Towards the search for Charged Lepton Flavour Violation with The Mu3e Experiment

Sean Paul Hughes
201111055

Thesis submitted in accordance with the requirements of the University of Liverpool for the degree of Doctor in Philosophy by Sean Paul Hughes.

Under the supervision of
Dr. Nikolaos Rompotis and Prof. Joost Vossebeld

At the
Faculty of Science and Engineering
Department of Physics

September 2023

Towards the search for Charged Lepton Flavour Violation with The Mu3e Experiment

Sean Paul Hughes

Abstract

The Mu3e experiment aims to search for the lepton flavour violating decay $\mu^+ \to e^+e^+e^-$ to a sensitivity of one in 10^{16} muon decays, with zero background. The decay is effectively forbidden within the Standard Model. However, many models that go beyond the Standard Model predict the decay to occur at much higher branching fractions. The experiment faces many challenges, namely from numerous background sources and the high event rate. These challenges are minimised through the use of a novel detector design, which is comprised of a highly granular, thin silicon pixel detector, scintillators and a high intensity muon beam. The use of simulation is vital, and within this work, the software analysis framework used by Mu3e will be described in detail.

Within the context of this thesis, several simulation studies are conducted and the results are shown. These include Geant4 simulations of the Mu3e detector, accompanied by tracking studies involving only two layers. A study into a possible Phase-2 design is also explored, whereby an extra silicon layer is added at high radius. New tracks are defined, and their improved momentum resolutions are shown. A simulation of the main $\mu^+ \to e^+e^+e^-$ search using the Phase-1 geometry is also outlined. It is also shown that the large number of muon decays will allow Mu3e to be sensitive to new physics in the dark sector. Simulation studies of muons decaying to dark photons and axion-like particle are conducted, where it is shown that Mu3e will have a competitive sensitivity to new physics alongside the $\mu^+ \to e^+e^+e^-$ search.

Acknowledgements

This work would not have been possible without the help of a great number of people. I'd like to thank my supervisors, Nikos and Joost, both whom I greatly admire. Their guidance and inspiration throughout my PhD was incredibly vital and was much appreciated, and I will be forever thankful. I'd like to thank the Mu3e collaboration for all the support and help they have given me. I also thank the many inspiring people in the Oliver Lodge who made the PhD such an enjoyable time.

I'd especially also like to thank Rebecca for her indefatigable support. She is an incredible person and I could not possibly have succeeded without her. I'd finally like to thank my family, for their unwavering support throughout this endeavour.

Contents

Chapter 1

An Introduction to Muon Physics

At currently observed energy scales, the Standard Model of Particle Physics (SM) provides an excellent description of the fundamental matter particles and the interactions among them. Despite its many experimental successes however, it is incomplete, since it fails to include gravity, dark matter and dark energy, as well as a number of other physical processes that are believed to occur in nature.

This chapter provides an introduction to muon physics, showcasing the muon's place within the SM. The history of the muon is discussed, with particular emphasis on the many experimental observations involving the muon that have had vital impact on the development of the SM. An overview of the SM is outlined, alongside its many achievements and several of its perceived shortfalls. The chapter closes with how the muon can be used as a probe for new physics, through searches for charged lepton violating processes.

1.1 Biography of the Muon

The story of the muon begins with its discovery in 1937, in the study of cosmic rays. Using cloud chambers, a highly penetrating particle with properties distinct from the electron was discovered by Anderson and Neddermeyer in California [1]. The amount of ionisation present within the chamber inferred the presence of a particle with mass larger than an electron, and lower than a proton - but with the same electric charge. An early hypothesis suggested that this particle could be represented in some form as an excited state of the electron. The existence of the particle was confirmed in the same year [2], in an experiment which measured its mass to be around 130 times the electron mass. In 1946, the first measurement of the muon lifetime was measured to be (2.33 ± 0.15) μs [3].

The discovery of this particle coincided with work published in 1935 by Yukawa [4] hypothesising the existence of a force carrier to strong interactions, the pi-meson. Yukawa had theorised that if the pi-meson did exist, it would have a mass of around 100 MeV, which was similar to that of the newly discovered muon. The muon found by Anderson and Neddermeyer interacted feebly with matter, and was later confirmed to not interact at all with the 'strong' interaction [5]. The search for Yukawa's particle instead lead to the discovery of charged pions in the Peruvian Andes in 1947, by a University of Bristol team, lead by Powell [6].

The muon had so far been observed to decay exclusively to electrons and neutral particles. This observation lead to the hypothesis that the muon could indeed be a higher mass state of the electron. If true, the muon could de-excite, producing a ground state electron and emitting a photon. The first search for lepton flavour violation was thus conducted, in the search for the decay $\mu \to e\gamma$ in 1947 by Hincks and Pontecorvo [7]. The experiment concluded there was no photon emission, setting an upper limit of 10% to the branching fraction, establishing that the muon was indeed its own particle, disfavouring the possibility that it was an excited electron state.

The measured continuous energy spectrum of the electron emitted from the muon decay inferred the presence of a three-body decay, where two of the particles had a neutral charge [8], and had therefore not ionised the gas within the cloud chamber. This lead to the first neutrinoless muon capture experiment, which found no observable conversion of a muon into an electron [9]. The support to the idea that at least one of the two neutral particles was massive was waning, as the maximum energy of the electron spectrum was experimentally found to be close to 55 MeV [10]. This lead to the hypothesis that the two neutral particles were both neutrinos. The energy spectrum of the electron was calculated by L. Michel [11], assuming a four-fermion contact interaction. The decay has henceforth been named the 'Michel decay'.

The development of powerful cyclotrons [12] allowed for the production of muons at increasingly higher levels of energy and intensity. This allowed for more stringent limits to be set on, then unknown, flavour violating processes. Upper limits to the processes $\mu \to e\gamma$ and $\mu\mathrm{Cu} \to e\mathrm{Cu}$ were set to be $< 2 \times 10^{-5}$ [13] and $< 5 \times 10^{-4}$ [14] respectively, in 1955 by Steinberger. In Liverpool, the synchrocyclotron (1954 - 1968) allowed for the production of 383 MeV protons [15], which were used to set the most stringent upper limit at the time on flavour violating processes $\mu \to e\gamma$, $< 2.2 \times 10^{-7}$ in 1962 [16], and $\mu\mathrm{Cu} \to e\mathrm{Cu}$, $< 7.5 \times 10^{-6}$ in 1964 [17]. Figure 1.1 summarises the upper limits to charged lepton flavour violating decays published as a function of time.

Crucial for the development of the theory of nuclear β decays, Lee and Yang proposed that nature did not conserve helicity in weak interactions [18, 19], as a means of explaining the θ-τ anomaly [20, 21], and developing the theory of two-component neutrinos[1] in muon decays. In 1952, parity violation was observed by Wu in nuclear β decays using Co^{60} [22], which was confirmed two pages later, in the same volume of PRL, by Lederman [23]. The theory of the 'weak' interaction around this time was fast developing. The interaction was hypothesised to be mediated by a single massive vector boson [24]. As a consequence, this would lead to the branching fraction of $\mu \to e\gamma$ to be of the order 10^{-4}, and $\mu \to eee$ to 10^{-6} [25], the former of which had been previously excluded - thus excluding the theory. The decays $\mu \to e\gamma$ and

---

[1] Where fermions are represented by two-component spinors which are acted upon by helicity operators outputting arbitrary 'left handed' and 'right handed' states.

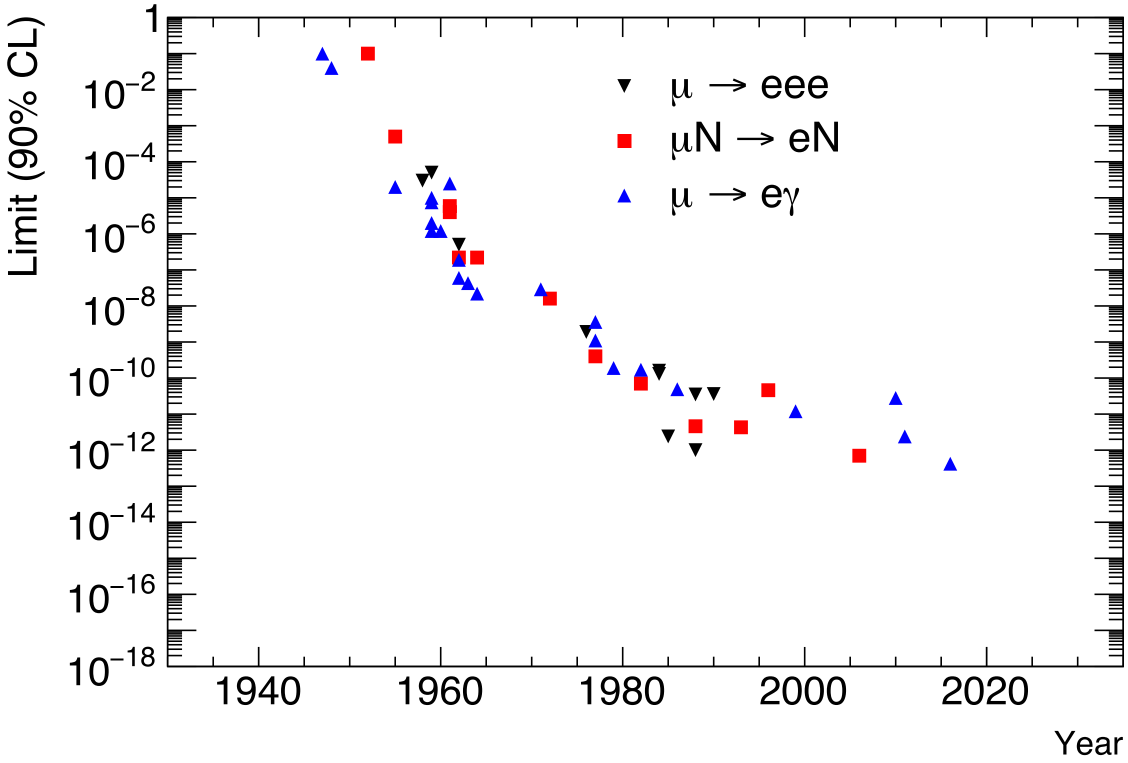

Figure 1.1: History of CLFV searches.

$\mu \rightarrow eee$ were concluded to be forbidden. The two-component neutrino theory prevailed, where from the muon decay two neutrinos differing by their helicity are produced [26, 27]. The consequence of which was the observation of a fundamental symmetry, the conservation of leptons. In Liverpool, this law was resoundingly confirmed through measurements of the polarisation of electrons emitted through muon decays by a team lead by Holt [28]. A later experiment by the same team measured the asymmetry between decaying positive and negative muons, proving that charge conjugation was also violated [29]. This was vital for the development of understanding for the weak interaction, as these experimental results established that the weak interaction is mediated by a vector and axial-vector (V-A), and not by a pseudoscalar tensor. The V-A nature has been recently tested by the TWIST collaboration, who have most recently conclusively demonstrated the V-A nature of the weak interaction [30, 31].

Beyond fundamental particle physics, investigations into how the muon could be used in material science was ongoing. In 1958, the parity violating nature of the weak interaction allowed for muon decays to be used to measure depolarisation factors of materials [32]. This would eventually lead to the development of a technique that is known today as muon-spin relaxation, rotation or resonance (μSR) [33]. The operating principle involves depositing polarised muons into a material, and measuring the asymmetry of the positron decay products. This allows for the analysis of a material's bulk structural properties with atomic precision. Also in 1958, the first measurement of the magnetic moment of the muon was conducted in Chicago, using a similar technique of depositing polarised muons into a target material [34]. The anomalous magnetic moment of the muon ($g-2$) [35] was measured further in experiments at CERN [36, 37] and Brookhaven National Laboratory (BNL) [38, 39, 40, 41, 42]. A discrepancy with the theoretical value calculated using radiative corrections is reported, and the experiment is now conducted at Fermilab National Accelerator Laboratory (Fermilab) [43], using the same detector apparatus as used at BNL. The most recent measurement used data collected by Fermilab, and a 5.0 σ discrepancy between the measured and theoretical value has been shown [44].

Aside from $g-2$ measurements, BNL also made further contributions to fundamental electroweak theory. In 1962, a neutrino beam was produced from a beam of decaying pions at BNL. The beam was directed to pass through matter, and the decay products from the neutrino interactions with matter were measured. It was observed that muons, and not electrons, were generated from the interaction between the neutrino beam and matter. This confirmed the existence of more than one type of neutrino [45], with the interaction occurring through a charged current. The experiment would have otherwise seen electrons and muons produced in equal abundance for the examined phase space. This extended the symmetry of the conservation of leptons. Leptons were observed to be conserved within their respective 'families'. Within each family there exists a quantum property known as 'flavour', and this is conserved in all interactions. Lepton flavour violation has since been conclusively demonstrated with the discovery of neutrino oscillations by Super-Kamiokande and SNO collaborations in 1998 and 2001 [46, 47]. The discovery of the W and Z bosons by UA1/2 experiments [48, 49] allowed for tests of lepton universality in the weak interaction, which was achieved by both the UA1/2 collaborations, but with UA1 using electrons and muons. Electroweak measurements were further tested using muons at LEP [50], where it was observed that measurements were mostly consistent with electroweak theory. This investigation found the deviation to be consistent with the SM [51]. The muon is also utilised in Higgs boson physics. The Yukawa coupling between the Higgs boson and second generation fermions can be measured through the $H \to \mu\mu$ channel, which has been measured by both the ATLAS and CMS collaborations [52, 53].

The muon is used as a tool for probing beyond the standard model physics in searches for charged lepton flavour violation in recent experiments, like MEG [54]. The muon has played a vital role in the development of understanding of nature, and in particular the SM. The development of particle physics therefore follows closely with the history of the muon. It remains to this day a tool with which we can achieve a greater understanding of nature.

1.2 The Standard Model of Particle Physics

The Standard Model of Particle Physics (SM) is a quantum field theory describing the fundamental constituents of matter, modelled as elementary fields, and their interactions. This is achieved using the gauge group $\mathrm{SU}(3)_C \times \mathrm{SU}(2)_L \times \mathrm{U}(1)_Y$, where C is for colour, L specifies left-handed spinors, and Y is the hypercharge. The theory describes three of the four known forces of nature: the weak, electromagnetic ($\mathrm{U}(2)_L \times \mathrm{U}(1)_Y$) and the strong ($\mathrm{SU}(3)_c$). The weak and electromagnetic forces are understood through what is known as the electroweak theory. The strong force is described using quantum chromodynamics (QCD).

Shown pictorially in Figure 1.2 are the elementary fields that compose the SM. The elementary fields are grouped according to their quantum numbers. A quantum number is an intrinsic property of the field, and often has no analogue in classical physics. For example: intrinsic angular momentum, which is referred to as 'spin'. One feature of quantum field theories is the ability to distinguish between two types of elementary field, differing only by its value of spin. A 'boson' is defined as a field with an integer value of spin, whilst a 'fermion' has a half integer value. Fermions are subject to Fermi-Dirac statistics, disallowing the same fermion to occupy identical quantum states. Bosons are subject to Bose-Einstein statistics, which on the contrary allow for the occupation of identical quantum states. The fermions present within the SM can be separated further into 'leptons' and 'quarks', depending on how they interact

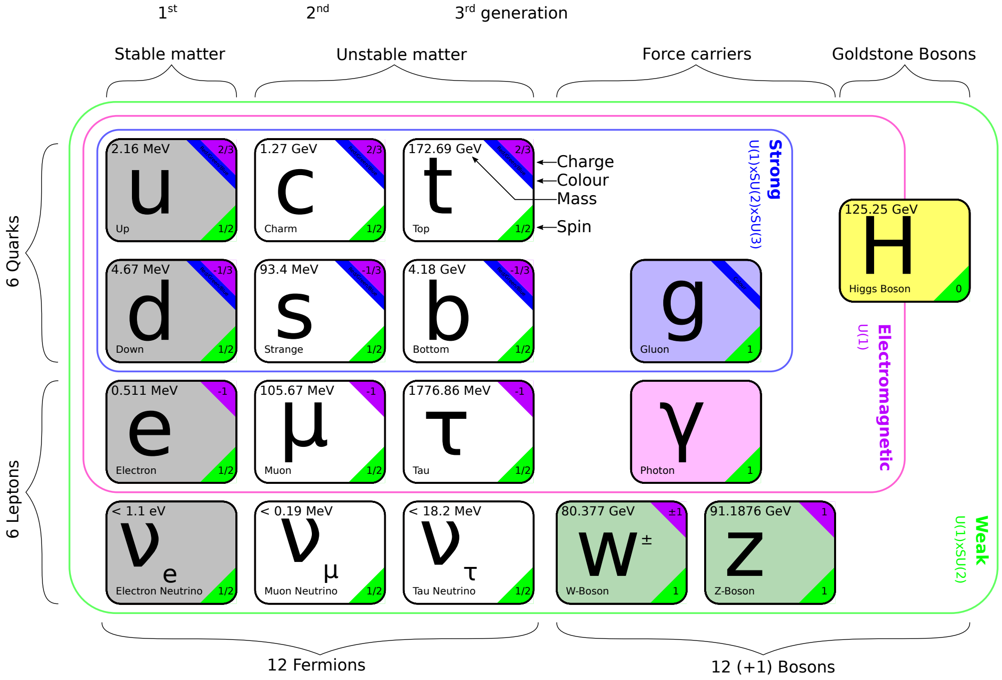

Figure 1.2: A pictorial view of the elementary particles within the Standard Model.

Within the SM, there are six 'quarks' [55, 56]: up (u), down (d), strange (s) [55], charm (c) [57, 58, 59], bottom (b) [60], and top (t) [61, 62]. There are also the 6 leptons: the electron (e), muon (μ) and the tau (τ), with their corresponding neutrinos: the electron neutrino (ν_e), muon neutrino (ν_μ), and the tau neutrino (ν_τ). Each fermion has also corresponding anti-particles, the quantum numbers of which are equal to the ordinary-matter value, multiplied by -1. There are four gauge bosons: the photon (γ), the gluon (g), and $W^\pm$ and Z particles, which are the force carriers of the electromagnetic, the strong, and the weak forces (charged and neutral current) respectively. In the model there is also one spin-zero (scalar) boson: the Higgs boson (H). Quarks and leptons are further separated into generations, of which the SM has only three, and are represented in columns in Figure 1.2.

Charged leptons like the electron, muon and tau particles interact with the electromagnetic and weak forces. The neutrinos do not carry electric charge, and are thus invisible to the electromagnetic force; they can only interact through weak currents, with other elementary fields. Quarks are fermions which carry electric and weak charge, akin to their lepton counterparts. Quarks interact with the strong interaction, as each quark carries ‘colour’ charge. Quantum chromodynamics (QCD) is the theory of the strong force [63, 64]. Only colourless objects have experimentally been observed in nature. Isolated quarks are not observed on their own, which in QCD is attributed to a property known as colour confinement. Colourless states can be formed with particular combinations of quarks. For example: baryons (composed of three quarks), and mesons (composed of quark-anti-quark pairs).

1.2.1 Electroweak theory and the Muon

The electroweak theory is formulated through the Glashow-Weinberg-Salam theory of weak interactions [65, 66, 67]. It combines both quantum electrodynamics and the weak interaction [63, 64, 68], introducing four vector bosons, three with mass and one without, and two coupling constants. The former introduces the photon as a force carrier to electromagnetism, whilst the latter introduces massive Z and $W^{\pm}$ bosons, responsible for mediating the weak interaction. The theory fully incorporates the weak interaction’s parity violating nature, through V-A currents. One feature of the electroweak interaction is the concept of lepton universality. From the observed decay widths of both the tau and muon, one can define the ratio of the relative coupling of both tau and muon particles to the electroweak gauge bosons. Experimentally, the ratio of the coupling strengths is consistent with one, showing that the weak interaction’s strength is independent of lepton flavour [63, 68]. This is a concept known as lepton universality, and is extensively tested by experiment [50, 69].

The muon has an electric charge, and so can interact through the electromagnetic force with other electrically charged particles. This can lead to the formation of muonium, which is a bound state of an electron and an anti-muon ($e^-\mu^+$) [70, 71]. One can calculate using quantum electrodynamics high precision features of the fundamental properties of the muon, which can be corroborated to great accuracy by experiment, thus providing a stringent test of the SM. One such feature is the anomalous magnetic moment of the muon, which has recently been measured to be 4.2 standard deviations away from the best estimate of this quantity from the SM [72].

1.2.2 Flavour

Formally, for leptons, one defines leptonic flavours for electrons (L_e), muons (L_μ) and tau (L_τ) particles. Using electron-type leptons as an example, the electron and its corresponding neutrino is given the value of L_e to be +1, and is said to belong to the same 'generation' or 'family'. The corresponding anti-particles are given -1. The sum of L_e (or L_μ or L_τ) is experimentally found, thus far, to be unchanged before and after an interaction. The weak interaction is responsible for the decay of particles, and thus for flavour change. However, ultimately the sum of each lepton quantum number is conserved.

The observation of neutrino oscillations solves the solar neutrino problem [73], and showed that not only do neutrinos carry mass, but that lepton flavour appears to only be an approximate symmetry. The neutrino flavours are found to be a superposition of three different mass eigenstates [74]. Furthermore, the eigenstates of the weak interaction are not the mass eigenstates, but are 'weak' or 'flavour' eigenstates. Mass eigenstates are written as a superposition of weak eigenstates, and for unique, non-zero neutrino masses, this allows for the neutrino to change phase and oscillate between flavours as it travels. The Pontecorvo-Maki-Nakagawa-Sakata (PMNS) [71, 75, 76] matrix describes the relationship between mass and flavour eigenstates of the three SM neutrinos.

For quarks, a similar phenomenon is observed for hadronic weak interactions. Flavour number is defined for each type of quark, and is conserved in the SM through strong and electromagnetic interactions, but not through the weak interaction. For example, one can define a 'strangeness' flavour number (S) for strange quarks. For normal-matter strange quarks, S is equal to +1, for anti-strange, S is -1. And likewise for charm (C), bottom (B), top (T) etc. In weak interactions, only transitions of type $\Delta \mathrm{N} = \pm 1$, where N is the quark flavour number, is allowed. Ultimately, the total quark number (baryon number) is conserved. One feature of hadronic weak interactions is the uniqueness to which the weak gauge bosons interact with each quark [63]. It is understood that the flavour eigenstates, not the mass eigenstates, interact with the weak interaction. This is known as the Cabibbo hypothesis, and the relationship between mass and flavour eigenstates is represented by the Cabibbo-Kobayashi-Maskawa (CKM) matrix [63], which is analagous to the lepton PMNS matrix.

1.2.3 Decay modes of the Muon

The weak force is known primarily for its role in nuclear β decay. However, it also is the force that mediates the decay of the muon. The muon has three main decay modes: the Michel decay, radiative decay of the muon, and what is known as 'internal conversion'. These decay modes are summarised in Table 1.1, where their branching fraction is also listed. For the radiative decay, the branching fraction is shown for $E_\gamma > 10$ MeV. Examples of lowest order Feynman diagrams are shown in Figure 1.3.

Decay mode		Branching Fraction
Michel decay	$\mu \to e\overline{\nu_e}\nu_\mu$	$\approx 100\%$
Radiative decay ($E_\gamma > 10$ MeV)	$\mu \to e\overline{\nu_e}\nu_\mu\gamma$	$(6.0 \pm 0.5) \times 10^{-8}$
Radiative decay with internal conversion	$\mu \to e\overline{\nu_e}\nu_\mu\ e^+e^-$	$(3.4 \pm 0.4) \times 10^{-5}$

Table 1.1: Branching fractions of muon decay modes. Charge of leptons is omitted as CP invariance is assumed. Obtained from Ref. [77].

The observation of neutrino oscillations by Super-Kamiokande [46] and SNO [47] allows for the calculation of the SM branching fraction of the muon decay $\mu \to eee$. These decays are suppressed by the standard model as they are mediated through oscillating neutrinos. The decay is suppressed by the factor $\frac{\Delta m_\nu^4}{m_W^4}$ [78], where m_W is the mass of the W boson, and Δm_ν^4 is the mass squared difference between muon and electron neutrinos, leading to a branching ratio $\mathcal{O}(10^{-54})$, rendering the SM decay completely unobservable.

1.3 Beyond the Standard Model

The SM is highly successful and capable of both explaining many of the observable phenomena in particle physics, and making incredibly precise predictions that are corroborated successfully by experiment. The success of discovering the Higgs boson [79], and the remarkable accuracy to which one can calculate properties using QED [80] are just some examples. The SM however has not incorporated gravity, dark matter or dark energy. The SM also lacks an explanation for baryogenesis [63] - no available mechanism accounts for the amount of matter-antimatter asymmetry in the universe, even with CP violation in the weak quark sector. The SM has also some aesthetic issues that they are believed by many to indicate physics beyond the Standard Model (BSM), such as the hierarchy problem, which remains unsolved, as well as the strong CP problem, which could be solved with the introduction of a new field, resulting in the axion [81, 82], which are dark matter candidates, through the Peccei-Quinn theory [83]. Few tensions and discrepancies are found, in a few cases, between theoretical and measured values of the model. More specifically to the muon, the existence of a 5.0 standard deviation discrepancy for the anomalous magnetic moment of the muon [44]. The development of BSM theories is thus vital in order to address the shortcomings.

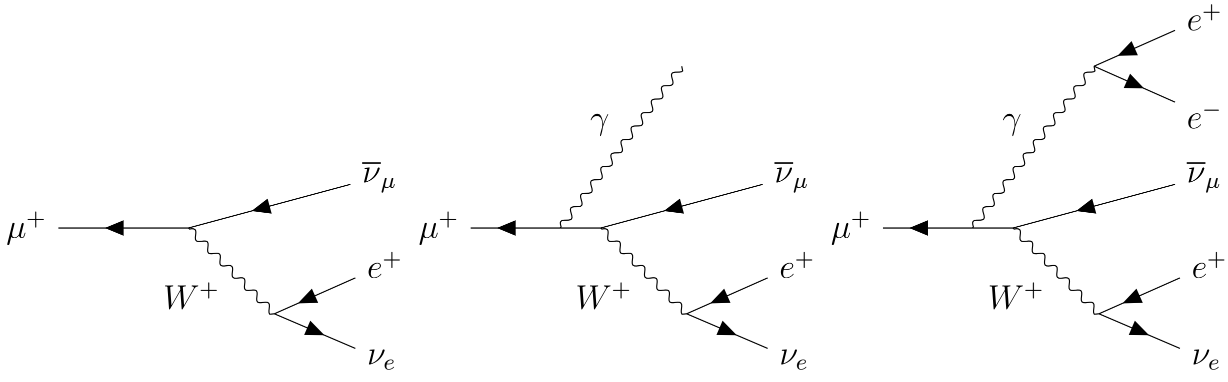

(a) Michel decay (b) Radiative decay (c) Internal conversion

Figure 1.3: Example leading order Feynman diagrams of the decay of the muon.

The history of particle physics has shown that the muon has often been on the frontier of new discoveries - muon physics experiments were vital in the development of the weak interaction. The muon can still be used to advance our understanding of these fundamental theories. Various supersymmetric models [84] predict charged lepton flavour violation to occur at a much higher branching fraction than what is calculated using SM processes. Models that hypothesise the existence of axion-like particles or dark photons could be explored with high intensity muon beams, which are also viable dark matter candidates. The mass reach of muon precision experiments far surpasses the physics reach attainable by any existing or planned particle collider, when unit couplings are assumed.

1.3.1 Lepton Flavour Violation

In the SM with massless neutrinos, lepton flavour is a conserved quantity. However, the observation of neutrino oscillations has established that lepton flavour is not conserved for neutral leptons. This is extended to the charged lepton sector through neutrino mixing, but at unobservable rates. Various theories [84] predict the branching fraction of charged lepton flavour violating processes to occur at much higher branching fractions than predicted by the SM, through neutrino mixing. Historically, many experiments have explored charged lepton flavour violating processes as a means of testing the SM. Figure 1.1 showed the development of upper limits to the processes $\mu \rightarrow e\gamma$, $\mu\mathrm{N} \rightarrow e\mathrm{N}$, and $\mu \rightarrow eee$. The results shown there correspond to the various eras of muon experiments, such as cosmic ray searches through to the development of more advanced accelerators. The most stringent results are shown in Table 1.2. The Mu3e experiment aims to search for the charged lepton flavour violating decay $\mu \rightarrow eee$. Data taking will begin in 2025, and by the end of Phase-2, Mu3e aims to reach a sensitivity of 1 in 10^{16} muon decays.

In the neutrino sector, lepton flavour violation is explored through searches of neutrinoless double beta decay ($0\nu\beta\beta$) [90]. This is a decay process forbidden within the standard model, and its observation would not only be a violation of lepton number ($\Delta\mathrm{L} = 2$), but also would confirm that neutrinos are majorana particles, i.e, they are their own antiparticle.

Decay mode		Fraction $[\Gamma_i/\Gamma]$ (90% CL)	Ref.
Muon to electron-gamma	$\mu \to e\gamma$	$< 4.2 \times 10^{-13}$	[85]
Muon-Electron conversion	$\mu\mathrm{N} \to e\mathrm{N}$		
	$\mu\mathrm{Ti} \to e\mathrm{Ti}$	$< 6.1 \times 10^{-13}$	[86]
	$\mu\mathrm{Pb} \to e\mathrm{Pb}$	$< 4.6 \times 10^{-11}$	[87]
	$\mu\mathrm{Au} \to e\mathrm{Au}$	$< 7.0 \times 10^{-13}$	[88]
Muon to three electrons	$\mu \to eee$	$< 1 \times 10^{-12}$	[89]

Table 1.2: Upper limit to branching fractions of CLFV muon decay modes. Charge of leptons is omitted as CP invariance is assumed.

This process is the transition from a parent nucleus to a daughter nucleus, with the emission of two electrons but no neutrinos. The most stringent limit is currently set at $> 2.3 \times 10^{26}$ [years] by KamLAND-Zen [91].

The search for lepton flavour violation can also be conducted with tau leptons. It is however more challenging, as tau leptons have a much shorter lifetime than the muon, and are produced in relatively small amounts within collider experiments [92]. However, due to the large mass of the tau, the predicted branching fractions for BSM lepton flavour violating processes are much higher than its' muon counterpart. The most stringent limits for lepton flavour violation in tau leptons has been set by BaBar [93] and Belle [94] collaborations, and are $\mathcal{O}(10^{-8})$. For $\tau \to e\gamma$, a lower limit of 3.3×10^{-8} is set at 90% CL [95] and for $\tau \to \mu\gamma$, 4.2×10^{-8} [96] is set. For lepton flavour violating tau decays to three leptons, these limits are dominated by the Belle experiment. For example, lower limits to 90% CL are set for processes $\tau \to \mu\mu\mu$ and $\tau \to eee$ respectively at 2.1×10^{-8} and 2.7×10^{-8} [97].

Lepton flavour violation can also be explored at B meson factories. A 2.6σ [98] tension between data and the standard model prediction for lepton universality in B decays could be explained through BSM processes that allow lepton flavour violation in B meson decays [92]. Lepton flavour violation with B mesons is thus explored at LHCb, Belle and BaBar. The best limits for $B^0 \to e^{\pm}\mu^{\mp}$ and $B^0 \to \tau^{\pm}\mu^{\mp}$ is set by LHCb at 1.3×10^{-9} and 1.4×10^{-5} at 90% and 95% CL respectively [99, 100].

1.3.2 Axion-Like Particles

The large dataset that will be acquired by experiments akin to Mu3e will allow the possibility of exploring exotic LFV decays of the muon. One such decay is the decay of a muon to a light, long-lived, pseudoscalar particle that will ultimately decay within the detector. For Mu3e, this was first explored using a simple analysis [101], where axion-like particles (ALPs) a, are coupled to charged leptons, and is parameterised by the Lagrangian $\mathcal{L}_{\mathrm{a}}$,

$$\mathcal{L}_a \supset \frac{1}{\Lambda}\partial_\mu a \bar{\ell}_a \gamma^\mu g_{\alpha\beta}\gamma^5 \ell_\beta, \tag{1.1}$$

where the effective energy scale is represented by Λ, charged leptons are represented by ℓ_α (ℓ_β), where α (β) $= e$, μ and τ, and the hermitian and anti-hermitian current structure matrix is represented by g. For convenience, the energy scales can be defined as $\Lambda_{\alpha\beta} \equiv \Lambda/g_{\alpha\beta}$. A diagram for the decay process is shown in Figure 1.4. The decay process results ultimately in a three-electron final state, and so the decay can be studied with the Mu3e Experiment's detector apparatus. The decay process is studied in Chapter 6. The study is conducted using the Mu3e software framework, using the standard tracking and vertex algorithms.

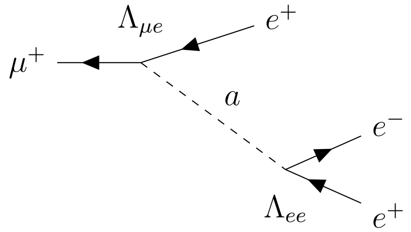

Figure 1.4: Diagram of a muon decaying to an ALP and a positron, with the ALP (either promptly or not) decaying into a dielectron pair. The lepton flavour violating conserving and violating couplings are also shown in the diagram.

From the Lagrangian presented in Eq. 1.1, the decay width of the muon to the LFV two body decay is:

$$\begin{aligned}\Gamma_{\mu\to ea} = \frac{m_\mu^3}{16\pi\Lambda^2}&\sqrt{(1-r_a^2)^2+r_e^4-2r_e^2(1+r_a^2)}\\ &\times[(g_{\mu a})^2(1-r_a^2)^2-(g_{\mu a})^2((1+r_e)^2)r_a^2]\end{aligned} \tag{1.2}$$

where $r_{e,a} \equiv m_{e,a}/m_\mu$, which, for $m_{e,a} \ll m_\mu$ is $m_\mu/(16\pi\Lambda_{\mu e}^2)$, (with $g_{\mu e}/\Lambda = \Lambda_{\mu e}$) and m_e is the electron mass, m_a is the mass of the ALP. The decay width of the ALP to a positron-electron pair is:

$$\Gamma_{a\to e^+e^-} = \frac{m_a}{2\pi}(g_{ee})^2\frac{m_a^2}{\Lambda^2}\sqrt{1-\frac{4m_e^2}{m_a^2}}. \tag{1.3}$$

The boosted decay length is an important parameter for the process, and is given by:

$$\gamma c\tau = \frac{c|P_a|}{m_a\Gamma_a}, \tag{1.4}$$

where the momentum of the ALP (P_a) is given by:

$$P_a = \frac{\sqrt{\left(m_\mu^2-(m_e+m_a)^2\right)\left(m_\mu^2-(m_e-m_a)^2\right)}}{2m_\mu}. \tag{1.5}$$

The lepton flavour conserving scale g_{ee}/Λ is proportional to the decay width and therefore lifetime of the ALP, and is given by,

$$\Lambda_{ee}^2 \simeq \frac{(\gamma c\tau)m_a^2 m_\mu m_e^2}{c\pi(m_\mu^2-m_a^2)}. \tag{1.6}$$

Figure 1.5 is an adaptation of the current calculated limits for Λ_{ee} for ALP masses m_a. Lines of constant decay width are drawn on the plot as a function of m_a, and are given by Eq. 1.6. The lines chosen are of geometrical significance for the experiment. The 'prompt' line is a constant decay width of 0.01 cm, the 0.3 cm constant decay width is chosen as a cut of 3

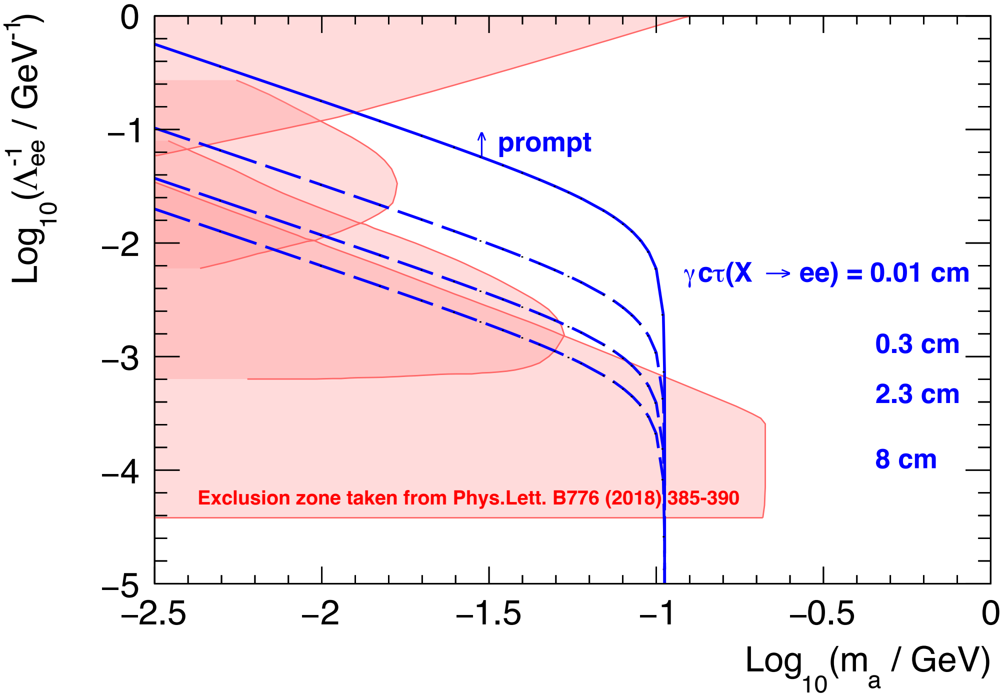

Figure 1.5: Phase space plot adapted from Ref. [101]. Lepton flavour conserving coupling plotted against mass of the ALP 'm_a'. The shaded regions are spaces covered by existing limits due to beam dump experiments, g-2, and cosmological events (SN1987A). The blue lines are lines of constant decay length.

mm could be used in a future $\mu \to eee$ analysis with Mu3e. The first silicon layer is located at 23.3 mm, and the fourth layer is beyond 80 mm, therefore the lines of constant width of 2.3 and 8 cm are chosen. All are found to be within the region of unexplored phase space. The Mu3e Experiment's sensitivity to ALPs are presented in Chapter 6.

1.3.3 Dark Photons

Dark matter is inferred to exist through a plethora of astronomical observations. Gravitational lensing of galaxies can only be explained through the addition of additional 'hidden masses' [102, 103]. Stellar rotation curves require additional 'invisible' mass to explain their observations [103]. This 'dark matter' differs from ordinary matter as it appears to only interact primarily gravitationally, in contrast to all other matter observed so far. The SM does not provide a viable DM candidate, and many BSM models attempt to answer this shortcoming. Several investigations into the dark sector are conducted through feeble channels known as portals, which connect SM particles to dark matter. A vector portal involves a boson mediator with a spin of 1, and results in the 'kinetic mixing' of two Abelian bosons (dark and visible) [104]. This small form of mixing with regular matter, could present itself alongside SM interactions [105], which could be detected by experiment. The resulting Dark Photon is a viable dark matter candidate [105].

To produce dark photons, in addition to $U(1)$ electromagnetism, another $U(1)$ gauge symmetry is added. This corresponds, in total, to two gauge fields A_1^μ and A_2^μ [106]. The Lagrangian $\mathcal{L}$ is the sum of $\mathcal{L}_0$ and $\mathcal{L}_1$. The possible kinetic terms, $\mathcal{L}_0$, is given by Eq. 1.7,

$$\mathcal{L}_0 \supset \frac{\epsilon}{2} F'_{\mu\nu} F^{\mu\nu} - \frac{1}{4} F_{\mu\nu} F^{\mu\nu} - \frac{1}{4} F'_{\mu\nu} F'^{\mu\nu}, \tag{1.7}$$

where $F^{\mu\nu}$ is the field strength tensor, and Eq. 1.8 below represents the interaction term,

$$\mathcal{L}_1 \supset e J_\mu A_2^\mu + e J'_\mu A_1^\mu. \tag{1.8}$$

where e is the electric charge, and J is the electric current. The kinetic part of the Lagrangian can be diagonalised, and simplified into its canonical form. For a massive dark photon, the interaction term of the Lagrangian is given by Eq. 1.9,

$$\mathcal{L}_{A'} \supset e\epsilon J_\mu A_2^\mu, \tag{1.9}$$

where ϵ is the kinetic mixing parameter, and is representative of the strength of the interaction between the dark photon and the electromagnetic current. The mass term of the dark photon is given by:

$$\mathcal{L}_{\text{mass}} \supset -\frac{1}{2} m_{A'}^2 A_{A'\mu} A_{A'}^\mu. \tag{1.10}$$

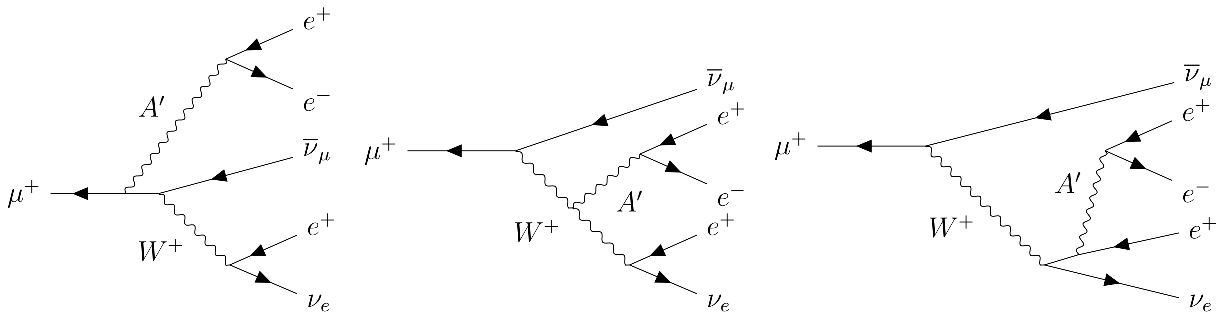

Figure 1.6: Leading order Feynman diagrams of the dark photon decay process.

The Feynman diagrams for the leading order processes are shown in Figure 1.6. Oscillations of a massive dark matter field, at low temperatures, behave like cold dark matter and constitute a viable dark matter candidate [107, 108]. In Chapter 3, The Mu3e experiment's sensitivity to dark photons via a complementary search to the main analysis of $\mu \rightarrow eee$ will be shown, where between the ranges of 10 - 70 MeV, Mu3e will have a competitive sensitivity after the first few years of running of the experiment.

1.4 Muon Experiments Searching for Charged Lepton Flavour Violation

The observation of charged lepton flavour violation would be an indisputable sign of physics beyond the SM. Several experiments, from the 20^{th} century onwards, have attempted to discover charged lepton violating processes to no avail. With the development of higher intensity muon beams, and better detector technologies, the existing limits on the processes can be further improved. Several experiments, such as SINDRUM, MEG and others are described below. The SINDRUM experiment set the most stringent limit on the decay $\mu \rightarrow eee$, which is characterised by three electrons emerging from a vertex with invariant mass equal to the muon mass. The MEG Experiment has set the most stringent limit to the muon decay $\mu \rightarrow e\gamma$, characterised by monochromatic positron and photon. For muon capture experiments like Mu2e and COMET, muons replace electrons in the atomic shells of a target material, forming muonic atoms. The decaying muon will then produce a monochromatic electron which is detected. The fundamentals of muon experiments are similar in each case: a large number of muons is generated from a pion beam, and the beam (pulsed/continuous) is directed upon a target. A beam of positive muons is more likely to be used as it is easier to reach a high beam current.

