\gamma\gamma \to ZZ$ production in proton-proton collisions at $\sqrt{s} = 13$ TeV"; JHEP **2023** (2023). Cited on page 134

[196] A. Alloul *et al.*; "FeynRules 2.0 - A complete toolbox for tree-level phenomenology"; Comput. Phys. Commun. **185**, pp. 2250–2300 (2014). Cited on page 144

[197] C. Degrande *et al.*; "UFO - The Universal FeynRules Output"; Comput. Phys. Commun. **183**, pp. 1201–1214 (2012). Cited on page 144

[198] "Tagging and suppression of pileup jets with the ATLAS detector"; Technical report; CERN; Geneva (2014). https://cds.cern.ch/record/1700870. Cited on page 150

[199] S. Pagan Griso *et al.*; "Low-p_{T} tracking in high pile-up environments with EventIndex"; Technical report; CERN; Geneva (2024). https://cds.cern.ch/record/2897513. 2 citations on pages 162 and 204

[200] ATLAS Collaboration; "Luminosity determination in pp collisions at $\sqrt{s} = 13$ TeV using the ATLAS detector at the LHC"; (2022). Cited on page 172

[201] ATLAS Collaboration; "Measurement of the exclusive $\gamma\gamma \to \mu^+\mu^-$ process in proton-proton collisions at $\sqrt{s} = 13$ TeV with the ATLAS detector"; Phys. Lett. B **777**, pp. 303–323 (2018). 3 citations on pages 183, 184, and 247

[202] K. Cranmer; "Practical Statistics for the LHC"; (2015). https://arxiv.org/abs/1503.07622; 1503.07622. Cited on page 188

[203] G. Cowan *et al.*; "Asymptotic formulae for likelihood-based tests of new physics"; Eur. Phys. J. C **71** (2011). 2 citations on pages 188 and 191

[204] K. Cranmer *et al.*; "HistFactory: A tool for creating statistical models for use with RooFit and RooStats"; Technical report; New York U.; New York (2012). https://cds.cern.ch/record/1456844. Cited on page 189

[205] W. Verkerke & D. Kirkby; "The RooFit toolkit for data modeling"; (2003). https://arxiv.org/abs/physics/0306116; physics/0306116. Cited on page 189

[206] L. Moneta *et al.*; "The RooStats Project"; (2011). https://arxiv.org/abs/1009.1003; 1009.1003. Cited on page 189

[207] X. Qian *et al.*; "The Gaussian CL_s method for searches of new physics"; Nucl. Instrum. Meth. A **827**, pp. 63–78 (2016). Cited on page 191

[208] ATLAS Collaboration; "Modelling and computational improvements to the simulation of single vector-boson plus jet processes for the ATLAS experiment"; JHEP **2022** (2022). Cited on page 212

List of Tables

4.1 Naming conventions, locations relative to the ATLAS IP and acceptances of each AFP station during 2017. The acceptance range corresponds to values of ξ with at least 80% proton reconstruction efficiency. 73

4.2 The RP position parameters $x_{\mathrm{RP}}\,(s,t)$ used at different points throughout 2017 data-taking. The RMS beam width σ is around 200 (100) μm at the position of the Near (Far) station [156]. 88

5.1 Results of R_{int} measurements made in Birmingham and Toronto for test chips from the same irradiation, at different temperatures. 123

5.2 Results of R_{int} measurements of the test chip VPX37415-W174 in normal conditions and after extensive drying and annealing. 124

6.1 Overview of the simulated signal samples considered in the analysis, the corresponding programs used to perform the generation and UEPS simulation steps, applied generator-level filters and the generated ranges of hypothesised signal masses. 145

6.2 Summary of object pre-selection for candidate events in the electron and muon channels. 151

6.3 Summary of the selection for signal events in each signal region, in addition to the pre-selection detailed in Table 6.2. 154
6.4 Summary of fiducial volume selection criteria for signal events. 158
6.5 Summary of all considered systematic uncertainties for the signal and background models, and their effects on the event yields of the corresponding samples. For the signal samples each generator is shown separately, and the range of absolute observed changes in event yield across all generated signal masses is given. 185
6.6 Binning used for signal models with each considered hypothesised signal mass, in each signal region. Wider bins are used in the low+high-p_{T} track veto signal region to compensate for reduce statistics. 190
6.7 Summary of fiducial volume selection criteria for signal events in the CMS missing mass analysis [185]. 201