1.4.1 SINDRUM

The SINDRUM experiment was based at the Swiss Institute for Nuclear Research[2] (SIN). SINDRUM used a beamline capable of producing a continuous muon beam with current 7×10^6 μ^+/s. The experiment searched for the lepton flavour violating decay $\mu \to eee$, and has set the limit to the decay, at $1 < 10^{-12}$ at 90% CL in 1985 [109, 89, 110], which is currently the strongest limit.

A schematic of the experimental apparatus is presented in Figure 1.7. The solenoid magnet is represented by S, and focuses the muon beam onto the hollow double-cone target T, upon the surface of which the muons will decay. The magnet, M, produces a homogeneous solenoidal magnetic field up to 0.6 T. The particles emitted from the muon decay are measured by a set of 5 concentric multiwire proportional chambers, represented by C. The cylindrical hodoscope, H, contains 64 scintillators, and is responsible for finding track timing information, and is complemented by photomultipliers, P. A total of 7.3×10^{12} muons were stopped upon the target [109]. No $\mu \to eee$ signal event was observed, producing an upper limit of 1×10^{-12} to 90% CL.

1.4.2 MEG Experiment

The Mu-to-Electron-Gamma (MEG) experiment is based at the Paul-Scherrer institut(PSI) in Switzerland. The aim of the MEG experiment is the search for the lepton flavour violating decay $\mu \to e\gamma$. This is made possible through the intense muon beam at PSI, which is capable of producing a muon current of more than $10^8\mu$/s. A schematic of the experimental

---

[2] Now known as the Paul Scherrer Institut (PSI)

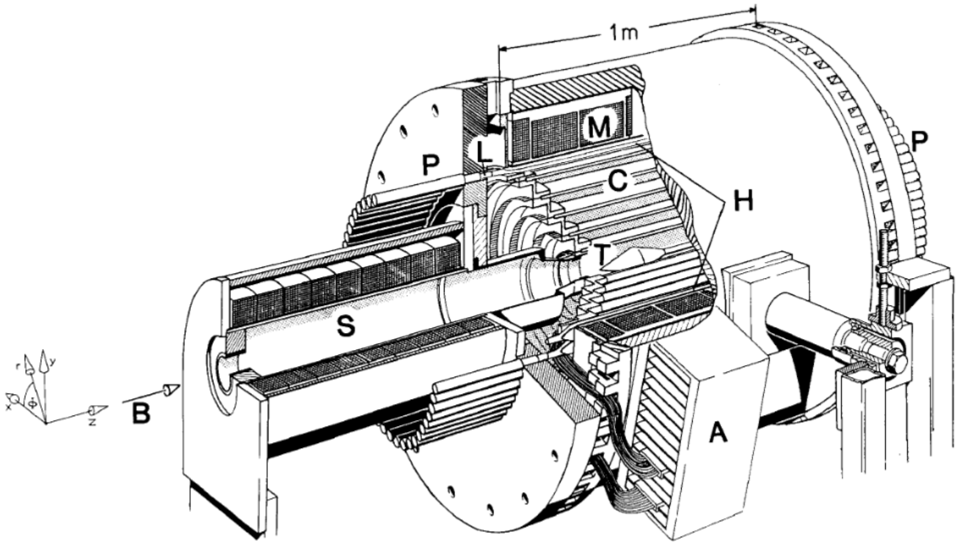

Figure 1.7: Schematic of the experimental apparatus of the SINDRUM detector. Taken from Ref. [110]

apparatus is shown in Figure 1.8. The beam of muons from the muon beam line is focused onto the thin stopping target, where the muons will stop and decay. The target is located within the centre of the COBRA magnet, which surrounds the detector, and is responsible for providing a gradient magnetic field, which only allows high momentum positrons to reach the drift chambers. The track of the positron is measured by the drift chambers, and timing information is provided by the timing counter, which consists of scintillator tubes. To measure the photon, an innovative design is used which consists of using liquid xenon as a scintillator material, the light of which is measured by surrounding photomultipliers. The decay process is of a unique signature, and is defined by the production of both a monochromatic positron and photon, each with energy equal to exactly half of the muon mass (52.83 MeV), that are reconstructed to be emitted from the same point and moving in opposite directions. The experiment has found no evidence of $\mu \to e\gamma$ decays, and has set the most stringent upper limit to the branching ratio at 4.2×10^{-13} in 2016.

The Phase-2 upgrade of the MEG experiment aims to push the upper limit to the charged lepton flavour violating decay $\mu \to e\gamma$ further to 6×10^{-14}. This is made possible through the planned experiment upgrades. The schematic is presented in Figure 1.9. Upgrades to the drift chamber, to a single volume drift chamber, with higher density drift cells, allows the

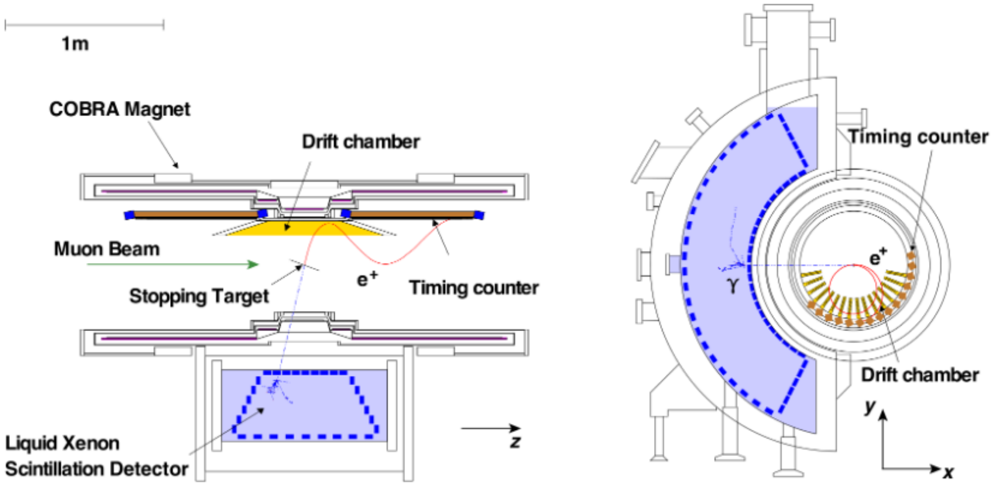

Figure 1.8: Schematic of the detector apparatus of the MEG experiment. Taken from Ref. [54]

detector to cope with higher muon rates whilst also improving the momentum resolution of positron tracks. Timing information is improved through the pixelated timing counter, which utilises scintillator tiles to measure timing information of the tracks [111]. This improved design increases acceptance of the detector by a factor of 2 [111]. Data taking is ongoing [112] as of the current year (2023).

1.4.3 Muon-Electron Conversion Experiments

Looking towards the near future, several experiments are also searching for muon to electron conversion within the radius of a nucleus in a target material. One example of such an experiment is the COherent-Muon-to-Electron-Transition (COMET) experiment, which is based at the Japan Proton Accelerator Research Complex (J-PARC), and aims to search for the lepton flavour violating muon-to-electron conversion process $\mu \mathrm{N} \to e \mathrm{N}$, utilising an aluminium nucleus [114]. The experiment aims for a Phase-1 sensitivity of 3.1×10^{-15}, and a Phase-2 sensitivity of 2.6×10^{-17}. This will be achieved using the detector apparatus outlined in Figure 1.10. Muons are generated by aiming a proton beam at a graphite target, and the muons are then stopped on an aluminium target, where they form a muonic aluminium atom before decaying. The emitted electrons are measured by the scintillators in the hodoscope

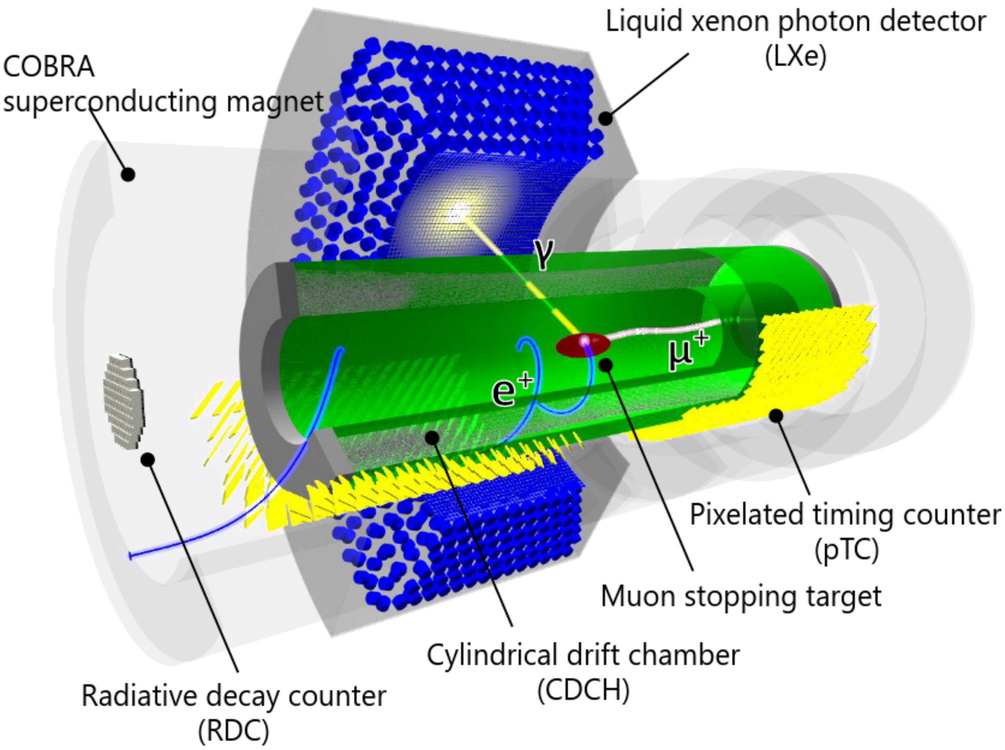

Figure 1.9: Schematic of the detector apparatus of the MEG II experiment. Taken from Ref. [113]

and the drift chamber. The Phase-2 apparatus will feature an extra C shape to the detector design. Improvements to the detector will allow for an increase in the achievable upper limit by two orders of magnitude [114]. Phase-1 of the experiment is currently undergoing commissioning runs, which will shortly be followed by Phase-2 running of the experiment.

The Muon-to-Electron (Mu2e) experiment is based at Fermilab, and follows a similar principle to COMET. A schematic of the detector apparatus is shown in Figure 1.11. A muon beam is generated by a pulsed proton beam colliding with a tungsten target and is carried by the transport solenoid towards an aluminium target. Muons will then form muonic aluminium atoms before decaying. The signal of the search will be an electron of energy equal to the mass of the muon, minus the binding energy and energy of the nuclear recoil. This electron will be detected by the tracker and calorimeter, which will accurately measure the energy of the produced electron [116]. A single-event sensitivity of 3×10^{-17} is projected to be reached [116]. As of 2023, commissioning is ongoing, and data taking is expected to begin in the mid 2020's.

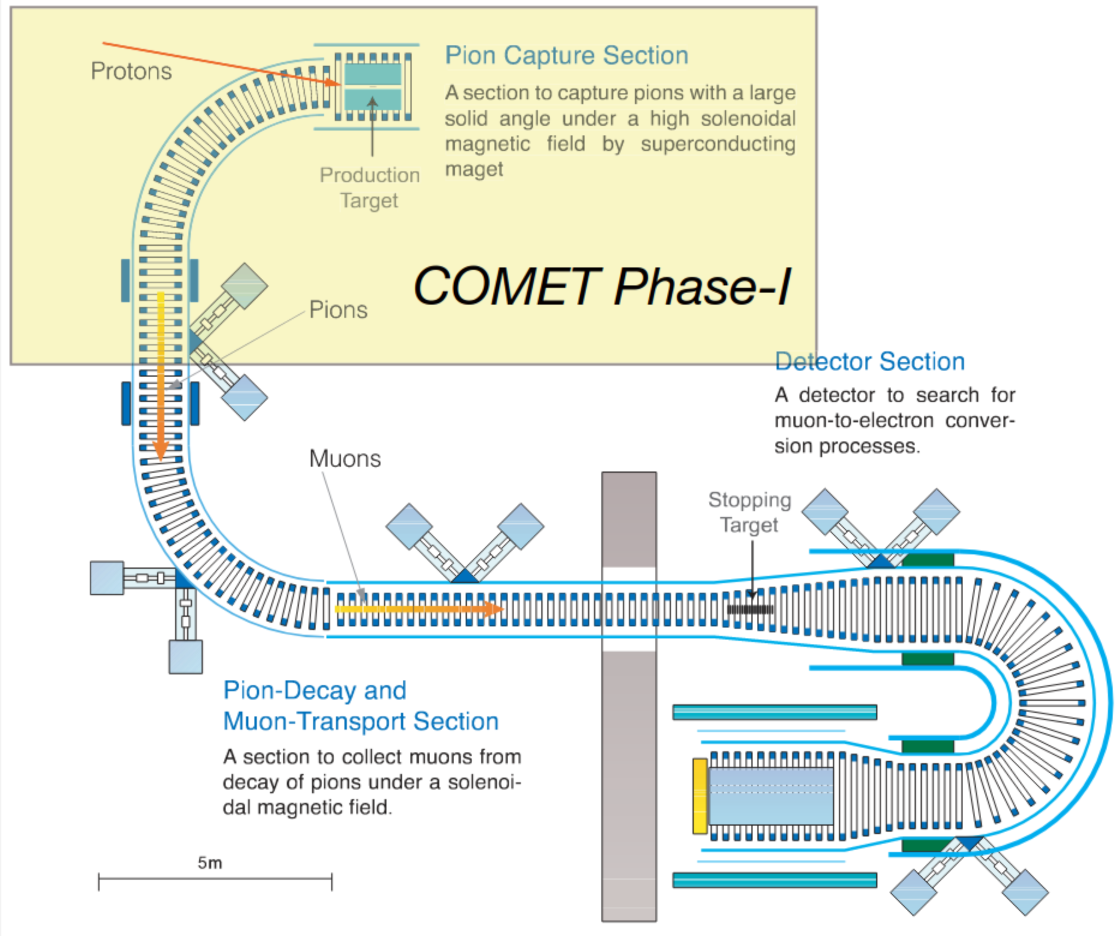

Figure 1.10: Schematic of the detector apparatus of both phases of the COMET experiment. Figure taken from Ref. [115].

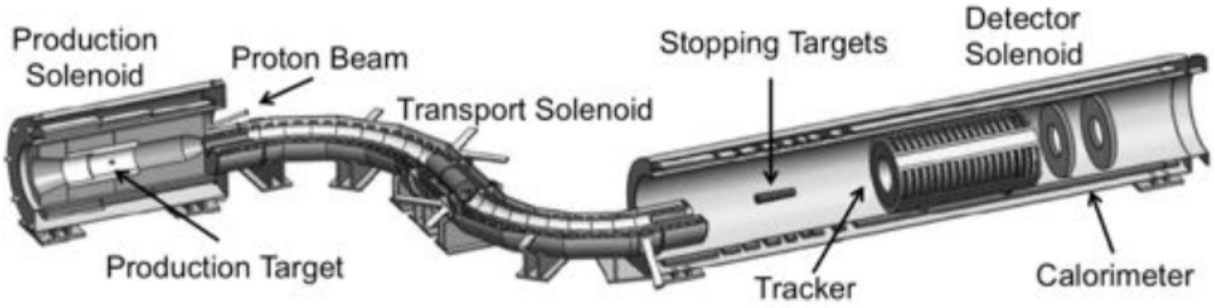

Figure 1.11: Schematic of the Mu2e design. Figure taken from Ref. [116].

1.4.4 Summary of Future Projections

Many experiments described here are currently either under construction, or are currently running. In Figure 1.12, the current era of muon experiments and their expected limits are summarised. The 2020's will be an exciting time for muon physics and could rule out several BSM theories.

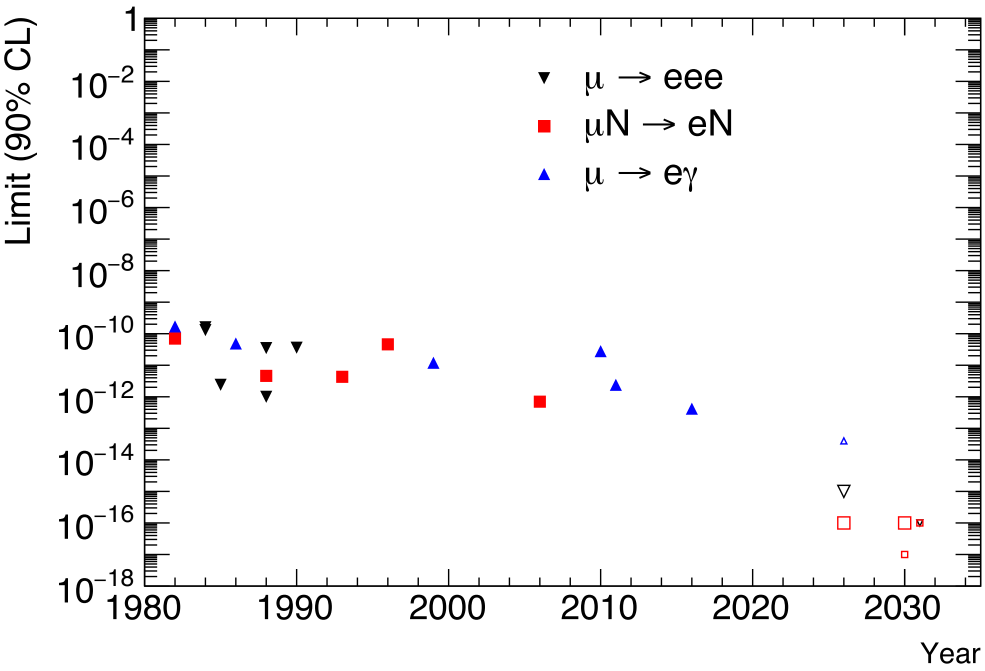

Figure 1.12: Current era of charged lepton flavour violating experiments included into the historical plot of muon experiments. Non-filled markers indicate experiments with projected results. MEG Phase-2 projection taken from [117]. Mu3e is the unfilled black triangle in the Figure.

1.5 Conclusion

This chapter provided a brief overview of muon physics, with a particular emphasis on lepton flavour violation in the intensity frontier. It began with its surprising discovery in a study of cosmic rays, leading to its usefulness in understanding nuclear β decay and the theory of electroweak interactions. It was shown that the muon's properties have found multiple uses outside of the realm of particle physics, in applications in solid state physics, through the μSR technique. Ultimately, the SM was introduced, implying the conservation of charged lepton flavour. The SM's inability to provide a candidate for dark matter was explained, and it was shown how the muon could be used as a tool for exploring dark photon and axion-like particle phase space, both of which would be a viable dark matter candidate. The charged lepton sector is an exciting sector to search for new physics, as many theories beyond the SM propose mechanisms for lepton flavour violation to occur at branching fractions reachable by experiment. An exciting summary of experiments searching for these golden decay channels was given, showing many experiments will begin to run (or have already begun) starting in the mid 2020's, pushing the most stringent limits to these channels down by several orders of magnitude.

Chapter 2

The Mu3e Experiment: Detector and Software

The Mu3e experiment aims to search for the lepton flavour violating muon decay $\mu \to eee$ to a sensitivity of 1 in 10^{16} muon decays with zero background. To achieve this, strict requirements on the momentum, vertex and time resolutions are necessary. This will be achieved through the use of several modern advancements in silicon and scintillator detector technology, as well as the eventual use of the High-Intensity Muon Beam (HIMB) upgrade at PSI.

This chapter describes the main challenges facing the Mu3e experiment. The solutions of which are provided through the design of the Mu3e detector. A description of the track and vertex reconstruction algorithms that are used will follow, with an introduction to the Mu3e software framework. A discussion of possible future upgrade scenarios concludes the chapter.

2.1 Experimental Concept

The Mu3e detector is a highly granular cylindrical detector specialising in the reconstruction of low momentum (10 - 55 MeV) electrons that originate from stopped muons delivered through the πE5 beamline at PSI. A longitudinal schematic of the baseline design is shown in Figure 2.1. In its entirety, the detector is placed within a solenoidal magnet that generates a homogeneous magnetic field of 1 T. The beam of positive muons, delivered by the πE5 beamline, is stopped by a thin mylar hollow double-cone target, where the muons will decay at rest. The central detector is comprised of a thin silicon double layer pixel tracker consisting of sensors made in High-Voltage Monolithic Active Pixel (HV-MAPS) technology. Timing information is provided by the two timing detector systems, the fibre and tile detectors which are responsible for reducing combinatorial background and track ambiguities. The momentum measurement relies primarily on the outer layer pixel layers. The central detector is extended by upstream and downstream recurl stations which comprise of two thin silicon pixel tracker layers and a tile detector.

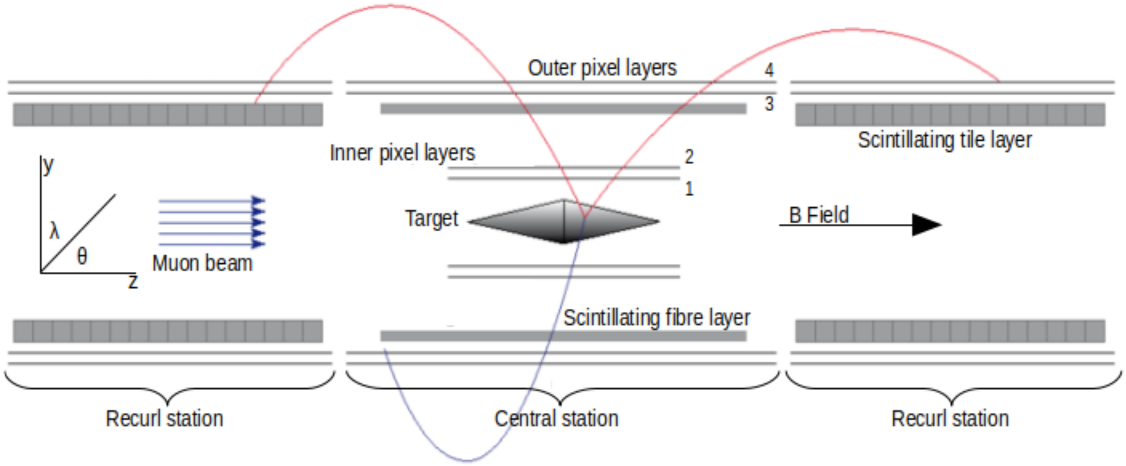

Figure 2.1: Schematic of the Phase-1 design of the Mu3e Detector. Modified from Ref. [118]. 'Upstream' is defined as beyond the Mu3e target, increasing in z, whilst 'Downstream' is defined as being in the opposite direction.

2.1.1 Coordinate System

The righthanded coordinate system used in Mu3e is shown below in Figure 2.2. The target is drawn as reference. The positive z axis is in the direction of the beam. The y-axis is chosen to point upwards. The inclination angle is represented by λ, and the azimuthal is represented by φ.

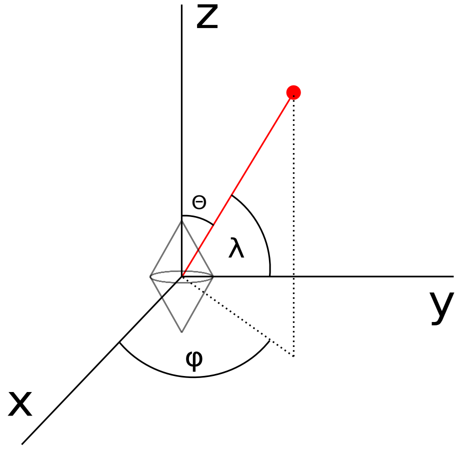

Figure 2.2: The Mu3e co-ordinate system is shown, with the target as reference. The direction of positive z-axis is the direction of the beam, and is defined as downstream, the opposite direction is upstream. The azimuthal is represented by φ, and the inclination angle is represented by λ.

2.1.2 Multiple Coulomb Scattering

Elastic Coulomb scattering alters the path of a particle as it traverses through material and is deflected by nuclei. Due to these deflections the particle exits the material with a different direction to how it entered. Assuming the layers traversed are thin, these deflections are described by a Gaussian distribution with a mean of zero, and a width given by [119]:

$$\sigma_{MS} = \frac{13.6\ \mathrm{MeV}}{p\beta c} q \sqrt{\frac{x}{X_0}} \left[1 + 0.038 \ln \frac{x}{X_0}\right], \tag{2.1}$$

where p represents the momentum, βc represents the velocity of the particle, and the charge (in units of electron charge) is represented by q. Radiation length is given by X_0, and the total distance travelled by the particle within the material is given by x. Therefore, it can be seen that the width of the formula increases for low momentum particles, and for material with high radiation length. Increasing the distance travelled within the material, x, also induces a larger width σ_{MS}. The more scattering occurs, will limit the momentum resolution. A low material budget is thus of paramount importance in Mu3e.

2.1.3 Momentum Measurement

The Mu3e pixel tracker takes advantage of the highly granular pixel sensors. Multiple scattering effects dominate over the position resolution of the pixel tracker. The precision of the measurement therefore depends critically on the multiple scattering angle Θ_{MS}, and the track curvature Ω in the magnetic field B. To first order it holds that:

$$\frac{\sigma_p}{p} \propto \frac{\Theta_{MS}}{\Omega}. \tag{2.2}$$

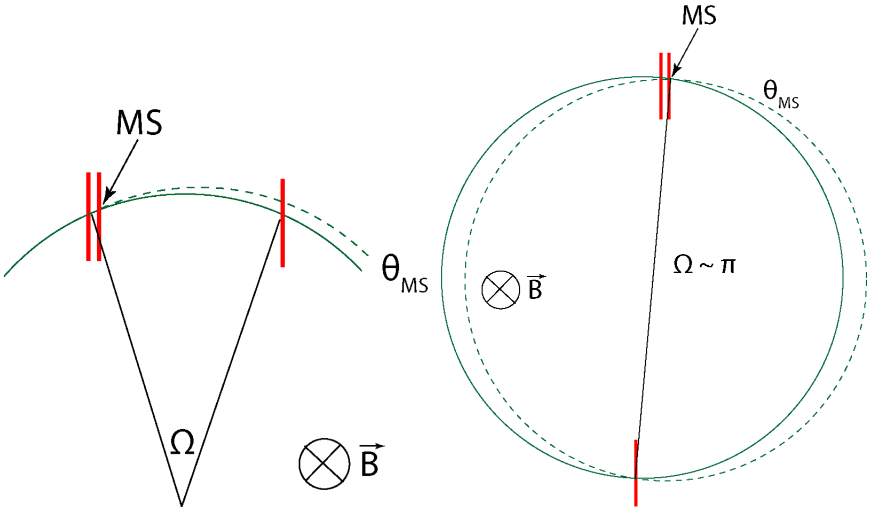

Figure 2.3: Figure showing the effect of increasing curvature has on the momentum resolution, taken from Ref. [118].

Where σ_p is the momentum resolution and p is the momentum. Figure 2.3 describes pictorially the effect of strategically placing the detectors within a multiple-scattering dominated experiment. At half a turn ($\Omega \approx \pi$), there is an overlap between the trajectory of a particle with and without multiple scattering. The scattering uncertainty vanishes at these points. This fact is what drives the design of the Mu3e experiment. 'Recurling' tracks that travel back into the detector are desired, as these tracks will obtain a superior momentum resolution from measurements in the pixel tracker.

2.2 Signal and Background

Signal

The aim of The Mu3e Experiment is to search for an anti-muon decaying into two positrons and an electron ($\mu^+ \rightarrow e^+e^+e^-$ [1]). With the assumption that the muon decays at rest, the decay has many unique and interesting properties. All of the decay products emerge from a single point in space (the vertex) and in time. Using the assumption that the muon decays at rest, the sum of the vector momenta of the three electron system vanishes:

$$|\vec{P}_\mu| = |\vec{P}_{eee}| = |\sum_e \vec{P}| = 0, \tag{2.3}$$

and the invariant mass of the decay electrons are equal to the muon mass:

$$\mathrm{E}_{tot} = \sum_e \mathrm{E} = m_\mu. \tag{2.4}$$

These kinematic properties make the decay unique, and their exploitation allows for signal to be distinguished from the various types of background. Mu3e, thanks to the large number of expected muon decays, can also be used to explore other exciting beyond the standard model physics phenomena including, the lepton flavour violating decay of a muon to an axion-like particle, or even the radiative decay of the muon, with the radiation of a dark

---

[1] For simplicity, the anti-muon is referred to as a muon, and both positrons and electrons are referred to as electrons.

matter candidate: the dark photon.

Background

As previously discussed in Chapter 1, the muon is observed experimentally to decay primarily through the Michel decay (BR: ∼100%) process. The decays of the muon then increase in rarity, as shown in Chapter 1. These decays, or overlaps of multiple decays, if the final decay topology has three outgoing electrons, can appear to be 'signal-like' when measured by the Mu3e detector. The signal along with the dominant background topologies are shown in Figure 2.4. The two dominant background sources are discussed in the following.

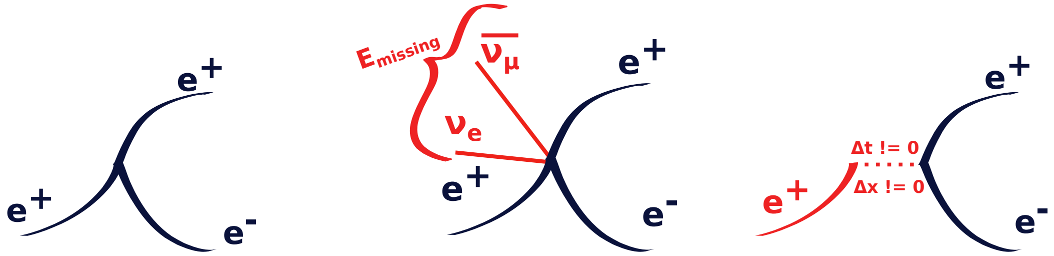

(a) (b) (c)

Figure 2.4: Sketches of the various signal and signal-like topologies are shown.

Irreducible Background

An irreducible background originates from the reconstruction of an internal conversion ($\mu^+ \to e^+\overline{\nu_e}\nu_\mu\, e^+e^-$), shown in Figure 2.4b. This decay has a branching fraction of $(3.5 \pm 0.4) \times 10^{-5}$ [77]. The electrons from this decay form a vertex. The only difference between this and a vertex formed from $\mu \to eee$ is the energy carried away by the two neutrinos. Assuming the total momentum of the event is zero, and the invariant mass is equal to the muon mass, the energy difference is the only way one can distinguish this decay from $\mu \to eee$ signal through interpreting this difference as inference of a neutrino pair.

Excellent momentum resolution is thus of paramount importance when dealing with the internal conversion background, as this is directly related to the reconstructed mass resolu-

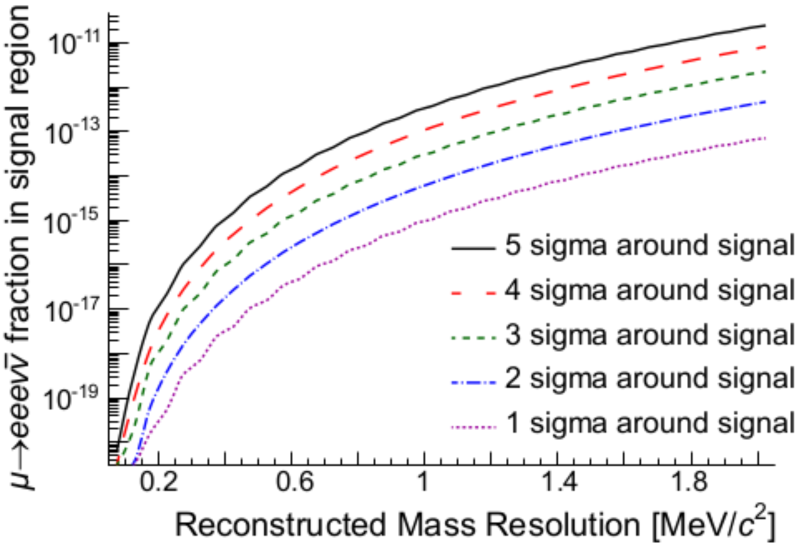

Figure 2.5: Amount of internal conversion in signal region (defined as the invariant mass region of 103 - 115 MeV) as a function of the momentum resolution. Plot taken from Ref. [118].

tion. In Figure 2.5, the contamination of the signal region is shown as a function of the mass resolution. A mass resolution better than 1 MeV is required for the experiment to obtain a sensitivity of 2×10^{-15} with a 2 σ cut on the reconstructed invariant mass of the vertex.

Combinatorial Background

Combinatorial background is the result of the accidental combination of decay products from separate muon decays and interactions with material. Decays such as the Michel decay, radiative decay with internal conversion, or interactions such as Bhabha scattering of a positron can all contribute. In this case some or all decay electrons originate from separate vertices, but are accidentally reconstructed as coming from a single vertex, with properties similar to that of a $\mu \to eee$ decay.

The πE5 beamline emits a beam consisting of positive muons, and so one of the most important constituents to combinatorial background is the Michel decay of the muon $\mu^+ \to e^+\nu\nu$. This decay will only produce positrons, so to create a candidate signal event, a Michel positron must have its charge misidentified or Michel decays must be combined with an electron producing process. Charge misidentification of a positron could cause a track to be accidentally mis-reconstructed as an electron, which can occur for tracks that recurl back into the target (which are an example of ‘recurling’ tracks).

The radiative decay of the muon $\mu^+ \to e^+\gamma\nu\nu$ also contributes to the combinatorial background. The photon in the radiative decay can convert to a positron-electron pair when passing through material, such as the Mu3e target or the silicon layers. A positron-electron pair produced from the conversion can be combined with a positron from another muon decay, producing a topology similar to that of a signal event

Positrons passing through material can also undergo Bhabha scattering, which produces a positron-electron pair, originating from a single point. These, combined with a reconstructed positron from another decay process, can also create a vertex. This process is one of the main sources of combinatorial background.

The reduction of combinatorial background involves understanding whether particles originate from a single point in space and time. The sources of different decay processes originate from different positions in space and time. In order to reduce the background with spatial constraints, extremely precise vertex and timing resolution is required. Vertex resolution is minimised by decreasing the amount of multiple scattering, and so detector material. Accurate timing information will be provided by the scintillating fibre and tile detectors, (<250 ps and <100 ps resolution respectively). The mass distribution of combinatorial background extends beyond the muon mass, which poses a challenge to the experiment.

2.3 Baseline Design (Phase-1)

As shown previously in Figure 2.1, there are two types of stations in the Phase-1 design of the Mu3e experiment: the 'central' and 'recurl' stations. The Mu3e target sits surrounded by the central station. Around which are four silicon layers with an overlapping, cylindrical ladder design located at increasing radii (at 23.3 mm, 29.8 mm, 73.9 mm and 86.3 mm), leading to a very high detector acceptance, which completely surrounds the target allowing for complete track reconstruction for 38.1% of simulated $\mu \to eee$ muon decays [118] assuming a four-Fermion contact interaction. Between layers 2 and 3 in the central station, is the cylindrical fibre detector, which provides timing information for all tracks. To take advantage of the superior momentum resolution of tracks that recurl back toward the detector, two recurl stations are placed upstream and downstream of the central detector. The recurl stations each contain two of the silicon layers of highest radii, (layers 3 and 4 at radii of 73.9 mm and 86.3 mm), as well as housing the tile detector, which provides the most precise of timing information. Table 2.1 provides a detailed description of the various specifications for the Mu3e detector. Figure 2.1 is a longitudinal view of the Mu3e detector, showing each of the stations and the various components of both tracking and timing detectors. In the following sections, the various parts of the Mu3e detector are described in more detail.

Detector	Total active length [mm]	Radius [mm]	Resolution
Target	100	19	-
First Si layer	124.7	23.3	< 10 ns / < 30 μm
Second Si layer	124.7	29.8	< 10 ns / < 30 μm
Fibre layer	280	61	250 ps / 100 μm
Tile layer	342	64	60 ps
Third Si layer	351.9	73.9	< 10 ns / < 30 μm
Fourth Si layer	372.6	86.3	< 10 ns / < 30 μm

Table 2.1: Specifications of the various active parts of the Mu3e Detector. Information taken from Ref. [118], and are shown per station.

2.3.1 The PSI Facility and Muon Beam

To reach the Phase-1 sensitivity goal, a high intensity muon beam is required. It is also favourable for the beam to be continuous, as opposed to being pulsed, as this allows the muon decays to be spread evenly in time allowing combinatorial background to be more easily suppressed. The Phase-1 Mu3e sensitivity goals necessitates a beam current of over 1×10^8 muons/s. This is achieved using PSI's High Intensity Proton Accelerator [120], which accelerates protons from ionised hydrogen up to energies of 590 MeV, with a beam power of 1.4 MW. Mu3e uses specifically the πE5 channel, which is a secondary beamline which uses a graphite target to produce pions from the protons colliding with the target. The pions then decay to muons, which have an average momentum of 28 MeV. The CAD model of which is shown in Figure 2.6.

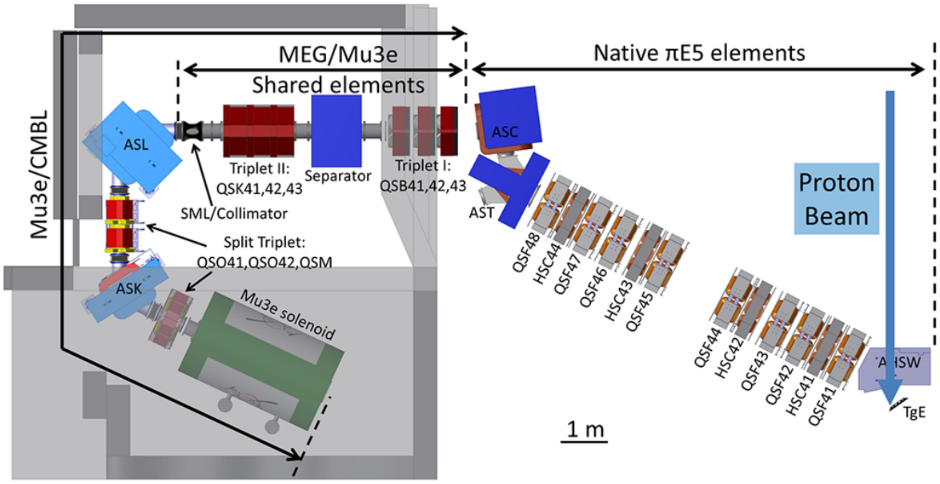

Figure 2.6: Figure shows a CAD model of both the πE5 beamline at PSI, and the entire CMBL. Plot taken from Ref. [118].

The πE5 beamline will be shared between the MEG-II and Mu3e experiments. This is allowed by the Compact Muon Beam Line (CMBL), which allows for the 3.2 m Mu3e solenoid to be placed in the πE5 area. Switching between detectors (from MEG II to Mu3e) is made possible by replacing MEG-II's superconducting beam transport solenoid with a dipole magnet (see Figure 2.6: ASL). The continuous muon beam from πE5 that satisfies Mu3e's requirements begins with the generation of 'surface muons'. These muons are produced from the decay of a stopped pion at the primary production target's surface. These muons are then transported through πE5 and focused on the Mu3e target.

2.3.2 Magnet

A homogeneous solenoidal magnetic field is required in the Mu3e experiment to allow the precise determination of momenta of the muon decay products. The magnet must also be able to focus the muon beam onto the hollow double-cone target, as well as have a maximum length of 3.2 m to fit within the tight constraints of the πE5 area. It is also required for field inhomogeneities to not stray above 10^{-3} within ±60 cm around the centre of the magnet. Over a period of 100 days of data-taking, the stability of the magnet should be very high, ($\Delta B/B \leq 10^{-4}$ from Ref. [118]). These strict requirements are satisfied by the Mu3e magnet, which was developed and constructed by Cryogenic Ltd.

2.3.3 Target

The beam of positive muons delivered by the πE5 beamline is directed towards the Mu3e target. The Mu3e target is a hollow, double-cone constructed from 'mylar', which is a synthetic polyester material. The double-cone has a maximal radius of 19 mm, and is 100 mm in total length. The thickness of the target is 70 μm on the upstream part of the cone and 80 μm on the downstream section. A schematic of the design is shown in Figure 2.7. Several considerations were taken into account in the design of the target: one aim was to maximise the stopping power of the muon beam, whilst also minimising the multiple scattering of the outgoing electrons, which would impair the momentum resolution of reconstructed tracks. Multiple designs were studied, but it was found that these goals are best achieved by the double-cone design [118].

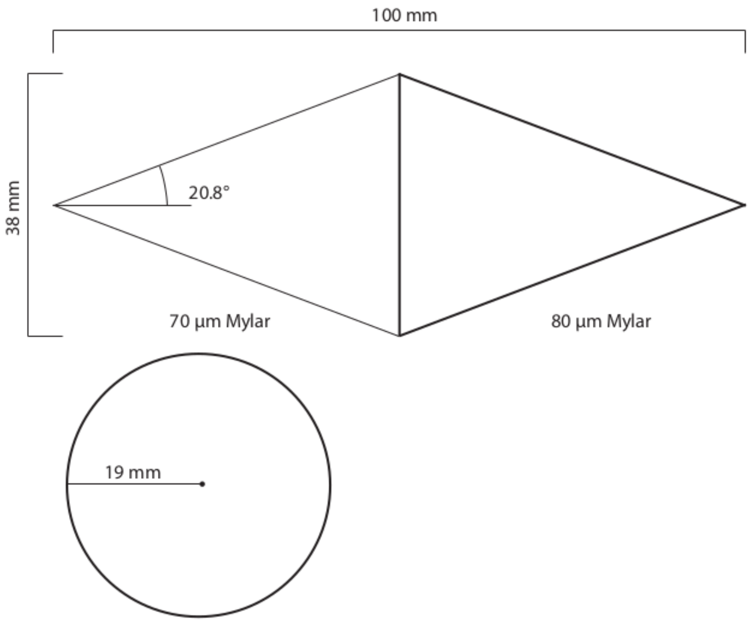

Figure 2.7: Schematic of the double-cone mylar target used in the Mu3e detector. Plot taken from Ref. [118].

2.3.4 Pixel Tracker

The Mu3e pixel tracker is a vital component to the Mu3e detector as it will provide highly precise hit information of the decay electrons that will be utilised for track, and ultimately, vertex reconstruction. The pixel tracker design is optimised for a minimal detector material budget, whilst also maximising the ability of reconstructed tracks to have superior momentum resolution. This is achieved by the detector design shown in Figure 2.8.

Primary features of the detector, namely, the existence of a 'central' and a 'recurl' station, have already been introduced in an earlier section discussing the experimental concept (shown in Figure 2.1). The success of the Mu3e experiment depends heavily on the ability of the detector to reconstruct vertices and the suppression of background. This design philosophy has lead to the creation of a pixel tracking detector consisting of multiple pixel layers of several cylinders at multiple radii around the mylar double-cone target. This will allow for a maximal acceptance, with the additional ability of having a momentum resolution of less than 1 MeV for recurling tracks, which is a requirement for background suppression. Within

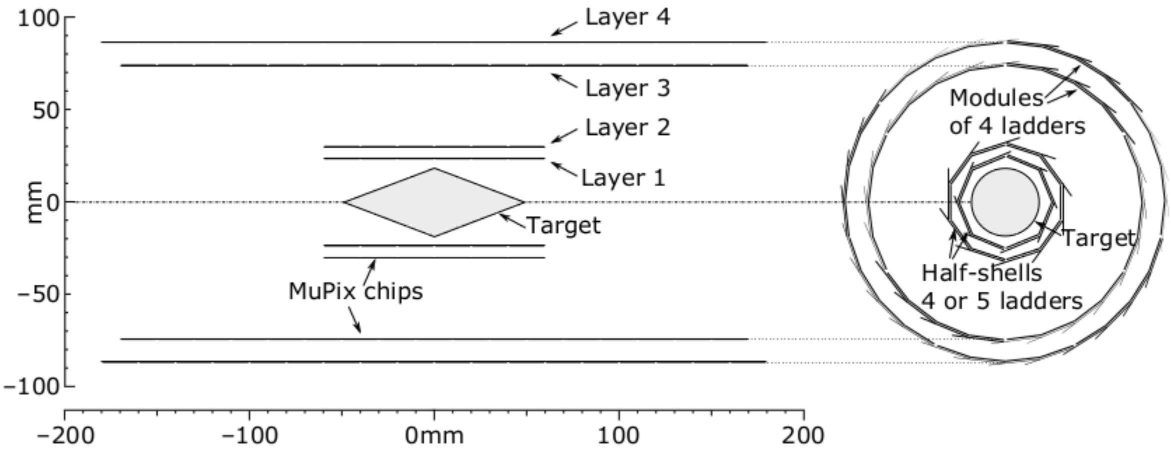

Figure 2.8: Schematic of the transverse and longitudinal points of view of the Mu3e pixel tracker, with target included. Figure taken from Ref. [118].

the central station, two layers are positioned similarly in maximum radius and length to the target. The positions of these layers close to the target allow for the precise spatial reconstruction of a vertex. It also allows for the heavy suppression of combinatorial background, as this precision can be used to resolve the locations of multiple separate muon decay vertices.

In both the central and recurl stations, as shown in Figure 2.8 two 'outer' pixel layers are located at radii 76 mm and 83 mm. The distance between these two layers and the vertex layers is optimised for momentum resolution. A particle is required to register hits across all four layers of the detector. With a magnetic field strength of 1 T, this limits the reconstruction to particles with a transverse momentum above 10 MeV, which corresponds to an acceptance of 38.1% of muon decays for $\mu \to eee$ decays simulated according to a four-fermion contact interaction.

Tracks that recurl back into the detector recurl either in the central or upstream/downstream recurl stations, where they will again be registered by the outer pixel layers. This increases the acceptance of the detector, providing the basis of hits required for the reconstruction of tracks with excellent momentum resolution. The acceptance is reduced however for tracks that recurl into the areas between the recurl and central stations, as these areas have no active sensors. Tracks may also have enough momentum along z to recurl past either the upstream or downstream detectors, which is observed for tracks of momentum higher than

layer	1	2	3	4
number of modules	2	2	6	7
number of ladders	8	10	24	28
number of chips per ladder	6	6	17	18
length [mm]	124.7	124.7	351.9	372.6
minimum radius [mm]	23.3	29.8	73.9	86.3

Table 2.2: Table shows a summary of the parameters of the various pixel layers.

40 MeV, at inclination angles of greater than 0.7 rad.

The fundamental unit of a layer in the pixel detector is the silicon sensor. These sensors, depending on the layer number, make up a ‘ladder’, which consist of 6 to 18 silicon sensors electrically bonded to a High-Density Interconnect circuit. When attached to two polyetherimide (PEI) endpieces (one upstream and one downstream), these ladders form a mechanically sound object defined as a module. Focusing on the inner detector, an example 3D-model is shown in Figure 2.9. Table 2.2 outlines the parameters of the various layers.

Two parameters dictate the possible momentum resolution that can be obtained with use of the pixel tracker. The first, relating to the total curvature of the tracks, is accomplished with the aforementioned cylindrical design. The other is the material budget, as this effects the multiple scattering angle Θ_{MS}. Materials with a long radiation length are preferred when considering detector material. A thin silicon sensor is thus of paramount importance, and is supplied by the MuPix 11 sensors [121], in High-Voltage Monolithic Active Pixel Sensor (HV-MAPS) technology. A schematic of the pixel sensor concept is shown in Figure 2.10. The support structure must also be light in material, and this is provided by the High-Density Interconnects upon which the silicon sensors are bonded. These High-Density Interconnects are constructed from thin layers of polyimide, glue, and aluminium layers [118]. A single pixel layer corresponds to $\sim 0.11\%$ radiation lengths.

During their use, the silicon chips will generate heat, and thus cooling is needed for the pixel tracker. This will be provided by a gaseous helium flow throughout the whole detector. A flow of 5 g/s of helium will be required to supply and cool the vertex detectors, whilst 16 g/s will be required for the outer layers. The mechanical stability of the former has been

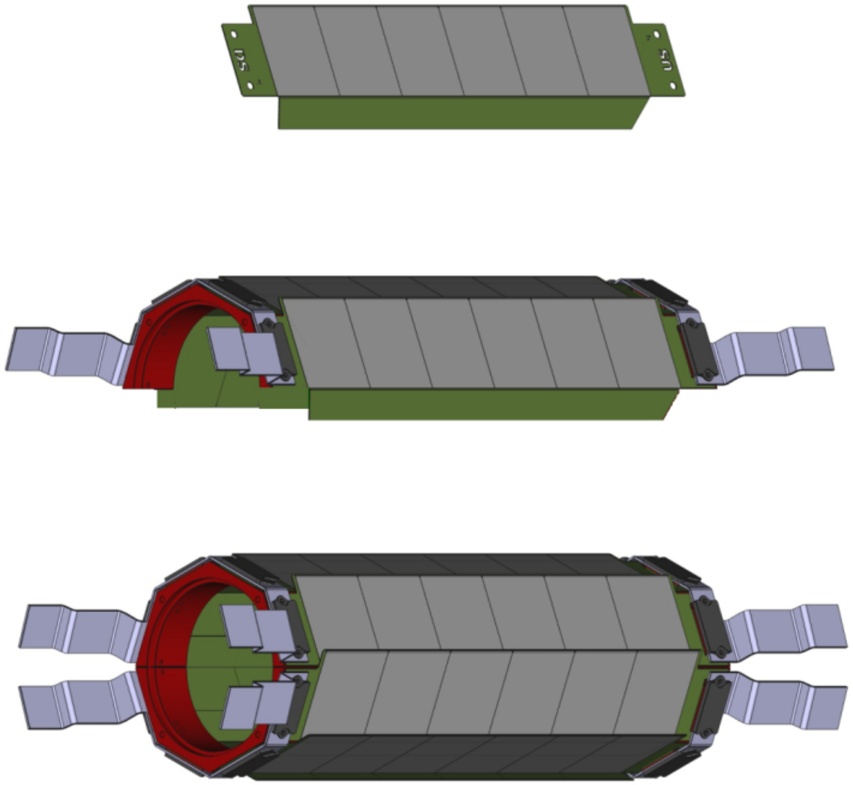

Figure 2.9: A 3 Dimensional CAD model of the deconstruction of a layer. Top to down shows a ladder, a module consisting of multiple ladders, and a layer consisting of multiple modules. Figure taken from Ref. [118].

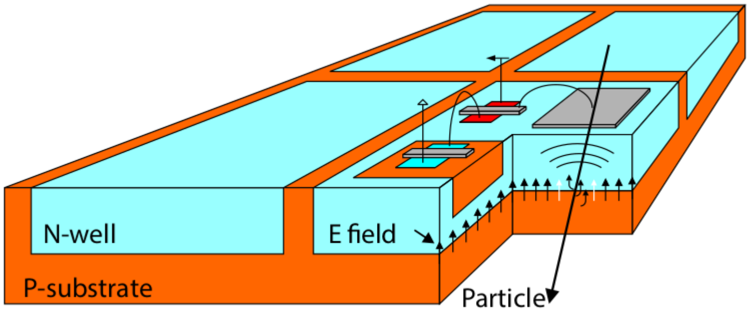

Figure 2.10: Schematic of a HV-MAPS detector. Figure is taken from Ref. [122].

tested, whilst the latter still requires study.

2.3.5 Timing Detectors

Timing information of tracks is an effective tool for suppressing combinatorial background. If timing measurements are available, they can be used in correcting tracks with a misidentified charge as it can be used to correctly identify a particle's direction of movement. Within the Mu3e detector, there are two scintillating timing detectors that provide timing information: the fibre and tile detectors.

The fibre detector has a radius of 61 mm and is 300 mm in length. As such, the roughly cylindrical shape is placed 5 mm below the outer pixel layers. This cylinder is composed of 12 ribbons, which each consist of three layers of scintillating fibre material, which are each 250 μm in diameter. Each ribbon is 32.5 mm wide and 300 mm in length. The fibre ribbons themselves are connected to multi-channel silicon photomultipliers at both ends of the detector. The full CAD drawing is shown in Figure 2.11, where an individual ribbon is also shown. A ribbon of three layers corresponds to a radiation length ($X/X_0 \approx 0.2\%$).

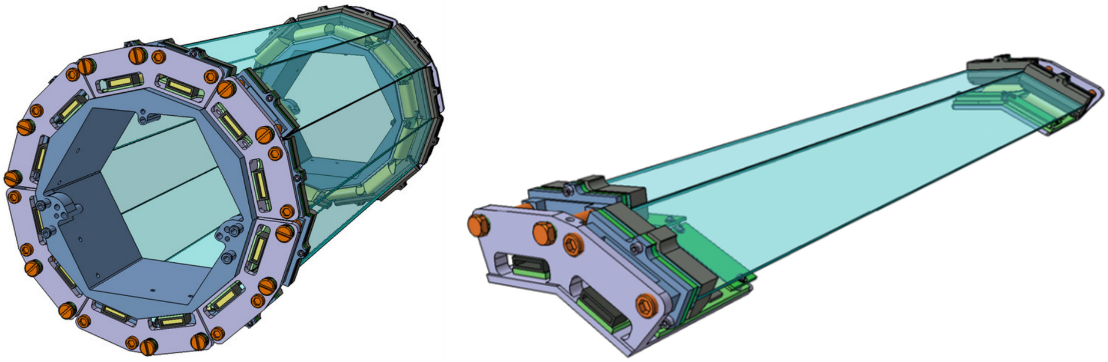

Figure 2.11: CAD drawing of the fibre detector. Drawing taken from Ref. [118].

The most precise timing information in the Mu3e detector is provided by the tile detector. Its position within the detector is essentially at the end of each track recurling into one of the recurl stations, and so material considerations are no longer relevant. However the detector must sit inside the recurl station barrel, within the outer pixel layers.

The aim of the time resolution of the tile detector is to be much better than 100 ps, with a 100% detection efficiency. The detector itself is split between the two recurl stations, with an outer radius of 62 mm and of length 342 mm. Figure 2.12 shows a CAD drawing with a partially expanded view of the detector, with the module and endrings labelled. The detector is segmented into cubes of plastic scintillator material, of dimension $0.5 \times 0.5 \times 0.5$ cm^3 which ensures high granularity of the detector, which can be used to suppress accidental background.

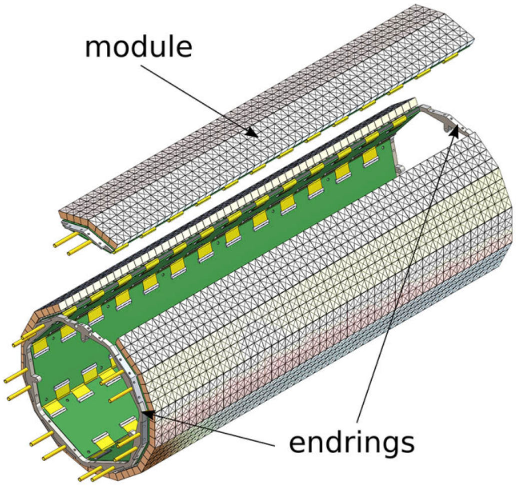

Figure 2.12: CAD drawing of the tile detector. Drawing taken from Ref. [118].

2.3.6 Data Acquisition

With regards to the measurement of muon decay products, the Mu3e detector will operate in a continuous mode when data is taken. Each of the sub-detectors will send a continuous stream of zero-suppressed hit-carrying information. The readout scheme is described pictorially in Figure 2.13, which shows the three clear levels of the scheme, which involve Field Programmable Gate Arrays (FPGAs), switching boards and ultimately a filter farm. This unique trigger-less approach forces the experiment to define ‘time-slices’ in which the Mu3e data is partitioned. These time slices are 64 ns in length, and are known as ‘frames’. This length can change based on the muon beam intensity. For Phase-1 analysis, it allows for an average of one muon decay per frame.

Data is continuously sent from the MuPix/MuTRiG chips and is collected by the front-end FPGA boards. This data contains timestamped hit information, which is sorted on the

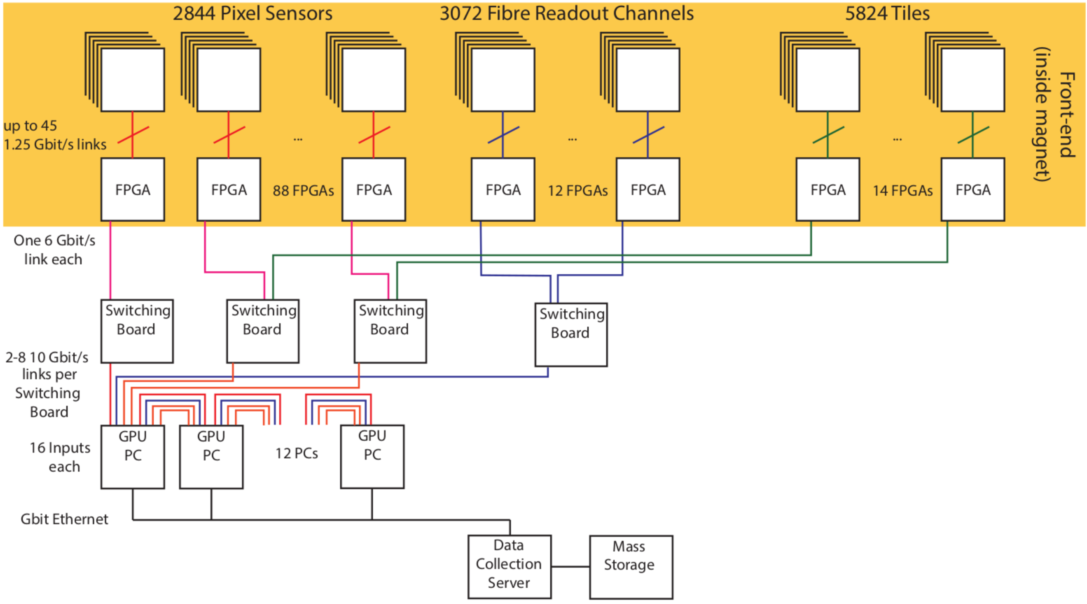

Figure 2.13: DAQ scheme of the Mu3e experiment. Drawing taken from Ref. [118].

FPGAs to ensure that all information reaches the rest of the readout system in a time-ordered format. The FPGA boards process this data and send the output onwards to the switching boards using optical links, whereby the switching board FPGA time-aligns and merges the multiple data streams into frames for input into the filter farm. The control of each sensor and pixel is achieved through a control link which connects the front end board to the chips. This allows temperature and other monitoring information to be measured and acted upon. A correction is applied on the hit time stamps on data supplied from the pixel detector, to account for the variation of the rise time with the deposited charge. This correction is based on the measured time over threshold. Regarding the scintillator layers, a particle passing can trigger multiple Silicon Photomultipliers. These form a cluster in time and space on the FPGA. Single hit clusters are suppressed, as these would otherwise dominate the registered hit [123].

The links from the switching boards are connected to the 12 GPUs that are connected via daisy-chain in the filter farm. If one buffer is filled, the data is forwarded along to the next PC in the chain. The location of the hits are extracted from the bitwise data, and if three hits in sequential pixel detectors are present, originating from the target region, a fast triplet fit is performed. If the fit is successful, this is followed by extrapolation to the fourth layer in search of a further hit. If a fourth hit is found within a set window, this defines a track. If three tracks are found in the corresponding frame, a simple vertex fit is then performed and, if this is successful, the frame is saved for readout [123].

2.4 Mu3e Software Framework

The features of the Mu3e Software framework and analysis methods are described in more detail in Chapter 3, with an accompanied analysis of a $\mu \rightarrow eee$ sample. The Mu3e simulation software is based on Geant4. The following sections will provide an overview of aspects of the Mu3e frame generation, simulation and reconstruction.

2.4.1 Reconstruction of Tracks and Vertices

The path of a charged particle within a homogeneous magnetic field is described by a helix. The reconstruction of such a path is difficult, and attempts have been made to solve the problem of their reconstruction through Kalman filters and broken line fits. However, both methods are computationally intensive due to complex matrix inversions. In Mu3e, track fitting is performed on the filter farm, so this demands the use of a high-speed, non computationally intensive algorithm. This algorithm must also take into account the main source of uncertainty in the experiment: multiple scattering due to the detector layers. The tracking algorithm used in Mu3e to fit tracks is based on a novel 'fast three-dimensional triplet fit' [124], which aims to provide a solution of reconstructing the helical path of a particle through the use of three dimensional triplet fits, the outputs of which can then be used to reconstruct tracks and vertices. The algorithm assumes the presence of a homogeneous magnetic field.

The algorithm begins by selecting three hits in consecutive layers to form the seed of the tracking algorithm in Mu3e, the triplet. Multiples of these triplets are combined to form track segments. Within the magnetic field, the track travels in a helical form. Assuming little or no energy loss, ten parameters [125] are required to describe the triplet. These parameters come from the distance between hits, the direction, and the initial position, as well as the two multiple scattering angles at the location of the scatter and curvature. As a result of the theory, the two multiple scattering angles have to be minimised, and these angles are dependant on the curvature. The algorithm aims to estimate the three dimensional radius R_{3D} that minimises the multiple scattering angle. More specifically, the χ^2 function:

$$\chi^2(R_{3D}) = \frac{\Phi_{MS}(R_{3D}^2)}{\sigma_\Phi^2} + \frac{\Theta_{MS}(R_{3D}^2)}{\sigma_\Theta^2}, \tag{2.5}$$

where Φ_{MS} is the azimuthal scattering angle, and Θ_{MS} is the polar scattering angle. The widths Φ_{MS} and Θ_{MS} can be estimated using Eq. 2.1. A fast multiple scattering fit is performed on a triplet, and if the track passes a χ^2 threshold, the track is extrapolated to find a fourth hit. Track segments are created from triplets by selecting two triplets that share the same two hits, and then forming a segment of four hits. In Mu3e, a segment of four hits is a 'short' track. Up to 6 triplets can be used to form tracks of 6 or 8 hits, known as 'long-6' and 'long-8' tracks respectively. An example long-8 track is shown in Figure 2.14. Long tracks have superior momentum resolution, as shown in Figure 2.15, whereby a minimum at roughly 20 MeV is observed, which corresponds to a track that recurls back towards the detector, giving a curvature angle of roughly π.

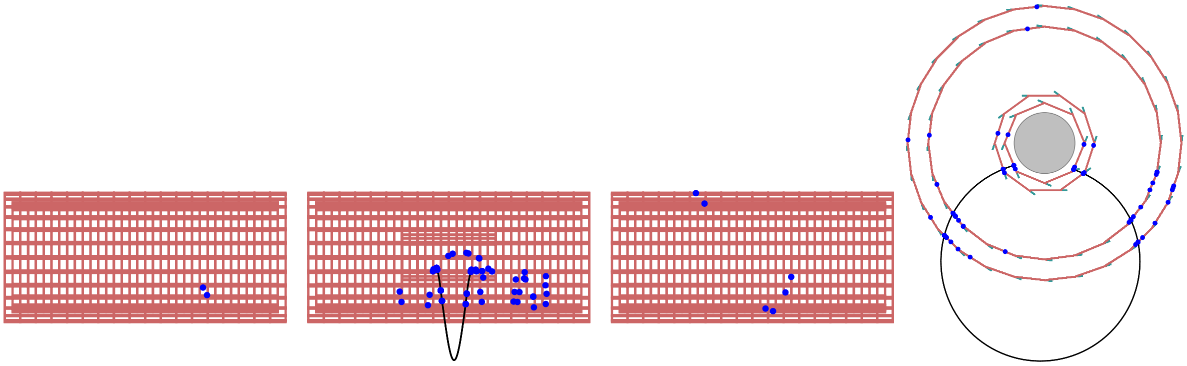

Figure 2.14: Longitudinal and transverse view of a reconstructed track in the Mu3e event display.

Any combination of three tracks, (two positron candidates and one electron candidate) is then used in an attempt to fit a vertex. In one method, the vertex location is estimated by

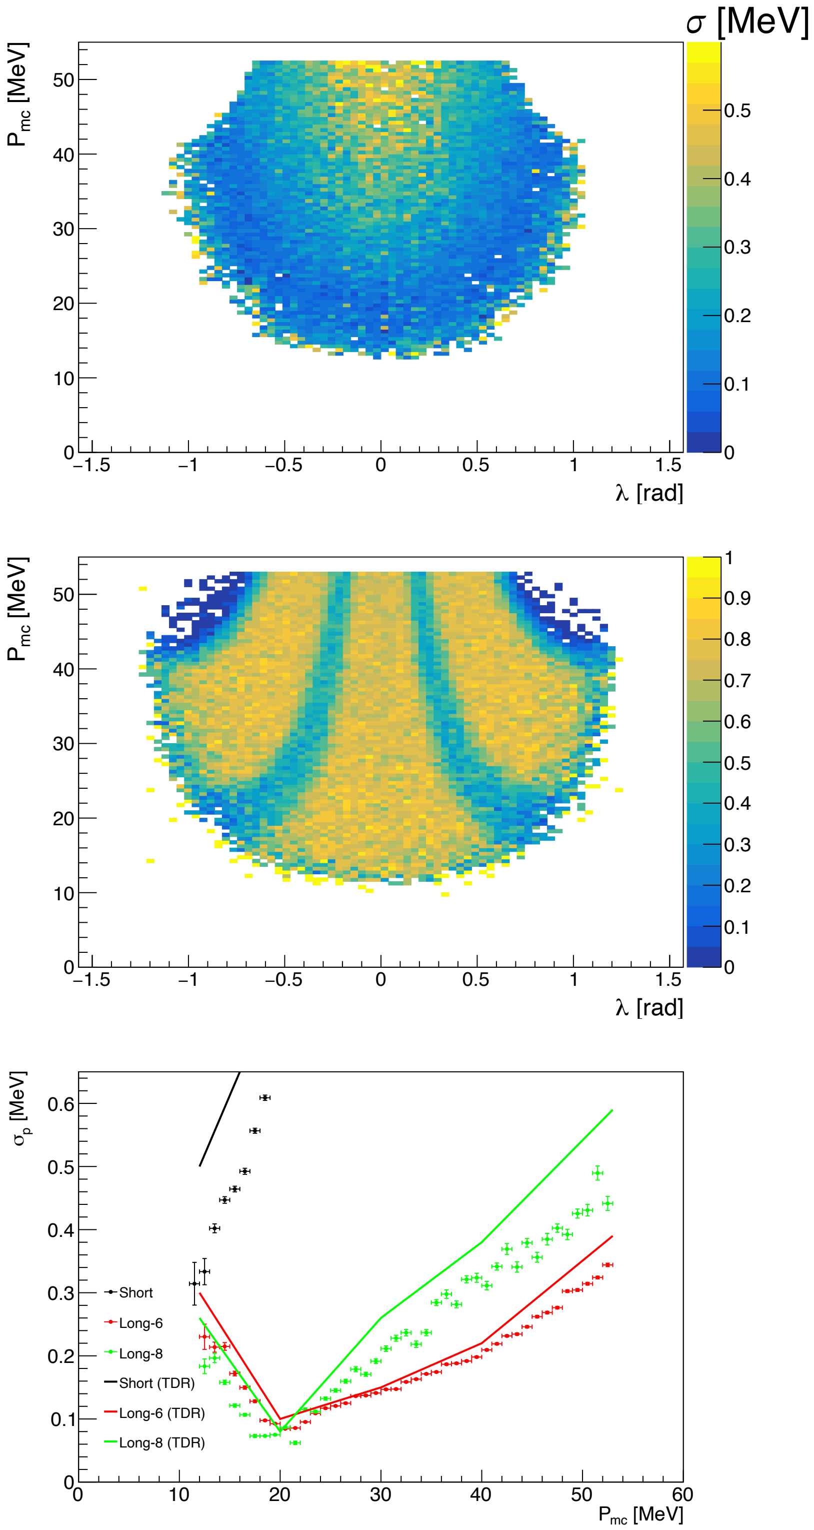

Figure 2.15: The acceptance, momentum as a function of inclination angle λ and the momentum resolutions of short, long-6 and long-8 tracks.

minimising the scattering angles of each track at the innermost vertex layer. A new, updated vertex fitter has also now been written for Mu3e, which also takes into account the multiple scattering in the target, as well as the pixel resolution of the silicon sensors.

2.4.2 Simulation

The Mu3e software framework is based on the Geant4 [126] simulation toolkit, which accurately models the behaviour and interactions of particles as they pass through matter. Simulated with the toolkit are each of the main Phase-1 subdetector components, as well as several of the services that will hold the subdetectors in place. The dimensions and location of each of the detectors is implemented as described by the Mu3e Technical Design Report [118], and is as close as possible to the modelled CAD drawings of the detector. The detector modelled in Geant4 is shown in Figure 2.16. The muon beam is not fully implemented. The beam of muons is currently generated with Geant4, 1 m upstream from the Mu3e target. It is possible to input an arbitrary beam profile that would more accurately describe the muon beam, to suit the needs of the simulation.

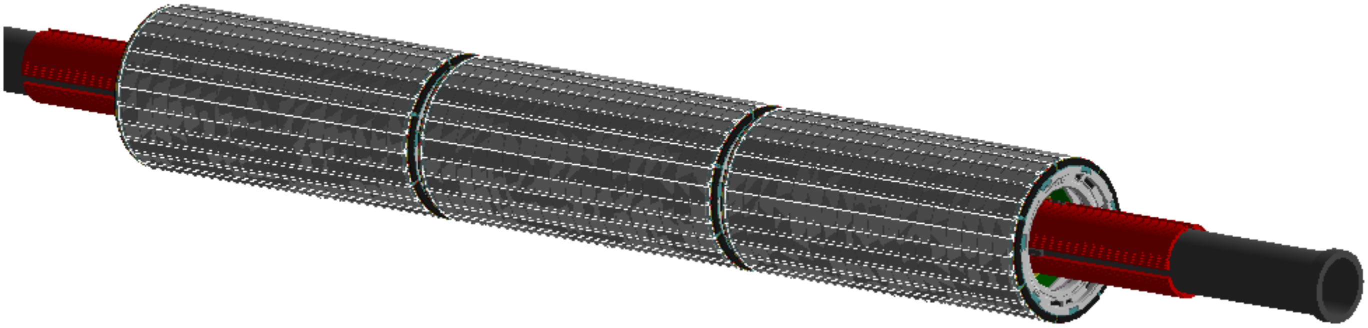

Figure 2.16: Geant4 modelling of the Phase-1 Mu3e detector.

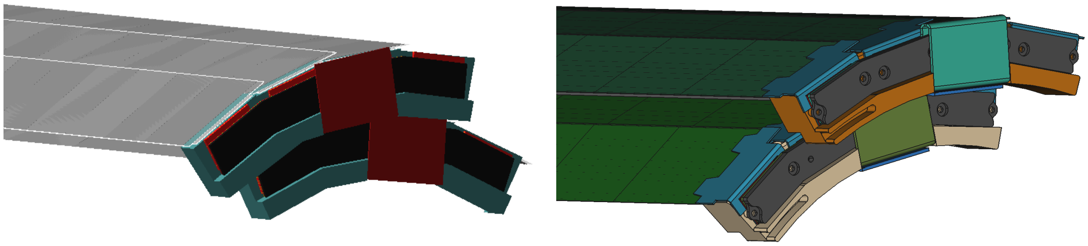

Figure 2.17: Side by side comparison of Geant4 and CAD model of the current outer pixel layer endpiece.

The simulation can also very easily be modified to either reflect changes in the detector design or to experiment with potential Phase-2 scenarios. In Figure 2.17, the Geant4 model outlines the result of modifying the PEI endpieces (light blue) to more closely reflect the CAD detector design. Carbon fibre plates (in black) also required modification. These changes were due to a change of angle of the step in the endpiece from $0.5°$ to $1°$. Other changes can be made to evaluate possible future detector designs, such as the addition of an extra outer layer at a high radius.

The use of Geant4 in the simulation allows for the accurate modelling of a particles path throughout the modelled detector, and accurately simulates detector responses and particle behaviour through the material. Physical processes such as Bhabha scattering, photon conversion, multiple Coulomb scattering, bremsstrahlung and ionisation can all induce a loss of energy to any of the decay electrons, and are all modelled within the framework.

2.5 Phase-2 Mu3e

The muon beam at PSI will undergo a long shutdown in 2027-2028 to implement the High Intensity Muon Beam (HIMB) upgrade. The HIMB will provide Mu3e with a muon beam rate greater than 10^9 muons/s. This presents the current experiment with a number of challenges to handle; not only the increase in rate, but also in the combinatorial background. This increase in rate will render the fibre detector obsolete due to high occupancy, and will therefore need to be replaced.

Several upgrades to the Mu3e experiment have been proposed. To minimise the gap in the acceptance of Mu3e (see Figure 2.15, extra recurl stations could be placed upstream and downstream of the initial Phase-1 design. This would allow Mu3e to obtain a higher sensitivity. This is shown pictorially in Figure 2.18. A search for $\mu \rightarrow e\gamma$ could also be undertaken with the addition of a photoconversion layer at high radius. The photons would photoconvert in the photoconversion layer and the resulting hits from the electron products are measured by additional two pixel layers at high radii. This is shown by the schematic in Figure 2.19.

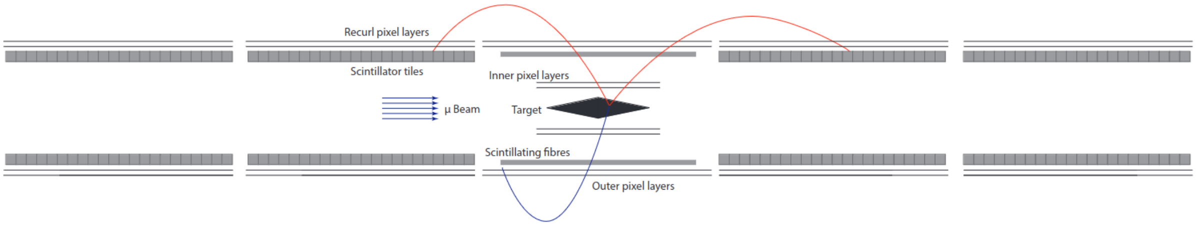

Figure 2.18: Proposed Phase-2 schematic of Mu3e, an additional two recurl stations are placed upstream and downstream of the original Phase-1 design.

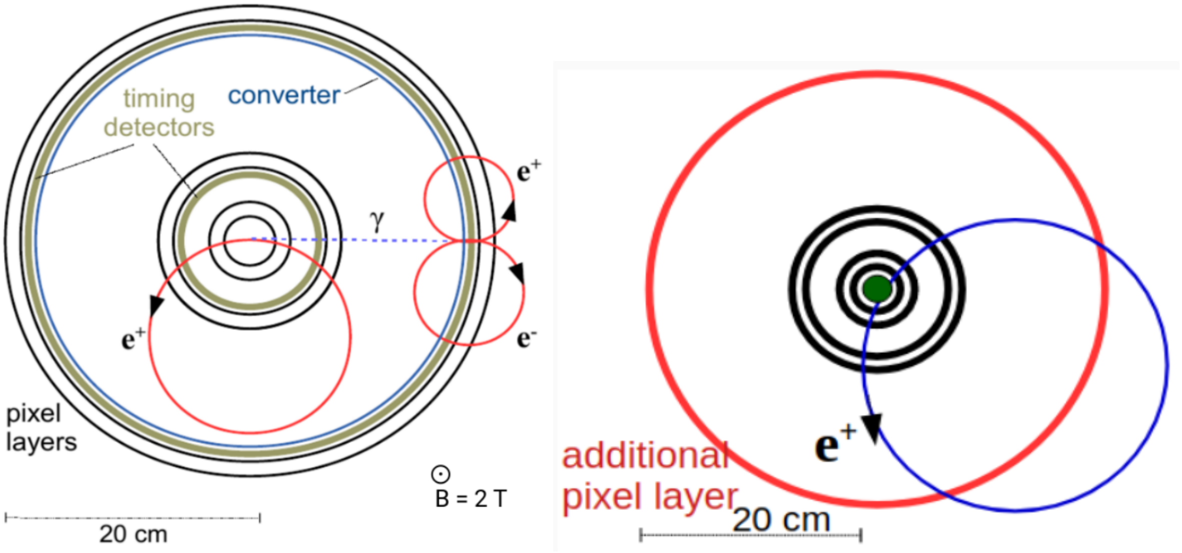

Figure 2.19: Proposed Phase-2 Mu3e gamma schematic, an additional layer at 20 cm is proposed. Alongside it, a sketch of the extra outer pixel layer design is shown.

A third proposal is the sole addition of a single pixel layer at a high radius. This would allow the study of high momentum particles with a resolution of less than 100 keV, which would be very interesting when studying exotic decays of the muon. A study into this Phase-2 design has been undertaken within the context of this thesis, and the results of which is described in Chapter 6.

Chapter 3

Mu3e Simulation Framework and Reconstruction

This chapter provides a detailed description of the Mu3e software and analysis framework. It will begin with the simulation of the Mu3e detector, as well as the track reconstruction algorithms and cuts used for tracks in Phase-1 Mu3e. A description of the vertex reconstruction algorithms will follow. Finally, the sensitivity to $\mu \rightarrow eee$ will be studied using the latest available version of the Mu3e simulation package, validating the results against those published in the Mu3e technical design report [118].

The presented offline software workflow is the basis for all Phase-1 software and performance analysis with Mu3e. The simulation and offline reconstruction is used in the dark photon and axion-like particle analyses, the latter of which which will be expanded upon in Chapter 6. The former is discussed in this chapter.

3.1 Event Generation

The Mu3e software framework accurately models the performance of each of the subdetectors using the Geant4 simulation toolkit. The simulation is the basis for all sensitivity studies undertaken by the collaboration, and it is used in all analyses contained within this work. Figure 3.1 shows the Geant4 modelling for the full Phase-1 detector, as well as the Fibre detector and the vertex detector.

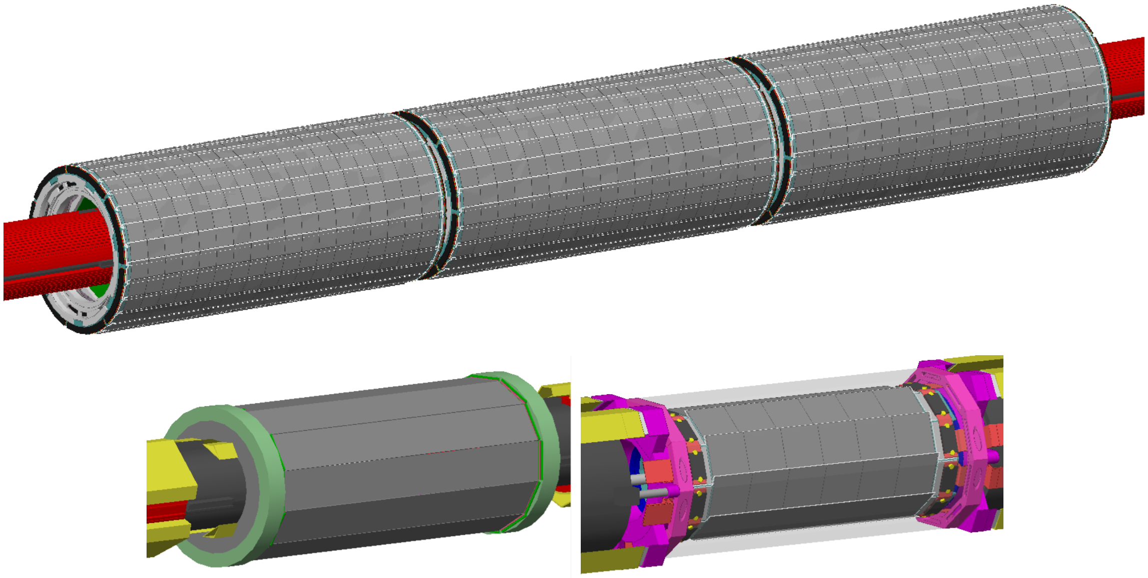

Figure 3.1: Figure shows different components of the simulated Mu3e detector.

The simulation software emulates the behaviour of particles as they pass through material, and the responses of each of the subdetectors as particles pass through the silicon layers or fibre material. The data acquisition is also ‘simulated’, which is conducted before the simulation is ran through event generation. As such, the output of the simulation is recorded in ‘frames’, as it would be in the final experiment. One consequence is that the subdetector readout does not always produce pixel hit information in the ‘frames’ in which they are produced. Therefore, after running the simulation, a ‘frame sorting’ is conducted. This sorting emulates the time ordering that would be required in the final experiment. At least 116 frames are required to correctly sort the simulated output.

The simulation software used for the work presented in this thesis has several shortcomings. For example, timing information for both the fibre and tile scintillators has not been fully implemented. When reconstructed variables are considered, a 90% cut on the overall efficiency due to timing is assumed. The simulation also, as standard, does not take into account the effect of silicon pixel clustering. If used in the simulation, the number of fake tracks will increase, due to overlapping clusters.

3.2 Triplet, Track and Vertex Construction

The methodology of track and vertex reconstruction has been introduced in Chapter 2. Regarding nomenclature, tracks are categorised into three main types for the Phase-1 analysis: short, long-6 and long-8 tracks. Their appearance for transverse and longitudinal planes of the detector is shown in Figures 3.2 and 3.3. These figures show clearly the recurling nature of the tracks due to the presence of the homogeneous magnetic field. The curvature of tracks is used to determine its charge. For long-8 tracks, timing information from the fibre detector is required to determine the direction of the particle, and thus its charge.

The long-8 track can be seen in Figure 3.3 to be associated with many additional hits in the silicon layers, as tracks in the central detector pass through the detector multiple times. These hits cause confusion for the track reconstruction algorithms, and so may produce fake reconstructed tracks.

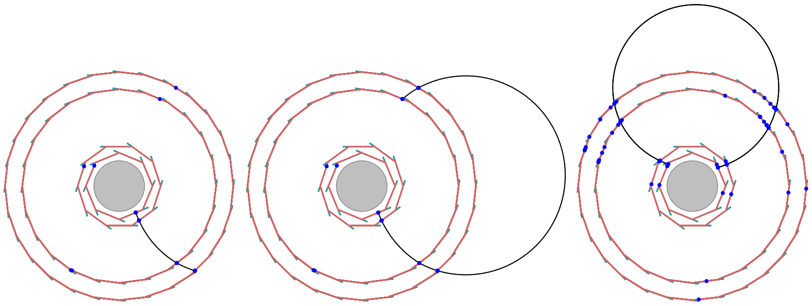

Figure 3.2: Different track topologies on the x-y plane for short, long-6 and long-8 tracks.

To select good quality tracks, cuts are applied on the difference in ϕ angles between hits in

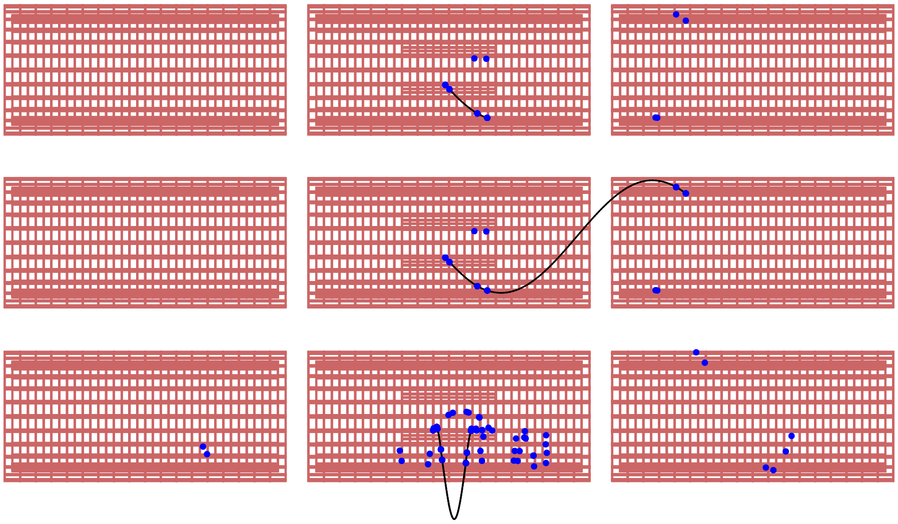

Figure 3.3: Different track topologies on the r-z plane for short, long-6 and long-8 tracks.

sequential layers, z-coordinate locations of hits and the radii of the triplets that form the track. The χ^2 of the track fit is used to determine the quality of the track fit. Summaries of the cuts used for specific tracks are shown in Tables 3.1, 3.2 and 3.3 below, and are used in the $\mu \to eee$ sensitivity study. The symbols $z_{x\to y}$ and $\phi_{x\to y}$ refer to the window that is opened for the coordinate between layer x and layer y.

Requirement	Value
Triplet transverse radius [mm]	> 30
$\phi_{1\to 2}$ [rad]	$\Delta\phi < 0.8$
$\phi_{2\to 3}$ [rad]	$\Delta\phi < 0.05$
$\phi_{3\to 4}$ [rad]	$\Delta\phi < 1.25$
$z_{1\to 2}$ [mm]	$\Delta z < 30$
$z_{2\to 3}$ [mm]	$\Delta z < 180$
$z_{3\to 4}$ [mm]	$\Delta z < 3$
χ^2	< 32

Table 3.1: Table of variable windows used to reconstruct short tracks.