A.1 Overview of the simulated samples used to model the central dilepton component of the combinatorial background for the analysis, the corresponding programs used to perform the generation and UEPS simulation steps, and applied generator-level filters. 206

List of Figures

2.1 Fundamental particles of the Standard Model (SM) of particle physics [4]. 5

2.2 Example (a) Leading Order (LO) and (b) and (c) Next-to-Leading Order (NLO) Feynman diagram outlines. 13

2.3 Compilation of exclusion limits at 95% CL in the ALP-photon coupling ($1/\Lambda_a$) versus ALP mass (m_a) plane obtained by different experiments, adapted from [39], assuming a 100% ALP decay branching fraction into photons. The phase space probed by the analysis presented in this thesis is shown in red. Recent results from measurements of light-by-light scattering in collisions between heavy nuclei (Pb) are shown from both the ATLAS experiment, the experiment used in this thesis, and the CMS experiment, an equivalent experiment at the same facility. Additionally, results from proton-proton (pp) collisions measured by the same experiments are shown, in a similar phase space to this analysis. 18

2.4 Summary of several SM total and fiducial production cross-section measurements [72]. 22

2.5 Illustration of the principle behind the Equivalent Photon Approximation (EPA). The electric field around a proton at ultra-relativistic speeds becomes compressed such that it resembles a coherent flux of photons. Adapted from [3]. 24

2.6 Illustration of the three scenarios for photon-induced central exclusive production of some central state C, in the (a) elastic, (b) single-dissociative and (c) double-dissociative channels. 27

2.7 Representation of $t\bar{t}H$ production simulated in a Monte-Carlo (MC) event generator, showing the initial hard interaction (central red blob) and resulting quark and Higgs decays (smaller red blobs) and QCD radiative parton showers (red). Overlaid secondary interactions are also shown (purple), with all final state partons below the energy threshold hadronising (light green), before subsequently decaying (dark green). Additional photon radiation occurring at various stages is also shown (yellow) [86]. 30

3.1 The full CERN accelerator complex as of 2022 [100]. 34

3.2 Luminosity-weighted distribution of the mean number of interactions per crossing for the 2017 pp collision data at 13 TeV centre-of-mass energy. [109] 37

3.3 Long term schedule for the LHC and future HL-LHC as of January 2025 [110]. 38
3.4 Layout of the ATLAS detector [111]. 40
3.5 Layout of the ATLAS Inner Tracking Detector composed of the Pixel Detector, SCT and TRT [112]. 42
3.6 Layout of the ATLAS calorimeters composed of the LAr and Tile Calorimeters [116]. 45
3.7 Cross section of a barrel module from the LAr calorimeter [117], with three layers of decreasing granularity. 47
3.8 Schematic of a barrel module from the Tile calorimeter [119]. 48
3.9 Layout of the ATLAS Muon Spectrometer [120]. 50
3.10 The positions of the four forward detectors of the ATLAS experiment, used to measure particles with high-$|\eta|$ and make luminosity measurements [123]. 52
3.11 The ATLAS TDAQ system in Run 2 showing the components relevant for triggering as well as the detector read-out and data flow [129]. . . . 54
3.12 Summary of ATLAS experiment reconstruction principles for each particle type [133]. 56
3.13 Global track parameters with respect to the perigee representation [134]. 57
3.14 Illustration of the track reconstruction process, with red layers for the pixel detector, blue layers for the SCT and red circles showing Silicon layer hits in the ID. Taken from [134]. 58
3.15 Track reconstruction efficiency for each available working point as a function of (a) p_{T} and (b) η [137]. 60
3.16 Electron reconstruction efficiencies for each (a) identification and (b) isolation working point, determined from $Z \to e^+e^-$ decays as a function of transverse energy [144]. 65
3.17 Muon reconstruction categories in the ATLAS detector [145]. 66
3.18 Muon reconstruction efficiencies for (a) each identification working point and (b) the combined identification and isolation working points used in this thesis, determined from $J/\psi \to \mu^+\mu^-$ and $Z \to \mu^+\mu^-$ decays as a function of transverse momentum [147]. 68

4.1 General scheme of the AFP spectrometer detectors [153]. 73
4.2 Diagram of a Far station Roman Pot (RP) module with SiT and ToF detectors mounted [153]. 74
4.3 Photo of the SiT and ToF detectors for a single Far station mounted on the RP. Adapted from [153]. 76
4.4 Simulation of radiation fluence through the SiT at 212 m [127]. 76
4.5 Diagram of the 3D silicon pixel design used for the ATLAS IBL and AFP SiT detectors [127]. 77
4.6 (a) Demonstration of the proton vertex z position reconstruction using ToF measurements on both sides of the AFP spectrometer and (b) diagram of a single LQ-bar used in ToF [153]. 78