If there are three reconstructed long tracks (two positively charged, one negatively charged) within a framelength of 64 ns, these tracks are then inputted into the vertex fitter in an attempt to reconstruct a vertex. As discussed in Chapter 2, from the first silicon layer, the tracks are extrapolated back towards the target. If the three tracks meet at the same point in space and time, a vertex is reconstructed. Several quantities are used to evaluate the quality of the vertex fit, and Mu3e uses several selection criteria in order to select signal

Requirement	Value
Triplet transverse radius [mm]	> 30
$\phi_{4\rightarrow5}$ [rad]	$\Delta\phi < 0.35$
$\phi_{5\rightarrow6}$ [rad]	$\Delta\phi < 0.035$
$z_{4\rightarrow5}$ [mm]	$\Delta z < 100$
$z_{5\rightarrow6}$ [mm]	$\Delta z < 3$
χ^2	< 48

Table 3.2: Table of variable windows used to reconstruct long-6 tracks.

Requirement	Value
Triplet transverse radius [mm]	> 30
$\phi_{4\rightarrow5}$ [rad]	$\Delta\phi < 0.35$
$z_{4\rightarrow5}$ [mm]	$\Delta z < 100$
χ^2	< 48

Table 3.3: Table of variable windows used to reconstruct long-8 tracks.

vertices whilst also completely rejecting background.

3.3 Reconstruction

3.3.1 Selection Criteria

The event selection criteria are chosen such that all background is minimised within the signal region, whilst still retaining a significant amount of signal vertices. This is achieved through the following selection, which is identical to that used in the Mu3e technical design report [118].

In monte-carlo (MC) studies, the selection criteria use a number of generator level and reconstructed variables. The use of generator level information in MC studies ensures that the efficiency quoted for a particular decay process is purely for that process. In some cases, the use of generator level information hides an inadequacy of the current reconstruction algorithm, such as the existence of recurling tracks within the detector, which could contribute to background. Recurling tracks can be removed in MC studies by only choosing tracks on their first recurl, which is information that can only be obtained using generator level information. When running the actual experiment, no generator level information can be used, and analyses are currently ongoing to optimise the selection criteria using solely reconstructed variables.

Requirements at Generator Level

The first selection evaluates the location of the truth vertex, otherwise known as a 'target region' requirement. This ensures that the true vertex location originates from within a predefined region within the detector, which is a cylinder around the Mu3e target. This cylinder completely surrounds the Mu3e target, with a radius of 19 mm and a length of 50 mm. The aim of this requirement in particular is to remove vertices of muons that decay in flight.

The remaining generator level requirements apply to each of the three tracks that form the vertex. The first of these enforces the tracks that form a vertex to be on their first recurl, removing troublesome tracks that recurl many times in the detector. This is represented by a hit sequence number of hits in a reconstructed track, where the reconstructed hits are matched to generator level hits (truth matched). For tracks radially moving outwards, the hit sequence will be a positive integer. For tracks radially moving inwards, the value will be negative. The cut is pre-defined to remove tracks with a hit sequence of less than ± 3.

Another requirement ensures that the vertex is reconstructed using tracks that originate from the same mother particle. This ensures that, when looking at processes that originate from only one vertex, tracks that constitute a vertex are not originating from other sources. This would be modified in processes that occur through displaced vertices, or with a study into combinatorial background.

Vertices that are reconstructed from tracks on their N^{th} ($N > 1$) recurl through the detector are ignored within the MC studies of Mu3e. Tracks that would contribute to these cases are removed using a requirement that hits from a reconstructed track must match a track created using generator level hits. This reconstructed track must also be on its first recurl, moving radially outward.

Often, the analysis only examines reconstructed vertices from a specific decay process. A requirement using generator level information enforces only a particular process to be analysed. Information carried from the Geant4 simulation contains within it the information of the decay process from which a track originated. Using this, it is possible to filter out decay processes that are irrelevant to an analysis.

Each track in the simulation is given a track identification number. This is carried from the Geant4 simulation up toward the reconstruction. To ensure vertices are not reconstructed with the same track, or with tracks that are duplications of the same track (created from a track that recurls back into the central station), each track that forms a vertex is required to have a unique track identification number (differing from the other two tracks that make the vertex), and has also not been used in the reconstruction of another vertex.

Reconstructed Information

To assess the quality of a reconstructed vertex, several quantities are used that refer to the properties of the reconstructed muon decay. There are various requirements on these quantities which are used to ensure background minimisation and signal maximisation. These requirements will be defined in the following: firstly, a requirement on the χ^2 of the vertex fit[1] is used. This aims to reduce the number of poorly-fitted vertices. It also reduces vertices due to accidental background, as not all tracks involved in these vertex reconstructions have truly a common origin. Mu3e, due to the beam current, expects one muon decay every 10 ns in Phase-1. If more than one vertex is reconstructed within the frame, this is defined as vertex multiplicity. If multiple vertices are reconstructed, the vertex with the lowest spatial χ^2 of the vertex fit is kept. In the simulation, if generating a particular process, only one muon per frame is generated to decay through a desired process, unless otherwise needed for combinatorial studies.

---

[1] Defined as spatial χ^2 henceforth.

Further requirements exploit that the muon is decaying at rest. As mentioned in Chapter 2, the total momentum of the decay electrons from the $\mu \rightarrow eee$ process vanishes. This is not the case for scenarios involving decays with undetected particles, like the internal conversion process. A requirement on the total reconstructed momentum therefore will severely reduce the amount of background.

The properties of the reconstructed positron-electron pairs that formed the vertex can also be examined. As there are two positrons and one electron that form a vertex, two positron-electron pairs can be formed. In order to reduce the amount of combinatorial background due to Bhabha scattering, the invariant mass of the pair can be constrained. The requirement can apply to either the lower (M_{ee}^{Low}), or higher (M_{ee}^{High}). This greatly reduces the number of combinatorial background events due to Bhabha scattering, with the side effect of also reducing the signal efficiency.

Finally, there is a requirement on the invariant mass of the three electron system. For $\mu \rightarrow eee$ events, a peak is expected at the muon mass. The tail of the internal conversion events will leak into the signal region. This cut greatly reduces the number of background events.

The signal region is defined as events passing all of the above requirements. Further requirements that can in the future be added to the section include: a cut on the distance between the reconstructed vertex and the Mu3e target; information from the fibre and tile detectors. A cut of 90% due to timing is assumed in every final efficiency. Example plots from a simulated $\mu \rightarrow eee$ sample are shown in Figure 3.4.

3.3.2 Normalisation

Measuring the efficiency of selecting decay processes within Mu3e is computed with respect to the number of muon stops upon the surface of the target, as it is not possible to generate the total massive amount of muons decays that Mu3e will observe. One muon decay is expected every 10 ns with the Phase-1 Mu3e beam (with a 100% stopping rate). In the simulation, the number of generated events for a given process can be verified using generator level information, as tracks may exist beyond the frame length, leading to double counting. For example, if at least one track from a muon decay is present within a frame, and the track originates from the correct process, and the track is also unique and the vertex originates from the target region, the muon decay is qualified as good, and is counted in the normalisation.

A recurling particle could exist in multiple frames. Using the method outlined in the previous paragraph, such 'duplicate' frames are ignored. In the simulation also, if a particle from a previous frame exists in the following frame, no new muon decay is generated. This leads to the correct calculation of the normalisation, as you remove all duplicate frames. In the Mu3e simulation, it is observed that 82±1% of frames contain a muon that decays on the surface of the target. All muons that reach the target will decay at rest on its surface.

3.4 $\mu \rightarrow eee$ Study

A sample of $\mu \rightarrow eee$ is generated using the Mu3e software framework. The decay is simulated assuming a four-fermion contact interaction. The standard track and vertex reconstruction algorithms are implemented, and an analysis identical to that conducted within the Mu3e technical design report is implemented using the ROOT software framework [127]. A cutflow of the cuts described above is shown in Table 3.4, where, after applying the final timing cut, a final efficiency of 12.8% is observed, compatible with that reported in the TDR (13%), which was based on an earlier version of the software [128]. Before the application of any cuts, plots of the main reconstructed properties of the events are plotted, and are shown in Figure 3.4.

The cuts in Table 3.4 are split into two sections. The first section involves the use of truth information to remove undesired vertices. The second section uses measured information obtained from the use of the Mu3e detector. When looking at particular samples, such as Bhabha scattering or internal conversion, the vertices are weighted to a differential decay rate, for signal events each vertex is weighted as 1. The requirements were briefly described

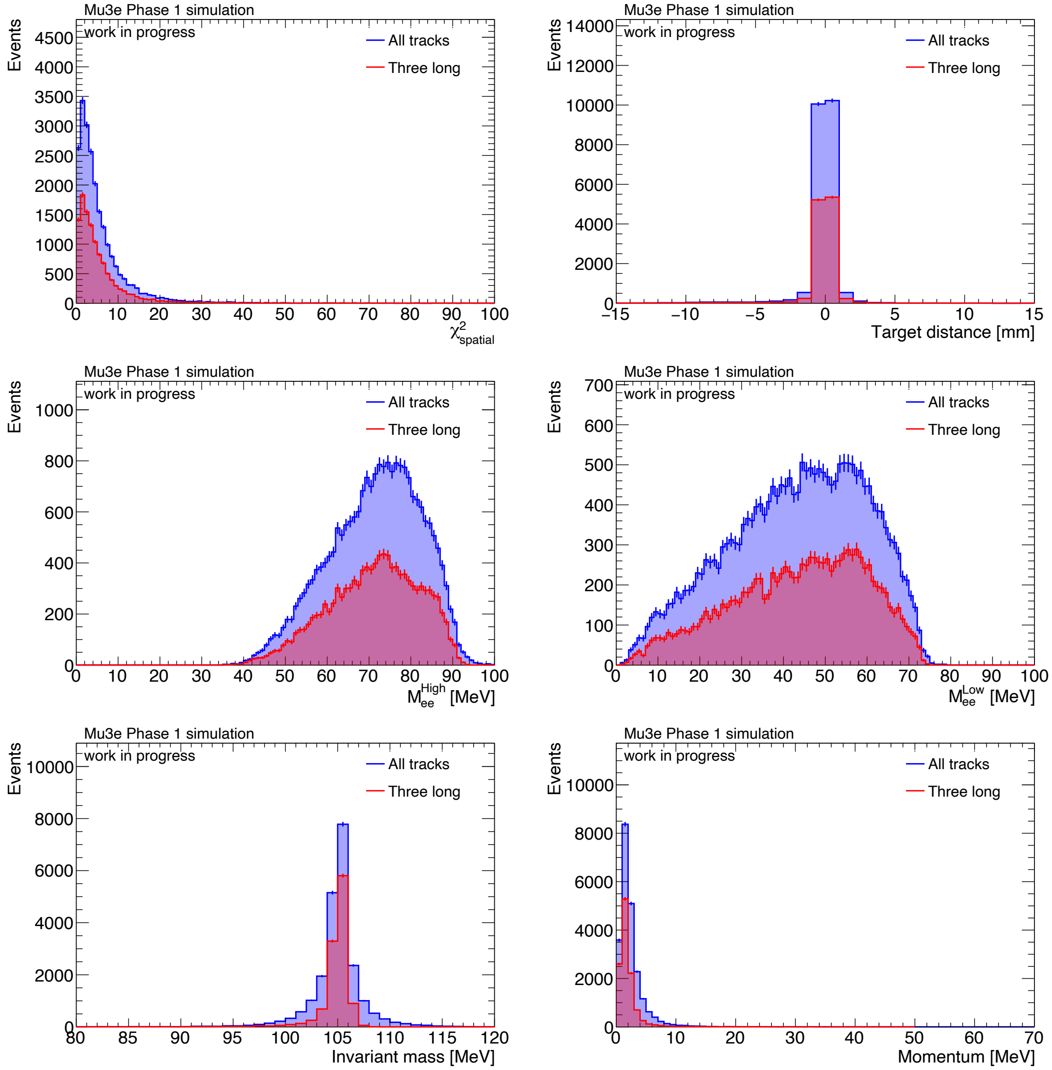

Figure 3.4: Signal sample variables for reconstructed signal vertices with all and 3 long tracks.

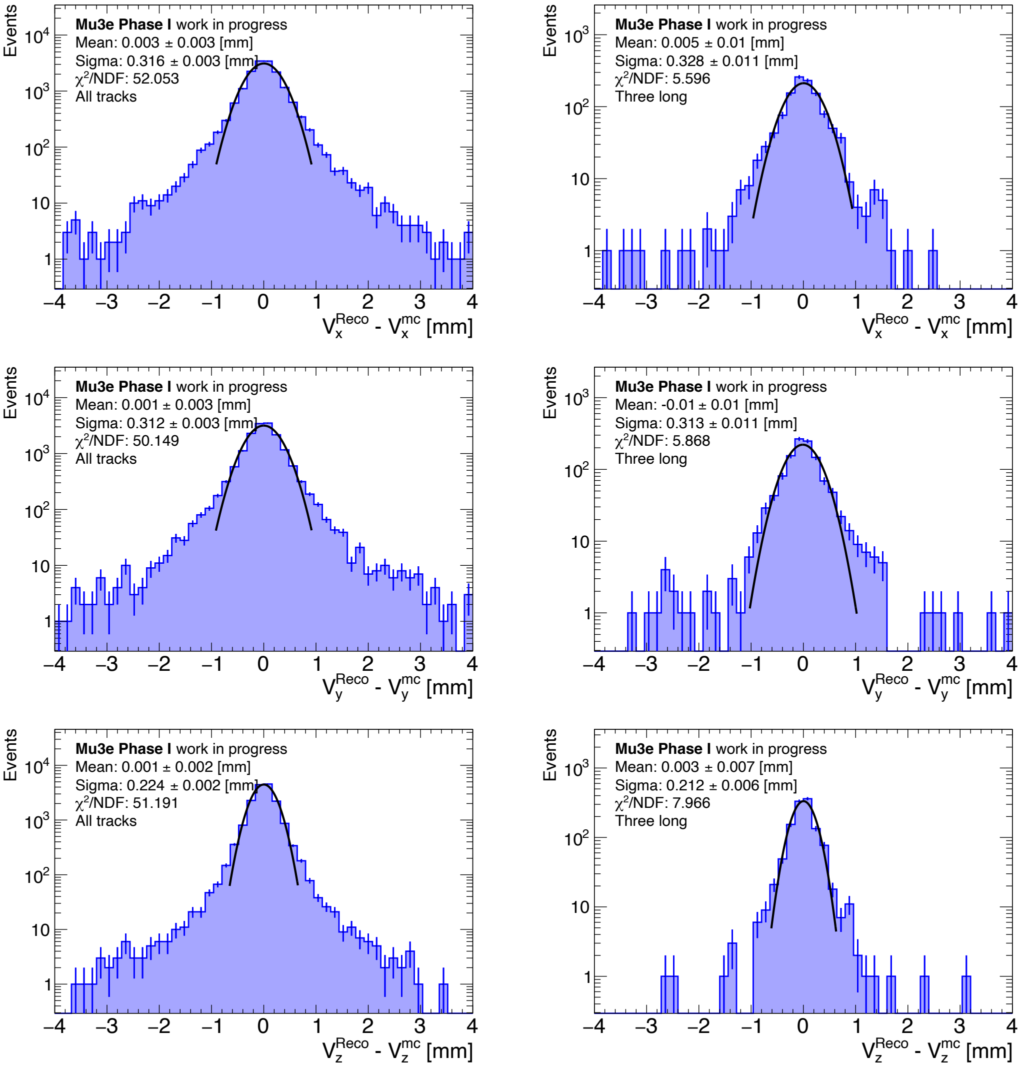

Figure 3.5: Vertex resolutions for each Cartesian coordinate for a sample of reconstructed signal vertices in the signal region. All track vertices and 3 long track vertices are considered. All cuts are applied, a Gaussian fit is performed to estimate the mean and sigma of the distributions.

Selection	N_{Vertices}	Σw	ϵ [%]	$\sqrt{\Sigma w^2}$/norm [%]	ϵ_{step} [%]
Initial reconstructed vertices	24569	24569	-	-	-
Target region	23840	23840	35.25	0.23	97.03
Track hit number < 3	11842	11842	17.51	0.16	49.67
Track same mother ID	11710	11710	17.31	0.16	98.89
Track 'prime'	11502	11502	17.01	0.16	98.22
Track type	11502	11502	17.01	0.16	100.00
Track unique	11502	11502	17.01	0.16	100.00
Vertex Multiplicity	10714	10714	15.84	0.15	93.15
Vertex $\chi^2_{\text{spatial}} < 15$	10708	10708	15.83	0.15	99.94
$P_{\text{eee}} < 4$	10268	10268	15.18	0.15	95.88
$M^{\text{Low}}_{\text{ee}} < 5$ or $M^{\text{Low}}_{\text{ee}} > 10$	10056	10056	14.87	0.15	97.94
$103.00 > M_{\text{eee}} < 110.00$	9762	9762	14.43	0.15	97.08

Table 3.4: Cutflow for a $\mu \rightarrow eee$ sample. Only three long tracks are considered for the vertex. Final efficiency is comparable to the efficiency outlined in the Mu3e technical design report. Timing cuts and a cut on the vertex to target distance are not used. 80,000 frames are simulated, with one $\mu \rightarrow eee$ decay expected per frame, with 67,634 muons found to decay on the surface of the target. The denominator is the number of muons decaying on the target.

in the previous sections. The target region cut removes vertices outside of the target region, the track hit number requirement ensures that the tracks are moving radially outwards. The same mother requirement ensures that all tracks originate from the same muon, track prime ensures that the tracks are matched correctly to their monte-carlo truth counterparts, track type ensures that there are two positrons and an electron that forms the vertex, and track unique ensures that all three tracks do not have the same truth counterpart. These cuts are used in all analyses. The remainder of the cuts use reconstructed information from the three track vertices. The final efficiencies in analyses will be quoted from after reconstructed mass requirement is applied on a sample. Studies are ongoing to create an analysis that uses no generator level information. Analysis of the events after the full selection is applied produces the vertex resolution plots shown in Figure 3.5. Using the reconstructed efficiency shown in the previous paragraph, with the assumption that there is no background in the signal region, the Poisson upper limit can be estimated. The final sensitivity as a function of beam time is shown in Figure 3.6. Here, the sensitivity is calculated with respect to three significances: single-event sensitivity, 90% CL and 95% CL upper limit, referring to

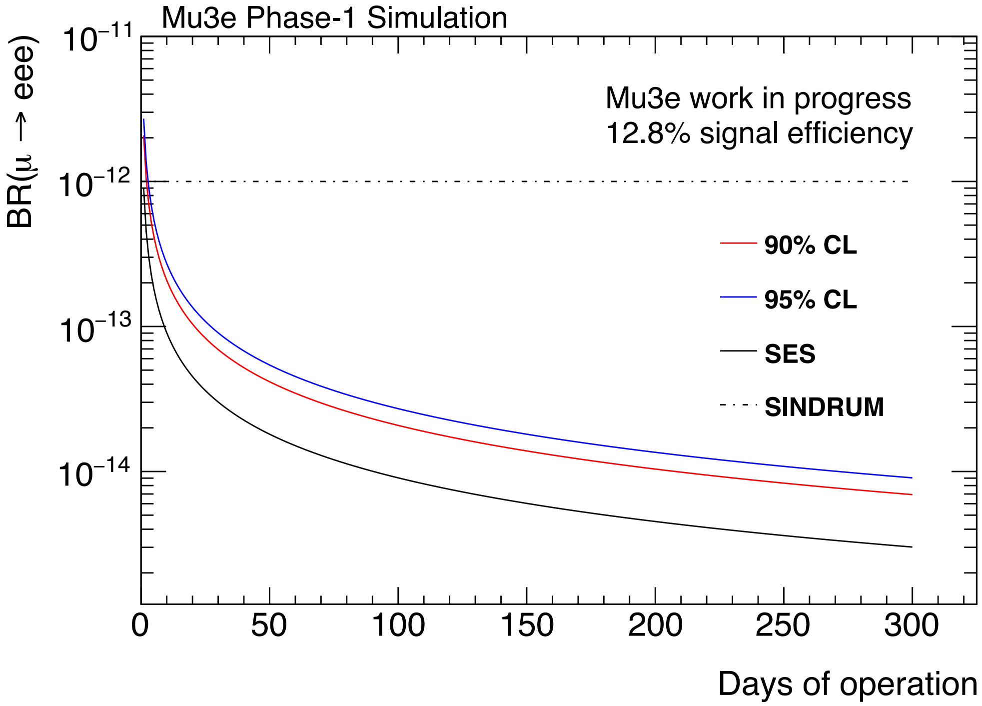

Figure 3.6: Single-event sensitivity for the $\mu \to eee$ signal at 90% CL and 95% CL as a function of the number of days of operation of the experiment.

three confidence intervals found via inversion of the Neyman belt [129]. The sensitivity with respect to the significance is calculated as an upper limit to the branching ratio as calculated by Eq. 3.1:

$$\mathrm{BR}_{\mathrm{UL}} = \frac{Z}{\epsilon \times \mathrm{N}_{\mathrm{muons}}}, \tag{3.1}$$

where $\mathrm{BR}_{\mathrm{UL}}$ is the upper limit on the branching fraction, Z is the significance, and is equal 1, 2.3 and 3 for single-event sensitivity, 90% CL and 95% CL respectively, ϵ is the signal efficiency, and $\mathrm{N}_{\mathrm{muons}}$ is the total number of muons stopped on the target. The denominator is effectively the amount of signal events that the detector will be sensitive to. Superimposed onto the plot is the previous upper limit set by the SINDRUM collaboration [110].

3.5 Dark Photon Analysis

The Mu3e experiments' sensitivity to the dark photon model introduced in Chapter 1 was first estimated by Ref. [130] using the full Mu3e detector and software. The study was first introduced by Ref. [131], and revisited later by Ref. [132] using toy MC. These studies showed that during Phase-1, the Mu3e detector has a competitive sensitivity at low dark photon masses. The study is repeated here using the latest version, which incorporates the latest Mu3e detector design. Differences between the two Mu3e studies includes a decrease in the number of fibre layers in the scintillator from 4 to 3, and generally improved reconstruction algorithms and updating the silicon detector design in Geant4. Asymptotic formulae [133] are used in estimating the sensitivity.

The dark photon is produced during the decay of the muon and is assumed to decay promptly into a positron-electron pair. With the Michel positron, reconstructed tracks from the pair can be inputted into the vertex fitter. Decays following this process would not pass the baseline $\mu \to eee$ selection. Therefore, only a partial event selection scheme is used: only applying the requirements on the vertex χ^2 and vertex to target distance. As the full selection is not being used, a large internal conversion background will be present within the study. The signal will appear as a resonance in the positron-electron invariant mass spectrum at the mass of the dark photon, against a background dominated by internal conversion and combinatorial processes. Only internal conversion is considered for this study. The amount of combinatorial background expected is very small.

3.5.1 Simulation and Reconstruction

Dark photons with masses ranging from 10 to 70 MeV are generated with MadGraph5@NLO [128], and inputted into the Phase-1 Mu3e simulation framework through the Les Houches event [134] file format. For each mass, 80,000 frames are generated, with one dark photon producing muon decay in each frame. For the internal conversion background, 14 million frames are generated to NLO with no generator cuts applied. The dark photons are assumed to decay promptly within the Mu3e detector, thus allowing for the standard vertex fitter to be used. The standard track algorithms are used, and a semi-standard event selection is applied, which will be described in the following.

The selection criteria for both dark photon and internal conversion samples use the same truth cuts as explained above for the $\mu \rightarrow eee$ analysis. Regarding reconstructed variables however, only two cuts are taken into account to prevent the removal of the signal. The two cuts are a cut on the vertex spatial χ^2, and the distance of the reconstructed vertex to the target. The signal (masses between 10 and 70 MeV) and background (internal conversion) efficiencies obtained for both old and new versions of the Mu3e software are plotted in Figure 3.7. Distributions of the invariant masses of the positron-electron pair are plotted in Figure 3.8. It is impossible to know which positron originated from the dark photon, and so two track pairings are considered for each event. The higher of the invariant masses is denoted M_{ee}^{high}, whilst the lower is denoted M_{ee}^{low}. Also considered is the case where the distribution of both of these masses is plotted, this is denoted M_{ee}^{both}. An example cutflow is shown in Table 3.5, where the effects of the full selection can be seen.

Selection	$\mathrm{N_{Vertices}}$	Σw	ϵ [%]	$\sqrt{\Sigma \mathrm{w}^2}$/norm [%]	ϵ_{step} [%]
Initial reconstructed vertices	5790	5790	-	-	-
Target region	5581	5581	22.01	0.29	96.39
Track hit number < 3	3044	3044	12.00	0.22	54.54
Track same mother ID	2925	2925	11.53	0.21	96.09
Track 'prime'	1955	1955	7.71	0.25	66.84
Track type	1955	1955	7.71	0.17	100.00
Track unique	1955	1955	7.71	0.17	100.00
Vertex Multiplicity	1920	1920	7.57	0.17	98.21
Vertex $\chi^2_{\mathrm{spatial}} < 30$	1913	1913	7.54	0.17	99.64
Vertex distance from target < 3	1851	1851	7.30	0.17	96.76

Table 3.5: Dark photon cutflow for a dark photon of mass 70 MeV, vertices with any type of track are considered. Vertices are weighted ('w') as 1 in this sample.

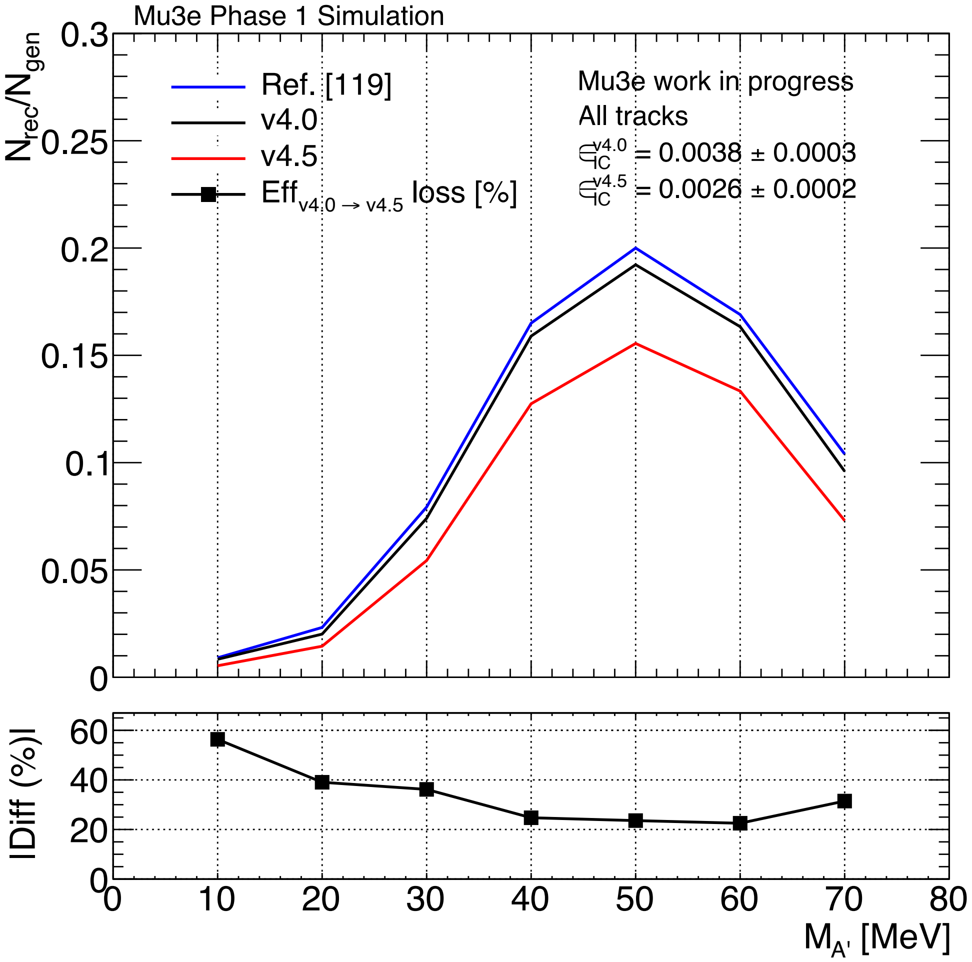

Figure 3.7: Dark photon (the solid libes) and internal conversion background efficiencies (text overlay on plot) are plotted. They are compared with the results from the previous study. Vertices with any type of track are considered. The previous dark photon study was undertaken with v4.0 of the Mu3e software. This study is conducted using the most recent version, v4.5.

3.5.2 Sensitivity

The distributions shown in Figure 3.8 are entered into a likelihood fit to estimate the upper limit. The test statistic used in the analysis is a one sided profile likelihood. To construct it, asymptotic formulae from Ref. [133] are used. The formulae estimates the median sensitivity to a particular confidence level through a profile likelihood ratio, estimating the upper limit via a one sided test [129]. The background distribution is smoothed taking the median of

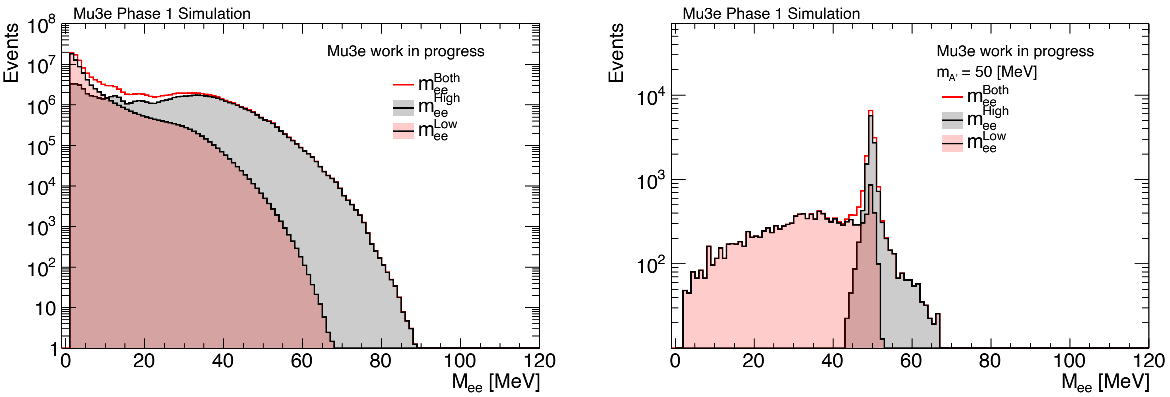

(a) Internal Conversion background. (b) Dark photon of mass 50 MeV.

Figure 3.8: Sample distributions before fitting, normalised to 300 days of data taking with $2{\times}10^8$ muons/s decaying on target.

three bins. Shown in the figure is the internal conversion background and an example signal resonance of a dark photon of mass 50 MeV. The final sensitivity is shown in Figure 3.9. At low dark photon masses it can be seen that the Mu3e detector is sensitive to regions of unexplored phase space. The region around 17 MeV is of particular interest, as Mu3e is sensitive to a region that could explain the Atomki anomaly [135]. It can be observed that a higher sensitivity could be obtained at low masses through the use of the M_{ee}^{Low} variable. The most stringent limits present in this area are due to beam dumps and experiments such as NA48 and BaBar [136, 137].

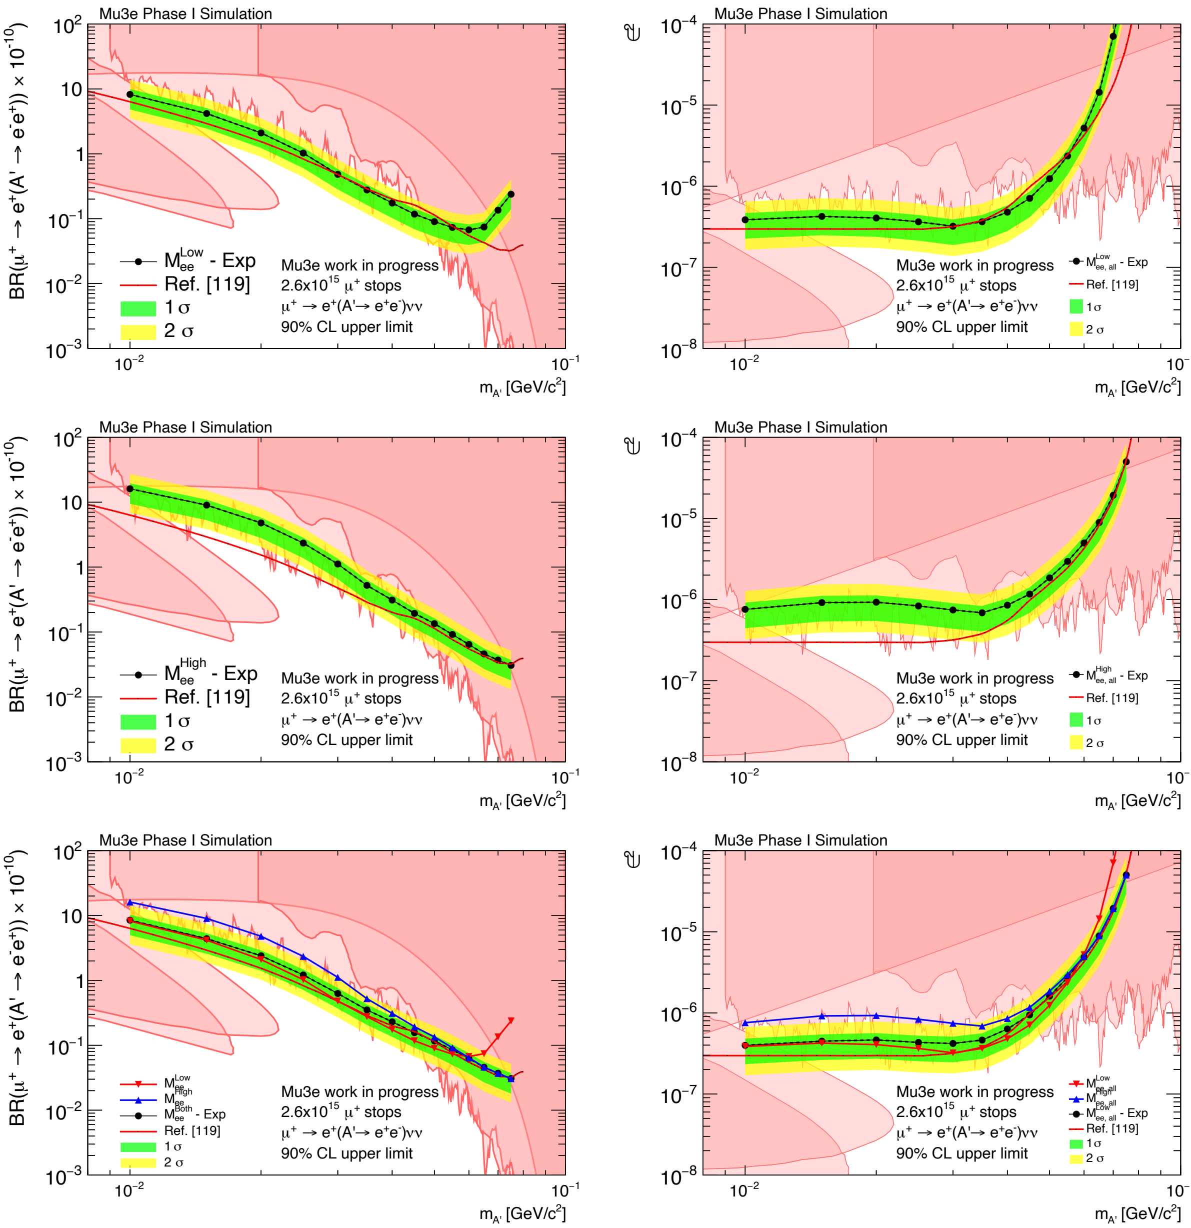

Figure 3.9: Dark photon sensitivity to the branching ratio of the dark photon decay, and to the kinetic mixing parameter as a function of mass. The shaded regions show areas where previous experiments have placed upper limits on the dark photon to 90% CL. The effect of choosing different track pairings is also shown.

3.6 Conclusion

This chapter first provided a description of the Phase-1 Mu3e simulation and analysis framework. The track and vertex requirements that are used in all Phase-1 analyses was described, and was utilised in studies relating to the search for $\mu \rightarrow eee$ and a search for dark photons with Mu3e. The results from the former showed agreement with the published results that can be found in the Mu3e technical design report. The dark photon study concluded that with the latest version of the simulation framework, incorporating the full Phase-1 detector design, Mu3e is sensitive at low masses to unexplored dark photon phase space.

Chapter 4

Two-layer tracking and integration run simulation

For the preparation of data taking for Phase-1 Mu3e, a test beam was undertaken at PSI with the purpose of demonstrating a fully functioning data acquisition system that could simultaneously record hits from the various subdetectors. To this aim, a prototype vertex detector was constructed. Whilst the test beam was ongoing, simulation studies were conducted in parallel, exploring the possibility of using only two layers to construct tracks with Mu3e.

In this chapter, the bespoke simulation of the prototype vertex detector is presented, alongside descriptions of the tracking algorithms used to reconstruct tracks using only two of the silicon layers. Studies into the fake rate and reconstruction efficiency estimated as a function of the track momentum are shown. The momentum resolution of the new tracks is also estimated.

4.1 The Prototype Vertex Detector

The vertex detector, first discussed in Chapter 2, consists of two layers at radii of 23.3 m, and 29.8 mm of MuPix 11 chips that fully surround the target. At the time of the test beam, only MuPix 10 chips were available for use. For the construction of a prototype vertex detector, substantial changes in the design of the detector were required in the running of the chip. Mainly, the use of a thicker printed circuit board (PCB) instead of a high density interconnect (HDI) forced the prototype's design to differ massively from the design presented in Chapter 2. The work in this thesis incorporated this alternate design into the Mu3e software framework.

The building block of the prototype is the ladder, which itself is built from a thick U-shaped PCB. Glued onto the PCB is polyimide foil, upon which 6 MuPix 10 chips are glued. This design was implemented as it allowed for the slow-control of the MuPix 10 chips, which, due to a design flaw, was impossible with the original HDI design. A CAD model of the prototype is shown in Figure 4.1. The thicker PCB's, and the additional services required for the ladder to function, forced the ladders to be placed at radii much larger than the Phase-1 vertex detector radii specifications outlined in Chapter 2. This leads to the emergence of gaps between the active silicon sensors, which is detrimental to the detector acceptance. Table 4.1 outlines the dimensions of the prototype vertex detector.

layer	1	2
number of modules	-	-
number of ladders	8	10
number of chips per ladder	6	6
length [mm]	124.7	124.7
minimum radius [mm]	31.7	44.6

Table 4.1: Table shows a summary of the parameters of the pixel layers in the prototype vertex detector. MuPix 10 was used instead of MuPix 11.

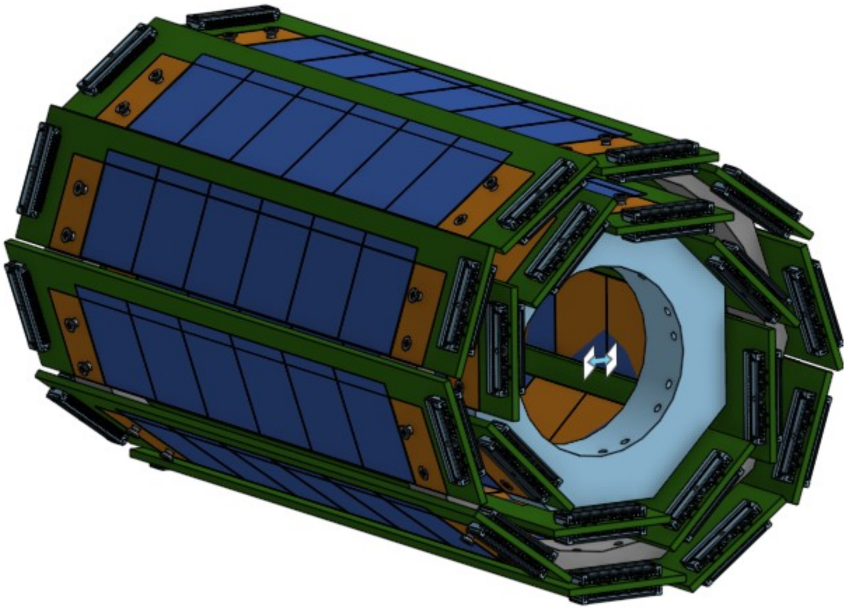

Figure 4.1: CAD model of the prototype vertex detector.

4.2 Simulation

The Geant4 simulation of the prototype vertex detector is built upon the existing Phase-1 Mu3e software framework. The simulation begins with modelling the thick U-shaped PCB, the dimensions of which are known from the CAD model in Figure 4.1, and can be inputted into Geant4. The thickness of the PCB is 1.55 mm, the width is 28.18 mm, and the length is 176.46 mm. The empty space beneath the MuPix chips is 124.7 mm in length and 20.88 mm in width, which is the size of the active area of the chips. The MuPix 10 and MuPix 11 chips have identical dimensions, and so for the purpose of the simulation, they can be treated interchangeably.

A model of one of the ladders that constitute the prototype vertex detector is shown in Figure 4.2. The dimensions of the ladder are inputted to be exactly what is depicted in the CAD model. Once simulated, the ladder itself can be placed at the correct locations as shown in Figure 4.1. The resulting Geant4 model is shown in Figure 4.3.

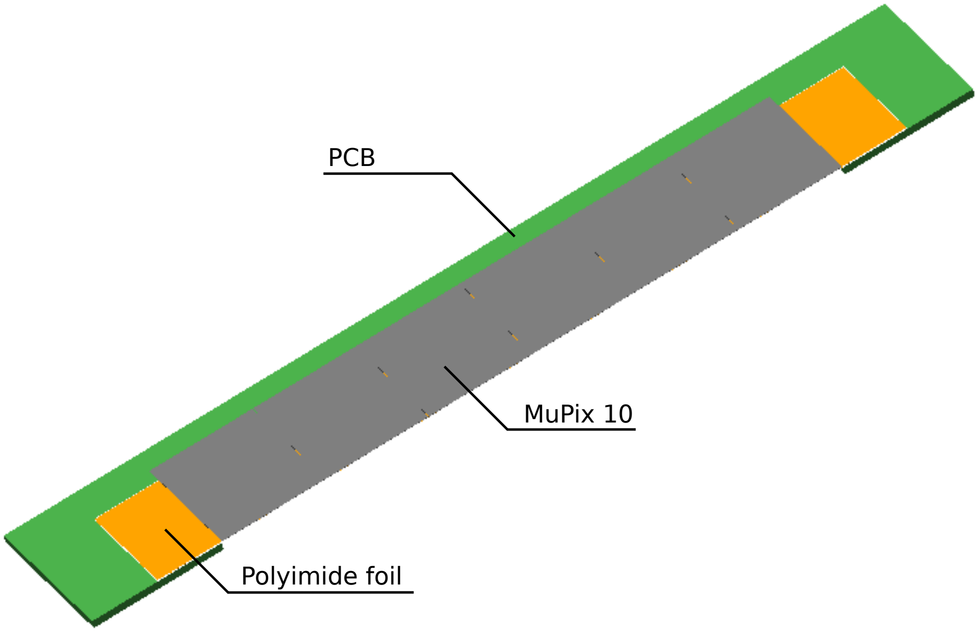

Figure 4.2: Geant4 model of a simulated ladder of the prototype vertex detector.