4.7 Simulated dispersion of proton x and y positions in the AFP SiT planes for particular values of proton ξ and p_{T} (a) for a particular station and (b) compared between the Near and Far stations on a given side, for equally spaced values of azimuthal scattering angle [156]. 82
4.8 Simulated difference Δx between proton x positions measured in the Near and Far stations on AFP Side A and the x coordinate measured in the Near station, as a function of proton ξ and p_x, demonstrating the unique mapping existing between these properties [156]. . 83
4.9 ξ and p_{T} acceptance of the AFP spectrometer in standard running conditions [160]. 84
4.10 Local coordinate system used for the AFP SiT planes, where the origin is defined as the corner of the first SiT plane in a station. Adapted from [3]. 85
4.11 Evolution of the alignment parameter corresponding to translation in the y axis over 20 iterations of the inter-plane alignment procedure, with final values for each plane in μm. Each offset is calculated with respect to plane 0 to remove dependence on the global alignment [156]. 86
4.12 Example of the method for determining the in-situ correction for global alignment, showing the $x_{\mathrm{AFP}} - x_{\mu\mu}$ distribution in the AFar station for exclusive dimuon events (a) before and (b) after applying the correction. The raw signal distribution (left) is fitted to a Gaussian and the fitted mean is taken as the correction [156]. 89
4.13 Luminosity recorded over time by the ATLAS detector and AFP spectrometer in 2017 [163]. 90
4.14 Proton reconstruction efficiencies for each AFP station determined from the "tag-and-probe" study, throughout 2017 data-taking as a function of ATLAS run number. The uncertainties shown are statistical [156]. 91
4.15 Distributions of the difference between ξ measurements from protons in the AFP spectrometer and from the dilepton system in the ATLAS detector in exclusive dilepton production events in the muon channel [156]. 92
4.16 The z_{ATLAS}–z_{ToF} distribution in double-tagged events measured in the ATLAS detector in Run 341419. The excess of signal (purple) over the background (yellow) is due to double-tagged events with both protons originating in the same signal process [166]. 93
4.17 ToF train efficiencies determined using AFP calibration stream data in ATLAS Run 331020 in the (a) AFar and (b) CFar stations, in events with exactly one reconstructed SiT track [166]. 93

5.1 Diagram showing the charge carrier concentration across a p–n junction, and the resulting electric field and built-in voltage [169]. 97
5.2 The I–V behaviour of a diode for a large range of forward and reverse bias voltages. Adapted from [173]. 99
5.3 A basic silicon strip sensor using a reverse biased p–n junction with (a) DC and (b) AC coupling. Based on [174]. 99

5.4 (a) Defects which form in the silicon lattice due to bulk damage. [Based on Figure 2.1 from [176]] and (b) the effects due to the resulting new energy levels in the silicon band gap [176, 177]. 102

5.5 Simulation of the ITk layout, comprised of silicon pixel and strip layers, with a barrel and endcap structure [178]. 105

5.6 Schematic of the ITk layout, with pixel layers shown in red and strip layers shown in blue [180]. 106

5.7 Test chip and MD8 layout. 108

5.8 (a) Test chip bench with the different measuring instruments labelled and (b) inside of the climate chamber where the box holding a wire-bonded test chip is placed for measurements. 109

5.9 Example MD8 I–V curve and corresponding k-factor of an unirradiated test chip, showing early breakdown at $V_{\text{bias}} = 434$ V and relatively high leakage current of 0.58 μA/cm^2 at 500 V, failing both specifications for unirradiated test chips. 110

5.10 Example MD8 I–V curve and corresponding k-factor of an irradiated test chip, showing no breakdown below $V_{\text{bias}} = 700V$ and relatively low leakage current of 20.64 μA/cm^2 at 500 V, passing both specifications for irradiated test chips. 110

5.11 Leakage current measured at $V_{\text{bias}} = 500$ V for all irradiated MD8s investigated by the author, plotted against irradiation date. Combined results are shown for batches of test chips which were irradiated together, with the number of test chips in each batch indicated alongside the date of each irradiation. The locations of each irradiation are also indicated, with the CERN IRRAD and Birmingham MC40 irradiating using protons and Ljubljana using a combination of neutrons and photons intended to be equivalent. For each batch, the result is presented as a box plot, with the mean, minimum and maximum measurements and the standard deviation all shown, as demonstrated to the right of the plot. 111

5.12 MD8 C–V showing two linear fits of the regions above and below V_{FD}, with their intersection giving a measurement of V_{FD}. 112

5.13 (a) Example I–V plots for each bias resistor on an irradiated test chip, with extremely similar results obtained for each structure, such that the resulting distributions overlap and not all are visible. (b) R_{bias} measurements for every test chip investigated by the author, as a function of irradiation date. Combined results giving the mean, minimum and maximum measurements and the standard deviation are shown for batches of test chips which were irradiated together, with the number of test chips in each batch indicated alongside the date of each irradiation. The location and particle type of each irradiation are also indicated. 114

5.14 Diagram of an interdigitated structure, with two sets of isolated n-type strip implants arranged on top of a p-type bulk such that each strip is neighboured on either side by strips from the other set. The structure is surrounded by a bias ring to allow a voltage to be applied to the bulk, and then a guard ring to isolate the structure from other test chip components. Connecting pads (purple, numbered) are used to make electrical connection to each set of strips through Pads 1-3 and 10-12, respectively, and to the bias ring via Pads 5-7. 115

5.15 (a) Example inter-strip I-V curves for each interdigitated structure on an irradiated test chip and (b) R_{int} measurements for every test chip investigated by the author, split by structure type, as a function of irradiation date. Combined results giving the mean, minimum and maximum measurements and the standard deviation are shown for batches of test chips which were irradiated together, with the number of test chips in each batch indicated alongside the date of each irradiation. The location and particle type of each irradiation are also indicated. 116

5.16 (a) Example inter-strip C–V plots for each interdigitated structure on an irradiated test chip and (b) C_{int} measurements for every test chip investigated by the author, split by structure type, as a function of irradiation date. Combined results giving the mean, minimum and maximum measurements and the standard deviation are shown for batches of test chips which were irradiated together, with the number of test chips in each batch indicated alongside the date of each irradiation. The location and particle type of each irradiation are also indicated. 117

5.17 (a) Example leakage current I-V curve for the coupling capacitor on an irradiated test chip and (b) I_{CPL} measurements for every test chip investigated by the author, as a function of irradiation date. Combined results giving the mean, minimum and maximum measurements and the standard deviation are shown for batches of test chips which were irradiated together, with the number of test chips in each batch indicated alongside the date of each irradiation. The location and particle type of each irradiation are also indicated. 119

5.18 Diagram of the coupling capacitor, showing the n-type implant deposited on top of the p-type bulk. Connecting pads (purple, numbered) are used to make electrical connections, with the p and n-type sides of the capacitor connected to Pads 10 and 11, respectively, and additional unconnected reference pads provided to allow a reference measurement of the capacitance of the measurement setup to be obtained. 120

5.19 C_{CPL} measurements for every test chip investigated by the author, as a function of irradiation date. Combined results giving the mean, minimum and maximum measurements and the standard deviation are shown for batches of test chips which were irradiated together, with the number of test chips in each batch indicated alongside the date of each irradiation. The location and particle type of each irradiation are also indicated. 120

5.20 Diagram of the PTP structure, composed of n^+-type strip implants on top of a p-type bulk, surrounded by a bias ring which allows a voltage to be applied to the bulk. Each strip is connected to the bias ring via a bias resistor (red), with connecting pads (purple, numbered) to allow electrical connections to the bias ring and each individual strip. The whole structure is surrounded by a guard ring to isolate it from other test chip components. 121

5.21 (a) Example I–V plots and (b) calculated R_{eff} for each PTP structure measured on an irradiated test chip. Extremely similar results are obtained for each structure, such that the plotted distributions overlap and not all are visible. 122

5.22 V_{PTP} measurements for every test chip investigated by the author, as a function of irradiation date. Combined results giving the mean, minimum and maximum measurements and the standard deviation are shown for batches of test chips which were irradiated together, with the number of test chips in each batch indicated alongside the date of each irradiation. The location and particle type of each irradiation are also indicated. 123

5.23 Results of R_{int} measurements after a series of total annealing times at room temperature (21°C) and low humidity (5%). 126

5.24 Mini sensor measurement setup, with the AliBaVa daughter board placed in the freezer, along with a ^{90}Sr β^- source, with the AliBaVa motherboard also connected to a scintillator and high voltage sourcemeter. 128