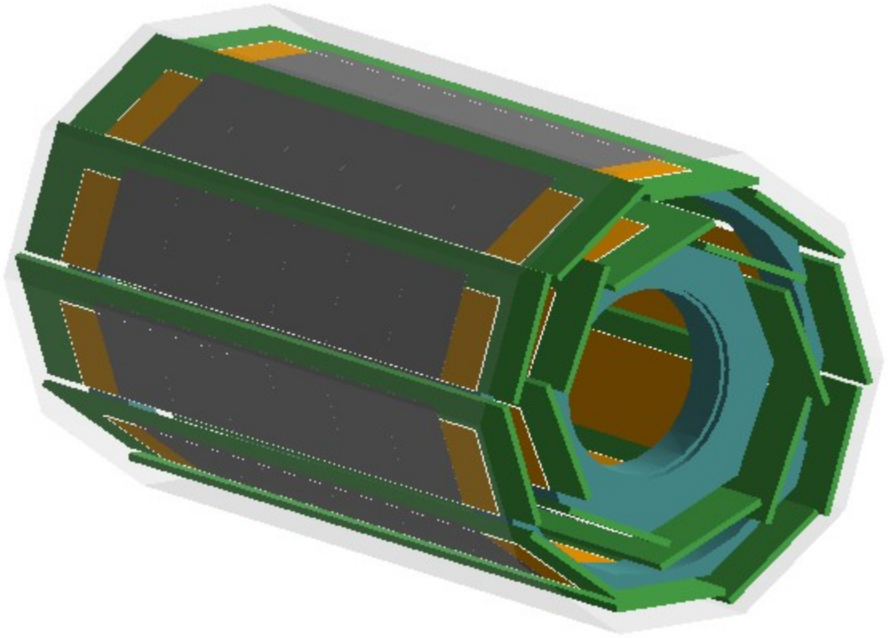

Figure 4.3: Complete Geant4 model of the prototype vertex detector.

4.3 Two-Layer Tracking

The Phase-1 tracking algorithms have previously been discussed in Chapters 2 and 3. In the Phase-1 analysis of Mu3e, the current implementations of the tracking algorithms are only able to reconstruct tracks from particles that pass through all four of the silicon layers. The momentum required for tracks to at least reach the most outer layer is 10 MeV. Particles below this threshold, although generating hits in the silicon layers that they do pass through, remain unseen by the Mu3e detector, and is a loss in the acceptance for Mu3e. The possibility of reconstructing particle tracks using only two layers therefore would only increase the acceptance of the Mu3e detector. To maximise the use of the prototype vertex detector, a study into two-layer tracking is conducted.

Long-4

The basis of all tracks in Mu3e is the triplet. With only two layers, a triplet is constructed from two outgoing hits in layers 1 and 2, and an incoming hit in layer 2. Combined with a further triplet from an outgoing hit in layer 2, and ingoing hits in layers 2 and 1, a track with four hits can be constructed. The ϕ and z windows used to find the hits have to be wide enough to find the hits in the other layers, whilst simultaneously small enough to reject the construction of fake tracks. The large curvature of these 4-hit tracks make them distinguishable from short tracks in Phase-1, and are instead named ‘long-4’ tracks. Due to the special curvature of these tracks, the momentum resolution can be expected to be much better than what is estimated for short tracks in Phase-1. To construct a long-4 track, the specialised requirements are presented in Table 4.2. The variables used in the requirements for both geometries are plotted in Figure 4.4.

Events displays of reconstructed long-4 tracks can be seen in Figure 4.5 and in Figure 4.6. The tracks are shown with respect to the x-y plane and the r-z plane. The former shows the event display of the simulated prototype vertex detector, with the ladders being placed at a higher radius and gaps appearing between the silicon sensors. Despite this, tracks could be

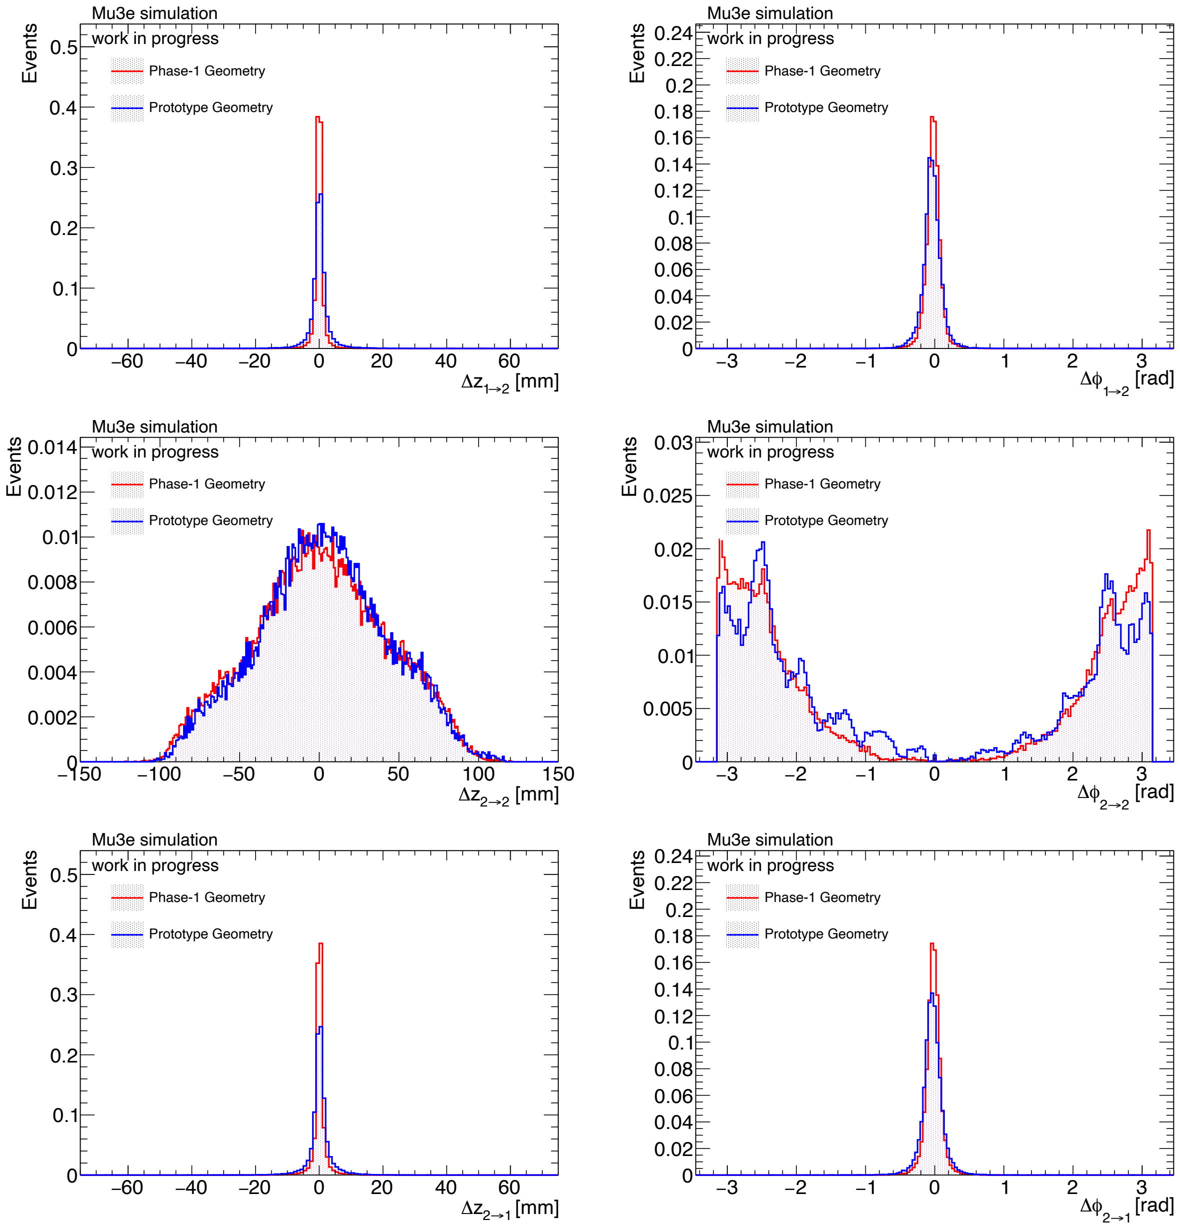

Figure 4.4: Difference in ϕ and z co-ordinates for various silicon hits for long-4 tracks in the Phase-1 geometry of the experiment, compared with the prototype geometry.

constructed with the detector setup. Figure 4.6 presents an event display of a long-4 track reconstructed with the Phase-1 Mu3e geometry. To obtain these event displays, the Mu3e simulation is ran with just the regular muon beam simulated with no special muon decays. The detected particles in these displays are generated from muon decays on the surface of the target.

Requirement	Value
Triplet transverse radius [mm]	> 30
$\phi_{0\to1}$ [rad]	$\Delta\phi < 0.7$
$\phi_{1\to2}$ [rad]	$\Delta\phi < \pi$
$\phi_{2\to3}$ [rad]	$\Delta\phi < 0.7$
$z_{0\to1}$ [mm]	$\Delta z < 15$
$z_{1\to2}$ [mm]	$\Delta z < 150$
$z_{2\to3}$ [mm]	$\Delta z < 15$
χ^2	< 32

Table 4.2: Table of variable windows used to reconstruct long-4 tracks.

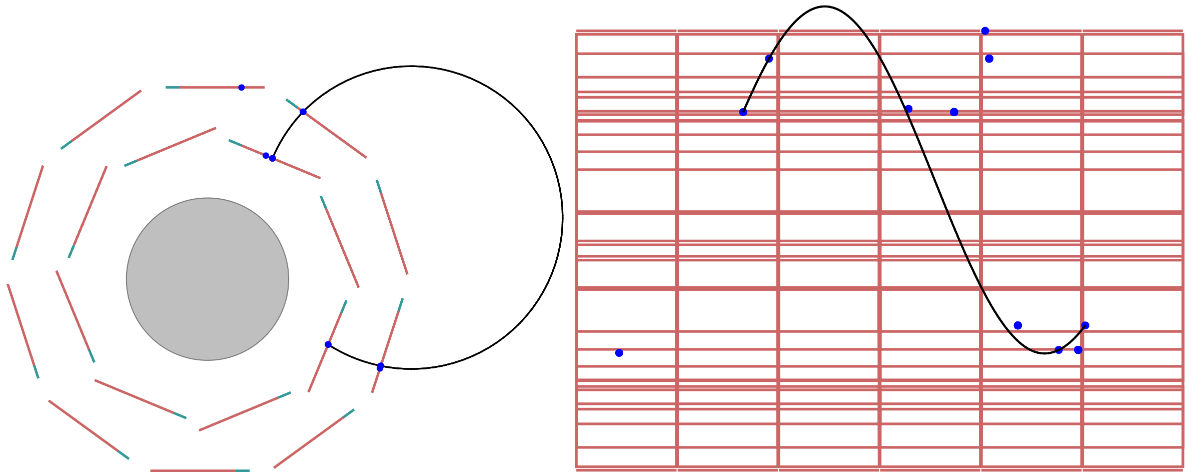

Figure 4.5: Longitudinal (Right: r-z axes) and transverse (Left: x-y axes) view of the event display of a reconstructed long-4 track using the simulated prototype vertex detector. The ladders are placed at a higher radius, and gaps between the active silicon sensors can be seen.

Efficiency

The efficiency of the track reconstruction is estimated, using the requirements in Table 4.2, as a function of the generator-level momenta of the reconstructed tracks. This efficiency is estimated with the prototype vertex detector. The efficiency is defined as the ratio between reconstructed tracks and tracks formed using generator level information. The tracks must be on their first recurl and register four sequential hits in the silicon layers. With the requirements provided in Table 4.2, a high efficiency of 80% can be observed for the prototype vertex detector.

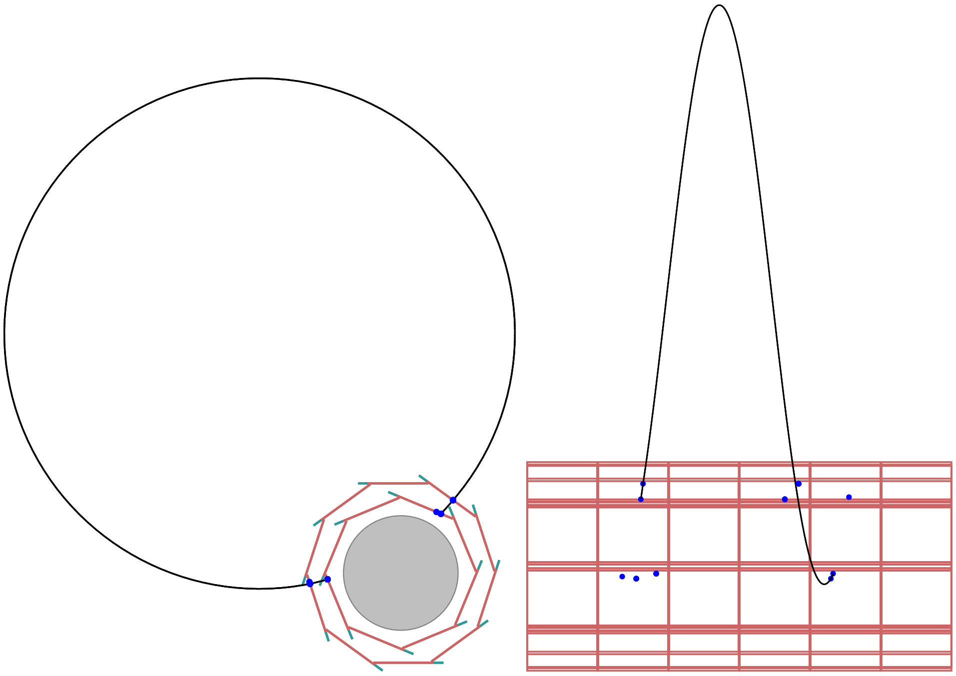

Figure 4.6: Longitudinal (Right: r-z axes) and transverse (Left: x-y axes) view of the event display of a reconstructed long-4 track using the simulated prototype vertex detector. The ladders are placed at a higher radius, and gaps between the active silicon sensors can be seen.

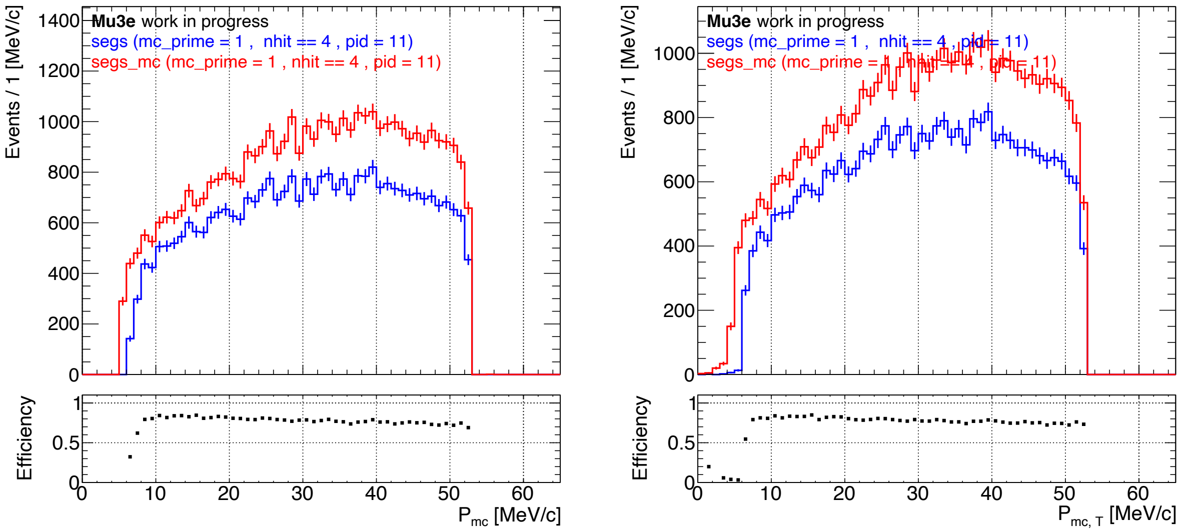

Figure 4.7: Efficiency of long-4 tracks as a function of momentum.

Fake-rate

A study into the fake-rate is performed using the simulation of the prototype vertex detector. It is estimated as a function of the reconstructed track momenta, and is estimated for both geometries. Fake-rate is defined as the ratio between reconstructed tracks that pass the requirements in Table 4.2, and are not matched with a generator-level track, and have no generator level counterpart. A moderately high fake-rate is observed for long-4 tracks, but it reaches a maximum at low transverse momentum as it is difficult to correctly identify the track's direction.

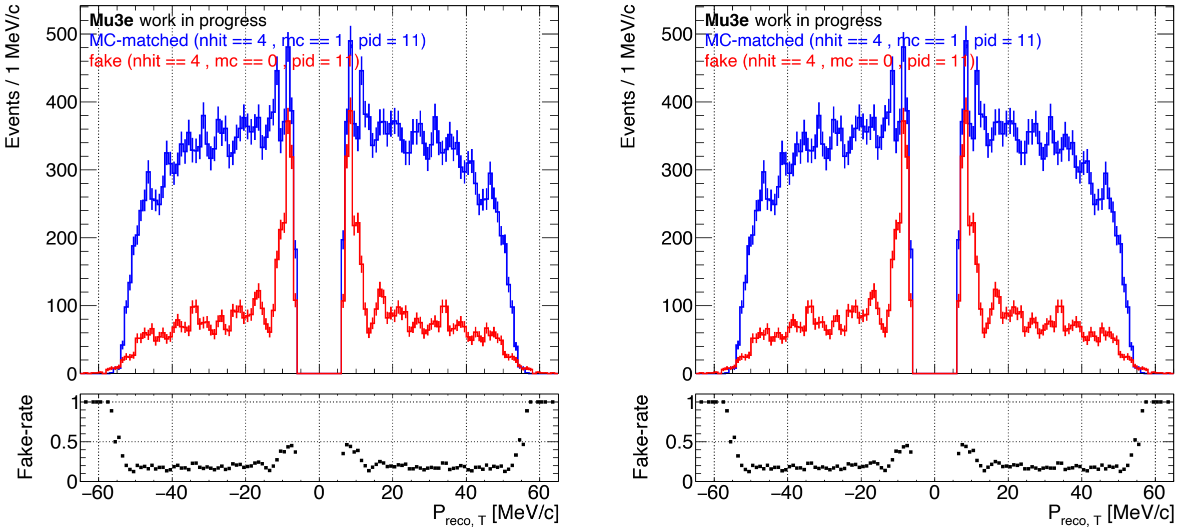

Figure 4.8: Fake-rate of long-4 tracks as a function of transverse momentum.

4.4 Momentum Resolution

The momentum resolution of long-4 tracks is estimated as a function of generator-level momentum. Reconstructed tracks are matched with their monte-carlo generator level counterparts, and the difference between the generator level and reconstructed momentum is taken, and a double sided crystal ball fit is performed on the difference. The sigma on the fit is taken to be the resolution. Example fits are shown in Figure 4.9.

The momentum resolution is plotted as a function of the generator level momentum in Figure 4.10. The momentum resolution of tracks in the Phase-1 analysis is plotted alongside the estimated momentum resolution of the novel long-4 tracks. At low momenta, due to the large curvature of long-4 tracks, the momentum resolution of long-4's are comparable to

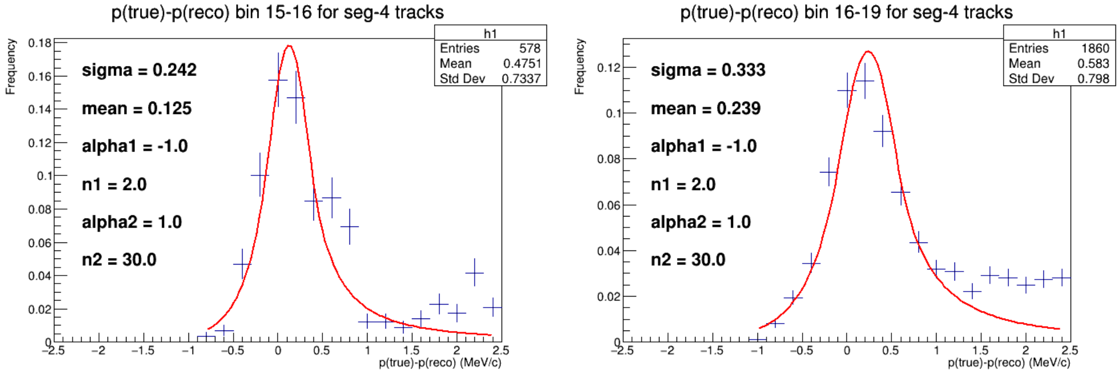

Figure 4.9: Example fit of long-4 track momentum difference.

long-6's and long-8's. Any improvement to the momentum resolution is incredibly useful for improving the physics reach of Mu3e. If long-4 tracks are considered in the final analysis, the acceptance and sensitivity of Mu3e could be improved.

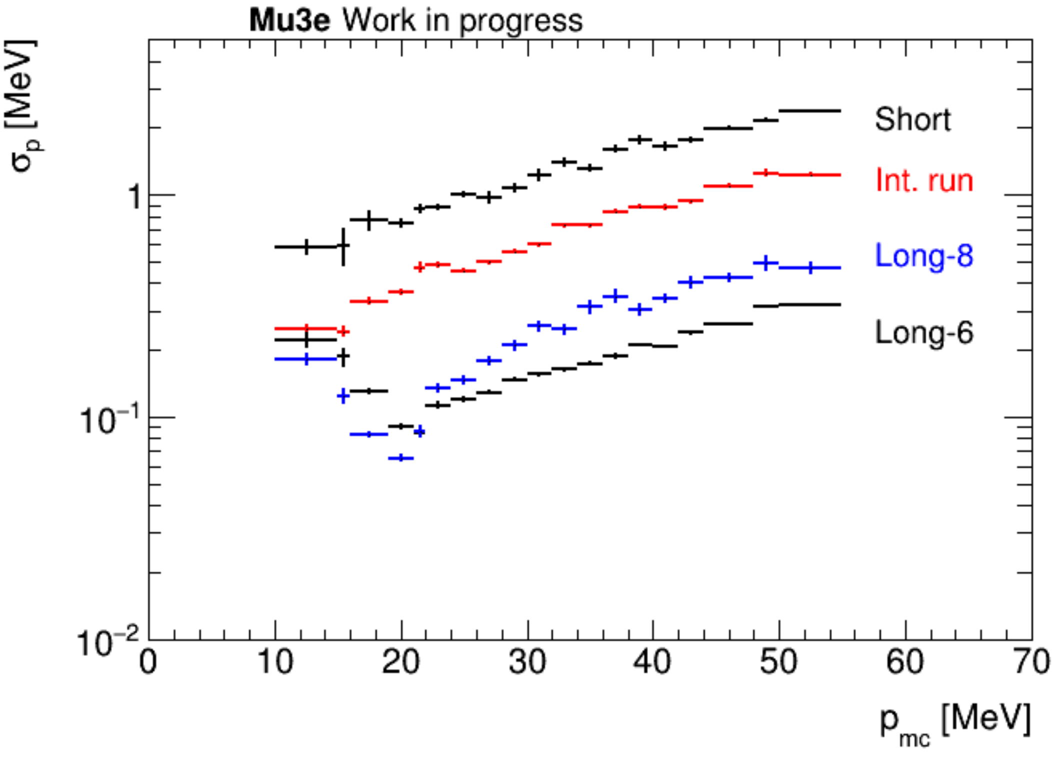

Figure 4.10: Momentum resolution of two layer tracks compared with the regular Phase-1 track momentum resolution.

4.5 Acceptance and Efficiency of Prototype Vertex Detector

The acceptance for long-4 tracks in the prototype vertex detector geometry is estimated with respect to ϕ and λ in Figure 4.11. This acceptance in ϕ shows the loss of efficiency due to the gaps between ladders of the detector, which is a consequence of using the PCB design. Tracks are also only reconstructed at small λ angles, and these tracks in Mu3e are tracks that are difficult to correctly identify the direction of rotation without the use of timing information. The efficiency is plotted in Figure 4.12, where the moderately high efficiency can be seen for lower momentum tracks.

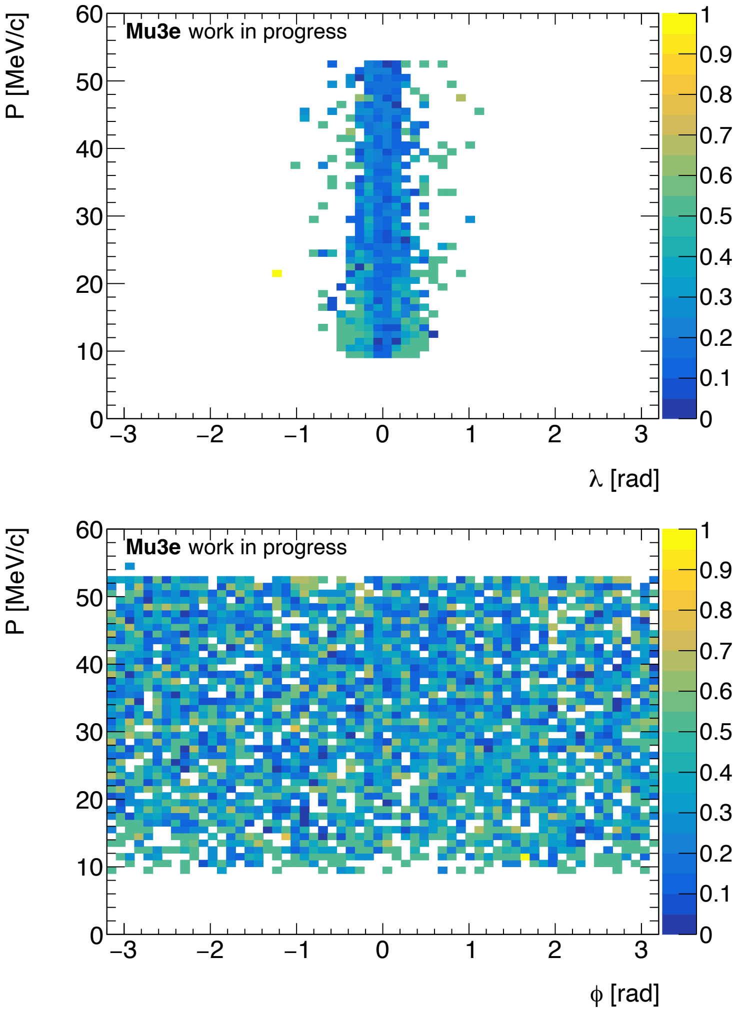

Figure 4.11: Acceptance of long-4 tracks as a function of ϕ and inclination angle λ. The acceptance is defined as the ratio of reconstructed tracks to monte-carlo truth tracks (z axis).

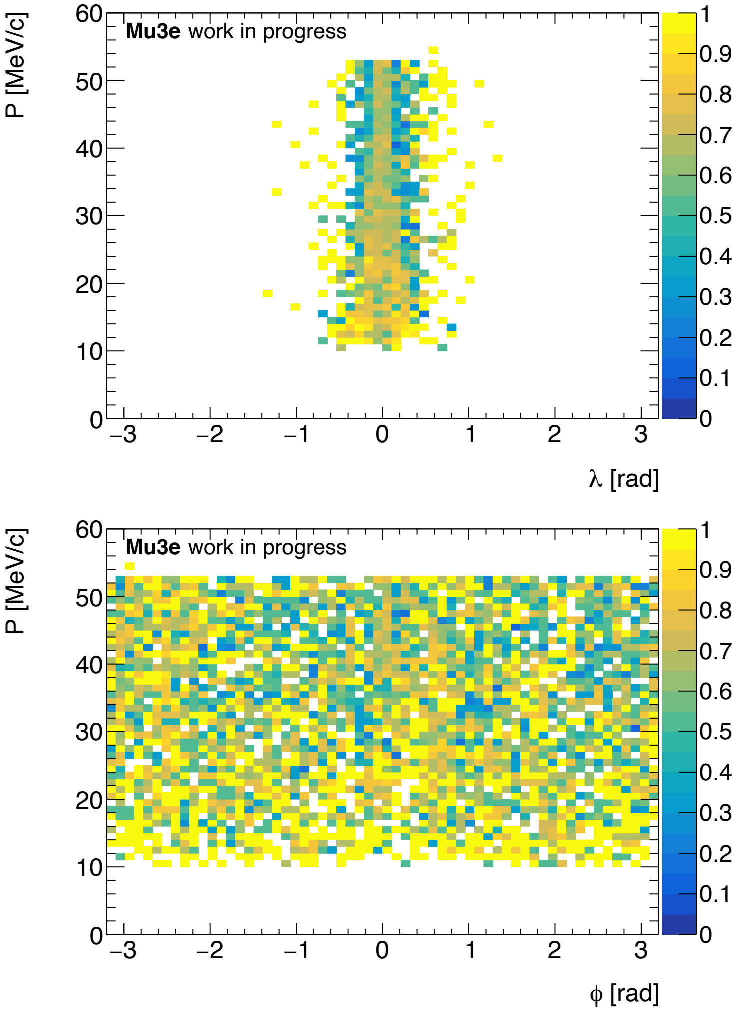

Figure 4.12: Efficiency of long-4 tracks as a function of inclination angle λ

4.6 Conclusion

This chapter presented a Geant4 simulation of the prototype vertex detector. Long-4 tracks were defined, extending existing nomenclature, and the track efficiency and fake-rate was estimated for the prototype vertex detector. The momentum resolution was estimated for both the nominal Phase-1 geometry and the prototype geometry. Although the reconstruction efficiency was found to be high, the fake-rate was also relatively high and would require further study if the tracks were to be used in the Phase-1 analysis to improve the performance of the Phase-1 detector.

Chapter 5

Phase-2 scenario: Extra Outer Layer

The High Intensity Muon Beam (HIMB) upgrade at PSI will require substantial updates to the Mu3e detector. In order to make use of the high muon rate ($>10^9$ muons/s) several of the Phase-1 detector components, such as the fibre detector, will either need to be removed or upgraded. New detector designs and technologies need to be tested and optimised to help suppress the now 400 times increase to the accidental background that the improved muon beam will provide. Several detector designs provide Mu3e with a means of testing more exotic phenomena with regards to muon decays, and will need to be optimised in the simulation.

An example of such a design is presented within this chapter. To improve the momentum resolution of tracks of high momentum, an extra outer layer is added to the Mu3e detector at high radius. The tracks will register hits at locations that are crucial for improving the momentum resolution. It will be shown that several of the current Phase-1 tracking algorithms will require changes in order to take advantage of the hits in this fifth layer, which will in turn lead to an improvement to the momentum resolution for tracks with a transverse momentum greater than 34 MeV.

5.1 Motivation

The High Intensity Muon Beam upgrade will force the πE5 beamline to shutdown for several years. During this shutdown, the design of the Mu3e detector will be altered to make full use of the upgraded beams' properties. Physics cases for a different implementation of the detector concept must be studied. One such concept has been simulated within this thesis. It's main feature is the addition of an extra outer layer at a high radius of 230 mm. As discussed in Chapter 2, the momentum resolution of a track does not necessarily improve with increasing numbers of hits. To first order, uncertainties in the curvature of a track cancel if two detector hits are located at a curvature of roughly π from one another in the transverse plane. This is demonstrated by the minimum in the momentum resolution for long-6 and long-8 tracks shown again in Figure 5.1 as a function of momentum and transverse momentum. Short tracks have a linear relationship between momentum and momentum resolution due to the lack of a recurling section of the track. Long-6 and long-8 tracks are seen to reach a minimum resolution of less than 100 keV at a momentum of roughly 20 MeV.

To obtain Figure 5.1, the Phase-1 detector is simulated using the existing Mu3e software framework. 100,000 frames are generated, 64 ns in length, with one $\mu \rightarrow eee$ decay in each frame. Tracks are reconstructed using the standard Phase-1 tracking algorithms, and each track (short, long-6 and long-8) has the difference of the truth momentum and reconstructed momentum plotted, and is fitted with a double sided crystal ball function. The standard deviation of the fit is taken as the momentum resolution. Examples of the fits are shown in Figure 5.2.

The best momentum resolution is observed for tracks of roughly 20 MeV. For tracks of larger momenta the resolution becomes worse, all the way up until the Michel edge. Although satisfactory for Phase-1 requirements, for the second phase of the experiment, momentum resolution will need to be improved upon to reduce the amount of accidental background. The solution discussed here would involve placing a silicon detector at a radius of 230 mm, which would provide the tracks of high momentum with a detector hit located at a curvature

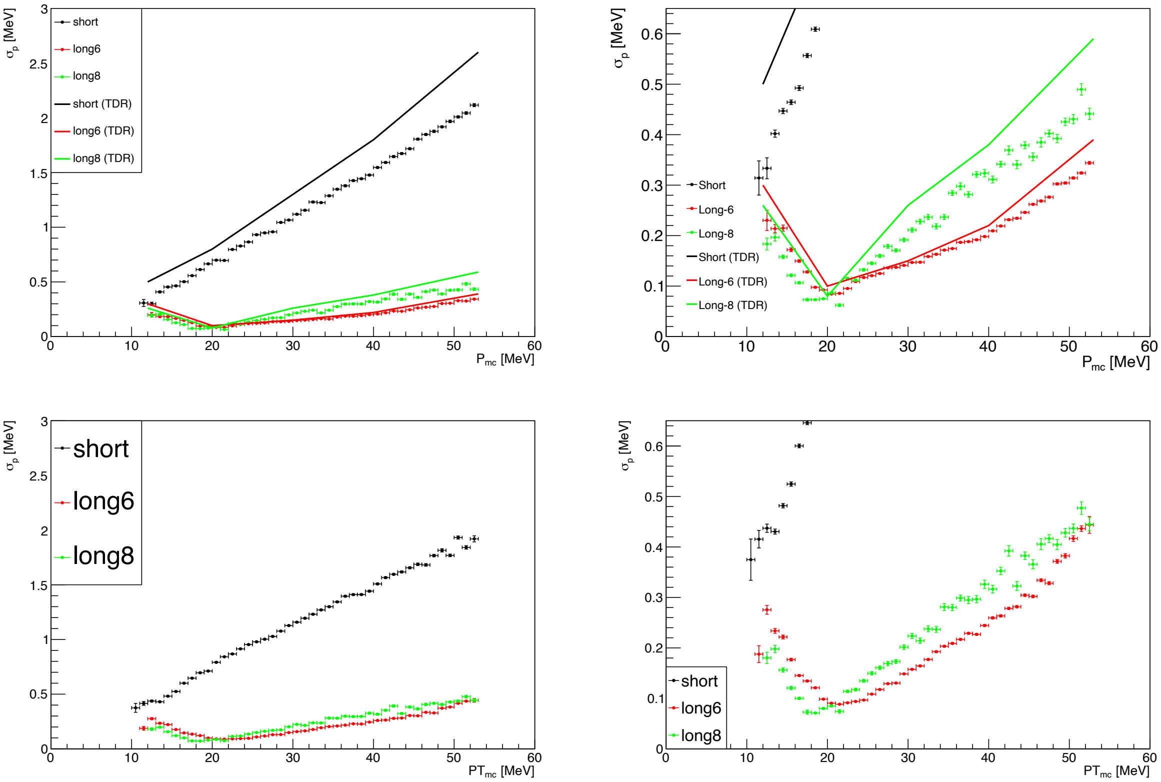

Figure 5.1: Plot of the Phase-1 momentum resolutions for Phase-1 tracks (short, long-6 and long-8) as a function of momentum and transverse momentum. Scales are changed for each track to show the range of resolutions for particular momentum bins.

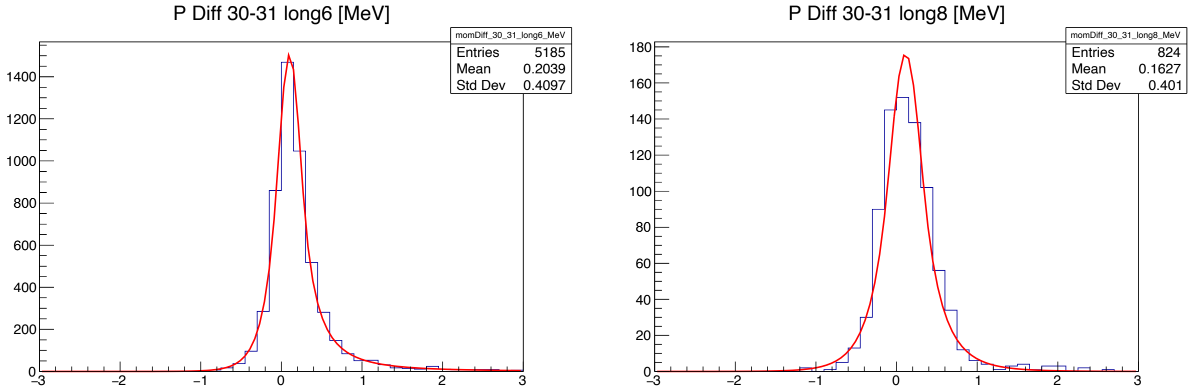

Figure 5.2: Example fits of the double-sided crystal ball function on momentum differences for long-6 and long-8 tracks.

of π, which would improve the overall performance of the detector. This extra silicon layer would be based primarily on the design of the existing outer layers: a basic ladder into module design, which would aim to be cylindrical, surrounding primarily the central station.

The simulation design will be based on the existing Phase-1 architecture, i.e, the Geant4 framework is utilised. The study is a demonstration of the flexibility of the simulation framework to add and remove sections of the detector, which will be used to simulate other detector configurations.

5.2 Current Outer Layer Pixel design

The Phase-1 outer layer pixel design has been previously introduced in Chapter 2. The silicon layers 3 and 4 are known as the outer pixel layers. A summary of their properties is shown in Table 5.1, and the Geant4 implementation is presented in Figure 5.3.

Layer	3	4
Number of modules	6	7
Number of ladders	24	28
Number of MuPix sensors per ladder	17	18
Active length [mm]	351.9	372.6
Minimum radius [mm]	73.9	86.3

Table 5.1: Summary of the properties of pixel layers 3 and 4.

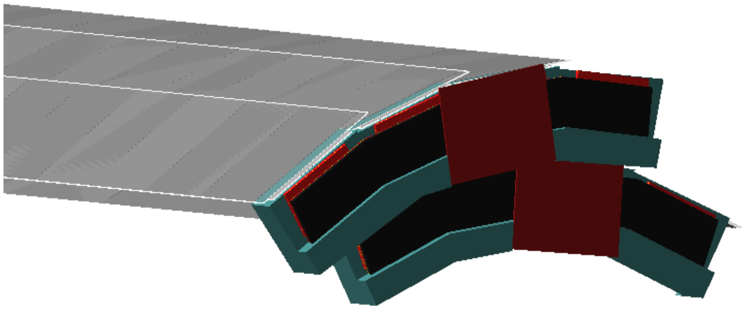

Figure 5.3: Geant4 model of the existing layer 3 and layer 4 outer layer modules.

The CAD design of the outer pixel layer is accurately modelled, using Geant4, within the Mu3e software framework. The ladders themselves, (silicon MuPix chips ontop of a thin support structure) is modelled alongside the PEI endpieces that hold the modules in place, alongside the v-folds which are implemented for mechanical stability.

5.3 Proposed Extra Outer Layer design

The proposed design is based on the sketch of the extra outer layer as shown previously in Figure 2.19. The full azimuthal angle is covered by the modules, which are placed at a minimum radius of 230 mm. This design could be implemented using the existing Phase-1 design philosophy. Ladders with the same numbers of chips and ladder length as the layer-4 module ladders could be utilised. With the current width of the ladders, to complete a circumference of a circle with radius 230 mm, it would require, assuming a staggered design like the current Phase-1 layers, 70 ladders, as shown by the following:

$$\text{circumference} = \pi d = 460\pi[\text{mm}], \tag{5.1}$$

$$\text{N}_{Ladders} = \frac{\text{circumference}}{\text{ladder width}} = \frac{460\pi}{20.66} \approx 70. \tag{5.2}$$

The summary of such a design, in comparison with the existing outer layer design, is shown in Table 5.2. Assumed is that the tilt of the ladders is 1°, similar to what is currently implemented for the existing outer pixel layers. MuPix 11 chips are also implemented into the simulation, and are used on the ladders, as well as the kapton-aluminium HDI support that is also currently used in Phase-1. The number of silicon pixel chips used in the ladder is limited by the mechanical stability of the ladder. In terms of the extra outer layer, as the existing from Phase-1 is being used, only 18 chips are used, limiting the ladder length to the Phase-1 ladder length. For simplicity, the simulated extra outer layer will not have a PEI endpiece. This is not necessary for the studies that will take place, but can easily be added at a later date.

5.4 Modelling the Extra Outer Layer with Geant4

The simulation of the detector design is implemented as described in the previous section. The model in Geant4 is shown in Figure 5.4, where the full Phase-1 detector can be seen

Layer	3	4	**5**
Number of modules	6	7	17.5
Number of ladders	24	28	70
Number of MuPix sensors per ladder	17	18	18
Active length [mm]	351.9	372.6	372.6
Minimum radius [mm]	73.9	86.3	230

Table 5.2: Summary of the properties of pixel layers 3 and 4, and a potential layer 5 as implemented in the simulation for studies presented in this chapter.

alongside the Phase-2 extra outer layer.

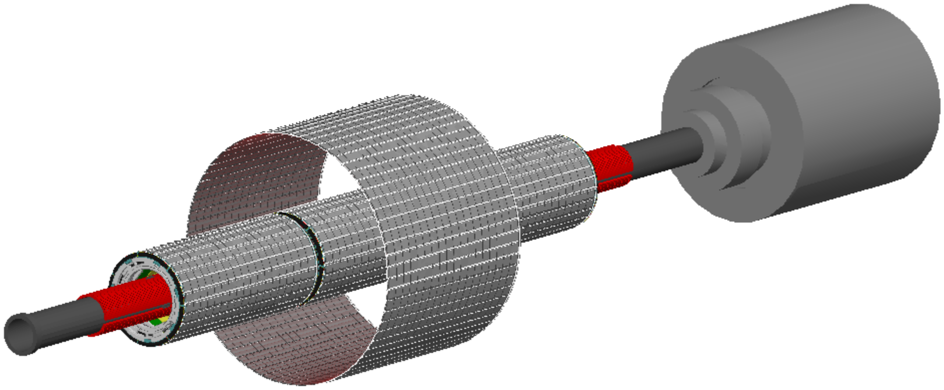

Figure 5.4: Geant4 model of the Phase-2 scenario of an additional outer layer.