5.25 Collected charge spectra obtained for 100,000 electron hits on a mini sensor with the backplane reverse biased to (a) 300 V and (b) 1000 V, with fitted Landau ⊗ Gaussian curves shown in red, and the corresponding best-fit parameters given. 129

5.26 Linear fit of uncorrected peak collected charge in ADC counts as outputted by the read-out chip as a function of the measurement temperature. 130

5.27 Temperature-corrected most probable collected charge in electron units plotted as a function of voltage for three irradiated mini sensors A, B and C, before and after annealing. A and B were irradiated with protons at Birmingham and C was irradiated using neutrons at the Jožef Stefan Institute (JSI) in Ljubljana. 130

5.28 Most probable collected charge for every mini sensor investigated by the author in Birmingham, with a reverse bias voltage of 500 V applied. The irradiation site and particle type of each sensor is shown, with most irradiated using 27 MeV protons in Birmingham, and the rest irradiated using neutrons at JSI in Ljubljana. The specification requirements are highlighted, with all measured sensors passing the requirements. 131

6.1 Feynman diagram of the signal process $pp \to p\,(\gamma\gamma \to \ell\ell + X)\,p$ considered in this analysis. 134
6.2 Missing mass acceptance region of the AFP spectrometer, where the red line is a rough estimate of the maximum missing mass value which can be obtained as a function of the highest proton ξ measured in an event, allowing the upper limit on m_X acceptance of the AFP spectrometer to be extrapolated from its ξ acceptance. The corresponding limit for the tightened ξ range considered in this analysis is additionally shown. 137
6.3 Event mixing procedure used to produce the data-driven model of the combinatorial background for this analysis, with an example event shift of $i = 2$. 138
6.4 Feynman diagram for $Z + X$ production via a four-point photon interaction. 141
6.5 Effect of turning soft-survival effects on and off in a SuperChic simulated sample of exclusive dilepton production $\gamma\gamma \to \ell\ell$, as a function of the mass of the central system $m_{\ell\ell}$. 142
6.6 (a) The simulated dependence on the soft-survival probability for exclusive dilepton production as a function of the central rapidity $y_{\ell\ell}$ [80], equivalent to y_{ZX} in the current analysis, for EE events as used here as well as SD and DD events. (b) The rapidity distribution for the SuperChic $Z+X$ signal model with $m_X = 700$ GeV in the muon channel. 142
6.7 Representative loop-induced Feynman diagrams contributing to $Z + H'$ production. 143
6.8 Feynman diagram for di-ALP production via photon fusion. 144
6.9 Unit normalised comparison of generator-level kinematic distributions generated for each signal model with a hypothesised signal mass of $m_X = 500$ GeV. For the MadGraph di-ALP model, m_{S_1} is set to the Z boson mass. 146
6.10 The size of the luminous region in the ATLAS detector during $\sqrt{s} = 13$ TeV pp collisions in Run 2. The data points are the hourly average of results of a maximum likelihood fit to the spatial distribution of primary vertices collected over a two minute period. The luminosity weighted average size is provided for each year. 147
6.11 Width of the reconstructed missing mass distribution in simulated signal samples as a function of hypothesised signal mass, for each signal model and in each lepton channel. 148

6.12 Examples of the different possible scenarios for a given event of how many protons are detected on each side of the AFP spectrometer. . . 152
6.13 Example event with more than one loose proton per side, but exactly one tight proton per side, which passes the signal selection but fails a pre-selection requiring exactly one loose proton per side. 153
6.14 Comparison of the ratio of the expected cross section upper limits obtained for each signal model between progressively increasing dilepton p_{T} cuts. 155
6.15 (a) Missing mass distributions and (b) average missing mass (m_X) versus average interactions per bunch crossing (μ) separated by data-taking period, with the loose $0.02 < \xi < 0.12$ proton selection applied. The different normalisations in (a) are expected due to the total integrated luminosity differing between each data-taking period. The error bars in (b) correspond to errors on the mean values. 156
6.16 (a) Missing mass distributions and (b) average missing mass (m_X) versus average interactions per bunch crossing (μ) separated by data-taking period, with the tightened $0.035 < \xi < 0.08$ proton selection applied. 157
6.17 Fiducial selection efficiency of each signal model in the (a) muon and (b) electron channel. Efficiencies calculated from additional SuperChic samples without generator level filters applied are overlaid to verify that the filter efficiency is correctly accounted for. 159
6.18 The estimated efficiency for signal and background processes in data with respect to the track veto selection with different considered window sizes (corresponding to the minimum allowed track distance from the dilepton vertex). 160
6.19 Track veto signal efficiency observed directly in simulated signal samples as a function of signal mass, for each model and lepton channel. . 161
6.20 Number of data events as a function of the maximum number of additional high-p_{T} tracks with 0.5 mm of the dilepton vertex. 163
6.21 Comparison of background models produced using normal data versus filtered data with a cut of $N^{0.5\ \mathrm{mm}}_{p_{\mathrm{T}}>500\ \mathrm{MeV}} \leqslant 15$ imposed on the number of additional high-p_{T} tracks within 0.5 mm of the dilepton vertex, using different numbers of summed event-mixed samples. Results are shown for the muon channel. The red and blue lines on the ratio plots show the statistical uncertainty ($\sqrt{N}$) of the corresponding distributions in the top panels. The same level of agreement is observed in the electron channel. 164
6.22 Comparison of estimated background and signal event yields in simulated Z+jets events obtained when applying high-p_{T} versus low+high-p_{T} track vetoes for different considered window sizes. The signal event yield is scaled to an arbitrary cross section, as only the ratio between the two veto yields is considered. 166
6.23 Comparison of signal events to the dominant combinatorial background process for the analysis. 167

6.24 Reconstructed missing mass m_X distributions from the combined simulated background model produced with all considered background contributions, after all signal selections are applied except for the track veto and dilepton p_{T} cut, in the (a) muon and (b) electron channels. 169
6.25 Example of an event with both signal protons correctly measured by the AFP spectrometer (left) and an event where one of the signal protons is missed by the AFP spectrometer and a single pile-up proton from a background interaction is measured in its place (right). 170
6.26 Comparison of lepton kinematic distributions in events with positive versus negative reconstructed missing mass. 171
6.27 Comparison of track veto signal efficiencies estimated as a function of the mean number of interactions per bunch crossing μ, using a pile-up based approach and calculated directly from the lepton vertex, in simulated data in the (a) muon and (b) electron channels. The ratio is fitted to a quadratic polynomial to reduce statistical fluctuations, and the fit range is limited to within $\pm 2\sigma$ of the mean value of μ to remove outliers. 173
6.28 Number of reconstructed clusters per proton track in each AFP SiT station for double-station reconstructed protons across the nominal dataset. 177
6.29 Missing mass distributions in blinded data for each AFP systematic which affects proton reconstruction methodology (a) before applying the track veto selection and (b) after applying the high-p_{T} track veto. 178
6.30 Bin-by-bin statistical uncertainty determination for the event-mixed background model estimated using bootstrapping, for different numbers of averaged orthogonal event-mixed samples N. The uncertainties are shown for each bin in the top panel, with the fractional uncertainty (divided by the absolute value) shown in the bottom panel. 179
6.31 Difference $|\Delta\mu|$ between the average number of interactions per bunch crossing for the central event and the proton event used in the event-mixing procedure, for different total numbers of combined samples. . 180
6.32 Comparison between unit-normalised missing mass distributions (normalised) (a) before and (b) after track veto using events with low $|\Delta\mu| < 5$ and high $|\Delta\mu| > 10$. The expected statistical fluctuations in each case are shown in red/blue on the ratio plot. 181
6.33 Comparison between missing mass distributions (a) before and (b) after track veto when mixing both protons per event (double-mixing) and only on one side (single-mixing). The expected statistical fluctuations in each case are shown in red/blue on the ratio plot. 182
6.34 (a) Exclusive $\gamma\gamma \to \mu^+\mu^-$ differential fiducial cross-section measurements as a function of dimuon invariant mass $m_{\mu\mu}$. (b) Comparison of the ratios of measured and predicted cross-sections to the bare EPA calculations as a function of the average dimuon invariant mass scaled to the pp centre-of-mass energy used [201]. 184

6.35 Comparison of the missing mass distribution between blinded data and the data-driven background model in the (a) muon and (b) electron channels with all signal selections applied. 192

6.36 Upper limits on signal cross section set using blinded data with an event-shift of $i = 1$ for the SuperChic $Z + X$ signal model with the standard track veto applied in the (a) muon and (b) electron channels. 192

6.37 Comparison between unblinded data and the data-driven background model in the muon channel with only pre-selection cuts applied for distributions of (a) dilepton pair p_{T}, (b) proton ξ and (c) missing mass m_X. The expectations for a signal with a hypothesised mass of $m_X = 500$ GeV using the SuperChic $Z + X$ model are overlaid and normalised to a cross section of 10 fb. 193