The model depicted is partially complete, but not too dissimilar to what is used in the current Phase-1 simulation. Shortfalls of the simulation include the lack of PEI endpiece support structure, or any simulation of the services that would be required to run the detector. However, for the purposes of understanding what possible gains a pixel detector at this radius would provide Mu3e, the partially complete simulation is sufficient. Further improvements would involve adding the missing non-sensitive volumes to the simulation.

5.5 Conducting a Simulation of the Extra Outer Layer scenario

As described briefly in Chapter 2, hit information is transported from the subdetectors to the filterfarm. This hit information is simulated within the software framework. The information is contained within a 32 bit number, where the first 16 bits represent the location of the hit in the detector, and the last 16 representing the column/hit location on the silicon chip where the a hit from a particle is registered. This is shown pictorially in Figure 5.5.

15	14-12	11-10	9-5	4-0
NULL	Station	Layer	Ladder	Chip

Figure 5.5: Bit scheme for Mu3e Phase-1.

The hit information contains within it information that describes in which part of the detector the hit has been registered. Bit 15 is null, bits 14-12 represent the station (recurl or central), 11-10 represent the layer, 9-5 represent the ladder number and the remainder represent the chip number. This scheme is sufficient for the requirements imposed by the Phase-1 geometry of Mu3e. However, with the addition of an extra layer the 2-bit ladder scheme currently implemented will not work for a layer with number ’5’. Furthermore, the Phase-1 scheme can only represent a layer with up to 32 ladders, and the layer 5 that is proposed will have up to 70 ladders. Therefore a new scheme is required. This scheme is shown below in Figure 5.6:

15-13	12-10	9-5	4-0
Station	Layer	Ladder	Chip

Figure 5.6: Alternative bit scheme for Mu3e Phase-2 studies.

This new scheme allows hits from full ‘Phase-2’ detector to be described within the hit information. After a simulation is run, the hits can be seamlessly inputted into the triplet fit for analysis, which could then be used in a vertex fit if required.

5.5.1 Detector Performance

The study of the detector’s performance is conducted utilising the special muon decay mode provided by the Mu3e framework. The simulation is set to force a muon to decay into three electrons per frame, and 100,000 frames are generated. This would provide enough tracks to conduct analysis on the performance of a fifth layer. Figure 5.7 shows example tracks constructed using Phase-1 algorithms that have enough transverse momentum to pass through the fifth layer. The momentum required is given by:

$$\mathrm{p} = 0.3 \times B \times r, \tag{5.3}$$

where p is momentum in GeV, B is the strength of the magnetic field in T, and r is the radius in metres. From this, it can be understood that the electron decay products require a transverse momentum of roughly 34 MeV to reach the fifth layer.

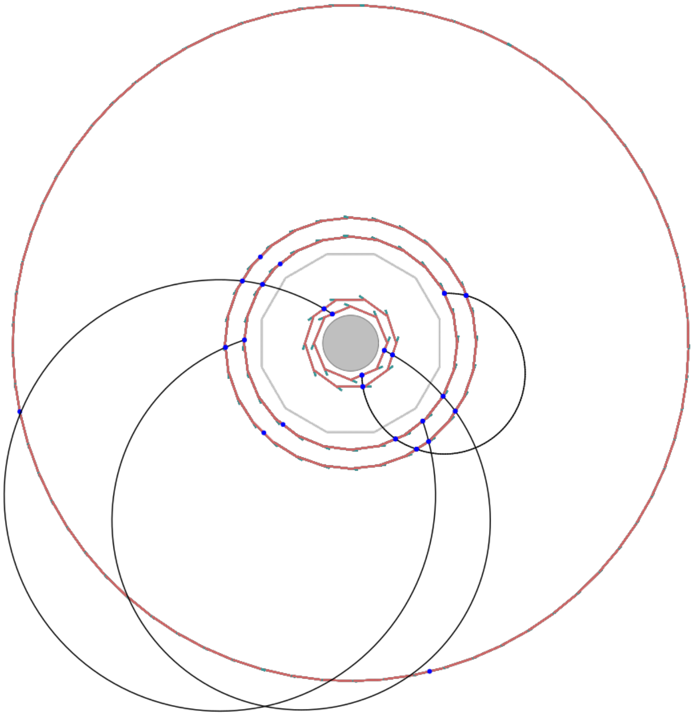

Figure 5.7: Transverse plane of the Mu3e detector with long-6 tracks shown. Tracks pass through the fifth layer, and are detected. Hits on the silicon layer are shown in blue.

5.5.2 Extending Existing Track Algorithms

As can be seen from Figure 5.7, tracks reach the fifth layer but no algorithm exists to make use of the extra hits. To take advantage of these extra hits, and thus the improvement to the momentum resolution for tracks, it is necessary to distinguish between long tracks that do and do not reach this extra outer layer, and to implement new algorithms that construct tracks utilising the extra hits offered by the fifth layer. As an extension to existing track nomenclature, a long track that reaches the extra outer layer is defined as an 'extra-long' track with 'N' hits. (Abbreviated as EL-N). Therefore, a track with 5 hits that reaches this extra outer layer will henceforth referred to as an EL-5 track.

The existing algorithms implemented for the Phase-1 geometry are briefly described in Chapter 2, and the requirements for a track to be constructed are described in Chapter 3. It involves a fast triplet fit which can easily be extended to reconstruct a track of an arbitrary number of triplets. Therefore, with the addition of this extra layer, it is trivial to add conditions that would allow the software to search and include a hit from a fifth layer. In this thesis, three novel tracks are studied and presented: the EL-5, El-8 and EL-10 tracks. EL-7 tracks are ignored as these tracks would hit the services that are required to power the layer, the model of which is ignored within the simulation.

EL-5

When constructing the EL-5 track, the input to the algorithm is the standard Phase-1 short track, which is a track consisting of 4 hits that crosses from the inner-most vertex layer (layer 1) to the outer-most pixel layer of the central station (layer 4). The newly-implemented reconstruction propagates the reconstructed track to the next layer (the fifth layer) in search of a hit. A ϕ window is opened, and if a hit is found within this window, the hit is added to a triplet consisting of itself and the previous two layers, and the triplet fit begins. A window in the z-coordinate is also opened, and if the hit is outside of the window it is removed. Once the fit is conducted, a cut on the χ^2 of the track fit is applied. For short tracks this cut is 32, but for an EL-5, due to the extra hit, this cut is defined to be 40. An example of such a track's reconstruction is shown by the event display in Figure 5.8.

The ϕ and z windows that are chosen require analysis to reduce the numbers of fake tracks. However, a study into the fake rate is not conducted within the scope of this thesis, and

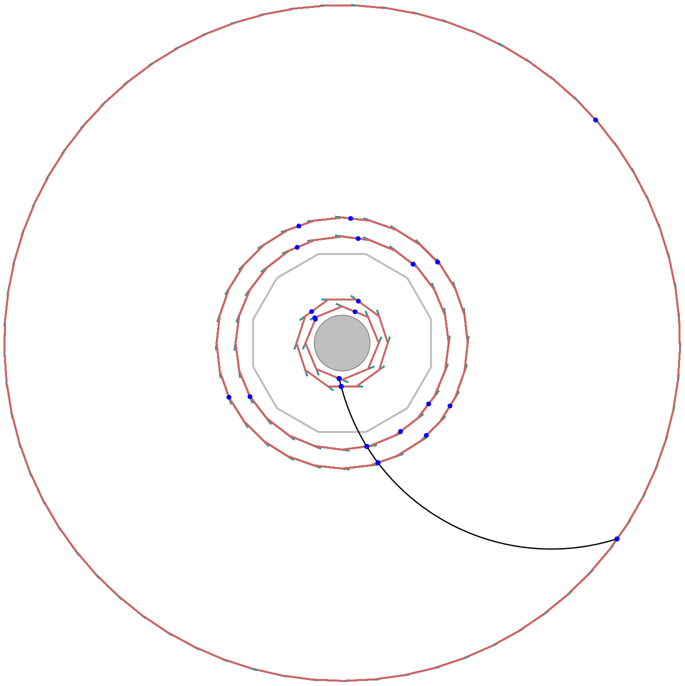

Figure 5.8: Transverse event display of a reconstructed EL-5 track that is successfully matched with a truth track.

the windows chosen are based solely on simulation, without considering the effect of any combination of hits that would give fake tracks. The windows are estimated by examining the locations of truth hits from the simulation, a plot of the ϕ and z-coordinate differences between the fourth and fifth hits of a track are shown in Figure 5.9, where in the figure it is examined that a ϕ window of ± 1.25 rad would be satisfactory. A table of the windows selected for an EL-5 track is shown in Table 5.3. These cuts are also used in the beginning steps of constructing EL-8 or EL-10 tracks.

Requirement	Value
$\phi_{4\rightarrow5}$ [rad]	$\phi < 1.25$
$z_{4\rightarrow5}$ [mm]	$z < 180$
χ^2	< 40

Table 5.3: Table of variable windows used to reconstruct EL-5 tracks.

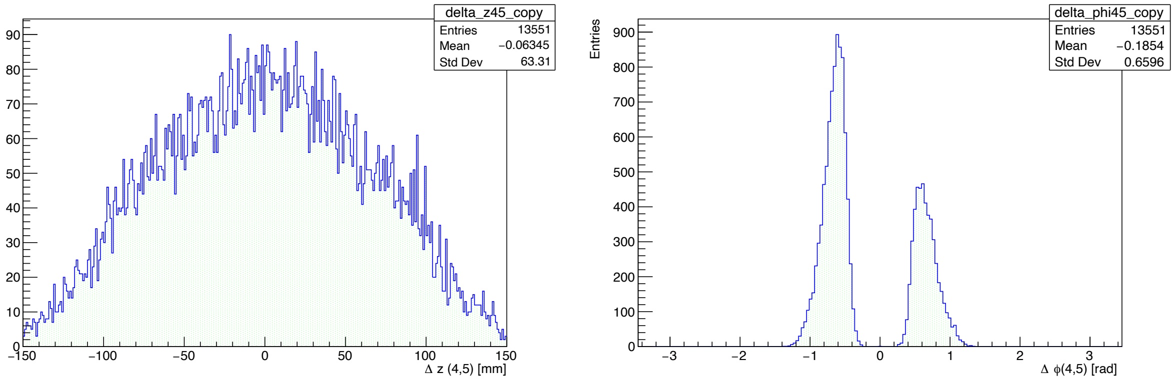

Figure 5.9: Figure shows the ϕ and z differences between a fourth and a fifth hit for tracks that reach the fifth layer.

EL-8

Constructing an EL-8 track begins with a short track as an input. The steps to reconstruct the EL-5 are repeated, and once the fit for an EL-5 track is satisfactory, the reconstructed track is propagated back towards the fifth layer searching for an incoming sixth hit. This is similar to how long-6 tracks are constructed with the standard reconstruction algorithms for Phase-1. A window in ϕ is opened, and if a hit is found a triplet fit is attempted, assuming the hit also passes the z window cut. If the fit is successful, and the χ^2 cut is passed (defined as 48, identical to long-6 and long-8 tracks), the track is propagated back towards layer 4 in search of another hit. Another ϕ and z window is opened, and, if successful, the hit again is used to form a triplet and the triplet fit is repeated. This method continues again with the third layer. Shown in Figure 5.10 are EL-8 tracks constructed with the previous algorithm. The tracks themselves are correctly matched with their truth counterparts, and can be used in the analysis to estimate their momentum resolution. The various ϕ and z window cuts are shown in Table 5.4, whilst their origins can be determined from the truth distributions from Figure 5.11.

Requirement	Value
$\phi_{5\to6}$ [rad]	$\phi < 2$
$z_{5\to6}$ [mm]	$z < 180$
$\phi_{6\to7}$ [rad]	$\phi < 1.25$
$z_{6\to7}$ [mm]	$z < 180$
$\phi_{7\to8}$ [rad]	$\phi < 0.35$
$z_{7\to8}$ [mm]	$z < 50$
χ^2	< 48

Table 5.4: Table of variable windows used to reconstruct EL-8 tracks.

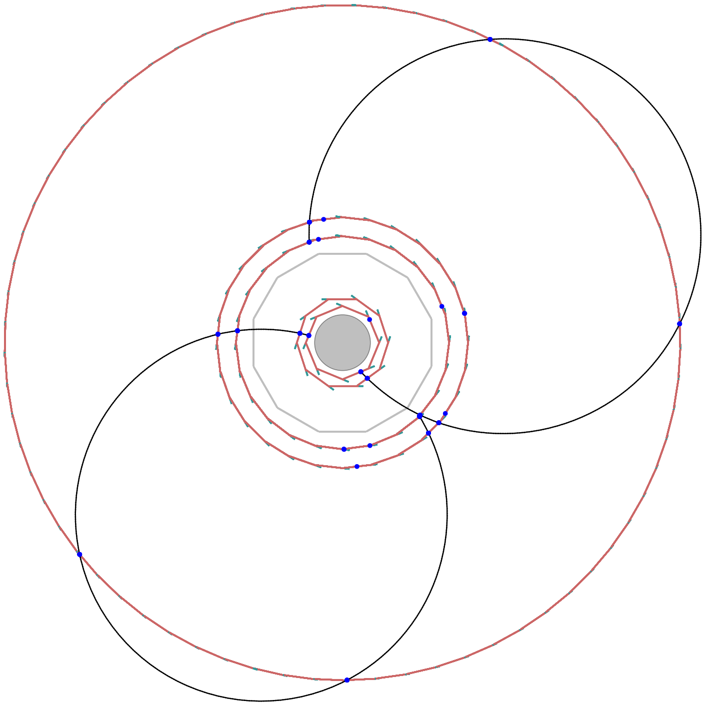

Figure 5.10: Transverse view of a frame in the Mu3e simulation of two successfully reconstructed EL-8 tracks that are correctly matched with a truth track.

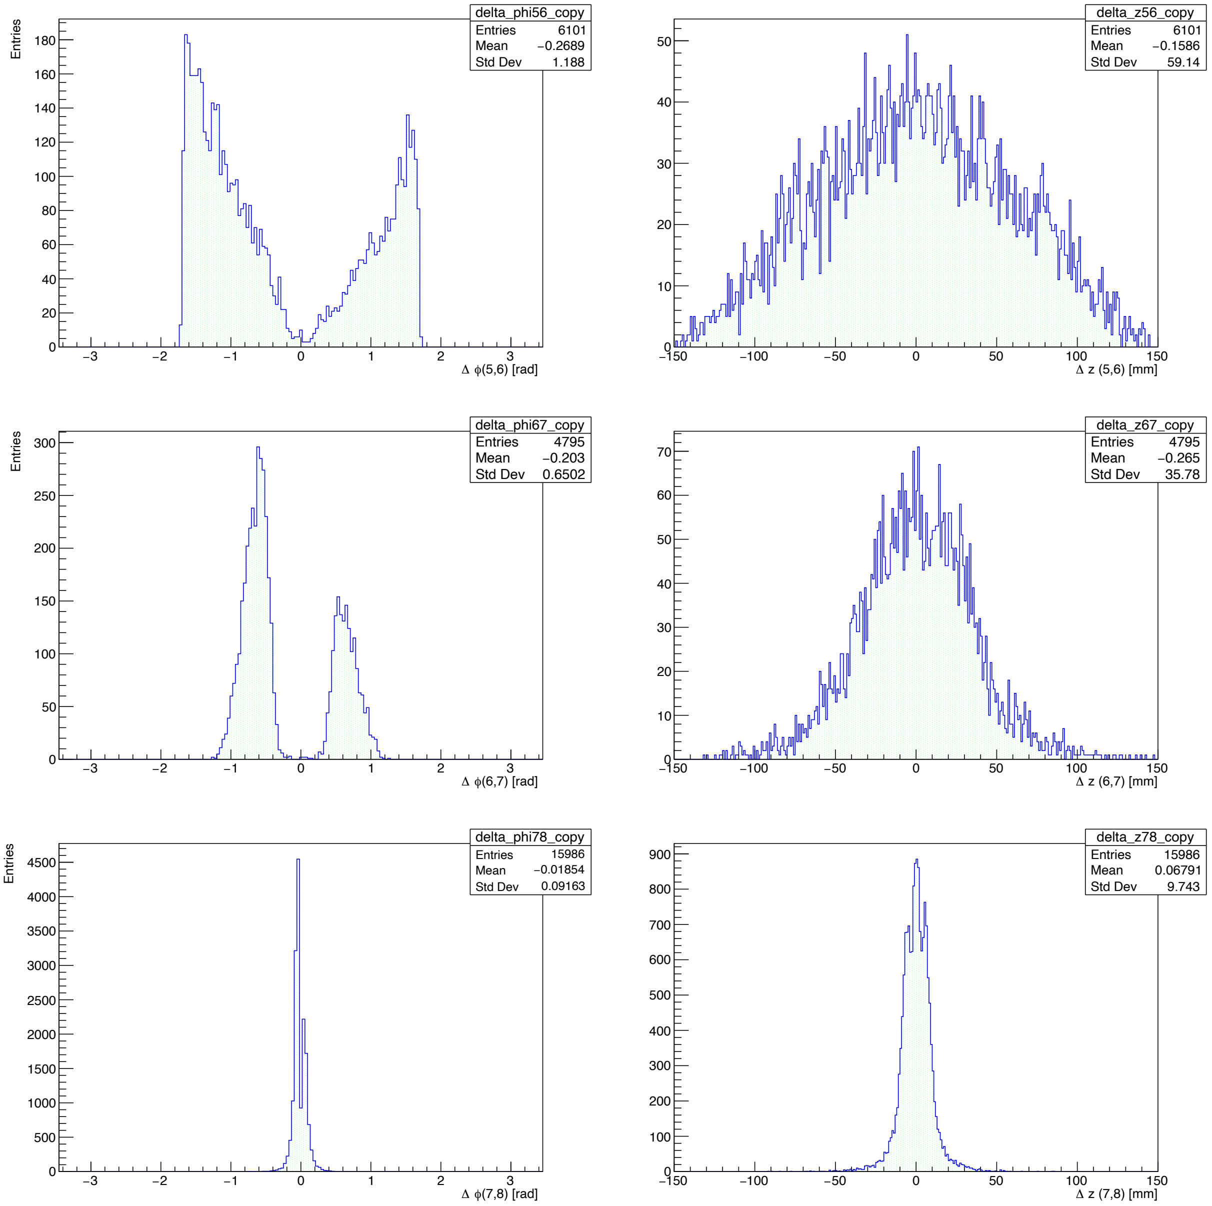

Figure 5.11: Figures showing the difference between ϕ and z of various hits of an EL-8 track. The difference of the fourth hit to the fifth hit is not shown, as it is the same as the EL-5 distributions.

EL-10

The construction of EL-10 tracks begins with a successfully constructed EL-5 track. However, all constructed EL-5 tracks are looped over within a frame, and the sixth hit that is searched for is one that will belong to an already existing EL-5 track that is ingoing. This

Requirement	Value
$\phi_{5\to5}$ [rad]	$\phi < 2$
$z_{5\to5}$ [mm]	$z < 100$
χ^2	< 48

Table 5.5: Table of variable windows used to reconstruct EL-10 tracks.

reconstruction method is identical to how long-8 tracks are constructed within the Phase-1 reconstruction. A figure of a successfully constructed EL-10 is shown in Figure 5.12. A window in ϕ and z is essentially left open in an attempt to locate another EL-5 track. A summary of the cuts are shown in Table 5.5.

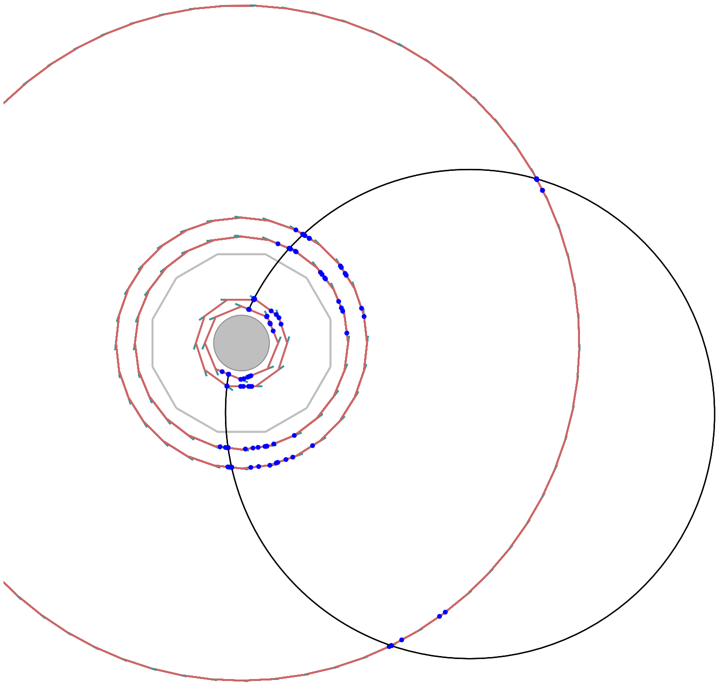

Figure 5.12: Transverse plane of a successfully reconstructed and matched EL-10 track.

5.6 Efficiency of Reconstruction

The reconstruction efficiency distributions as a function of the momentum of the novel extra-long tracks are explored. The truth tracks are found by utilising generator level information regarding hits detected by the pixel sensors. If the trajectory of a particle passes through a silicon layer, this information is known at the generator level, and is saved. This information will be stored as a number, and is equal to the number of the silicon layer. The number will be signed depending on whether the hit originates from a track that is ingoing or outgoing. For layer 1, the inner most vertex layer, an outgoing sensor hit will be equal to +1, an ingoing sensor hit will be -1. For layer 5, an outgoing sensor hit will be +5, and an ingoing hit will be -5. In the Phase-1 simulation, if 6 sequential hits follow the order: +1, +2, +3, +4, -4, -3, a 'truth' long-6 track will be predicted. A triplet fit using the hits will be attempted, and the resulting truth track will be saved. In this work, this is extended now for any track of an arbitrary number of hits, including EL-N, where N is 5, 8 and 10. For example, if sequential hits of +1, +2, +3, +4, +5, -5, -4, -3 are found, an EL-8 truth track is saved, and a triplet fit is performed using a track constructed from only generator level information.

Using the reconstruction algorithms previously outlined in this chapter, EL-N tracks from a signal sample are reconstructed from silicon sensor hits and matched with their generator-level counterparts. A reconstructed track is 'matched' with a truth track if all hits from the reconstructed track originate from the same particle, and that track has a corresponding truth track. The distributions of momentum p, inclination angle λ and polar angle ϕ of the tracks are presented in Figure 5.13, alongside the efficiency, which is defined by the following:

$$\epsilon = \frac{N_{\text{Reconstructed}}}{N_{\text{Truth}}}, \tag{5.4}$$

where both reconstructed and truth tracks have a $\chi^2 < 40$ for EL-5, and < 48 for EL-8 and EL-10. From the plots in 5.13, we observe a perfect efficiency. Every generator-level track which have a χ^2 below 40 is constructed by the algorithms described above. Studies of tracks that are not matched with their generator-level counterparts are not completed within the thesis, so perhaps without the use of truth information fake tracks could be produced with the ϕ and z windows described above. The uncertainty on the efficiency is obtained using the 1σ Clopper-Pearson confidence interval.

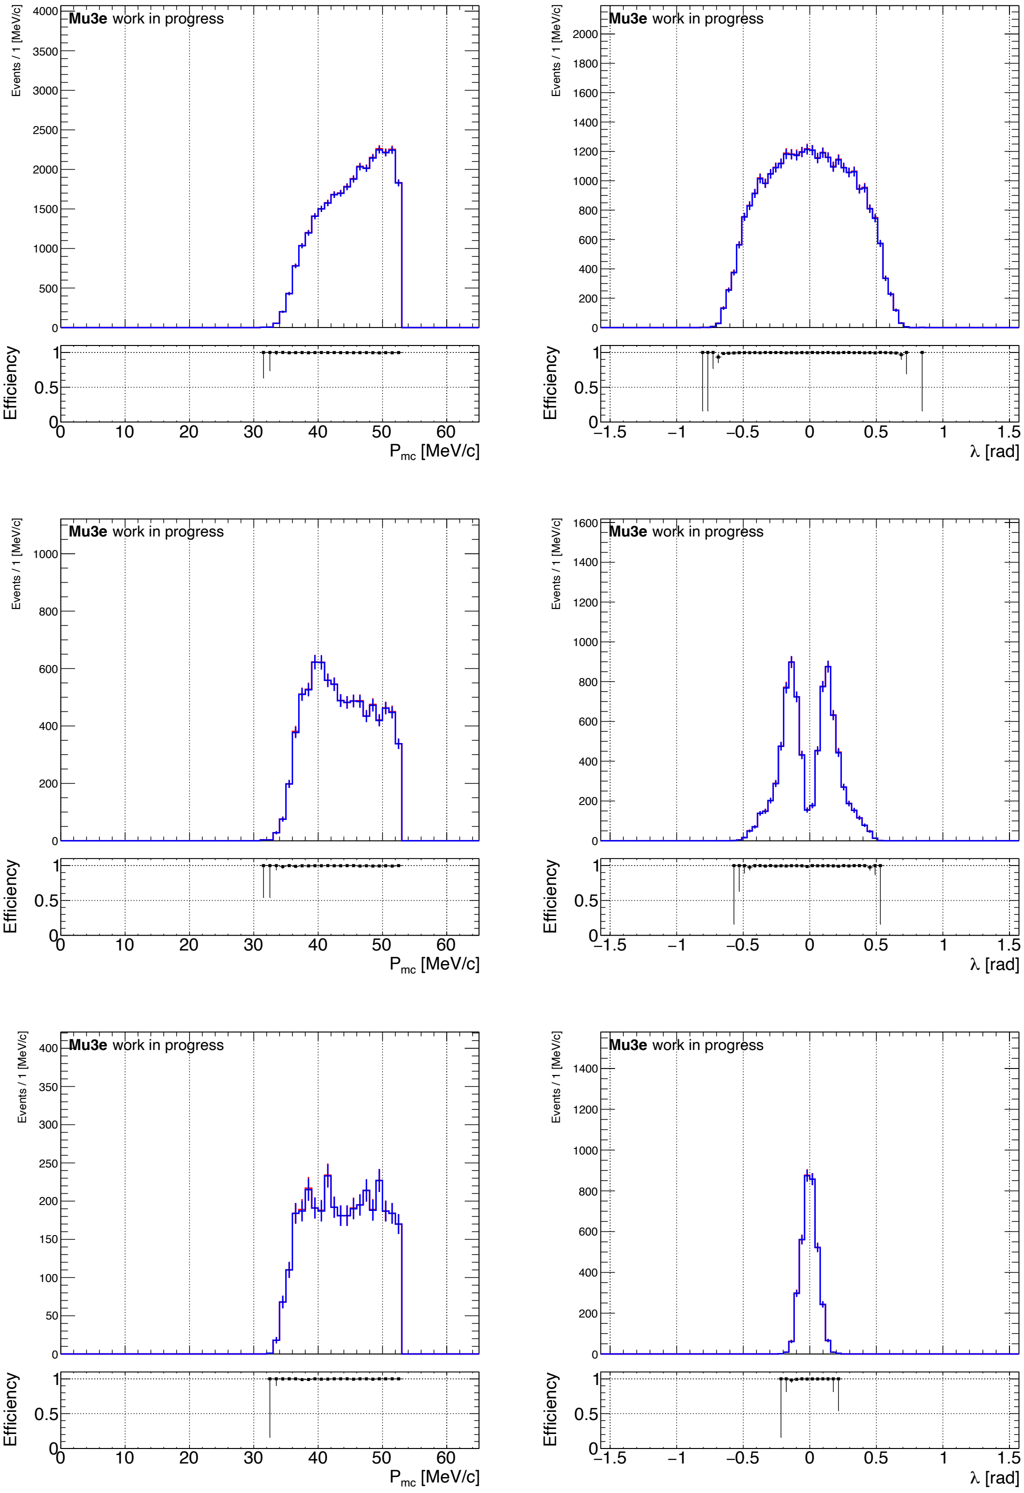

Figure 5.13: Distributions of truth momentum and inclination angle for each type of track explored within this thesis. The ratios of the number of reconstructed tracks with respect to the number of monte-carlo truth tracks are also plotted. No fake track analysis has been conducted.

5.7 Detector Acceptance

The Phase-1 acceptance of the Mu3e experiment was shown in Chapter 2 by Figure 2.15. The acceptance is calculated per bin by the ratio of reconstructed long tracks to reconstructed short tracks. Although Mu3e is a high acceptance detector, there are portions of the momentum - inclination angle phase space where the acceptance is low. Sources of low acceptance within Phase-1 Mu3e include areas where there are no sensitive detectors - either due to the presence of the services of detectors, gaps between the recurl station and the central station, or due to the lack of presence of a detector system altogether. This affects decay products with high momentum, and a low inclination angle. These tracks will travel very far downstream/upstream, way beyond the existing recurl stations. The latter source is answered by a possible Phase-2 scenario of adding extra recurl stations. The Phase-2 scenario of adding an extra layer effectively covers the high momentum, low inclination angle source of acceptance but providing an alternative type of long track to reconstruct that may be destined otherwise to hit the services between the stations.

With the addition of the extra outer layer for the Phase-2 scenario, an improvement to the acceptance is observed in Figure 5.14. In the figure, the acceptance of the Phase-1 scenario is plotted, and gaps in the acceptance can be seen. The length of the outer detector allows for new types of long track to be constructed, that would otherwise hit the non-sensitive services between the central station and the recurl station. These tracks would otherwise only be a short track, and are not considered in the analysis. This leads to an acceptance gain in those areas of 100%, which can be observed in Figure 5.14.

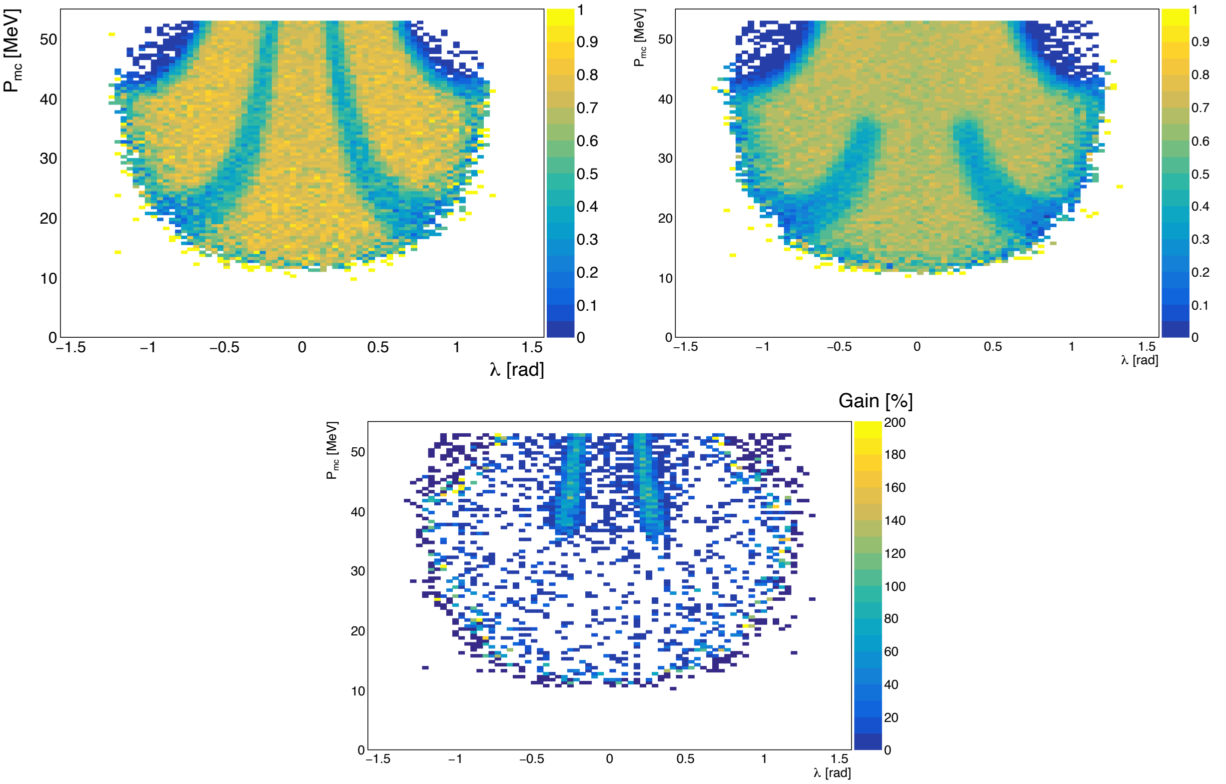

Figure 5.14: Acceptance comparison between Phase-1 and Phase-2 geometries. The gain in the acceptance is also shown in the bottom plot (z-axis is percentage), which is calculated through the percentage increase from the top left plot to the top right.

5.8 Momentum Resolution

As stated in Chapter 2, excellent momentum resolution is of utmost importance to Mu3e as it can be used to effectively suppress combinatorial background and the tail of the internal conversion distribution. The purpose of adding a new detector layer and extending the existing track collections is to improve the momentum resolution of tracks at high momentum. Mu3e would also be sensitive to exotic physics for high momentum tracks with excellent momentum resolution.

Each track reconstructed by the previous algorithms is also matched with a generator-level. From this, the momentum resolution of the detector is estimated. For each of the tracks, the difference between truth and reconstructed momentum is calculated, and a double-sided crystal ball fit is used to estimate the resolution. Bins of truth momentum 1 MeV wide are taken. The σ parameter of the double-sided crystal ball fit is taken to be the resolution. Example fits are shown in Figure 5.16. The momentum resolution is then plotted as a function of truth momentum, and truth transverse momentum, as shown in Figure 5.15. It can be observed that with increasing momentum (>34 MeV), for EL-8 and EL-10 tracks, a new best resolution of 0.1 MeV is obtained at 48 MeV. This is due to the curvature of the track between registered hits reducing the uncertainty on the momentum measurement, as described in Chapter 2. Without this curvature, the extra hits would have little effect, as shown by the resolution estimated with the EL-5 tracks. The curvature of EL-5 tracks isn't sufficient to obtain a resolution better than 0.2 MeV.

A temporary worsening to the resolution is observed at around 35 MeV for the existing Phase-1 long-6 and long-8 tracks. This behaviour is caused by the increase of multiple Coulomb scattering at this momenta, as this is the momentum required for a particle to only just reach the fifth layer. The particles then travel through the fifth layer, without benefiting from the extra hit that the detector provides, which is ultimately detrimental to the overall resolution.

The momentum resolution is also estimated on a two dimensional plot of inclination angle

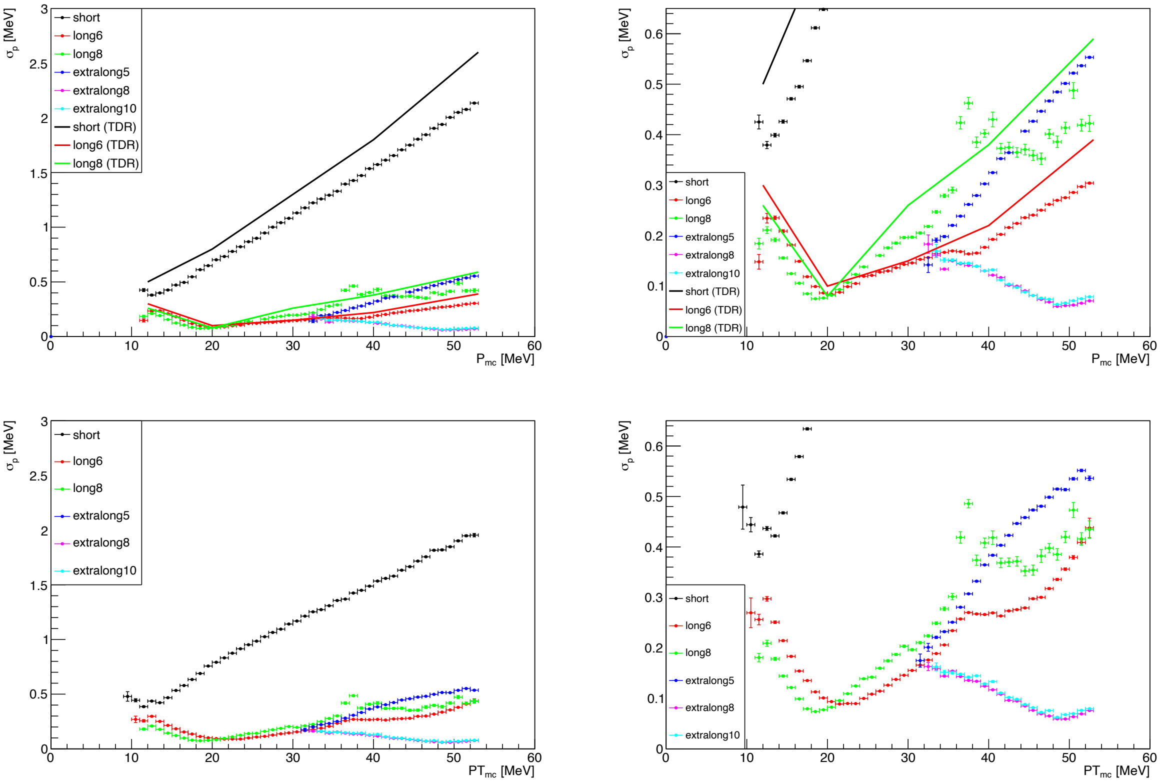

Figure 5.15: Momentum resolution plots for Phase-2 scenario. Double sided crystal ball fit is used.

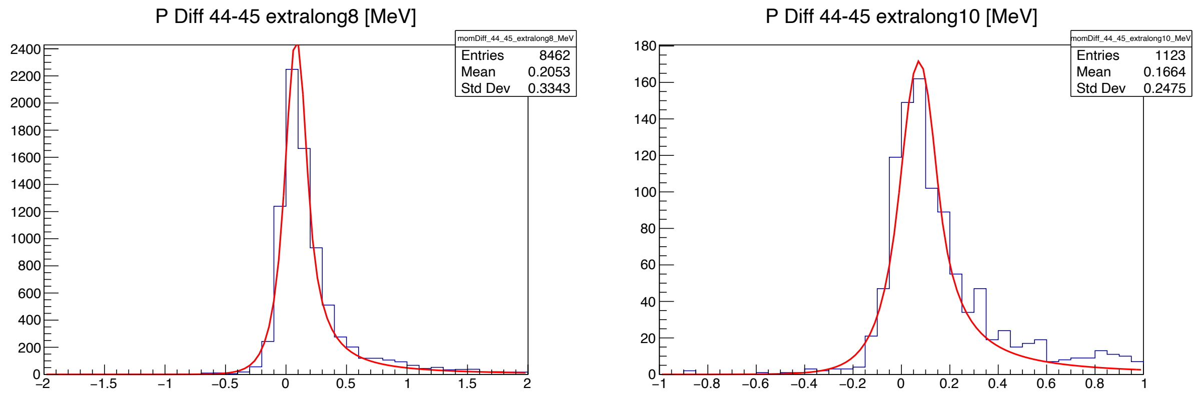

Figure 5.16: Example double-sided crystal ball fit for EL-8 and EL-10 tracks.

λ and momentum in Figure 5.17. Each bin has a number of resolutions resulting from the fits of different types of Phase-1 and Phase-2 tracks, and the average of the sigmas of the fits are taken. An improvement to the momentum resolution is clearly seen for tracks of low inclination angle and high momentum when the Phase-2 geometry iss considered, which is a phase space of importance for exotic physics.

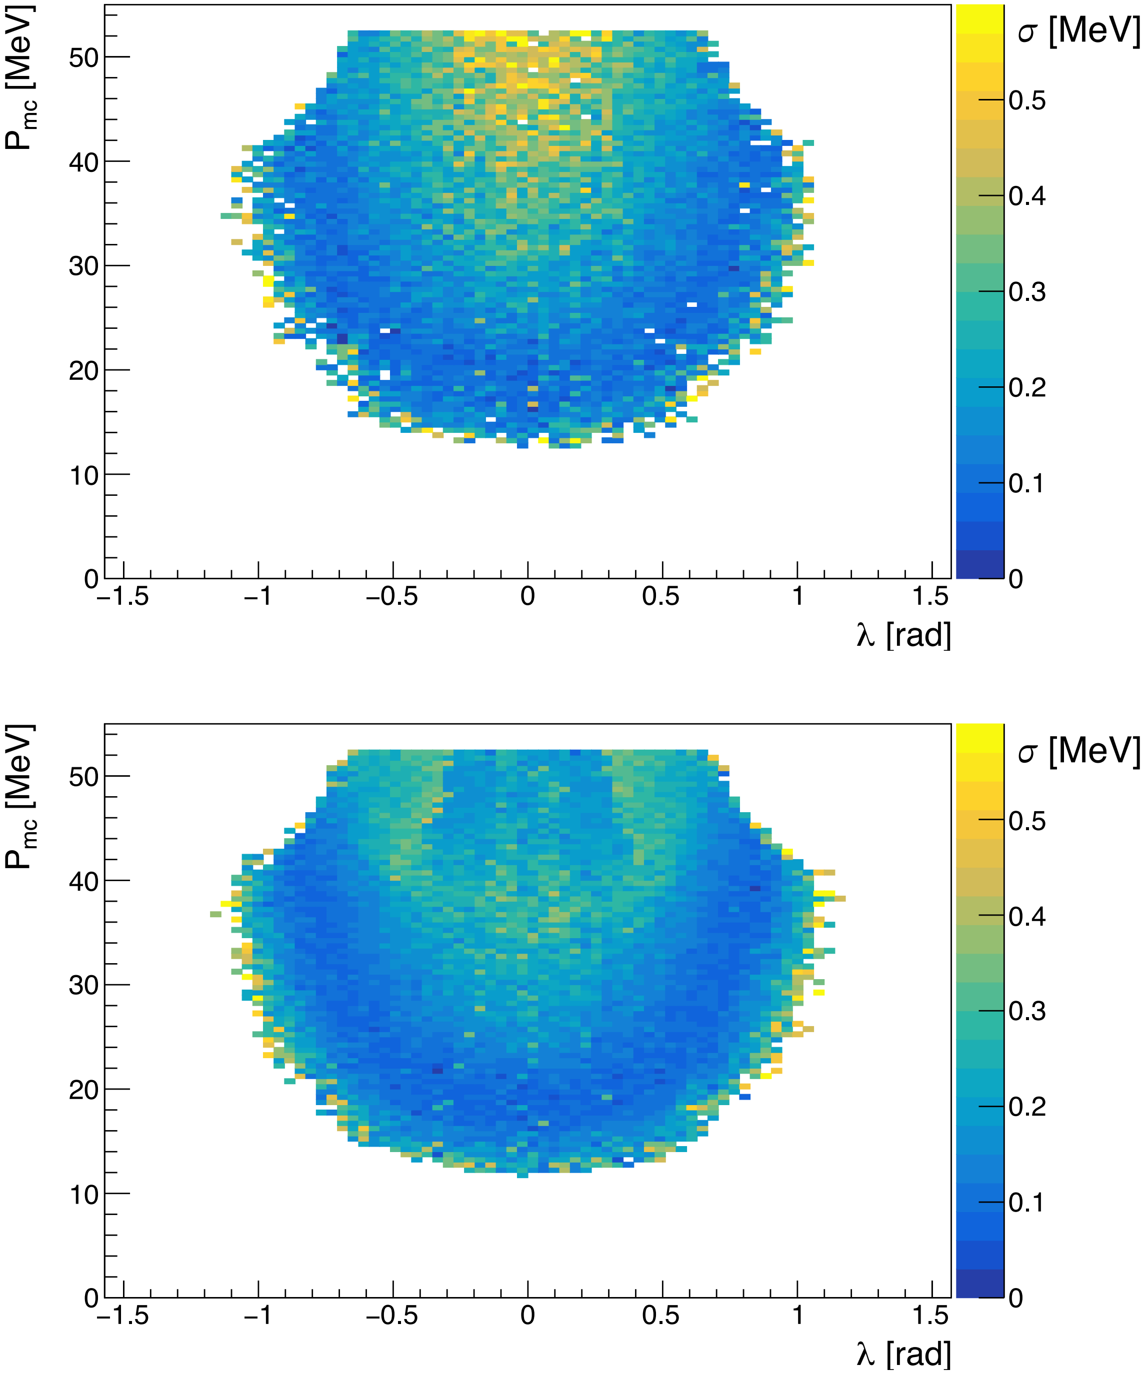

Figure 5.17: Two-dimensional plot of the momentum resolution. Z axis is the momentum resolution.