6.38 Comparison between unblinded data and the data-driven background model in the combined lepton channel with the pre-selection and all signal region cuts applied, for distributions of (a) dilepton pair p_{T}, (b) proton ξ and (c) missing mass m_X. The expectations for a signal with a hypothesised mass of $m_X = 500$ GeV using the SuperChic $Z + X$ model are overlaid and normalised to a cross section of 10 fb. . 194

6.39 Summary of all pre and post-fit distributions for the signal+background model for each tested signal mass, in the combined lepton channel, for the (a) SuperChic $Z + X$, (b) MadGraph di-ALP and (c)MadGraph $Z + H'$ models. Fits use a mass window of 100 GeV either side of the hypothesised signal mass, with the exception of 100 GeV and 900 GeV models, which use larger windows of $0 \leq m_X \leq 300$ GeV and $700 \leq m_X \leq 1000$ GeV, respectively. All pre-fit signals are shown normalised to 25 fb. 195

6.40 Systematic pulls for the 500 GeV signal for each signal model in each channel in the low+high-p_{T} track veto signal region. 196

6.41 Systematic ranking plots for each signal model with $m_X = 500$ GeV, in the combined channel. The highest ranked systematics are those having the largest impact on the final value of signal normalisation, with the soft-survival uncertainty, track veto signal efficiency and AFP spectrometer optics and alignment consistently being the highest ranked. In the case of MadGraph models, the soft-survival uncertainty is applied as a downwards-only variation, resulting in an asymmetric pull. 197

6.42 Observed and expected upper limits on the signal cross section set for each signal model with the standard track veto applied, with lepton channels shown separated and combined, and with all systematics and scale factors included. 199

6.43 Ratio of the expected limits obtained using only high-p_{T} tracks ($p_{\mathrm{T}} >$ 500 MeV) to those obtained with both high- and low-p_{T} tracks ($p_{\mathrm{T}} >$ 100 MeV) included in the track veto. Ratios greater than one indicate improved sensitivity when low-p_{T} tracks are included. Results are shown separately for the (a) muon, (b) electron, and (c) combined channels. 200

6.44 Ratio between the expected limits obtained obtained with the high and low+high-p_{T} track vetoes applied, with both results using the low+high-p_{T} track veto binning. Ratios above 1 indicate an improvement with the low+high-p_{T} track veto applied. 201

6.45 Comparison of the observed limits for the SuperChic $Z + X$ model between this analysis and the CMS results from [185]. The limits from this analysis are scaled to match the fiducial region used in the CMS study, and soft-survival factors are not included. The expected limits obtained in this analysis with no track veto selection are overlaid. The comparison is shown for the (a) muon, (b) electron, and (c) combined channels. 202

A.1 Kinematic distributions after signal selection, except for track veto, from a misidentified lepton background model created using a same-sign lepton selection in data. 207

A.2 Feynman diagrams showing dilepton production via photon fusion in (a) exclusive $pp \to p(\gamma\gamma \to \ell\ell)p$ (b) single dissociative semi-exclusive $pp \to p^*(\gamma\gamma \to \ell\ell)p$ and (c) double dissociative $pp \to p^*(\gamma\gamma \to \ell\ell)p^*$ topologies [85]. 208

A.3 Generator-level ξ distributions of protons which remain intact following Elastic-elastic (EE) and Single Dissociative (SD) photon-induced dilepton production in the (a) muon and (b) electron channels. Distributions are shown following pre-selection and without the fiducial selection applied. 209

A.4 Kinematic distributions from the combined simulated background model produced with all significant background contributions included, after all signal selections are applied except for the track veto and dilepton p_{T} cut, in the muon channel. 210

A.5 Kinematic distributions from the combined simulated background model produced with all significant background contributions included, after all signal selections are applied except for the track veto and dilepton p_{T} cut, in the electron channel. 211

A.6 Kinematic distributions after the high-p_{T} track veto is applied, from the total simulated background model produced with all significant background contributions included and compared to data-driven model, in the muon channel. 213

A.7 Distributions after high-p_{T} track veto is applied, from total simulated background model produced with all considered background contributions and compared to data-driven model, in the electron channel. 214

B.1 Proportion of events with $m_X > 0$ with (a) both reconstructed protons matched to the truth-level signal protons and (b) at least one reconstructed proton mismatched to a pile-up proton, for each signal model as a function of mass. 216

B.2 Reconstructed missing mass distributions for SuperChic signal model at selected mass points for (a) all events, (b) only events with both reconstructed protons matched to the truth level signal protons, (c) only events with exactly one reconstructed proton which is not matched to a truth level signal proton and (d) only events with both reconstructed protons not matched to a truth level signal proton. . . 217
B.3 Proportion of events with $m_X > 0$ with (left) both reconstructed protons matched to the generator-level signal protons and (right) at least one reconstructed proton mismatched to a pile-up proton, when a mass window is applied 50 GeV either side of the hypothesised signal mass for a given model, for each signal model as a function of mass in each lepton channel. 218
B.4 Proportion of events with $m_X > 0$ with (left) both reconstructed protons matched to the truth-level signal protons and (right) at least one reconstructed proton mismatched to a pile-up proton, when the fiducial selection is applied, for each signal model as a function of mass in each lepton channel. 219

Definitions of acronyms

AC Alternating Current

ADC Analogue-to-Digital Conversion

AFP ATLAS Forward Proton

ALFA Absolute Luminosity For ATLAS

ALICE A Large Ion Collider Experiment

ALP Axion-like Particle

ATLAS A Toroidal LHC ApparatuS

BBA Beam-Based Alignment

BEH Brout-Englert-Higgs

BLM Beam Loss Monitor

BPM Beam Position Monitor

BSM Beyond the Standard Model

CB Combined

CCE Charge Collection Efficiency

CD Central Diffractive

CEP Central Exclusive Production

CERN Centre Européen pour la Recherche Nucléaire
European Center for Nuclear Research

CL Confidence Level

CMS Compact Muon Solenoid

CP Combined Performance

CSC Cathode Strip Chamber

CT Calorimeter-Tagged

CTP Central Trigger Processor

DAQ Data Acquisition

DC Direct Current

DD Double Dissociative

DPE Double Pomeron Exchange

DIS Deep Inelastic Scattering

DM Dark Matter

ECal Electromagnetic Calorimeter

EE Elastic-elastic

EM Electromagnetic

EPA Equivalent Photon Approximation

FCal Forward Calorimeter

FD Full Depletion

GRL Good Run List

HCal Hadronic Calorimeter

HEC Hadronic End Cap

HL-LHC High Luminosity LHC

HLT High-Level Trigger

IBL Insertable B-Layer

ID Inner Detector

IP Interaction Point

IRRAD Proton Irradiation Facility

ITk Inner Tracker

JVT Jet-Vertex Tagger

LAr Liquid Argon

LEIR Low Energy Ion Ring

LEP Large Electron-Positron

LHC Large Hadron Collider

LHCb Large Hadron Collider beauty

LINAC2 Linear Accelerator 2

LINAC3 Linear Accelerator 3

LINAC4 Linear Accelerator 4

L1 Level 1

LO Leading Order

LQ L-shaped Quartz

LS Long Shutdown

LUCID Luminosity Cherenkov Integrating Detector

MC Monte-Carlo

MCP-PMT Microchannnel Plate Photomultiplier Tube

MD8 Monitor Diode

MDT Monitored Drift Tubes

ME Muon Spectrometer Extrapolated

MET Missing Transverse Energy

MPI Multi-Parton Interactions

MIP Minimally Ionising Particle

MPI Multi-Parton Interaction

MS Muon Spectrometer

NIEL Non-Ionising Energy Loss

NLO Next-to-Leading Order

NP Nuisance Parameter

PCB Printed Circuit Board

PDF Parton Distribution Function

PID Particle Identification

PPS Precision Proton Spectrometer

PS Proton Synchrotron

PTP Punch-Through Protection

PV Primary Vertex

QA Quality Assurance

QCD Quantum Chromodynamics

QED Quantum Electrodynamics

QFT Quantum Field Theory

RoI Region of Interest

RF Radio Frequency

RMS Root Mean Square

RP Roman Pot

RPC Resistive-Plate Chamber

SA Stand-alone

SCR Space Charge Region

SCT Semi-conductor Tracker

SD Single Dissociative

SiT Silicon Tracker

SM Standard Model

SPS Super Proton Synchrotron

SSB Spontaneous Symmetry Breaking

ST Segment-Tagged

SUSY Supersymmetry

TDAQ Trigger and Data Acquisition

TGC Thin-Gap Chamber

TID Total Ionising Dose

ToF Time-of-Flight

TOTEM TOTal and Elastic Measurement

TRT Transition Radiation Tracker

UE Underlying Event

UEPS Underlying Event and Parton Shower

UV Ultraviolet

VEV Vacuum Expectation Value

YETS Year-end Technical Stop

ZDC Zero Degree Calorimeter