5.9 Improvements to $\mu \rightarrow eee$ Sensitivity

The improvements to the sensitivity of $\mu \rightarrow eee$ signal is also studied. With the addition of the extra layer, extra long tracks can be constructed that have a superior momentum resolution to default Phase-1 tracks. This leads to an improvement to the mass resolution, which is vital for suppression of accidental background. This is shown in Figure 5.18, where it is shown that the addition of extra long tracks improves the mass resolution by 10%. In the Phase-1 Mu3e detector design, particles can escape detection by travelling into the space between stations. With the addition of an extra outer layer, tracks that have a high enough momentum would first hit the extra outer layer twice before passing inbetween the stations, thus gaining in efficiency when extra long tracks are constructed. The overall change to the sensitivity is shown in Figure 5.19, where the addition of extra long tracks improves the efficiency, which is shown to increase to 12.96%, but unfortunately has only a marginal effect on the final sensitivity.

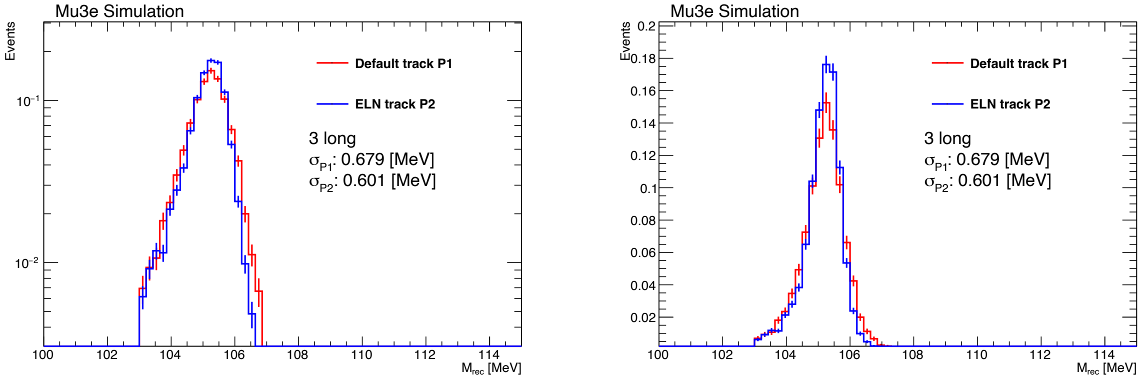

Figure 5.18: Improvements to the mass resolution for a $\mu \rightarrow eee$ signal sample, using extra long tracks, compared with the default Phase-1 tracks. The resolution is taken to be the RMS of the histograms.

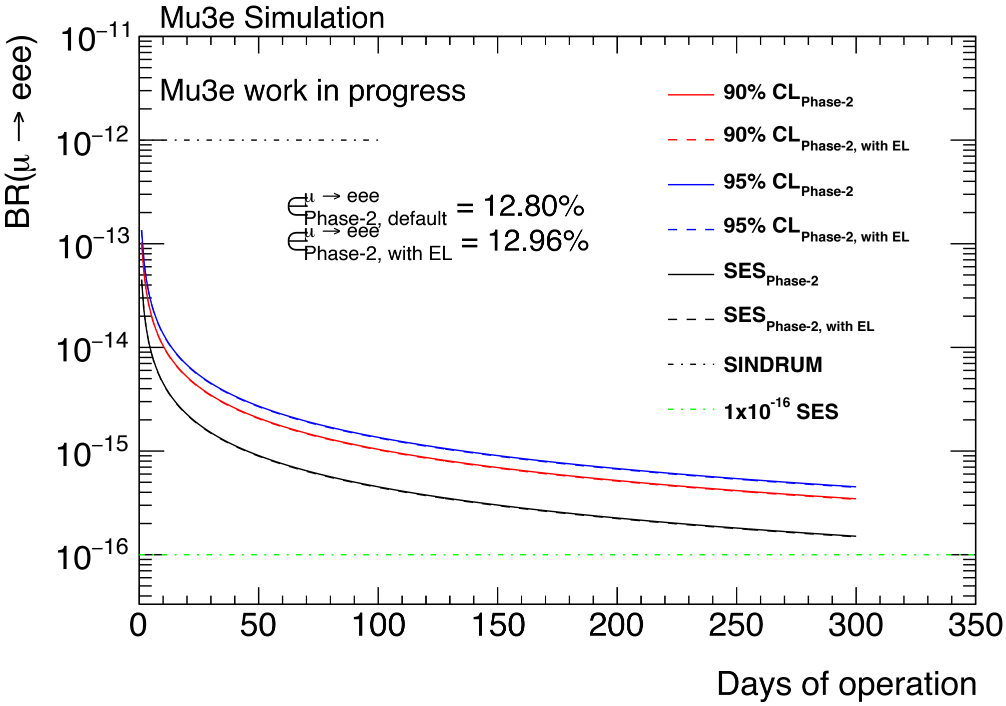

Figure 5.19: Estimated Sensitivity to $\mu \to eee$ for Phase-2 Mu3e detector. Nominal Phase-2 efficiency is assumed to be the same as Phase-1. Dashed lines are the sensitivity obtained by using extra long tracks, and are just below the original solid lines.

5.10 Conclusion

A modification of the Phase-1 simulation was undertaken with the goal of exploring the addition of an extra outer layer for the second phase of the Mu3e experiment. Algorithms for new tracks, EL-5, EL-8 and EL-10 were added seamlessly to the track reconstruction. Reasonable windows were defined for the reconstruction, and the reconstruction efficiency was estimated to be very high, however this could change with a study into the rate of track fakes. Ultimately the momentum resolution of these tracks were estimated, granting a second minimum in the momentum resolution plot at around 48 MeV of 0.1 MeV. The additional layer however adds only a very marginal improvement to the sensitivity of the primary $\mu \to eee$ search via an increase to the number of vertices reconstructed. It could however improve background rejection, but this will need further study. The study could serve as a basis for exploring alternative Phase-2 scenarios, the existing Phase-1 simulation has been found to be flexible in the implementation of new innovative designs. These new designs need to be explored in order to effectively address the challenges Mu3e will face with the incoming HIMB upgrade.

Chapter 6

Searching for Axion-Like Particles with Mu3e

The Mu3e Experiment's primary goal is to search for the lepton flavour violating decay $\mu \rightarrow eee$. However, due to Mu3e's large detector acceptance, and the large number of expected muon decays, other aspects of muon physics could be probed as well. The Mu3e detector's ability to probe the dark sector (through a search for the dark photon) has already been discussed in Chapter 3. This chapter will discuss how a search for Axion-Like Particles (ALPs) with Mu3e would lead to a significant improvement on the sensitivity to lepton flavour violating ALP models.

In this chapter, a study motivating the search for ALPs using the Mu3e detector is discussed using simulated samples of muon decays with an ALP in the final state. The decay process is $\mu \rightarrow ea$, where the ALP, denoted by a, will subsequently decay into an electron-positron pair. The results from the sensitivity study were contributions to the 2022 Snowmass process [138].

6.1 Searching for Axion-Like Particles (ALPs) with Mu3e

In the ALP model introduced in Chapter 1, ALPs have two couplings: a lepton flavour violating coupling, $g_{\mu e}/\Lambda$, and a lepton flavour conserving coupling, g_{ee}/Λ, where Λ is the energy scale and g is a dimensionless constant. These are written as $\Lambda_{\mu e}$ and Λ_{ee}, respectively. The lepton flavour violating coupling, $\Lambda_{\mu e}$, is proportional to the branching ratio of the decay process $\mu \to ea$. The lepton flavour conserving coupling Λ_{ee} is proportional to the lifetime of the ALP, $a \to ee$. The efficiency of reconstructing an ALP of a finite lifetime is estimated using the Geant4 simulation of the Mu3e detector, described in Chapter 2. These lifetimes are proportional to Λ_{ee}, and so by simulating many ALPs of varying masses and lifetimes, a 2-dimensional plane of m_a- lifetime (and so Λ_{ee}) will be constructed. The branching fraction that can be excluded from transforming these efficiencies to an upper limit on the branching fraction will exclude a particular scale $\Lambda_{\mu e}$. These efficiencies will then be interpreted on the 2-dimensional plane $m_a - \Lambda_{ee}$ for different values of the lepton flavour violating coupling $\Lambda_{\mu e}$, producing a contour plot.

The $\mu \to ea(a \to ee)$ process occurs in two steps: first the muon comes to rest upon the surface of the target, and then it decays. The decay of the muon will lead to the generation of a positron and the ALP will appear roughly back-to-back. At a next stage, the second process occurs, whereby the ALP decays to an electron-positron pair. This is significant when vertex fitting is concerned, as the backwards extrapolation of the Michel positron track fit will follow roughly the straight-line path of the ALP, which will lead towards the electron-positron pair resulting from the decay of the ALP. This then allows for the reconstruction of $\mu \to ea(a \to ee)$ processes using a prompt reconstruction algorithm, even those decays where the ALP could travel up to a few centimetres away from the muon decay vertex. Figure 6.1 shows a correlation plot between the reconstructed momentum of the three-electron system P_{CMS} and the distance travelled by an ALP before it decays, for an ALP of mass 40 MeV and lifetime 0.01 ns. It can be observed from the figure that for higher ALP travel distances, due to the backwards extrapolation of the Michel positron, and its curvature due to the magnetic field, the centre of mass system momentum will stray from zero. A bias appears at 0 mm in the plot, which is an artefact caused by using the centre of mass momentum and not the x, y and z components. The extrapolation of the positron track is assumed to be a helix, and the length of the helix, and so total curvature, will increase as the ALP travels a further distance before it decays. This increase in the distance the ALP travels will result in an increase to the curvature of a positron track that is used to reconstruct the vertex, and this will negatively impact the reconstructed momentum of the three electron system, as one of the momentum vectors will become increasingly distorted the further the ALP travels before it decays.

ALPs of masses 10 to 90 MeV/c^2, and of finite lifetimes of up to 1 ∼ ns were inputted into the Mu3e experiment's software framework. Muon decays to ALPs are generated by MadGraph5@NLO[1]. The kinematics of the decay of the ALP, and the resulting momenta

---

[1] Modifying code supplied by Andrea Thamm. Thank you for providing help with the MadGraph5@NLO cards for the study

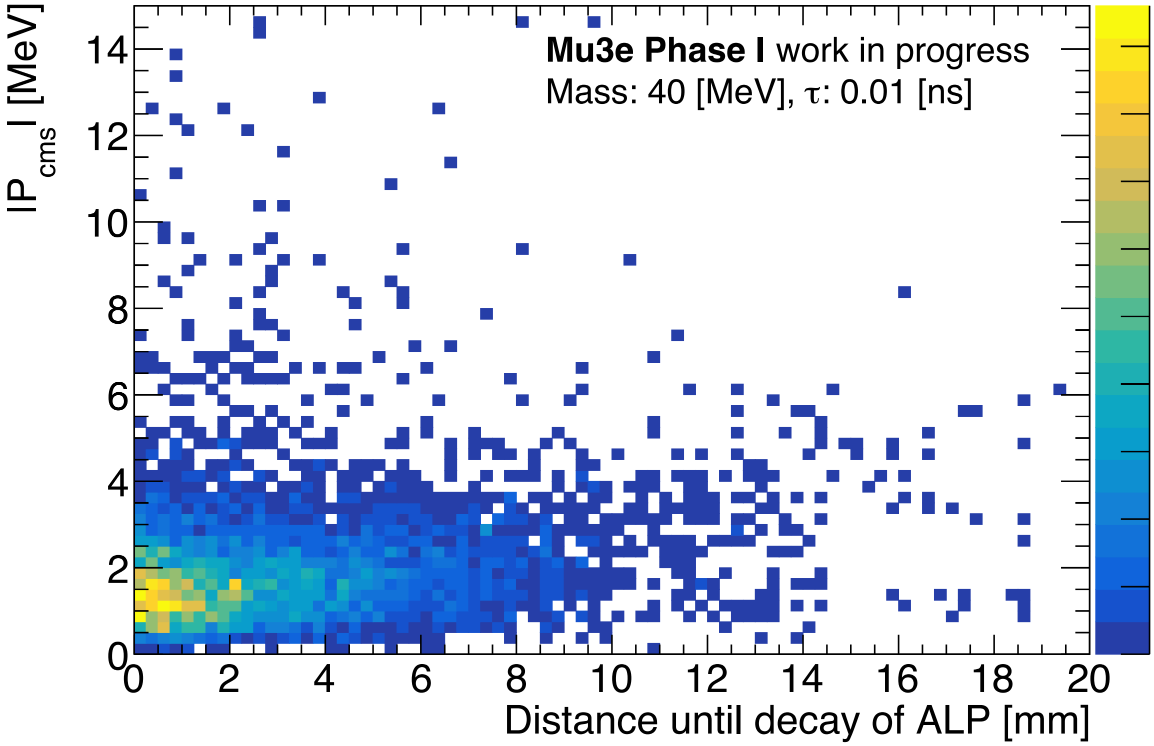

Figure 6.1: Correlation plot between the reconstructed muon's centre-of-mass system momentum and the ALP travel distance.

of the electron-positron daughters, was inputted into the Mu3e simulation. The simulation is performed using the Phase-1 setup described in Chapter 2. These tracks are then used with the standard three-track, prompt, vertex reconstruction. There are also possibilities within Mu3e to add further vertex algorithms that will reconstruct particle decays with displaced vertices. This is not explored here, but their importance is stressed in several phenomenological papers, such as the search for displaced dark photons [132].

Fortunately, the final topology of the decay process $\mu \rightarrow ea(a \rightarrow ee)$ has exclusively three electrons. This allows us to search for ALPs with no change to the online reconstruction and analysis chain. When requiring three long tracks to form vertices, it also implies that the signals will be background free as the same selection requirements are applied as in the search for $\mu \rightarrow eee$.

6.2 Generation of $\mu \rightarrow ea(a \rightarrow ee)$ events

A Universal FeynRules Output (UFO) model incorporating the ALP coupling is implemented, using the FeynRules2.3.43 [139] software framework. The output files are then inputted into MadGraph5@NLO 3.4.1 [128] (MG5). The $\mu \rightarrow ea(a \rightarrow ee)$ decays are outputted into Les Houches Event [134] file format. The Mu3e Experiment's Software framework takes these files as an input to run the simulation. The simulation is set to generate one muon per frame which will decay into the ALP. This ALP will then decay into an e^+e^- pair.

Due to the momentum acceptance of the Mu3e detector (tracks lower than 10 MeV are outside the acceptance), as well as the kinematics of the decay, ALP masses of 10 to 90 MeV/c^2 were generated with MG5, and inputted into the Mu3e Experiment's software framework. For each mass a number of lifetimes were generated between prompt (0 ns lifetime) and displaced decays (1 ns lifetime) of the ALP. For each mass, 100,000 signal events were generated, which was necessary as the efficiency for some ALPs of a high lifetime was observed to be very low ($< 0.01\%$). No background is generated, as the full Mu3e signal selection is applied, which is assumed to allow for a background-free signal if we exclusively select vertices constructed from three-long tracks.

6.3 Generator Level ALP Distributions

Truth distributions for the e^+e^- pair generated from the decay of a multitude of ALPs are shown in Figure 6.2. Plotted are the momenta of the electron and positron decay products and their angles between one another. The positron that results from the muon decay is defined as 'electron one', whilst 'positron two' is emitted from the decay of the ALP.

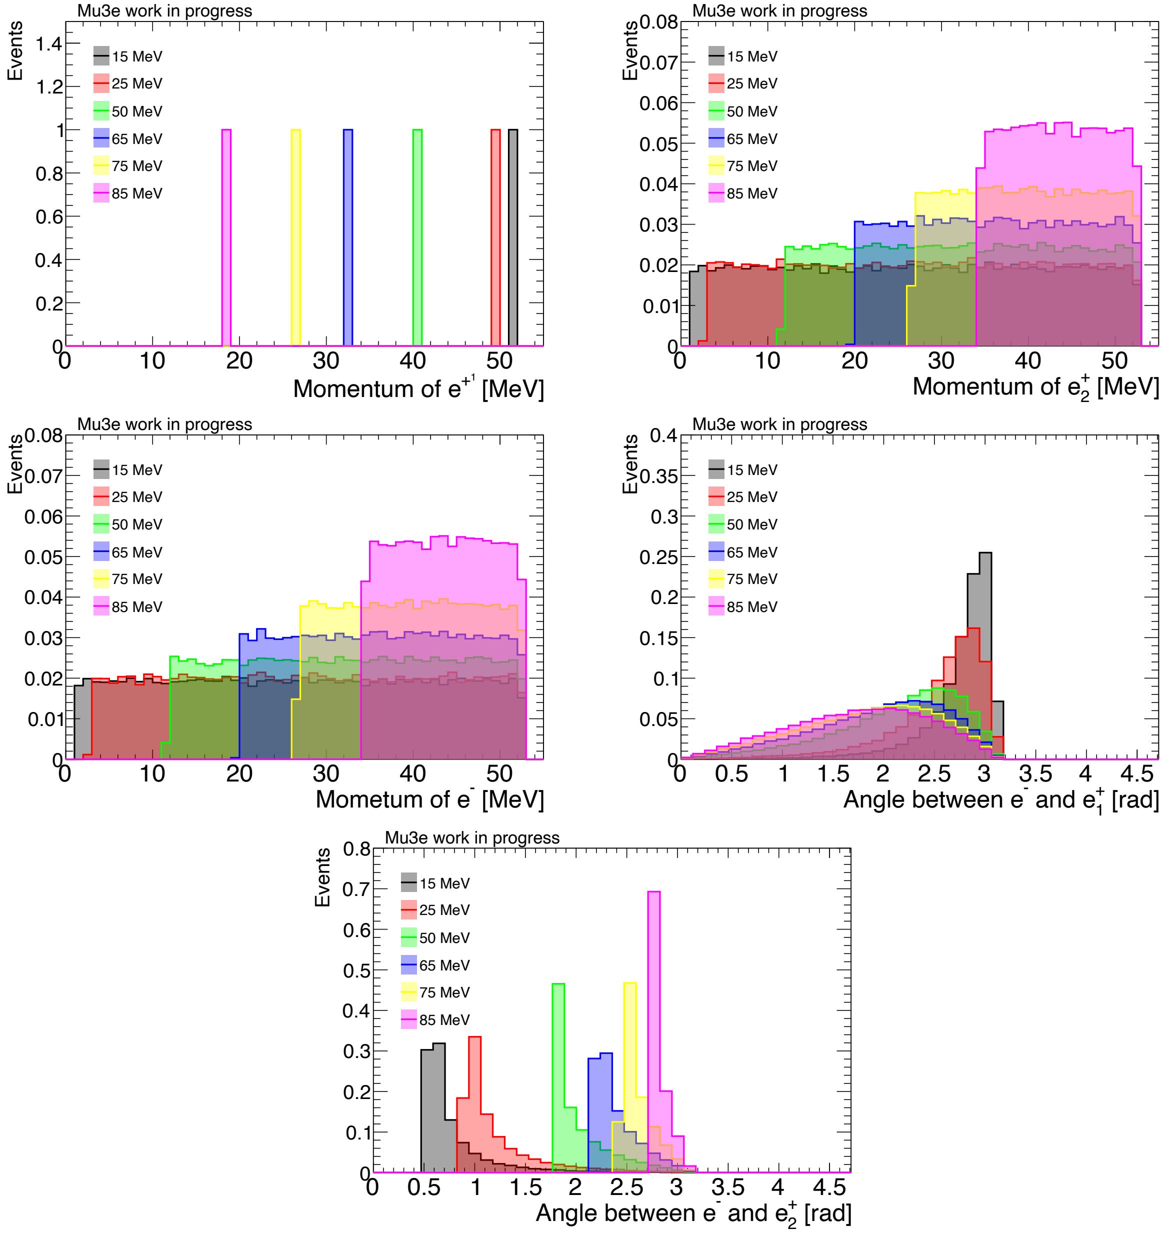

Figure 6.2: Truth distributions for a multitude of ALP masses.

6.4 Sensitivity Estimation

The sensitivity of Mu3e to ALP couplings is obtained by first estimating the efficiencies of the Mu3e Phase-1 detector to reconstruct ALPs of various masses and lifetimes. All the requirements applied to the reconstructed ALP vertices are also applied to analysis of $\mu \rightarrow eee$, so by enforcing that three long tracks form a vertex, our analysis is assumed to be free from background. With this assumption in mind, the efficiencies can then be transformed into a 90% confidence level upper limit assuming Poisson statistics, in the same way as in the $\mu \rightarrow eee$ analysis.

ALPs of masses 10 to 90 MeV/c^2 and of lifetimes up to 1 ns are simulated using the Mu3e software framework. These efficiencies are shown in Figure 6.3, where the efficiencies are also compared with the efficiency to reconstruct $\mu \rightarrow eee$ signal. We observe that higher ALP masses tend to have a higher vertex reconstruction efficiency. Even higher than the efficiency observed to reconstruct a $\mu \rightarrow eee$ event using the same selection criterion.

Two one-dimensional splines are introduced to interpolate between each mass and lifetime point on the efficiency plane. Each spline is then interpreted into an upper limit on the branching fraction, which can then be utilised to exclude branching fractions of the decay $\mu \rightarrow ea(a \rightarrow ee)$ for a particular scale $\Lambda_{\mu e}$ per mass of the ALP, m_a.

6.5 Reconstruction

The search for an ALP signal with Mu3e can be performed using three-track prompt event candidates. The candidates used in the analysis will appear identical to those searched for in the $\mu \rightarrow eee$ signal analysis, so no change in the data acquisition and online reconstruction will be necessary. The vertices in the search will also have properties similar to those Mu3e vertices in the signal region, therefore there will be an absence of background in our search if we exclusively use vertices with three long tracks.

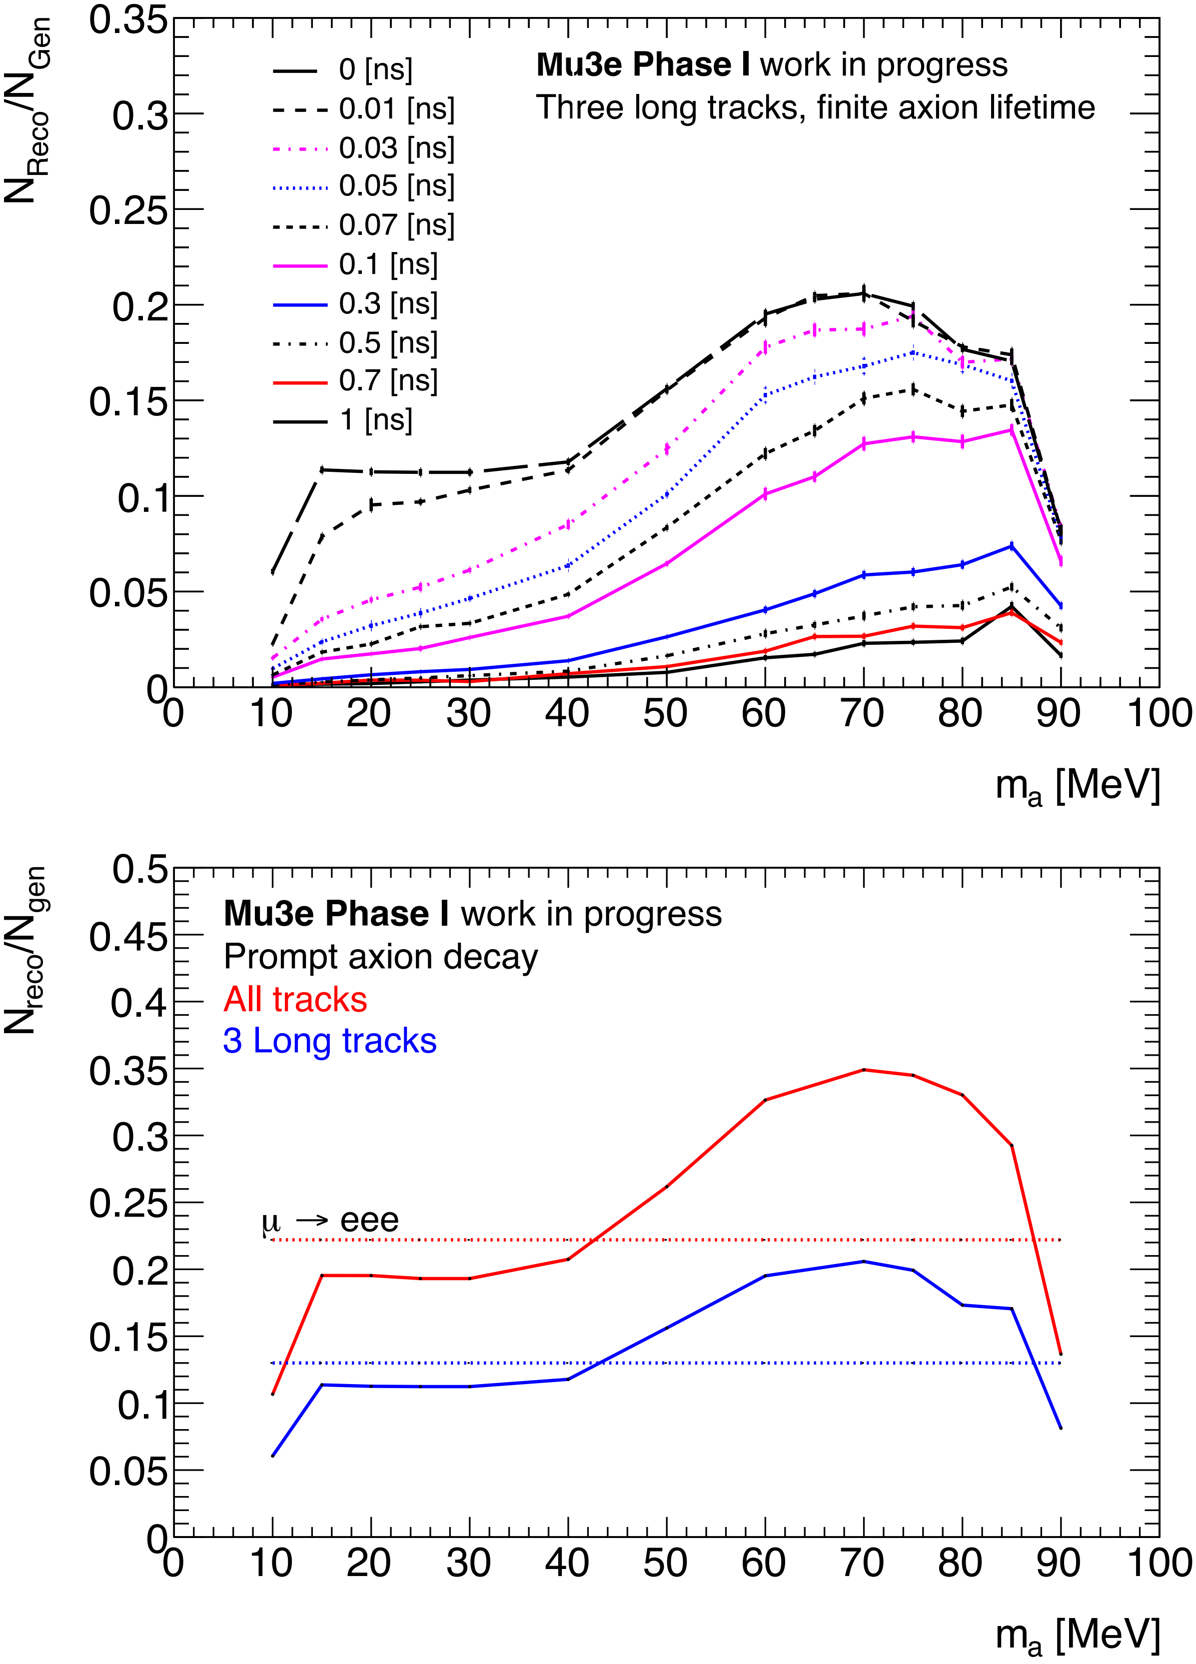

Figure 6.3: Vertex reconstruction efficiency for a multitude of ALP masses and lifetimes. Three long tracks are required to form a vertex.

6.6 Selection Criteria

The selection criterion applied to reconstructed vertices from the $\mu \to ea(a \to ee)$ process are identical to those of the $\mu \to eee$ analysis shown in Chapter 3. In the online analysis, as previously described in Chapter 2, track reconstruction uses hits from the silicon layers to form tracks, with timing information from scintillator layers added later. These tracks are then extrapolated towards the target in an attempt to reconstruct a vertex. These reconstructed vertices then must match the rigorous $\mu \to eee$ selection criteria.

The quantities used to assess the quality of a reconstructed vertex in this analysis are identical to those that maximises $\mu \to eee$ signal efficiency and completely rejects background. There are cuts on the following quantities: vertex fit χ^2, momentum and invariant mass of three electron system, positron-electron mass, and timing. These distributions are plotted in Figure 6.4 for an ALP mass of 40 MeV. Further requirments could be applied in the future, for example, a requirement that has previously been explored in the $\mu \to eee$ analysis is the distance between the surface of the target and the reconstructed vertex. Figure 6.5 shows the vertex reconstruction efficiency as a function of the radius of decay of the ALP, and the total distance until decay of the ALP. An example cutflow of the selection used is shown in Table 6.6.

It can be observed that no vertices are reconstructed where the radius of the ALP decays past the first silicon layer of Mu3e's detector (23.3 mm), as no tracks will be constructed without using this first silicon layer hit (and so there will be no vertex that can be reconstructed). The figure also shows that it is possible for the ALP to travel a distance of roughly 2 cm, and still be reconstructed using the current prompt three-track vertex reconstruction algorithm that is employed by Mu3e. However as the ALP travels to larger distances, the momentum cut begins to reduce the signal reconstruction efficiency. This is due to the physical properties of the ALP decay process: there are essentially two back-to-back decays, (ALP and positron, and then positron-electron). When the positron is extrapolated back towards the positron-electron pair that results from the ALP, it is quite easy to reconstruct a vertex. This becomes

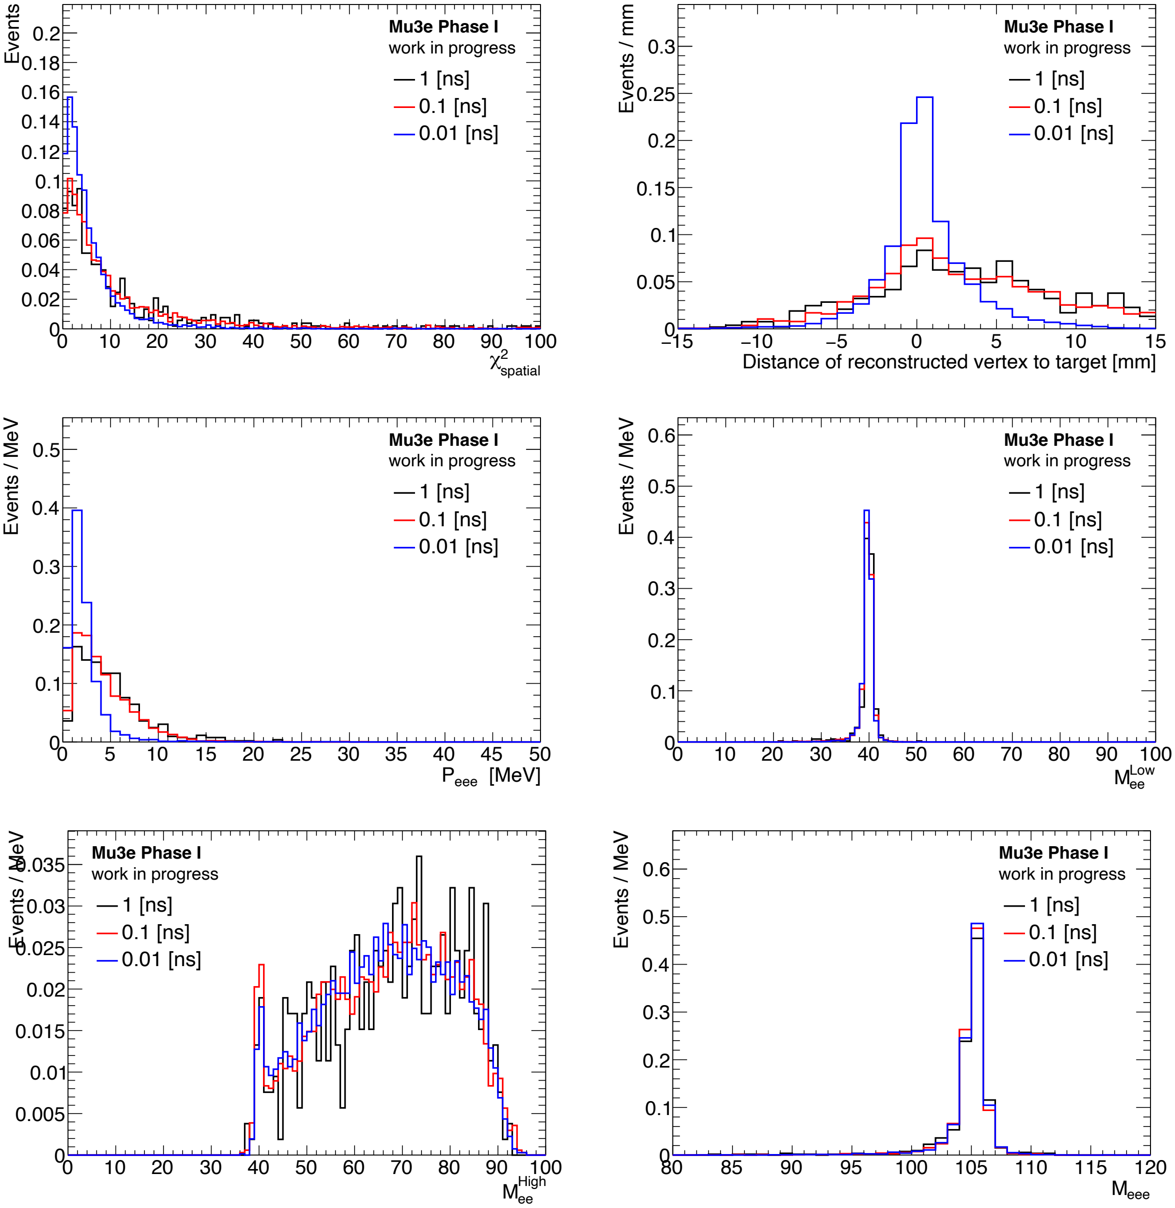

Figure 6.4: Several reconstructed variables (Vertex spatial χ^2, distance of reconstructed vertex to surface of target, three-electron system momentum, dielectron invariant mass and three-electron system invariant mass) are plotted for a 40 MeV ALP for varying lifetimes. Distance of the reconstructed vertex to the target surface is the only variable currently not cut for Mu3e signal.

more difficult the further the ALP decay distance becomes, as observed by Figure 6.1 - the momentum of the three electron system increases as the ALP distance increases.

Figure 6.6 shows that the vertices that are reconstructed are within ∼2 mm of the ALP

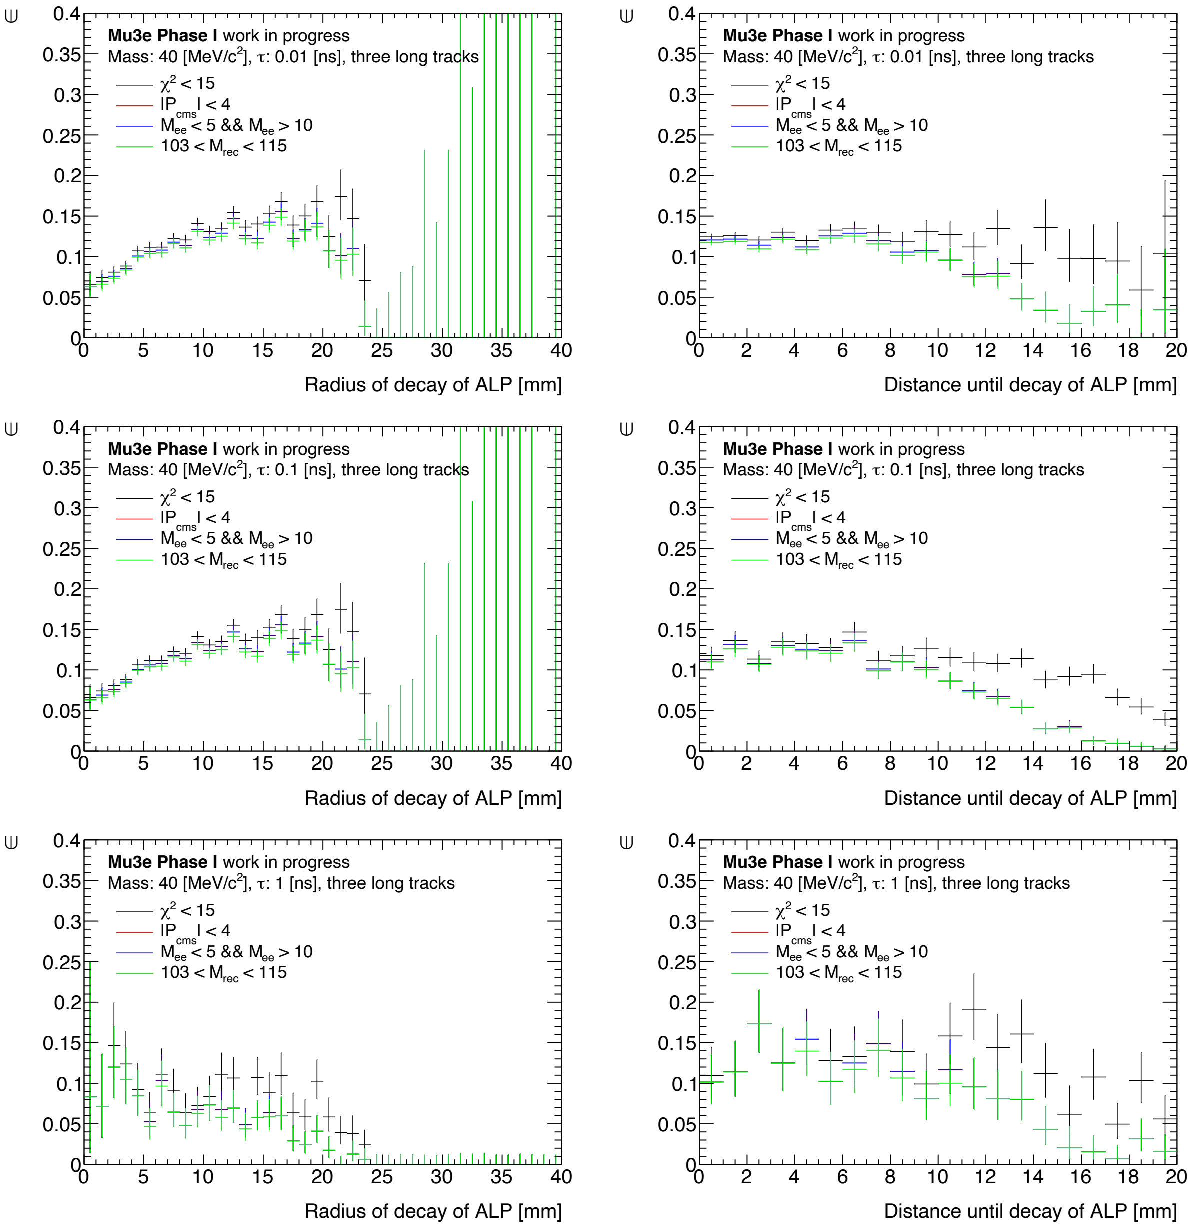

Figure 6.5: Efficiencies for varying cuts on the quality of the reconstructed vertex. Vertex reconstruction efficiency is plotted as a function of the radius of decay ALP and distance until decay of ALP. Radius efficiency drops to zero past 23.3 mm, as this is the distance at which the first silicon layer is located, and reconstructed tracks require this initial hit in the reconstruction.

decay vertex. The vertex reconstruction algorithms are therefore reconstructing the ALP decay location rather than the muon. If the ALP travels a far enough distance from the surface of the target, a cut on this reconstructed vertex to the surface of the target will reduce the signal efficiency.

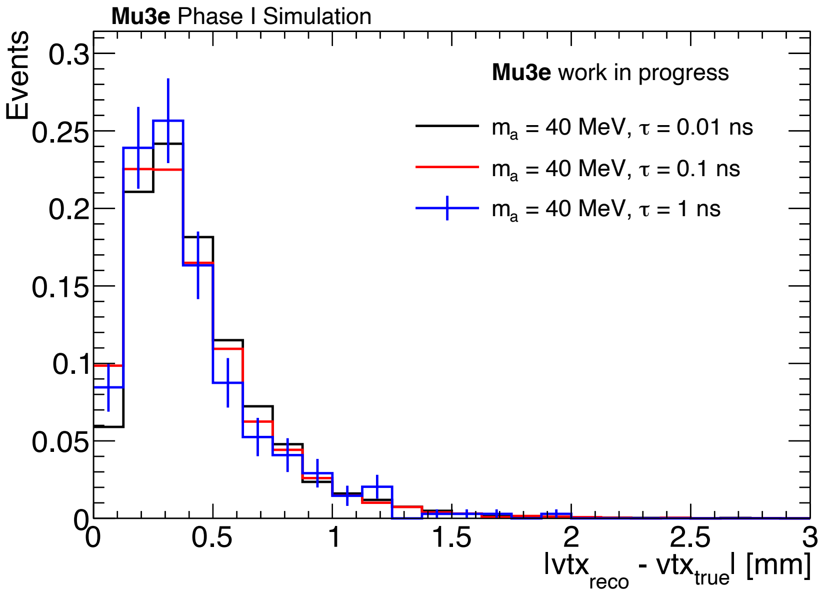

Figure 6.6: Distance between monte-carlo reconstructed and truth vertex positions of the ALP. The errors are the square root of the sum of the square of the weights.

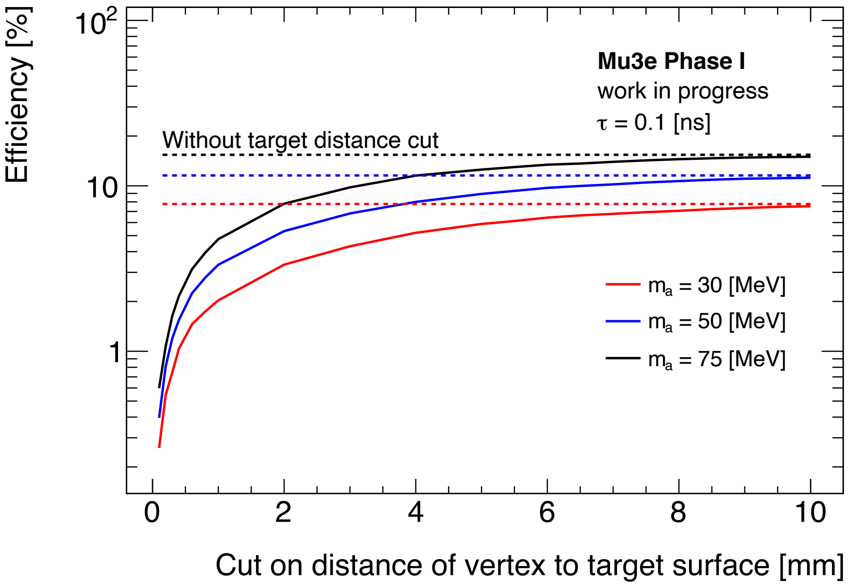

Figure 6.7: Signal efficiency for several ALP masses after applying all selection cuts and adding additional cuts on the distance between the surface of the target and the reconstructed vertex.

Selection	$\mathrm{N_{Vertices}}$	Σw	ϵ [%]	$\sqrt{\Sigma w^2}$/norm [%]	ϵ_{step}[%]
Initial reconstructed vertices	10442	10442	-	-	-
Target region	9245	9245	22.81	0.24	88.54
Track hit number < 3	4327	4327	10.67	0.16	46.80
Track 'prime'	4227	4227	10.43	0.16	97.69
Track type	4227	4227	10.43	0.16	100.00
Track unique	4227	4227	10.43	0.16	100.00
Vertex Multiplicity	3385	3385	8.35	0.14	80.08
Vertex $\chi^2_{\mathrm{spatial}} < 15$	3385	3385	8.35	0.14	100.00
$\mathrm{P_{eee}} < 4$	2587	2587	6.38	0.13	76.43
$\mathrm{M^{Low}_{ee}} < 5.00$ or $\mathrm{M^{Low}_{ee}} > 10.00$	2586	2586	6.38	0.13	99.96
$103.00 > \mathrm{M_{eee}} < 110.00$	2525	2525	6.23	0.12	97.64

Table 6.1: Table of selection cuts and their total efficiencies for an ALP of mass 50 MeV and lifetime 0.1 ns. The denominator is the number of muons decaying on the target. Cuts are described in Chapter 3, and in section 3.4. Vertices are weighted ('w') as 1 in this sample.

6.7 Sensitivity

The expected sensitivity will be estimated assuming a given number of muon stops on the target. The number of muon stops chosen in this analysis is the same number of stops expected in Phase-1 Mu3e to reach a branching ratio upper limit of 1×10^{-15} for $\mu \to eee$ signal, which will take 450 days of beam time at 2×10^{8} muons per second.

Using the method described in the earlier section to estimate a branching ratio upper limit, estimates of the range of sensitivities that could be obtained by Mu3e is shown in Figure 6.9. The upper limit to the lepton flavour conserving coupling for a multitude of ALP masses (10 to 100 MeV) is estimated at 90% confidence level for a variety of lepton flavour violating couplings. Using the Mu3e detector, it can be observed that to completely exclude this region, an energy scale of 5×10^{13} GeV is necessary for the lepton flavour violating coupling. Small exclusion regions begin to appear with an increase to the energy scale, which is a consequence of the shape of our upper limit distribution. However it can also be noted that the highest energy scale available to be probed by Mu3e is roughly 8.2×10^{13} ~GeV, and Mu3e is able to probe ALP masses ranging between 10 MeV up to the muon mass. This is an order of two increase from the SINDRUM experiment, which results from the four-order increase in the number of total muon stops. The sensitivity of the SINDRUM experiment to ALP couplings is shown in Figure 6.8.

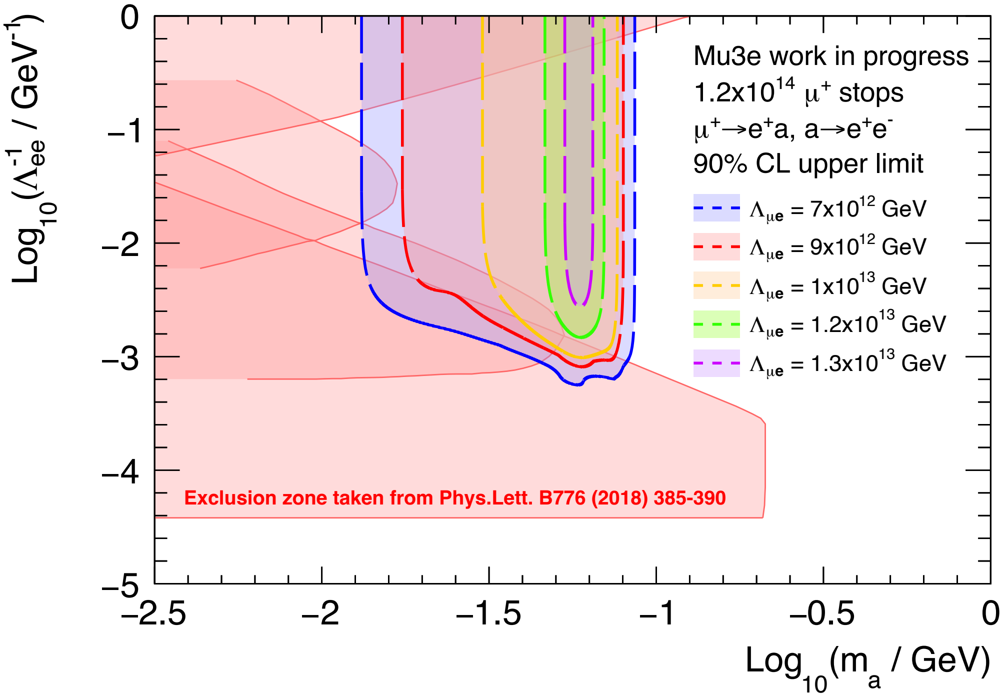

Figure 6.8: Sensitivities to ALPs estimated using SINDRUM's total number of muon stops (1.2×10^{14}), but using Mu3e's simulated Phase-1 detector and reconstruction algorithms.

The effect of adding a further cut on the accepted distance between the surface of the target and the reconstructed vertex is also studied, where it is found that an added requirement for vertices located less than 3 mm from the target would slightly negatively impact Mu3e's sensitivity to ALP masses. In Figure 6.7 a further cut on the distance of the reconstructed vertex to the target is applied: if the vertex is reconstructed further than 3 mm away from the target, the vertex is removed. The result of the sensitivity plot for various values of the target distance cut is shown in Figure 6.11, where it can be seen that the sensitivity is only marginally affected with a requirement of 3 mm.

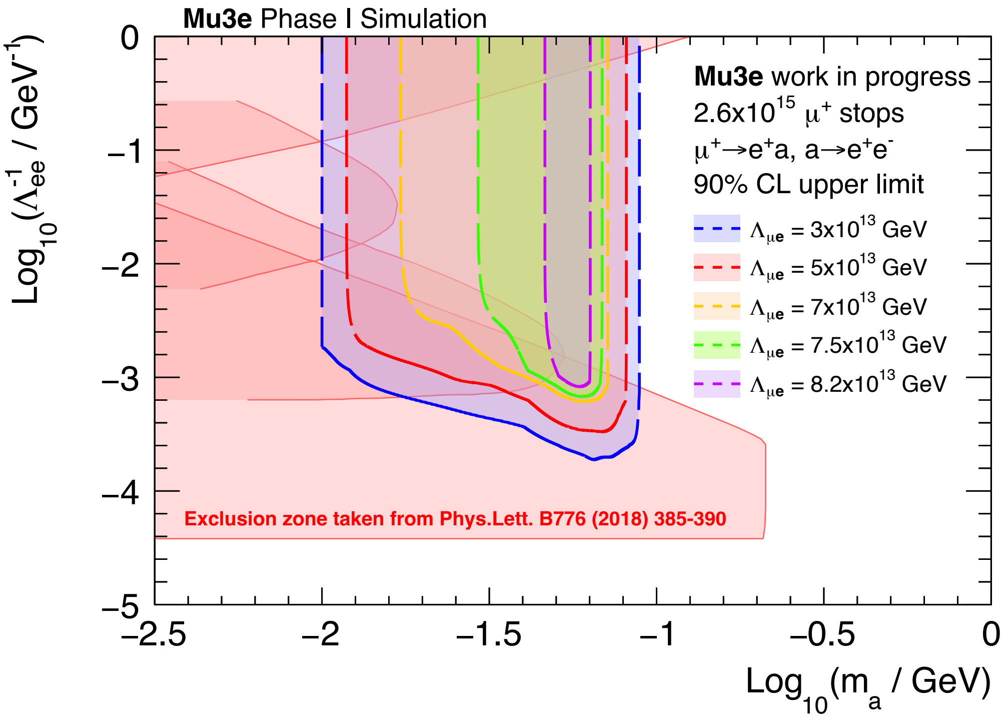

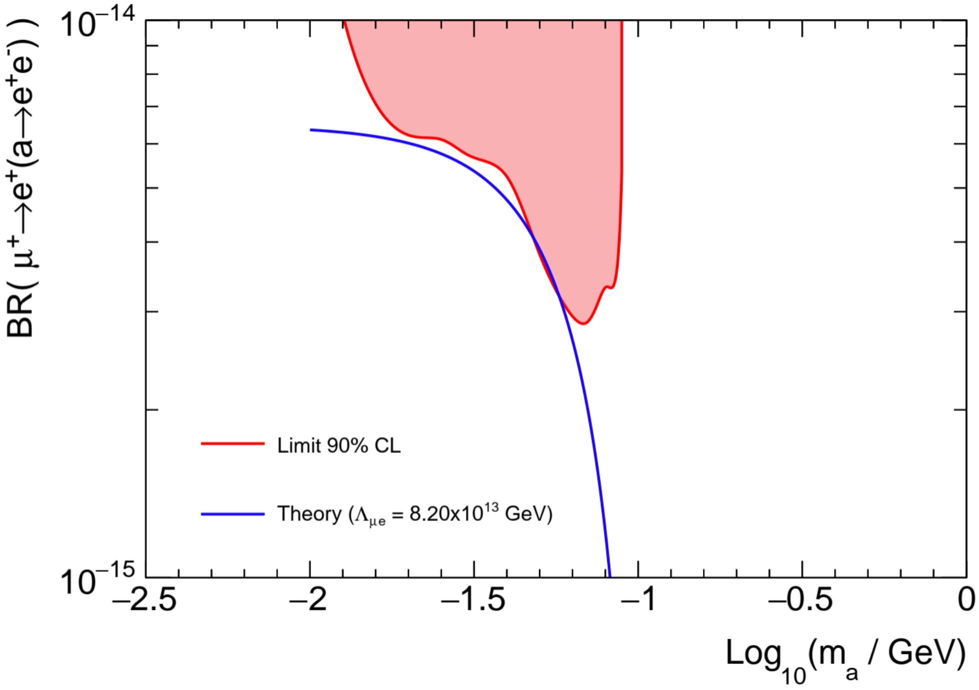

Figure 6.9: Expected Mu3e sensitivities to ALPs of masses 10 to 90 MeV/c to 90% CL for various $\Lambda_{e\mu}$ values.

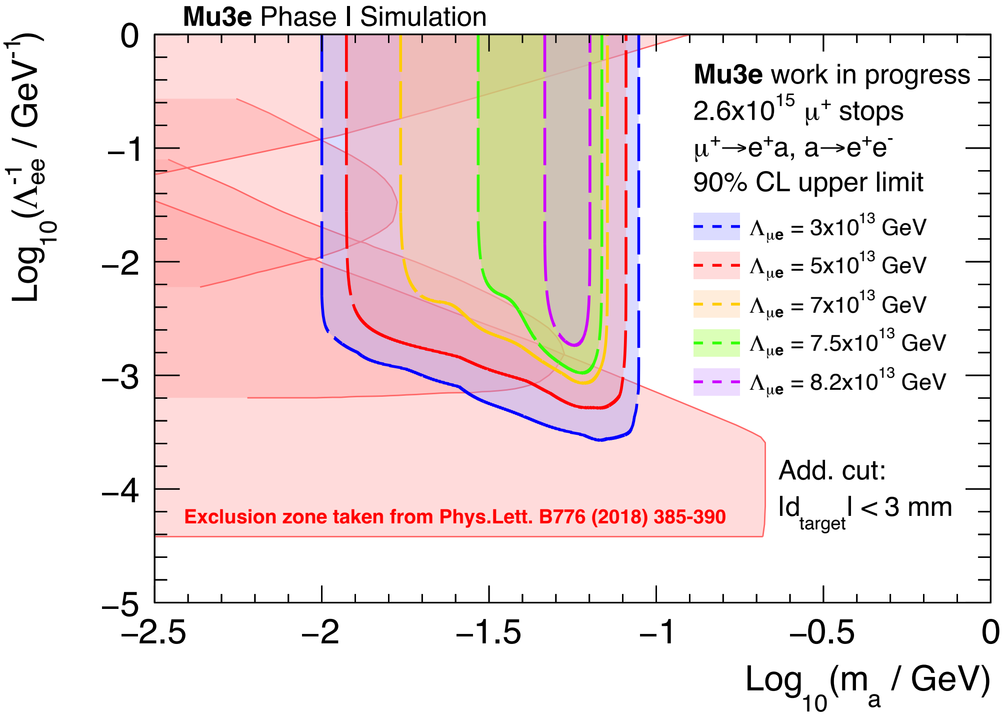

Figure 6.10: Expected Mu3e sensitivities to ALPs of masses 10 to 90 MeV/c to 90% CL for various $\Lambda_{e\mu}$ values. With a cut on the distance from the reconstructed vertex to the target.

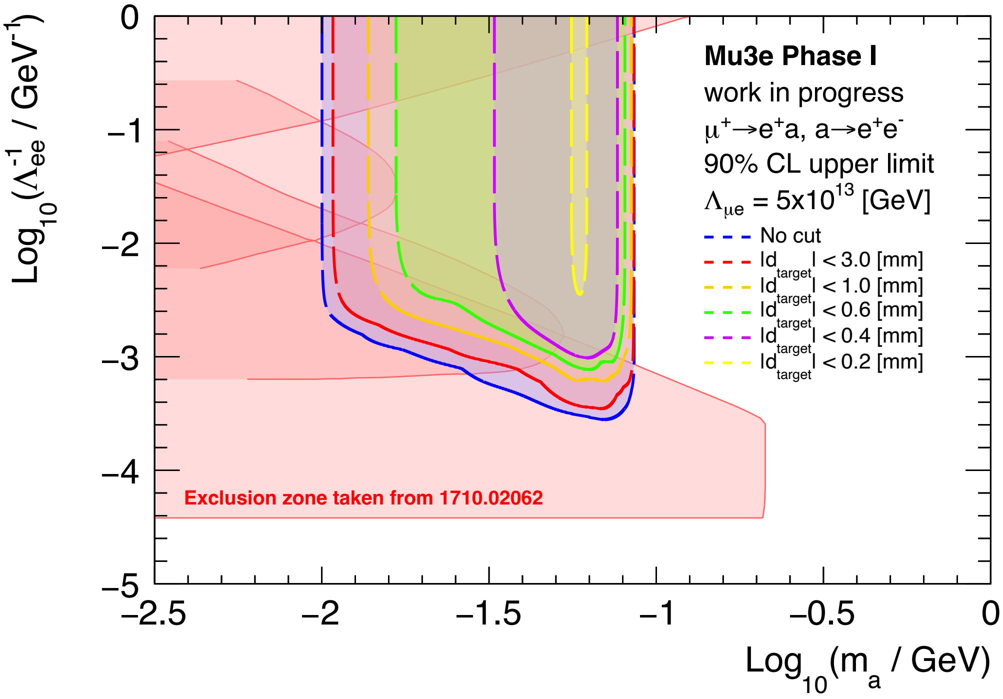

Figure 6.11: Sensitivity plot with varying target distance cuts.

6.8 Conclusions and discussion

This chapter aims to introduce and motivate a new physics study within Mu3e that required no change to data collection. Due to the unprecedented number of muon stops Mu3e aims to analyse, it can be shown that the Mu3e experiment can set competitive limits to current ALP couplings using prompt reconstruction algorithms.

Future work may involve the reconstruction of displaced vertices with Mu3e, as a means of increasing the ALP signal efficiency. However, this is not within the scope of this work. This is nonetheless important, as further cuts on the Mu3e signal are likely to be added, namely, requirements on the distance of the reconstructed vertex and the target, which would only negatively affect the current ALP signal efficiency. The Mu3e Phase-1 detector and beam both nonetheless have impressive sensitivity to new physics.

Chapter 7

Conclusion

The Mu3e experiment aims to search for the charged lepton flavour violating decay $\mu^+ \to e^+e^+e^-$ to an unprecedented sensitivity of one in 10^{16} muon decays. The decay is heavily suppressed within the standard model. Several BSM models however greatly enhance the branching fraction, making the search for the decay an excellent probe for BSM physics. The experiment faces many challenges, mainly due to the various backgrounds and the high rate from the incoming muon beam. These challenges are mitigated by a novel experiment design, which utilises the highest intensity muon beam in the world, with a detector consisting of highly granular, thin silicon pixel detectors, as well as scintillators for the purpose of timing information.

The use of simulations to meet these various challenges is vital, and in Chapter 2, the Mu3e simulation framework was described in detail. The reconstruction and vertex algorithms were introduced, accompanied by a $\mu \to eee$ sensitivity study, where an efficiency of 12.8% is reported. The physics reach of Mu3e beyond the search for $\mu \to eee$ is also shown, with a study into dark photons, where it is shown that, for masses between 10 and 70 MeV, Mu3e has a competitive sensitivity to previously unexplored phase space after 300 days of data taking. The physics reach of Mu3e was extended further by motivating a study into axion-like particles with no change to data collection. Mu3e would have a sensitivity to previously unexplored phase space, reaching a maximum energy scale of 8.2×10^{13} for axion-like particles.

The Mu3e experiment is approaching its commissioning phase, and integration runs have become more frequent. The software simulation work that accompanied an integration run involving a prototype vertex detector is showcased. The prototype was simulated using Geant4 within this thesis. The study explored the possibility of using only hits in the vertex layers to construct a Long-4 track. It was however found that the fake-rate at low momentum was very high, and the use of the tracks for Phase-I analysis requires further study.

The Mu3e experiment will begin commissioning runs in 2024, and will begin to take data in 2025, where it will run for 300 days before the scheduled shutdown for the high intensity muon beam upgrade. Resulting from this upgrade, the Mu3e target will face an order higher ($\mathcal{O}(10^9)$) muon beam intensity, which will greatly increase the amount of accidental background. Geant4 simulations of new and innovative detector designs will be required to face this challenge. One such configuration was explored within the thesis, and involved the placement of an extra outer layer at a high radius. This configuration did not improve the sensitivity to the primary $\mu \to eee$ search, but could help improve the background rejection with the improved momentum resolution. Furthermore, whilst not solving the issues with the accidental background that Mu3e will face, it showed the flexibility of the Mu3e software to modify the detector design and run simulations. This study defined new extra-long (EL) tracks, and calculated their momentum resolutions. This resulted in a new minimum momentum resolution for roughly 50 MeV tracks.

To conclude, the Mu3e experiment offers an exciting physics program that is certain to push the boundaries of the standard model. Data taking for Phase-I is to begin in 2026. After the completion of the high intensity muon beam upgrade, data taking for Phase-II will begin in 2028 and beyond, aiming for a sensitivity of 1 in 10^{16} muon decays.

Bibliography

[1] S. H. Neddermeyer and C. D. Anderson. Note on the nature of cosmic-ray particles. *Phys. Rev.*, 51:884–886, 1937.

[2] J. C. Street and E. C. Stevenson. New evidence for the existence of a particle of mass intermediate between the proton and electron. *Phys. Rev.*, 52:1003–1004, 1937.

[3] M. Conversi and O. Piccioni. On the mean life of slow mesons. *Phys. Rev.*, 70:859–873, 1946.

[4] H. Yukawa. On the Interaction of Elementary Particles I. *Proc. Phys. Math. Soc. Jap.*, 17:48–57, 1935.

[5] M. Conversi et al. On the disintegration of negative mesons. *Phys. Rev.*, 71:209–210, 1947.

[6] C. M. G. Lattes et al. Processes Involving Charged Mesons. *Nature*, 159:694–697, 1947.

[7] E. P. Hincks and B. Pontecorvo. Search for gamma-radiation in the 2.2-microsecond meson decay process. *Phys. Rev.*, 73:257–258, 1948.

[8] J. Steinberger. On the range of the electrons in meson decay. *Phys. Rev.*, 74:500–501, 1948.

[9] A. Lagarrigue and C. Peyrou. Capture of negative μ mesons in copper and tin. *Compt. rend.*, 234, 1952.

[10] R. B. Leighton et al. The energy spectrum of the decay particles and the mass and spin of the mesotron. *Phys. Rev.*, 75:1432–1437, 1949.

[11] L. Michel. Interaction between four half-spin particles and the decay of the -meson. *Proceedings of the Physical Society. Section A*, 63(5):514, 1950.

[12] E. O. Lawrence and M. S. Livingston. The Production of High Speed Light Ions Without the Use of High Voltages. *Phys. Rev.*, 40:19–35, 1932.

[13] Anonymous. Minutes of the 1954 Winter Meeting at the University of California, Berkeley, California, December 28-30, 1954. *Phys. Rev.*, 98:240, 1955.

[14] J. Steinberger and H. B. Wolfe. Electrons from Muon Capture. *Phys. Rev.*, 100(5):1490, 1955.

[15] A. V. Crewe et al. The extraction of the beam from the Liverpool synchrocyclotron. II. Experimental work. *Proceedings of the Royal Society of London. Series A. Mathematical and Physical Sciences*, 232(1189):242–251, 1955.

[16] A.W. Merrison. Recent advances in pion and muon physics. I. *Advances in Physics*, 11(41):1–37, 1962.

[17] J. H. Bartley et al. A search for neutrinoless decay modes of the negative muon. *Physics Letters*, 13(3):258–259, 1964.

[18] T. D. Lee and C. N. Yang. Question of Parity Conservation in Weak Interactions. *Phys. Rev.*, 104:254–258, 1956.

[19] T. D. Lee and C. N. Yang. Parity Nonconservation and a Two-Component Theory of the Neutrino. *Phys. Rev.*, 105:1671–1675, 1957.

[20] G. Harris et al. Lifetimes of the τ^+ and ${K_L}^+$-mesons. *Phys. Rev.*, 100:932–932, 1955.

[21] T. D. Lee and J. Orear. Speculations of Heavy Mesons. *Phys. Rev.*, 100:932–933, 1955.

[22] C. S. Wu, E. Ambler, R. W. Hayward, D. D. Hoppes, and R. P. Hudson. Experimental Test of Parity Conservation in β Decay. *Phys. Rev.*, 105:1413–1414, 1957.

[23] R. L. Garwin et al. Observations of the Failure of Conservation of Parity and Charge Conjugation in Meson Decays: the Magnetic Moment of the Free Muon. *Phys. Rev.*, 105:1415–1417, 1957.

[24] R. P. Feynman and M. Gell-Mann. Theory of the Fermi Interaction. *Phys. Rev.*, 109:193–198, 1958.

[25] G. Feinberg. Decays of the μ Meson in the Intermediate-Meson Theory. *Phys. Rev.*, 110:1482–1483, 1958.

[26] K. Nishijima. Vanishing of the Neutrino Rest Mass. *Phys. Rev.*, 108:907–908, 1957.

[27] J. S. Schwinger. A Theory of the Fundamental Interactions. *Annals Phys.*, 2:407–434, 1957.

[28] G. Culligan et al. Longitudinal Polarization of the Positrons from the Decay of Unpolarized Positive Muons. *Nature*, 180(4589):751–752, 1957.

[29] G. Culligan et al. Longitudinal polarization of the electrons from the decay of unpolarized positive and negative muons. *Proceedings of the Physical Society*, 73(2):169, 1959.

[30] R. Bayes et al. Experimental Constraints on Left-Right Symmetric Models from Muon Decay. *Phys. Rev. Lett.*, 106:041804, 2011.

[31] J. F. Bueno et al. Precise measurement of parity violation in polarized muon decay. *Phys. Rev. D*, 84:032005, 2011.

[32] R. A. Swanson. Depolarization of Positive Muons in Condensed Matter. *Phys. Rev.*, 112:580–586, 1958.

[33] S. J. Blundell et al. *Muon Spectroscopy: An Introduction*. Oxford University Press, 2021.

[34] T. Coffin et al. Magnetic moment of the free muon. *Phys. Rev.*, 109:973–979, 1958.

[35] F. Jegerlehner. *The Anomalous Magnetic Moment of the Muon*, volume 274. Springer, Cham, 2017.

[36] G. Charpak et al. Measurement of the anomalous magnetic moment of the muon. *Phys. Rev. Lett.*, 6:128–132, 1961.

[37] J. Bailey et al. Final report on the CERN muon storage ring including the anomalous magnetic moment and the electric dipole moment of the muon, and a direct test of relativistic time dilation. *Nuclear Physics B*, 150:1–75, 1979.

[38] R. M. Carey et al. New measurement of the anomalous magnetic moment of the positive muon. *Phys. Rev. Lett.*, 82:1632–1635, 1999.

[39] H. N. Brown et al. Improved measurement of the positive muon anomalous magnetic moment. *Phys. Rev. D*, 62:091101.

[40] H. N. Brown et al. Precise measurement of the positive muon anomalous magnetic moment. *Phys. Rev. Lett.*, 86:2227–2231, 2001.

[41] G. W. Bennett et al. Measurement of the positive muon anomalous magnetic moment to 0.7 ppm. *Phys. Rev. Lett.*, 89:101804, 2002.

[42] G. W. Bennett et al. Measurement of the negative muon anomalous magnetic moment to 0.7 ppm. *Phys. Rev. Lett.*, 92:161802, 2004.

[43] J. Grange et al. Muon (g-2) Technical Design Report. 1 2015.

[44] D. P. Aguillard et al. Measurement of the Positive Muon Anomalous Magnetic Moment to 0.20 ppm. 2023.

[45] G. Danby et al. Observation of high-energy neutrino reactions and the existence of two kinds of neutrinos. *Phys. Rev. Lett.*, 9:36–44, 1962.

[46] Y. Fukuda et al. Evidence for oscillation of atmospheric neutrinos. *Phys. Rev. Lett.*, 81:1562–1567, 1998.

[47] Q. R. Ahmad et al. Measurement of the rate of $\nu_e + d \to p + p + e^-$ interactions produced by ^8b solar neutrinos at the sudbury neutrino observatory. *Phys. Rev. Lett.*, 87:071301, 2001.

[48] G. Arnison et al. Results on W^{+-} and Z^0 Physics from the UA1 Collaboration. In *21st Rencontres de Moriond: Perspectives in Electroweak Interactions and Unified Theories*, pages 63–72, 1986.

[49] R. Ansari et al. Measurement of W and Z production properties at the CERN pp collider. *Physics Letters B*, 194(1):158–166, 1987.

[50] Electroweak measurements in electron–positron collisions at W-boson-pair energies at LEP. *Physics Reports*, 532(4):119–244, 2013.

[51] Test of the universality of τ and μ lepton couplings in W boson decays from $t\bar{t}$ events at 13 TeV with the ATLAS detector. 2020.

[52] Measurement of Higgs boson decay to a pair of muons in proton-proton collisions at $\sqrt{s} = 13\,\mathrm{TeV}$. 2020.

[53] G. Aad et al. A search for the dimuon decay of the Standard Model Higgs boson with the ATLAS detector. *Phys. Lett. B*, 812:135980, 2021.

[54] A. M. Baldini et al. Search for the lepton flavour violating decay $\mu^+ \to \mathrm{e}^+\gamma$ with the full dataset of the MEG experiment. *Eur. Phys. J. C*, 76(8):434, 2016.

[55] M. Gell-Mann. A schematic model of baryons and mesons. *Physics Letters*, 8(3):214–215, 1964.

[56] G. Zweig. *An SU(3) model for strong interaction symmetry and its breaking. Version 2*, pages 22–101. 1964.

[57] S. L. Glashow et al. Weak interactions with lepton-hadron symmetry. *Phys. Rev. D*, 2:1285–1292, 1970.

[58] J. J. Aubert et al. Experimental observation of a heavy particle j. *Phys. Rev. Lett.*, 33:1404–1406, 1974.

[59] J. E. Augustin et al. Discovery of a narrow resonance in e^+e^- annihilation. *Phys. Rev. Lett.*, 33:1406–1408, 1974.

[60] J. H. Christenson et al. Observation of massive muon pairs in hadron collisions. *Phys. Rev. Lett.*, 25:1523–1526, 1970.

[61] F. Abe et al. Observation of Top Quark Production in $\bar{p}p$ Collisions with the Collider Detector at Fermilab. *Phys. Rev. Lett.*, 74(14):2626–2631, 1995.

[62] S. Abachi et al. Observation of the Top Quark. *Phys. Rev. Lett.*, 74(14):2632–2637, 1995.

[63] M. Thomson. *Modern particle physics / Mark Thomson, University of Cambridge.* Cambridge University Press, Cambridge, United Kingdom ;, 2013.

[64] D. H. Perkins. *Introduction to High Energy Physics.* Cambridge University Press, 4 edition, 2000.

[65] S. L. Glashow. Partial Symmetries of Weak Interactions. *Nucl. Phys.*, 22:579–588, 1961.

[66] S. Weinberg. A model of leptons. *Phys. Rev. Lett.*, 19:1264–1266, 1967.

[67] A. Salam. Weak and Electromagnetic Interactions. *Conf. Proc. C*, 680519:367–377, 1968.

[68] A. Bettini. *Introduction to elementary particle physics / Alessandro Bettini, University of Padua, Italy.* Cambridge University Press, Cambridge, second edition. edition, 2014.

[69] R. Aaij et al. Test of lepton universality in beauty-quark decays. *Nature Physics*, 18(3):277–282, 2022.

[70] V. W. Hughes et al. Formation of muonium and observation of its larmor precession. *Phys. Rev. Lett.*, 5:63–65, 1960.

[71] B. Pontecorvo. Mesonium and anti-mesonium. *Sov. Phys. JETP*, 6:429, 1957.

[72] B. Abi et al. Measurement of the positive muon anomalous magnetic moment to 0.46 ppm. *Phys. Rev. Lett.*, 126:141801, 2021.

[73] B. T. Cleveland et al. Measurement of the solar electron neutrino flux with the homestake chlorine detector. *The Astrophysical Journal*, 496(1):505, 1998.

[74] Y. Farzan and M. Tortola. Neutrino oscillations and Non-Standard Interactions. *Front. in Phys.*, 6:10, 2018.

[75] B. Pontecorvo. Inverse beta processes and nonconservation of lepton charge. *Zh. Eksp. Teor. Fiz.*, 34:247, 1957.

[76] Z. Maki et al. Remarks on the Unified Model of Elementary Particles. *Progress of Theoretical Physics*, 28(5):870–880, 1962.

[77] R. L. Workman et al. Review of Particle Physics. *PTEP*, 2022:083C01, 2022.

[78] S. T. Petcov. The Processes $\mu \to e + \gamma$, $\mu \to e + \overline{e}$, $\nu' \to \nu + \gamma$ in the Weinberg-Salam Model with Neutrino Mixing. *Sov. J. Nucl. Phys.*, 25:340, 1977. [Erratum: Sov.J.Nucl.Phys. 25, 698 (1977), Erratum: Yad.Fiz. 25, 1336 (1977)].

[79] G. Aad et al. Observation of a new particle in the search for the Standard Model Higgs boson with the ATLAS detector at the LHC. *Phys. Lett. B*, 716:1–29, 2012.

[80] A. Crivellin et al. Fermi Constant from Muon Decay Versus Electroweak Fits and Cabibbo-Kobayashi-Maskawa Unitarity. *Phys. Rev. Lett.*, 127(7):071801, 2021.

[81] F. Wilczek. Problem of strong p and t invariance in the presence of instantons. *Phys. Rev. Lett.*, 40:279–282, 1978.

[82] S. Weinberg. A new light boson? *Phys. Rev. Lett.*, 40(4):223–226, 1978.

[83] R. D. Peccei and H. R. Quinn. CP conservation in the presence of pseudoparticles. *Phys. Rev. Lett.*, 38(25):1440–1443, 1977.

[84] Y. Kuno and Y. Okada. Muon decay and physics beyond the standard model. *Rev. Mod. Phys.*, 73:151–202, 2001.

[85] T. Mori. Final Results of the MEG Experiment. *Nuovo Cim. C*, 39(4):325, 2017.

[86] J. Kaulard et al. Improved limit on the branching ratio of $\mu^+ \to e^+$ conversion on titanium. *Phys. Lett. B*, 422:334–338, 1998.

[87] W. Honecker et al. Improved limit on the branching ratio of $\mu \to e$ conversion on lead. *Phys. Rev. Lett.*, 76:200–203, 1996.

[88] W. H. Bertl et al. A Search for muon to electron conversion in muonic gold. *Eur. Phys. J. C*, 47:337–346, 2006.

[89] U. Bellgardt et al. Search for the Decay $\mu^+ \to e^+$e$^+$e. *Nucl. Phys. B*, 299:1–6, 1988.

[90] Michelle J. Dolinski, Alan W. P. Poon, and Werner Rodejohann. Neutrinoless Double-Beta Decay: Status and Prospects. *Ann. Rev. Nucl. Part. Sci.*, 69:219–251, 2019.

[91] S. Abe et al. Search for the Majorana Nature of Neutrinos in the Inverted Mass Ordering Region with KamLAND-Zen. *Phys. Rev. Lett.*, 130(5):051801, 2023.

[92] Michel Hernández Villanueva. Experimental review of Lepton Flavor Violation searches. In *20th Conference on Flavor Physics and CP Violation* , 8 2022.

[93] Bernard Aubert et al. The BaBar detector. *Nucl. Instrum. Meth. A*, 479:1–116, 2002.

[94] A. Abashian et al. The belle detector. *Nuclear Instruments and Methods in Physics Research Section A: Accelerators, Spectrometers, Detectors and Associated Equipment*, 479(1):117–232, 2002. Detectors for Asymmetric B-factories.

[95] Bernard Aubert et al. Searches for Lepton Flavor Violation in the Decays tau+- —> e+- gamma and tau+- —> mu+- gamma. *Phys. Rev. Lett.*, 104:021802, 2010.

[96] A. Abdesselam et al. Search for lepton-flavor-violating tau-lepton decays to $\ell\gamma$ at Belle. *JHEP*, 10:19, 2021.

[97] K. Hayasaka et al. Search for Lepton Flavor Violating Tau Decays into Three Leptons with 719 Million Produced Tau+Tau- Pairs. *Phys. Lett. B*, 687:139–143, 2010.

[98] Gregory Ciezarek, Manuel Franco Sevilla, P. M. Hamilton, Robert Kowalewski, Thomas Kuhr, Vera Lüth, and Yutaro Sato. A Challenge to Lepton Universality in B Meson Decays. *Nature*, 546:227–233, 2017.

[99] Roel Aaij et al. Search for the lepton-flavour violating decays $\mathrm{B}^0_{(s)} \to e^\pm\mu^\mp$. *JHEP*, 03:078, 2018.

[100] Roel Aaij et al. Search for the lepton-flavour-violating decays $B^0_s \to \tau^\pm\mu^\mp$ and $B^0 \to \tau^\pm\mu^\mp$. *Phys. Rev. Lett.*, 123(21):211801, 2019.

[101] J. Heeck and W. Rodejohann. Lepton flavor violation with displaced vertices. *Physics Letters B*, 776:385–390, 2018.

[102] H. Andernach and F. Zwicky. English and Spanish Translation of Zwicky's (1933) The Redshift of Extragalactic Nebulae. *arXiv e-prints*, page arXiv:1711.01693, 2017.

[103] G. Bertone and D. Hooper. History of dark matter. *Rev. Mod. Phys.*, 90:045002, 2018.

[104] B. Holdom. Two U(1)'s and ϵ charge shifts. *Physics Letters B*, 166(2):196–198, 1986.

[105] M. Fabbrichesi, E. Gabrielli, and G. Lanfranchi. *The Physics of the Dark Photon*. Springer International Publishing, 2021.

[106] D. Feldman et al. Stueckelberg extension with kinetic mixing and millicharged dark matter from the hidden sector. *Physical Review D*, 75(11), 2007.

[107] C. Bœhm and P. Fayet. Scalar dark matter candidates. *Nuclear Physics B*, 683(1-2):219–263, 2004.

[108] T. Hambye, M. H. G. Tytgat, J. Vandecasteele, and L. Vanderheyden. Dark matter from dark photons: A taxonomy of dark matter production. *Physical Review D*, 100(9), 2019.

[109] W. Bertl et al. Search for the decay $\mu^+ \to e^+\mathrm{e}^+\mathrm{e}$. *Nuclear Physics B*, 260(1):1–31, 1985.

[110] R. Eichler and C. Grab. The SINDRUM-I Experiment. *SciPost Phys. Proc.*, 5:007, 2021.

[111] A. M. Baldini et al. The design of the MEG II experiment. *Eur. Phys. J. C*, 78(5):380, 2018.

[112] M. Meucci. MEG II experiment status and prospect. *PoS*, NuFact2021:120, 2022.

[113] Alessandro M. B. et al. The search for $\mu^+ \to e^+\gamma$ with 10^{-14} sensitivity: The upgrade of the MEG experiment. *Symmetry*, 13(9):1591, 2021.

[114] R. Abramishvili et al. COMET Phase-I Technical Design Report. *PTEP*, 2020(3):033C01, 2020.

[115] R. Abramishvili et al. COMET phase-i technical design report. *Progress of Theoretical and Experimental Physics*, 2020(3), 2020.

[116] L. Bartoszek et al. Mu2e technical design report, 2015.

[117] D. Palo. Status of the meg ii experiment and performance results from the first year's data taking. *Physical Sciences Forum*, 8(1), 2023.

[118] K. Arndt et al. Technical design of the phase i mu3e experiment. *Nuclear Instruments and Methods in Physics Research Section A: Accelerators, Spectrometers, Detectors and Associated Equipment*, 1014:165679, 2021.

[119] V. L. Highland. Some practical remarks on multiple scattering. *Nuclear Instruments and Methods*, 129(2):497–499, 1975.

[120] Daniela Kiselev, Christian Baumgarten, Rudolf Dölling, Pierre Duperrex, Dietmar Götz, Joachim Grillenberger, Davide Reggiani, Markus Schneider, Marco Schippers, Mike Seidel, and Hui Zhang. Status and Future Projects of the PSI High Intensity Proton Accelerator. In *3rd J-PARC Symposium (J-PARC2019)*, page 011004, January 2021.

[121] H. Augustin et al. The MuPix sensor for The Mu3e Experiment. *Nuclear Instruments and Methods in Physics Research Section A: Accelerators, Spectrometers, Detectors and Associated Equipment*, 979:164441, 2020.

[122] I. Perić. A novel monolithic pixelated particle detector implemented in high-voltage CMOS technology. *Nuclear Instruments and Methods in Physics Research Section A: Accelerators, Spectrometers, Detectors and Associated Equipment*, 582(3):876–885, 2007. VERTEX 2006.

[123] H. Augustin et al. The Mu3e data acquisition. *IEEE Transactions on Nuclear Science*, 68(8):1833–1840, 2021.

[124] N. Berger et al. A new three-dimensional track fit with multiple scattering. *Nuclear Instruments and Methods in Physics Research Section A: Accelerators, Spectrometers, Detectors and Associated Equipment*, 844:135–140, 2017.

[125] A. Kozlinskiy, A. Schöning, M. Kiehn, N. Berger, and S. Schenk. A new track reconstruction algorithm for the Mu3e experiment based on a fast multiple scattering fit. *JINST*, 9(12):C12012, 2014.

[126] S. Agostinelli et al. Geant4—a simulation toolkit. *Nuclear Instruments and Methods in Physics Research Section A: Accelerators, Spectrometers, Detectors and Associated Equipment*, 506(3):250–303, 2003.

[127] R. Brun et al. root-project/root: v6.18/02, 2019.

[128] J. Alwall et al. The automated computation of tree-level and next-to-leading order differential cross sections, and their matching to parton shower simulations. *Journal of High Energy Physics*, 2014(7), 2014.

[129] Luca Lista. *Statistical Methods for Data Analysis in Particle Physics*, volume 941. Springer, 2017.

[130] A. K. Perrevoort. *Sensitivity Studies on New Physics in the Mu3e Experiment and Development of Firmware for the Front-End of the Mu3e Pixel Detector*. PhD thesis, U. Heidelberg (main), 2018.

[131] B. E. et al. Projections for dark photon searches at Mu3e. *Journal of High Energy Physics*, 2015(1), 2015.

[132] I. Fernando et al. Targeted Exploration of Dark Photon Parameter Space at Mu3e. 2020.

[133] G. Cowan, K. Cranmer, E. Gross, and O. Vitells. Asymptotic formulae for likelihood-based tests of new physics. *The European Physical Journal C*, 71(2), 2011.

[134] J. Alwall et al. A standard format for Les Houches Event files. *Computer Physics Communications*, 176(4):300–304, 2007.

[135] A. J. Krasznahorkay et al. Observation of anomalous internal pair creation in ^{8}Be: A possible indication of a light, neutral boson. *Phys. Rev. Lett.*, 116:042501, 2016.

[136] J. R. Batley et al. Search for the dark photon in π^0 decays. *Phys. Lett. B*, 746:178–185, 2015.

[137] J. P. Lees et al. Search for a Dark Photon in e^+e^- Collisions at BaBar. *Phys. Rev. Lett.*, 113(20):201801, 2014.

[138] G. Hesketh et al. The Mu3e Experiment, 2022.

[139] C. Degrande et al. UFO – the universal FeynRules output. *Computer Physics Communications*, 183(6):1201–1214, 2012.